Partitioned Pricing and Consumer Welfare

Kevin Ducbao Tran[1]

DIW Berlin, Technische Universität Berlin, and Berlin School of Economics

ktran@diw.de

[Job Market Paper. Please find the most current version *here*.[2]]

November 14, 2019

Abstract

In online commerce, obfuscation strategies by sellers are hypothesized to mislead consumers to their detriment and to the profit of sellers. One such obfuscation strategy is partitioned pricing in which the price is split into a base price and add-on fees. While empirical evidence suggests that partitioned pricing impacts consumer decisions through salience effects, its consumer welfare consequences are largely unexplored. Therefore, I provide a quantification of the welfare impact of the behavioral response to partitioned pricing. To do so, I derive a discrete choice model that jointly allows for differences in the reaction to marginal changes in add-on fees and the base price as well as a discontinuous effect of a zero fee. The model is based on a framework on limited attention and I estimate it using web scraped data of posted price transactions on eBay Germany. My results suggest under-reaction to marginal changes in the shipping fee, consistent with previous results in the literature. However, I also document a discontinuous positive effect of free shipping on consumer demand, which is novel to the literature. The combined impact of these effects on consumer welfare is small, less than six percent of consumer surplus. This small welfare impact occurs because the maximum shipping fee on eBay is capped and the free shipping effect partly counteracts the under-reaction to shipping fees in expectation.

[1] I thank Tomaso Duso and Hannes Ullrich for all of their invaluable support and advice. I am also grateful to Aaron Bodoh-Creed, Stefano DellaVigna, Ruben Durante, Paul Heidhues, Vicente Lagos, Christian Peukert, Imke Reimers, Maximilian Schäfer, Renke Schmacker, 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 in 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, and the 17th Annual International Industrial Organization Conference (IIOC) for useful comments and suggestions. I thank Adam Lederer for editorial support. All remaining errors are my own.

[2] The most current version of the paper is available at `http://jmp.kevintran.eu`.

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). Classical theory predicts rational consumers will 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), it is unclear whether consumers exhibit a 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 based on 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.[3]

Further, the welfare impact of partitioned pricing on consumer welfare in the online context has been 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.

Starting from a framework of limited attention proposed by DellaVigna (2009), I derive 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. Through transformations of the estimated parameters, I then recover the deep behavioral parameters of the DellaVigna (2009) framework.

To obtain the data necessary for the analysis, I automatically web scrape active listings on eBay Germany for various products 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.

[3] 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, for most products, consumers indeed behave as if they ignore part of the shipping fee if it is positive. For some products, consumers behave as if they do not consider the shipping fee at all, while for others they show only partial reaction to the shipping fee. For some products, however, the results suggest that consumers behave rationally with regard to marginal changes in the shipping 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 40 to 100 percent for the majority of the products 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.

In extensions of the base model, I allow for unobserved heterogeneity in consumers' consideration sets, potential endogeneity of the listing price, and unobserved heterogeneity in consumers' price sensitivity. The results are robust for half of the analyzed products, while for the others the results are affected by some of the extensions. For the welfare calculations, I use the results of the base model.

The welfare impact of these behavioral patterns is small. The relative loss in consumer surplus compared to fully rational behavior ranges from less than one to five percent. Two main reasons underlie this result: First, the size of the shipping fee is capped at 9.50 Euros by eBay. Welfare calculations show that if sellers were to charge higher shipping fees, the welfare loss could be higher. Second, the free shipping effect on demand partly offsets the under-reaction to shipping fees in expectation.

Finally, this paper adds to the 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. (2018) 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. Their setting 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 two price components are actually shown in different steps of the transaction process.

I proceed as follows. In Section 2, I present an empirical discrete choice model based on a framework proposed by DellaVigna (2009). 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 my base model and its extensions. In Section 7, I examine the welfare implications of my results. Section 8 concludes.

2 Model

In this section, I derive my empirical model based on the theoretical framework proposed by DellaVigna (2009). I then present three extensions to the base model accounting for unobserved consideration set heterogeneity, price endogeneity, and unobserved heterogeneity in price sensitivity.

2.1 The Base Model

DellaVigna (2009) proposes a framework to analyze what he terms "limited attention." I build on this framework to derive my econometric model. Assume the value that consumer i receives from good j is given by:

$$V_{ij} = v_{ij} + o_{ij} \, ,$$

where v_{ij} is a visible component and o_{ij} is an opaque component. If the consumer perceives o_{ij} differently than v_{ij}, let the *perceived* value of the good be denoted as:

$$\hat{V}_{ij} = v_{ij} + (1 - \theta) o_{ij} \, . \tag{2.1}$$

Applying this framework to posted prices and shipping fees, v_{ij} can be interpreted as the value of a good while $o_{ij} \equiv -c_{ij}$ represents the shipping fees. Considering the findings of Shampanier et al. (2007) and Einav et al. (2015), I also allow for a discontinuous effect of free shipping on consumer's 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.

Let the consumer's willingness-to-pay net of the shipping fees be denoted as

$$\hat{V}_{ij} = v_{ij} - (1-\theta)c_{ij} + \gamma_f f_{ij}\,, \tag{2.2}$$

where $f_{ij} \equiv \mathbb{1}(c_{ij} = 0)$. Because I am focusing on posted price transactions, no measure of willingness-to-pay is as readily available as in the case of auctions. Therefore I use a discrete choice framework to estimate the parameters θ and γ_f.

The perceived consumer surplus that consumer i receives from buying good j at price p_{ij} in DellaVigna (2009)'s framework is given by

$$\hat{CS}_{ij} = \hat{V}_{ij} - p_{ij} = v_{ij} - (1-\theta)c_{ij} + \gamma_f f_{ij} - p_{ij}\,.$$

Assuming that utility is linear in income,[4] the (conditional) indirect utility observable to the econometrician is given as

$$W(\hat{CS}_{ij}) = \beta\hat{CS}_{ij} = \beta\left[v_{ij} - (1-\theta)c_{ij} + \gamma_f f_{ij} - p_{ij}\right]\,. \tag{2.3}$$

I let $v_{ij} \equiv x'_{ij}\gamma$ be a linear function of observable non-financial product characteristics x_{ij} and corresponding coefficients γ. I discuss the selection of non-financial product characteristics in Section 3.1. Further, define the total price as $tp_{ij} = p_{ij} + c_{ij}$. Then the consumer's observable perceived indirect utility is

$$W(x_{ij}, tp_{ij}, c_{ij}; \delta, \beta, \theta) = x'_{ij}\gamma\beta - \beta tp_{ij} + \beta\theta c_{ij} + \beta\gamma_f f_{ij}\,.$$

Next, let ϵ_{ij} be the part of the utility that is observable only to the consumer and not to

[4] As Taubinsky and Rees-Jones (2018) discuss, the assumption that utility is linear in income is non-problematic for products whose prices are small relative to income. Another interpretation of the assumption is that around small value changes, any utility function can be approximated by a linear function.

the econometrician. Included this error term and rewriting the estimated parameters, the perceived indirect utility is given as

$$U(x_{ij}, tp_{ij}, c_{ij}, f_{ij}; \tilde{\gamma}, \tilde{\gamma}_f, \tilde{\beta}, \tilde{\theta}) = x'_{ij}\tilde{\gamma} + \tilde{\beta}tp_{ij} + \tilde{\theta}c_{ij} + \tilde{\gamma}_f f_{ij} + \epsilon_{ij}\,. \tag{2.4}$$

ϵ_{ij} could be non-zero, for example, because of differences in search behavior or distractions during the purchasing process.

Finally, I assume that consumers maximize their utility by choosing the one product in their choice set that yields the highest utility. Thus, to be consistent with this assumption, I exclude observations in which consumers buy multiple units of a product. These occasions are, however, rare. Assuming that ϵ_{ij} is extreme value type I distributed allows making use of the analytical logit choice probabilities following McFadden (1974):

$$P_{ij} = \frac{exp(W(X_{ij}, \Theta))}{\sum_{k \in S_i} exp(W(X_{ik}, \Theta))}\,, \tag{2.5}$$

where $W(X_{ij}, \Theta) = x'_{ij}\tilde{\gamma} + \tilde{\beta}tp_{ij} + \tilde{\theta}c_{ij} + \tilde{\gamma}_f f_{ij}$ and S_i is the choice set that consumer i is facing. Using these choice probabilities, I estimate the model parameters using maximum likelihood estimation. The original parameters of the DellaVigna (2009) model can then be retrieved from the estimated parameters by noting that $\beta \equiv -\tilde{\beta}$, $\theta \equiv -\frac{\tilde{\theta}}{\tilde{\beta}}$, and $\gamma_f \equiv -\frac{\tilde{\gamma}_f}{\tilde{\beta}}$.

DellaVigna (2009) assumes that $\theta \in [0, 1]$ and interprets it as the "inattention" parameter. For fully attentive consumers, $\theta = 0$, while for fully myopic 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 v_{ij} and o_{ij} (Taubinsky and Rees-Jones, 2018).[5]

Past evidence, however, suggests a θ between zero and one (e.g. Morwitz et al., 1998; Hossain and Morgan, 2006; Einav et al., 2015). Furthermore, research on the effects of a zero price suggest that $\gamma_f > 0$.

[5] 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.

2.2 Extensions of the Base Model

I now consider several extensions to the base model presented in Equation (2.4). In particular, I propose three extensions: I allow for consumers to randomly not consider all products in their choice sets, I use a control function approach to allow for endogeneity of the total price, and I allow for unobserved consumer heterogeneity in consumers' price sensitivity.

2.2.1 Unobserved Consideration Set Heterogeneity

One assumption I make in the base model is that all consumers consider all choices available on eBay at the time of purchase. This full information assumption is typical for discrete choice models. However, the number of choices can be large in the eBay setting. Figure 1 shows the number of choices for each of the analyzed products over time. At some points, for some of the products, more than one hundred choices were available. Therefore, one concern is that the full information assumption is unrealistic in this setting.

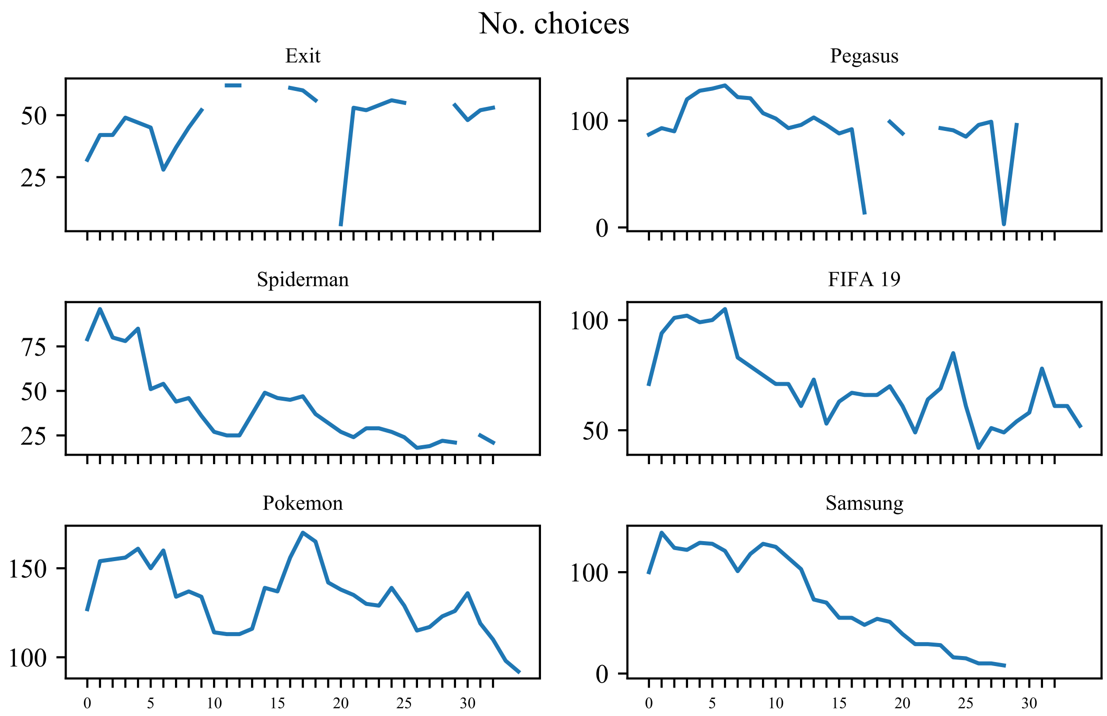

Figure 1: Average choice set size by week. Gaps indicate weeks during which no sale was observed.

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.

While being a limitation, the crucial issue for the estimation of θ and γ_f is whether the probability of consideration is correlated with the total price and the shipping fee. In other words, 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 that my web scraper sees and a listing's price or shipping fee. If the rank is indeed not impacted by the shipping fee and prices, then the unobserved consideration set heterogeneity will increase the noise of the estimates but will not cause any bias.

Nevertheless, if, for example, some of the consumers sort by price or filter by free shipping, the probability of consideration might be correlated with the price or the shipping fee. To address this issue, I employ 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. A functional form for the consideration probabilities is then assumed and the parameters determining these probabilities are jointly estimated together with the utility parameters. The probability that consumer i purchases a product j now becomes

$$P_{ij} = \sum_{C \in S_j} \prod_{l \in C} \pi_{il} \prod_{k \notin C} (1 - \pi_{ik}) \frac{exp(W(X_ij, \Theta))}{\sum_{k \in C} exp(W(X_{ik}, \Theta))}, \tag{2.6}$$

where S_j is the set of all consideration sets that include choice j and π_{ij} is the probability that consumer i considers choice j. The π_{il} are functions of listing characteristics that impact the probability that a listing is considered. I estimate the model using simulation. For more details on the procedure, please refer to Goeree (2008) or Appendix A.5.

I include, among the characteristics affecting the consideration probability, 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, and the total price of the listing. 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.

Blake et al. (2016) further show that, over time, consumers are quickly able to narrow down their search results to listings with, on average, lower prices. Therefore, I include the total price to account for consumers sorting on the price. 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.

2.2.2 Price Endogeneity

A common concern in demand estimation is that the price of the product might be correlated with unobservable factors that also affect demand. If that is the case, then the estimated price coefficient is biased. In my setting, I expect this problem to be less of an issue. I purposefully chose the products 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, should not be as much of 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.

Nevertheless, I present an extension here in which I use instruments in the spirit of Berry et al. (1995) to explain the total price in a first stage regression. However, using an instrumental variables approach in a non-linear estimation, such as the logit estimation, is not as well explored as for the linear case.

Berry et al. (1995) propose an approach that requires estimation of a set of listing fixed effects. Directly estimating the listing fixed effects is infeasible due to the large number of unique listings and varying choice sets. Berry et al. (1995) suggest a "contraction" procedure to calculate the fixed effects conditional on parameter values inside the estimation procedure. This removes the need to estimate fixed effects directly. However, this procedure does not work for choices that have choice probabilities of zero, which does occur in my data.

Instead, I apply the so-called control function approach suggested by Petrin and Train (2010). The approach requires additional assumptions on the structure of the price endogeneity. Then, the main idea is to condition on the source of the endogeneity such that the remaining error term in the indirect utility function is independent from the potentially endogenous variable by construction.

Applying the control function approach, the choice probability becomes

$$P_{ij} = \int \frac{exp(W(X_{ij}, \Theta) + \sigma_{CF}\eta_{ij})}{\sum_{k \in S_i} exp(W(X_{ik}, \Theta) + \sigma_{CF}\eta_{ik})} \prod_{l \in S_i} \phi(\eta_{il}) d\eta_{il}\,, \tag{2.7}$$

where $W(X_{ij}, \Theta) = x'_{ij}\tilde{\gamma} + \tilde{\beta}tp_{ij} + \tilde{\theta}c_{ij} + \tilde{\gamma}_f f_{ij} + \lambda\hat{\mu}_{ij}$ and $\phi(\eta)$ is the standard normal probability density function. $\hat{\mu}_{ij}$ is obtained from the residuals of a first stage price regression. Because this integral does not have an analytical solution, I use simulated choice probabilities for the estimation. For more details, please refer to Petrin and Train (2010) or Appendix A.6.

2.2.3 Unobserved Heterogeneity in Price Sensitivity

A common critique of the multinomial logit model is that its assumptions imply the so-called independence of irrelevant alternatives (IIA) (McFadden, 1974). Berry et al. (1995) propose the so-called mixed logit model to relax the IIA assumption. The mixed logit model allows some or all of the parameters in the indirect utility function to vary across consumers according to some mixing distribution. Because of this non-linearity in the indirect utility function, more flexible substitution patterns can be created.

In the eBay setting, in particular when focusing on listings that offer very homogeneous products, as I do, the IIA assumption seems less problematic. If all listings are very similar, it does indeed seem plausible that relative choice probabilities should not be affected by the exit or entry of any other listing. Nevertheless, I propose allowing for unobserved consumer heterogeneity with regard to price sensitivity β here. One interpretation for this is that consumers have different price sensitivities because of unobserved differences in income.[6]

Note that I assume unobserved consumer heterogeneity in the coefficient β of the model derived from the DellaVigna (2009) framework (see Equation (2.3)). This means that, in terms of the estimated parameters, the variation in β also translates to variation in all other estimated coefficients. To illustrate this, consider again Equation (2.4). Letting $\beta \equiv \beta_i$ vary across consumers, this equation becomes

$$U(x_{ij}, tp_{ij}, c_{ij}, f_{ij}; \tilde{\gamma}_i, \tilde{\gamma}_{fi}, \tilde{\beta}_i, \tilde{\theta}_i) = x'_{ij}\tilde{\gamma}_i + \tilde{\beta}_i tp_{ij} + \tilde{\theta}_i c_{ij} + \tilde{\gamma}_{fi} f_{ij} + \epsilon_{ij}\,, \tag{2.8}$$

where $\tilde{\gamma}_i = \beta_i\gamma$, $\tilde{\gamma}_{fi} = \beta_i\gamma_f$, $\tilde{\beta} = -\beta_i$, and $\tilde{\theta}_i = \beta_i\theta$. Because I assume that utility is linear in the monetary consumer surplus, the heterogeneity of the β_i transmits to all the other estimated coefficients as well. However, this is not equivalent to estimating a mixed logit model in which all coefficients 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 all other coefficients.

I assume that $\beta_i \sim N(\mu_\beta, \sigma^2_\beta)$. This in turn implies that $\tilde{\gamma}_i \sim N(\mu_\beta\gamma, \sigma^2_\beta\gamma^2)$, $\tilde{\gamma}_{fi} \sim N(\mu_\beta\gamma_f, \sigma^2_\beta\gamma^2_f)$, $\tilde{\beta}_i \sim N(-\mu_\beta, \sigma^2_\beta)$, and $\tilde{\theta}_i \sim N(\mu_\beta\theta, \sigma^2_\beta\theta^2)$. θ can therefore be recovered from the estimated coefficients as $\theta = \frac{\mu_{\tilde{\theta}}}{\mu_\beta}$, where $\mu_{\tilde{\theta}}$ is the estimated mean of the distribution of $\tilde{\theta}_i$. Equivalently, γ_f can be recovered as $\gamma_f = \frac{\mu_{\tilde{\gamma}_f}}{\mu_\beta}$. With the normal mixing distribution, the expected choice probability for each individual no longer has an analytical solution. Therefore, I approximate it using simulation.

3 Identification

Because, in contrast to previous literature, 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 δ_f. This section discusses these challenges and how I propose to overcome them.

[6] As discussed in Section 2, the assumption that the indirect utility is linear in income preference β can be seen as a linear approximation of non-linear income preferences around small price changes (Taubinsky and Rees-Jones, 2018). Here, the assumption of linearity in β remains, however, I now allow for the levels of β to differ across consumers.

3.1 v_{ij}: Making Choices Comparable

Consider again the expression for consumers' willingness-to-pay net of shipping fees presented in Equation (2.2). DellaVigna (2009) proposes to estimate θ by keeping v_{ij} constant while exogeneously varying c_{ij}. 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 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 keep v_{ij} exactly constant for all j. The structural assumptions on consumer decision making help overcome the issue of unobserved willingness-to-pay by assuming a functional form for it. However, the idea for identification of θ (and γ_f) remains the same except that I try to keep v_{ij} constant conditional on observable characteristics. Therefore, my identification of θ and γ_f relies on comparing products for which $v_{ij} = x_{ij}'\delta$ is 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 products are. Therefore, I chose to analyze products that are homogeneous, leaving the relevant variation in non-financial characteristics to 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, to fulfill the assumptions of the discrete choice model, I need to ensure that all relevant choices are included 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 can be seen as substitutes. Therefore, I need to include products for which a set of alternatives 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 chose 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 the beginning of January 2019 according to Amazon.de.

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 $v_{ij} = x_{ij}'\tilde{\gamma}$ to include all remaining relevant non-financial listing characteristics is important to obtain unbiased estimates of $\tilde{\theta}$, $\tilde{\gamma_f}$, and $\tilde{\beta}$. 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 6.1.

3.2 The Effect of Free Shipping

Note that few of the cited studies consider the discontinuous free shipping effect γ_f. Most studies only analyze θ. Morwitz et al. (1998) compare the no-fee case to only one level of fee and, thusly, could not 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, assume consumers' willingness-to-pay follows Equation (2.2). For clarity, Equation (2.2) can be rewritten as

$$\hat{V}_j = \begin{cases} v_j - (1-\theta)c_j, & \text{if } c_j > 0 \\ v_j + \gamma_f, & \text{otherwise} \end{cases} \tag{3.1}$$

where I omit the consumer suffix i for the exposition. Furthermore, let $\tilde{V}_j = v_j - (1-\theta)c_j$ be the functional form for willingness-to-pay ignoring a potential effect of γ_f.

Now consider two listings $j \in 1,2$ for which $c_2 > c_1 > 0$ and $v_1 = v_2 = v$. This representation corresponds to the "High Reserve Treatments" in Hossain and Morgan (2006) as well as the treatments in Brown et al. (2010). Let V_j be the observed willingness-to-pay under regime j. θ is identified using Equation (3.1) as

$$
\begin{aligned}
V_1 &= v - (1-\theta)c_1 \\
V_2 &= v - (1-\theta)c_2 \\
\Leftrightarrow \theta &= 1 - \frac{V_1 - V_2}{c_2 - c_1}\,.
\end{aligned}
$$

Note that, in this situation, θ can be correctly identified from two non-zero values of c_j, even if falsely using $\tilde{V}_j$ because γ_f is irrelevant for $c_j > 0$ and, therefore, $\tilde{V}_j = \hat{V}_j$. This is exactly what Hossain and Morgan (2006) and Brown et al. (2010) do by considering treatments with different non-zero shipping fee listings.

However, it is not possible to identify both θ and γ_f separately using only two different treatments. Consider the two treatments $c_1 = 0$ and $c_2 > 0$. This representation corresponds to, for example, the treatments in Morwitz et al. (1998) and in the "Low Reserve Treatments" of Hossain and Morgan (2006). Let $\tilde{\theta}$ be the inattention parameter obtained from using $\tilde{V}_j$. Ignoring γ_f and using $\tilde{V}_j$, one would identify $\tilde{\theta}$ as $\tilde{\theta} = 1 - \frac{V_1 - V_2}{c_2}$. However, if actually consumer utility took the form of Equation (3.1), this would imply that

$$
\begin{aligned}
V_1 &= v + \gamma_f \\
V_2 &= v - (1-\theta)c_2 \\
\Leftrightarrow \theta - \frac{\gamma_f}{c_2} &= 1 - \frac{V_1 - V_2}{c_2} = \tilde{\theta}\,.
\end{aligned}
$$

This shows that, first, using just these two treatments, θ and γ_f cannot be disentangled and second, ignoring γ_f and using $\tilde{V}_j$ results in a biased estimate of θ. In particular, if $\gamma_f > 0$, $\tilde{\theta} < \theta$ 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 estimation of θ which subsequently would affect the welfare calculations.

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 various products. 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 suggests 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 copies. 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 that sell their product on eBay.

Sellers can also choose whether to offer free shipping or set a shipping fee for their listings. For many product categories, eBay Germany caps the shipping fee at 9.50 Euros for national shipping.

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 sum of positive reviews minus the sum 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.

4.2 Choice Set Creation

For the estimation of the empirical model, I collect publicly available data on listings on eBay Germany using a web scraper. The web scraper searches for a range of products and subsequently visits all the listing pages that are found as a result. The eBay website offers well-suited data availability 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

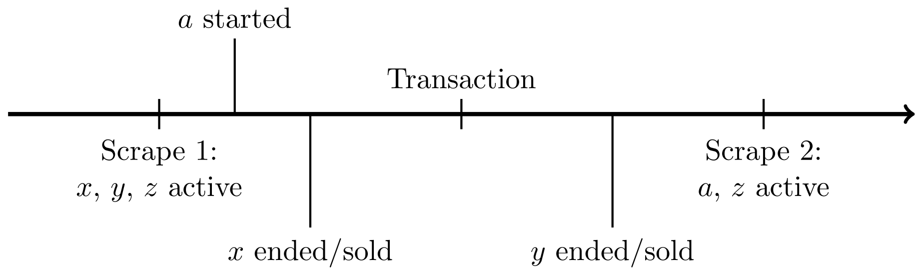

Figure 2: An illustration of the choice set reconstruction.

were facing. To do so, I searched eBay Germany for each product 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 (given sample restrictions discussed in Section 4.3). To reconstruct the potential choice set for each choice situation, I match all listings that I observed 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 2 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.[7] 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 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.

[7] The 48 hours tolerance window is chosen to strike a balance between being strict and allowing some flexibility for potential misses by the scraper. With about three to four scrapes a day, the chance that a listing is missed in each repetition should be fairly low.

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. Finally, 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.

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

5.1 Descriptive Statistics

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 were available over longer time periods enter the averages multiple times. I assume that each transaction represents one individual.

Table 1: Descriptive Statistics

	Exit	Azul	Spiderman	FIFA 19	Pokemon	Duos
Total price	16.40	41.93	48.49	49.30	69.09	179.90
	[14.89, 17.89]	[38.47, 44.89]	[37.90, 59.59]	[38.84, 61.95]	[49.99, 79.99]	[166.79, 183.99]
Product price	14.62	40.46	47.75	48.29	67.65	179.17
	[12.99, 15.99]	[37.89, 41.47]	[36.88, 59.59]	[36.46, 61.95]	[49.90, 77.49]	[166.78, 182.09]
Shipping fee	1.78	1.47	0.74	1.01	1.44	0.72
	[0.00, 3.99]	[0.00, 3.95]	[0.00, 1.45]	[0.00, 1.99]	[0.00, 2.95]	[0.00, 0.00]
Shipping fee (> 0)	4.46	4.95	2.74	3.11	3.72	3.92
	[3.00, 5.89]	[4.90, 5.50]	[1.99, 3.79]	[1.99, 4.10]	[2.00, 4.99]	[1.99, 4.99]
Share of shipping in total price	0.10	0.03	0.02	0.02	0.02	0.00
	[0.00, 0.23]	[0.00, 0.08]	[0.00, 0.03]	[0.00, 0.04]	[0.00, 0.05]	[0.00, 0.00]
Share of shipping in total price (> 0)	0.25	0.11	0.07	0.07	0.06	0.02
	[0.21, 0.30]	[0.09, 0.12]	[0.05, 0.08]	[0.04, 0.10]	[0.04, 0.07]	[0.01, 0.03]
Free shipping	0.60	0.70	0.73	0.67	0.61	0.82
	[0.00, 1.00]	[0.00, 1.00]	[0.00, 1.00]	[0.00, 1.00]	[0.00, 1.00]	[1.00, 1.00]
Pos. reviews (%)	98.61	98.33	99.40	98.67	99.27	99.23
	[99.50, 100.00]	[99.30, 100.00]	[99.40, 100.00]	[99.40, 100.00]	[99.40, 100.00]	[98.90, 99.80]
Seller score (K)	61.66	32.98	61.63	73.27	85.25	126.34
	[0.18, 21.98]	[0.18, 5.53]	[0.51, 36.78]	[0.49, 36.63]	[0.66, 45.93]	[4.50, 151.09]
Commercial seller	0.96	0.95	0.76	0.80	0.82	0.99
	[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	0.70	0.73	0.53	0.58	0.64	0.67
	[0.00, 1.00]	[0.00, 1.00]	[0.00, 1.00]	[0.00, 1.00]	[0.00, 1.00]	[0.00, 1.00]
Payment: Paypal	0.99	0.99	0.92	0.91	0.90	0.99
	[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		0.09				
		[0.00, 0.00]				
Pikachu edition					0.57	
					[0.00, 1.00]	
Pokeball bundle					0.26	
					[0.00, 1.00]	
Blue						0.19
						[0.00, 0.00]
Gold						0.27
						[0.00, 1.00]
N	2202	9374	36183	87648	120193	49555

Notes: Sample means with 25 and 75 quantiles in brackets. Each observation represents one listing in one choice situation.

Note that the products are sorted in ascending mean total and product price. However, the mean shipping fee does not increase proportionally with the mean product price, therefore the share of the shipping fee in the total price is decreasing with an increasing total price. While for "Exit," given that the fee is positive, it makes up about a quarter of the total price on average, for "Duos" that share is only two percent. A majority of the listings offer free shipping (i.e. do not used partitioned prices) with the shares ranging from 60 to 82 percent across the different products.

The two main indicators through which eBay tries to mitigate asymmetric information between buyers and sellers and increase trust 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, a 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 76 percent to 99 percent. The variables in the lower panel of table 1 are product-specific variables that I include for some of the products to control for different models and editions that were available. These variables are all indicator variables.

5.2 Shipping Fees and Prices

Table 1 shows that the share of listings with free shipping ranges from 60 to 82 percent across products. However, past research suggests that $\theta \in (0,1)$ and consumers do not react to changes in the add-on fee as much as they do to changes in the product price.[8] If this were the case, sellers that do set a positive shipping fee should be able to charge higher total prices with higher shipping fees, all else equal.

To assess whether such patterns can be observed, 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. In a market with homogeneous goods, perfect competition, and fully rational consumers, one would expect a one-to-one decrease in the product price for each additional Euro in the shipping fee, keeping the total price constant (at marginal cost). Although the market on eBay is most likely not perfectly competitive (e.g. due to search frictions) and the listings are not perfectly homogeneous, controlling for other covariates, one would not expect a change in the total price if the shipping fee changes.

However, the point estimates of the shipping fee for five of the six products are positive, although it is not statistically significantly different from zero for the "Duos" smart phone. Only for the "Azul" board game does the estimate suggest that the total price does not increase with higher shipping fees.

These results suggest that the total price of a listing, given that it does not offer free shipping, are, on average, indeed higher for higher shipping fees. This evidence is in line with at least some sellers trying to exploit a $\theta \in (0,1)$.

[8] For non-salient taxes, for example, Chetty et al. (2009) find $\theta = 0.75$ and Taubinsky and Rees-Jones (2018) find an average θ of 0.65.

Table 2: Estimated coefficients of a linear regression explaining total price

	Exit	Azul	Spiderman	FIFA 19	Pokemon	Duos
Dependent variable	Total price	Total price	Total price	Total price	Total price	Total price
Shipping fee	0.837***	-0.241	2.186*	1.821**	3.898***	2.179
	[0.511, 1.163]	[-1.465, 0.983]	[-0.185, 4.557]	[0.020, 3.622]	[1.740, 6.056]	[-0.511, 4.869]
No. individual-choice pairs	372	1434	879	13333	12306	2291
R^2	0.827	0.540	0.562	0.504	0.841	0.690

Notes: I restrict the sample to 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. The full estimation results are shown in Appendix A.2.

5.3 Correlates of the shipping fee

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 the shipping fee. Table 3 shows the main results of a linear regression of the free shipping indicator variable on various covariates.

Table 3: Estimated coefficients of a linear regression explaining free shipping

	Exit	Azul	Spiderman	FIFA 19	Pokemon	Duos
Dependent variable	Free shipping	Free shipping	Free shipping	Free shipping	Free shipping	Free shipping
Commercial seller	-0.020	-0.003	0.256**	0.505***	0.435***	0.657***
	[-0.390, 0.351]	[-0.235, 0.228]	[0.058, 0.454]	[0.360, 0.650]	[0.313, 0.557]	[0.420, 0.894]
Seller score (K)	-0.001***	-0.000	-0.001***	-0.001***	-0.001***	-0.000**
	[-0.001, -0.001]	[-0.000, 0.000]	[-0.001, -0.001]	[-0.001, -0.001]	[-0.001, -0.001]	[-0.000, 0.000]
Payment: Paypal	1.009***	0.738***	0.544***	0.149	0.400***	-0.477***
	[0.421, 1.597]	[0.385, 1.091]	[0.262, 0.826]	[-0.086, 0.384]	[0.210, 0.590]	[-0.826, -0.128]
No. individual-choice pairs	2202	9374	36183	87648	120193	49555
R^2	0.438	0.509	0.497	0.328	0.419	0.345

Notes: ***, **, and * indicate statistical significance at the one, five, and ten percent levels, respectively. The brackets show 95 percent confidence intervals. The full estimation results are shown in Appendix A.3.

A clear pattern that holds true for all products is difficult to determine, but the results suggest that, at least for the four most expensive products, commercial sellers are positively associated with free shipping. The linear probability regressions suggests that commercial sellers in these product categories have a 26 to 66 percentage point higher probability of setting a zero fee. At the same time, more experienced sellers, as measured by the eBay seller score, seem to be less likely to offer free shipping, but the effect is economically small at

0.1 percentage points or below. For five of the products, listings that also offer PayPal as a payment method are more likely to offer free shipping, with estimates suggesting a 15 to 100 percentage point increase in the free shipping probability. Only for the "Duos" smartphone is the correlation reversed. For other covariates, no clear patterns across all products can be seen.

A different perspective is to analyze shipping fees for listings with a positive shipping fee. Table 4 shows the main coefficients from such a 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 "Spiderman," "FIFA 19," and "Duos," this conditional correlation is the largest, suggesting 0.99 to 1.36 Euros higher shipping fees for listings with an inventory. If consumers value buying from listings with an inventory (e.g. because they seem more professional and trustworthy), this positive correlation might result in an upward bias of the shipping fee coefficient estimate ($\tilde{\theta}$) if an indicator for listings with an inventory is omitted. This bias in turn would directly translate into an upward bias of θ, suggesting a larger ignorance of shipping fees than there might be in reality.

Table 4: Regressions results of a linear regression explaining the shipping fee

	Exit	Azul	Spiderman	FIFA 19	Pokemon	Duos
Dependent variable	Shipping fee	Shipping fee	Shipping fee	Shipping fee	Shipping fee	Shipping fee
Multiple units	0.031	0.236	0.989*	1.318***	0.465**	1.375***
	[-0.781, 0.844]	[-0.262, 0.734]	[-0.017, 1.995]	[0.740, 1.896]	[0.057, 0.873]	[0.430, 2.320]
No. individual-choice pairs	877	2787	9834	28518	46457	9156
R^2	0.637	0.396	0.358	0.434	0.553	0.823

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. The full estimation results are shown in Appendix A.4.

6 Discrete Choice Estimation Results

This section discusses the selected covariates and provides the estimates from estimation of the base model specified in Equation (2.4) as well as its extensions. For clarity of exposition, I show only the estimates for the main coefficients of interest θ and γ_f in Table 5. The full estimation results are shown in Appendices A.7 to A.10.

6.1 Selected Covariates

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. Including the share of positive reviews does not impact the results. 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.

Table 3 further suggests that listings that accept PayPal as a payment method are also more likely to offer free shipping. I therefore also include an indicator variable for listings that accept PayPal. Further, as Table 4 indicates, including a fixed effect for listings with an inventory 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. An upward-biased shipping fee coefficient implies an upward biased θ. Therefore, I include an indicator for an inventory in my estimations.

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 indicators for different phone colors.

6.2 Results

Table 5 shows the estimation results from the base model and its three extensions, as discussed in Section 2, for all six products. The results are robust for all four models for three of the products: "Exit," "Spiderman," and "Duos." To varying degrees, the estimates of inattention, θ, suggest that consumers partly or fully ignore the shipping fee for these three products. In fact, the estimates for "Exit" and "Duos" are statistically not significantly different from one. $\theta = 1$ implies that consumers fully ignore the shipping fee. It should be noted, however, that the estimates are quite noisy in some specifications. The point estimates of θ for "Exit" range from 0.9 to 1.3, while the point estimates for "Duos" range from 0.6 to 1.7. Regardless, in all specifications, the point estimates of θ for both products are statistically significantly different from zero. Therefore, I can reject the hypothesis of full attention to shipping fees for the consumers of these two products.

For "Spiderman," the estimates suggest that consumers partially ignore the shipping fee but to a lesser extent than consumers buying "Exit" or "Duos." In fact, in the base model, the random consideration model, and the control function model, the null hypothesis of full attention ($\theta = 0$) cannot be rejected. Only in the specification with heterogeneous price sensitivity β_i, is the estimate precise enough to statistically distinguish it from zero. Thus, consumers buying the "Spiderman" video game seem to pay more attention to the shipping fees and might, in fact, behave rationally.

The estimates of γ_f are robustly positive for all three products in all four specifications. Indeed, these robust results suggest that there is a discontinuous effect of free shipping on utility.

My data does not allow me to make clear statements about the sources of the differences in the estimates of θ for the three products. 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 "Spiderman" video game 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. However, without additional data on consumers, these statements are only speculative.

The results for "Azul," "FIFA 19," and "Pokemon" are not robust for all extensions. The results for "Azul" in the base model and in the model including the control function suggest that consumers ignore large parts of the shipping fee. Further, the point estimates of γ_f suggest a positive effect of free shipping, although it is not statistically significantly different from zero. However, in the model with random consideration and heterogeneous price sensitivity, the estimates of θ are statistically not different from zero and the point estimates of γ_f are negative. A potential explanation for these differing results is that the random consideration model and the heterogeneous price sensitivity model require more variation in the data to be estimated because they both incorporate non-linearity in parameters in consumers' utility functions. Because of this problem, I focus my interpretation on the results of the base model and the model with the control function, keeping in mind that the results are not fully robust to all extensions.

The results for "FIFA 19" suggest that consumers are fully rational with regard to positive shipping fees $\theta = 0$ while the estimates suggest that free shipping might indeed have a negative effect, although the estimate is only precise enough in the base model to be statistically distinguished from zero. These results are robust for the base model, the random consideration model, and the control function model. The estimate of θ in the model with heterogeneous price sensitivity differs from the other three because it suggests that consumers only pay limited attention to the shipping fee. However, because the other three models are consistent, I focus my interpretation of the results on those results.

The results for "Pokemon" are consistent in the base model and the model with heterogeneous price sensitivity. Those estimates suggest that consumers partly ignore a positive shipping fee but react positively to free shipping. The estimates with random consideration still suggest a positive free shipping effect γ_f but the estimate of θ is negative. The estimates using the control function suggest an even larger (in absolute values) negative θ and a negative γ_f. For my interpretation, I focus on the base model because it is the least demanding on the data. However, especially for "Pokemon," the results should be interpreted with care because they are not robust to two of the extensions.

For the welfare calculations in Section 7, I make use of the results of the base model. For most products, these results are robust for most or all of the extensions. Further, using the base model for the welfare calculations avoids the need for simulation and allows for calculating the welfare impacts using analytical solutions. When interpreting the results, keep in mind that some of them are not robust to all extensions of the base model.

Table 5: Estimates of θ and γ_f for the base model and its extensions.

	Exit	Azul	Spiderman	FIFA 19	Pokemon	Duos
Base Model						
Inattention (θ)	1.065***	0.920**	0.177	0.039	0.390**	0.830**
	[0.335, 1.795]	[0.210, 1.630]	[-0.106, 0.460]	[-0.271, 0.349]	[0.045, 0.735]	[0.176, 1.484]
Free shipping effect (γ_f)	4.797***	1.789	2.264***	-2.649***	2.655***	8.250***
	[1.845, 7.749]	[-1.432, 5.010]	[1.067, 3.461]	[-3.845, -1.453]	[1.511, 3.799]	[4.188, 12.312]
Random Consideration						
Inattention (θ)	1.225***	-0.120	0.111	-0.145	-0.240**	0.720**
	[0.436, 2.014]	[-1.003, 0.763]	[-0.045, 0.267]	[-0.472, 0.182]	[-0.475, -0.005]	[0.132, 1.308]
Free shipping effect (γ_f)	5.299***	-1.866	1.869***	-0.797	1.073**	6.697***
	[2.184, 8.415]	[-5.430, 1.697]	[1.547, 2.191]	[-1.899, 0.305]	[0.180, 1.967]	[3.353, 10.041]
Control Function						
Inattention (θ)	1.282***	0.899***	0.043	-0.015	-0.932**	1.681**
	[0.110, 2.945]	[0.158, 1.447]	[-0.187, 0.283]	[-0.255, 0.225]	[-1.731, -0.131]	[0.585, 3.532]
Free shipping effect (γ_f)	7.236	1.689	1.305	-2.500	-13.487	12.939
	[.,.]	[.,.]	[.,.]	[.,.]	[.,.]	[.,.]
Heterogeneous Price Sensitivity						
Inattention (θ)	0.993***	0.000	0.246***	0.383***	0.408***	0.643***
	[0.509, 1.477]	[-0.011, 0.011]	[0.134, 0.358]	[0.205, 0.561]	[0.217, 0.599]	[0.159, 1.127]
Free shipping effect (γ_f)	4.590***	-1.750***	2.222***	-0.178	2.665***	6.553***
	[2.604, 6.576]	[-2.682, -0.817]	[1.859, 2.586]	[-0.860, 0.503]	[1.885, 3.444]	[3.650, 9.455]
No. individuals	53	108	933	1904	1236	740
No. unique choices	228	398	364	793	1040	617
No. individual-choice pairs	2202	9374	36176	87648	120193	49552

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 $\theta = -\frac{\tilde{\theta}}{\tilde{\beta}}$. γ_f is calculated as $\gamma_f = -\frac{\tilde{\gamma}_f}{\tilde{\beta}}$. Standard errors for θ and γ_f are calculated using the Delta method in all models, except in the control function approach. Standard errors in the control function approach are based on bootstrapping.

7 Consumer Welfare Implications

In this section, I use the estimates presented in Section 6 to assess the impact of partitioned pricing on consumer welfare. Specifically, I use the results of the base model shown in Table 5. To do so, first, I fix ideas on how to define consumer welfare in this context.

I follow the approach proposed by Bernheim and Rangel (2009) and described in Bernheim and Taubinsky (2018). I differentiate 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 is the domain in which consumers make decisions as they do in the real world, 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. $\theta = \delta_f = 0$. 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.[9]

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 structural estimates of my demand estimation to assess counterfactual choices in the welfare-relevant domain.

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.

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 probabilities now depend on the perceived utilities while the consumer surplus from each choice depends on the welfare-relevant utility. Train (2015) shows that, in cases like these, a term can be added to account for this discrepancy. I outline his approach here.

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

Let U_{ij} be the welfare-relevant utility that a consumer i receives from product j and $\hat{U}_{ij}$ be the perceived utility. Define the difference between the two as $d_{ij} = U_{ij} - \hat{U}_{ij}$. Let $ij*$ be the alternative that the consumer chooses based on $\hat{U}_{ij}$. 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 $E(\hat{U}_{ij*})$, i.e. the expected *perceived* utility if choosing based on $\hat{U}_{ij}$.

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

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

First, consider $E(\hat{U}_{ij*})$. 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 $E(\hat{U}_{ij*}) = ln\left(\sum_{ij} e^{\hat{W}_{ij}}\right)$, where $\hat{W}_{ij}$ 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 $U_{ij} - \hat{U}_{ij}$ which can be calculated, given the data and parameters.

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

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

As stated above, I assume that the actual indirect utility differs from the perceived indirect utility described in Equation (2.4) only because $\theta = 0 \rightarrow \tilde{\theta} = 0$ and $\delta_f = 0 \rightarrow \tilde{\delta_f} = 0$. Therefore, d_{ij} is given as

$$d_{ij} = U_{ij} - \hat{U}_{ij} = -\tilde{\delta}_f f_{ij} - \tilde{\theta} c_{ij} \,. \tag{7.4}$$

Let 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 consumer surplus due to not using the welfare-relevant utility for decision-making as $\Delta CS_i = \hat{CS}_i - CS_i$. Note that this is the expected loss in consumer surplus for any consumer who faces the same choice situation as consumer i. I then calculate the mean of this statistic for all observations in my sample to obtain the mean expected loss in consumer surplus.

Table 6: Mean expected loss in consumer surplus per transaction due to partitioned pricing

	Exit	Pegasus	Spiderman	FIFA 19	Pokemon	Duos
$\frac{1}{N}\sum_i \Delta CS_i$	-0.5710	-0.1464	-0.0621	-0.1442	-0.0598	-0.1685
$\frac{1}{N}\sum_i (\Delta CS_i / CS_i)$	0.0533	0.0054	0.0024	0.0281	0.0022	0.0013

Notes: $\Delta CS_i = \hat{CS}_i - CS_i$ is the loss in consumer surplus of consumer i due to not using the welfare-relevant utility for decision-making. $\Delta CS_i / CS_i$ 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 6, I show the mean expected loss per purchase that the consumers in the sample incurred due to not choosing according to their welfare-relevant utility. For all products, the expected loss is economically small on average, ranging from six to 57 cents per purchase. As a percentage of the consumer surplus under fully rational decisions, this amounts to relative losses of less than one percent to up to five 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 automatically displaying the total price already in the search results or removing the option for sellers to set a separate shipping fee.

7.2 Welfare by Shipping Fee Level

The small welfare impact of θ can be explained by two reasons. First, eBay caps the shipping fee at 9.50 Euros. Therefore, even if sellers wanted to exploit consumers more, there is a cap. Table 7 shows the mean expected welfare loss for a counterfactual in which sellers that set a positive shipping fee were to shift the entire price into the shipping fee. In this stylized counterfactual, the expected welfare loss becomes substantial for some of the products with relative losses of up to 34 percent of the consumer surplus under rational decision making.

Table 7: Mean expected loss in consumer surplus with high shipping fees

	Exit	Pegasus	Spiderman	FIFA 19	Pokemon	Duos
$\frac{1}{N}\sum_i \Delta CS_i$	-2.1051	-7.078	-0.7619	-0.3066	-9.0174	-20.6855
$\frac{1}{N}\sum_i (\Delta CS_i / CS_i)$	0.1916	0.2792	0.0295	0.0544	0.3395	0.1822

Notes: Consumer welfare losses shown are losses incurred due to partitioned pricing in a scenario in which sellers that do not offer free shipping set their shipping fee to the total price and the product price to zero. $\Delta CS_i = \hat{CS}_i - CS_i$ is the loss in consumer surplus of consumer i due to not using the welfare-relevant utility for decision-making. $\Delta CS_i/CS_i$ is that loss relative to the level of consumer surplus under rational decision-making. Numbers shown here are means over all consumers $i \in 1, ..., N$.

Such a scenario is, however, not possible as eBay caps the level of the shipping fee exogenously. Further, in the data, not many sellers actually set the shipping fee at the cap but rather at intermediate levels.

A second explanation is that demand seems to react discontinuously positively to an offer of free shipping for all products except "FIFA 19." This effect tends to reduce the impact that limited consideration of the shipping fee might have on consumer choices. 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 cap. 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 biases is not linear in the shipping fee. Rather, it seems that there is an optimal level of the shipping fee for each product at which the mean expected loss in consumer surplus from deviating from the welfare-relevant utility is closest to zero. At these values, the choices made with the perceived utility are closest to those in the fully rational scenario on average.

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

	Exit	Pegasus	Spiderman	FIFA 19	Pokemon	Duos
Shipping fee = 1	-1.0456	-0.025	-0.0827	-0.1362	-0.1199	-0.4332
Shipping fee = 2	-0.5567	-0.0001	-0.0714	-0.1403	-0.0845	-0.3522
Shipping fee = 3	-0.1766	-0.031	-0.0606	-0.1444	-0.0544	-0.277
Shipping fee = 4	-0.0149	-0.1146	-0.0504	-0.1487	-0.0302	-0.2087
Shipping fee = 5	-0.0102	-0.2428	-0.0407	-0.1529	-0.0127	-0.1483
Shipping fee = 6	-0.0665	-0.4055	-0.0319	-0.1573	-0.0025	-0.0968
Shipping fee = 7	-0.1481	-0.5941	-0.0239	-0.1617	-0.0002	-0.0553
Shipping fee = 8	-0.2563	-0.8032	-0.0169	-0.1661	-0.0063	-0.0247
Shipping fee = 9	-0.3876	-1.0314	-0.0109	-0.1707	-0.0213	-0.006

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 cap while keeping the total price constant. If the total price is smaller than the cap, 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.

These optimal 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.

Table 9: Indifference shipping fee levels

	Exit	Pegasus	Spiderman	FIFA 19	Pokemon	Duos
Indifference shipping fee	4.50	1.95	12.82	-67.36	6.78	9.95

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.

The indifference shipping fee corresponds to the minima of the mean expected losses shown in Table 8. “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.

8 Conclusion

Past 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 derive an empirical discrete choice model based on a theoretical framework suggested by DellaVigna (2009). Using the estimates of the discrete choice model, I apply the framework of Bernheim and Rangel (2009) and the method of Train (2015) to calculate the impact such behavioral patterns have on consumer welfare.

My results suggest that, for most products, consumers indeed behave as if they ignore part of the shipping fee. However, my results also suggest that consumer demand tends to react discontinuously positively to the offer of free shipping. This is a result that past research could not capture 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 results are robust to extensions that allow for unobserved consideration sets, endogeneous prices, and unobserved heterogeneity in price sensitivity for half of the analyzed products. For the other products, the results are affected by some of the extensions.

The welfare impact of the behavioral patterns identified in the data is small with average losses in consumer surplus below six percent of the absolute level of consumer surplus under rational decision making. Two main reasons underlie this result: First, the size of the shipping fee is capped at 9.50 Euros by eBay. Welfare calculations show that if sellers were to charge higher shipping fees, the welfare loss could be higher. This cap puts a bound on how much sellers can exploit potential biases of consumers. Second, the positive demand effect of free shipping partly counteracts the under-reaction to shipping fees in expectation.

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. (2018). Price Salience and Product Choice. *NBER Working Paper Series*, 25186.

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.

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.

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.

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.

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. (1974). Conditional logit analysis of qualitative choice behavior. In Zarembka, P., editor, *Frontiers in Econometrics*, chapter 4, pages 105–142. Academic Press, New York.

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.

Petrin, A. and Train, K. (2010). A Control Function Approach to Endogeneity in Consumer Choice Models. *Journal of Marketing Research*, 47(1):3–13.

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, 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.” For the active listings, I restrict the results to only those that are still active. Similarly for 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 active 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 3 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 4 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 5.

Figure 3: 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.

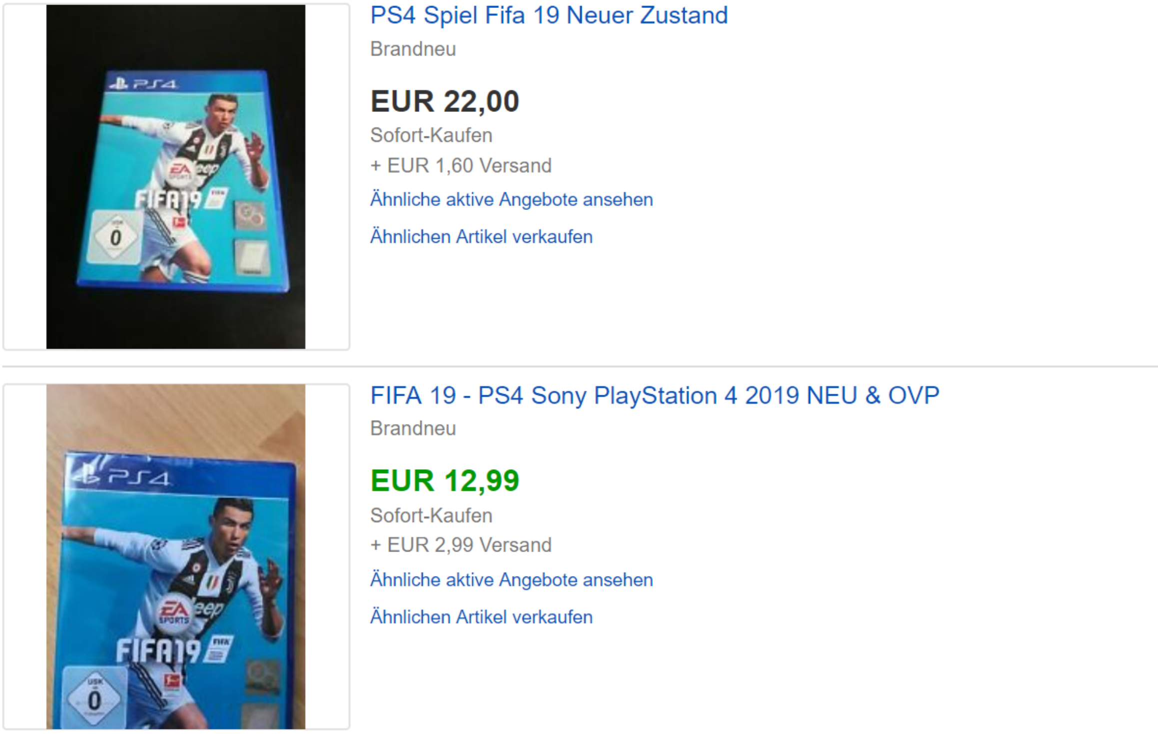

Figure 4: An active listing with an inventory for sale. A click on "6 verkauft" (6 sold) opens a list of past transactions.

Figure 5: The list of past transactions for a listing on eBay Germany.

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

A.2 Full OLS Results for Total Price

Table 10: Estimated coefficients of a linear regression explaining total price

	Exit	Azul	Spiderman	FIFA 19	Pokemon	Duos
Dependent variable	Total price	Total price	Total price	Total price	Total price	Total price
Shipping fee	0.837***	-0.241	2.186*	1.821**	3.898***	2.179
	[0.511, 1.163]	[-1.465, 0.983]	[-0.185, 4.557]	[0.020, 3.622]	[1.740, 6.056]	[-0.511, 4.869]
Seller score (K)	-0.002*	-0.009**	0.009	0.008	-0.003	0.014**
	[-0.004, 0.000]	[-0.017, -0.001]	[-0.003, 0.021]	[-0.007, 0.023]	[-0.010, 0.004]	[0.004, 0.024]
Pos. reviews (%)	-0.066	-0.002	0.120**	0.067	0.156***	0.154
	[-0.416, 0.284]	[-0.074, 0.071]	[0.001, 0.239]	[-0.025, 0.158]	[0.090, 0.222]	[-1.342, 1.650]
Commercial seller	-0.667	8.817***	9.815***	11.270***	18.280***	19.580***
	[-1.601, 0.267]	[3.630, 14.004]	[4.013, 15.617]	[5.025, 17.515]	[14.334, 22.226]	[12.638, 26.522]
Multiple units	0.160	1.919	-5.068	0.690	-3.339*	-0.630
	[-0.648, 0.968]	[-0.535, 4.373]	[-11.757, 1.621]	[-5.126, 6.506]	[-6.671, -0.007]	[-10.737, 9.477]
Payment: Bill	2.142***	2.575	-7.847**	0.310	-6.700*	7.749*
	[1.289, 2.995]	[-2.564, 7.714]	[-15.452, -0.242]	[-6.799, 7.419]	[-14.202, 0.802]	[-1.068, 16.566]
Payment: Cash on delivery	3.348***	-5.129*	0.000	-9.974	0.000	0.000
	[1.574, 5.122]	[-11.016, 0.758]	[.,.]	[-21.974, 2.026]	[.,.]	[.,.]
Payment: Cash on pickup	0.608	-0.316	-4.029	-2.846	-1.879	0.188
	[-0.344, 1.560]	[-2.852, 2.220]	[-9.750, 1.692]	[-6.974, 1.282]	[-4.835, 1.077]	[-5.646, 6.022]
Payment: Credit Card	0.248	0.365	-10.070**	7.029**	-5.259***	-58.830***
	[-0.246, 0.742]	[-1.391, 2.121]	[-18.029, -2.111]	[0.067, 13.991]	[-9.207, -1.311]	[-80.028, -37.632]
Payment: Other	0.245	-0.902	0.000	-24.960***	0.327	26.600***
	[-0.553, 1.043]	[-6.079, 4.275]	[.,.]	[-33.334, -16.586]	[-7.533, 8.187]	[13.653, 39.547]
Payment: Paypal	0.435	-2.256	8.514*	-11.860***	2.834	72.990***
	[-1.327, 2.197]	[-9.112, 4.600]	[-1.274, 18.302]	[-20.146, -3.574]	[-2.257, 7.925]	[50.380, 95.600]
Payment: Bank transfer	-0.612	0.670	-4.789*	-3.471	2.113	-7.280
	[-1.505, 0.281]	[-3.700, 5.040]	[-10.033, 0.455]	[-8.718, 1.776]	[-0.719, 4.945]	[-17.450, 2.890]
Intercept	20.850	33.470***	23.250***	16.730**	7.888	98.050
	[-14.525, 56.225]	[24.327, 42.613]	[6.238, 40.262]	[3.430, 30.030]	[-5.616, 21.392]	[-48.868, 244.968]
New edition		-0.114				
		[-1.803, 1.575]				
Eevee edition					0.762	
					[-1.955, 3.479]	
Pokeball bundle					37.040***	
					[33.618, 40.462]	
Black						-4.078**
						[-7.333, -0.823]
Blue						5.250**
						[0.555, 9.945]
Results page	Yes	Yes	Yes	Yes	Yes	Yes
Ranking on results page	Yes	Yes	Yes	Yes	Yes	Yes
Rage x Ranking	Yes	Yes	Yes	Yes	Yes	Yes
No. individual-choice pairs	372	1434	879	13333	12306	2291
R^2	0.827	0.540	0.562	0.504	0.841	0.690

Notes: Full estimations results corresponding to Table 2. I restrict the sample to 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.

A.3 Full OLS Results for Free Shipping

Table 11: Estimated coefficients of a linear regression explaining free shipping

	Exit	Azul	Spiderman	FIFA 19	Pokemon	Duos
Dependent variable	Free shipping	Free shipping	Free shipping	Free shipping	Free shipping	Free shipping
Commercial seller	-0.020	-0.003	0.256**	0.505***	0.435***	0.657***
	[-0.390, 0.351]	[-0.235, 0.228]	[0.058, 0.454]	[0.360, 0.650]	[0.313, 0.557]	[0.420, 0.894]
Seller score (K)	-0.001***	-0.000	-0.001***	-0.001***	-0.001***	-0.000**
	[-0.001, -0.001]	[-0.000, 0.000]	[-0.001, -0.001]	[-0.001, -0.001]	[-0.001, -0.001]	[-0.000, 0.000]
Pos. reviews (%)	-0.000	0.001	0.000	0.000	-0.003**	-0.080**
	[-0.002, 0.002]	[-0.004, 0.006]	[-0.006, 0.006]	[-0.006, 0.007]	[-0.005, -0.001]	[-0.146, -0.014]
Multiple units	-0.124	0.013	-0.040	-0.116	0.012	-0.061
	[-0.299, 0.051]	[-0.082, 0.109]	[-0.174, 0.094]	[-0.257, 0.025]	[-0.096, 0.120]	[-0.174, 0.053]
Payment: Bill	0.277	0.198*	0.174**	0.011	0.222***	0.120**
	[-0.119, 0.673]	[-0.025, 0.421]	[0.015, 0.333]	[-0.214, 0.236]	[0.098, 0.346]	[0.025, 0.215]
Payment: Cash on delivery	0.090	-0.188	0.218***	0.378***	0.237***	0.118**
	[-0.398, 0.578]	[-0.490, 0.114]	[0.083, 0.353]	[0.219, 0.537]	[0.095, 0.379]	[0.011, 0.225]
Payment: Cash on pickup	-0.021	0.067	0.137**	-0.138*	-0.175***	-0.193***
	[-0.329, 0.287]	[-0.176, 0.310]	[0.006, 0.268]	[-0.277, 0.001]	[-0.295, -0.055]	[-0.301, -0.085]
Payment: Credit Card	-0.041	0.070*	-0.102	-0.005	-0.059	0.513***
	[-0.177, 0.096]	[-0.013, 0.152]	[-0.271, 0.067]	[-0.123, 0.113]	[-0.148, 0.030]	[0.235, 0.791]
Payment: Other	-0.191	0.263*	0.212**	0.098	0.276***	-0.094
	[-0.713, 0.331]	[-0.009, 0.535]	[0.026, 0.398]	[-0.202, 0.397]	[0.115, 0.437]	[-0.234, 0.045]
Payment: Paypal	1.009***	0.738***	0.544***	0.149	0.400***	-0.477***
	[0.421, 1.597]	[0.385, 1.091]	[0.262, 0.826]	[-0.086, 0.384]	[0.210, 0.590]	[-0.826, -0.128]
Payment: Bank transfer	0.080	-0.086	-0.043	-0.029	-0.078	0.254***
	[-0.301, 0.460]	[-0.313, 0.141]	[-0.179, 0.093]	[-0.197, 0.138]	[-0.191, 0.034]	[0.123, 0.385]
Intercept	-0.037	-0.087	0.152	0.133	0.589***	7.923**
	[-0.580, 0.507]	[-0.654, 0.479]	[-0.375, 0.679]	[-0.602, 0.868]	[0.303, 0.875]	[1.449, 14.397]
New edition		-0.296***				
		[-0.459, -0.133]				
Eevee edition					0.011	
					[-0.084, 0.105]	
Pikachu edition					0.000	
					[.,.]	
Pokeball bundle					-0.063	
					[-0.158, 0.032]	
Black						0.086
						[-0.031, 0.203]
Blue						0.116*
						[-0.009, 0.241]
Results page	Yes	Yes	Yes	Yes	Yes	Yes
Ranking on results page	Yes	Yes	Yes	Yes	Yes	Yes
Rage x Ranking	Yes	Yes	Yes	Yes	Yes	Yes
No. individual-choice pairs	2202	9374	36183	87648	120193	49555
R^2	0.438	0.509	0.497	0.328	0.419	0.345

Notes: Full estimations results corresponding to Table 3. ***, **, and * indicate statistical significance at the one, five, and ten percent levels, respectively. The brackets show 95 percent confidence intervals.

A.4 Full OLS Results for Shipping Fee

Table 12: Regressions results of a linear regression explaining the shipping fee

	Exit	Azul	Spiderman	FIFA 19	Pokemon	Duos
Dependent variable	Shipping fee	Shipping fee	Shipping fee	Shipping fee	Shipping fee	Shipping fee
Commercial seller	0.188	-0.511	0.106	-0.168	0.435**	-0.124
	[-0.587, 0.963]	[-1.795, 0.773]	[-1.054, 1.266]	[-0.899, 0.563]	[0.090, 0.780]	[-0.916, 0.668]
Seller score (K)	-0.003***	-0.003***	-0.002***	-0.003***	-0.003***	-0.004***
	[-0.004, -0.002]	[-0.004, -0.002]	[-0.003, -0.001]	[-0.004, -0.002]	[-0.004, -0.002]	[-0.005, -0.003]
Pos. reviews (%)	-0.003	-0.020	-0.017***	0.002	-0.014	-0.013
	[-0.009, 0.003]	[-0.055, 0.015]	[-0.024, -0.009]	[-0.006, 0.010]	[-0.038, 0.010]	[-0.270, 0.244]
Multiple units	0.031	0.236	0.989*	1.318***	0.465**	1.375***
	[-0.781, 0.844]	[-0.262, 0.734]	[-0.017, 1.995]	[0.740, 1.896]	[0.057, 0.873]	[0.430, 2.320]
Payment: Bill	-1.741***	0.070	-0.151	-0.023	-2.411***	-1.484***
	[-2.232, -1.250]	[-0.316, 0.456]	[-2.566, 2.264]	[-0.770, 0.724]	[-3.015, -1.807]	[-2.280, -0.688]
Payment: Cash on delivery	3.356***	0.147	0.000	-0.013	0.000	0.000
	[2.153, 4.559]	[-0.563, 0.857]	[.,.]	[-0.963, 0.938]	[.,.]	[.,.]
Payment: Cash on pickup	0.860**	-0.215	-0.407	-0.673**	-0.344	-0.318
	[0.110, 1.610]	[-0.735, 0.305]	[-0.972, 0.158]	[-1.194, -0.152]	[-0.836, 0.148]	[-1.010, 0.374]
Payment: Credit Card	0.034	0.372**	0.122	-0.004	-0.221	-0.509
	[-0.371, 0.438]	[0.023, 0.721]	[-0.529, 0.773]	[-0.639, 0.631]	[-0.684, 0.242]	[-1.536, 0.518]
Payment: Other	1.171***	-0.506**	0.000	0.264	1.164**	0.476
	[0.378, 1.964]	[-0.896, -0.116]	[.,.]	[-0.745, 1.273]	[0.190, 2.138]	[-0.504, 1.456]
Payment: Paypal	-0.088	0.223	0.230	0.690	0.970***	0.563
	[-1.234, 1.058]	[-0.944, 1.390]	[-0.732, 1.192]	[-0.339, 1.719]	[0.276, 1.664]	[-0.980, 2.106]
Payment: Bank transfer	-1.426***	-0.450**	-0.311	-0.098	0.121	1.704***
	[-2.203, -0.649]	[-0.883, -0.017]	[-1.364, 0.742]	[-0.764, 0.569]	[-0.479, 0.721]	[0.855, 2.553]
Min. days of shipping	0.112**	0.117	-0.074	-0.008	0.090***	0.012
	[0.025, 0.199]	[-0.024, 0.258]	[-0.213, 0.065]	[-0.087, 0.071]	[0.056, 0.124]	[-0.092, 0.116]
Intercept	4.661***	6.533***	4.549***	2.524***	3.870***	5.234
	[3.674, 5.648]	[3.319, 9.747]	[3.296, 5.802]	[1.581, 3.467]	[1.414, 6.326]	[-19.935, 30.403]
New edition		-0.223				
		[-0.548, 0.102]				
Pikachu edition					-0.022	
					[-0.353, 0.309]	
Pokeball bundle					0.898***	
					[0.549, 1.247]	
Blue						-0.151
						[-0.551, 0.249]
Gold						-0.085
						[-0.432, 0.262]
No. individual-choice pairs	877	2787	9834	28518	46457	9156
R^2	0.637	0.396	0.358	0.434	0.553	0.823

Notes: Full estimations results corresponding to Table 4. 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.

A.5 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 (2.6):

$$P_{ij} = \sum_{C \in S_j} \prod_{l \in C} \pi_{il} \prod_{k \notin C} (1 - \pi_{ik}) \frac{exp(W(X_ij, \Theta))}{\sum_{k \in C} exp(W(X_{ik}, \Theta))} .$$

Following Goeree (2008), I specify π_{ij} as

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

where $\kappa_{ij} = \varphi + K'_{ij}\rho$. 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 (2.6) 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 equation (2.5) exists. However, as Goeree (2008) already notes, to calculate Equation (2.6) 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^0_{ij} 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

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

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 $\Pi^0_{ir} = \prod_{l \in C_{ir}} \pi^0_{il} \prod_{k \notin C_{ir}} (1 - \pi^0_{ik})$. 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^s_{ij} 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

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

The weight $\frac{1}{\Pi^0_{ir}}$ 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

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

A.6 Estimation Procedure for the Control Function Approach

I outline the procedure to estimate the base model with a control function more formally here. I follow the description in Petrin and Train (2010) closely. Assume that the total price tp_{ij} can be explained by some function

$$tp_{ij} = T(x_{ij}, c_{ij}, f_{ij}, z_{ij}) + \mu_{ij} \, .$$

x_{ij} is the same vector of non-financial characteristics that enters the indirect utility. c_{ij} are the shipping fees, f_{ij} is the free shipping indicator, and z_{ij} is a vector of instrumental variables that are excluded from the linear utility function. Finally, μ_{ij} is a vector of factors that are unobserved to the researcher.

Further assume that the error term ϵ_{ij} of the indirect utility and μ_{ij} are independent of x_{ij}, c_{ij}, f_{ij}, and z_{ij}. However, the two error terms are not independent of each other. Because of the correlation between ϵ_{ij} and μ_{ij}, the total price tp_{ij} is endogenous in the utility specification. The idea of the control function approach is to control for the part of ϵ_{ij} that depends on μ_{ij}.

Petrin and Train (2010) suggest to decompose ϵ_{ij} into a mean conditional on μ_{ij} and the remaining deviation around this conditional mean: $\epsilon_{ij} = E(\epsilon_{ij}|\mu_{ij}) + \tilde{\epsilon}_{ij}$. By construction, $\tilde{\epsilon}_{ij}$ is independent of μ_{ij}. $E(\epsilon_{ij}|\mu_{ij}) \equiv CF(\mu_{ij}; \lambda)$ is the so-called control function.

To implement the approach, estimates for μ_{ij} are needed as well as a functional form for $CF(\mu_{ij}; \lambda)$. To estimate μ_{ij}, I estimate

$$tp_{ij} = Z'_{ij}\eta + \mu_{ij}$$

using ordinary least squares regression. I then calculate $\hat{\mu}_{ij} = tp_{ij} - Z'_{ij}\hat{\eta}$ as an estimate for μ_{ij}. Z_{ij} is a vector containing x_{ij}, c_{ij}, f_{ij}, as well as z_{ij}.

z_{ij} contains instruments in the spirit of Berry et al. (1995). More specifically, for each listing I include the sum across all other listings in the choice set for each characteristic. Therefore, for a choice j in choice set S_i and characteristic k, the instrument is

$$\sum_{r \neq j, r \in S_i} x_{irk} \, .$$

To obtain an expression for the control function, assume that $\epsilon_{ij} = \epsilon^1_{ij} + \epsilon^2_{ij}$, where ϵ^1_{ij} and

μ_{ij} are jointly normally distributed and ϵ^2_{ij} is independent of μ_{ij}.[10] Petrin and Train (2010) show that this results in an indirect utility function of the form

$$U(x_{ij}, tp_{ij}, c_{ij}, f_{ij}; \delta, \tilde{\beta}, \tilde{\theta}) = x'_{ij}\tilde{\gamma} + \tilde{\beta}tp_{ij} + \tilde{\theta}c_{ij} + \tilde{\gamma}_f f_{ij} + \lambda\mu_{ij} + \sigma_{CF}\eta_{ij} + \epsilon^2_{ij} \, ,$$

where $\eta_{ij} \sim N(0,1)$. I plug in $\hat{\mu}_{ij}$ from the first stage price regression as an estimate of μ_{ij}. Assuming that ϵ^2_{ij} follow an extreme value type I distribution, I obtain the familiar logit choice probabilities conditional on the realization of η_{ij}. The actual choice probability is then

$$P_{ij} = \int \frac{exp(W(X_{ij}, \Theta) + \sigma_{CF}\eta_{ij})}{\sum_{k \in S_i} exp(W(X_{ik}, \Theta) + \sigma_{CF}\eta_{ik})} \prod_{l \in S_i} \phi(\eta_{il}) d\eta_{il} \, ,$$

where $W(X_{ij}, \Theta) = x'_{ij}\tilde{\gamma} + \tilde{\beta}tp_{ij} + \tilde{\theta}c_{ij} + \tilde{\gamma}_f f_{ij} + \lambda\hat{\mu}_{ij}$ and $\phi(\eta)$ is the standard normal probability density function. Because this integral does not have an analytical solution, I use simulated choice probabilities for the estimation.

[10] Petrin and Train (2010) show that these assumptions are consistent with sellers pricing at marginal costs or at a constant markup.

A.7 Full Estimation Results of Base Model

Table 13: Logit estimation results of the base model described in Equation (2.4)

	Exit	Azul	Spiderman	FIFA 19	Pokemon	Duos
Inattention (θ)	1.065***	0.920**	0.177	0.039	0.390**	0.830**
	[0.335, 1.795]	[0.210, 1.630]	[-0.106, 0.460]	[-0.271, 0.349]	[0.045, 0.735]	[0.176, 1.484]
Free shipping effect (γ_f)	4.797***	1.789	2.264***	-2.649***	2.655***	8.250***
	[1.845, 7.749]	[-1.432, 5.010]	[1.067, 3.461]	[-3.845, -1.453]	[1.511, 3.799]	[4.188, 12.312]
Total price ($\tilde{\beta}$)	-0.965***	-0.606***	-0.561***	-0.208***	-0.304***	-0.176***
	[-1.346, -0.584]	[-0.747, -0.465]	[-0.616, -0.506]	[-0.218, -0.198]	[-0.325, -0.283]	[-0.188, -0.164]
Free shipping ($\tilde{\gamma}_f$)	4.629***	1.084	1.270***	-0.551***	0.807***	1.452***
	[1.559, 7.699]	[-0.849, 3.017]	[0.616, 1.924]	[-0.796, -0.306]	[0.460, 1.154]	[0.735, 2.169]
Shipping fee ($\tilde{\theta}$)	1.028**	0.557**	0.099	0.008	0.119**	0.146**
	[0.223, 1.833]	[0.115, 0.999]	[-0.059, 0.258]	[-0.056, 0.073]	[0.013, 0.225]	[0.030, 0.262]
Commercial seller	-1.083	0.449	0.562**	0.531***	0.775***	1.444*
	[-2.767, 0.601]	[-0.905, 1.803]	[0.105, 1.019]	[0.327, 0.735]	[0.494, 1.056]	[-0.233, 3.121]
Seller score (K)	0.004***	0.003***	-0.000	0.001***	-0.001***	-0.001***
	[0.001, 0.006]	[0.002, 0.005]	[-0.001, 0.001]	[0.000, 0.001]	[-0.001, -0.000]	[-0.002, -0.000]
Payment: Paypal	-1.426	-0.118	0.801***	0.951***	0.102	1.393**
	[-4.153, 1.301]	[-2.079, 1.843]	[0.291, 1.311]	[0.625, 1.277]	[-0.208, 0.412]	[0.215, 2.571]
Multiple units	0.553	1.291***	1.939***	1.821***	0.894***	1.058***
	[-0.569, 1.675]	[0.502, 2.080]	[1.590, 2.288]	[1.646, 1.996]	[0.725, 1.063]	[0.722, 1.394]
New edition		1.042				
		[-0.439, 2.523]				
Pikachu edition					0.035	
					[-0.102, 0.173]	
Pokeball bundle					8.777***	
					[8.200, 9.354]	
Blue						-0.525***
						[-0.880, -0.170]
Gold						-0.352**
						[-0.678, -0.026]
No. individuals	53	108	933	1904	1236	740
No. unique choices	228	398	364	793	1040	617
No. individual-choice pairs	2202	9374	36176	87648	120193	49552

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

A.8 Full Estimation Results of Random Consideration Model

Table 14: Estimation results with random consideration

	Exit	Azul	Spiderman	FIFA 19	Pokemon	Duos
Inattention (θ)	1.225***	-0.120	0.111	-0.145	-0.240**	0.720**
	[0.436, 2.014]	[-1.003, 0.763]	[-0.045, 0.267]	[-0.472, 0.182]	[-0.475, -0.005]	[0.132, 1.308]
Free shipping effect (γ_f)	5.299***	-1.866	1.869***	-0.797	1.073**	6.697***
	[2.184, 8.415]	[-5.430, 1.697]	[1.547, 2.191]	[-1.899, 0.305]	[0.180, 1.967]	[3.353, 10.041]
Total price ($\tilde{\beta}$)	-1.030***	-0.929***	-0.857***	-0.396***	-0.424***	-0.210***
	[-1.380, -0.681]	[-1.167, -0.692]	[-0.965, -0.749]	[-0.418, -0.374]	[-0.444, -0.404]	[-0.226, -0.193]
Shipping fee ($\tilde{\theta}$)	1.262***	-0.111	0.095	-0.058	-0.102**	0.151**
	[0.439, 2.084]	[-0.937, 0.714]	[-0.038, 0.227]	[-0.187, 0.072]	[-0.202, -0.002]	[0.028, 0.274]
Free shipping ($\tilde{\gamma}_f$)	5.460***	-1.735	1.601***	-0.315	0.455**	1.405***
	[2.370, 8.551]	[-5.153, 1.684]	[1.329, 1.872]	[-0.750, 0.120]	[0.079, 0.831]	[0.694, 2.115]
Commercial seller	0.352	1.737	1.300***	0.386***	0.904***	-0.007
	[-16.589, 17.293]	[-0.656, 4.131]	[0.721, 1.878]	[0.167, 0.606]	[0.657, 1.152]	[-0.790, 0.777]
Seller score (K)	0.005***	0.003**	-0.002***	0.001***	-0.000	-0.001***
	[0.002, 0.007]	[0.000, 0.005]	[-0.003, -0.001]	[0.001, 0.001]	[-0.001, 0.000]	[-0.002, -0.000]
Payment: Paypal	-3.601	-1.941	1.100***	1.376***	0.504***	2.539***
	[-22.147, 14.944]	[-4.482, 0.600]	[0.832, 1.369]	[0.995, 1.757]	[0.213, 0.794]	[1.626, 3.452]
Multiple units	0.490	1.661**	2.625***	1.972***	1.060***	1.070***
	[-0.715, 1.695]	[0.219, 3.104]	[1.969, 3.280]	[1.801, 2.144]	[0.890, 1.230]	[0.704, 1.435]
π: Total price	-0.013	-0.003	-0.001***	-0.002***	-0.000***	-0.000***
	[-0.054, 0.028]	[-0.008, 0.003]	[-0.002, -0.001]	[-0.003, -0.001]	[-0.000, -0.000]	[-0.001, -0.000]
π: Rank on results page	-0.003	-0.001	-0.001***	-0.004***	-0.003***	-0.001*
	[-0.007, 0.001]	[-0.003, 0.001]	[-0.002, -0.000]	[-0.005, -0.002]	[-0.005, -0.001]	[-0.002, 0.000]
π: No. choices	-0.002	-0.000	-0.000	-0.001	-0.000	-0.001***
	[-0.008, 0.003]	[-0.002, 0.001]	[-0.001, 0.000]	[-0.002, 0.000]	[-0.001, 0.000]	[-0.002, -0.001]
π: Results page	-0.167	-0.016	-0.009***	-0.028***	-0.038***	-0.023**
	[-0.664, 0.331]	[-0.054, 0.022]	[-0.015, -0.002]	[-0.037, -0.019]	[-0.058, -0.018]	[-0.042, -0.003]
π: Constant	0.547	0.202	0.128***	0.320***	0.249***	0.273***
	[-0.364, 1.457]	[-0.138, 0.542]	[0.103, 0.154]	[0.253, 0.386]	[0.186, 0.312]	[0.184, 0.362]
New edition		2.040**				
		[0.270, 3.811]				
Pokeball bundle					13.853***	
					[13.127, 14.579]	
Eevee edition					-0.047	
					[-0.191, 0.097]	
Gold						-0.124
						[-0.433, 0.185]
Blue						-0.395**
						[-0.757, -0.033]

Notes: ***, **, and * indicate statistical significance at the one, five, and ten percent levels, respectively. The brackets show 95 percent confidence intervals. Details of the estimation procedure are outlined in Appendix A.5.

A.9 Full Estimation Results of Control Function Approach

Table 15: Estimation results from logit estimation with control function

	Exit	Azul	Spiderman	FIFA 19	Pokemon	Duos
Inattention (θ)	1.282***	0.899***	0.043	-0.015	-0.932**	1.681**
	[0.110, 2.945]	[0.158, 1.447]	[-0.187, 0.283]	[-0.255, 0.225]	[-1.731, -0.131]	[0.585, 3.532]
Free shipping effect (γ_f)	7.236	1.689	1.305	-2.500	-13.487	12.939
	[.,.]	[.,.]	[.,.]	[.,.]	[.,.]	[.,.]
Total price ($\tilde{\beta}$)	-0.567	-0.644**	-0.503***	-0.268***	-0.163***	-0.114***
	[-13.811, 0.362]	[-1.196, -0.441]	[-0.576, -0.441]	[-0.298, -0.242]	[-0.177, -0.147]	[-0.135, -0.077]
Shipping fee ($\tilde{\theta}$)	0.727	0.579*	0.022	-0.004	-0.152**	0.192***
	[-0.290, 10.517]	[0.088, 1.268]	[-0.104, 0.144]	[-0.070, 0.060]	[-0.274, -0.023]	[0.072, 0.308]
Free shipping ($\tilde{\gamma}_f$)	4.102	1.088	0.656***	-0.669***	-2.202***	1.478***
	[1.054, 23.861]	[-1.320, 3.039]	[0.217, 1.094]	[-0.900, -0.443]	[-2.564, -1.812]	[0.879, 2.138]
Commercial seller	-0.356	0.521	0.280	0.584***	2.147***	-0.128
	[-1.627, 110.689]	[-0.655, 3.121]	[-0.199, 0.756]	[0.409, 0.782]	[1.572, 2.657]	[-1.230, 1.126]
Seller score (K)	0.005	0.003	-0.001***	0.001***	-0.001***	-0.002***
	[-0.034, 0.010]	[0.000, 0.006]	[-0.002, -0.000]	[0.001, 0.001]	[-0.001, -0.000]	[-0.002, -0.001]
Payment: Paypal	-2.660	-0.097	0.525**	1.150***	-0.994***	0.234
	[-11.102, 48.079]	[-1.149, 1.050]	[0.023, 1.047]	[0.797, 1.629]	[-1.307, -0.657]	[-1.461, 1.597]
Multiple units	0.595	1.312**	2.251***	2.132***	1.317***	1.181***
	[-2.173, 2.312]	[0.640, 2.860]	[1.980, 2.583]	[1.940, 2.382]	[1.101, 1.481]	[0.865, 1.524]
Control Function: λ	-0.398	0.019	-0.059***	-0.009	-0.162***	-0.062***
	[-1.553, 12.517]	[-0.239, 0.446]	[-0.100, -0.021]	[-0.030, 0.011]	[-0.174, -0.149]	[-0.103, -0.043]
Control Function: σ_{CF}	0.054	0.577	0.002	1.637***	1.066***	0.010
	[0.000, 1.680]	[0.007, 2.159]	[0.000, 0.007]	[1.427, 1.897]	[0.708, 1.247]	[0.000, 0.163]
New edition		1.072				
		[-3.668, 2.542]				
Pokeball bundle					1.344***	
					[1.087, 1.613]	
Eevee edition					-0.372***	
					[-0.548, -0.267]	
Gold						-0.092
						[-0.440, 0.310]
Blue						-0.384*
						[-0.772, -0.027]
Bootstrap repetitions	1000	1000	329	134	50	135

Notes: The choice probabilities described in Equation (2.7) are approximated using simulation. For the simulation, I use 100 antithetic draws. ***, **, and * indicate statistical significance at the one, five, and ten percent levels, respectively. The brackets show 95 percent confidence intervals. The standard errors and confidence intervals are calculated using bootstrapping. The first stage regression results for the original sample are shown in Table 16.

Table 16: First stage price regression results for the original sample.

	Exit	Azul	Spiderman	FIFA 19	Pokemon	Duos
Shipping fee ($\tilde{\theta}$)	0.828***	0.716***	1.447***	2.701***	1.221***	-0.754***
	[0.736, 0.920]	[0.545, 0.887]	[1.224, 1.670]	[2.588, 2.814]	[0.991, 1.451]	[-0.953, -0.555]
Free shipping ($\tilde{\gamma}_f$)	1.526***	-1.500***	10.830***	10.315***	18.148***	-1.279***
	[1.057, 1.994]	[-2.433, -0.568]	[10.155, 11.505]	[9.895, 10.734]	[17.045, 19.250]	[-2.065, -0.493]
Payment: Paypal	2.763***	5.539***	4.997***	-5.141***	7.651***	18.279***
	[1.952, 3.573]	[4.458, 6.620]	[4.606, 5.389]	[-5.464, -4.818]	[6.808, 8.495]	[16.762, 19.796]
Sum of Payment: Paypal	0.217***	-0.413***	0.351***	-1.040***	0.793***	-0.331**
	[0.085, 0.349]	[-0.585, -0.240]	[0.165, 0.537]	[-1.149, -0.931]	[0.684, 0.902]	[-0.644, -0.019]
Commercial seller	-1.769***	1.769***	5.022***	11.059***	-8.339***	26.274***
	[-2.375, -1.163]	[1.250, 2.288]	[4.576, 5.469]	[10.787, 11.330]	[-10.536, -6.143]	[25.086, 27.462]
Sum of Commercial seller	-0.157***	0.263***	0.497***	1.123***	-0.366***	0.871***
	[-0.266, -0.047]	[0.175, 0.351]	[0.331, 0.663]	[1.050, 1.195]	[-0.443, -0.289]	[0.426, 1.316]
Sum of Free shipping ($\tilde{\gamma}_f$)	-0.060	0.152*	-0.119	0.532***	-0.545***	-0.804***
	[-0.166, 0.046]	[-0.001, 0.306]	[-0.278, 0.040]	[0.451, 0.614]	[-0.734, -0.356]	[-1.093, -0.514]
Sum of Shipping fee ($\tilde{\theta}$)	-0.005	0.020	-0.004	0.189***	-0.391***	-0.009
	[-0.030, 0.020]	[-0.010, 0.050]	[-0.042, 0.034]	[0.164, 0.214]	[-0.443, -0.338]	[-0.045, 0.026]
Multiple units	-0.154*	-0.734***	-7.042***	1.833***	-1.846***	-1.898***
	[-0.317, 0.008]	[-0.954, -0.514]	[-7.379, -6.705]	[1.574, 2.092]	[-2.197, -1.495]	[-2.322, -1.475]
Sum of Multiple units	-0.017	0.021	-1.705***	-0.290***	0.226***	0.087**
	[-0.062, 0.028]	[-0.016, 0.058]	[-1.815, -1.596]	[-0.343, -0.238]	[0.096, 0.357]	[0.018, 0.157]
New edition		1.875***				
		[1.538, 2.213]				
Sum of New edition		0.036***				
		[0.013, 0.060]				
Pokeball bundle					51.942***	
					[50.869, 53.015]	
Sum of Pokeball bundle					0.711***	
					[0.512, 0.910]	
Eevee edition					3.053***	
					[2.397, 3.709]	
Sum of Eevee edition					0.942***	
					[0.770, 1.114]	
Gold						-5.059***
						[-5.444, -4.675]
Sum of Gold						-0.870***
						[-1.049, -0.690]
Blue						-2.462***
						[-3.032, -1.893]
Sum of Blue						-0.188***
						[-0.298, -0.078]
Intercept	13.423***	35.711***	35.773***	23.994***	30.910***	159.361***
	[12.589, 14.257]	[34.444, 36.977]	[35.059, 36.486]	[23.347, 24.642]	[28.165, 33.654]	[156.670, 162.052]
F	9.975	7.349	256.751	1437.917	64.778	205.583
p	0.000	0.000	0.000	0.000	0.000	0.000

A.10 Full Estimation Results of Mixed Logit

Table 17: Estimation results for a mixed logit based on the indirect utility specified in Equation (2.8)

	Exit	Azul	Spiderman	FIFA 19	Pokemon	Duos
Inattention (θ)	0.993***	0.000	0.246***	0.383***	0.408***	0.643***
	[0.509, 1.477]	[-0.011, 0.011]	[0.134, 0.358]	[0.205, 0.561]	[0.217, 0.599]	[0.159, 1.127]
Free shipping effect (γ_f)	4.590***	-1.750***	2.222***	-0.178	2.665***	6.553***
	[2.604, 6.576]	[-2.682, -0.817]	[1.859, 2.586]	[-0.860, 0.503]	[1.885, 3.444]	[3.650, 9.455]
Total price ($\tilde{\beta}$)	-0.968***	-0.892***	-0.757***	-0.270***	-0.498***	-0.232***
	[-1.350, -0.585]	[-1.140, -0.643]	[-0.826, -0.687]	[-0.284, -0.256]	[-0.522, -0.473]	[-0.251, -0.212]
Shipping fee ($\tilde{\theta}$)	0.961***	0.000	0.186***	0.103***	0.203***	0.149***
	[0.575, 1.347]	[-0.010, 0.010]	[0.100, 0.272]	[0.056, 0.151]	[0.109, 0.297]	[0.037, 0.261]
Free shipping ($\tilde{\gamma}_f$)	4.442***	-1.560***	1.681***	-0.048	1.326***	1.518***
	[3.450, 5.435]	[-2.493, -0.628]	[1.405, 1.958]	[-0.232, 0.136]	[0.945, 1.707]	[0.853, 2.183]
Commercial seller	-1.035	2.301	0.766***	0.266***	1.257***	3.385***
	[-2.398, 0.329]	[-0.615, 5.217]	[0.532, 1.000]	[0.095, 0.438]	[1.037, 1.478]	[1.365, 5.405]
Seller score (K)	0.004***	0.001**	-0.000	0.001***	0.000**	-0.001***
	[0.002, 0.005]	[0.000, 0.002]	[-0.001, 0.001]	[0.001, 0.001]	[0.000, 0.001]	[-0.002, -0.000]
Payment: Paypal	-1.347**	-0.879	1.819***	1.812***	0.401***	2.771***
	[-2.540, -0.154]	[-3.914, 2.156]	[1.104, 2.534]	[1.438, 2.187]	[0.146, 0.655]	[1.325, 4.217]
Multiple units	0.604	1.743**	2.086***	2.087***	1.156***	0.916***
	[-0.316, 1.525]	[0.303, 3.183]	[1.834, 2.338]	[1.905, 2.268]	[0.992, 1.321]	[0.558, 1.274]
Std.: Total price (σ_β)	0.141	0.408***	0.277***	0.097***	0.153***	0.088***
	[-2.931, 3.212]	[-0.073, 0.889]	[0.112, 0.442]	[-0.018, 0.211]	[0.056, 0.249]	[-0.051, 0.226]
New edition		0.773				
		[-0.167, 1.712]				
Pokeball bundle					9.691***	
					[9.263, 10.119]	
Eevee edition					-0.191***	
					[-0.313, -0.070]	
Gold						-0.432**
						[-0.832, -0.032]
Blue						-0.782***
						[-1.176, -0.387]

Notes: ***, **, and * indicate statistical significance at the one, five, and ten percent levels, respectively. The brackets show 95 percent confidence intervals. I use 100 antithetic draws per individual to simulate the choice probabilities.