Partitioned Pricing and Consumer Welfare

Kevin Ducbao Tran[1]

University of Bristol

kevin.tran@bristol.ac.uk

September 4, 2024

Abstract

Partitioned pricing is a price obfuscation strategy in which the price is split into a base price and add-on fees. While empirical evidence suggests that partitioned pricing affects consumer decisions through salience effects, its consumer welfare consequences are largely unexplored. I provide a quantification of the welfare impact of partitioned pricing on eBay Germany. I employ a discrete choice model that jointly allows for behavioral reactions to marginal changes in the add-on fees as well as a discontinuous effect of a zero fee. I estimate the model parameters using web scraped data of posted price transactions on eBay Germany. The results suggest some under-reaction to marginal changes in the shipping fee and a discontinuous positive effect of free shipping on consumer demand. The combined impact of these effects on consumer welfare is not larger than six percent of consumer surplus.

Keywords: partitioned pricing, limited attention, consumer welfare, shipping fees, eBay

JEL Codes: D12, D60, D83, L11

[1] I thank Tomaso Duso and Hannes Ullrich for all of their invaluable support and advice. I am also grateful to comments by Luis Abreu, Vicente Lagos, Imke Reimers, Maximilian Schäfer, and Renke Schmacker. This paper has also benefited from discussions with Aaron Bodoh-Creed, Stefano DellaVigna, Ruben Durante, Paul Heidhues, Christian Peukert, Steve Tadelis, and participants of the data workshop at the University of Zürich, the 2nd Doctoral Workshop on the Economics of Digitization at Télécom ParisTech, the 23rd Spring Meeting of Young Economists at the University of the Balearic Islands, the 20th Summer Workshop for Young Economists at the ZEW Mannheim, the IO Brown Bag Seminar at DIW Berlin, the Annual Meeting of the Verein für Socialpolitik in Freiburg, the Berlin Behavioral Economics (BBE) Workshop, the 17th Annual International Industrial Organization Conference (IIOC), and the 2022 MaCCI Annual Conference. I thank Adam Lederer for editorial support. All remaining errors are my own.

1 Introduction

The separation of the price for a good or service into a base part and one or more smaller parts is called partitioned pricing (Morwitz et al., 1998). With the rise of online retail, this practice has become ever more prominent (Greenleaf et al., 2016). Fully rational consumers should only consider the total price of a good and not its division into smaller parts. Empirical evidence suggests, however, that consumers do react to partitioned pricing by not fully considering the add-on fees (Greenleaf et al. (2016) and Voester et al. (2017) provide comprehensive reviews on the topic). While this result is documented for auctions in the online shopping context (Hossain and Morgan, 2006; Brown et al., 2010; Einav et al., 2015), there is less evidence on whether consumers exhibit similar behavior when purchasing at posted prices. However, posted price transactions make up the majority of transactions nowadays, even on eBay, an online platform that at its inception only featured auctions (Einav et al., 2018). Furthermore, the consumer welfare implications of such behavioral reactions to partitioned pricing remain unexplored. Understanding such behavioral patterns and their impact on consumer welfare is relevant both for online platforms designing their marketplaces as well as consumer protection agencies considering policies to protect consumers from potential harm.

My paper provides an analysis of the consumer welfare consequences of partitioned pricing. Another novelty is that I consider a discontinuous reaction to a zero fee in addition to an under- or over-reaction to marginal changes in the fee as compared to the product price. Furthermore, I focus on posted price transactions rather than auctions. To correctly estimate consumer reaction to partitioned pricing in settings that also include choices with a zero fee, it is important to include the discontinuous effect of free shipping. Obtaining unbiased parameter estimates are in turn important to calculate the welfare implications of partitioned pricing. To address these issues, I derive an empirical discrete choice model that can be interpreted within a theoretical framework on limited attention, as proposed by DellaVigna (2009). I use web scraped data from eBay Germany to reconstruct potential choice sets available to consumers and estimate the behavioral parameters. Following the framework of Bernheim and Rangel (2009), I then apply an approach proposed by Train (2015) for consumer welfare calculations when the choice-relevant utility function differs from the welfare-relevant utility function to calculate the consumer welfare implications of the observed behavior.

The joint analysis of a differential reaction to marginal changes in add-on fees and the base price as well as a potentially discontinuous effect of a zero fee is new to the literature. Prior research cannot disentangle the two effects because they either lack variation in the shipping fee (e.g. Morwitz et al. (1998) and the “low reserve treatment” in Hossain and Morgan (2006)) or do not consider listings with free shipping (e.g. Brown et al. (2010) and the “high reserve treatment” in Hossain and Morgan (2006)). Such a discontinuous effect of free shipping might, however, be relevant. Indeed, as Shampanier et al. (2007) show, demand increases discontinuously for goods that are sold at a price of zero. Einav et al. (2015) provide, to the best of my knowledge, the only other evidence in this direction. They find a discontinuous effect of free shipping. Listings with free shipping are, on average, associated with higher auction revenues conditional on a sale. In a separate analysis, the authors further show that conditional on a positive shipping fee, larger shipping fees are associated with larger revenues.[2]

Further, the welfare impact of partitioned pricing on consumer welfare in the online context is largely unexplored. Most relatedly, Chetty et al. (2009) and Taubinsky and Rees-Jones (2018) analyze the effect that limited attention to non-salient taxes has on the welfare impact of taxes.

I employ an empirical discrete choice model that allows for differential consumer reactions to variation in the total price of a good and the associated shipping fee as well as a discontinuous effect of free shipping. The estimated parameters can be interpreted as measures of limited attention following the DellaVigna (2009) framework.

To obtain the data necessary for the analysis, I automatically web scrape active listings on eBay Germany for various product categories several times a day. The publicly available data from eBay allows me to observe the exact time and price at which transactions occurred. Through my repeated web scrapes, I can reconstruct the potential choice set that each consumer was facing at the time of transaction.

Given the estimated coefficients, I calculate the expected loss in consumer surplus that occurs because of consumers’ reaction to partitioned pricing. To do so, I apply the framework of Bernheim and Rangel (2009) and assume that consumers would optimize perfectly in a world without partitioned pricing. I then apply an approach proposed by Train (2015) and based on Small and Rosen (1981) to calculate the loss in expected consumer surplus due to partitioned pricing.

[2] Furthermore, Frischmann et al. (2012) find that sellers listed on an online price comparison website tend to either offer free shipping or high shipping fees with no mass at smaller shipping fee values.

My results suggest that the degree to which consumers consider the shipping fee in their decisions varies across product categories. For the board games and smartphone in the sample, consumers behave as if they do not consider large parts of the shipping fee. For the video games in the sample, consumers seem to be fairly attentive to the fee. Hossain and Morgan (2006) find results suggesting that consumers in their auctions tend to ignore 18 to 45 percent of the shipping fee on average. The results of Chetty et al. (2009) even suggest behavior consistent with ignorance of 75 up to 94 percent of non-salient taxes. My estimates suggest ignorance of approximately five to 53 percent across the different product categories analyzed.

Additionally, my findings document a novel result concerning add-on fees: Consumer demand tends to react discontinuously positively to the offer of free shipping. This finding is in line with the findings of the research on consumer reaction to zero prices.

The relative loss in consumer surplus compared to fully rational behavior ranges from less than one to six percent. Three factors diminish the effect of partitioned pricing on consumer welfare in this setting: First, the average size of the shipping fee is relatively low compared to the total price. If sellers were to charge higher shipping fees, the welfare loss could be higher. Second, the estimates suggest that consumers are fairly attentive to the shipping fee for many product categories. Third, the free shipping effect on demand partly offsets the under-reaction to shipping fees in expectation.

This paper adds to the empirical literature on attention to add-on fees by focusing on posted price transactions. Previous research mainly analyses consumer behavior in auctions (Morwitz et al., 1998; Hossain and Morgan, 2006; Brown et al., 2010; Einav et al., 2015). These studies document that auctions with larger shipping fees tend to attract more bidders and receive earlier first bids than auctions with lower shipping fees. Conditional on a sale, the auctions with higher shipping fees generate higher revenues on average. However, while auctions were more popular in the early years of eBay, posted price purchases are now more common (Einav et al., 2018). Blake et al. (2021) provide one of the few studies analyzing attention in posted price transactions. Using data from a field experiment on StubHub, an online ticket resale platform, the authors show that revealing fees later in the purchasing process results in an average of 21 percent higher revenue. Their analysis suggests that at least 28 percent of this revenue increase results from consumers not only being more likely to purchase but also choosing higher quality products conditional on purchase. Similarly, Dertwinkel-Kalt et al. (2020) randomise the display of 3D surcharges when consumers are buying movie tickets on the website of a large German cinema. Interestingly, they do not find that revealing the fee at a later stage in the checkout process affects overall demand. While consumers do start the purchasing process more frequently if the fee is hidden, they are also more likely to end the process without buying once the fee is revealed. The authors argue that this null result is due to the fee being relatively large compared to the base price and that the cost of exiting is relatively low. The setting in Blake et al. (2021) as well as Dertwinkel-Kalt et al. (2020) is different from mine insofar as the difference in salience between the product price and the add-on fees in their setting is arguably larger because the hidden price component is only revealed in a later step of the transaction process.

More broadly, this paper also relates to the theoretical literature on price obfuscation and limited attention. Theoretical work has analyzed if and when firms can benefit from hiding product attributes from consumers if they are not fully attentive to them (e.g. Gabaix and Laibson, 2006; Heidhues et al., 2017; Johnen, 2020). More recent work also highlights how incentives to hide product attributes may differ between firms and two-sided platforms (Johnen and Somogyi, 2024).

I proceed as follows. In Section 2, I present an empirical discrete choice model. In Section 3, I discuss identification of the model parameters. In Section 4, I describe the eBay platform and my data collection procedure. In Section 5, I show some descriptive statistics and evidence from preliminary regressions. In Section 6, I provide the results from estimation of the demand model. In Section 7, I examine the welfare implications of my results. Section 8 concludes.

2 Model

Suppose each consumer i chooses to buy exactly one good from a choice set C_i. Assume that the indirect utility that a consumer i receives from purchasing a good j is given by

$$U_{ij} = x_{ij}'\gamma - \beta_i tp_{ij} + \epsilon_{ij}\ . \tag{2.1}$$

x_{ij} denotes non-financial characteristics that the consumer cares about and tp_{ij} denotes the full price (inclusive of fees) of the good. In the eBay setting with product prices and shipping fees, the fee-inclusive price can be denoted as $tp_{ij} = p_{ij} + c_{ij}$, where p_{ij} represents the base product price and c_{ij} represents the shipping fee. I discuss the selection of non-financial product characteristics in Section 3.1. ϵ_{ij} is the part of the utility that is observable only to the consumer and not to the econometrician. ϵ_{ij} could be non-zero, for example, because of differences in search behavior or distractions during the purchasing process.

Previous research suggests that consumers may react differently to variation in p_{ij} compared to variation in c_{ij}. If that is the case, the *perceived* indirect utility may differ from the true indirect utility. Write the perceived indirect utility as

$$\begin{aligned} \tilde{U}_{ij} &= x'_{ij}\gamma - \beta_i \tilde{tp}_{ij} + \epsilon_{ij} \\ &= x'_{ij}\gamma - \beta_i(p_{ij} + (1-\theta)c_{ij}) + \epsilon_{ij}\,. \end{aligned} \tag{2.2}$$

In line with DellaVigna (2009), θ can be interpreted as the degree of consumer inattention. In the framework of DellaVigna (2009), θ is bounded between zero and one. For fully attentive consumers, $\theta = 0$, while for fully inattentive consumers, $\theta = 1$. In my estimation I do not restrict the values of θ. Following DellaVigna (2009), I refer to θ as the inattention parameter, but note that due to the general form of the framework, θ can in fact capture mechanisms other than limited attention that can result in differential reactions to p_{ij} and c_{ij} (Taubinsky and Rees-Jones, 2018).[3]

Considering the findings of Shampanier et al. (2007) and Einav et al. (2015), I also allow for a discontinuous effect of free shipping on consumers' perceived utility, denoted as γ_f. When analyzing the welfare implications of partitioned pricing in a setting with listings that offer free shipping, ignoring such an effect could potentially bias the estimate of θ and, therefore, also the welfare calculations. I discuss this insight in more detail in Section 3. Making this discontinuity at a zero fee explicit, denote the perceived indirect utility as

$$\tilde{U}_{ij} = x'_{ij}\gamma - \beta_i(p_{ij} + (1-\theta)c_{ij}) + \gamma_f \mathbb{I}(c_{ij} = 0) + \epsilon_{ij}\,. \tag{2.3}$$

Assume that ϵ_{ij} is extreme value type I distributed and that consumers maximize their utility by choosing the one product in their choice set that yields the highest perceived utility. These assumptions yield a mixed logit model (McFadden and Train, 2000). To allow for more flexible substitution patterns, I allow the price coefficient β_i to vary across consumers (Berry

[3] Other mechanisms relevant in the online commerce context could be, for example, rounding or a left-digit bias, similar to what Lacetera et al. (2012) found in the used cars market. In both cases the resulting θ would be unclear, as the reaction to the price components would depend on the decimals in either of the components.

et al., 1995). One interpretation for this heterogeneity is that consumers have different price sensitivities because of unobserved differences in income.[4]

Note that I assume unobserved consumer heterogeneity in the price coefficient β_i. Because I assume that inattention affects the perceived total price, the heterogeneity of β_i interacts with consumer inattention. Note that this assumption is not equivalent to specifying a mixed logit model in which both the price coefficent and the shipping fee coefficient are random, because the heterogeneity of the coefficients in my model is coupled to the heterogeneity of β_i. In other words, each consumer takes one draw from the distribution of β_i that then transmits to the shipping fee coefficient. I parametrize β_i as $\beta_i \sim N(\mu_\beta, \sigma_\beta^2)$.

3 Identification

Because, in contrast to prior research, I am using observational data and focusing on posted price transactions, I am faced with several obstacles to identification of the parameters of interest, θ and γ_f. This section discusses these challenges and how I propose to overcome them.

3.1 Making Choices Comparable

DellaVigna (2009) proposes to estimate θ by keeping the visible part of the utility constant while exogeneously varying the opaque part (the shipping fee in this setting). With a measure of consumers' willingness-to-pay, it is then possible to identify θ. In second-price auctions, assuming rational bidding, the final price is the willingness-to-pay of the second highest bidder. Therefore, conducting experiments using second-price auctions is a natural path to identifying inattention θ. The majority of the early literature has used exactly this idea by auctioning identical goods while varying the add-on fee (e.g. Morwitz et al., 1998; Hossain and Morgan, 2006; Brown et al., 2010).

Because I focus on posted price transactions and because I use observational data, I encounter two obstacles to implementing this identification strategy. First, no measure of willingness-to-pay is observable. Second, I cannot only vary the shipping fee across products while keeping everything else constant. The structural assumptions on consumer decision-making help overcome the issue of unobserved willingness-to-pay by assuming a functional

[4]The assumption that the indirect utility is linear in income preference β_i can be seen as a linear approximation of non-linear income preferences around small price changes (Taubinsky and Rees-Jones, 2018).

form for it. However, the idea for identification of θ (and γ_f) remains the same except that I need to make different choices comparable conditional on observable characteristics. Therefore, my identification of θ and γ_f relies on comparing products for which observable characteristics are similar. In order to implement this strategy, I need to condition on all x_{ij}'s that might impact consumer demand. This conditioning is more difficult, the more heterogeneous the choices are. Therefore, I analyze product categories in which the different choices are arguably fairly homogeneous. As a result, the relevant variation in non-financial characteristics should come from observable seller and listing characteristics, not the characteristics of the products themselves. For the same reason, I restrict the analysis to products in new condition and exclude used or defunct ones.

The selection of products for the analysis was motivated by three additional considerations. First, correct estimation of the discrete choice model requires including all relevant choices in the estimation. For example, solely analyzing one particular kind of pencil would mean a very homogeneous product but would likely exclude various different kinds of pencils that are relevant substitutes. Therefore, I am focusing on product groups for which the relevant substitutes can be plausibly defined without introducing too much heterogeneity. I argue that board and video games, as well as specific smart phone models, are well-suited product categories in this regard. Second, I want to analyze consumer reaction to partitioned pricing for products of different price levels. Therefore, I use product categories that likely cover a wide range of product prices. Third, to maximize the expected number of observations, I include the most popular products in each category. Because eBay does not provide details about the popularity of individual items, I chose the most popular products in each category in early January 2019 according to Amazon Germany.

With these requirements in mind, I saved data on two board games ("Exit - Der versunkene Schatz" ("Exit") and "Azul"), three video games ("FIFA 19" for Playstation 4, "Spider Man" for Playstation 4, and "Pokémon Let's Go" for Nintendo Switch ("Pokemon")), and the "Samsung Galaxy J5 Duos" smart phone ("Duos").

However, even with such homogeneous products, a plausible specification of x_{ij} to include all remaining relevant non-financial listing characteristics is important to obtain unbiased estimates of θ, γ_f, and β. Although the choice of these products reduces the need to worry about unobserved differences in product quality, there is still variation in listing and seller characteristics that might affect demand. I discuss the choice of non-financial characteristics to include in the estimation in more detail in Section 6.

3.2 The Free Shipping Discontinuity

Note that few of the cited studies consider the discontinuous free shipping effect γ_f. Most studies only analyze θ. Morwitz et al. (1998), Blake et al. (2021), and Dertwinkel-Kalt et al. (2020) analyse experiments that compare the zero-fee case to only one level of fee and, thus, cannot distinguish between the effects of γ_f and θ. Hossain and Morgan (2006) and Brown et al. (2010) vary the amount of shipping required in their auctions but do not consider a case in which the shipping fee is zero. If the interest lies in obtaining an estimate of θ or in the net effect of partitioned pricing versus non-partitioned pricing, ignoring γ_f is reasonable. To identify θ, restricting the analysis to listings with a positive shipping fee is sufficient if willingness-to-pay is observable, even if there is a non-zero γ_f in reality. To see why, note that based on Equation (2.3), consumer i's expected willingness-to-pay (net of the shipping fee) can be written as $V_{ij} = \frac{\tilde{U}_{ij}}{\beta_i} + p_{ij}$. For clarity, V_{ij} can be rewritten as

$$V_{ij} = \begin{cases} x'_{ij}\frac{\gamma}{\beta_i} - (1-\theta)c_{ij}, & \text{if } c_j > 0 \\ x'_{ij}\frac{\gamma}{\beta_i} + \frac{\gamma_f}{\beta_i}, & \text{otherwise}. \end{cases} \tag{3.1}$$

Now consider two choices $j \in 1, 2$ for which $c_{i2} > c_{i1} > 0$ and $x_{i1} = x_{i2} = x_i$. This representation corresponds to the "High Reserve Treatments" in Hossain and Morgan (2006) as well as the treatments in Brown et al. (2010). Suppose V_{ij} is observable or can be estimated. Then, θ can be identified using Equation (3.1) as

$$\begin{aligned} V_{i1} &= x'_i\frac{\gamma}{\beta_i} - (1-\theta)c_{i1} \\ V_{i2} &= x'_i\frac{\gamma}{\beta_i} - (1-\theta)c_{i2} \\ \Rightarrow \theta &= 1 - \frac{V_{i1} - V_{i2}}{c_{i2} - c_{i1}}. \end{aligned} \tag{3.2}$$

Equation (3.2) shows that, in this situation, θ can be correctly identified from two non-zero values of c_{ij}, even in the presence of a zero-fee discontinuity because γ_f is irrelevant for $c_{ij} > 0$. This is exactly what Hossain and Morgan (2006) and Brown et al. (2010) do by considering treatments with different non-zero shipping fees.

However, it is not possible to identify both θ and γ_f separately using only two different treatments. Consider two treatments $c_{i1} = 0$ and $c_{i2} > 0$. This representation corresponds to, for example, the treatments in Morwitz et al. (1998) and in the "Low Reserve Treatments" of Hossain and Morgan (2006). Supposing that V_{ij} is observable or can be estimated, one can again use Equation (3.1) to rearrange as follows:

$$\begin{aligned} V_{i1} &= x_i' \frac{\gamma}{\beta_i} + \frac{\gamma_f}{\beta_i} \\ V_{i2} &= x_i' \frac{\gamma}{\beta_i} - (1-\theta) c_{i2} \\ \Leftrightarrow \theta - \frac{\gamma_f/\beta_i}{c_{i2}} &= 1 - \frac{V_{i1} - V_{i2}}{c_{i2}} \, . \end{aligned} \tag{3.3}$$

Equation (3.3) shows that θ and γ_f cannot be separately identified from these two treatments, even if willingness-to-pay were observable. Additionally, Equation (3.3) shows that with these two treatments, ignoring the presence of a zero-fee discontinuity γ_f, results in a biased estimate of θ. To see this, suppose I ignore γ_f. Then, in this setting, θ would be estimated as $\hat{\theta} = 1 - \frac{V_{i1} - V_{i2}}{c_{i2}}$. However, as Equation (3.3) shows, if $\gamma_f \neq 0$, then $\hat{\theta} \neq \theta$ and the estimator is biased. In particular, if $\gamma_f > 0$, $\hat{\theta} < \theta$ and inattention is underestimated. Intuitively, this result shows that if there is a positive effect of free shipping on demand, comparing a listing with free shipping to a listing with a positive shipping fee and ignoring the free shipping effect assigns the drop in demand entirely to the shipping fee, even though a part of it might be due to the drop caused by moving from the free shipping regime to any positive shipping fee.

For the results of Morwitz et al. (1998), these insights imply that θ and γ_f cannot be distinguished. However, this does not devalue their work, as first, they are only interested in showing a net effect of partitioned pricing on demand. Second, in their setting, it seems less likely that there is an effect of a fee of zero. The reason is that in their zero fee treatment, there is no mention of the fee at all. Thus, subjects are probably completely unaware that the other group is charged a fee. In the eBay setting, this is different because consumers see listings with both positive shipping fees as well as free shipping. Furthermore, free shipping is made slightly more salient than the shipping fee with a bold font. Ignoring γ_f in the eBay setting would, thus, likely lead to a biased estimate of θ which subsequently would affect the welfare calculations.

3.3 Additional Concerns

There are two additional potential concerns for identification. First, consumers may only consider subsets of the potential choice sets I observe. Second, there may be unobservable characteristics affecting both price as well as demand, resulting in price endogeneity.

One assumption of the demand model is that all consumers consider all choices included in the choice set. This full information assumption is typical for discrete choice models. However, the number of choices can be large in the eBay setting. For some of the products, more than one hundred choices were available at some points in time. Therefore, one concern is that the full information assumption is unrealistic in this setting.

If one were to observe the search behavior of every consumer, modeling the search or explicitly only using those listings that a consumer looked at would be a natural approach. However, I do not observe which listings consumers considered during their search. The crucial issue for the identification of θ and γ_f is whether the probability that a consumer considers a given listing is correlated with the total price and the shipping fee. Because part of the consideration process enters the estimation error ϵ, it needs to be independent of the total price and the shipping fee. At least in preliminary regressions, I do not find correlations between the rank or the search page on which my web scraper encounters a listing and its price or shipping fee. At least when considering that the rank is likely correlated with consideration probability, the fact that it is not correlated with price or shipping fee implies that the potential unobserved consideration set heterogeneity increases the noise of the estimates but should not cause any bias.

To address this issue further, I include a robustness exercise based on the approach suggested by Goeree (2008). The basic idea is that each available choice enters the consideration set of the consumer with a certain probability. Assuming a functional form for the consideration probabilities, the parameters determining these probabilities are jointly estimated together with the utility parameters. Denote as $P_{ij}(C, \beta_i)$ the probability that consumer i chooses product j out of a consideration set C. Given a realization of β_i, this choice probability is given by the logit choice probability. To obtain the probability that consumer i chooses product j, $P_{ij}(C, \beta_i)$ needs to be integrated over all consideration sets that contain product j, weighting these by the probability that consumer i uses that consideration set. Denote as π_{ij} the probability that consumer i considers product j. For any consideration set C, the probability that consumer i uses that consideration set is $\prod_{l \in C} \pi_{il} \prod_{k \notin C}(1 - \pi_{ik})$.

Finally, one needs to integrate over the distribution of β_i to obtain the choice probability for consumer i and product j:

$$P_{ij} = \int \sum_{C \in S_j} \prod_{l \in C} \pi_{il} \prod_{k \notin C} (1 - \pi_{ik}) P_{ij}(C, \beta_i) f(\beta_i) d\beta_i \,. \tag{3.4}$$

S_j is the set of all consideration sets that include choice j. Identification of π_{il} requires variables that impact the probability that a listing is considered but not consumers' utility. I estimate the model using simulation. For more details on the procedure, please refer to Goeree (2008) or Appendix A.2.

In the characteristics affecting the consideration probability but not utility, I include the rank and the page of the search results on which the listings appeared for my web scraper at the time closest to the purchase. Further, I include the total size of the choice set. This specification of the consideration probability can be regarded as a reduced form approximation of consumers' actual search and consideration processes.

Baye et al. (2009) show that the ranking on a search results page has a large impact on the clickthrough rate. The ranking and search page results that my web scraping program encounters are imperfect measures of the ranking and search page that the consumer sees. In particular, it depends on the exact search term that the consumer uses as well as the filters and sorting that they apply. Nevertheless, it is likely that the ranking and search page that my web scraping program sees is correlated with the probability that a consumer searching at the same time considers a listing. My web scraper observes the ranking and search page results sorted by eBay's default sorting algorithm, which is also what consumers observe first. Blake et al. (2016) show that almost 85 percent of consumers on eBay use the default sorting at first. Further, the authors show that, on average, eBay users start with a more general search (i.e. using fewer words) and refine that over time. I programmed my web scraper to search for rather general terms as well. Dinerstein et al. (2018) state that eBay's default ranking is not personalized for individual buyers. Therefore, it is likely that at least at the start of consumers' search on eBay, the ranking and search page results that the consumer sees are similar to those found by my web scraper. Furthermore, I include the total size of the choice set with the idea that more choices might result in the probability of consideration for each single choice decreasing.

Regarding price endogeneity, in my setting, I expect it to be less of an issue. I purposefully choose the product categories to be as homogeneous as possible. The sold products are all in new condition and basically identical. Therefore, unobserved differences in product quality, which can be a concern when analyzing consumer choices over different brands or car makes, is not an issue here. The non-financial variation that is relevant for choices in my context is that across different sellers and listings. However, because my web scraper sees the same information that consumers see when browsing through the listings, I expect that I can, at least in theory, observe all the relevant characteristics of a listing. I include fixed effects for commercial sellers as well as the seller score to account for observable differences in seller quality. It could still be that sellers or listings differ in unobserved characteristics, e.g. because some listings are more nicely designed. If nicer listings charged higher prices, then this would result in an underestimation of β. However, this correlation of unobserved characteristics and prices would have to occur *conditional* on the included observable seller characteristics. At least from own and anecdotal experience, it seems as if consumers base their decisions a lot more on these more salient, easily observable seller characteristics on eBay.

4 Data and Setting

For the analysis, I collect choice-level data on transactions on eBay Germany by automatically web scraping publicly available information. eBay's publicly available data is well-suited for discrete choice estimation because individual transactions can be observed. Another advantage for the assessment of consumer reaction to partitioned pricing is that sellers set their own shipping fees. This leads to the variation in shipping fees that is needed for the estimation. To reconstruct the potential choice sets faced by each consumer, I continuously save information on active and finished listings for each product category. I can then match observed transaction to those listings that were available at the time of purchase.

4.1 About eBay

eBay is an online marketplace that has been active since 1995. In the beginning, eBay only featured auctions. In 2002, eBay also introduced posted price purchases (so-called Buy-it-Now (BiN) listings). Since then, the BiN format has become increasingly popular. In recent years, the majority of listings on eBay worldwide use the BiN format, although this differs across product categories (Hasker and Sickles, 2010; Einav et al., 2018). Einav et al. (2018) further document that auctions are more popular among less experienced sellers, for used goods, and for more heterogeneous goods. Their estimates suggest that the decrease in popularity of auctions cannot be explained by a change in the composition of products sold but rather by a decrease in the demand for auctions, and, to a lesser extent, by an increase in competition on eBay.

Sellers on eBay Germany can choose whether to list their product as an auction, an auction with BiN option, or a BiN listing. For BiN listings, there is also the possibility to list an inventory of a product to sell multiple units of the same item. This is often used by commercial sellers who use eBay as a platform for their retail business. Sellers on eBay range from private sellers to smaller commercial sellers to traditional brick-and-mortar stores.

Sellers can also choose whether to offer free shipping or set a shipping fee for their listings. At the time the data were collected, eBay Germany capped the shipping fee at 9.50 euro for many product categories.

To build trust, eBay includes a system of ratings in which sellers and buyers rate each other after successful transactions. The reputation of sellers on eBay depends mainly on two measures: the eBay seller score and the percentage of positive reviews in total reviews received. The eBay seller score is calculated as the number of positive reviews minus the number of negative reviews. Further, the eBay seller score is also presented as discretized values in the form of eBay stars. These eBay stars are small icons that are shown next to each seller's eBay score. In total, there are 12 different icons that sort sellers into different brackets according to their eBay score.

When searching for an item on eBay Germany, consumers have various choices of how to sort their search results. The default sorting is an algorithm that is supposed to maximize eBay's expected income (Blake et al., 2016). According to the eBay website, the algorithm takes into account the completeness of the product description, the competitiveness of the listing price, and the seller's services (e.g. return policy, speed of delivery, past reviews). Consumers can also sort by the geographical distance to their location, the time until the end of a listing, and newly advertised listings. Further, consumers can sort with regard to price, both including and excluding shipping. There is also the possibility to save searches and receive notifications whenever a relevant listing is added. Here, however, only the product price can be set as a relevant parameter, but not the shipping or total price.

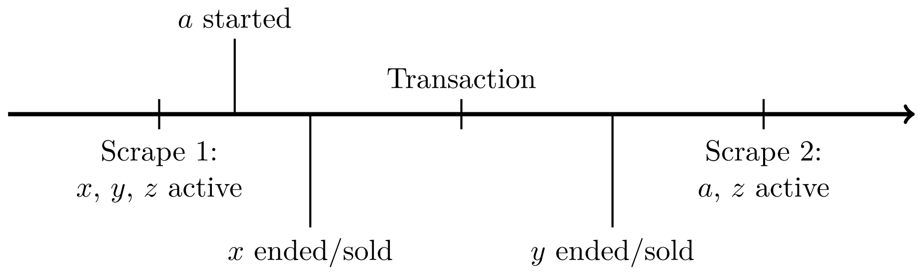

Figure 1: An illustration of the choice set reconstruction

4.2 Choice Set Creation

To collect the data, the web scraper searches for each product category and subsequently visits all the listing pages that are found as a result. The eBay website offers well-suited data for the estimation because individual transactions can be observed. Appendix A.1 provides more details on the web scraping procedure and how to identify transactions on eBay Germany.

In order to estimate the discrete choice model, I reconstruct the choice sets that consumers were facing. To do so, I searched eBay Germany for each product category and saved information on all active listings that were shown as search results multiple times a day. In addition, less frequently, I saved the results for all finished listings that matched the search.

Each observed purchase is then used as the base for one choice situation. To reconstruct the potential choice set for each choice situation, I match all listings that I observe being active before the time of purchase to the listings that are either active or ended after the purchase. Because I save all active listings for each search term multiple times a day, the reconstruction works plausibly precisely.

Figure 1 shows an illustration of this process. For each transaction, I compare the set of listings that were active in any of the scrapes up to 24 hours prior to the transaction to those that were active in any of the scrapes up to 24 hours after the transaction.[5] All listings that were active both before and after the transaction I consider to be in the choice set (listing z in the example). For listings that were active before but not after the transaction, I compare the time the listing ended to the time of the transaction. If the listing ended before the transaction (such as x in the example), I do not include it in the choice set. If the listing ended after the transaction (y in the example), I include it in the choice set. Because I cannot observe the exact time a listing was first activated, I do not know if listings that first

[5]The 48 hours tolerance window is chosen to allow some flexibility for potential misses by the scraper. With about three to four scrapes a day, it is unlikely that an active listing is missed in each repetition.

appeared in the scrape after the transaction were activated before or after the transaction. Therefore, I do not include these listings. This means that I wrongfully exclude listings that were activated between the last scrape before the transaction and the time of transaction, such as a. However, since I scrape new data once every few hours, this should not be a big issue.

4.3 Sample

In order to reduce the probability of missing relevant listings, I search for listings in a rather broad fashion. To obtain the final sample for the analysis, I subsequently exclude all listings that are not posted price listings. Further, I exclude all listings with items that are not in new condition. Additionally, I only include listings that are located in Germany. The results often also include listings shipped from outside Germany and, as a result, have comparatively high shipping fees. While this could introduce interesting variation in the shipping fee, the main problem is that I do not observe the location of the buyer. Therefore, if I see a transaction I need to assume that the buyer is from Germany, because the shipping fees that I observe are those that apply to shipping to Germany. When including listings located outside Germany, the likelihood is high that I actually observe a transaction with a buyer outside of Germany for whom the assigned shipping fee as well as the choice set would be incorrect. While there are also listings in Germany that ship to other countries, I expect that the probability that I actually observe an order from outside of Germany on a German listing is low. Finally, I only include listings in the analysis for which I see at least one purchase at any point. There are many listings that are active on eBay that are never purchased. The implicit assumption is that these listings that are never purchased are not relevant substitutes. On average, these listings are more expensive which suggests that indeed these listings are not competitive.

Next, I need to make sure that I only include listings that are actually relevant substitutes. As an example, when searching for the “Duos,” usually a large part of the search results are actually cases or other accessories for the phone. Excluding these irrelevant results is complicated because entering product information is not mandatory for the sellers on eBay. If those details are available, I use them to determine whether a listings should be part of the sample. For listings where such information is not readily available, I use the title of the listing to infer its relevance. Further, I use seller-entered product characteristics as well as the title to infer product-specific characteristics such as the color of the phone.

Sometimes, a listing offers to sell different models of a product. For the "Duos," for example, some listings would have different colors available. In these cases, I treat each of the different models as a separate observation (given they are a relevant choice).

5 Descriptive Statistics and Regressions

Table 1 shows summary statistics for the final sample for selected variables. Each observation is one listing in one particular choice situation. This means that listings that are part of multiple choice sets enter the averages multiple times.

Table 1: Descriptive statistics

	Exit	Azul	FIFA 19	Spiderman	Pokemon	Duos
Total price	14.32	38.23	41.51	41.69	58.16	175.79
	[12.98, 14.95]	[36.29, 39.99]	[32.89, 44.99]	[34.98, 44.99]	[43.88, 74.88]	[159.89, 184.98]
Product price	13.08	37.77	40.51	40.59	56.83	174.88
	[10.99, 14.69]	[35.98, 39.99]	[29.99, 44.95]	[34.00, 44.90]	[42.90, 69.99]	[159.89, 179.99]
Shipping fee	1.24	0.46	1.00	1.10	1.33	0.91
	[0.00, 1.99]	[0.00, 0.00]	[0.00, 1.99]	[0.00, 1.99]	[0.00, 1.99]	[0.00, 0.00]
Shipping fee (> 0)	3.10	3.02	2.92	2.78	3.16	4.26
	[1.99, 4.95]	[1.99, 3.00]	[1.99, 3.79]	[1.99, 3.90]	[1.99, 4.90]	[1.99, 4.99]
Share of shipping in total price	0.08	0.01	0.03	0.03	0.02	0.01
	[0.00, 0.17]	[0.00, 0.00]	[0.00, 0.04]	[0.00, 0.05]	[0.00, 0.04]	[0.00, 0.00]
Share of shipping in total price (> 0)	0.21	0.08	0.08	0.07	0.05	0.03
	[0.15, 0.27]	[0.06, 0.08]	[0.04, 0.11]	[0.04, 0.10]	[0.04, 0.07]	[0.01, 0.03]
Free shipping	0.60	0.85	0.66	0.60	0.58	0.79
	[0.00, 1.00]	[1.00, 1.00]	[0.00, 1.00]	[0.00, 1.00]	[0.00, 1.00]	[1.00, 1.00]
Seller score	251122.66	124708.15	107094.03	90883.36	166150.26	148842.67
	[17096.00, 552311.00]	[1673.00, 47443.00]	[628.00, 27819.00]	[1065.00, 83570.00]	[1190.00, 71471.00]	[4391.00, 245819.00]
Pos. reviews (%)	99.27	98.86	98.34	99.41	99.33	99.46
	[99.60, 99.90]	[99.30, 100.00]	[99.50, 100.00]	[99.60, 100.00]	[99.60, 100.00]	[99.40, 100.00]
Commercial seller	0.98	0.97	0.86	0.77	0.80	0.97
	[1.00, 1.00]	[1.00, 1.00]	[1.00, 1.00]	[1.00, 1.00]	[1.00, 1.00]	[1.00, 1.00]
Multiple units	0.89	0.92	0.69	0.64	0.69	0.80
	[1.00, 1.00]	[1.00, 1.00]	[0.00, 1.00]	[0.00, 1.00]	[0.00, 1.00]	[1.00, 1.00]
Payment: Paypal	0.98	0.99	0.97	0.91	0.91	0.98
	[1.00, 1.00]	[1.00, 1.00]	[1.00, 1.00]	[1.00, 1.00]	[1.00, 1.00]	[1.00, 1.00]
New edition		0.08				
		[0.00, 0.00]				
Pokeball bundle					0.26	
					[0.00, 1.00]	
Eevee edition					0.47	
					[0.00, 1.00]	
Gold						0.25
						[0.00, 0.00]
Blue						0.21
						[0.00, 0.00]
Observations	444	1347	34642	15488	41289	13631

Notes: Sample means with lower and upper quartiles in brackets. Each observation represents one listing in one choice situation.

Note that the products are sorted in ascending mean total and product price. The average shipping fee does not increase proportionally with the average product price, therefore the average share of the shipping fee in the total price is smaller in product categories with a higher total price. Focusing on listings that set a non-zero shipping fee, the fee makes up about 21 percent of the total price on average in the "Exit" product category. In the "Duos" product category that share is only three percent. A majority of the listings offer free shipping (i.e. do not use partitioned prices) with the shares ranging from 58 to 85 percent across the different product categories.

The two main indicators through which eBay tries to build trust between buyers and sellers are the share of positive reviews a seller received and the seller's eBay score. There is little variation in the share of positive seller reviews and almost all sellers have a very high rating. Einav et al. (2015) also document this pattern. There is, however, larger variation in the eBay seller score, which is the number of positive reviews a seller received minus the number of negative reviews. Thus, it can be interpreted as a measure of seller experience. The vast majority of listings in the sample are offered by commercial sellers. The share of commercial sellers ranges from 77 percent to 98 percent. The variables in the lower panel of Table 1 are product-specific variables that I include for some of the products to account for different consumer valuation of different models and editions that were available. These variables are all indicator variables.

To assess whether there is a relationship between the shipping fee and the total price, Table 2 reports the main results of an OLS regression of total prices on the shipping fee and other covariates. Here, I restrict the sample to only those listings that have a positive shipping fee. Note that the product price is not included in the regression. With fully attentive consumers, one would expect an average one-to-one decrease in the product price for each additional euro in the shipping fee, keeping the total price constant. Therefore, if consumers fully internalized the fee, one would not expect a change in the total price if the shipping fee changes.

However, the estimated shipping fee coefficient for four of the six products is positive (albeit too noisy to statistically distinguish from zero for "Azul"). These results suggest that the total price of a listing, given that it does not offer free shipping, are, on average, higher with higher shipping fees for these product cetagories. This evidence is in line with at least some sellers trying to exploit consumer inattention to fees. The point estimate for the "Pokemon" product category suggests a negative relationship between the shipping fee and the total price.

In order to avoid omitted variable bias, ideally, I need to include all variables that correlate with the shipping fee as well as consumer utility. While I cannot directly measure correlation with consumer utility, I can explore correlation with whether a seller sets a seperate shipping fee at all and, if so, the size of the shipping fee. Table 3 shows the results of a linear regression of whether a listings features free shipping on various covariates. This regression

Table 2: Total price and shipping fee

	Exit	Azul	FIFA 19	Spiderman	Pokemon	Duos
Dependent variable	Total price	Total price	Total price	Total price	Total price	Total price
Shipping fee	3.121***	2.045	3.257***	-0.112	-1.092**	6.374*
	[2.791, 3.451]	[-0.573, 4.663]	[0.866, 5.648]	[-1.761, 1.537]	[-1.943, -0.241]	[-0.020, 12.768]
Commercial seller	0.000	-132.800**	10.430***	6.354***	3.331**	11.680
	[.,.]	[-232.495, -33.105]	[4.606, 16.254]	[2.108, 10.600]	[0.577, 6.085]	[-12.828, 36.188]
Seller score	0.000***	0.000	0.000	0.000***	0.000	0.000***
	[0.000, 0.000]	[-0.000, 0.000]	[-0.000, 0.000]	[0.000, 0.000]	[-0.000, 0.000]	[0.000, 0.000]
Pos. reviews (%)	8.931***	6.651***	0.053	0.016	0.024	-6.334
	[7.781, 10.081]	[2.552, 10.750]	[-0.019, 0.126]	[-0.018, 0.051]	[-0.080, 0.128]	[-14.959, 2.291]
Multiple units	-0.695***	1.443	-2.449	-2.159	6.131***	-2.801
	[-0.977, -0.413]	[-1.073, 3.959]	[-7.055, 2.157]	[-5.887, 1.569]	[3.665, 8.597]	[-26.603, 21.001]
Payment: Cash on delivery	0.000	132.400**	-8.574**	0.000	0.000	0.000
	[.,.]	[33.631, 231.169]	[-15.466, -1.682]	[.,.]	[.,.]	[.,.]
Payment: Cash on pickup	0.000	-1.693***	-3.554	0.304	0.289	14.400**
	[.,.]	[-2.393, -0.993]	[-9.250, 2.142]	[-4.093, 4.701]	[-2.622, 3.200]	[2.468, 26.332]
Payment: Credit card	0.255	-0.634	3.861	-9.006***	0.320	-48.250***
	[-0.643, 1.153]	[-1.691, 0.423]	[-5.420, 13.142]	[-10.755, -7.257]	[-1.524, 2.164]	[-75.307, -21.193]
Payment: Check	0.000	0.000	0.000	0.000	0.000	0.000
	[.,.]	[.,.]	[.,.]	[.,.]	[.,.]	[.,.]
Payment: Debit	0.000	0.000	0.000	0.000	0.000	0.000
	[.,.]	[.,.]	[.,.]	[.,.]	[.,.]	[.,.]
Payment: Other	0.000	0.000	0.000	0.000	-5.790***	20.600***
	[.,.]	[.,.]	[.,.]	[.,.]	[-10.010, -1.570]	[7.036, 34.164]
Payment: Paypal	-2.333***	2.552	-4.819	10.170***	-3.877**	68.370***
	[-3.229, -1.437]	[-1.921, 7.025]	[-13.408, 3.770]	[4.987, 15.353]	[-7.076, -0.678]	[36.686, 100.054]
Payment: Receipt	0.000	0.000	2.626	-3.835	-4.780***	15.910
	[.,.]	[.,.]	[-4.801, 10.053]	[-10.137, 2.467]	[-7.922, -1.638]	[-9.617, 41.437]
Payment: Transfer	0.000	-134.300**	-0.475	-2.690	-0.391	-14.370
	[.,.]	[-232.359, -36.241]	[-10.915, 9.965]	[-7.939, 2.559]	[-3.809, 3.027]	[-38.250, 9.510]
Intercept	-888.600***	-510.200***	18.850***	36.470***	40.690***	728.500*
	[-1004.085, -773.115]	[-833.515, -186.885]	[7.379, 30.321]	[27.165, 45.775]	[29.463, 51.917]	[-112.801, 1569.801]
New edition		7.184***				
		[7.025, 7.343]				
Eevee edition					1.217	
					[-0.529, 2.963]	
Pokeball edition					40.650***	
					[38.335, 42.965]	
Blue						16.440
						[-7.205, 40.085]
Gold						7.614
						[-5.079, 20.307]
Observations	183	209	14848	9932	20724	3628
R^2	0.705	0.943	0.349	0.566	0.921	0.617
Mean total price	14.553	36.964	41.843	44.572	62.062	173.052

Notes: Includes only listings with a positive shipping fee. ***, **, and * indicate statistical significance at the one, five, and ten percent level, respectively. The brackets show 95 percent confidence intervals.

helps understanding whether listings that feature free shipping are systematically different from those that do not in some observable characteristics.

The results show that commercial sellers are more likely to set a zero shipping fee. At the same time, more experienced sellers, as measured by the eBay seller score, seem to be less likely to offer free shipping. For other covariates, no patterns emerge that apply systematically across all product categories.

Given that a seller does set a separate shipping fee, there may be systematic differences between listings that set a higher versus lower fee. To understand which observable characteristics correlate with the size of the shipping fee, I analyze shipping fees focusing only on listings with a positive shipping fee. Table 4 shows the main coefficients from such a

Table 3: Covariates of free shipping

	Exit	Azul	FIFA 19	Spiderman	Pokemon	Duos
Dependent variable	Free shipping	Free shipping	Free shipping	Free shipping	Free shipping	Free shipping
Commercial seller	1.292***	0.373***	0.509***	0.284	0.563***	0.604***
	[0.842, 1.742]	[0.167, 0.579]	[0.268, 0.750]	[-0.139, 0.707]	[0.419, 0.707]	[0.314, 0.894]
Seller score	-0.000**	-0.000***	-0.000***	-0.000***	-0.000***	-0.000***
	[-0.000, -0.000]	[-0.000, -0.000]	[-0.000, -0.000]	[-0.000, -0.000]	[-0.000, -0.000]	[-0.000, -0.000]
Pos. reviews (%)	-0.002	0.004	-0.005**	-0.001	0.000	-0.145**
	[-0.005, 0.001]	[-0.004, 0.012]	[-0.009, -0.000]	[-0.005, 0.002]	[-0.003, 0.003]	[-0.262, -0.028]
Multiple units	-0.142	-0.021	-0.035	0.065	-0.051	-0.110
	[-0.535, 0.251]	[-0.082, 0.039]	[-0.237, 0.166]	[-0.222, 0.351]	[-0.154, 0.051]	[-0.353, 0.133]
Payment: Cash on delivery	0.208**	-1.025***	0.270**	0.148	0.000	0.221*
	[0.061, 0.355]	[-1.088, -0.962]	[0.013, 0.527]	[-0.126, 0.422]	[.,.]	[-0.032, 0.474]
Payment: Cash on pickup	-0.003**	0.005	-0.071	0.186	0.093*	-0.409***
	[-0.005, -0.000]	[-0.023, 0.033]	[-0.298, 0.157]	[-0.079, 0.451]	[-0.017, 0.202]	[-0.644, -0.174]
Payment: Credit card	-0.239	0.025	-0.052	-0.227*	-0.023	0.679***
	[-0.571, 0.093]	[-0.044, 0.094]	[-0.178, 0.073]	[-0.492, 0.038]	[-0.122, 0.076]	[0.303, 1.055]
Payment: Other	0.000	0.033***	0.311***	0.270**	0.088	-0.109
	[.,.]	[0.014, 0.052]	[0.111, 0.511]	[0.021, 0.519]	[-0.187, 0.362]	[-0.340, 0.122]
Payment: Paypal	0.000	0.553***	-0.075	0.546***	0.260*	-0.410
	[.,.]	[0.255, 0.851]	[-0.390, 0.241]	[0.156, 0.936]	[-0.010, 0.530]	[-0.953, 0.133]
Payment: Receipt	0.105	0.095	0.053	0.203	0.006	0.228***
	[-0.156, 0.366]	[-0.046, 0.236]	[-0.232, 0.337]	[-0.062, 0.468]	[-0.210, 0.221]	[0.078, 0.378]
Payment: Transfer	0.230	-0.011	-0.151	-0.039	-0.102*	0.063
	[-0.071, 0.531]	[-0.041, 0.018]	[-0.386, 0.084]	[-0.351, 0.273]	[-0.213, 0.009]	[-0.168, 0.294]
Intercept	-0.030	-0.304	1.003***	0.137	0.174	14.510**
	[-0.211, 0.150]	[-1.091, 0.483]	[0.564, 1.442]	[-0.237, 0.511]	[-0.047, 0.395]	[3.010, 26.010]
New edition		-0.005				
		[-0.103, 0.094]				
Eevee edition					-0.032	
					[-0.128, 0.063]	
Pokeball bundle					-0.194***	
					[-0.313, -0.075]	
Blue						0.206**
						[0.047, 0.365]
Gold						0.037
						[-0.150, 0.225]
Observations	448	1350	34646	15500	40990	13635
R^2	0.530	0.730	0.340	0.384	0.597	0.482
Mean free shipping share	.141	.089	.055	.063	.034	.055

Notes: ***, **, and * indicate statistical significance at the one, five, and ten percent levels, respectively. The brackets show 95 percent confidence intervals.

regression. A striking result is that listings that have an inventory and sell multiple units of a product tend to charge higher shipping fees on average. For “FIFA 19” and “Duos,” this conditional correlation is the largest, suggesting more than one euro more in shipping fees for listings with an inventory.

These preliminary regressions are important in determining which observable characteristics should be included in the specification for the demand estimation. In Section 6, I discuss how I use these preliminary regressions to inform the selection of characteristics to include in the demand estimation.

Table 4: Covariates of the shipping fee

	Exit	Azul	FIFA 19	Spiderman	Pokemon	Duos
Dependent variable	Shipping fee	Shipping fee	Shipping fee	Shipping fee	Shipping fee	Shipping fee
Commercial seller	1.748***	22.570***	-0.770**	0.563	-0.184	-0.580
	[1.623, 1.873]	[16.921, 28.219]	[-1.483, -0.057]	[-1.060, 2.186]	[-0.631, 0.263]	[-1.647, 0.487]
Seller score	-0.000***	-0.000***	-0.000***	-0.000***	-0.000***	-0.000***
	[-0.000, -0.000]	[-0.000, -0.000]	[-0.000, -0.000]	[-0.000, -0.000]	[-0.000, -0.000]	[-0.000, -0.000]
Pos. reviews (%)	-1.432	-0.948***	0.014**	-0.020***	-0.016	0.299
	[-3.224, 0.360]	[-1.181, -0.715]	[0.001, 0.028]	[-0.032, -0.008]	[-0.049, 0.018]	[-0.452, 1.050]
Multiple units	0.314**	0.310***	1.503***	0.439	0.416*	1.992***
	[0.122, 0.506]	[0.200, 0.420]	[0.944, 2.062]	[-0.904, 1.782]	[-0.015, 0.847]	[1.516, 2.468]
Payment: Receipt	0.000	0.000	-0.287	-0.427	-1.734***	-0.582
	[.,.]	[.,.]	[-1.108, 0.534]	[-2.952, 2.098]	[-2.197, -1.271]	[-1.437, 0.273]
Payment: Cash on delivery	0.000	-22.420***	1.757**	0.000	0.000	0.000
	[.,.]	[-28.138, -16.702]	[0.273, 3.241]	[.,.]	[.,.]	[.,.]
Payment: Cash on pickup	0.000	-0.224	-0.196	-0.260	-0.371	0.181
	[.,.]	[-0.735, 0.287]	[-1.047, 0.655]	[-1.234, 0.714]	[-0.969, 0.227]	[-0.894, 1.256]
Payment: Credit card	-0.136**	0.021	0.463	0.062	-0.349**	-1.497*
	[-0.215, -0.057]	[-0.145, 0.187]	[-0.141, 1.067]	[-0.338, 0.462]	[-0.667, -0.031]	[-3.228, 0.234]
Payment: Other	0.000	0.000	0.000	0.000	0.681	-0.042
	[.,.]	[.,.]	[.,.]	[.,.]	[-0.209, 1.571]	[-0.822, 0.738]
Payment: Paypal	0.000	-1.271***	0.117	0.372	0.182	0.640
	[.,.]	[-1.482, -1.060]	[-0.851, 1.085]	[-0.730, 1.474]	[-0.426, 0.790]	[-1.370, 2.650]
Payment: Transfer	0.000	22.620***	-0.122	-0.194	0.080	2.503***
	[.,.]	[17.129, 28.111]	[-0.941, 0.697]	[-1.664, 1.276]	[-0.491, 0.650]	[0.897, 4.109]
Intercept	146.000	78.150***	1.138	4.730***	4.750***	-26.260
	[-33.210, 325.210]	[60.368, 95.932]	[-0.287, 2.563]	[2.942, 6.518]	[1.343, 8.157]	[-99.456, 46.936]
New edition		0.035				
		[-0.033, 0.104]				
Eevee edition					-0.139	
					[-0.594, 0.316]	
Pokeball bundle					0.970***	
					[0.480, 1.460]	
Blue						0.038
						[-0.776, 0.851]
Gold						0.336
						[-0.162, 0.834]
Observations	178	205	11862	6165	17368	2919
R^2	0.966	0.979	0.419	0.322	0.668	0.896
Mean shipping fee	.141	.089	.055	.063	.034	.055

Notes: Includes only listings with a positive shipping fee. ***, **, and * indicate statistical significance at the one, five, and ten percent levels, respectively. The brackets show 95 percent confidence intervals.

6 Discrete Choice Estimation Results

This section discusses the selected covariates and provides the parameter estimates from estimation of the demand model for each of the product categories. The covariate selection builds on the results of the analyses shown in Section 5.

6.1 Included characteristics

A major issue in online marketplaces is the asymmetric information between sellers and buyers. Trust issues concern, for example, the condition of the product and the speed of processing. Thus, building trust is an important task for online platforms (Tadelis, 2016). To capture the effect of trustworthiness of a seller, I include the eBay seller score, which is the number of positive reviews minus the number of negative reviews that a seller has received. I do not include the share of positive reviews because this share is mostly either zero (if a seller does not have sufficiently many reviews) or very high and does not vary much. The seller score is also correlated with whether a listing offers free shipping and the size of the shipping fee if a listing does not feature free shipping.[6] Furthermore, one can expect that commercial sellers are viewed as more trustworthy on average. Therefore, I include an indicator variable for commercial sellers. Including this commercial seller fixed effect is also important, as Table 3 shows that commercial sellers seem to offer free shipping more frequently. When consumers pay their purchases using Paypal, there is an additional layer of consumer protection in case something goes wrong. Availability of Paypal can therefore also increase trust. Table 3 further suggests that listings that accept PayPal as a payment method are also more likely to offer free shipping, at least in some of the product categories. Therefore, I also include an indicator variable for listings that accept PayPal for all product categories except for "Exit" and "Azul". For "Exit", the Paypal variable is perfectly collinear with the commercial seller variable. For "Azul", the share of observations accepting Paypal is 99 percent. Further, as Table 4 indicates, including a fixed effect for listings that sell multiple units is important. Otherwise, the positive correlation of having an inventory and charging a higher shipping fee will result in an upward-biased estimate of the shipping fee coefficient if consumers value buying from inventory listings.

Finally, I include product-specific characteristics where necessary. Specifically, for "Azul," I include a fixed effect if the listing sells the second edition. For "Pokemon," two versions of the game exist: the "Pikachu" and the "Eevee" edition. I include an indicator for the "Eevee" edition, leaving the other as the reference category. Additionally, I include an indicator for bundles that include a "Pokeball" controller that can be used with the game. For "Duos," I include fixed effects for different phone colors.

6.2 Main results

Table 5 shows the estimation results for all six product categories. The estimates suggest limited under-reaction to the shipping fee in the eBay setting. Only for "Azul" and "Duos", the estimates of θ are statistically distinguishable from zero, i.e. the case of full attention. For example, consumers buying the "Azul" board game behave as if they ignore 53 percent

[6] See Tables 3 and 4, respectively.

of the shipping fee. The estimates of θ for the other four product categories are positive but not statistically distinguishable from zero at conventional confidence levels. The results also suggest a positive effect of free shipping on demand for five of the product categories. Only for “FIFA 19” is the free shipping parameter not statistically significantly different from zero.

In general, consumers buying the board games seem to be less attentive to the shipping fee. Consumers buying the three video games seem to be most attentive to the shipping fee. The data do not allow to make clear statements about the sources of the differences in the estimates of θ for the different product categories. One explanation could be that the consumers buying the different products are inherently different with regard to their attention to shipping fees. Drivers of such differences could be, for example, that consumers buying the video games have lower incomes or are more tech-savvy and used to buying products online. Lower income might result in more attention to the price components in general while experience in online shopping might increase attention to common online practices like partitioned pricing. While consumers buying a smart phone online may be tech-savvy as well, the lower average share that the shipping fee has in the total price in the “Duos” category may explain why consumers exhibit less attention to it compared to in the video games categories. However, without additional data on consumers, these statements are only speculative.

6.3 Additional results

Table 6 shows the estimation results for all six product categories when allowing for random consideration as proposed by Goeree (2008). While less precisely estimated, the estimates are mostly consistent with those in Table 5. In particular, the three video game categories still show the least degree of limited attention, while there is some evidence for limited attention in the board games and smartphone categories. The main difference is that the degree to which consumers seem to ignore the shipping fee is much larger for “Exit” than in the baseline results. In fact, the estimates suggest that consumers fully ignore the shipping fee.

The chosen variables to include in the consideration probabilities seem to work well in the sense that the signs are mostly in the expected direction. A larger rank (i.e. being further down the page) or being found on a later page result in lower consideration probabilities. Similarly, being in a larger potential choice set reduces consideration probabilities.

Table 5: Demand model estimates

	Exit	Azul	FIFA 19	Spiderman	Pokemon	Duos
Inattention (θ)	0.167	0.530***	0.053	0.051	0.102	0.171***
	[-0.249, 0.583]	[0.340, 0.721]	[-0.107, 0.213]	[-0.016, 0.118]	[-0.046, 0.251]	[0.079, 0.263]
Free shipping effect (γ_f)	2.340***	0.465***	-0.073	1.226***	0.667***	1.036***
	[1.514, 3.166]	[0.155, 0.776]	[-0.436, 0.290]	[1.026, 1.427]	[0.354, 0.980]	[0.887, 1.185]
Price coefficient: Mean ($-\beta$)	-0.535***	-0.630***	-0.263***	-0.714***	-0.424***	-0.193***
	[-0.910, -0.161]	[-0.668, -0.593]	[-0.276, -0.250]	[-0.725, -0.704]	[-0.440, -0.408]	[-0.201, -0.184]
Price coefficient: SD (σ_β)	0.552***	0.367***	0.040	0.309***	0.180***	0.084***
	[0.189, 0.916]	[0.319, 0.415]	[-0.018, 0.099]	[0.301, 0.318]	[0.175, 0.186]	[0.077, 0.091]
Shipping fee coefficient ($-\beta\theta$)	0.089	0.334***	0.014	0.037*	0.043	0.033
	[-0.154, 0.333]	[0.169, 0.500]	[-0.089, 0.117]	[-0.005, 0.079]	[-0.038, 0.125]	[-0.013, 0.079]
Free shipping coefficient	2.340***	0.465***	-0.073	1.226***	0.667***	1.036***
	[1.514, 3.166]	[0.155, 0.776]	[-0.436, 0.290]	[1.026, 1.427]	[0.354, 0.980]	[0.887, 1.185]
Commercial seller	-1.963	1.487***	-0.001	0.650***	1.032***	1.224***
	[-5.109, 1.182]	[1.132, 1.842]	[-0.155, 0.152]	[0.405, 0.894]	[0.827, 1.236]	[0.733, 1.716]
Seller score (K)	0.003***	0.001***	0.001***	-0.001**	-0.000	-0.001**
	[0.002, 0.004]	[0.001, 0.002]	[0.001, 0.001]	[-0.002, -0.000]	[-0.000, 0.000]	[-0.002, -0.000]
Multiple units	-0.559***	0.259	1.299***	2.021***	0.825***	0.814***
	[-0.924, -0.194]	[-0.069, 0.588]	[1.156, 1.441]	[1.948, 2.094]	[0.675, 0.975]	[0.705, 0.923]
Payment: Paypal			0.947***	1.236***	0.281***	2.157***
			[0.662, 1.232]	[1.191, 1.281]	[0.097, 0.465]	[1.618, 2.696]
New edition		0.759***				
		[0.223, 1.295]				
Pokeball bundle					8.588***	
					[8.318, 8.858]	
Eevee edition					-0.157***	
					[-0.276, -0.039]	
Gold						-0.188*
						[-0.391, 0.015]
Blue						-0.356***
						[-0.575, -0.137]
Draws	100	100	100	100	100	100
Observations	444	1347	34642	15488	41289	13631
Transactions	59	117	1918	940	1396	744
Unique listings	19	47	454	179	535	158
McFadden's R^2	0.20	0.36	0.34	0.57	0.37	0.48

Notes: ***, **, and * indicate statistical significance at the one, five, and ten percent levels, respectively. The brackets show 95 percent confidence intervals. θ is calculated as $\theta = \frac{\hat{\theta}}{\beta}$. Standard errors for θ are calculated using the Delta method.

7 Consumer Welfare Implications

In this section, I use the point estimates presented in Section 6 to assess the impact of partitioned pricing on consumer welfare. To do so, first, I fix ideas on how to define consumer welfare in a context in which consumers make decisions based on a notion of perceived welfare that may differ from their realized welfare.

I follow the framework proposed by Bernheim and Rangel (2009) and described in Bernheim and Taubinsky (2018). This framework differentiates between the naturally occurring domain and the welfare-relevant domain. The welfare-relevant domain is that in which consumers make decisions based on fully rational welfare maximization. In contrast, the naturally occurring domain describes how consumers make decisions in the real data, including potential mistakes. In my setting, I assume that behavior according to the estimated (perceived) indirect utility function is how consumer behave in the naturally occurring domain. Instead, I assume that the welfare-relevant domain is that in which consumers do not care about price partitioning but only consider the total price, i.e. $\theta = \gamma_f = 0$. Therefore, I argue that consumers do not receive any true utility from how the total price is divided into shipping fee and product price. I assume that all other parameters are the same in both domains. For the welfare calculations, I assume that consumers choose according to what I call perceived utility while the consumer surplus they experience is based on what I call the welfare-relevant utility. Another way to describe these assumptions is that I assume that consumers would optimize perfectly if there was no partitioned pricing.[7]

While, ideally, the researcher would want to analyze choices in both domains, often, only choices under the naturally occurring domain are observable, as is also the case in my setting. Therefore, I use the demand estimates to assess counterfactual choices in the welfare-relevant domain. This calculation allows quantification of the welfare impact of consumers' decision-making based on perceived instead of welfare-relevant utility.

7.1 Estimated Expected Loss in Consumer Welfare

Note that because of this discrepancy between what I call the perceived utility, which is relevant for consumers' choices, and the welfare-relevant utility, which is relevant for consumer surplus, the formulas typically used to calculate consumer surplus slightly change. For cases in which the welfare-relevant and perceived utility are the same, Small and Rosen (1981) show that the expected consumer surplus takes on an analytical form in the logit case, known as the log-sum:

$$E(CS) = \frac{1}{\beta} E(max_j W_j + \epsilon_j) = \frac{1}{\beta} ln \left(\sum_j e^{W_j} \right) , \tag{7.1}$$

where β is the estimated income coefficient, W_j is the deterministic part of the indirect utility, and ϵ_j is extreme value type I distributed.[8]

However, if the welfare-relevant and perceived utilities are not equal, then the proof in Small and Rosen (1981) no longer holds. The reason is that the choice probability now depends on the perceived utility while the consumer surplus from each choice depends on the

[7] This assumption is closely related to that of Taubinsky and Rees-Jones (2018), who assume that the welfare-relevant domain is that without taxes.

[8] For this exposition, I abstract from the unobserved heterogeneity in β. However, the results include the estimated heterogeneity. The formulas discussed in this section are readily extended to the mixed logit by integrating over the distribution of β. For the implementation, this integration requires simulating the expected consumer surplus.

welfare-relevant utility. Train (2015) shows that, in cases like these, a term can be added to account for this discrepancy. I outline this approach here.

Let U_{ij} be the welfare-relevant utility that a consumer i receives from product j and let $\tilde{U}_{ij}$ be the perceived utility. Define the difference between the two as $d_{ij} = U_{ij} - \tilde{U}_{ij}$. Let $ij*$ be the alternative that the consumer chooses based on $\tilde{U}_{ij}$. Let $ik*$ denote the alternative that the consumer would have chosen based on U_{ij}. Note that if $ij* = ik*$, then consumer i incurs no loss in consumer surplus from deciding based on the perceived utility. Further note that using the log-sum in equation (7.1), I could calculate both $E(U_{ik*})$, i.e. the expected welfare-relevant utility if choosing based on U_{ij}, as well as $E(\tilde{U}_{ij*})$, i.e. the expected *perceived* utility if choosing based on $\tilde{U}_{ij}$.

The problem is that I am interested in the expected true utility that consumer i obtains if choosing based on the perceived utility. Denote this value as $\tilde{CS}_i = \frac{1}{\beta}E(U_{ij*})$. Using the definition of d_{ij}, I rewrite this expression as

$$E(\tilde{CS}_i) = \frac{1}{\beta}E(\tilde{U}_{ij*} + d_{ij*}) = \frac{1}{\beta}\left[E(\tilde{U}_{ij*}) + E(d_{ij*})\right]. \tag{7.2}$$

First, consider $E(\tilde{U}_{ij*})$. This term is the expected *perceived* utility if choosing based on perceived utilities. As noted above, this expression can be evaluated with the regular log-sum expression in Equation (7.1) and using the perceived utility function because the same utility function applies for the choice probabilities as well as the consumer surplus calculation. Therefore, I know that $E(\tilde{U}_{ij*}) = ln\left(\sum_{ij} e^{\tilde{W}_{ij}}\right)$, where $\tilde{W}_{ij}$ denotes the deterministic part of the perceived indirect utility.

Next consider $E(d_{ij*})$. This expression denotes the expected difference between actual and perceived utility if consumer i chooses according to their perceived utility. This expectation can simply be evaluated as a weighted average of this utility discrepancy for each product, weighted by the product's choice probability based on the perceived utility. Therefore, I can write $E(d_{ij*}) = \sum_j P_{ij}d_{ij}$, where P_{ij} is the choice probability of product j based on the perceived indirect utility. d_{ij} is simply $U_{ij} - \tilde{U}_{ij}$ which can be calculated, given the data and estimated parameters.

Thus, the expected consumer surplus I am interested in can be calculated as

$$E(\tilde{CS}_i) = \frac{1}{\beta}\left[ln\left(\sum_{j\in S_i} e^{\tilde{W}_{ij}}\right) + \sum_{j\in S_i} P_{ij}d_{ij}\right]. \tag{7.3}$$

As stated above, I assume that the actual indirect utility differs from the perceived indirect utility described in Equation (2.3) only because $\theta = 0$ and $\gamma_f = 0$. Therefore, d_{ij} is given as

$$d_{ij} = U_{ij} - \tilde{U}_{ij} = -\beta\theta c_{ij} - \gamma_f \mathbb{I}(c_{ij} = 0)\,. \tag{7.4}$$

Let $E(CS_i)$ be the expected consumer surplus had consumer i chosen based on the welfare-relevant utility U_{ij}. Then I denote the loss in expected consumer surplus due to not using the welfare-relevant utility for decision-making as $\Delta CS_i = E(\tilde{CS}_i) - E(CS_i)$. Note that this is the expected loss in consumer surplus for any consumer who faces the same choice situation as consumer i. Then, I calculate the mean of this statistic for all observations in my sample to obtain the mean expected loss in consumer surplus.

Table 7: Mean expected loss in consumer surplus per transaction due to partitioned pricing

	Exit	Azul	FIFA 19	Spiderman	Pokemon	Duos
$\frac{1}{N}\sum_i \Delta CS_i$	-0.7027	-0.0464	-0.0042	-0.0478	-0.0485	-0.1875
$\frac{1}{N}\sum_i (\Delta CS_i / E(CS_i))$	0.0622	0.0016	0.0003	0.0017	0.0016	0.0015

Notes: $\Delta CS_i = E(\hat{CS}_i) - E(CS_i)$ is the expected loss in consumer surplus of consumer i due to not using the welfare-relevant utility for decision-making. $\Delta CS_i / E(CS_i)$ is that loss relative to the level of consumer surplus under rational decision-making. Numbers shown here are means over all consumers $i \in 1, ..., N$.

In Table 7, I show the mean expected loss per purchase that the consumers in the sample incur due to not choosing according to their welfare-relevant utility. The expected loss ranges from 0.4 to 70 cents per purchase across the different products. Relative to the consumer surplus under fully rational decision-making, this amounts to relative losses of less than one percent to up to six percent. The absolute values of these figures can be interpreted as the consumer welfare that the average consumer in the sample would gain if eBay were to implement measures to ensure that consumers react identically to shipping fees and product prices. Such measures could include, for example, displaying the total price in the search results or removing the option for sellers to set a separate shipping fee.

7.2 Additional Considerations

The welfare impact of θ is reduced by three factors in this setting. First, the average share of the shipping fee in the total price is relatively low for most product categories. Higher fees might result in a larger welfare impact.[9] Second, the demand estimates suggest that consumers are actually quite attentive to the shipping fee at least in the video game categories. Finally, the discontinuous demand increase from free shipping counteracts a potential impact of limited attention to the shipping fee. One way to illustrate this idea is to calculate the mean expected welfare loss for different counterfactual scenarios in which those sellers that set a positive shipping fee in the data set it at an exogenously given amount. I let those sellers that set free shipping originally continue to have free shipping. The total prices remain unchanged. Table 8 shows the results of such an exercise.

As Table 8 illustrates, the welfare impact of the shipping fee is not monotonously increasing as it would be if the only bias was limited attention to the shipping fee. Rather, for some product categories, the welfare loss from partly ignoring the fee as well as having a preference for free shipping increases with higher fees, while for other product categories it decreases with higher fees. For each product category, one can calculate a level of the shipping fee at which the mean expected loss in consumer surplus from deviating from the welfare-relevant utility is zero. At these values, the choices made with the perceived utility are closest to those in the fully rational scenario on average.

These minimum-loss shipping fee levels depend on the proportion of the shipping fee coefficient γ_f and the inattention parameter θ. Intuitively, given that $\gamma_f > 0$, if a seller decides to move from a shipping fee of zero to a shipping fee of one cent while keeping the total price constant, they would incur a discontinuously large drop in demand. However, given that $\theta \in (0, 1)$, the seller could now increase the shipping fee, while decreasing the product price by the same amount, keeping the total price constant. This would then increase demand again. If the seller increases the shipping fee far enough, they can offset the loss of the free shipping premium. This level of the shipping fee, at which the average consumers are, *ceteris paribus*, indifferent between a listing with free shipping and a listing with the given shipping fee, can be calculated as $\frac{\gamma_f}{\theta}$. Table 9 shows this indifference shipping fee for the different products.

[9]It is unlikely though that the linear nature of θ assumed in this paper would hold with large variations in the fee. For example, Morrison and Taubinsky (2023) show that consumers pay more attention to hidden taxes with higher stakes.

Table 9: Indifference shipping fee levels

	Exit	Azul	FIFA 19	Spiderman	Pokemon	Duos
Indifference shipping fee	26.15	1.39	-5.25	33.42	15.38	31.49

Notes: Shipping fee levels at which consumers were indifferent between a listing with free shipping and a listing with this shipping fee, all else equal.

At these indifference shipping fees, the mean expected welfare losses as calculated in Table 8 would be lowest. "FIFA 19" is an exception as the estimates suggest a negative effect of free shipping for this product category. Therefore, the indifference shipping fee is negative. The mean expected loss is lowest at these values of the shipping fee because the two sources of bias γ_f and θ cancel each other out and, thus, consumers ignore the partitioned pricing in their decision-making in expectation.

8 Conclusion

Prior research shows that consumers participating in auctions seem to pay limited attention to add-on fees (Morwitz et al., 1998; Hossain and Morgan, 2006; Brown et al., 2010). A similar effect is found for the reaction to non-salient taxes (Chetty et al., 2009; Taubinsky and Rees-Jones, 2018). However, the consumer welfare consequences of such behavioral reactions to partitioned pricing have been largely unexplored.

My paper provides a quantification of the welfare calculations of partitioned pricing in the context of posted price transactions online. More specifically, I consider the example of the splitting of prices for goods on eBay into a product price and a shipping fee. For the analysis, I also include a discontinuous effect of free shipping in addition to an over- or under-reaction to marginal changes in the shipping fee. Such a discontinuity is consistent with the results of Shampanier et al. (2007), who show that a price of zero has a discontinuously positive demand effect. Including this discontinuity is important to make correct assessments of the impact that limited attention has on consumer welfare in situations where consumers can choose from listings with free shipping as well as listings with different levels of shipping fees.

To conduct my analysis, I web scrape publicly available transactions data for different products from eBay Germany. To obtain a measure for consumer surplus, I propose an empirical discrete choice model that can be interpreted within a theoretical framework on limited attention suggested by DellaVigna (2009). Using the estimates of the discrete choice model, I apply the framework of Bernheim and Rangel (2009) and a method described in Train (2015) to calculate the impact such behavioral patterns have on consumer welfare.

My main results suggest that the degree to which consumers ignore the shipping fee depends on the product category. For board games as well as the smart phone category, consumer choose as if they ignore 17 to 53 percent of the fee. For video games instead, consumers behave as if they only ignore five to ten percent of the shipping fee. My results also suggest that consumer demand reacts discontinuously positively if a listing features free shipping. This is a result that past research could not capture separately because they did not analyze listings with free shipping (Hossain and Morgan, 2006; Brown et al., 2010) or did not have sufficient variation in the add-on fee (Morwitz et al., 1998).

The behavioral patterns identified in the data suggest average losses in consumer surplus not larger than six percent of the absolute level of consumer surplus under rational decision making. Three main factors attenuate the welfare impact: First, for most product categories, the share of the fee in the total price is quite low on average. Second, the estimates suggest that many consumers are actually quite attentive to the fee. Third, the positive demand effect of free shipping partly counteracts the under-reaction to shipping fees in expectation.

When considering the policy implications of the results, one needs to keep in mind that the analysis is conducted entirely from the perspective of consumers. For an evaluation of whether a policy to regulate partitioned pricing might be necessary, total welfare needs to be considered. The losses in consumer welfare estimated here would then be the potential benefit from such a regulation. This benefit needs to be compared to the costs of such a policy which would likely lie with the platform and/or the sellers.

A caveat of my paper is that all interpretations and welfare statements in my paper are conditional on actually purchasing on eBay. This implies that my welfare calculations do not include a potential expansion or contraction of the eBay market size due to changes in the transparency of the shipping fees. By analyzing different products in different product categories and price ranges, I am able to increase the scope of external validity compared to the previous literature. However, in order to make well-founded statements about the general population, further research is needed.

Another interesting question that these results on consumer reaction to partitioned pricing raise is whether or not sellers are aware of this behavior and optimize accordingly. As I am only considering a demand model, this question is outside the scope of this paper and I leave it for future research.

References

Baye, M. R., Gatti, J. R. J., Kattuman, P., and Morgan, J. (2009). Clicks, discontinuities, and firm demand online. *Journal of Economics and Management Strategy*, 18(4):935–975.

Bernheim, B. D. and Rangel, A. (2009). Beyond Revealed Preference: Choice-Theoretic Foundations for Behavioral Welfare Economics. *The Quarterly Journal of Economics*, 124(1):51–104.

Bernheim, B. D. and Taubinsky, D. (2018). Behavioral Public Economics. *Handbook of Behavioral Economics*, 1:381–516.

Berry, S., Levinsohn, J., and Pakes, A. (1995). Automobile Prices in Market Equilibrium. *Econometrica*, 63(4):841–890.

Blake, T., Moshary, S., Sweeney, K., and Tadelis, S. (2021). Price Salience and Product Choice. *Marketing Science*, 40(4):619–636.

Blake, T., Nosko, C., and Tadelis, S. (2016). Returns to Consumer Search: Evidence from eBay. *17th ACM Conference on Electronic Commerce*, pages 531–545.

Brown, J., Hossain, T., and Morgan, J. (2010). Shrouded Attributes and Information Suppression: Evidence from the Field. *The Quarterly Journal of Economics*, 125(2):859–876.

Chetty, R., Looney, A., and Kroft, K. (2009). Salience and Taxation : Theory and Evidence. *American Economic Review*, 99(4):1145–1177.

DellaVigna, S. (2009). Psychology and Economics: Evidence From the Field. *Journal of Economic Literature*, 47(2):315–372.

Dertwinkel-Kalt, M., Köster, M., and Sutter, M. (2020). To buy or not to buy? Price salience in an online shopping field experiment. *European Economic Review*, 130:103593.

Dinerstein, M., Einav, L., Levin, J., and Sundaresan, N. (2018). Consumer Price Search and Platform Design in Internet Commerce. *American Economic Review*, 108(7):1820–1859.

Einav, L., Farronato, C., Levin, J., and Sundaresan, N. (2018). Auctions versus Posted Prices in Online Markets. *Journal of Political Economy*, 126(1):178–215.

Einav, L., Kuchler, T., Levin, J., and Sundaresan, N. (2015). Assessing Sale Strategies in Online Markets Using Matched Listings. *American Economic Journal: Microeconomics*, 7(2):215–247.

Frischmann, T., Hinz, O., and Skiera, B. (2012). Retailers' Use of Shipping Cost Strategies: Free Shipping or Partitioned Prices? *International Journal of Electronic Commerce*, 16(3):65–88.

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

Goeree, M. S. (2008). Limited Information and Advertising in the U.S. Personal Computer Industry. *Econometrica*, 76(5):1017–1074.

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

Hasker, K. and Sickles, R. (2010). eBay in the Economic Literature: Analysis of an Auction Marketplace. *Review of Industrial Organization*, 37(1):3–42.

Heidhues, P., Koszegi, B., and Murooka, T. (2017). Inferior products and profitable deception. *Review of Economic Studies*, 84:323–356.

Hossain, T. and Morgan, J. (2006). ...Plus Shipping and Handling: Revenue (Non) Equivalence in Field Experiments on eBay. *Advances in Economic Analysis and Policy*, 6(2):1–27.

Johnen, J. (2020). Dynamic competition in deceptive markets. *RAND Journal of Economics*, 51:375–401.

Johnen, J. and Somogyi, R. (2024). Deceptive features on platforms. *The Economic Journal*, 134.

Lacetera, N., Pope, D. G., and Sydnor, J. R. (2012). Heuristic Thinking and Limited Attention in the Car Market. *The American Economic Review*, 102(5):2206–2236.

McFadden, D. and Train, K. (2000). Mixed mnl models for discrete response. *Journal of Applied Econometrics*, 15(5):447–470.

Morrison, W. and Taubinsky, D. (2023). Rules of thumb and attention elasticities: Evidence from under- and overreaction to taxes. *Review of Economics and Statistics*, 105:1110–1127.

Morwitz, V. G., Greenleaf, E. A., and Johnson, E. J. (1998). Divide and Prosper: Consumers' Reactions to Partitioned Prices. *Journal of Marketing Research*, 35(4):453–463.

Shampanier, K., Mazar, N., and Ariely, D. (2007). Zero as a Special Price: The True Value of Free Products. *Marketing Science*, 26(6):742–757.

Small, K. A. and Rosen, H. S. (1981). Applied Welfare Economics with Discrete Choice Models. *Econometrica*, 49(1):105–130.

Tadelis, S. (2016). Reputation and Feedback Systems in Online Platform Markets. *Annual Review of Economics*, 8:321–340.

Taubinsky, D. and Rees-Jones, A. (2018). Attention Variation and Welfare: Theory and Evidence from a Tax Salience Experiment. *Review of Economic Studies*, 85(4):2462–2496.

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

Voester, J., Ivens, B., and Leischnig, A. (2017). Partitioned pricing: review of the literature and directions for further research. *Review of Managerial Science*, 11(4):879–931.

A Appendix

A.1 Details on the Web Scraping Procedure

To ensure that I find all relevant listings for each product category, I let my web scraping program search for rather broad terms. More specifically, I search for "exit der versunkene schatz", "pegasus azul," "spiderman ps4," "fifa 19 ps4," "pokemon lets go," and "samsung galaxy j5 duos." I conduct the searches separately for active as well as finished listings. I do not restrict the search further, meaning that I also save auctions and products that are not in a new condition. However, I remove these from the sample afterwards. For each scraping iteration, I first search eBay for the respective search term and save all results that I find on the search results pages. After having saved all listings shown in the results, I load each individual listing page to save the details for each listing. I loop through the different search terms and infinitely repeat this without pause for the active listings. For the finished listings, I pause several days between each loop through the searches.

On eBay Germany, the exact time and date of transactions can be observed. There are two different ways of identifying successful transactions depending on the type of listings. For listings that sell exactly one unit of a product, transactions can be observed by searching only for finished listings. Figure 2 shows a screenshot of search results of finished listings on eBay Germany. A price in green indicates that a listing was sold while a black price indicates that a listing ended without having been purchased. For listings that have an inventory of products and sell multiple copies, transaction can already be observed while the listing is still active. On the page of the listing, if copies have already been sold, a link leads to a list of past transactions including exact date and time, price, and model of the product, if applicable. Figure 3 shows a screenshot from the page of a listing with an inventory of products. A click on "6 verkauft" (6 sold) opens a list of past transactions such as the one shown in Figure 4.

Figure 2: Search results for finished listings on eBay Germany. Prices in a green font indicate that a listing was sold and prices in black indicate that it was not.

Figure 3: An active listing with an inventory for sale. A click on "6 verkauft" (6 sold) opens a list of past transactions.

Mitgliedsname	Preis	Stückzahl	Kaufdatum
1***8 (22)	EUR 26,90	1	01.09.19 23:28:59 MESZ
k***_ (1)	EUR 26,90	1	08.08.19 04:22:51 MESZ
s***t (0)	EUR 26,90	1	07.08.19 11:29:42 MESZ
a***1 (172)	EUR 26,90	1	06.08.19 22:31:05 MESZ
2***2 (49)	EUR 26,90	1	04.08.19 10:39:22 MESZ
i***e (91)	EUR 26,90	1	30.07.19 06:55:35 MESZ

Figure 4: The list of past transactions for a listing on eBay Germany

A.2 Estimation Procedure for Random Consideration Sets

This section provides details about the estimation procedure for the logit estimation allowing for random variation in consideration sets. The discussion follows Goeree (2008) closely but adapts it to the setting of my paper.

Consider again the probability that consumer i chooses listing j specified in Equation (3.4):

$$P_{ij} = \int \sum_{C \in S_j} \prod_{l \in C} \pi_{il} \prod_{k \notin C} (1 - \pi_{ik}) P_{ij}(C, \beta_i) f(\beta_i) d\beta_i \,.$$

For this exposition, I abstract from the additional unobserved heterogeneity across consumers induced by the heterogeneity in β_i. In the actual estimation, I include this heterogeneity, which effectively requires an additional simulation layer. Following Goeree (2008), I specify π_{ij} as

$$\pi_{ij}(\theta_\pi) = \frac{exp(\kappa_{ij})}{1 + exp(\kappa_{ij})} \,,$$

where $\kappa_{ij} = \varphi + K\prime_{ij}\rho$. φ is a constant and K_{ij} contains a vector of characteristics that might be correlated with a consumers probability to consider a product.

This weighted sum in Equation (3.4) is an expectation over all possible subsets of the full choice set of i that contain listing j. The term $\prod_{l \in C} \pi_{il} \prod_{k \notin C}(1 - \pi_{ik})$ is the probability that a given consideration set C is realized for consumer i. For all listings in the consideration set, the probability of being considered (π_{ij}) is multiplied, while for all others, the probability of not being considered is used $(1 - \pi_{ij})$.

An analytical solution for $\sum_{C \in S_j} \prod_{l \in C} \pi_{il} \prod_{k \notin C}(1 - \pi_{ik}) P_{ij}(C, \beta_i)$ exists. However, as Goeree (2008) notes, to calculate this expression analytically, for J choices, for each individual, $2^{(J-1)}$ different consideration sets would need to be considered. For ten choices, this already implies calculating consideration and choice probabilities for 512 different consideration sets for each individual and product. Thus, to limit computational burden, I simulate consideration sets similar to Goeree (2008).

I follow the following steps for the estimation:

1. Before starting the estimation:

    (a) For each individual i and available choice j draw R draws from a uniform distribution. Denote the draw r for consumer i and choice j as u_{ijr}.

2. In the first iteration of the maximization algorithm:

(a) First, calculate the consideration probability π_{ij}^0 given initial parameter values for each consumer and choice.

(b) Next, for each draw r, define an indicator for consideration of a choice j by consumer i by

$$b_{ijr}^0 = \begin{cases} 1, & \text{if } \pi_{ij}^0 > u_{ijr} \\ 0, & \text{otherwise} \end{cases} .$$

This binary variable fixes the simulated consideration set. Denote this consideration set as C_{ir}. Calculate the probability of this consideration set given the initial parameter values as $\Pi_{ir}^0 = \prod_{l \in C_{ir}} \pi_{il}^0 \prod_{k \notin C_{ir}} (1 - \pi_{ik}^0)$. The consideration set remains fixed for the next iterations to reduce variance.

3. In each step s of the maximization algorithm:

(a) Given the set of parameters, first calculate the consideration probability π_{ij}^s for each consumer and choice.

(b) Then, given the consideration sets determined in the initial step, I calculate the simulated choice probability for consumer i, listing j, draw r, and iteration s as

$$P_{ijrs} = \prod_{l \in C_{ir}} \pi_{il}^s \prod_{k \notin C_{ir}} (1 - \pi_{ik}^s) \frac{exp(W(X_ij, \Theta_s))}{\sum_{k \in C_{ij}} exp(W(X_{ik}, \Theta_s))} \frac{1}{\Pi_{ir}^0} .$$

The weight $\frac{1}{\Pi_{ir}^0}$ accounts for the fact that I fixed the consideration set based on the distribution of consideration sets based on the initial parameter values.

(c) For each individual i and listing j in estimation step s, the simulated choice probability is then

$$\hat{P}_{ijs} = \frac{1}{R} \sum_r P_{ijrs} .$$

Table 6: Demand model estimates with random consideration

	Exit	Azul	FIFA 19	Spiderman	Pokemon	Duos
Inattention (θ)	1.693**	0.464	-0.016	0.010	0.000	0.276
	[0.220, 3.167]	[-0.783, 1.712]	[-0.087, 0.054]	[-0.370, 0.390]	[-0.000, 0.000]	[-0.175, 0.726]
Free shipping effect (γ_f)	3.715**	0.961	-0.126	1.831***	22.707***	1.889***
	[0.190, 7.241]	[-3.412, 5.333]	[-0.307, 0.055]	[0.791, 2.872]	[19.313, 26.100]	[0.594, 3.185]
Price coefficient: Mean ($-\beta$)	-0.366**	-1.201***	-0.420***	-1.141***	-16.181***	-0.323***
	[-0.702, -0.030]	[-2.005, -0.398]	[-0.440, -0.400]	[-1.421, -0.861]	[-16.284, -16.078]	[-0.397, -0.250]
Price coefficient: SD (σ_β)	0.000	0.521***	0.137***	0.324***	3.829***	0.143***
	[-8.131, 8.132]	[0.288, 0.755]	[0.128, 0.145]	[0.274, 0.375]	[3.828, 3.829]	[0.128, 0.159]
Shipping fee coefficient ($-\beta\theta$)	0.620	0.558	-0.007	0.011	0.000***	0.089
	[-0.241, 1.482]	[-0.525, 1.641]	[-0.052, 0.039]	[-0.227, 0.250]	[0.000, 0.000]	[-0.137, 0.316]
Free shipping coefficient	3.715**	0.961	-0.126	1.831***	22.707***	1.889***
	[0.190, 7.241]	[-3.412, 5.333]	[-0.307, 0.055]	[0.791, 2.872]	[19.313, 26.100]	[0.594, 3.185]
Commercial seller	-2.274	3.274	0.421***	1.583***	33.790***	0.998*
	[-5.669, 1.121]	[-0.908, 7.457]	[0.200, 0.642]	[0.444, 2.723]	[31.132, 36.449]	[-0.098, 2.093]
Seller score (K)	0.004***	0.002	0.001***	-0.003***	0.021***	-0.001**
	[0.001, 0.006]	[-0.001, 0.006]	[0.001, 0.002]	[-0.005, -0.001]	[0.017, 0.026]	[-0.002, -0.000]
Multiple units	-0.564	-0.017	1.777***	3.369***	28.443***	0.926**
	[-1.943, 0.816]	[-1.883, 1.848]	[1.636, 1.918]	[2.316, 4.423]	[26.134, 30.751]	[0.166, 1.687]
Payment: Paypal			0.948***	1.387**	25.895***	4.265***
			[0.690, 1.206]	[0.263, 2.511]	[22.423, 29.367]	[2.852, 5.677]
π: Rank on results page	-0.011*	-0.001	-0.003***	-0.000	-0.000	-0.000
	[-0.024, 0.002]	[-0.003, 0.002]	[-0.003, -0.003]	[-0.001, 0.000]	[-0.000, 0.000]	[-0.001, 0.000]
π: Choice set size	-0.022	-0.002*	-0.006***	0.000	-0.000	-0.000
	[-0.114, 0.069]	[-0.005, 0.000]	[-0.007, -0.005]	[-0.002, 0.002]	[-0.000, 0.000]	[-0.006, 0.006]
π: Results page	-0.386	-0.009	-0.052***	-0.003	-0.000	-0.005
	[-1.438, 0.665]	[-0.141, 0.122]	[-0.065, -0.039]	[-0.013, 0.007]	[-0.000, 0.000]	[-0.021, 0.011]
π: Constant	0.857	0.065	0.292***	0.016	0.000	0.021
	[-0.225, 1.939]	[-0.169, 0.300]	[0.260, 0.324]	[-0.033, 0.066]	[-0.000, 0.000]	[-0.105, 0.147]
New edition		2.285				
		[-0.646, 5.216]				
Pokeball bundle					388.894***	
					[383.186, 394.601]	
Eevee edition					-1.746***	
					[-2.397, -1.095]	
Gold						0.164
						[-0.217, 0.545]
Blue						-0.201
						[-0.713, 0.311]
Draws	100	100	100	100	100.0	100
Observations	444	1347	34642	15488	41289.0	13631
Transactions	59	117	1918	940	1396.0	744
Unique listings	19	47	454	179	535	158
McFadden's R^2	0.25	0.52	0.46	0.70	0.45	0.58

Notes: ***, **, and * indicate statistical significance at the one, five, and ten percent levels, respectively. The brackets show 95 percent confidence intervals. θ is calculated as $\theta = \frac{\tilde{\theta}}{\beta}$. Standard errors for θ are calculated using the Delta method.

Table 8: Mean expected loss in consumer surplus at different shipping fee values

	Exit	Azul	FIFA 19	Spiderman	Pokemon	Duos
Shipping fee = 1	-0.7725	-0.0014	-0.0024	-0.053	-0.0626	-0.2459
Shipping fee = 2	-0.7128	-0.0034	-0.0033	-0.0502	-0.0548	-0.2309
Shipping fee = 3	-0.6546	-0.023	-0.0042	-0.0474	-0.0473	-0.2164
Shipping fee = 4	-0.5982	-0.0588	-0.0053	-0.0447	-0.0404	-0.2022
Shipping fee = 5	-0.5436	-0.1087	-0.0065	-0.042	-0.0339	-0.1884
Shipping fee = 6	-0.4912	-0.1702	-0.0079	-0.0394	-0.028	-0.1751
Shipping fee = 7	-0.441	-0.2405	-0.0093	-0.0369	-0.0225	-0.1622
Shipping fee = 8	-0.3933	-0.3171	-0.0109	-0.0344	-0.0176	-0.1497
Shipping fee = 9	-0.3481	-0.3978	-0.0126	-0.032	-0.0133	-0.1377
Shipping fee = 10	-0.3056	-0.4813	-0.0145	-0.0297	-0.0095	-0.1262

Notes: Consumer welfare losses shown are losses incurred due to partitioned pricing in scenarios in which sellers that do not offer free shipping set their shipping fee at an exogenously given fee while keeping the total price constant. If the total price is smaller than the imposed fee, the entire price is shifted to the shipping fee. The values shown are the means of ΔCS_i, i.e. the mean absolute loss in consumer surplus due to not using the welfare-relevant utility for decision-making.