Platform Design, Consumer Behavior, and Market Power*

Adrian Düll † Florian Heiss‡ Florian Kerzenmacher§

Heiner Schumacher¶ Joachim Winter‖

Version: March 17, 2026

Abstract

We analyze how the design of an online retail platform shapes the market power of sellers competing on that platform. To this end, we study consumer behavior on a generic price comparison webpage and derive pricing equilibria from a structural demand model estimated on consumers' choices. Consumers exhibit heterogeneous degrees of price sensitivity, even in settings without physical search costs and straightforward price comparisons. As a consequence, the pure-strategy pricing equilibrium features price dispersion among identical firms. Design features of the platform – the presentation of a default through a buy box, the number of transaction dimensions, and the ranking algorithm – matter for price dispersion and sellers' markups, sometimes in interdependent and non-monotonic ways. Simplifying price comparisons can increase equilibrium markups as it may cause a subset of sellers to focus on price-insensitive consumers.

Keywords: Platforms, Internet Commerce, Market Power

JEL Classification: C90, D43, D12, D83

---

*We thank Emmanuelle Auriol, Markus Dertwinkel-Kalt, José Moraga-González, and Laura Grigolon as well as seminar audiences at the University of Cologne and University of Hamburg for valuable comments and suggestions. Financial support from the OeNB Anniversary Fund (Project Grant 19009) and the German Science Foundation via CRC TRR 190 (Project Number 280092119) is gratefully acknowledged.

†University of Innsbruck. E-Mail: adrian.duell@uibk.ac.at.

‡University of Düsseldorf. E-Mail: florian.heiss@hhu.de.

§University of Innsbruck. E-Mail: florian.kerzenmacher@uibk.ac.at.

¶University of Innsbruck and KU Leuven. E-Mail: heiner.schumacher@uibk.ac.at.

‖University of Munich. E-Mail: winter@lmu.de.

1 Introduction

Today online retail platforms and price comparison websites organize the trade of large volumes of goods.[1] In order to grow and to operate profitably, these intermediaries must "get both sides on board" (Rochet and Tirole, 2003): On the one hand, they have to attract consumers with the promise to offer the desired products and to facilitate price comparisons. On the other hand, they must provide sellers access to consumers willing to pay positive markups for their products. How can these conflicting demands be served simultaneously?

In this paper, we provide an answer to this question by studying consumer behavior on a generic price comparison webpage and deriving pricing equilibria from a structural demand model estimated on consumers' choices. We consider a setting with homogeneous products and essentially identical firms where the pricing equilibrium would be the Bertrand-Nash equilibrium with zero markups if consumers were fully rational. We find that even when there are virtually no search costs and price comparisons are straightforward, there is a share of consumers who exhibit price-insensitive choices. Simple means of making price comparisons more difficult – such as splitting up the total price into base price and shipping fee – substantially reduce the consumers' price sensitivity. The pure-strategy equilibrium features price dispersion due to (a) heterogeneity in price sensitivity across consumers and (b) within-consumer heterogeneity in price sensitivity across different dimensions of the transaction. Design features of the platform – the provision of a default through a buy box, the number of transaction dimensions, and the ranking algorithm – matter for sellers' markups, sometimes in interdependent and non-monotonic ways.

There are several mechanisms that potentially affect the consumers' price sensitivity on retail platforms and price comparison webpages: Some consumers may ignore or underweight transaction dimensions beyond the base price (Gabaix and Laibson, 2006, Heidhues et al., 2018, Johnen, 2020). Others may be confused and choose among options in a random manner (Carlin, 2009, Piccione and Spiegler, 2012, Chioveanu and Zhou, 2013). Some may exhibit a tendency to consider mostly top-ranked options (De Los Santos and Koulayev, 2017, Ursu, 2018, Fong et al., 2025). Finally, a fraction of consumers may be inattentive and avoid comparing options altogether if they have a default or if one option is presented as default to them (Hortaçsu et al., 2017, Heiss et al., 2021, Lee and Musolff, 2025). It cannot be ruled out that all of these mechanisms matter for consumers' choices on price comparison websites.

If we want to examine consumer behavior on existing retail platforms and price comparison websites, we face a major identification problem: For many choices, we cannot distinguish

[1] In the second quarter of 2025, total e-commerce sales in the United States reached 304.2 billion USD, which corresponds to 16.3 percent of total retail sales (US Census Bureau, Quarterly E-Commerce Report, Q2 2025, https://www2.census.gov/retail/releases/historical/ecomm/25q2.pdf).

between genuine choices and mistakes made by the consumer. Suppose, for example, that a consumer chooses an option with low base price, but also long delivery time. Without further information, it is unclear whether the consumer chose this option because delivery time is unimportant to her or because she failed to notice it. Alternatively, suppose that a consumer chooses the top-ranked option from a list. If the option is not strictly dominated by another, we cannot determine whether the consumer selected it because it best matches her needs or because she tends to pick the top-ranked option regardless of its features.

To circumvent this obstacle, we run an online experiment on *Prolific* in which subjects have to purchase a (hypothetical) book at the lowest price on a shopping platform that we create. This platform resembles a generic price comparison website so that subjects can navigate it intuitively (in Section 3, Figure 1 and Figure 2 show screenshots). All payoff relevant dimensions – base price, shipping fee, delivery time, and seller rating – are tied to a money-metric that is made transparent to subjects (i.e., higher delivery time and lower seller rating imply higher costs to the consumer). This allows us to distinguish between consumer mistakes and preferences, and it enables us to disentangle the different behavioral mechanisms mentioned above. On the platform, several sellers offer the book at varying prices. The subjects' payoff in the experiment is the realized price savings. Our treatment variation is the complexity of price comparisons and the availability of a buy box that suggests a default option. Regarding complexity, there is either one dimension (base price), or two dimensions (base price and shipping fee), or four dimensions (base price, shipping fee, delivery time, and seller rating). Regarding the availability of a buy box, there is either the option to purchase from the top-ranked seller directly without observing any other offers or this option does not exist (so that subjects observe all offers before making a choice).

In total, 3438 individuals participate in the experiment. Nearly all of them shop online in real life, and the vast majority purchases online at least once a month. We observe that a substantial fraction of consumers do not pick the cheapest seller on our platform, despite significant price differences. We detect all the behaviors mentioned above: A significant group of consumers remains inattentive and uses the buy box, even when price comparisons are straightforward. There is a group of individuals who choose expensive options with low base prices as well as a (small) share of consumers who pick the top-ranked seller even when this seller is the most expensive one. Some choices seem to be made in random manner. We find that overpayments (the price difference between the chosen seller and the cheapest seller) increase significantly in both the complexity of price comparisons and buy box availability.

To examine the link between consumer behavior and sellers' equilibrium markups on the platform, we estimate a structural model of consumer behavior and compute the equilibrium of the corresponding Bertrand-Nash pricing game. In the model, consumer choice is organized as a two-stage process as in Heiss et al. (2021). In the attention stage, consumers choose between purchasing from the top-ranked seller through the buy box when it exists (i.e., remaining inattentive) and examining alternative offers. If they are attentive, they reach the choice stage where they can select between varying options including the offer of the top-ranked seller. The model features two consumer types, a price-sensitive and a price-insensitive type. The price-sensitive type considers (when attentive) all options and prices, but may attach different weights to different transaction dimensions. For example, she may underweight the shipping fee relative to the base price. The price-insensitive type chooses (when attentive) the top-ranked option with some probability and otherwise selects a random offer. Thus, the model allows for a variety of behavioral mechanisms that affect consumers' price sensitivity on the price comparison webpage.

We find that the full model provides a better fit to the data than any restricted version of it. In particular, a model with only one consumer type is statistically rejected. The probability to pay attention (i.e., to not use the buy box) decreases significantly in the complexity of price comparisons. The price-sensitive type constitutes the majority of consumers (around 93 percent) while the price-insensitive type comprises only a minority (around 7 percent). Further, the price-sensitive type attaches significantly less weight to the shipping fee and delivery time costs than to base price and seller rating costs. The price coefficients of this type decrease substantially in the complexity of price comparisons.

The simulations of the pricing equilibrium generate three main results. First, the consumers' behavior on the platform implies that the pure-strategy equilibrium exhibits price dispersion and positive markups for all considered platform designs and ranking algorithms. Two factors induce price dispersion: heterogeneity in consumers' price sensitivity and the varying weights of the price-sensitive consumers on the different transaction dimensions. The first factor is relatively important if price comparisons are simple. In this case, only few sellers compete for price-sensitive consumers (who then choose nearly optimally) and charge fairly low prices. The remaining sellers charge high markups to serve a few price-insensitive consumers. The second factor is relatively important if price comparisons are complex. In that case, most sellers compete for price-sensitive consumers. Sellers with high prices in dimensions that are underweighted relative to the base price (shipping fee and deliver time in our case) then charge higher markups than sellers with relatively low prices in these dimensions. Platform design and ranking algorithms govern the relative strength of these two forces and hence price dispersion and markups.

The second main result is that the complexity of price comparisons has opposing implications for the markups that sellers charge (henceforth *unweighted markups*) and the markups that consumers pay (henceforth *weighted markups*). If price comparisons are complex, most sellers compete for price-sensitive consumers and weighted and unweighted markups tend to be similar. If price comparisons are straightforward, weighted markups are relatively small while unweighted markups are relatively large. In this case, the platform shows many unattractive offers that are chosen by a few (price-insensitive) consumers. Accordingly, price-sensitive consumers benefit while price-insensitive consumers are hurt when price comparisons become more transparent.

Third, the consumers' behavior on the platform implies that the ranking algorithm – which determines the top-ranked seller (who owns the buy box when it exists) – has a major impact on markups. It does so due to consumer inattention and the small share of consumers who trade with the top-ranked seller even if its offer is inferior to that of other sellers. The ranking algorithm increases markups if it overweights some transaction dimensions that sellers cannot change in the short-term (such as the rating) and underweights the base price. In this way, the ranking algorithm favors a few sellers who can charge high markups while also enjoying the benefits from being the buy box seller who serves all inattentive consumers by default. Markups are lowest if the ranking algorithm puts the same weight on all transaction dimensions (as a rational consumer would do). Both consumer types benefit from such a ranking algorithm. Interestingly, price-insensitive consumers may then pay even lower markups than price-sensitive consumers.

Related Literature. The paper contributes to several strands of the industrial organization literature. Its focus on consumer behavior is motivated by the literature that examines the influence of cognitive limitations and behavioral biases on market outcomes. In particular, this literature has considered add-on price neglect and price obfuscation. The effects of different forms of add-on price neglect have been examined, for example, in Ellison (2005), Gabaix and Laibson (2006), Armstrong and Vickers (2012), Heidhues et al. (2018), Johnen (2020), and Inderst and Obradovits (2023). Heidhues et al. (2021) consider a model where consumers have to choose between fully understanding all transaction dimensions of an option and browsing different options without understanding the full extent of their financial consequences (which is a choice that some consumers face also in our experiment). Markets with price obfuscation – where firms confuse consumers with their pricing schedules so that consumers choose between options randomly – have been studied, for example in Spiegler (2006), Carlin (2009), Piccione and Spiegler (2012), Chioveanu and Zhou (2013), and Hefti et al. (2022). Our contribution to this literature is to capture different behavioral mechanisms in a multinomial logit framework that can be estimated using choice data from an online platform. This allows us study the link between platform design, consumer behavior, and market outcomes.

Further, the paper complements the growing literature on platform design, position effects, and market outcomes, see, for example, Ellison and Ellison (2009), Dinerstein et al. (2018), Lam (2023), Yu (2025), and Lee and Musolff (2025). The most closely related paper to ours is Lee and Musolff (2025). They reverse-engineer the algorithm that *Amazon* uses to select the seller who is presented in the buy box. Then they estimate consumer preferences, leveraging sales ranks and the fact that for some products the buy box is empty so that consumers are forced to compare the offers of different sellers. They find that *Amazon* provides "search guidance" and promotes offers in the buy box that consumers often prefer to alternative offers. We follow their approach and model ranking algorithms as a logit model. Our contribution to this literature is that we can examine in detail how the consumers' behavior and platform design jointly shape market outcomes. In particular, we can clarify what types of consumer behavior drive price dispersion on an online retail platform.

More generally, the paper contributes to the theoretical and empirical literature on search, consideration sets, and price dispersion in online markets. To generate price dispersion among identical sellers of homogeneous products, this literature assumes search costs (Stahl, 1989, Janssen et al., 2005) or fixed consideration sets (Varian, 1980, Armstrong and Vickers, 2022). In both cases, the resulting equilibrium is typically a mixed-strategy equilibrium where sellers randomize over an interval of prices. One concern is that price dispersion is not stable as some sellers wish to undercut their rivals once prices are realized. Myatt and Ronayne (2025) resolve this issue by restricting sellers' ability to raise prices. The present paper provides an alternative explanation for price dispersion among essentially identical sellers by assuming that consumers exhibit bounded and heterogeneous degrees of price sensitivity.

Finally, the paper extends the literature that combines structural models and experimental methods. Salz and Vespa (2020) evaluate models of dynamic competition using experiments. Karle et al. (2025) use experiments to study to what extent consumers take scale effects into account when choosing their search effort. Fong et al. (2025) consider a setting that is similar to ours in order to identify search costs and belief biases in a setting with prominence effects. The innovation in this paper is that we build a generic price comparison website (without any monetary or physical search costs) to examine consumer behavior.

The remainder of the paper is organized as follows. In Section 2, we introduce a model of oligopolistic competition on a platform that captures different types of consumer behavior. In Section 3, we explain the experimental setting and reduced-form results from the experiment. In Section 4, we estimate a structural model of consumer behavior and present the estimation results. In Section 5, we simulate the equilibria of the Bertrand-Nash pricing game and derive our main findings. Finally, Section 6 concludes. The appendix contains mathematical proofs, instructions, and further empirical results.

2 A Model of Oligopolistic Competition on a Platform

We model an oligopolistic market where trade takes place via an online platform. There are N consumers $i \in \boldsymbol{N} = \{1, ..., N\}$. Each of them has to purchase one unit of a homogeneous product. They can choose between M sellers $j \in \boldsymbol{M} = \{1, ..., M\}$ who offer this product through the platform. The consumers' choices – and hence the sellers' pricing decisions – potentially depend on the design of the platform. In Subsection 2.1, we describe the consumers' behavior on the platform. In Subsection 2.2, we examine the pricing game and market equilibrium, keeping the sellers' ranking on the platform fixed. Finally, in Subsection 2.3, we extend the pricing game by including the platform ranking algorithm into the model.

2.1 Platform Design and Consumer Behavior

The platform presents the sellers' offers on two webpages that are linked to each other. Consumers first see the *overview webpage*. On this page, they find the offer of the top-ranked seller. They can purchase the product from this seller directly through a buy box (without further price comparisons) or click on the link to a *price comparison webpage*. In the latter case, consumers see a list of all seller offers (including the offer of the top-ranked seller) and their features. They can choose any option from this list. Following Heiss et al. (2021), we model the consumers' behavior as a two-stage process with attention and choice stage. For convenience, we start with the consumers' behavior on the price comparison webpage (choice stage) and then describe their decision on the overview webpage (attention stage).

Price Comparison Webpage (Choice Stage). The offer of seller j on the platform is characterized by a price vector[2] $\boldsymbol{p}_j = (p_j^{[1]}, ..., p_j^{[T]})$ and a design vector $\boldsymbol{d}_j = (d_j^{[1]}, ..., d_j^{[T']})$. The price vector contains the base-price $p_j^{[1]}$ as well as other seller-related transfers such as shipping costs. The design vector captures how the seller's offer is represented on the platform's price comparison webpage. For simplicity, we assume that the seller's label j also indicates its rank on the platform (we introduce flexible rankings at the end of this section).

We distinguish between a consumer's decision-utility and her experienced utility.[3] Consumer i's experienced utility from choosing option j is $v - \sum_{t=1}^{T} p_j^{[t]}$, i.e., her utility from the product v minus the sum of partial prices. On the price comparison webpage, the decision-utility of consumer i from seller j's offer equals

$$u_{ij} = v - \boldsymbol{p}_j \boldsymbol{\beta}_i^{[p]} + \boldsymbol{d}_j \boldsymbol{\beta}_i^{[d]} + \varepsilon_{ij}, \tag{1}$$

[2]Throughout, we write vectors in bold letters. Further, the product of any two vectors in our paper is a scalar. We therefore suppress the notation for the transpose.

[3]See also, for example, Allcott (2013) or Train (2015).

where ε_{ij} is distributed according to the extreme value type I distribution. The vector $\boldsymbol{\beta}_i^{[p]}$ captures the extent to which consumer i takes the different entries of the price vector p_j into account and the vector $\boldsymbol{\beta}_i^{[d]}$ describes to what extent the design variables matter for the consumer's choice. In the following, we use the notation $\boldsymbol{\beta}_i^{[p]} = (\beta_i^{[1,p]}, ..., \beta_i^{[T,p]})$ and $\boldsymbol{\beta}_i^{[d]} = (\beta_i^{[1,d]}, ..., \beta_i^{[T',d]})$. The formulation in equation (1) allows for different consumer behaviors. For example, we can have consumers who put different weights on the elements in the price vector or consumers who react to how an offer is represented on the price comparison webpage, e.g., whether it is the top-ranked offer or not.

The error terms ε_{ij} are distributed i.i.d. for all consumers and sellers. The probability with which consumer i chooses seller j on the price comparison webpage is then given by

$$P_{ij} = \frac{\exp(-\boldsymbol{p}_j\, \boldsymbol{\beta}_i^{[p]} + \boldsymbol{d}_j\, \boldsymbol{\beta}_i^{[d]})}{\sum_{j'=1}^{M} \exp(-\boldsymbol{p}_{j'}\, \boldsymbol{\beta}_i^{[p]} + \boldsymbol{d}_{j'}\, \boldsymbol{\beta}_i^{[d]})}. \tag{2}$$

Overview Webpage (Attention Stage). Consumers first encounter the overview webpage where they observe the price vector $\boldsymbol{p}_1$ of the top-ranked seller. They can purchase the product from this seller directly or continue on the price comparison webpage. The decision of consumer i on the overview webpage is determined by the latent propensity a_i. This propensity equals

$$a_i = q_i + \boldsymbol{p}_1\, \boldsymbol{\beta}_i^{[p,o]} + \boldsymbol{d}\, \boldsymbol{\beta}_i^{[d,o]} + e_i, \tag{3}$$

where $\boldsymbol{d}$ is a vector that captures the design of the platform (e.g., the number of price dimensions), q_i a consumer-specific constant, and e_i a random variable with logistic distribution and zero mean. For the vector of price component parameters, we use the notation $\boldsymbol{\beta}_i^{[p,o]} = (\beta_i^{[1,p,o]}, ..., \beta_i^{[T,p,o]})$. Consumer i continues on the price comparison webpage if $a_i > 0$ and purchases from seller 1 on the overview webpage through the buy box if $a_i \leq 0$. Hence, she enters the price comparison webpage with probability

$$\alpha_i = \Pr(a_i > 0) = \frac{1}{1 + \exp(-q_i - \boldsymbol{p}_1\, \boldsymbol{\beta}_i^{[p,o]} - \boldsymbol{d}\, \boldsymbol{\beta}_i^{[d,o]})}. \tag{4}$$

The general idea here is that an increase in the price variables of seller 1 or a drop in the complexity of the price comparison webpage may have a positive impact on the chance that consumer i visits the price comparison webpage. Each consumer i is characterized by how she reacts to the price and design variables on the two webpages. We summarize consumer i's characteristics in the vector $\boldsymbol{\Gamma}_i = (\boldsymbol{\beta}_i^{[p,o]}, \boldsymbol{\beta}_i^{[d,o]}, \boldsymbol{\beta}_i^{[p]}, \boldsymbol{\beta}_i^{[d]}, q_i)$.

2.2 Pricing Game and Market Equilibrium

The Pricing Game. Seller j produces the product at constant marginal costs $c_j < v$ and chooses its base price $p_j^{[1]}$. All other elements of the price vector, $p_j^{[2]}, ..., p_j^{[T]}$, as well as the design vectors $\boldsymbol{d}$ and $\boldsymbol{d}_j$ are exogenously given (sellers cannot change partial prices other than the base price in the short run). Define by $\bar{p}_j = \sum_{t=1}^{T} p_j^{[t]}$ the total price of seller j. Sellers know the consumers' characteristics, $\boldsymbol{\Gamma}_i$ for all $i \in \boldsymbol{N}$, the design of the platform $\boldsymbol{d}$, and how offers are represented on the platform, $\boldsymbol{d}_j$ for all $j \in \boldsymbol{M}$. We say that there is a representative consumer if a vector $\boldsymbol{\Gamma} = (\boldsymbol{\beta}^{[p,o]}, \boldsymbol{\beta}^{[d,o]}, \boldsymbol{\beta}^{[p]}, \boldsymbol{\beta}^{[d]}, q)$ exists so that we have $\boldsymbol{\Gamma}_i = \boldsymbol{\Gamma}$ for all consumers $i \in \boldsymbol{N}$ (in which case we drop the i-subscripts in the formulas).

Sellers choose their base prices simultaneously with the objective to maximize their profit from trade on the platform. Seller j's total price $\bar{p}_j$ must be in the interval $[c_j, \bar{p}_{\max}]$ for some $\bar{p}_{\max}$ with $c_j < \bar{p}_{\max} \leq v$ for all $j \in \boldsymbol{M}$. The upper bound on the total price may represent that the platform excludes sellers whose offers are too unfavorable for consumers. Further, we assume that $p_j^{[2]} + ... + p_j^{[T]} \leq c_j$ for each seller j, so that there is no artificial restriction of total prices in equilibrium. Let $\bar{\boldsymbol{p}} = (\bar{p}_1, ..., \bar{p}_M)$ be the vector of total prices and $\bar{\boldsymbol{p}}_{-j}$ the vector of total prices of seller j's rivals. Denote by $\mathcal{P}$ the set of feasible total price vectors in the pricing game. The log-profit function of the top-ranked seller 1 is

$$h_1(\bar{\boldsymbol{p}}) = \log(\bar{p}_1 - c_1) + \log(\sum_{i=1}^{N}(1 - \alpha_i + \alpha_i P_{i1})), \tag{5}$$

and the log-profit function of any seller $j \neq 1$ equals

$$h_j(\bar{\boldsymbol{p}}) = \log(\bar{p}_j - c_j) + \log(\sum_{i=1}^{N} \alpha_i P_{ij}). \tag{6}$$

Let $\mathrm{br}_j(\cdot)$ be seller j's best reply (in terms of total price) to its rivals' total prices $\bar{\boldsymbol{p}}_{-j}$ and $\mathbf{br}(\cdot) = (\mathrm{br}_1(\cdot), ..., \mathrm{br}_M(\cdot))$ the best reply map. Assume for a moment that $\mathrm{br}_j(\cdot)$ is single-valued for each $\bar{\boldsymbol{p}} \in \mathcal{P}$ and every seller $j \in \boldsymbol{M}$. For seller 1 it is then either implicitly defined by the first-order condition

$$\frac{\partial h_1}{\partial p_1^{[1]}} = \frac{1}{\bar{p}_1 - c_1} - \frac{1}{\sum_{i=1}^{N}(1 - \alpha_i(1 - P_{i1}))}\left(\sum_{i=1}^{N} \alpha_i(1 - P_{i1})(\beta_i^{[1,p,o]}(1 - \alpha_i) + \beta_i^{[1,p]} P_{i1})\right) = 0, \tag{7}$$

or it is given by the maximal total price $\bar{p}_{\max}$. For a seller $j \neq 1$ the best reply is either implicitly defined by the first-order condition

$$\frac{\partial h_j}{\partial p_j^{[1]}} = \frac{1}{\bar{p}_j - c_j} - \frac{1}{\sum_{i=1}^N \alpha_i P_{ij}} \left(\sum_{i=1}^N \alpha_i \beta_i^{[1,p]} P_{ij}(1 - P_{ij}) \right) = 0, \tag{8}$$

or given by the maximal total price $\bar{p}_{\max}$. An equilibrium of the pricing game refers to a Nash equilibrium. By Glicksberg's theorem, at least one (possibly mixed) market equilibrium exists. We examine this equilibrium in two steps. First, we consider the *restricted game-$\tilde{p}$*. In the restricted game-$\tilde{p}$, the sellers $j \in \{2, ..., M\}$ choose their base prices (within the limits imposed by $\mathcal{P}$) while seller 1 is bound to charge $\bar{p}_1 = \tilde{p} \in [c_1, \bar{p}_{\max}]$. Second, we characterize the market equilibrium in the full game where all sellers can choose their base prices.

Market Equilibrium in the Restricted Game-$\tilde{p}$. The restricted game-$\tilde{p}$ is potentially supermodular (Milgrom and Roberts, 1990, Topkis, 1998, Vives, 1990, 1999). In our context, this is a fairly desirable property since it implies that the market equilibrium is unique and can be computed through a variety of methods. Seller j's objective function exhibits the increasing differences property if we have

$$\frac{\partial^2 h_j}{\partial p_j^{[1]} \partial p_{j'}^{[1]}} \geq 0 \ \text{ for all } \ j' = \{2, ..., M\} \setminus \{j\} \ \text{ and at all } \ \boldsymbol{\bar{p}} \in \mathcal{P}. \tag{A1}$$

The intuition behind this property is that it becomes more profitable for seller j to charge a higher base price if its rival j' increases its base price. Since the sellers' strategies are one-dimensional, the restricted game-$\tilde{p}$ is supermodular if condition (A1) is satisfied for all sellers $j \in \{2, ..., M\}$. We then get the following result (the proof is in Appendix A.1).

Proposition 1 (Market Equilibrium in the Restricted Game-$\tilde{p}$)**.** *For any given $\tilde{p} \in [c_1, \bar{p}_{\max}]$ consider the restricted game-$\tilde{p}$. If condition* (A1) *is satisfied for all sellers $j \in \{2, ..., M\}$, then this game has a unique equilibrium. This equilibrium can be approximated by iterating the best reply map.*

Unfortunately, the restricted game-$\tilde{p}$ is not supermodular for all configurations of consumer characteristics $\{\boldsymbol{\Gamma}_i\}_{i \in N}$. However, for some important configurations assumption (A1) is satisfied for all sellers $j \in \{2, ..., M\}$. First, this is the case if we have a representative consumer. We can verify this by computing, for $j \in \{2, ..., M\}$, the first derivative

$$\frac{\partial h_j}{\partial p_j^{[1]}} = \frac{1}{\bar{p}_j - c_j} - \beta^{[1,p]}(1 - P_j), \tag{9}$$

and, for $j' \in \{2, ..., M\} \setminus \{j\}$, the cross derivative

$$\frac{\partial^2 h_j}{\partial p_j^{[1]} \partial p_{j'}^{[1]}} = (\beta^{[1,p]})^2 P_j P_{j'}. \tag{10}$$

Note that the term on the right-hand side of this equality is strictly positive. Therefore, if there is a representative consumer, the pricing game has a unique equilibrium and this equilibrium can be approximated by iterating the best reply map. Next, suppose that (i) there is a subset of consumers $\boldsymbol{N}_1 \subset \boldsymbol{N}$ for which we have a representative consumer, $\boldsymbol{\Gamma}_i = \boldsymbol{\Gamma}$ for all $i \in \boldsymbol{N}_1$, and (ii) all other consumers do not react to base prices, $\beta_i^{[1,p]} = 0$ for all $i \notin \boldsymbol{N}_1$. In this case, we can show that assumption (A1) is satisfied for all sellers $j \in \{2, ..., M\}$ if, all else equal, the fraction of the first group of consumers $\frac{|\boldsymbol{N}_1|}{|\boldsymbol{N}|}$ is sufficiently large.

Market Equilibrium in the Pricing Game. We now consider the full pricing game where all sellers can choose their base prices. This game is in general not supermodular, even when the restricted game-$\tilde{p}$ has this property. Seller 1 may want to lower its total price when another seller $j \neq 1$ raises $\bar{p}_j$, so that assumption (A1) is violated.[4]

However, even when the pricing game is not supermodular, we can characterize under which circumstances we obtain a unique equilibrium. Note that it is optimal for seller 1 to charge the maximum total price $\bar{p}_1 = \bar{p}_{\max}$ if, at all total prices $\bar{p}_1 \in [c_1, \bar{p}_{\max}]$, there are sufficiently many consumers who trade with seller 1 through the buy box and few consumers who trade with it through the price comparison webpage. Suppose that assumption (A1) is satisfied for all sellers $j \in \{2, ..., M\}$. Define by $\bar{\boldsymbol{p}}_{-1}^*(\tilde{p})$ the vector of total prices of the sellers $j \in \{2, ..., M\}$ in the equilibrium of the restricted game-$\tilde{p}$. By Proposition 1, this vector is well-defined. We then obtain a unique equilibrium if

$$\frac{\partial h_1(\bar{p}_1, \bar{\boldsymbol{p}}_{-1}^*(\tilde{p}))}{\partial p_1^{[1]}} > 0 \text{ for all } \bar{p}_1, \tilde{p} \in [c_1, \bar{p}_{\max}]. \tag{A2}$$

[4]To see this, assume that there is a representative consumer (so that the restricted game-$\tilde{p}$ is supermodular). Seller 1's objective function then equals $h_1 = \log(\bar{p}_j - c_j) + \log(1 - \alpha(1 - P_1))$. We can calculate

$$\frac{\partial h_1}{\partial p_1^{[1]}} = \frac{1}{\bar{p}_1 - c_1} - \frac{\alpha(1 - P_1)}{1 - \alpha(1 - P_1)}(\beta^{[1,p,o]}(1 - \alpha) + \beta^{[1,p]} P_1),$$

and, for $j \in \{2, ..., M\}$, we obtain

$$\frac{\partial^2 h_1}{\partial p_1^{[1]} \partial p_j^{[1]}} = \frac{\alpha \beta^{[1,p]} P_1 P_j}{(1 - \alpha(1 - P_1))^2}(\beta^{[1,p,o]}(1 - \alpha) + \beta^{[1,p]} P_1) - \frac{\alpha(1 - P_1)}{1 - \alpha(1 - P_1)}(\beta^{[1,p]})^2 P_1 P_j.$$

The term on the right-hand side of this equation can be strictly negative, for example, if $\beta^{[1,p,o]} \approx 0$, $\alpha < 1$ and P_1 is sufficiently small. In this case, seller 1 mostly serves consumers who use the buy box, possibly at the maximal total price $\bar{p}_{\max}$. If another seller $j \neq 1$ now increases $\bar{p}_j$, then lowering the total price $\bar{p}_1$ to serve some more non-users of the buy box becomes more profitable for seller 1.

This condition ensures that at any $\bar{p}_1 \in [c_1, \bar{p}_{\max}]$ it pays off for seller 1 to increase the base price, provided that its rivals play a best response against this price. It is satisfied, for example, if sufficiently many consumers use the buy box – that is, the values α_i are sufficiently small for all total price vectors in $\mathcal{P}$ and all $i \in \boldsymbol{N}$. We summarize this result in the next proposition.

Proposition 2 (Market Equilibrium in the Pricing Game)**.** *Suppose that condition (A1) is satisfied for all sellers $j \in \{2, ..., M\}$. If additionally assumption (A2) is satisfied, then this game has a unique equilibrium. In this equilibrium, seller 1 charges $\bar{p}_1 = \bar{p}_{\max}$ and its rivals charge $\bar{\boldsymbol{p}}^*_{-1}(\bar{p}_{\max})$. This equilibrium can be approximated by iterating the best reply map.*

2.3 Pricing Game with Ranking Algorithm

So far, we assumed that the seller ranking is given. In order to relax this assumption and to capture alternative ranking algorithms, we follow Lee and Musolff (2025) and formulate the ranking algorithm as a logit model. Define by r_j the rank of seller j and by $\boldsymbol{r} = (r_1, ..., r_M)$ the sellers' ranking on the platform. Further, let $\boldsymbol{R}$ be the set of all possible rankings, $\boldsymbol{R}^{[j]}$ the subset of rankings in $\boldsymbol{R}$ in which seller j is ranked first, and $\boldsymbol{r}^{[j]}$ an element from $\boldsymbol{R}^{[j]}$. The *ranking utility* of seller j's offer is given by

$$u_j^{[r]} = -\boldsymbol{p}_j \boldsymbol{\beta}^{[r]} + \zeta_j, \tag{11}$$

where ζ_j is distributed according to the standard double exponential distribution and the vector $\boldsymbol{\beta}^{[r]} = (\beta^{[1,r]}, ..., \beta^{[T,r]})$ defines the algorithms' weights on each price component. By varying the values in the vector $\boldsymbol{\beta}^{[r]}$, we obtain ranking algorithms that may over- or under-weight certain price components. The probability that seller j's offer is ranked first depends on the ranking utilities of all offers and is given by[5]

$$P_j^{[r]} = \frac{\exp(-\boldsymbol{p}_j \boldsymbol{\beta}^{[r]})}{\sum_{j'=1}^{M} \exp(-\boldsymbol{p}_{j'} \boldsymbol{\beta}^{[r]})}. \tag{12}$$

We update the sellers' objective function based on this ranking probability, see Appendix A.1 for details. A seller's base price now influences the choice probability of its offer not only directly through price comparison, but also indirectly through its effect on the chance of being the first-ranked offer (which is presented in the buy box).

[5] One may also define the probability for any other position in the ranking. However, for our purpose, the probability of being ranked first is sufficient.

3 Experimental Design and Reduced-Form Results

We create an online experiment that allows us to examine consumers' shopping behavior under varying platform designs. An important objective of the experimental setup is that the treatment variation – platform design and seller prices – is rich enough so that we can estimate a structural model of consumer behavior based on the framework from Section 2. With this model we then can simulate the market equilibrium for different platform designs and ranking algorithms. In Subsection 3.1, we explain the experimental design and procedures. In Subsection 3.2, we describe the reduced-form results of this experiment.

3.1 Experimental Design

Experimental Design Baseline Treatment. We invite subjects on *Prolific* to an online experiment which consists of two parts. The first part is a survey in which we elicit demographic information (age, gender, education), online labor supply, online shopping experience, as well as measures of cognitive ability, risk tolerance, patience, and trust. After the survey, we describe the second part, which consists of an experimental shopping task. The instructions of the shopping task can be found in Appendix A.2.

In the shopping task, subjects have to "buy" the hypothetical book *Midnight's Witness* on a price comparison webpage that we create. They will not actually purchase a book, but their earnings in the shopping task increase in the price savings that they realize. There are $M = 12$ sellers that offer the book. If a subject purchases the book from seller j at the total price $\bar{p}_j$ USD, then her payoff in the shopping task is $15 - \bar{p}_j$ USD. The screenshots in Figure 1 and Figure 2 show the overview and the price comparison webpage in the baseline treatment.

In the baseline treatment, the total price $\bar{p}_j$ of seller j consists of four partial prices, so that we have $\boldsymbol{p}_j = (p_j^{[1]}, p_j^{[2]}, p_j^{[3]}, p_j^{[4]})$. Their labels are as follows: $p_j^{[1]}$ is the base-price, $p_j^{[2]}$ the shipping fee, $p_j^{[3]}$ delivery time costs, and $p_j^{[4]}$ seller rating costs. The first two entries are displayed as amounts of money in USD. Delivery time costs are represented as delivery time (from same day delivery to 3 days); each additional day costs 0.42 USD. Seller rating costs are displayed as stars (from one to five stars); each star missing from the five-star rating costs 0.93 USD. We choose these numbers so that they are at the same order of magnitude as realistic values of delivery time and seller rating costs.[6]

[6]Delivery time costs: According to Amazon (2025), US Amazon Prime members in 2024 paid a membership fee of 139 USD and ordered on average 100 items. If we assume that the membership fee represents the willingness to pay for two days faster delivery and the average item costs 25 USD, then obtaining a 15 USD item one day earlier is worth around 0.42 USD. Seller rating costs: Ba and Pavlou (2002) find that consumers on *eBay* are willing to pay a price premium of 11.9 percent for the most trusted seller and a premium of 5.7 percent for the seller in the trust category below. If the difference of 6.2 percent corresponds to a 1-star rating change, this would imply that one star is worth around 0.062×15 USD = 0.93 USD to consumers in our setting.

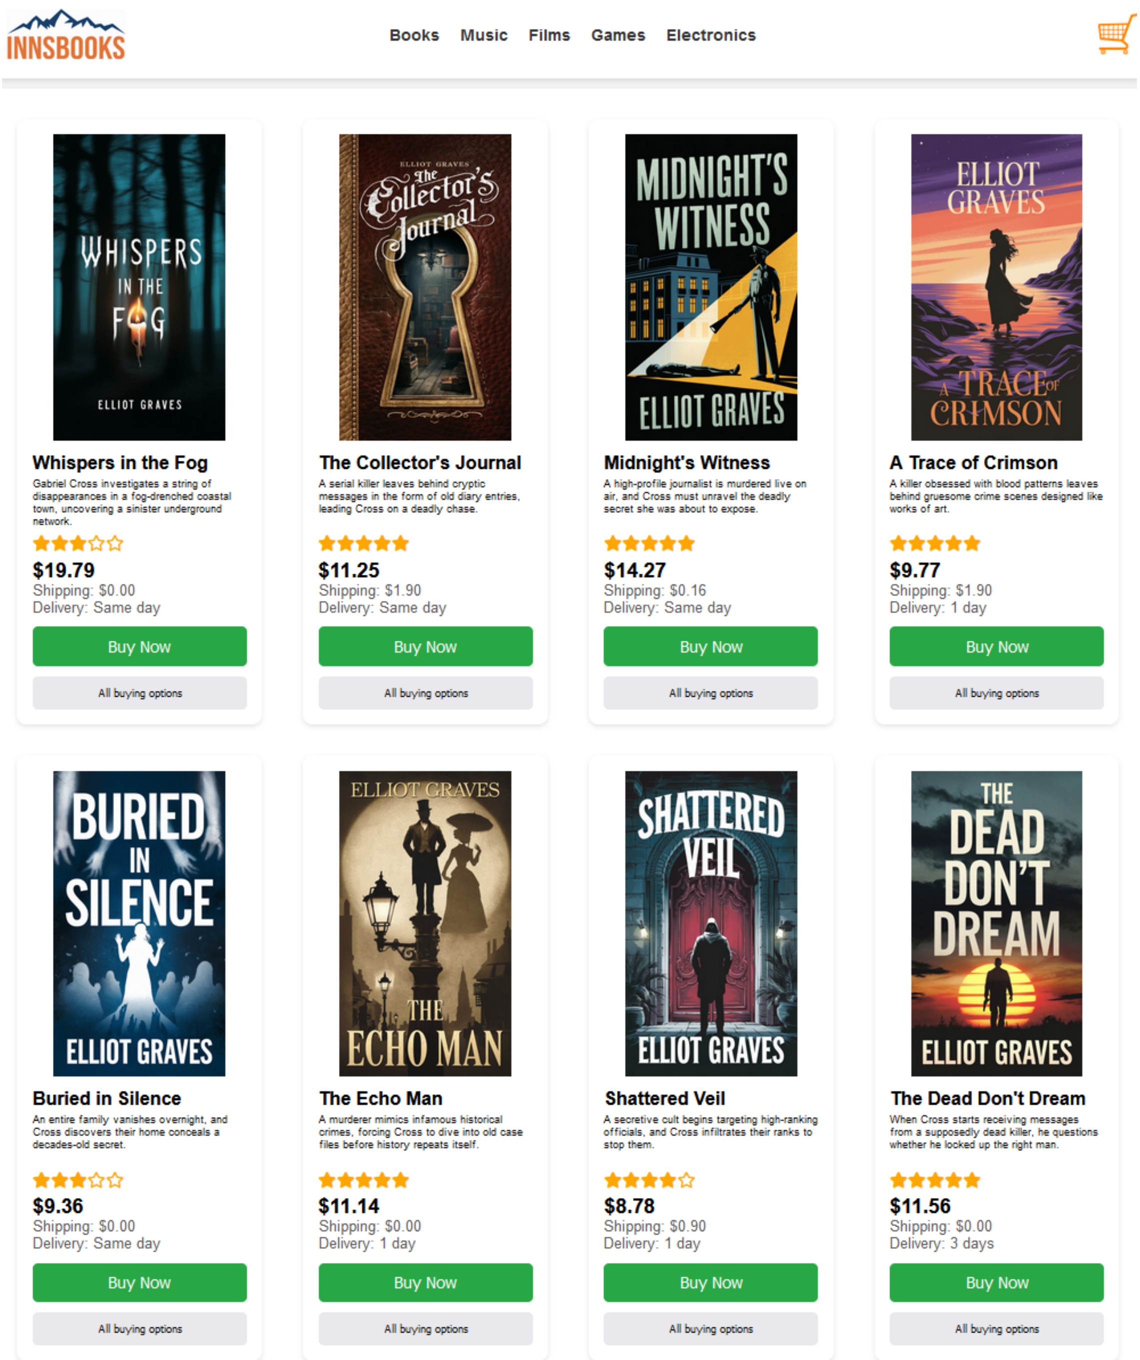

Figure 1: Overview of different products on our shopping website

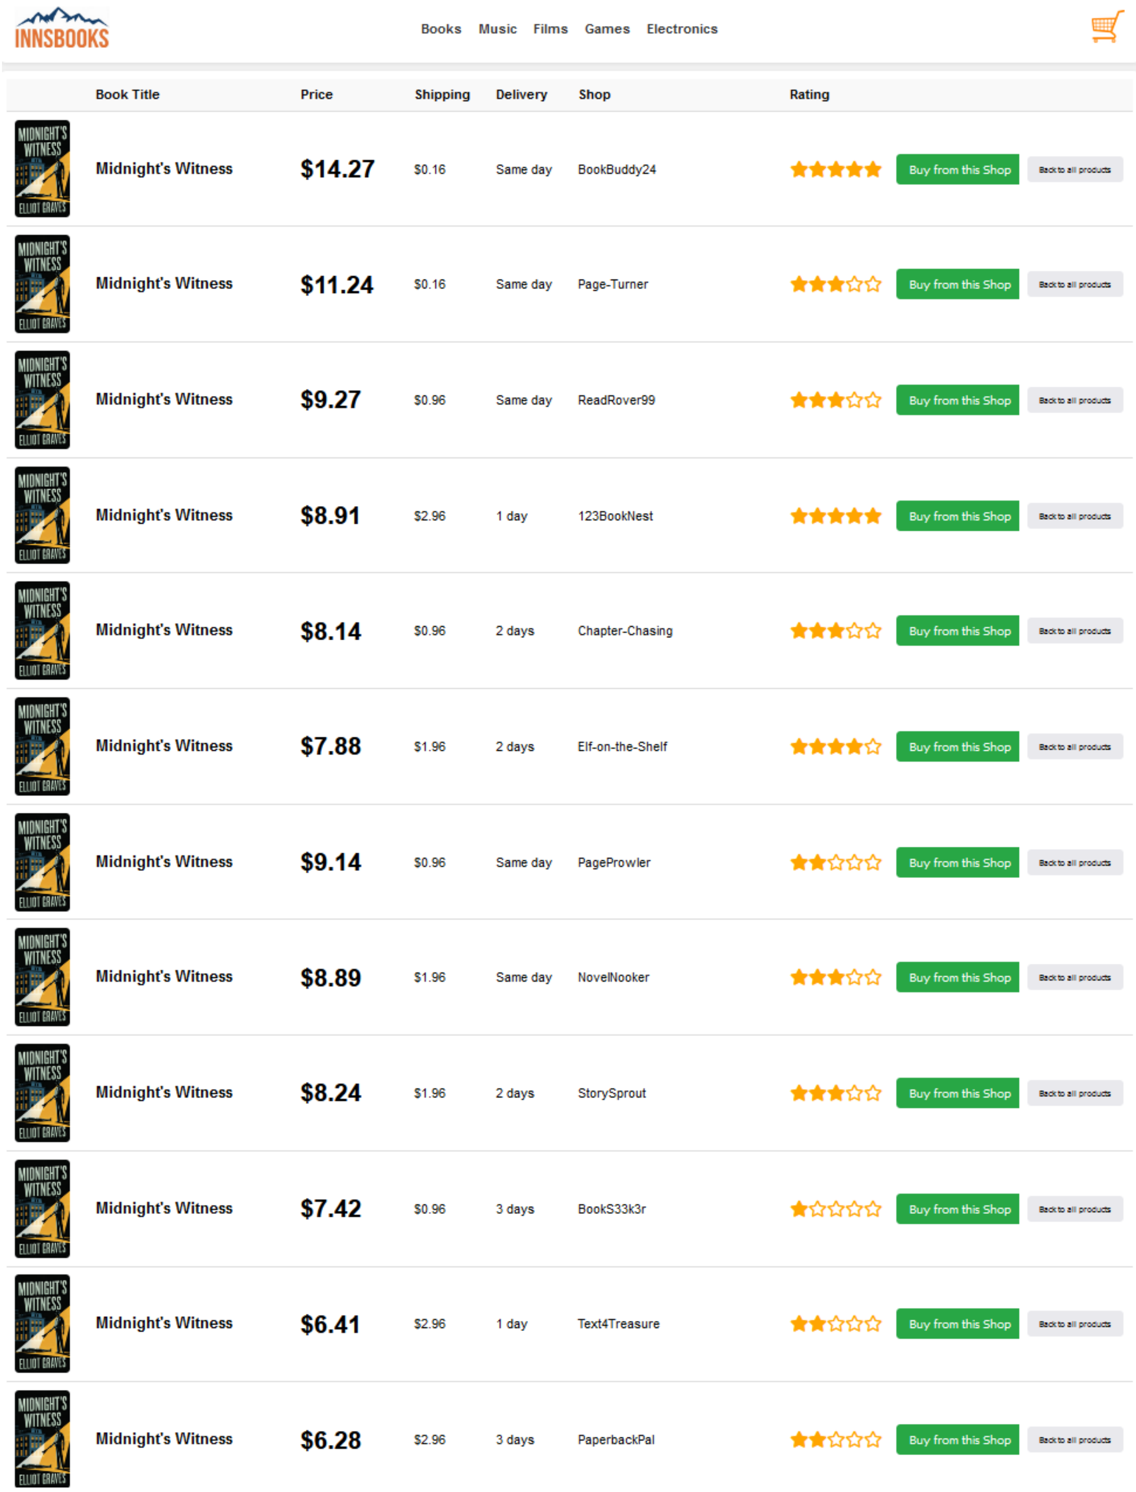

Figure 2: Price comparison for *Midnight's Witness* on our shopping website

Both the overview and the price comparison webpage are designed like a generic online retail platform so that subjects have an intuitive understanding of how they work.[7] At the same time, our representation of different dimensions of an online transaction allows us to disentangle choices that reflect preferences and choices that are driven by behavioral mechanisms. This is crucial to estimate a simple structural model that captures varying types of boundedly rational behavior. Further, this experimental setup ensures that we can use the Bertrand-Nash equilibrium with zero profits as the rational benchmark equilibrium.[8]

At the start of the shopping task, participants first see the overview webpage. This webpage does not only display *Midnight's Witness*, but also other books. By design, it is not possible to purchase any other book. In the baseline treatment, there is a buy box: Participants see the price vector of the top-ranked seller $\boldsymbol{p}_1$ and can directly purchase the book from it by clicking on the "Buy Now" button. Alternatively, they can click on the button "All buying options", in which case they enter the price comparison webpage. This page shows the offers of all sellers. They are ordered according to the sum of the last three entries in the price vector, $p_j^{[2]} + p_j^{[3]} + p_j^{[4]}$, from lowest to highest value. Table 1 displays the price vectors and total prices for all sellers in the baseline treatment and Figure 2 shows how participants see this information on the price comparison webpage. Seller 1 is the most expensive seller (and it is communicated explicitly to subjects that better deals may be available), seller 6 is the cheapest one, and seller 12 offers the product at the lowest base price, but it is actually the second most expensive seller.

Subjects can purchase from the top-ranked seller 1 on the overview webpage through the buy box or they can purchase from any seller 1 to 12 on the price comparison webpage by clicking on a green button next to the corresponding seller's price vector. Upon clicking on such a button, subjects need to confirm their choice. At this stage, the payoff consequences of buying from the chosen seller are displayed on the screen (i.e., the exact calculation that leads to the total price, but not the total price itself) and they still can change their mind. After confirming their choice, the shopping task is completed. If subjects do not trade with any seller, their payoff in the shopping task is zero.

---

[7]The imaginary book titles were generated with ChatGPT-4 of *OpenAI* (in March 2025) and adjusted so that they differ from existing book titles. The prompt used was: *Can you invent the name of an author and the titles of a crime book series that he or she wrote, with 10 to 12 book titles with the respective release years?* Similarly, the fictitious book shop names were generated with ChatGPT-4 to resemble the names of existing online book sellers. The prompt used was: *On Amazon, there are multiple sellers that offer books. Some of the shops are named "bookbot", "bookworm", "bookshopper" etc. Can you invent 10 to 15 imaginary names for book shops in that manner (that do not exist)?*

[8]The implicit assumption here is that sellers and consumers have preferences regarding delivery time and seller rating that are identical, with opposed signs: One day faster delivery increases the costs for the seller and decreases (by the same amount) the waiting costs for the consumer. One star more in the seller rating increases the seller's costs (as it has to spend more resources to implement a more reliable transaction), but decreases (by the same amount) the costs of annoyance on the side of the consumer.

Table 1: Price Vectors and Total Prices in the Baseline Treatment

Seller	(1) base price $p_j^{[1]}$	(2) shipping fee $p_j^{[2]}$	(3) delivery t. costs $p_j^{[3]}$	(4) rating costs $p_j^{[4]}$	(5) total price $\bar{p}_j$
1	14.27	0.16	0.00	0.00	14.43
2	11.24	0.16	0.00	1.86	13.26
3	9.27	0.96	0.00	1.86	12.09
4	8.91	2.96	0.42	0.00	12.29
5	8.14	0.96	0.84	1.86	11.80
6	7.88	1.96	0.84	0.93	11.61
7	9.14	0.96	0.00	2.79	12.89
8	8.89	1.96	0.00	1.86	12.71
9	8.24	1.96	0.84	1.86	12.90
10	7.42	0.16	1.26	3.72	12.56
11	6.41	2.96	0.42	2.79	12.58
12	6.23	2.96	1.26	2.79	13.24

Treatment Variation. We implement six main treatments that we call BB-complex, BB-medium, BB-simple, NBB-complex, NBB-medium, and NBB-simple. BB-complex is the baseline treatment described above. In all main treatments, the sellers' ranking and their total prices are identical, only the presentation of prices and the design of the platform change between treatments. BB-medium is identical to BB-complex except that the total price in this treatment consists of only two partial prices instead of four, base price and shipping fee. For each seller, the base price in BB-medium is identical to the base price in BB-complex, and the shipping fee in BB-medium equals the sum of shipping fee, delivery time, and rating costs from BB-complex. On both the overview and the price comparison webpage, only the base price and the shipping fee are displayed. BB-simple is identical to BB-complex except that there is only one price dimension, the base price. For each seller, the base price in BB-simple equals the sum of the four partial prices in BB-complex. On both the overview and the price comparison webpage, only the base price is presented. Finally, NBB-complex (NBB-medium, NBB-simple) are identical to BB-complex (BB-medium, BB-simple) except that there is no buy box: On the overview page, participants cannot purchase the book right away. Instead, they have to access the price comparison webpage in order to purchase the book.

Parallel to the main treatments, we run additional treatments in which we vary the sellers' price vectors, see Appendix A.3 for an overview of these treatments and the reduced-form results. In particular, we vary the total price of seller 1 so that we can identify to what extent subjects react to the total price of the buy box seller. In the reduced-form analysis, we focus on the six main treatments. For the structural estimation in Section 4, we use all data.

Experimental Procedures. We run the experiment on *Prolific*. In the first part of the experiment, we measure subjects' cognitive ability through a cognitive reflection test, which comprises three questions. We then elicit subjects' willingness to take risks, patience, and their level of trust.[9] Next, we measure online labor supply by asking subjects about their expected hourly earnings and about how many hours per week they spend working on Prolific. Finally, we ask subjects how often they shop online and whether they have memberships with the largest online shopping platforms. Before they start the shopping task, they have to respond to a comprehension check.[10]

We program the experimental software with oTree (Chen et al., 2016). The experiment took place in August and November 2025. Before starting the experiment, we obtained IRB approval from the Board for Ethical Questions in Science of the University of Innsbruck and we pre-registered it on the AEA RCT Registry (registry number AEARCTR-0016524). We recruited only US based subjects with an approval rate of 99 percent or higher, 100 or more previous submissions, and fluency in English. In total, we obtained complete data from 3438 subjects. The treatment samples are balanced in terms of gender, age, education, CRT score, risk tolerance, patience, trust, and shopping frequency, see Appendix A.4. To avoid responses that originate from non-human subjects, we implemented several measures (in addition to the internal quality control system of Prolific).[11]

[9] We elicit the general willingness to take risk as in Dohmen et al. (2011) by asking the question: *Do you generally try to avoid taking risks or are you a risk-taker?* This question has to be answered on a scale between zero (not at all willing to take risks) and ten (very willing to take risks). To elicit patience, we ask the question: *How willing are you to sacrifice something that benefits you today in order to benefit more from it in the future?* This question has to be answered on a scale between zero (not at all willing to sacrifice) and ten (very willing to sacrifice). To elicit trust, we asked the standard question from the world values survey: *Generally speaking, would you say that you need to be very careful in dealing with people or that most people can be trusted?* The response of this question has to be provided on a scale between zero (people can't be trusted at all) and ten (people can fully be trusted).

[10] This check asks about the book subjects have to purchase, the nature of the transaction (whether the book purchase is hypothetical or not), and how the total price of the chosen option influences subjects' payoff in the shopping task. If they provide a wrong answer, they get another chance to provide a correct answer. The shopping task starts as soon as all three questions are answered correctly.

[11] First, we use *Google's* reCAPTCHA. Second, we examine browser headers for patterns that are commonly associated with automated usage, such as terms like "chatgpt", "openai" or indicators of scripting tools. In addition, we check for the presence of unique browser signatures (meta-data automatically transmitted during a web request that may reveal whether the request originated from an AI agent or automated script rather than a standard web browser). Third, we register indicators of user activity, including keyboard presses, mouse clicks, mouse movement, or touchscreen use. Fourth, we implement a hidden "honeypot", i.e., a field that is not visible to a normal user and is accessible only by inspecting the source code of the webpage, allowing us to identify non-human interactions.

3.2 Reduced-form Results

The treatment variation allows us to examine the effects of complexity and buy box availability on subjects' shopping behavior on our platform. Table 2 shows, for each main treatment, the shares of subjects who (a) use the buy box, (b) choose the top-ranked seller (seller 1) on the price comparison webpage, (c) choose the seller with the lowest base-price (seller 12), and (d) choose the cheapest seller (seller 6). At the bottom, it shows subjects' overpayment, i.e., the difference between the total price of the chosen seller and the lowest total price available on the platform.[12] Figure 3 displays the distribution over subjects' choices by seller number.

Table 2: Descriptive Statistics Shopping Behavior (Main Treatments)

	(1)	(2)	(3)	(4)	(5)	(6)
	BB-complex	BB-medium	BB-simple	NBB-complex	NBB-medium	NBB-simple
Share top-ranked seller (through buy box)	0.168	0.086	0.103	–	–	–
Share top-ranked seller (price comparison)	0.018	0.007	0.004	0.044	0.028	0.033
Share lowest base price (price comparison)	0.062	0.050	[0.000]	0.077	0.077	[0.000]
Share lowest total price (price comparison)	0.321	0.695	0.848	0.337	0.700	0.908
Overpayment	0.92	0.42	0.32	0.67	0.32	0.12
(SD)	(1.02)	(0.87)	(0.87)	(0.70)	(0.66)	(0.52)
N	274	279	282	273	287	306

Notes: The shares in squared brackets refer to the seller who has the lowest base price in the complex and medium treatments (but not in the simple treatments where the base price equals the total price).

We find that a significant share of subjects use the buy box, even in the BB-simple treatment where price comparisons are straightforward. The respective shares are 10.3 percent in the BB-simple treatment, 8.6 percent in the BB-medium treatment, and 16.8 percent in the BB-complex treatment. The difference in the shares between the BB-simple and BB-complex treatment is statistically significant (Fisher's Exact Test p-value = 0.026). Thus, the level of complexity increases the share of subjects who purchase the book through the buy box. For all complexity levels, the availability of the buy box more than triples the share of subjects who trade with the top-ranked (and most expensive) seller.

[12]For the two alternative treatments, the corresponding tables can be found in Appendix A.3.

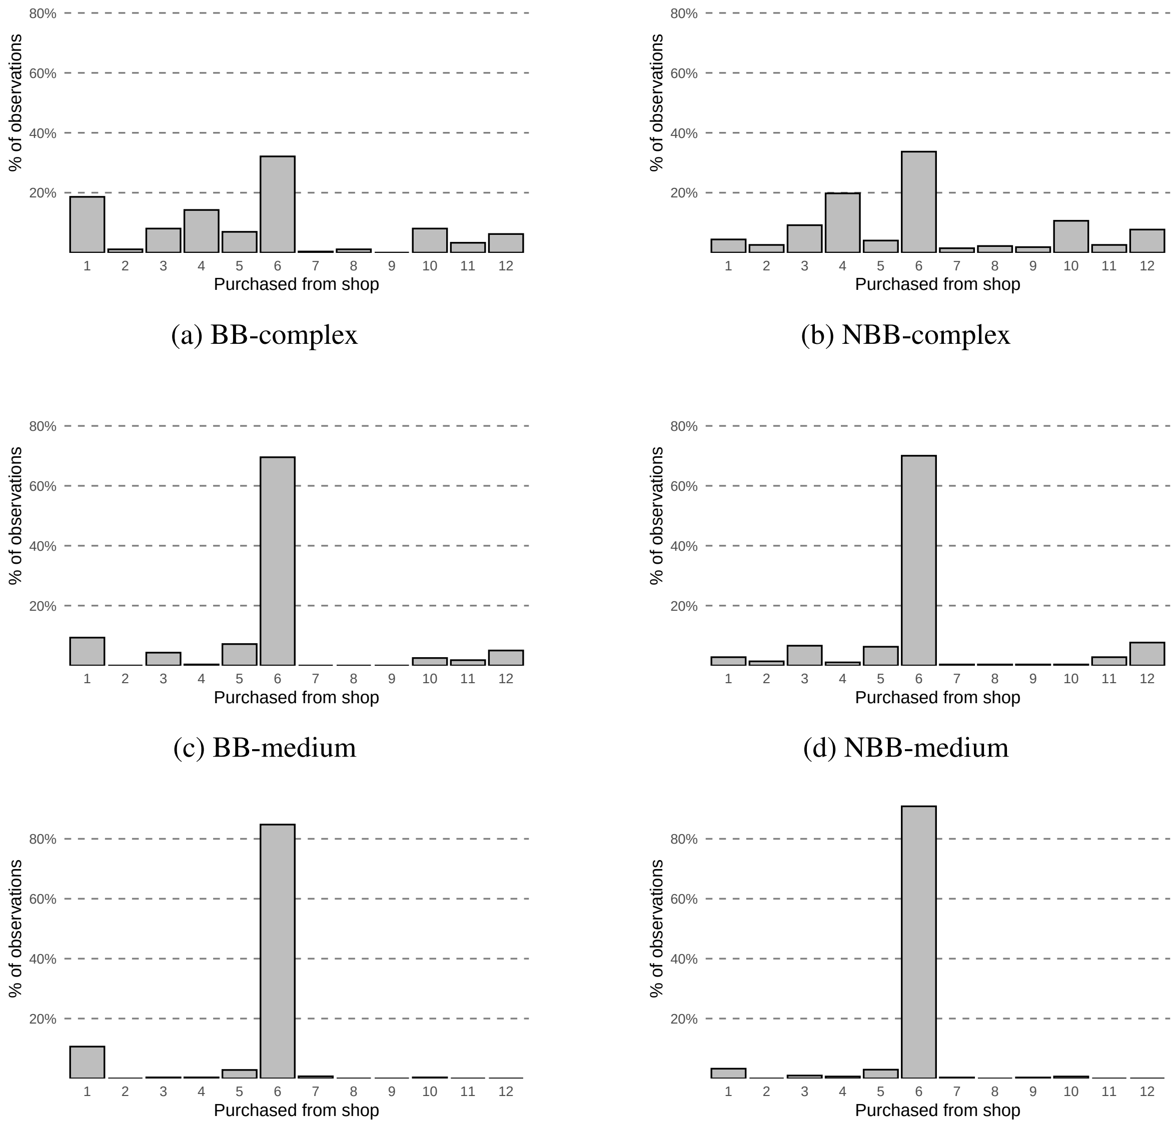

(a) BB-complex (b) NBB-complex

(c) BB-medium (d) NBB-medium

(e) BB-simple (f) NBB-simple

Figure 3: Distribution over purchases (main treatments), by seller number

When there is no buy box, the share of subjects who choose the top-ranked seller is substantially smaller, 3.3 percent in the NBB-simple treatment, 2.8 percent in the NBB-medium treatment, and 4.4 percent in the NBB-complex treatment. The difference between the NBB-simple and NBB-complex treatment is not significant (Fisher's Exact Test p-value = 0.519). The share of subjects who choose the seller with the lowest base price (seller 12) is 7.7 percent, both in the NBB-medium and the NBB-complex treatment. This seller is never chosen in the NBB-simple (and also not in the BB-simple) treatment where its high total price cannot be hidden in other components of the price vector.

Next, we observe that the share of subjects who choose the cheapest seller declines in the level of complexity. In treatments with buy box, this share drops from 84.8 percent in the BB-simple treatment to 32.1 percent in the BB-complex treatment (Fisher's Exact Test p-value < 0.001). In treatments without buy box, the share of subjects who choose optimally decreases from 90.8 percent to 33.7 percent (Fisher's Exact Test p-value < 0.001).

Both complexity and the availability of the buy box tend to increase overpayments: While in the NBB-simple treatment, the average overpayment is only 0.12 USD, this value is 0.92 USD in the BB-complex treatment (Mann-Whitney U test p-value < 0.001). These values correspond to 3.5 percent and 27.1 percent, respectively, of the maximal possible payoff. From Figure 3 we can spot conveniently what drives up overpayments. As price comparisons become more complex, the share of subjects who choose the cheapest seller declines and sales become spread out over different sellers, despite the differences in total prices.

Table 3: Platform Design and Shopping Behavior

	(1)	(2)	(3)	(4)	(5)	(6)
	Buy Box Purchase			Overpayment		
Treatments	Main	Main	All	Main	Main	All
Constant	0.103***	0.112	0.117	0.130***	0.110	0.066
	(0.019)	(0.063)	(0.244)	(0.037)	(0.107)	(0.077)
Total price $\bar{p}_1$			−0.004			
			(0.017)			
Buy Box				0.183***	0.188***	0.189***
				(0.038)	(0.038)	(0.027)
Medium	−0.017	−0.015	−0.012	0.151**	0.152***	0.143***
	(0.027)	(0.027)	(0.019)	(0.046)	(0.046)	(0.033)
Complex	0.065*	0.063*	0.050**	0.572***	0.564***	0.607***
	(0.027)	(0.027)	(0.019)	(0.047)	(0.046)	(0.033)
Controls		✓	✓		✓	✓
N	835	835	1736	1701	1701	3438
R^2	0.012	0.025	0.029	0.098	0.116	0.124

Notes: Results from OLS regressions. Standard errors are in brackets. Controls are gender, age, education, CRT score, risk tolerance, patience, trust, and shopping frequency. The base category in Column (1) to Column (3) is the BB-simple treatment. The base category in Column (4) to Column (6) is the NBB-simple treatment. Significance at * $p < 0.05$, ** $p < 0.01$, and *** $p < 0.001$.

We obtain these results also in a linear regression framework, see Table 3. In Columns (1) to (3) of this table, the dependent variable is a dummy that indicates whether a subject uses the buy box or not. In Columns (4) to (6), the dependent variable is the level of overpayment. In both cases, we first regress the dependent variable on the treatment dummies, using the data from the main treatments. Then we include demographic variables (age, gender, education), the CRT score, risk tolerance, patience, trust, as well as the online shopping frequency as controls in the regression. Finally, we run this regression with all treatments and, for buy box usage, also include the total price as independent variable. We find that complexity has a statistically significant effect on the probability of buying through the buy box. The share of subjects who use the buy box increases by 5.0 percentage points if there are four instead of only one transaction dimensions. Importantly, the total price in the buy box has no significant effect on buy box usage. The corresponding coefficient is close to zero and not statistically significant. Regarding overpayments, we find that both complexity and buy box availability significantly increase the amount of money that subjects leave on the table.[13]

In summary, we see that subjects exhibit different types of boundedly rational behavior. Some subjects avoid price comparisons (even when they are simple) and purchase the book through the buy box. Others choose the top-ranked seller on the price comparison webpage even though this seller offers the worst deal. Some subjects appear to overweight the base price relative to other transaction dimensions, resulting in purchases from sellers who offer a lower base price but a higher total price. Some subjects choose other, non-optimal options. The size of each of these groups is small or modest. In particular, the share of subjects who use the buy box is less than 20 percent – much less than the 90 percent that are reported for purchases on *Amazon* in the recent platform literature. Nevertheless, for the sellers' equilibrium prices all these groups and behaviors may play an important role.

In order to study the impact of heterogeneous consumer behaviors on the sellers' equilibrium prices, we estimate in the next section a structural model that is based on the framework from Section 2. This model is rich enough to capture the different behaviors in our setting and simple enough so that, in the corresponding Bertrand-Nash pricing game, an equilibrium exists that can be obtained in simulations. This allows us to examine the link between platform design and market power, in a way that is informed by subjects' behavior on our price comparison webpage.

[13]Some controls have a significant impact on buy box usage and overpayments. If we use the data from all treatments, we find that buy box usage is negatively correlated with the CRT score and online shopping frequency, and it is positively correlated with risk tolerance. Further, the CRT score and online shopping frequency are negatively associated with overpayments, while risk tolerance is positively correlated with this variable.

4 Empirical Model and Estimation Results

4.1 Empirical Model

To estimate the demand model from Section 2 using the experimental data, we assume the following functional forms. The propensity that subject i pays attention equals

$$a_i = q + \boldsymbol{p}_{1,[i]}\,\boldsymbol{\beta}^{[p,o]} + \boldsymbol{d}_{[i]}\,\boldsymbol{\beta}^{[d,o]} + e_i, \tag{13}$$

where q is a constant, $\boldsymbol{p}_{1,[i]}$ includes the price information from seller 1 given to subject i at the time of deciding about using the buy box. The vector $\boldsymbol{d}_{[i]}$ is the design vector given to subject i. It equals $(1, 0)$ if the treatment assigned to subject i is a simple treatment; it equals $(0, 1)$ if the assigned treatment is a medium treatment, and $(0, 0)$ if it is a complex treatment. Conditional on $\boldsymbol{p}_{1,[i]}$ and $\boldsymbol{d}_{[i]}$ the probability to pay attention (i.e., to not use the buy box) is

$$\alpha_i = \Pr(a_i > 0) = \frac{1}{1 + \exp(-q - \boldsymbol{p}_{1,[i]}\,\boldsymbol{\beta}^{[p,o]} - \boldsymbol{d}_{[i]}\,\boldsymbol{\beta}^{[d,o]})}. \tag{14}$$

We observed that subjects' attention does not depend on the total price of the buy box seller. Hence, we set $\boldsymbol{\beta}^{[p,o]} = \boldsymbol{0}$. Given a subject pays attention, she faces the choice stage where she selects between the different options. There are two subject types that we model as latent types $k = 1, 2$. We define these types below. The probability that a subject is of type k equals

$$\Pr(k) = \frac{\exp(\gamma_k)}{\sum_{k'=1,2} \exp(\gamma_{k'})}. \tag{15}$$

Define the vector of these constants by $\boldsymbol{\gamma} = (\gamma_1, \gamma_2)$. Each type has complexity-specific coefficients $\boldsymbol{\beta}^{[p]}_{k,\boldsymbol{d}}$ and $\boldsymbol{\beta}^{[d]}_{k,\boldsymbol{d}}$, i.e., the level of complexity potentially influences how the different subject types react to the price and the representation of an option. The decision-utility of subject i from seller j's offer if she is of type k equals

$$u_{ijk} = -\boldsymbol{p}_{j,[i]}\,\boldsymbol{\beta}^{[p]}_{k,\boldsymbol{d}_{[i]}} + \boldsymbol{d}_{j,[i]}\,\boldsymbol{\beta}^{[d]}_{k,\boldsymbol{d}_{[i]}} + \varepsilon_{ij}, \tag{16}$$

where ε_{ij} are i.i.d. extreme value type I error terms. The vectors $\boldsymbol{p}_{j,[i]}$ and $\boldsymbol{d}_{j,[i]}$ are the price and design vectors of seller j assigned to subject i. The observed outcome for each subject i is her choice $y_i \in Y$, where $Y = \{\text{bb}, 1, ..., M\}$ is the set of all possible choices; $y_i = \text{bb}$ means that subject i purchases the book from seller 1 through the buy box and $y_i = j \in \{1, ..., M\}$ means that subject i purchases the book from seller j on the price comparison website. On the price comparison webpage, the choice probability of seller j's offer for subject i if she is of type k is

$$\Pr(y_i = j \mid k) = \frac{\exp(-\boldsymbol{p}_{j,[i]}\,\boldsymbol{\beta}^{[p]}_{k,\boldsymbol{d}_{[i]}} + \boldsymbol{d}_{j,[i]}\,\boldsymbol{\beta}^{[d]}_{k,\boldsymbol{d}_{[i]}})}{\sum_{j'=1}^{M}\exp(-\boldsymbol{p}_{j',[i]}\,\boldsymbol{\beta}^{[p]}_{k,\boldsymbol{d}_{[i]}} + \boldsymbol{d}_{j',[i]}\,\boldsymbol{\beta}^{[d]}_{k,\boldsymbol{d}_{[i]}})}. \tag{17}$$

The vector $\boldsymbol{d}_{j,[i]}$ equals $(1,0)$ if $j = 1$ and it equals $(0,1)$ if $j \in \{2, ..., 12\}$. All price vectors have four entries, where in the simple (medium) treatment the last three (two) entries are equal to zero. The probability that subject i chooses seller j's offer given that she pays attention equals

$$\Pr(y_i = j) = \sum_{k=1}^{4} \Pr(k)\Pr(y_i = j \mid k). \tag{18}$$

Define by $\boldsymbol{\theta}$ the vector of all parameters that are estimated through the model, i.e., $\boldsymbol{\gamma}$, $\boldsymbol{\beta}^{[p,o]}$, $\boldsymbol{\beta}^{[d,o]}$ as well as $\boldsymbol{\beta}^{[p]}_{k,\boldsymbol{d}_{[i]}}$ and $\boldsymbol{\beta}^{[d]}_{k,\boldsymbol{d}_{[i]}}$ for all k and $\boldsymbol{d}_{[i]}$. The likelihood contribution of subject i is given by $P_i = \Pr(y_i \mid \boldsymbol{\theta})$ and the log-likelihood function is

$$\ln L = \sum_{i=1}^{N} \ln P_i. \tag{19}$$

With this function we can jointly estimate the behavioral parameters of the model through maximum likelihood estimation. We now define the two latent subject types $k \in \{1, 2\}$ for the choice stage. We call them *price-sensitive type* and *price-insensitive type*. The price-sensitive type ($k = 1$) has the following characteristics:

$$\beta^{[1,p]}_{1,\boldsymbol{d}_{[i]}}, ..., \beta^{[4,p]}_{1,\boldsymbol{d}_{[i]}} \text{ free and } \beta^{[1,d]}_{1,\boldsymbol{d}_{[i]}} = \beta^{[2,d]}_{1,\boldsymbol{d}_{[i]}} = 0. \tag{20}$$

Thus, the price-sensitive type may attach varying weights to the different price components (e.g., she may react more strongly to the base price than to the shipping fee) and she does not give higher weight to the top-ranked product. Next, the price-insensitive type ($k = 2$) features the following characteristics:

$$\beta^{[1,p]}_{2,\boldsymbol{d}_{[i]}} = ... = \beta^{[4,p]}_{2,\boldsymbol{d}_{[i]}} = 0,\ \beta^{[1,d]}_{2,\boldsymbol{d}_{[i]}} \text{ free, and } \beta^{[2,d]}_{2,\boldsymbol{d}_{[i]}} = 0. \tag{21}$$

The price-insensitive type therefore ignores the price components and only considers whether an offer is top-ranked or not. If $\beta^{[1,d]}_{2,\boldsymbol{d}_{[i]}} > 0$, she chooses the top-ranked product with higher probability than any other option; in case she does not choose the top-ranked product, she chooses each remaining option with equal probability.

4.2 Structural Estimates

Table 4 shows the estimation results. The estimated shares of the price-sensitive and price insensitive type are 93.2 percent and 6.8 percent, respectively. The coefficient of the price-insensitive type is statistically significant. The type shares reflect the fact that only a small fraction of subjects chooses the top-ranked product on the price comparison webpage and that most subjects seem to engage in some type of price comparisons.

Table 4: Main Estimation Results (Model V)

	(1) Coefficient	(2) SE	(3) Share
Type Distribution			
Type 1 (p-sensitive type)			93.2%
Type 2 (p-insensitive type)	-2.615***	(0.109)	6.8%
Attention Stage			
Constant	1.652***	(0.114)	
Simple treatment	0.441*	(0.174)	89.0%
Medium treatment	0.630***	(0.184)	90.7%
Complex treatment			83.9%
Choice Stage			
Type 1 (p-sensitive type)			
Base price, simple	18.809***	(1.527)	
Base price, medium	4.685***	(0.174)	
Shipping, medium	3.696***	(0.123)	
Base price, complex	1.636***	(0.064)	
Shipping, complex	1.313***	(0.066)	
Delivery time, complex	1.197***	(0.108)	
Seller rating, complex	1.616***	(0.060)	
Type 2 (p-insensitive type)			
Top-ranked seller, simple	1.405***	(0.311)	27.0%
Top-ranked seller, medium	1.340***	(0.279)	25.8%
Top-ranked seller, complex	3.021***	(0.456)	65.1%

Notes: Results from maximum likelihood estimation. Standard errors are in brackets. The last column shows the type shares (type distribution), the estimated probability with which subjects enter the price comparison webpage (attention stage), and the estimated probability with which the price-insensitive type chooses the top-ranked seller (choice stage). Significance at * $p < 0.05$, ** $p < 0.01$, and *** $p < 0.001$.

We first consider the price-sensitive type. The price coefficients drop significantly as price comparisons become more complex, from a coefficient of around 19 in the simple treatments to coefficients between 1 and 2 in the complex treatments (Wald test, p-values < 0.001). Thus, the complexity of price comparisons reduces the price-sensitive type's price sensitivity. Further, the price-sensitive type attaches different weights to different price components. In the medium treatments, the weight on the base price is significantly larger than the weight on the shipping fee (Wald test, p-value < 0.001). In the complex treatments, base price and seller rating costs receive significantly higher weights than shipping fee and delivery time costs (Wald test, p-values < 0.001). Interestingly, base price and seller rating costs receive roughly the same weight (Wald test, p-value = 0.639).

Next, we consider the price-insensitive type. According to our estimation results, this type attaches strictly positive weights on top-ranked offers. The corresponding coefficients are significant for all complexity levels. They indicate that the price-sensitive type chooses the top-ranked offer with 27.0 percent probability in the simple treatments, with 25.8 percent probability in the medium treatments, and with 65.1 percent probability in the complex treatments. The difference in the coefficients is statistically significant between the complex and the other treatments (Wald test, p-values < 0.003). Thus, the price-insensitive type's propensity to choose the top-ranked product increases in the complexity of price comparisons.

Table 5: Alternative Models

	(1) Model I	(2) Model II	(3) Model III	(4) Model IV	(5) Model V
Parameters	3	6	10	10	14
Model Components					
Design: Complexity	✓	✓	✓	✓	✓
Total Price	✓	✓		✓	
Attention Stage		✓	✓	✓	✓
Price Dimensions			✓		✓
P-Insensitive Type				✓	✓
Estimated Models					
Log likelihood	-5,509.8	-4,720.3	-4,682.2	-4,195.9	-4,150.4
LR test/Model V (df)	2,718.8 (11)	1,139.8 (8)	1,063.6 (4)	91 (4)	
AIC	11,025.5	9,452.6	9,384.4	8,411.8	8,328.7
BIC	11,043.9	9,489.5	9,445.8	8,473.2	8,414.7

Notes: All models are estimated on our analytic sample with 3,438 observation using maximum likelihood estimation. (df) refers to degrees of freedom.

We consider several simplified versions of our structural demand model. These simplifications are obtained by restricting the behavioral parameters. Model I only consists of the choice stage. Buy box purchases are treated as purchases of the top-ranked product on the price comparison webpage. The only free parameters are the total prices for each level of complexity. Next, in Model II, we add the attention stage and thus allow for inattention through buy box purchases. Model III is the same as Model II except that there are price coefficients for different price components. Model IV adds the price-insensitive type, but restricts the price-sensitive type to only care about the total price. Finally, Model V is the full model which comprises the attentions stage and two latent subject types. The price-sensitive type has price coefficients for different price components.

We estimate all models and compare their predictive power. Table 5 provides an overview of the five models and it shows the Akaïke and Bayesian information criterion as well as the likelihood-ratio test statistics for comparisons of each simplified model with Model V. We find that the more flexible models provide a better fit to the data. Model V turns out as best explaining the observed behavior. This result is confirmed by both information criteria.

5 Equilibrium Simulation and Counterfactuals

We simulate the equilibrium of the pricing game by iteratively computing each seller's best response based on their log-profit function. We set $c_j = 8$ for each seller $j \in \boldsymbol{M} = \{1, ..., 12\}$ and $\bar{p}_{\max} = 15$. The consumers' characteristics that govern choice probabilities are given by the structural estimates from Table 4. For each setting, we simulate the equilibrium for 1,000 independent markets where each market has 12 sellers with i.i.d. price components.[14] To evaluate the uniqueness of the equilibrium, we run the simulations for several different initial price vectors from the set $\mathcal{P}$ of feasible total prices. An equilibrium is found when the Euclidean distance between one best-reply vector $\bar{\boldsymbol{p}}_{[\kappa]}$ and the next best-reply vector $\bar{\boldsymbol{p}}_{[\kappa+1]}$ is small enough such that we have

$$\sqrt{\sum_{j=1}^{12}(\mathrm{br}_i(\bar{\boldsymbol{p}}_{[\kappa]}) - \mathrm{br}_i(\bar{\boldsymbol{p}}_{[\kappa+1]}))^2} < 0.1. \tag{22}$$

We now analyze how platform design influences sellers' markups and how this relationship depends on consumer behavior. To this end, we proceed as follows. In Subsection 5.1, we keep the seller ranking fixed as in the experiment. This helps us to build intuition how variations in

[14] Specifically, we draw the shipping fee, delivery time costs, and seller rating costs, respectively, from the sets $p_j^{[2]} \in \{0.16, 0.96, 1.96, 2.96\}$, $p_j^{[3]} \in \{0.00, 0.42, 0.84, 1.26\}$, and $p_j^{[4]} \in \{0.00, 0.93, 1.86, 2.79, 3.72\}$, according to the uniform distribution in each case. These values are the same as in our experimental treatments.

platform design influence market outcomes. In Subsection 5.2, we consider the pricing equilibrium with symmetric ranking algorithms, i.e., the algorithm weights each price component equally. Finally, in Subsection 5.3, we study how markups change if the ranking algorithm puts varying weights on different components of the price vector.

5.1 Fixed Seller Ranking

For each platform design, we calculate the average markup weighted by market shares (w-markup) as well as the average markup that sellers post on the platform (u-markup). Further, we consider the markup of the top-ranked seller, minimal and maximal u-markups, as well as the markups that the two consumer types pay on average.

Table 6 shows the results in Column (1) and Column (2). We observe three patterns. First, with a fixed ranking, the sellers' profits increase both in the complexity of price comparisons and the provision of a default through the buy box. The difference is substantial: In BB-complex, the weighted markup is roughly four times as large as in NBB-simple. The top-ranked seller exploits its position by charging the maximal price $\bar{p}_{\max}$ in all settings. Notably, this is true both in treatments with and without buy box.

Second, we observe that – in contrast to the weighted markup – the unweighted markup decreases in the level of complexity. As price comparisons become more simple, more sellers start charging the maximal price $\bar{p}_{\max}$ to exploit consumers who fail at comparing prices (mostly type 2 consumers). The expected market shares of these sellers are fairly small. Nevertheless, this strategy is for many sellers superior to competing for type 1 consumers. In the simple treatments, only two sellers compete for consumers while the rest serves a few type 2 consumers who make random choices. A consequence of the sellers' equilibrium strategies is substantial price dispersion akin to the Varian (1980) model of sales for all platform designs in our sample.

Third, while type 1 consumers benefit from less complexity, type 2 consumers pay on average more when price comparisons become more transparent. This reflects the change in the sellers' pricing strategies: In the complex treatments, most sellers compete for type 1 consumers so that consumers who choose randomly pay on average a relatively low price. In the medium and simple treatments, only few sellers compete for type 1 consumers so that a random choice leads to a relatively high average price.

Table 6: Platform Design and Market Power (Fixed Ranking)

	(1)	(2)	(3)	(4)	(5)	(6)
	Model V Estimates		Model V Estimates [Pr(k = 2) = 0]		Model III Estimates	
	BB-complex	NBB-complex	BB-complex	NBB-complex	BB-complex	NBB-complex
w-markup, avg.	1.946	0.977	1.696	0.674	1.876	0.886
u-markup, avg.	1.218	1.218	1.199	0.667	1.397	0.880
u-markup, min.	0.674	0.674	0.643	0.639	0.857	0.851
u-markup, max.	7.000	7.000	7.000	0.709	7.000	0.920
u-markup, seller 1	7.000	7.000	7.000	0.644	7.000	0.896
w-markup, type 1	1.712	0.698	1.696	0.674	1.877	0.886
w-markup, type 2	5.152	4.797	–	–	–	–
	BB-medium	NBB-medium	BB-medium	NBB-medium	BB-medium	NBB-medium
w-markup, avg.	1.271	0.686	0.903	0.280	0.999	0.384
u-markup, avg.	5.033	5.035	0.802	0.236	0.932	0.378
u-markup, min.	0.276	0.276	0.215	0.214	0.363	0.359
u-markup, max.	7.000	7.000	7.000	0.328	7.000	0.409
u-markup, seller 1	7.000	7.000	7.000	0.214	7.000	0.359
w-markup, type 1	0.957	0.340	0.903	0.280	0.999	0.384
w-markup, type 2	5.555	5.409	–	–	–	–
	BB-simple	NBB-simple	BB-simple	NBB-simple	BB-simple	NBB-simple
w-markup, avg.	1.223	0.510	0.821	0.058	0.967	0.222
u-markup, avg.	5.850	5.850	0.637	0.058	0.788	0.222
u-markup, min.	0.101	0.101	0.058	0.058	0.223	0.222
u-markup, max.	7.000	7.000	7.000	0.058	7.000	0.222
u-markup, seller 1	7.000	7.000	7.000	0.058	7.000	0.222
w-markup, type 1	0.860	0.102	0.821	0.058	0.967	0.222
w-markup, type 2	6.185	6.085	–	–	–	–

We conduct two comparisons to further illustrate these results. First, we use the estimates from our main model, but reduce the share of type 2 consumers to zero. This allows us to examine how the presence of these consumers affects the equilibrium outcome. The results are shown in Column (3) and Column (4) of Table 6. We observe that both sellers' markups and the level of price dispersion are smaller when type 2 consumers are no longer present in the market. In the NBB-simple treatment, all sellers charge essentially the same price. Next, we use the estimates from Model III (which assumes that there are no type 2 consumers) to simulate pricing equilibria. The Model III parameter estimates are displayed in Appendix A.5. Column (5) and Column (6) of Table 6 show the simulation results. Again, we observe that markups and price dispersion are smaller compared to the results for the main model with both consumer types. Therefore, if we do not take the possibility of price-insensitive consumers into account, we would potentially underestimate the sellers' market power on price comparison websites and obtain biased results regarding their equilibrium pricing strategies.

5.2 Symmetric Ranking Algorithms

We simulate pricing equilibria for the case where the top position in the ranking (and hence the seller who owns the buy box) is determined through a ranking algorithm that puts equal weight on all price dimensions. The higher this weight, the higher is the probability that the best offer is also the top-ranked offer. We consider three different weights $\beta^{[\tau,r]} \in \{5, 10, 15\}$. Table 7 shows the results for $\beta^{[\tau,r]} = 5$ in Columns (1) and (2), for $\beta^{[\tau,r]} = 10$ in Columns (3) and (4), and for $\beta^{[\tau,r]} = 15$ in Columns (5) and (6).

We make two important observations with respect to platform design and market power. First, with symmetric ranking algorithms, the buy box now provides "search guidance" as in Lee and Musolff (2025) and helps consumers. The markups that consumers pay are on average smaller in treatments with buy box compared to treatments without buy box. The effect is particularly pronounced for type 2 consumers in the medium and simple treatments. For these consumers, remaining inattentive is in expectation more beneficial than making an active choice on the price comparison webpage. For type 1 consumers, the absolute differences are smaller, but still positive in all considered settings. Interestingly, in the complex treatments, type 2 consumer pay on average smaller markups than type 1 consumers.

Second, the complexity of price comparisons again tends to increase weighted markups, but the sellers' pricing strategies constitute an important countervailing force. As in the case with fixed ranking, sellers on average increase their prices when price comparisons become more transparent. For some settings, the weighted markups are even larger in the medium treatments than in the complex treatments. Again, we observe that the range of posted prices decreases in complexity. This has a fairly adverse effect on type 2 consumers who on average pay substantially more in the medium and simple treatments than in the complex treatments.

Table 7: Platform Design and Market Power (Symmetric Ranking Algorithms)

	(1)	(2)	(3)	(4)	(5)	(6)
	Model V Estimates $[\beta^{[\tau,r]} = 5]$		Model V Estimates $[\beta^{[\tau,r]} = 10]$		Model V Estimates $[\beta^{[\tau,r]} = 15]$	
	BB-complex	NBB-complex	BB-complex	NBB-complex	BB-complex	NBB-complex
w-markup, avg.	0.495	0.636	0.449	0.593	0.499	0.619
u-markup, avg.	0.485	0.626	0.483	0.578	0.580	0.472
u-markup, min.	0.391	0.570	0.232	0.405	0.187	0.345
u-markup, max.	0.596	0.687	0.676	0.688	0.708	0.710
w-markup, type 1	0.497	0.637	0.457	0.598	0.512	0.631
w-markup, type 2	0.472	0.622	0.337	0.521	0.316	0.460
	BB-medium	NBB-medium	BB-medium	NBB-medium	BB-medium	NBB-medium
w-markup, avg.	0.547	0.596	0.526	0.592	0.510	0.587
u-markup, avg.	4.782	5.003	4.868	5.008	4.929	5.011
u-markup, min.	0.267	0.273	0.235	0.268	0.213	0.264
u-markup, max.	7.000	7.000	7.000	7.000	7.000	7.000
w-markup, type 1	0.324	0.339	0.298	0.335	0.277	0.329
w-markup, type 2	3.591	4.109	3.645	4.110	3.683	4.112
	BB-simple	NBB-simple	BB-simple	NBB-simple	BB-simple	NBB-simple
w-markup, avg.	0.342	0.415	0.344	0.415	0.375	0.414
u-markup, avg.	5.283	5.851	5.342	5.851	5.786	5.850
u-markup, min.	0.090	0.101	0.089	0.101	0.098	0.100
u-markup, max.	7.000	7.000	7.000	7.000	7.000	7.000
w-markup, type 1	0.091	0.103	0.090	0.103	0.100	0.102
w-markup, type 2	3.770	4.678	3.812	4.678	4.129	4.678

We examine the effect of the presence of type 2 by simulating the pricing equilibrium for the Model V estimates without type 2 and for the Model III estimates. The corresponding results are in Appendix A.5. As for the main model, we observe that providing a default through the buy box helps consumers. However, the sellers' reaction to more transparency is missing: Sellers' markups decrease substantially when price comparisons are made less complex. Further, both markups and price dispersion are significantly smaller than under the original Model V estimates – as we already observed in the case with fixed seller ranking.

5.3 Asymmetric Seller Ranking Algorithms

As final step in our analysis, we simulate pricing equilibria for asymmetric ranking algorithms. The algorithm puts a certain weight on the add-on prices and a weight $\beta^{[1,r]}$ on the base price that is smaller or larger than this value. As illustration, we choose $\beta^{[2,r]} = \beta^{[3,r]} = \beta^{[4,r]} = 5$ and $\beta^{[1,r]} \in \{1, 4, 10\}$. Table 8 shows the results for $\beta^{[1,r]} = 1$ in Columns (1) and (2), for $\beta^{[1,r]} = 4$ in Columns (3) and (4), and for $\beta^{[1,r]} = 10$ in Columns (5) and (6). Figure 4 shows the evolution of weighted markups as the weight on the base price changes. For the simple treatments, the asymmetry in the ranking algorithm is irrelevant since there is only a base price (for the sake of completeness, we nevertheless show all values also for these treatments).

With regard to complexity, we observe the same patterns as for the symmetric ranking algorithms: As price comparisons become less complex, weighted markups decrease while price dispersion and unweighted markups increase. Thus, more transparency causes a subgroup of sellers to increase their prices. In the medium and simple treatments, only few sellers charge low markups in order to compete for type 1 consumers, while many sellers charge high prices to serve a few type 2 consumers.

Importantly, we observe that providing a default through the buy box either helps or hurts consumers, depending on weight the algorithm places on the base price relative to the weight on the add-on prices. If $\beta^{[1,r]}$ is small, the algorithm favors sellers with low add-on prices who then can charge relatively high markups to buy box users; in Table 8, we see this from Columns (1) and (2) for the complex and medium treatments. If we increase $\beta^{[1,r]}$ (starting from low levels), weighted markups decrease as sellers charge lower markups to compete for the buy box; see Columns (3) and (4) of Table 8 as well as Figure 4. Interestingly, the association between the weight on the base price $\beta^{[1,r]}$ and weighted markups is non-monotonic: For some sufficiently large values of $\beta^{[1,r]}$ weighted markups are larger than for the symmetric ranking algorithm; to see this, compare Columns (1) and (2) of Table 7 to Columns (5) and (6) of Table 8 for the medium and complex treatments. The intuition behind this is that very high weights on the base price favor sellers with high add-on costs (which type 1 consumers underweight relative to the base price). These sellers can exploit this advantage to some extent by undercutting their rivals. Quantitatively, the effect is small though.

Overall, we observe that asymmetric ranking algorithms with relatively low weights on the base price increase the sellers' market power on the platform. Note that this also holds true if there is no buy box; see Figure 4 for the NBB-complex and NBB-medium treatments. This is due to the fact that type 2 consumers favor top-ranked offers. However, the effect is small compared to the case when there is a buy box.

Table 8: Platform Design and Market Power (Asymmetric Ranking Algorithms)

	(1)	(2)	(3)	(4)	(5)	(6)
	Model V Estimates $[\beta^{[1,r]} = 1]$		Model V Estimates $[\beta^{[1,r]} = 4]$		Model V Estimates $[\beta^{[1,r]} = 10]$	
	BB-complex	NBB-complex	BB-complex	NBB-complex	BB-complex	NBB-complex
w-markup, avg.	1.084	0.775	0.595	0.668	0.559	0.678
u-markup, avg.	0.929	0.863	0.602	0.660	0.630	0.678
u-markup, min.	0.677	0.675	0.463	0.589	0.337	0.590
u-markup, max.	3.086	2.672	0.707	0.718	0.703	0.788
w-markup, type 1	1.015	0.707	0.599	0.671	0.563	0.676
w-markup, type 2	2.029	1.706	0.543	0.635	0.504	0.703
	BB-medium	NBB-medium	BB-medium	NBB-medium	BB-medium	NBB-medium
w-markup, avg.	0.841	0.652	0.540	0.590	0.575	0.596
u-markup, avg.	4.247	4.781	4.121	4.953	4.995	5.010
u-markup, min.	0.274	0.274	0.268	0.274	0.271	0.273
u-markup, max.	7.000	7.000	7.000	7.000	7.000	7.000
w-markup, type 1	0.611	0.339	0.347	0.336	0.340	0.337
w-markup, type 2	3.993	4.922	3.177	4.066	3.774	4.129
	BB-simple	NBB-simple	BB-simple	NBB-simple	BB-simple	NBB-simple
w-markup, avg.	0.346	0.415	0.342	0.415	0.339	0.415
u-markup, avg.	5.273	5.850	5.273	5.851	5.272	5.851
u-markup, min.	0.092	0.101	0.090	0.102	0.087	0.102
u-markup, max.	7.000	7.000	7.000	7.000	7.000	7.000
w-markup, type 1	0.095	0.103	0.091	0.103	0.088	0.103
w-markup, type 2	3.770	4.685	3.763	4.678	3.761	4.678

The presence of type 2 consumers again increases markups and price dispersion, see the results for the Model V estimates without type 2 and for the Model III estimates A.5. The effects are especially pronounced in the medium and simple treatments. We also observe that, if we ignore type 2 consumers, then asymmetric ranking algorithms no longer increase the sellers' market power if there is no buy box.

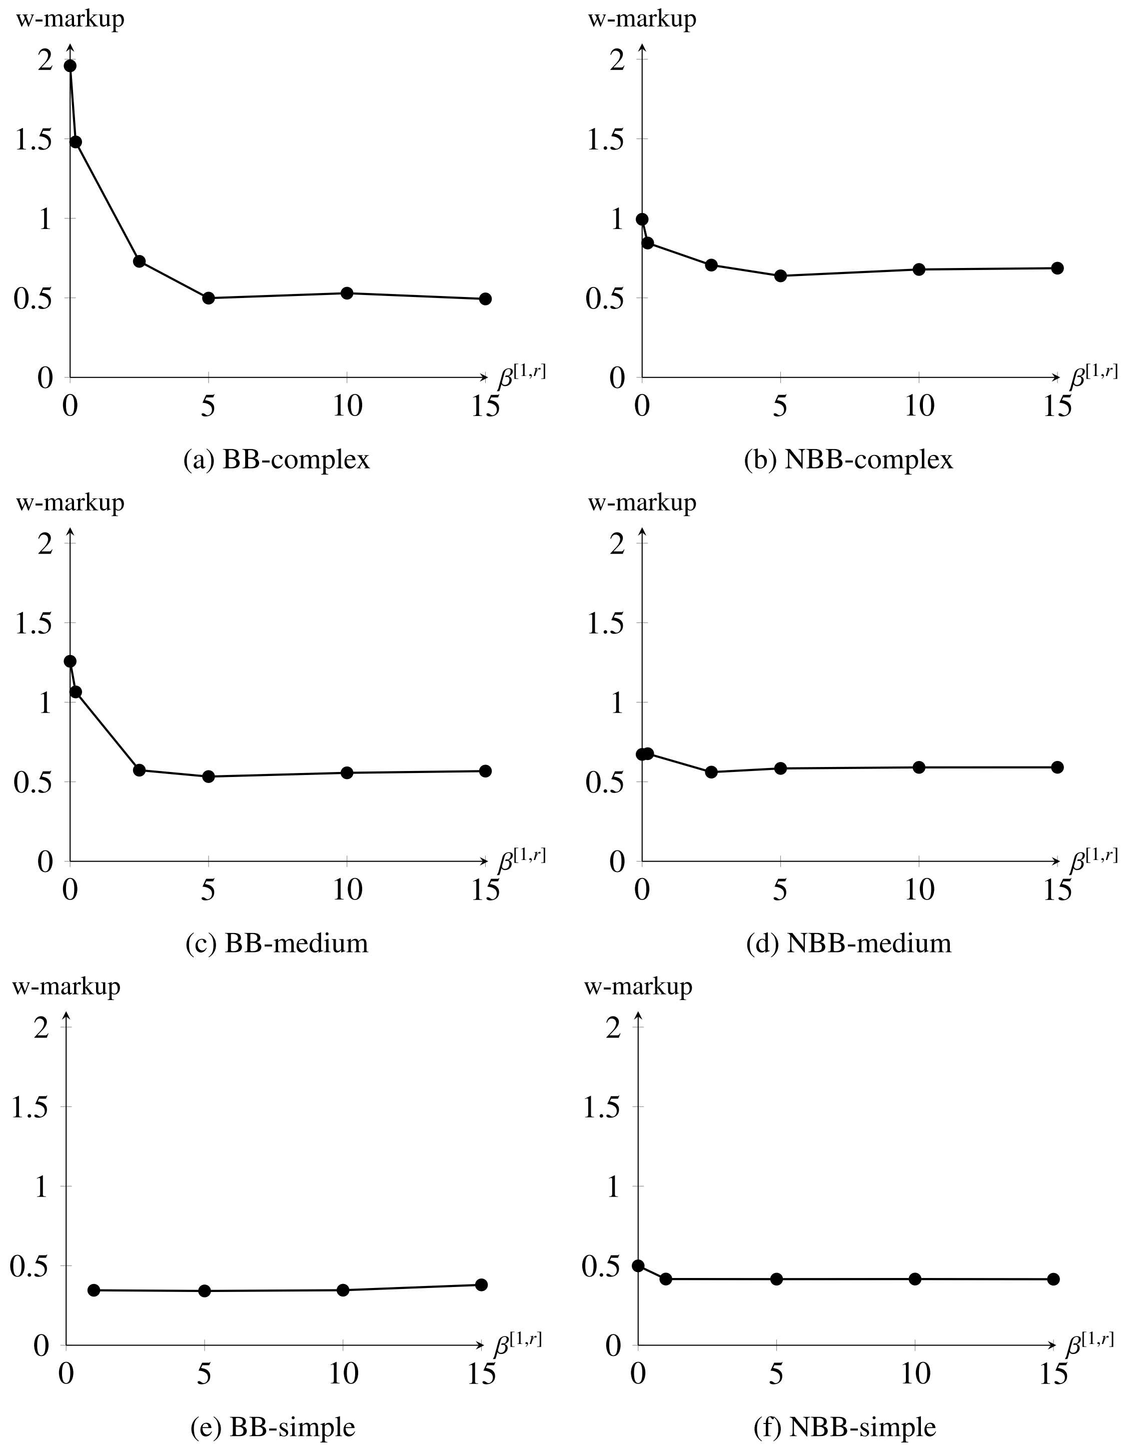

Figure 4: Weighted markups for different values of $\beta^{[1,r]}$

6 Conclusion

In this paper, we examined consumer behavior on a generic online retail platform and what this behavior implies for sellers who post offers on this platform. We observed that consumers exhibit a variety of behaviors that imply price-insensitive choices: Some consumers remain inattentive to the sellers' offers as soon as one option is presented as default to them. Others choose the top-ranked seller even if there are better alternatives. Some consumers ignore or underweight transaction dimensions beyond the base price. A few choices appear to have been made randomly.

To investigate how these behaviors affect the sellers' market power on the platform, we estimated a structural consumer demand model and simulated the pricing equilibrium based on the models' estimates. The model features two consumer types. The price-sensitive type reacts to all transaction dimension (e.g., base price, shipping fee, delivery time, and seller rating), but may attach biased weights to these dimensions. The price-insensitive type either chooses the top-ranked or a randomly chosen seller. With some positive probability, both types remain inattentive and avoid price comparisons.

The estimation results indicate that only a small fraction of consumers are price-insensitive types. Nevertheless, the presence of these consumers has a significant impact on the pricing equilibrium. It increases markups significantly and generates price dispersion among identical sellers, especially when the platform design is such that price comparisons are fairly straightforward. This leads to our first important result on the link between platform design and market power: Making price comparisons more transparent may not result in more competition and reduced markups. If price comparisons are relatively difficult, it pays off for sellers to compete for price-sensitive types who constitute the majority of consumers. The choices of these consumers in such a setting are sufficiently erratic to allow for significant markups even when there are many sellers. However, if price comparisons are relatively simple, the choices of price-sensitive types are smart enough so that only few sellers charge low prices to compete for these consumers while other sellers charge high prices to serve a few price-insensitive consumers. This means that the platform features a few attractive offers (that are chosen by the majority of consumers) and many unattractive offers (which are chosen only by a few customers).

We find that the effect of defaults through a buy box depends on the platform's ranking algorithm that selects the default from the set of offers. If the ranking algorithm places the same weight as consumers on the different transaction dimensions, it helps consumers and reduces markups. It may help both price-sensitive and price-insensitive consumers, especially when price comparisons are complex. However, if it underweights the base price relative to other transaction dimensions, it may favor some sellers (with advantageous outcomes in these transaction dimensions) who then can sell to inattentive consumers at substantial markups. The combination of buy box and ranking algorithms is therefore a further method to influence sellers' market power on a platform.

In order to disentangle consumers' preferences and behavioral biases, we focused on homogeneous products and implemented a money-metric in our experiment so that the transaction dimensions "delivery time" and "seller rating" are represented by risk-free payments. An important next step to better understand the link between platform design and market power would be to allow for differentiated products and uncertainty. This would make it necessary to allow for more general utility functions. Our approach may provide a useful starting-point: We increased the complexity of price comparisons sequentially through treatment variations and used the full dataset to identify different behavioral types. This should also be possible if we allow for differentiated products and uncertain transaction outcomes, though it may require substantial resources to obtain the data necessary to estimate the full model.

References

Allcott, Hunt (2013) "The welfare effects of misperceived product costs: Data and calibrations from the automobile market," *American Economic Journal: Economic Policy*, 5 (3), 30–66. 6

Amazon (2025) "Amazon Prime members enjoyed fastest-ever delivery speeds in 2024, and US members saved on average over $500 on delivery fees," https://www.aboutamazon.com/news/retail/prime-members-us-savings-fastest-delivery-2024, Accessed: 2025-07-30. 12

Armstrong, Mark and John Vickers (2012) "Consumer protection and contingent charges," *Journal of Economic Literature*, 50 (2), 477–493. 4

——— (2022) "Patterns of competitive interaction," *Econometrica*, 90 (1), 153–191. 5

Ba, Sulin and Paul Pavlou (2002) "Evidence of the effect of trust building technology in electronic markets: Price premiums and buyer behavior," *MIS Quarterly*, 26 (3), 243–268. 12

Carlin, Bruce (2009) "Strategic price complexity in retail financial markets," *Journal of Financial Economics*, 91 (3), 278–287. 1, 4

Chen, Daniel, Martin Schonger, and Chris Wickens (2016) "oTree – An open-source platform for laboratory, online, and field experiments," *Journal of Behavioral and Experimental Finance*, 9, 88–97. 17

Chioveanu, Ioana and Jidong Zhou (2013) "Price competition with consumer confusion," *Management Science*, 59 (11), 2450–2469. 1, 4

De Los Santos, Babur and Sergei Koulayev (2017) "Optimizing click-through in online rankings with endogenous search refinement," *Marketing Science*, 36 (4), 542–564. 1

Dinerstein, Michael, Liran Einav, Jonathan Levin, and Neel Sundaresan (2018) "Consumer price search and platform design in internet commerce," *American Economic Review*, 108 (7), 1820–1859. 4

Dohmen, Thomas, Armin Falk, David Huffman, Uwe Sunde, Jürgen Schupp, and Gert Wagner (2011) "Individual risk attitudes: Measurement, determinants, and behavioral consequences," *Journal of the European Economic Association*, 9 (2), 522–550. 17

Ellison, Glenn (2005) "A model of add-on pricing," *Quarterly Journal of Economics*, 120 (2), 585–637. 4

Ellison, Glenn and Sara Ellison (2009) "Search, obfuscation, and price elasticities on the Internet," *Econometrica*, 77 (2), 427–452. 4

Fong, Jessica, Olivia Natan, and Ranmit Pantle (2025) "Consumer inferences from Product rankings: The role of beliefs in search behavior," Working Paper, University of Michigan. 1, 5

Gabaix, Xavier and David Laibson (2006) "Shrouded attributes, consumer myopia, and information suppression in competitive markets," *Quarterly Journal of Economics*, 121 (2), 505–540. 1, 4

Hefti, Andreas, Shuo Liu, and Armin Schmutzler (2022) "Preferences, confusion and competition," *Economic Journal*, 132 (645), 1852–1881. 4

Heidhues, Paul, Johannes Johnen, and Botond Kőszegi (2021) "Browsing versus studying: A pro-market case for regulation," *Review of Economic Studies*, 88 (2), 708–729. 4

Heidhues, Paul, Botond Kőszegi, and Takeshi Murooka (2018) "Inferior products and profitable deception," *Review of Economic Studies*, 84 (1), 323–356. 1, 4

Heiss, Florian, Daniel McFadden, Joachim Winter, Amelie Wuppermann, and Bo Zhou (2021) "Inattention and switching costs as sources of inertia in Medicare Part D," *American Economic Review*, 111 (9), 2737–2781. 1, 3, 6

Hortaçsu, Ali, Seyed Ali Madanizadeh, and Steven Puller (2017) "Power to choose? An analysis of consumer inertia in the residential electricity market," *American Economic Journal: Economic Policy*, 9 (4), 192–226. 1

Inderst, Roman and Martin Obradovits (2023) "Excessive competition on headline prices," *International Economic Review*, 64 (2), 783–808. 4

Janssen, Maarten, José Luis Moraga-González, and Matthijs Wildenbreest (2005) "Truly costly sequential search and oligopolistic pricing," *International Journal of Industrial Organization*, 23, 451–466. 5

Johnen, Johannes (2020) "Dynamic competition in deceptive markets," *RAND Journal of Economics*, 51 (2), 375–401. 1, 4

Karle, Heiko, Florian Kerzenmacher, Heiner Schumacher, and Frank Verboven (2025) "Search costs and context effects," *American Economic Journal: Microeconomics*, 17 (2), 127–161. 5

Lam, H. Tai (2023) "Platform Search Design and Market Power," Working Paper, University of California, Los Angeles, Anderson School of Management. 5

Lee, Kwok Hao and Leon Musolff (2025) "Two-sided markets shaped by platform-guided search," Working Paper, National University of Singapore Business School. 1, 5, 11, 29

Milgrom, Paul and John Roberts (1990) "Rationalizability, learning, and equilibrium in games with strategic complementarities," *Econometrica*, 58 (6), 1255–1277. 9

Myatt, David and David Ronayne (2025) "Stable price dispersion under heterogeneous buyer consideration," *RAND Journal of Economics*, 0 (0), 0. 5

Piccione, Michele and Ran Spiegler (2012) "Price competition under limited comparability," *Quarterly Journal of Economics*, 127 (1), 97–135. 1, 4

Rochet, Jean-Charles and Jean Tirole (2003) "Platform competition in two-sided markets," *Journal of the European Economic Association*, 1 (4), 990–1029. 1

Salz, Tobias and Emanuel Vespa (2020) "Estimating dynamic games of oligopolistic competition: an experimental investigation," *RAND Journal of Economics*, 51 (2), 447–469. 5

Spiegler, Ran (2006) "Competition over agents with boundedly rational expectations," *Theoretical Economics*, 1 (2), 207–231. 4

Stahl, Dale (1989) "Oligopolistic pricing with sequential consumer search," *American Economic Review*, 79 (4), 700–712. 5

Topkis, Donald (1998) *Supermodularity and complementarity*, Princeton, New Jersey: Princeton University Press. 9

Train, Kenneth (2015) "Welfare calculations in discrete choice models when anticipated and experienced attributes differ: A guide with examples," *Journal of Choice Modelling*, 16, 15–22. 6

Ursu, Raluca (2018) "The power of rankings: Quantifying the effect of rankings on online consumer search and purchase decisions," *Marketing Science*, 37 (4), 530–552. 1

Varian, Hal (1980) "A model of sales," *American Economic Review*, 70 (4), 651–659. 5, 27

Vives, Xavier (1990) "Nash equilibrium with strategic complementarities," *Journal of Mathematical Economics*, 19 (3), 305–321. 9

——— (1999) *Oligopoly pricing: Old ideas and new tools*, Cambridge, Massachusetts: MIT Press. 9, 39, 40

Yu, Chuan (2025) "The welfare effects of sponsored product advertising," Working Paper, Harvard Business School. 5

A Online Appendix

A.1 Mathematical Details

Proof of Proposition 1. We prove the uniqueness of the equilibrium in the restricted game-$\tilde{p}$ with given fixed $\tilde{p} \in [c_1, \bar{p}_{\max}]$. To this end, we build on arguments from Vives (1999). Assume first that, for any seller $j \in \{2, ..., M\}$, the best reply $\text{br}_j(\cdot)$ is given by a function. We verify that assumption (A1) then implies that the best reply map $\mathbf{br}(\cdot)$ is a contraction. This is the case if the dominant diagonal condition is satisfied, i.e., we have

$$-\frac{\partial^2 h_j}{\partial (p_j^{[1]})^2} \geq \sum_{j' \in \{2,...,M\} \setminus \{j\}} \left| \frac{\partial^2 h_j}{\partial p_j^{[1]} \partial p_{j'}^{[1]}} \right| \tag{23}$$

for each seller $j \in \{2, ..., M\}$. Recall seller j's objective function

$$h_j = \log(\bar{p}_j - c_j) + \log(\sum_{i=1}^{N} \alpha_i P_{ij}). \tag{24}$$

We obtain the first derivative

$$\frac{\partial h_j}{\partial p_j^{[1]}} = \frac{1}{\bar{p}_j - c_j} - \frac{1}{\sum_{i=1}^N \alpha_i P_{ij}} \left(\sum_{i=1}^{N} \alpha_i \beta_i^{[1,p]} P_{ij} (1 - P_{ij}) \right), \tag{25}$$

and the second derivative

$$\begin{aligned} \frac{\partial^2 h_j}{\partial (p_j^{[1]})^2} = -\frac{1}{(\bar{p}_j - c_j)^2} - \frac{\sum_{i=1}^N \alpha_i \beta_i^{[1,p]} P_{ij} (1 - P_{ij})}{(\sum_{i=1}^N \alpha_i P_{ij})^2} \left(\sum_{i=1}^{N} \alpha_i \beta_i^{[1,p]} P_{ij} (1 - P_{ij}) \right) \\ - \frac{1}{\sum_{i=1}^N \alpha_i P_{ij}} \left(\sum_{i=1}^{N} \alpha_i (\beta_i^{[1,p]})^2 P_{ij} (2P_{ij} - 1)(1 - P_{ij}) \right), \end{aligned} \tag{26}$$

as well as the cross derivative

$$\begin{aligned} \frac{\partial^2 h_j}{\partial p_j^{[1]} \partial p_{j'}^{[1]}} = \frac{\sum_{i=1}^N \alpha_i \beta_i^{[1,p]} P_{ij} P_{ij'}}{(\sum_{i=1}^N \alpha_i P_{ij})^2} \left(\sum_{i=1}^{N} \alpha_i \beta_i^{[1,p]} P_{ij} (1 - P_{ij}) \right) \\ + \frac{1}{\sum_{i=1}^N \alpha_i P_{ij}} \left(\sum_{i=1}^{N} \alpha_i (\beta_i^{[1,p]})^2 P_{ij} (2P_{ij} - 1) P_{ij'} \right) \end{aligned} \tag{27}$$

for $j' \in \{2, ..., M\} \setminus \{j\}$. By assumption (A1) the values in the bracket on the right-hand side of the inequality in (23) are weakly positive. We therefore can write

$$\sum_{j' \in \{2,...,M\} \setminus \{j\}} \left| \frac{\partial^2 h_j}{\partial p_j^{[1]} \partial p_{j'}^{[1]}} \right| = \frac{\sum_{i=1}^N \alpha_i \beta_i^{[1,p]} P_{ij}(1 - P_{ij})}{(\sum_{i=1}^N \alpha_i P_{ij})^2} \left(\sum_{i=1}^N \alpha_i \beta_i^{[1,p]} P_{ij}(1 - P_{ij}) \right) + \frac{1}{\sum_{i=1}^N \alpha_i P_{ij}} \left(\sum_{i=1}^N \alpha_i (\beta_i^{[1,p]})^2 P_{ij}(2P_{ij} - 1)(1 - P_{ij}) \right). \tag{28}$$

This allows us to write

$$-\frac{\partial^2 h_j}{\partial (p_j^{[1]})^2} = \frac{1}{(\bar{p}_j - c_j)^2} + \sum_{j' \in \{2,...,M\} \setminus \{j\}} \left| \frac{\partial^2 h_j}{\partial p_j^{[1]} \partial p_{j'}^{[1]}} \right|, \tag{29}$$

which implies equation (23). The dominant diagonal condition further implies that we have

$$\frac{\partial^2 h_j}{\partial (p_j^{[1]})^2} < 0. \tag{30}$$

Hence, the objective function of seller $j \in \{2, ..., M\}$ is concave (and thus quasi-concave) in its own total price. Seller j's best reply function is thus well-defined and the initial assumption on $\text{br}_j(\cdot)$ correct. It follows that the game has a unique equilibrium and that this equilibrium can be approximated by iterating the best reply map (Vives, 1999). □

Ranking Algorithms – Formal Details. We characterize the equilibrium for alternative ranking algorithms. Let the price matrix $\boldsymbol{p} = (\boldsymbol{p}_1, ..., \boldsymbol{p}_M)$ be given. We find for given ranking $\boldsymbol{r}$ the choice and attention probabilities. The probability with which consumer i selects seller j's offer on the price comparison webpage now depends on the ranking $\boldsymbol{r}$ and is given by

$$P_{ij}(\boldsymbol{r}) = \frac{\exp(-\boldsymbol{p}_j \boldsymbol{\beta}_i^{[p]} + \boldsymbol{d}_j(\boldsymbol{r}) \boldsymbol{\beta}_i^{[d]})}{\sum_{j'=1}^M \exp(-\boldsymbol{p}_{j'} \boldsymbol{\beta}_i^{[p]} + \boldsymbol{d}_{j'}(\boldsymbol{r}) \boldsymbol{\beta}_i^{[d]})}, \tag{31}$$

where $\boldsymbol{d}_j(\boldsymbol{r}) = (1, 0)$ if $\boldsymbol{r}$ is such that seller j is top-ranked and $\boldsymbol{d}_j(\boldsymbol{r}) = (0, 1)$ otherwise. Let j_r be the inversion of r_j, i.e., the seller j that has rank r. Accordingly, j_1 is the top-ranked seller. With this, we obtain the attention probability of consumer i as

$$\alpha_i(\boldsymbol{r}) = \frac{1}{1 + \exp(-\boldsymbol{p}_{j_1} \boldsymbol{\beta}_i^{[p,o]} + \boldsymbol{d} \boldsymbol{\beta}_i^{[d,o]} - q_i)}. \tag{32}$$

Using the probability of being the top-ranked seller $P_j^{[r]}$ defined in equation (12) we obtain the expected log-profit of seller j as

$$\mathbb{E}(h_j \mid \boldsymbol{p}) = \log(\bar{p}_1 - c_1) + \log\Bigg(P_j^{[r]} \sum_{i=1}^{N} [1 - \alpha_i(\boldsymbol{r}^{[j]}) + \alpha_i(\boldsymbol{r}^{[j]}) P_{ij}(\boldsymbol{r}^{[j]})] + \sum_{j'=1, j' \neq j}^{M} [P_{j'}^{[r]} \sum_{i=1}^{N} \alpha_i(\boldsymbol{r}^{[j']}) P_{ij}(\boldsymbol{r}^{[j']})]\Bigg). \tag{33}$$

and the corresponding first-order condition

$$\begin{aligned}\frac{\partial \mathbb{E}(h_j \mid \boldsymbol{p})}{\partial p_j^{[1]}} = \frac{1}{\bar{p}_j - c_j} - \frac{1}{\Delta}\Bigg(&\beta^{[1,r]} P_j^{[r]}(1 - P_j^{[r]}) \sum_{i=1}^{N} [1 - \alpha_i(\boldsymbol{r}^{[j]})(1 - P_{ij}(\boldsymbol{r}^{[j]}))] \\ &+ P_j^{[r]} \sum_{i=1}^{N} \alpha_i(\boldsymbol{r}^{[j]})(1 - P_{ij}(\boldsymbol{r}^{[j]}))[\beta_i^{[1,p,o]}(1 - \alpha_i(\boldsymbol{r}^{[j]})) + \beta_i^{[1,p]} P_{ij}(\boldsymbol{r}^{[j]})] \\ &- \sum_{j'=1, j' \neq j}^{M} [\beta^{[1,r]} P_j^{[r]} P_{j'}^{[r]} \sum_{i=1}^{N} (\alpha_i(\boldsymbol{r}^{[j']}) P_{ij}(\boldsymbol{r}^{[j']}))] \\ &+ \sum_{j'=1, j' \neq j}^{M} [P_{j'}^{[r]} \sum_{i=1}^{N} \alpha_i(\boldsymbol{r}^{[j']}) \beta_i^{[1,p]} P_{ij}(\boldsymbol{r}^{[j']})(1 - P_{ij}(\boldsymbol{r}^{[j']}))]\Bigg) = 0. \end{aligned} \tag{34}$$

where we have

$$\Delta = P_j^{[r]} \sum_{i=1}^{N} [1 - \alpha_i(\boldsymbol{r}^{[j]})(1 - P_{ij}(\boldsymbol{r}^{[j]}))] + \sum_{j'=1, j' \neq j}^{M} [P_{j'}^{[r]} \sum_{i=1}^{N} \alpha_i(\boldsymbol{r}^{[j']}) P_{ij}(\boldsymbol{r}^{[j']})]. \tag{35}$$

With this condition, we can then again compute the equilibrium for each platform design. □

A.2 Instructions

This appendix shows the instructions for the main treatments.

A.2.1 Instructions BB-simple treatment

Instructions BB-simple Page 1/3

In the next part of the study, your task is to purchase a book from an online retail website. This website resembles platforms like *amazon.com* or *pricegrabber.com*. We are interested in peoples' purchase behavior on such websites.

The book you are asked to purchase is *Midnight's Witness*. You will not actually buy this book (it does not really exist), but your bonus increases in the price savings that you realize. Your budget will be 15 USD. If you purchase *Midnight's Witness* at a total price of, say 13 USD, then your bonus payment will be $15 - 13 = 2$ USD.

The earnings from this part of the study will be paid as a bonus on Prolific (in addition to your participation fee of 2 USD).

Instructions BB-simple Page 2/3

On the retail website, the total price of a book consists of its base price.

Here is an example (for a book other than *Midnight's Witness* that is more expensive):

The total price of this book at the shop "FictionFox" is 24.87 USD.

Instructions BB-simple Page 3/3

Once you enter the retail website, you will first see an overview of products. You have to purchase the book *Midnight's Witness*.

On the overview page, you can directly purchase *Midnight's Witness* from a pre-selected shop.

Alternatively, you can click on the button "All buying options" to see a list of 12 shops that all offer the book at varying prices. The pre-selected shop from the overview page appears at the top of this list. You may be able to find a better deal in this list.

After a brief comprehension check, the shopping task starts on the next page.

A.2.2 Instructions BB-medium treatment

Instructions BB-medium Page 1/3

In the next part of the study, your task is to purchase a book from an online retail website. This website resembles platforms like *amazon.com* or *pricegrabber.com*. We are interested in peoples' purchase behavior on such websites.

The book you are asked to purchase is *Midnight's Witness*. You will not actually buy this book (it does not really exist), but your bonus increases in the price savings that you realize. Your budget will be 15 USD. If you purchase *Midnight's Witness* at a total price of, say 13 USD, then your bonus payment will be 15 − 13 = 2 USD.

The earnings from this part of the study will be paid as a bonus on Prolific (in addition to your participation fee of 2 USD).

Instructions BB-medium Page 2/3

On the retail website, the total price of a book consists of two partial prices:

- base price,
- shipping fee.

The total price of a product is the sum of these partial prices. Here is an example (for a book other than *Midnight's Witness* that is more expensive):

The total price of this book at the shop "FictionFox" is 20.79 USD (base price) + 4.08 USD (shipping) = 24.87 USD.

Instructions BB-medium Page 3/3

Once you enter the retail website, you will first see an overview of products. You have to purchase the book *Midnight's Witness*.

On the overview page, you can directly purchase *Midnight's Witness* from a pre-selected shop (the one with the lowest shipping fee).

Alternatively, you can click on the button "All buying options" to see a list of 12 shops that all offer the book at varying prices. The pre-selected shop from the overview page appears at the top of this list. You may be able to find a better deal in this list.

After a brief comprehension check, the shopping task starts on the next page.

A.2.3 Instructions BB-complex treatment

Instructions BB-complex Page 1/3

In the next part of the study, your task is to purchase a book from an online retail website. This website resembles platforms like *amazon.com* or *pricegrabber.com*. We are interested in peoples' purchase behavior on such websites.

The book you are asked to purchase is *Midnight's Witness*. You will not actually buy this book (it does not really exist), but your bonus increases in the price savings that you realize. Your budget will be 15 USD. If you purchase *Midnight's Witness* at a total price of, say 13 USD, then your bonus payment will be 15 – 13 = 2 USD.

The earnings from this part of the study will be paid as a bonus on Prolific (in addition to your participation fee of 2 USD).

Instructions BB-complex Page 2/3

On the retail website, the total price of a book consists of four partial prices:

- base price,
- shipping fee,
- delivery time costs (each day more costs 0.42 USD),
- shop rating costs (each star below a 5-star rating costs 0.93 USD).

The total price of a product is the sum of these partial prices. Here is an example (for a book other than *Midnight's Witness* that is more expensive):

The total price of this book at the shop "FictionFox" is 20.79 USD (base price) + 0.96 USD (shipping) + 3 × 0.42 USD (three days delivery) + 2 × 0.93 USD (two stars missing from 5-star rating) = 24.87 USD.

Instructions BB-complex Page 3/3

Once you enter the retail website, you will first see an overview of products. You have to purchase the book *Midnight's Witness*.

On the overview page, you can directly purchase *Midnight's Witness* from a pre-selected shop (the one with the lowest sum of shipping fee, delivery time, and shop rating costs).

Alternatively, you can click on the button "All buying options" to see a list of 12 shops that all offer the book at varying prices. The pre-selected shop from the overview page appears at the top of this list. You may be able to find a better deal in this list.

After a brief comprehension check, the shopping task starts on the next page.

A.2.4 Instructions NBB-simple, NBB-medium, and NBB-complex treatment

The instructions to the NBB-simple, NBB-medium, and NBB-complex treatment are identical to those of the corresponding BB treatments, except on the last page. The text on this page is as follows.

Instructions NBB-simple, NBB-medium, and NBB-complex Page 3/3

Once you enter the retail website, you will first see an overview of products. You have to purchase the book *Midnight's Witness*.

To do so, click on the button "All buying options" to see a list of 12 shops that all offer the book at varying prices.

After a brief comprehension check, the shopping task starts on the next page.

A.3 Additional Treatments: Parameters and Results

We run two alternative versions of the six main treatments. In these sets of alternative treatments (labeled A1 and A2), we vary the sellers' prices. The price vectors and the descriptive statistics for subjects' shopping behavior in these treatments are shown in the tables below.

Table A1: Price Vectors and Total Prices in Alternative Treatments

Seller	(1) base price $p_j^{[1]}$	(2) shipping fee $p_j^{[2]}$	(3) delivery t. costs $p_j^{[3]}$	(4) rating costs $p_j^{[4]}$	(5) total price $\bar{p}_j$
			BB-complex A1		
1	13.17	0.16	0.42	0.00	13.75
2	11.32	0.96	0.42	1.86	14.56
3	10.05	1.96	0.84	0.93	13.78
4	9.11	0.16	0.84	2.79	12.90
5	8.35	0.16	1.26	2.79	12.56
6	7.98	2.96	0.42	0.93	12.29
7	7.11	0.96	0.84	2.79	11.70
8	6.84	2.96	0.00	1.86	11.66
9	5.85	1.96	0.42	2.79	11.02
10	7.42	2.96	0.84	1.86	13.08
11	6.41	2.96	1.26	1.86	12.49
12	5.24	2.96	1.26	2.79	12.25
			BB-complex A2		
1	12.80	0.16	0.42	0.00	13.38
2	11.10	0.96	0.84	0.00	12.90
3	10.20	0.96	0.84	0.00	12.00
4	8.40	0.96	0.84	0.93	11.13
5	11.45	0.16	0.84	1.86	14.31
6	10.90	0.96	0.00	2.79	14.65
7	9.92	2.96	0.84	0.00	13.72
8	8.72	0.16	0.00	3.72	12.60
9	8.11	1.96	0.42	1.86	12.35
10	7.42	1.96	1.26	1.86	12.50
11	8.23	1.96	0.84	2.79	13.82
12	7.82	2.96	0.42	3.72	14.92

Table A2: Descriptive Statistics Shopping Behavior (A1 and A2 Treatments)

	(1) BB-complex A1	(2) BB-medium A1	(3) BB-simple A1	(4) NBB-complex A1	(5) NBB-medium A1	(6) NBB-simple A1
Share top-ranked seller (through buy box)	0.122	0.086	0.156	–	–	–
Share top-ranked seller (price comparison)	0.064	0.007	0.000	0.081	0.058	0.015
Share lowest base price (price comparison)	0.109	0.053	[0.000]	0.097	0.029	[0.007]
Share lowest total price (price comparison)	0.378	0.803	0.823	0.476	0.813	0.941
Overpayment (SD)	1.04 (1.04)	0.37 (0.84)	0.46 (1.01)	0.76 (0.88)	0.30 (0.72)	0.09 (0.42)
N	156	152	147	124	139	135
	BB-complex A2	BB-medium A2	BB-simple A2	NBB-complex A2	NBB-medium A2	NBB-simple A2
Share top-ranked seller (through buy box)	0.190	0.113	0.080	–	–	–
Share top-ranked seller (price comparison)	0.042	0.014	0.006	0.066	0.013	0.008
Share lowest base price (price comparison)	0.106	0.021	[0.000]	0.086	0.065	[0.000]
Share lowest total price (price comparison)	0.444	0.801	0.896	0.572	0.779	0.939
Overpayment (SD)	0.96 (0.94)	0.39 (0.84)	0.22 (0.68)	0.69 (0.91)	0.31 (0.69)	0.08 (0.35)
N	142	141	163	152	154	132

Notes: The shares in squared brackets refer to the seller who has the lowest base price in the complex and medium treatments (but not in the simple treatments where the base price equals the total price).

A.4 Balancing Tables

Table A3: Balancing Table Demographic Variables and Personal Characteristics I

Treatment	(1) N	(2) Share Females	(3) Age	(4) Education Score	(5) CRT Score
BB-complex	274	0.500	42.6 (13.3)	2.8 (0.7)	1.9 (1.1)
BB-medium	279	0.516	42.9 (13.4)	2.8 (0.7)	2.0 (1.1)
BB-simple	282	0.582	43.1 (12.7)	2.9 (0.8)	1.8 (1.1)
NBB-complex	273	0.553	41.3 (13.5)	2.8 (0.7)	1.8 (1.1)
NBB-medium	287	0.477	43.4 (13.3)	2.9 (0.7)	2.1 (1.0)
NBB-simple	306	0.516	42.9 (13.7)	2.8 (0.7)	2.0 (1.1)
BB-complex-A1	156	0.500	43.1 (13.3)	2.9 (0.7)	2.0 (1.0)
BB-medium-A1	152	0.500	43.2 (13.7)	2.9 (0.7)	1.9 (1.1)
BB-simple-A1	147	0.599	43.4 (12.4)	2.8 (0.8)	1.9 (1.1)
NBB-complex-A1	124	0.581	41.0 (13.1)	2.9 (0.8)	1.9 (1.1)
NBB-medium-A1	139	0.489	42.7 (12.6)	2.8 (0.7)	1.9 (1.1)
NBB-simple-A1	135	0.556	43.6 (12.2)	2.9 (0.7)	2.0 (1.1)
BB-complex-A2	142	0.493	43.7 (12.8)	2.9 (0.8)	2.0 (1.1)
BB-medium-A2	141	0.518	42.3 (14.4)	2.8 (0.8)	1.9 (1.0)
BB-simple-A2	163	0.472	40.6 (12.3)	2.9 (0.8)	2.0 (1.1)
NBB-complex-A2	152	0.487	43.9 (13.2)	2.8 (0.8)	1.9 (1.1)
NBB-medium-A2	154	0.519	41.5 (12.9)	2.8 (0.7)	1.9 (1.1)
NBB-simple-A2	132	0.492	43.7 (13.6)	2.8 (0.8)	2.0 (1.1)
ANOVA p-value		0.377	0.518	0.886	0.393

Notes: Standard deviation in parentheses. Education is indicated on a scale of 0 to 4 (0 = No degree, 1 = Some high school, 2 = High school degree, 3 = Bachelor's degree, 4 = Master's degree or higher) The CRT is indicated on a scale from 0 to 3.

Table A4: Balancing Table Demographic Variables and Personal Characteristics II

Treatment	(1) N	(2) Risk Tolerance	(3) Patience	(4) Trust	(5) Shopping Frequency
BB-complex	274	4.9 (2.6)	7.0 (2.0)	4.9 (2.3)	3.4 (1.2)
BB-medium	279	4.9 (2.5)	6.9 (2.1)	4.7 (2.4)	3.4 (1.1)
BB-simple	282	4.9 (2.5)	7.0 (2.2)	4.8 (2.3)	3.4 (1.2)
NBB-complex	273	5.1 (2.6)	7.1 (2.1)	4.6 (2.3)	3.3 (1.2)
NBB-medium	287	5.0 (2.6)	7.0 (2.1)	5.0 (2.3)	3.3 (1.2)
NBB-simple	306	4.8 (2.4)	6.9 (2.0)	4.8 (2.2)	3.3 (1.2)
BB-complex-A1	156	5.3 (2.6)	7.4 (2.1)	4.8 (2.5)	3.3 (1.2)
BB-medium-A1	152	4.8 (2.4)	6.9 (2.1)	4.8 (2.4)	3.5 (1.3)
BB-simple-A1	147	5.1 (2.7)	7.1 (2.1)	4.8 (2.4)	3.5 (1.1)
NBB-complex-A1	124	5.1 (2.6)	7.1 (2.0)	4.8 (2.6)	3.5 (1.2)
NBB-medium-A1	139	4.9 (2.7)	7.0 (2.2)	4.6 (2.6)	3.4 (1.3)
NBB-simple-A1	135	4.8 (2.7)	6.9 (2.3)	4.8 (2.4)	3.5 (1.2)
BB-complex-A2	142	5.1 (2.4)	7.3 (1.8)	5.2 (2.3)	3.3 (1.1)
BB-medium-A2	141	4.7 (2.5)	7.2 (2.1)	4.8 (2.4)	3.2 (1.3)
BB-simple-A2	163	5.2 (2.6)	6.8 (2.2)	4.4 (2.4)	3.3 (1.2)
NBB-complex-A2	152	4.8 (2.6)	7.2 (1.8)	4.9 (2.4)	3.5 (1.2)
NBB-medium-A2	154	5.0 (2.7)	6.9 (2.2)	4.4 (2.4)	3.2 (1.2)
NBB-simple-A2	132	4.8 (2.7)	7.2 (2.0)	5.0 (2.5)	3.6 (1.1)
ANOVA p-value		0.781	0.632	0.218	0.283

Notes: Standard deviation in parentheses. Risk Tolerance, Patience, and Trust take on values on a scale between 0 and 10. Shopping Frequency is indicated on a scale of 0 to 5 (0 = Never, 1 = Once in a few months or longer, 2 = About once a month, 3 = Several times a month, 4 = About once per week, 5 = More than once a week).

A.5 Additional Tables

Table A5: Estimation Results Model III

	(1) Coefficient	(2) SE	(3) Share
Attention Stage			
Constant	1.652***	(0.114)	
Simple treatment	0.441*	(0.174)	89.0%
Medium treatment	0.630***	(0.184)	90.7%
Complex treatment			83.9%
Choice Stage			
Base price, simple	4.917***	(0.079)	
Base price, medium	2.888***	(0.051)	
Shipping, medium	2.609***	(0.057)	
Base price, complex	1.240***	(0.041)	
Shipping, complex	1.202***	(0.060)	
Delivery time, complex	0.972***	(0.094)	
Seller rating, complex	1.449***	(0.052)	

Notes: Results from maximum likelihood estimation. Standard errors are in brackets. The last column shows the estimated probability with which subjects enter the price comparison webpage (attention stage). Significance at * $p < 0.05$, ** $p < 0.01$, and *** $p < 0.001$.

Table A6: Symmetric Ranking Algorithms, Model V Estimates without Type 2

	(1)	(2)	(3)	(4)	(5)	(6)
	Model V Estimates [Pr(k = 2) = 0] [$\beta^{[\tau,r]}$ = 5]		Model V Estimates [Pr(k = 2) = 0] [$\beta^{[\tau,r]}$ = 10]		Model V Estimates [Pr(k = 2) = 0] [$\beta^{[\tau,r]}$ = 15]	
	BB-complex	NBB-complex	BB-complex	NBB-complex	BB-complex	NBB-complex
w-markup, avg.	0.513	0.674	0.462	0.674	0.512	0.674
u-markup, avg.	0.501	0.667	0.479	0.667	0.570	0.667
u-markup, min.	0.407	0.639	0.247	0.639	0.202	0.639
u-markup, max.	0.605	0.709	0.661	0.709	0.694	0.709
w-markup, type 1	0.513	0.674	0.462	0.674	0.512	0.674
w-markup, type 2	-	-	-	-	-	-
	BB-medium	NBB-medium	BB-medium	NBB-medium	BB-medium	NBB-medium
w-markup, avg.	0.274	0.280	0.253	0.280	0.244	0.280
u-markup, avg.	0.236	0.236	0.196	0.236	0.179	0.236
u-markup, min.	0.217	0.214	0.128	0.214	0.095	0.214
u-markup, max.	0.326	0.328	0.318	0.328	0.318	0.328
w-markup, type 1	0.274	0.280	0.253	0.280	0.244	0.280
w-markup, type 2	-	-	-	-	-	-
	BB-simple	NBB-simple	BB-simple	NBB-simple	BB-simple	NBB-simple
w-markup, avg.	0.063	0.058	0.061	0.058	0.059	0.058
u-markup, avg.	0.063	0.058	0.061	0.058	0.059	0.058
u-markup, min.	0.063	0.058	0.061	0.058	0.059	0.058
u-markup, max.	0.063	0.058	0.061	0.058	0.059	0.058
w-markup, type 1	0.063	0.058	0.061	0.058	0.059	0.058
w-markup, type 2	-	-	-	-	-	-

Table A7: Symmetric Ranking Algorithms, Model III Estimates

	(1)	(2)	(3)	(4)	(5)	(6)
	Model III Estimates $[\beta^{[\tau,r]} = 5]$		Model III Estimates $[\beta^{[\tau,r]} = 10]$		Model III Estimates $[\beta^{[\tau,r]} = 15]$	
	BB-complex	NBB-complex	BB-complex	NBB-complex	BB-complex	NBB-complex
w-markup, avg.	0.612	0.886	0.636	0.886	0.663	0.886
u-markup, avg.	0.605	0.880	0.716	0.880	0.769	0.880
u-markup, min.	0.466	0.851	0.291	0.851	0.238	0.851
u-markup, max.	0.755	0.920	0.889	0.920	0.906	0.920
w-markup, type 1	0.612	0.886	0.636	0.886	0.663	0.886
w-markup, type 2	-	-	-	-	-	-
	BB-medium	NBB-medium	BB-medium	NBB-medium	BB-medium	NBB-medium
w-markup, avg.	0.359	0.384	0.317	0.384	0.301	0.384
u-markup, avg.	0.352	0.378	0.306	0.378	0.290	0.378
u-markup, min.	0.319	0.359	0.236	0.359	0.180	0.359
u-markup, max.	0.391	0.409	0.373	0.409	0.380	0.409
w-markup, type 1	0.359	0.384	0.317	0.384	0.301	0.384
w-markup, type 2	-	-	-	-	-	-
	BB-simple	NBB-simple	BB-simple	NBB-simple	BB-simple	NBB-simple
w-markup, avg.	0.221	0.222	0.199	0.222	0.180	0.222
u-markup, avg.	0.221	0.222	0.199	0.222	0.180	0.222
u-markup, min.	0.221	0.221	0.199	0.221	0.178	0.221
u-markup, max.	0.221	0.222	0.199	0.222	0.181	0.222
w-markup, type 1	0.221	0.222	0.199	0.222	0.180	0.222
w-markup, type 2	-	-	-	-	-	-

Table A8: Asymmetric Ranking Algorithms, Model V Estimates without Type 2

	(1)	(2)	(3)	(4)	(5)	(6)
	Model V Estimates $\Pr(k=2)=0$ $[\beta^{[1,r]}=1]$		Model V Estimates $\Pr(k=2)=0$ $[\beta^{[1,r]}=4]$		Model V Estimates $\Pr(k=2)=0$ $[\beta^{[1,r]}=10]$	
	BB-complex	NBB-complex	BB-complex	NBB-complex	BB-complex	NBB-complex
w-markup, avg.	0.987	0.674	0.598	0.674	0.566	0.674
u-markup, avg.	0.899	0.667	0.598	0.667	0.615	0.667
u-markup, min.	0.650	0.639	0.473	0.639	0.365	0.639
u-markup, max.	3.006	0.709	0.696	0.709	0.697	0.709
w-markup, type 1	0.987	0.674	0.598	0.674	0.566	0.674
w-markup, type 2	-	-	-	-	-	-
	BB-medium	NBB-medium	BB-medium	NBB-medium	BB-medium	NBB-medium
w-markup, avg.	0.509	0.280	0.285	0.280	0.288	0.280
u-markup, avg.	0.632	0.236	0.265	0.236	0.237	0.236
u-markup, min.	0.221	0.214	0.225	0.214	0.214	0.214
u-markup, max.	3.259	0.328	0.386	0.328	0.354	0.328
w-markup, type 1	0.509	0.280	0.285	0.280	0.288	0.280
w-markup, type 2	-	-	-	-	-	-
	BB-simple	NBB-simple	BB-simple	NBB-simple	BB-simple	NBB-simple
w-markup, avg.	0.074	0.058	0.063	0.058	0.061	0.058
u-markup, avg.	0.239	0.058	0.063	0.058	0.061	0.058
u-markup, min.	0.066	0.058	0.063	0.058	0.061	0.058
u-markup, max.	1.037	0.058	0.064	0.058	0.061	0.058
w-markup, type 1	0.074	0.058	0.063	0.058	0.061	0.058
w-markup, type 2	-	-	-	-	-	-

Table A9: Asymmetric Ranking Algorithms, Model III Estimates

	(1)	(2)	(3)	(4)	(5)	(6)
	Model III Estimates $[\beta^{[1,r]} = 1]$		Model III Estimates $[\beta^{[1,r]} = 4]$		Model III Estimates $[\beta^{[1,r]} = 10]$	
	BB-complex	NBB-complex	BB-complex	NBB-complex	BB-complex	NBB-complex
w-markup, avg.	1.180	0.886	0.734	0.886	0.708	0.886
u-markup, avg.	1.070	0.880	0.761	0.880	0.790	0.880
u-markup, min.	0.856	0.851	0.596	0.851	0.332	0.851
u-markup, max.	2.944	0.920	0.867	0.920	0.909	0.920
w-markup, type 1	1.180	0.886	0.734	0.886	0.708	0.886
w-markup, type 2	-	-	-	-	-	-
	BB-medium	NBB-medium	BB-medium	NBB-medium	BB-medium	NBB-medium
w-markup, avg.	0.574	0.384	0.380	0.384	0.390	0.384
u-markup, avg.	0.634	0.378	0.375	0.378	0.381	0.378
u-markup, min.	0.373	0.359	0.347	0.359	0.337	0.359
u-markup, max.	2.936	0.409	0.423	0.409	0.531	0.409
w-markup, type 1	0.574	0.384	0.380	0.384	0.390	0.384
w-markup, type 2	-	-	-	-	-	-
	BB-simple	NBB-simple	BB-simple	NBB-simple	BB-simple	NBB-simple
w-markup, avg.	0.243	0.222	0.226	0.222	0.199	0.222
u-markup, avg.	0.243	0.222	0.226	0.222	0.199	0.222
u-markup, min.	0.243	0.221	0.226	0.221	0.199	0.221
u-markup, max.	0.243	0.222	0.226	0.222	0.199	0.222
w-markup, type 1	0.243	0.222	0.226	0.222	0.199	0.222
w-markup, type 2	-	-	-	-	-	-