Identifying Consumer-Welfare Changes when Online Search Platforms Change Their Lists of Search Results

Ryan James Martin*

February 19, 2021

Abstract

Online search platforms influence product demand through their choices of how to order search results in response to their users' queries. I study the identification of consumer-welfare changes in response to exogenous changes in these choices. I focus on the case of consumers engaging in costly searches for a single, indivisible (discrete) product among a collection of substitutes. I show that exact consumer-welfare changes—that is, compensating variation and equivalent variation—can be calculated with the use of straightforward integrals of the aggregate demand. I apply my results to shopping data provided by an online travel agency (OTA).

Keywords: consumer welfare, ecommerce, search, discrete choice, compensating and equivalent variation, antitrust

*Senior Economist at the Bank of Canada. Email at RMartin@bankofcanada.ca. The opinions expressed in this document are those of the authors alone. They do not purport to reflect the opinions or views of the Bank of Canada.

1 Introduction

Online, consumers often rely on search tools to help them find relevant products to buy. Typically, consumers type keywords into the query box of a search platform, and the platform delivers a list of products; I call this delivered list of products a consumer's "search-result list." The number of substitutes available online is often substantial, and an exhaustive search over all substitutes is impractical. Consumers often start their search from the beginning of the search-result list and stop well before the end. Thus, the order products appear in consumers' search-result lists affects the products consumers discover and ultimately purchase. When search platforms change the order of products that appear in their lists of search results, aggregate demand and consumer welfare are also changed. I study the identification of consumer-welfare changes in response to exogenous changes in search-result lists. I focus on the case of consumers shopping for a single, indivisible product. I show that, in this environment, exact consumer-welfare changes—that is, compensating variation and equivalent variation—can be calculated using straightforward integrals of aggregate demand.

I model consumers as having product knowledge limited to "consideration sets." Each consumer forms her consideration set as a function of her own characteristics and the search-result lists that are returned to her queries. Each consumer's demand is limited to the products that are in her consideration sets: consideration sets censor demand. Aggregate demand is a function of all prices and individual consideration sets. Changes in search-result lists cause discontinuous shifts in demand through their actions on consideration sets. I determine formulas that recover average welfare changes from these shifts in aggregate demand. My results do not require consideration sets to be observed. My model places minimal restrictions on how consumers form their consideration sets. The key restriction is that consideration sets do not depend directly on price. (They are allowed to depend on past prices or price beliefs.) This assumption ensures that demand lines reflect the tradeoffs between goods and their prices. I show that without this assumption, areas under demand curves and even simple utility differences across A/B settings need not reflect meaningful welfare measures. My context of an A/B experiment allows for rich welfare inferences under relatively weak restrictions on consumer primitives.

I consider two applications. In both, I use data from an online travel agency (OTA). For the first application, I measure the welfare consequences of the OTA's changing search-result rules (or algorithms). I take advantage of an experiment that the OTA ran listing hotel bookings for treated consumers in random order and hotel bookings for untreated consumers in a proprietary order. I find that the OTA's proprietary ordering improves average welfare by \$8.11 per user over the random ordering. In the second application, I estimate the welfare loss that would result if the OTA were to remove the five most popular products from all search-result lists. I estimate this would lower average welfare by $23.87 per person.

My welfare formulas focus on the utility consumers receive from their final product purchases under different search-result lists. That is, my welfare measures do not explicitly account for the psychic and time costs that arise from changes in search-result lists. My welfare measures account for these costs implicitly through their impact on consideration sets and ultimate product choice. This allows me to leave the consumer's search behavior relatively unrestricted.[1] My demand model makes minimal assumptions about consumer preferences. My most general results require that utility be linear in money but allow the remaining terms to be nonparametric and (potentially) nonseparable in the unobservables. For the special case of a change in search-result lists that causes the (probabilistic) removal or addition of a collection of products from consideration sets, I show that the weaker assumption of a utility that is monotonic in money (and otherwise unrestricted) is sufficient for bounding the resulting welfare changes.

My research is motivated by recent antitrust concerns over the growing concentration of online platforms.[2] Online, most consumers and sellers find each other through the services of these platforms. Without these platforms, consumers and sellers would have a hard time making new connections. Thus, a concentrated platform may have great influence over who buys what from whom online and at what price. This gives rise to several antitrust concerns over search-platform conduct. For example, a search platform that also sells its own products, such as Amazon, may be tempted to flex its influence over consumers' search results in negotiations with sellers in order to extract high proportions of seller revenue on their site and limit seller behavior off their site.[3] Alternatively, a search platform may be tempted to hide the search results that would lead consumers to other search platforms, in order to protect its own market share. This is especially concerning when the search platform represents its organic[4] search results as unbiased. European antitrust authorities fined Google $2.7 billion for this type of conduct.[5] To the extent that this platform conduct influences the distribution of sellers and the ultimate product choice of consumers, it may also have strong consumer-welfare implications. My model framework allows for a rigorous,

[1] To be clear, consumer search is allowed to be costly in my model. Consumers will take their own psychic and time costs into account when deciding how much of the product space to explore. The welfare effects of these costs will then be captured by how they change their final product choice.

[2] For example, the Federal Trade Commission (FTC) is holding ongoing hearings on competition and consumer protection in the 21st century that address this topic: https://www.ftc.gov/policy/hearings-competition-consumer-protection.

[3] See Khan (2017) for more discussion on evidence of this behavior.

[4] An organic search result is one whose position is not chosen through a payment to the search platform, but rather by the search platform's own algorithm.

[5] http://money.cnn.com/2017/06/27/technology/business/google-eu-antitrust-fine/index.html

minimally restricted study of these consumer-welfare consequences.

My research also highlights the important role A/B tests can play in online search settings. Economists are long familiar with the importance of exogenous variation in price for demand and welfare estimation in classic economic environments. Analogously, exogenous variation in platform listings (in the context of an A/B test or policy intervention) can be essential for accurately measuring welfare changes in search settings. Consider a platform that may choose to list products according to the strategy `list_exploit` or the strategy `list_encourage`. When the platform uses the strategy `list_exploit`, it places the products with the largest markup on the first page. When the platform uses the strategy `list_encourage`, it puts products that generate the most consumer surplus on the first page. Further suppose that consumers may be one of two types: *lazy* or *determined*. *Lazy* consumers only search the first page of results, while *determined* consumers know their surplus-maximizing product and will search all pages, if necessary, starting from the first and continuing until they find it. If the data show consumers only searching the first page of results, then this could be explained equally well by either *lazy* consumers and a `list_exploit` strategy or *determined* consumers and a `list_encourage` strategy. An A/B test lets us separate these two cases and draw welfare conclusions in ways data under one search-result-list algorithm cannot.

To the best of my knowledge, this paper is the first to identify compensating variation and equivalent variation from search-result listings as a function of changes in aggregate demand in a general search environment. However, this paper does draw insights and inspiration from several existing strands of literature.

First, this paper is related to that of Bhattacharya (2015). In particular, both his paper and my paper identify the compensating and equivalent variations in a discrete-choice environment where utility is monotonic in money but otherwise unrestricted. While Bhattacharya (2015) examines the welfare consequences of a single price increase, I focus on the welfare consequences of changing the search-result lists. In Bhattacharya (2015), consumer knowledge is perfect and there is no search; the demand and welfare consequences of consideration sets and limited consumer knowledge are not modeled. I show that the additional complications of the search environment require strengthening the utility assumptions to a utility that is linear in money in order to measure generic consumer welfare changes from changes in search-result lists. Further, while my environment is more complex, I find novel and relatively elegant proofs to achieve nonparametric results analogous to his.

Second, my results are related to the literature on the consumer welfare changes that occur from the introduction of new goods. This econometric literature goes back to Hausman (1981, 1996). The above papers and their extensions focus on identifying and estimating the exact consumer-welfare changes that result from the introduction of a single product or product category. All consumers are assumed to have perfect knowledge of all available products. The newness of the product is modeled as the consequence of technological innovation or regulatory decisions and is assumed to be exogenous to the consumer's final purchase decision. In contrast, my paper focuses on the role that search-result lists have in shaping product knowledge and welfare. Individual search behavior is heterogeneous and the consequences of a changing search-result list may be heterogeneous and unpredictable across consumers in my environment.

This paper is also related to Small and Rosen (1981), who developed tools to estimate consumer welfare changes in discrete-choice environments. Small and Rosen (1981) provide tools to estimate the consumer welfare changes that occur in response to a change in price, quality or any variable that varies continuously with indirect utility. In contrast, my paper focuses on changes in consideration sets that, by their very nature, provide discontinuous shifts to indirect utility functions. Thus, the results of this paper represent a significant extension of those of Small and Rosen (1981). Indeed, no simple adaption of the techniques used in their research will lead to correct welfare measures in an environment with changing, heterogeneous consideration sets.

Some recent situation-specific methods have been developed to estimate welfare changes that result from non-idiosyncratic product removal or exit. These include Nevo (2003), Gentzkow (2007), Quan and Williams (2018), and Petrin (2002). My paper provides a generalization of their results, allowing for simultaneous product entry and exit in a flexible utility environment. My results are also shown to be exactly equal to compensating variation or equivalent variation, rather than just being approximations. Finally, my methodology allows for recovery of the welfare changes that were caused by unobservable preference matching rather than just recovering the average welfare components.

In addition, a number of empirical and game-theoretic papers have studied consumer welfare in search markets. The topics studied include welfare changes as a result of search-ranking changes (e.g., Ursu (2018) and Athey and Ellison (2011)), the welfare effects of platform changes (e.g., Lewis and Wang (2013), Dinerstein et al. (2018), Fradkin (2018)), the welfare effects of advertising and search (e.g., Honka, Hortaçsu, and Vitorino (2017) and Seiler and Yao (2017)), and the welfare effects of changing search costs (e.g., Honka (2014), Ershov (2016) and Moraga-Gonzalez, Sándor, and Wildenbeest (2017)). These papers rely on strong modeling assumptions of the search process, parametric utility, and situation-specific measures of consumer welfare. My paper allows for less-restricted search behavior, less-restricted preferences, and welfare measures equivalent to the readily interpretable classical measures of compensating variation and equivalent variation.

This paper also relates to several papers within the growing literature on inattentive consumers. The literature on inattentive consumers goes back at least as far as Manski (1975). Papers in this literature classically study the economic consequences of having products missing at random from unobservable consideration sets.[6] This assumption is antithetical to consideration set formation in a search environment. However, recent advances in the literature on inattentive consumers have explored weaker restrictions on consideration set formation, making their results relevant to the results of this paper. In particular, recent advances in Barseghyan, Molinari, and Thirkettle (2019), Barseghyan, Coughlin, et al. (2019), Abaluck and Adams (2018), and Iaria, Crawford, and Griffith (2020) have weakened assumptions on how inattentive consumers form their consideration sets. Barseghyan, Coughlin, et al. (2019) find partial identification results for preferences and the distribution of consideration sets in an environment where consideration sets are essentially unrestricted. The other papers focus on the point identification of preferences and consideration set distributions, but with the additional assumption that consideration sets are independent of preferences conditional on observables.[7] Importantly, all these papers focus on identifying and estimating consumer preferences. Welfare measures in these papers, if explored, are based on differences in before-and-after-average utility and are not tied to compensating variation or equivalent variation. In contrast, this paper is—to the best of my knowledge—the first to study the identification of classically interpretable welfare measures in a setting with minimal assumptions on consideration sets.

Moreover, I provide examples in this paper that illustrate how before-and-after utility differences need not coincide with compensating variation or equivalent variation in a search environment, even when utility is linear in money. In order to identify compensating variation and equivalent variation, I assume that consideration sets are independent of prices. This restriction on consideration sets is, in general, stronger than those of Barseghyan, Coughlin, et al. (2019). While my modeling assumptions allow for a wide variety of search behavior that the other papers mentioned above cannot accommodate, the papers above generally allow consideration sets to depend on prices. Like the papers above, I do not require that consideration sets are observable for my main identification results. However, I do assume that demand is known to the researcher in my identification results and my empirical results leverage consideration set information to better estimate demand.

Finally, there is a growing body of research that is interested in assessing the value of

[6]The literature on inattentive consumers often refers to consideration sets as "choice sets." The difference between consideration sets and choice sets is only the standard assumptions under which consumers respectively form them.

[7]Abaluck and Adams (2018) explores a second environment in which consumers only shop a default product or shop all products. A thorough discussion comparing and contrasting Abaluck and Adams (2018), Barseghyan, Coughlin, et al. (2019) and other recent papers in the literature on inattentive consumers can be found at the end of Barseghyan, Coughlin, et al. (2019).

technology, the internet, and free (digital) goods and services in terms of their contributions to GDP. See, for example, Brynjolfsson, Collis, et al. (2019), Diewert and Feenstra (2017), Diewert, Fox, and Schreyer (2018), Feldstein (2017), Groshen et al. (2017), Syverson (2017), Brynjolfsson and Oh (2012) and Greenstein and McDevitt (2011). These studies use the welfare-analysis tools of Hausman (1996) or Small and Rosen (1981). Their models do not account for searching consumers. Instead, their focus is on measuring the aggregate welfare consequences of products that are available in the digital economy. In contrast, my paper allows for average welfare changes that occur as a result of idiosyncratic changes in shopping behavior.

The rest of this paper is organized as follows. In Section 2, I develop the notation and discuss the unique characteristics of welfare measures in search environments. The discussion provides an example, detailed in the appendix, where utility differences do not coincide with compensating variation or equivalent variation despite utility being linear in money. In Section 3, I present the results under a monotonicity constraint on preferences. In Section 4, I strengthen my assumptions on consumer preferences to quasi-linearity and obtain a general formula to measure the average welfare changes that occur from arbitrary search-result-list changes. I also present a simple example that illustrates the key ideas captured in the formula. In Section 5, I apply my main results to several simple search-result-list changes of practical and empirical interest. In Section 6, I apply my results to data provided by an OTA. Finally, in Section 7, I conclude. All proofs are left to the appendix.

2 Notation

In this section, I develop the notation required for the rest of the paper. I start by discussing the role of search platforms in Section 2.1. In Section 2.2, I discuss product and consumer preference notation. I develop notation and assumptions for consideration sets in Section 2.3. I define and develop notation for welfare measures in Section 2.4.

2.1 Search-Result Lists

A consumer searching for a product online will type keywords into a search platform's query box. From there, the search platform has a *listing rule* α that determines the order (or layout) in which the relevant products are listed for the consumer to review. The focus of this paper is not on how platforms determine this listing rule; the listing rule is free to depend on observable consumer characteristics and advertising concerns in addition to the consumer's keywords. Rather, this paper focuses on measuring consumer-welfare changes that are a response to changes in the listing rule, say from $\alpha = A$ to $\alpha = B$.

2.2 Preferences

My setup for products and preferences follows the multinomial choice framework with non-separable utility laid out in Bhattacharya (2015). There is an *observable*[8] *set of products* $\mathcal{J} = \{0, 1, \ldots, J\}$.[9] I denote the *observable vector of market prices* $p^m_{\mathcal{J}} = (0, p^m_1, \cdots, p^m_J)$. The price of the outside product is normalized to 0. When discussing product prices that may differ from their market-values, I drop the superscript m.

Consumer utility is affected by income y and *observable attributes* Ψ. Both y and Ψ are fixed for each individual. For readability, I suppress the notation for Ψ from utility. All identification results should be interpreted as conditional on Ψ.

Consumers have *unobservable preferences* η. I do not restrict the dimension of these unobservable preferences.[10] However, I restrict utility using the following assumptions:

Assumption 1.A (Monotonicity). *Utility for product j, $u_j(y - p_j, \eta) \in \mathbb{R}$, is strictly increasing and continuous in its first argument for all products $j \in \mathcal{J}$.*

Assumption 1.B (Quasi-linearity). *Utility for product j is linear in its first argument. That is,*

$$u_j(y - p_j, \eta) = y - p_j + \tilde{U}_j(\eta), \tag{1}$$

where $\tilde{U}_j(\eta) \in \mathbb{R}$ for all products $j \in \mathcal{J}$.

Monotonicity is a weak and intuitively appealing assumption. It rules out consumers remaining indifferent between two goods over any price interval and requires consumers to like a good less as its price increases. Quasi-Linearity is stronger but standard in empirical applications.[11]

I maintain monotonicity for the rest of the paper. Under monotonicity, income is assumed observable. Under quasi-linearity, income need not be observable.[12]

[8]The term "observable" is always used to mean observable to the researcher, not the consumer.

[9]It is without loss of generality (WLOG) to have the products invariant over the listing rules. For example, if product M becomes available after a change in the listing rule, then we can just disallow M from being in the consideration sets under the initial listing rule.

[10]See Bhattacharya (2015) for a discussion on the importance of leaving the heterogeneity dimension unrestricted in discrete-choice preferences.

[11]All formulas in this paper that hold under quasi-linearity also hold with the addition of a consumer-specific coefficient to money, a_η. That is, for my results, it is WLOG to normalize utility to the form in Equation (1) from a form where $u_j(y - p_j, \eta) = a_\eta(y - p_j) + \tilde{U}_j(\eta)$.

[12]A vector of observable product characteristics $(0, X_1, \ldots, X_J)$ can also be available to the researcher.

2.3 Consideration Sets

I assume that search is costly and thus consumers do not, in general, view all of the products returned to them in platform search lists. Instead, each consumer's product knowledge—and therefore demand—is limited to a sub-collection of $\mathcal{J}$, which is referred to as her consideration set.

Consideration sets are determined through a *consideration function* $\mathcal{C}$. The consideration function takes the following arguments: (1) a listing rule α; (2) observable consumer characteristics Ψ; (3) unobservable consumer preferences η; and (4) an unobservable non-preference characteristic vector ζ. For readability, I suppress Ψ from the consideration-set notation.

Definition 1 (Consideration Sets)**.** *A consumer with characteristics* (y, η, ζ) *has consideration set* $\{0\} \subseteq \mathcal{C}(\eta, \zeta, \alpha) \subseteq \mathcal{J}$ *under listing rule* α*.*

When discussing a fixed consumer (y, η, ζ), I abbreviate her consideration set under listing rules A and B by $\mathcal{C}_A$ and $\mathcal{C}_B$, respectively.

For a consumer (y, η, ζ), the components of ζ capture the unobservable factors that affect her consideration set but that do not enter her utility function. For example, a component of ζ captures a consumer's preference for the act of shopping, itself. Consumers who like to shop will likely have larger consideration sets than consumers who do not like to shop, even when their product preferences are identical. ζ may also capture product and price beliefs that guide the search over different keywords.

As an example, consider a consumer searching for a face moisturizer on *Amazon.com.* Suppose she does not like shopping for very long and chooses to either buy a product from the first page of search results that comes up after her keyword search or not buy anything (captured in ζ). She has a preference for branded products (captured in η) and thinks that "Biore" is likely to be a relatively inexpensive branded product (captured in ζ). Thus, she types "Biore face moisturizer" into the search box and hits the return key. Some of the products in the resulting list are organic and some are sponsored links. Her consideration set is exactly the products listed on this first page of search results, as well as the outside product.

To derive my main results, I rely on the following assumption:

Assumption 2 (Price Independence)**.** *For all products* $j \in \mathcal{J}$*, the consideration function* $\mathcal{C}$ *does not depend on prices* p_j *or income minus prices* $y - p_j$*.*

This information can be treated as being suppressed from utility for readability. If we were to suppress the product characteristics and observable individual characters from utility, we would write $u_j(y - p_j, X_j, \Psi, \eta)$ under monotonicity and $y - p_j + \tilde{U}_j(\eta; X_j, \Psi)$ under quasi-linearity.

Note that price independence does not preclude a consumer from considering products according to her *beliefs* about prices. It only requires that the prices she observes do not cause her to change her shopping behavior.[13]

Price independence rules out the possibility of a platform removing a product from its lists due to a change in its price. It also rules out certain consumer behaviors, such as a consumer expanding her consideration set after discovering that all the products in her initial consideration set are unexpectedly expensive. In practice, some price dependence can be tolerated: the key behavior needed is that, as a good's price increases, a consumer's demand for that good falls to 0 for preference reasons before the price increases causes the good to exit (or another good to enter) her consideration set. That is, as long as the consumer's preferences are more sensitive than the search listing rule, the main results should still hold.[14]

A consumer (y, η, ζ) purchases product j in $\mathcal{C}(\alpha, \zeta, \eta)$ if

$$u_j(y - p_j^m, \eta) > u_k(y - p_k^m, \eta) \text{ for all } k \in \mathcal{C}(\alpha, \zeta, \eta) \setminus \{ m \} .$$

Her *individual demand* is defined by

$$q_j(y, p_j, p_{-j}, \alpha, \eta, \zeta) := \begin{cases} 1 \text{ if } j = \arg\max_{\ell \in \mathcal{C}(\alpha,\eta,\zeta)} u_\ell(y - p_\ell, \eta) \\ 0 \text{ otherwise.} \end{cases} \tag{2}$$

If the price of all products except good j are fixed at their market level, then I denote her *individual demand for product j at market prices* by

$$q_j^m(y, p_j, \eta, \zeta, \alpha) := q_j(y, p_j, p_{-j}^m, \alpha, \eta, \zeta). \tag{3}$$

I will denote the joint distribution of η and ζ by F. Thus, the *average demand for product j* is

$$Q_j(y, p_j, p_{-j}, \alpha) = \int q_j(y, p_j, p_{-j}, \alpha, \eta, \zeta) dF \tag{4}$$

[13] Although not modeled explicitly, it is perfectly fine for a product's non-price characteristics X_j to affect consideration sets. In addition, it is perfectly fine for a consumer's consideration set to depend on her perceptions of her own income level or wealth level, as a part of Ψ. The only complication that might arise here is if a consumer's final product choice changes her perception of her own wealth level. By assumption, this is not allowed.

[14] Mathematically, if individual $i = (y, \eta, \zeta)$ no longer includes good j in her consideration set for prices above $\bar{p}_j^i$ and would not purchase good j at prices above $\tilde{p}_j^i$ even if good j were in her consideration set, then $\bar{p}_j^i > \tilde{p}_j^i$ is sufficient for the identification results to hold for consumer i in the case of the platform removing good j in an A/B test.

for all consumers with income y. Similarly, the *average demand for product j at market prices* is,

$$Q_j^m(y, p_j, \alpha) = \int q_j(y, p_j, p_{-j}^m, \alpha, \eta, \zeta) dF \tag{5}$$

for all consumers with income y.

For my identification results, I assume demand is known to the researcher. That is, I follow Bhattacharya (2015) in making the following assumption:

Assumption 3 (Known Demand)**.** *The researcher observes Equation* (4)*, and therefore also observes Equation* (5)*, over all prices and incomes.*

This assumption requires that the researcher observe demand for each good over all prices and incomes. It does not require that the researcher know individual's consideration sets.

This is a substantial assumption. Nonparametric demand identification is itself a challenging problem, particularly due to the problem of price endogeneity. See the corresponding discussion of nonparametric identification in Bhattacharya (2015), in Matzkin (2013), or Matzkin (2007) for the challenges of nonparametrically identifying demand in settings without consideration sets. See Koulayev (2014) and Abaluck and Adams (2018) for recent work on the nonparametric identification of demand in a setting with restricted consideration sets.

For the empirical application of my identification results, I drop the assumption of known demand and estimate a parametric demand function with separable errors. I take advantage of information on individual consideration sets in my demand estimation strategy, but the researcher need not observe consideration sets to apply the identification results in general.

2.4 Welfare Measures

Suppose a platform changes its search listing rule from A to B. The following definitions adapt the classic measures of welfare changes, *compensating variation* and *equivalent variation*, to this situation. For an individual consumer (y, η, ζ), her equivalent variation S^{EV} is the solution in S to

$$\max_{j \in \mathcal{C}_A} u_j(y - S - p_j^m, \eta) = \max_{j \in \mathcal{C}_B} u_j(y - p_j^m, \eta), \tag{6}$$

while her compensating variation S^{CV} is the solution in S to

$$\max_{j \in \mathcal{C}_A} u_j(y - p_j^m, \eta) = \max_{j \in \mathcal{C}_B} u_j(y + S - p_j^m, \eta). \tag{7}$$

S^{EV} is the income loss under the initial listing rule that would harm this consumer as much as the damage done by the new listing rule. S^{CV} is the increase in income under the new listing rule that would return a consumer to the utility level she would have had under the original listing rule. S^{EV} and S^{CV} are both positive if the consumer is harmed by the new listing rule, relative to the older rule, and negative otherwise.

Both S^{EV} and S^{CV} depend on market prices and both listing rules, as well as the individual's unobservable characteristics η and ζ. In the case of monotonicity, both S^{EV} and S^{CV} also depend on individual income. Thus, I denote the functions for equivalent variation and compensating variation by $S^{EV}(y, \eta, \zeta, A, B, p^m_{\mathcal{J}})$ and $S^{CV}(y, \eta, \zeta, A, B, p^m_{\mathcal{J}})$, respectively. Similarly, *average compensating variation* and *average equivalent variation* over all consumers are denoted by μ^{CV} and μ^{EV}. As functions, I write these as $\mu^{CV}(y, A, B, p^m_{\mathcal{J}})$ and $\mu^{EV}(y, A, B, p^m_{\mathcal{J}})$; the average is taken over unobservables η and ζ, while income is fixed.[15] When it is clear from the context, I suppress the arguments for the listing rules and prices from these welfare functions.

Under quasi-linearity and price independence, each individual (y, η, ζ) has $S^{CV} = S^{EV}$. In this case, I use S^W to represent both S^{CV} and S^{EV}. S^W is simply the difference between utility under the initial listing rule and utility under the final listing rule.

Importantly, quasi-linearity is not sufficient for differences in utility to be equal to compensating variation or equivalent variation. That is, for a generic search model where price independence does not hold,

$$S^{CV} \neq \max_{j \in \mathcal{C}_A} u_j(y - p^m_j, \eta) - \max_{j \in \mathcal{C}_B} u_j(y - p^m_j, \eta) \neq S^{EV}$$

in general.

When consideration sets depend on prices, traditional intuition about welfare and utility breaks down. For example, when consideration sets depend on prices, a small increase in a price could send the consumer searching for new products and, ultimately, finding a product that gives her higher utility than those in her original consideration set. That is, a consumer who experiences a small price *increase* may experience a welfare *increase* that would require a compensating income *reduction* to return her to her original utility level. A more detailed example is provided in Section B of the Appendix. This example, along with the other discussions in this paper, demonstrates the importance of reevaluating traditional welfare measures in search environments.

[15] Averages would still be functions of consumer observables Ψ, although the researcher could, of course, average over Ψ (and y) if she desired. Switching the integral order to accommodate earlier averaging is never a problem, either, since demand is nonnegative.

3 Measuring Welfare Under Monotonicity

This section considers welfare identification under monotonicity. Below, I present the bounds on compensating variation that result from a search-listing change that makes a collection of goods more likely to be considered.[16] To ease notation in the statement of the theorem, I introduce the following notation. Fix A and B and define operator Δ as follows:

$$\Delta Q(y, p_J^m) \equiv Q_j(y, p_J^m, A) - Q_j(y, p_J^m, B), \tag{8}$$

Theorem 1 (Probable Addition of Goods)**.** *Suppose the listing rule changes exogenously from A to B causing a collection of goods $\mathcal{R}$ to be more likley to enter consideration sets.*[17] *Then, the average compensating variation of this change in listing rule is such that*

$$0 \geq \mu^{CV} \geq \sum_{j \in \mathcal{R}} \lim_{p_k \to \infty \forall k \in \mathcal{R} \setminus \{j\}} \int_{p_j^m}^{\infty} \Delta Q_j(y, p_j, (p_{\mathcal{R}\setminus\{j\}}, p_{-\mathcal{R}}^m)) dp_j. \tag{9}$$

The appendix contains the proof, as well as the analogous bounds for equivalent variation under an exogenous search-result list change that causes a decrease in the probability that a collection of goods enters consideration sets.

Intuition for the bounds in Theorem 1 can be developed from the following observations. First, a consumer's preference for a good depends on which other goods are in her consideration set. To fix ideas, I focus on a consumer's demand for good 1. I will assume good 1 is initially her favorite among those in her consideration set. However, as the number of goods in her consideration set grows, her willingness to pay for good 1 falls weakly. In particular, if a new good 2 enters her consideration set and she prefers 2 to 1 or her next best choice after 1, her demand line for good 1 will shift downward. Otherwise, no change will occur. The same is true for each additional new good that enters her consideration set.

Second, as the price of a good goes to infinity, it becomes so undesirable that it is "as

[16] Bhattacharya (2018), which extends several results from Bhattacharya (2015), finds the distribution of welfare when a collection of goods have price increases or decreases. Sending price to infinity for a collection of goods in his setting finds the welfare consequences of removing this collection of products for *all* consumers. (His setting is not a search setting.) In constrast, Theorem 1 and Theorem 4 below allow an arbitrary collection of consumers to continue shopping the good while another collection of consumers has the good removed. That is, the results below are more general in that they allow for some some consumers to discover the new product and some consumers to fail to discover the new product, rather than have all consumers discover the new product. Improving the inequalities in Theorem 1 to equalities can be achieved if we increase the econometrician's knowledge to a level effectively the same as Bhattacharya (2018) in a search setting: if we assume all consumers' consideration sets are observable and demand conditional on these consideration sets is observable over all prices, in addition to maintaining my above assumptions on consideration sets and monotonicity of preferences in money.

[17] Specifically, I mean for every (y, η, ζ), $\mathcal{C}_A \subseteq \mathcal{C}_B \subseteq \mathcal{C}_A \cup \mathcal{R}$.

if" the good were not even included in the consumer's consideration se. Thus, we can simulate the removal of several goods from a consumer's consideration set by examining how her demand changes as the price of these goods goes to infinity. For more intuition in the simpler case of quasi-linearity, see the example in Section 4.

Collectively, the two intuitions above give us an interpretation of the lower bound of average compensating variation as follows. Fix a good j among those being summed over in the rightmost expression of Equation (9). The terms to the right of the sum for that good j then find the total value of increasing the probability of shopping that good j when all of the other goods from $\mathcal{R}$ are not considered. Thus, the sum over all goods in $\mathcal{R}$ in the right-most expression of Equation (9) sums the total welfare contribution of increasing the probability that each good in $\mathcal{R}$ is individually considered. Since the collection of goods in $\mathcal{R}$ are all substitutes, the value of each good added individually is less than the making the collection available.

When $\mathcal{R}$ is a singleton, the above bounds can be refined to a single point. In this case we have the following corollary:

Corollary 1 (Probable Addition of a Single Good)**.** *Suppose the listing rule changes exogenously from A to B in a way that makes each consumer more likely to shop for product M.*[18] *Then, the average compensating variation of this change in listing rule is*

$$\mu^{CV} = \int_{p_M^m}^{\infty} \Delta Q_M(y, p_M, p_{-M}^m) dp_M. \tag{10}$$

The proof of Corollary 1 is in the appendix. However, the results are intuitively appealing. They are similar to classic results on the welfare of a new good. In the appendix, I include analogous results that relate the average equivalent variation in the case of a probable removal of a good.

Note that the results of Theorem 1 and Corollary 1 hold regardless of how many consumers' consideration sets hold good from $\mathcal{R}$ or M respectively under listing rule A.

4 Measuring Welfare Under Quasi-linearity

In this section, I determine how to measure welfare changes as a response to search-listing changes. I start with my most general result: a formula that measures welfare changes from aggregate demand lines under quasi-linearity. I leave the proof for the appendix, but follow up with an example that shows the key ideas. I then present some simpler formulas that

[18]Specifically, I mean for every (y, η, ζ), $\mathcal{C}_B$ is equal to $\mathcal{C}_A$ or $\mathcal{C}_A \cup \{ M \}$.

can be applied to specific cases of welfare changes.

For succinctness, I first denote *total consumer welfare under listing rule* α by Ω_α and define it with the following formula:

$$\Omega_\alpha := \lim_{p_2,\ldots,p_J\to\infty} \int_{p_1^m}^{\infty} Q_1(y,p,p_{-1},\alpha)dp \;\; + \tag{11}$$
$$\lim_{p_3,\ldots,p_J\to\infty} \int_{p_2^m}^{\infty} Q_2(y,p,(p_1^m,p_{-(1,2)}),\alpha)dp + \cdots + \int_{p_J^m}^{\infty} Q_J(y,p,p_{-J}^m,\alpha_m)dp.$$

Total consumer welfare under listing rule α—hereafter abbreviated as total welfare—captures the total value to consumers of the products in their (heterogeneous) consideration sets under listing rule α. This total value is relative to the outside good—no purchase. More precisely, the first term in the sum that defines Ω_α calculates the average value of allowing product 1 to enter all consumers' consideration sets; if a consumer does not have product 1 in her consideration set under listing rule α, then this consumer's contribution to the average is 0. This added value is relative to consideration sets that only contain the outside product. The second term in the sum that defines Ω_α adds the average value consumers gain by having product 2 in their consideration sets, relative to consideration sets that contain (at most) product 0 and product 1; consumers without product 2 in their consideration set contribute nothing to this value. Consumers with product 2 but not product 1 in their consideration set will contribute average values to the 2nd term that reflect product 2's value relative to the outside product alone. This process of adding in the average value of one more product is continued from the third term until the Jth term in the sum. By the Jth term, all of the consideration sets will have reached their full size under listing rule α. All the terms together give Ω_α, the total value of sequentially allowing products 1 to J to enter all of the consideration sets.

Lemma 1 below provides some deeper intuition. Fix a consumer (y,η,ζ) and let ω_α denote her "total welfare." That is,

$$\omega_\alpha(y,\eta,\zeta) := \lim_{p_2,\ldots,p_J\to\infty} \int_{p_1^m}^{\infty} q_1(y,p,p_{-1},\alpha,\eta,\zeta)dp \;\; + \tag{12}$$
$$\lim_{p_3,\ldots,p_J\to\infty} \int_{p_2^m}^{\infty} q_2(y,p,(p_1^m,p_{-(1,2)}),\alpha,\eta,\zeta)dp \;\; + \cdots +$$
$$\int_{p_J^m}^{\infty} q_J(y,p,p_{-J}^m,\alpha_m,\eta,\zeta)dp.$$

Then we have the following result:

Lemma 1. *Under quasi-linearity,*

$$\omega_\alpha(y, \eta, \zeta) = \max_{j \in \mathcal{C}_\alpha} u_j(y - p_j^m, \eta) - u_0(y, \eta).$$

This result is proved in the appendix as part of my proof of Theorem 2. Theorem 2 is stated below. Mathematically, Equation (12) becomes a telescoping sum of the utility differences from the utility of the outside good up to the utility of the good the consumer most prefers in her consideration set. The utilities for all goods between the outside good and her most preferred good will difference out. Aggregating Lemma 1 into a statement about average demand lines is straightforward and yields the following conclusion.

Theorem 2. *Under quasi-linearity, the average welfare change μ^W that occurs as a result of a change in the platform listing rule from A to B is*

$$\mu^W = \Omega_A - \Omega_B.$$

In words, by looking at the difference in the total welfare created by a change in the listing rule, we can recover the exact average compensating variation. (The exact average compensating variation is also the exact average equivalent variation, since the terms coincide under quasi-linearity.) Of course, under quasi-linearity, the numbering of products 1 to J can be arbitrarily rearranged and the formula still holds.

Theorem 2 captures several key ideas about welfare analysis in a search environment. The first key is the use of prices. By raising a product's price high enough—beyond consumers' reservation prices—we can effectively turn off the value consumers gain from considering that product. The second key is that, under quasi-linearity, the total value consumers gain from their purchases under a given listing rule can be calculated by a sum across all demand lines. Thus, it is not necessary to know each consumer's idiosyncratic consideration set; it suffices to know the aggregate demand curves.

The final key is the importance of the reference product, product 0. Theorem 2 works by building up the utility around the outside product under the different listing rules. If the utility of the outside product is not comparable across listing rules, then we cannot hope to make meaningful welfare comparisons between the two outcomes. If there is greater homogeneity in the content of the consideration set across the listing rules, then a larger reference collection of products can be used and the welfare formula can be simplified, as illustrated in Section 5.

The proof of Theorem 2 can be found in the appendix. However, the following example captures many of the key ideas at work.

Example: Consider a market with a single consumer (y, η, ζ) who considers a single product, product 1, along with the outside product under listing rule A. That is, this consumer has consideration set $\{0, 1\}$ under listing rule A. Under listing rule B, her consideration set grows to $\{0, 1, 2, 3\}$; she gains *two* additional products in her consideration set. For simplicity, assume all market prices are 0 and that utility has the following form:

$$\begin{aligned} u_0(y, \eta) &= y \\ u_1(y, \eta) &= y + a \\ u_2(y, \eta) &= y + 10a \\ u_3(y, \eta) &= y + 10a + \epsilon, \end{aligned}$$

where a and ϵ are positive.

Then, we see that the consumer's product choice under listing rule A is 1 and her product choice under listing rule B is 3. Her change in utility is $S^W = u_1(y, \eta) - u_3(y, \eta) = -9a - \epsilon$.

Similarly, since $Q_1(y, p, p^m_{-1}, A) = \lim_{p_2, p_3 \to \infty} Q_1(y, p, p_{-1}, B)$ and $Q_2(y, p, p_{-2}, A) = 0 = Q_3(y, p, p_{-3}, A)$ for all price vectors,

$$\begin{aligned} \Omega_A - \Omega_B &= -\lim_{p_3 \to \infty} \int_{p_2^m}^{\infty} Q_2(y, p, (0, p_3), B) dp - \int_{p_3^m}^{\infty} Q_3(y, p, p^m_{-3}, B) dp \qquad (13) \\ &= -\int_0^{\infty} 1(9a > p) dp - \int_0^{\infty} 1(\epsilon > p) dp \\ &= -9a - \epsilon, \end{aligned}$$

as claimed in Theorem 2.

In the first integral of Equation (13) above, taking the price of product 3 to infinity makes the consumer behave *as if* product 3 were not in her consideration set. This allows us to measure the value of product 2's addition to her initial consideration set of $\{0, 1\}$. The second integral of Equation (13) then takes account of the value of adding product 3 to a consideration set of $\{0, 1, 2\}$. For more intuition, consider Figure 1 and Figure 2.[19]

[19]Note, if the pricing limit is not included, then we have $\int_0^{\infty} Q_2(y, p, \mathbf{0}, B) dp = 0 < \lim_{p_3 \to \infty} \int_0^{\infty} Q_2(y, p, (0, p_3), B) dp = 9a$. That is, the pricing limits are essential for obtaining the correct welfare conclusions.

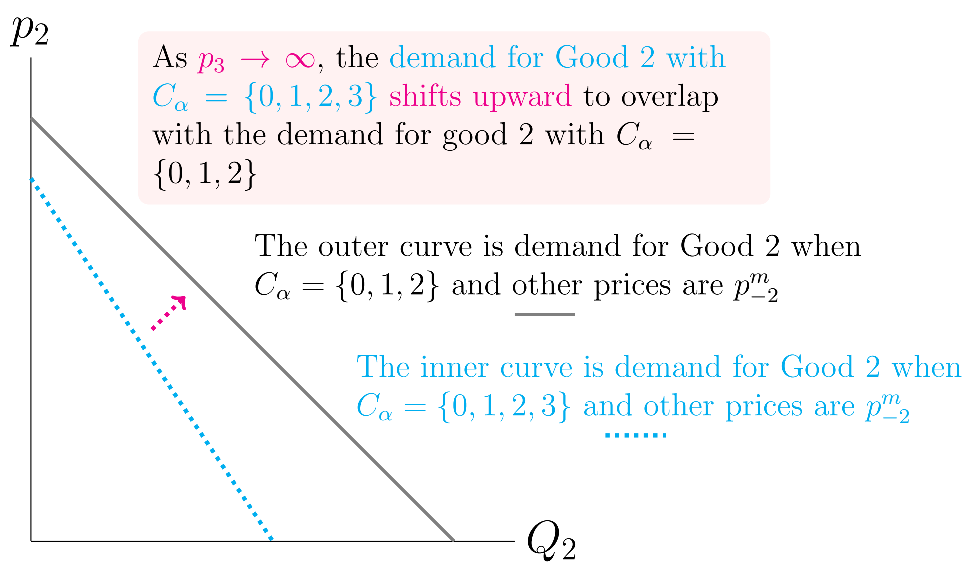

Figure 1: *As the Price of Good 3 Increases, the Demand Curve for Good 2 Shifts Outward Until it is as if no Consumer Considers Good 3.* Both the dotted, blue and solid, black lines show the relation between the average quantity demanded of product 2, Q_2, against the price of product 2, p_2, when the prices of all other products are at market prices. The blue line captures demand under search listing rule B, when the consideration sets are all $\{0,1,2,3\}$, whereas the black line captures the counterfactual demand such that consideration sets are all $\{0,1,2\}$. Necessarily, the blue line is (weakly) below the black line at all quantities. If the market price of good 3 went to infinity, then the blue line would shift up to become exactly equal to the black line, as pictured. This figure is the author's own diagram.

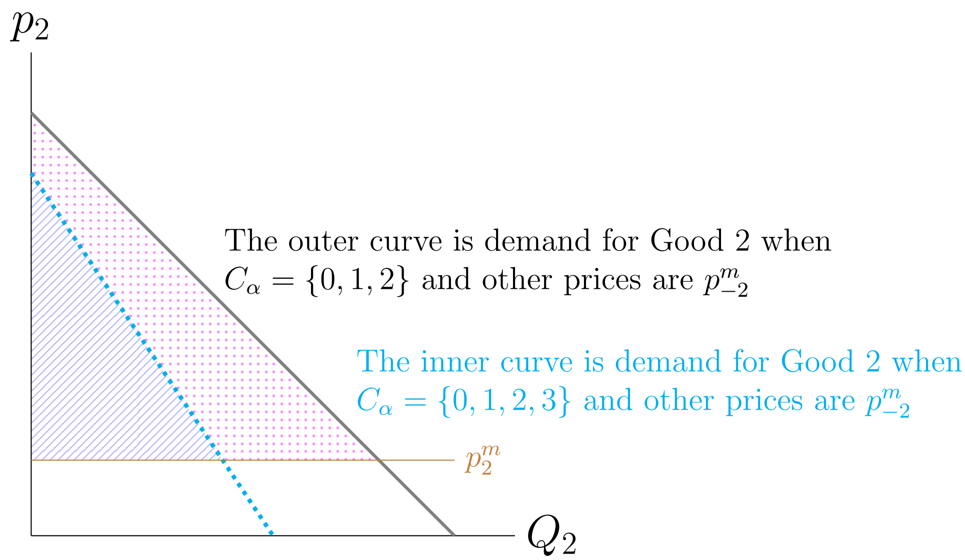

Figure 2: *Calculating the Welfare Benefit of Considering Good 2 and Good 3 Requires A Pricing Limit for Good 3.* The northeast, navy lines indicate the area between the market price for good 2 and the average demand curve for good 2 when all goods are at their market prices. The dotted, blue line is the demand curve for Good 2 when the platform listing rule B is in effect. Under search listing rule B, each consumer's consideration set is $\{0, 1, 2, 3\}$. The solid, black line is the counterfactual average demand for good 2 when all consumers consider $\{0, 1, 2\}$. Under listing rule A, all consumers consider $\{0, 1\}$. The average value that consumers gain in going from listing rule A to B would include the area indicated by the northeast, navy lines and the area indicated by the pink dots. The sum of these areas captures the average welfare gained by consumers who can now consider good 2; that is, the sum captures the gain in going from $\{0, 1\} \rightarrow \{0, 1, 2\}$. To add the welfare gained from going from $\{0, 1, 2\}$ to $\{0, 1, 2, 3\}$, the researcher would add the area under the average demand curve for good 3 when all other goods are constrained to their market prices (not pictured). This figure is the author's own diagram.

5 Welfare Results Under Simple Listing Rule Changes

In this section, I look at formulas for measuring the changes in welfare that result from simple changes in the listing rules. I demonstrate that the calculations required in Theorem 2 can be simplified in many situations of practical and counterfactual interest. I consider listing rules that (1) swap a single product in a search-result list with a new product and (2) add or remove a collection of products from a search-result list. While case (2) was already considered in Section 3, the stronger assumption of quasi-linearity allows me to strengthen the bounds of Theorem 1 into equalities.

5.1 Single Product Swap

In this section, I present measures for the welfare change that results from swapping two goods in the search-result list. That is, if good 1 is being swapped for good M, good M is not included in search-result lists under rule A while good 1 is not included in search-result lists under rule B.

Theorem 3 (Probable Product Swap)**.** *Under quasi-linearity, when the listing rule changes from A to B such that good 1 is swapped for good M, then the average welfare change is*

$$\mu^W = \int_{p_1^m}^{\infty} Q_1(y, p, p_{-1}^m, A)dp - \int_{p_M^m}^{\infty} Q_M(y, p_M, p_{-M}^m, B)dp_M.$$

In Theorem 3, good 1 and good M are not contemporaneously present in any individual's consideration set. The welfare comparison can still be made with the $\mathcal{C}_A \setminus \{1, M\} = \mathcal{C}_B \setminus \{1, M\}$ being the reference collection for each consumer. This is why quasi-linearity is essential for Theorem 3 or any time a change in the search-result list simultaneously increases some products' probability of being shopped while decreasing other products' probability of being shopped. The proof is provided in Section B.5.

5.2 Adding a Collection of Goods to Search-Result Lists

I conclude this section with a formula for measuring welfare changes that result when a collection of goods is added to search-result lists. These results are also applicable when the cost of searching products (weakly) decreases for each consumer. The results for the removal of a collection of goods from the search-result lists are analogous and included in the appendix.

First, for a collection of products $\mathcal{R} = \{ r_1, \ldots, r_R \}$, define the *total consumer value of products in* $\mathcal{R}$ by

$$\Gamma_\alpha(\mathcal{R}) := \lim_{p_{r_2},\ldots,p_{r_R}\to\infty} \int_{p^m_{r_1}}^{\infty} \Delta Q_{r_1}(y, p, p_{-r_1}, \alpha)dp \ + \lim_{p_{r_3},\ldots,p_{r_R}\to\infty} \int_{p^m_{r_2}}^{\infty} Q_{r_2}(y, p, (p^m_{r_1}, p_{-(r_1,r_2)}), \alpha)dp + \cdots + \int_{p^m_{r_R}}^{\infty} Q_{r_R}(y, p, p^m_{-r_R}, \alpha)dp.$$

Note that $\Gamma_\alpha(\mathcal{J}) = \Omega_\alpha$.

Then, the following formula can be used to calculate the exact welfare changes that result from changes in the search-result list that weakly increase the probability that each good in $\mathcal{R}$ enters the consideration sets.[20]

Corollary 2. *Let* $\mathcal{R} = \{ r_1, \ldots, r_R \}$. *Suppose that a change in the listing rule from A to B increases the probability that each of the products in* $\mathcal{R}$ *is considered. Then, under quasi-linearity,*

$$\mu^W = \Gamma_A(\mathcal{R}) - \Gamma_B(\mathcal{R}).$$

The proof is provided in the appendix.

6 Data Application

6.1 Data Overview

In this section, I estimate welfare changes from listing rule changes for a data set that details the click and purchase behavior of a collection of consumers booking hotels using an online travel agent (OTA). The data is from the 2013 data challenge for the IEEE's International Conference on Data Mining (ICDM).[21] The competition was open to the public through the online data science community Kaggle. Data for the contest was provided by Expedia.

The data is centered on a collection of *search impressions* that OTA users interacted with, primarily on *Expedia.com*. To understand the term search impression, first consider a consumer searching on *Expedia.com* for vacation accommodations in 2013. This consumer would initially face a page as pictured in Figure 5. Here, the consumer would enter her vacation destination, the days she planned to spend at her vacation destination, the number of rooms she wanted to book and the number of adults and children that she will be traveling

[20] This result does not allow for the displacement of other goods.

[21] IEEE is the Institute of Electrical and Electronics Engineers. The ICDM is considered the world's premier research conference in data mining. Data challenges are typically held annually.

with. All this information, pictured in the blue boxes in Figure 5, is collected by Expedia and used to produce a sequence of listings of available hotel rooms. The user is promptly directed to this listing sequence upon entering her information and clicking the button that says "Search for Hotels."

An example of a single hotel listing is given in Figure 6a. The blue boxes identify the information that Expedia collects and that is included in the data set for each hotel listing. Each search is likely to produce several listings. The number of individual listings will vary, depending upon the destination city, the availability of hotel rooms on the given date, and the presence of advertising. In the data, the number of listings on the first page of results varies between 1 and 34. A *search impression* is then defined to be the first page of search results for a given user query. For a large fraction of searches, the number of relevant listings is greater than the length of the search impression. This fact is suggested by the high frequency of long search distributions; see Figure 3.

When a consumer clicks on a listing, a new page opens and provides more details about the property's available rooms. In particular, when a user clicks on a hotel listing, she receives information on the available hotel such as the size of the beds, the parking fees, some pictures of the room interiors, the availability of free breakfast, the room amenities and any hidden fees. An example of this final page is shown in Figure 6b. From here, the consumer may choose to book with the hotel.

The data provides information on the entire first page of the search results each consumer faces. It also tells us the listings each consumer clicked on as well as the listing the consumer eventually booked (if any). Thus, if we assume the products that enter a consumer's consideration set are exactly the products in a consumer's search impression, then we can readily estimate demand. Moreover, Expedia provides us with data from one of their experiments: this data set includes search impressions where the hotel-listing order was determined by Expedia's proprietary ranking algorithm as well as search impressions that resulted in listings being ordered randomly over the pages of results. This provides us with an excellent opportunity to study the welfare consequences of moving from random rankings to Expedia's proprietary ranking system.

While the data set provides an excellent opportunity to study the relationship between listing rules and consumer welfare, a few important caveats should be pointed out. First, search impressions only list the first page of results for each user query. Thus, a consumer who searches beyond the first page of listing results (should there be additional listings) will not have her full consideration set observed in this study.[22] Her behavior on the second

[22] Of course, Expedia would have been able to tell this in their own data. This information had simply been left out of the competition data.

page of results would be treated as a separate (unassociated) search impression, if included at all. The same would be true for a consumer who searched over multiple start and end dates. Thus, to the extent that these consumers viewed multiple search-result pages or considered alternative booking dates, the results of this study will underestimate the size of the individual OTA user's consideration sets. Ursu (2018) provides some evidence from a companion data set[23] where more than 40% of *Expedia.com* users only look at the first page of results. Thus, for a large fraction of *Expedia.com* users, it would be reasonable to assume that each consumer has a consideration set that is exactly equal to her search impression. Second, for competitive reasons, Expedia would not verify how representative the sample was. However, Ursu (2018) was able to verify that the Expedia data set was representative of the largest shopping groups on Expedia, except for transactions that lead to sales being oversampled. Thus, the results should be interpreted as averages for this collection of typical consumers with more serious intent to purchase rather than averages over all consumers. Ursu (2018) provides extensive justification for studying and comparing the two groups; see the discussion therein for more details.

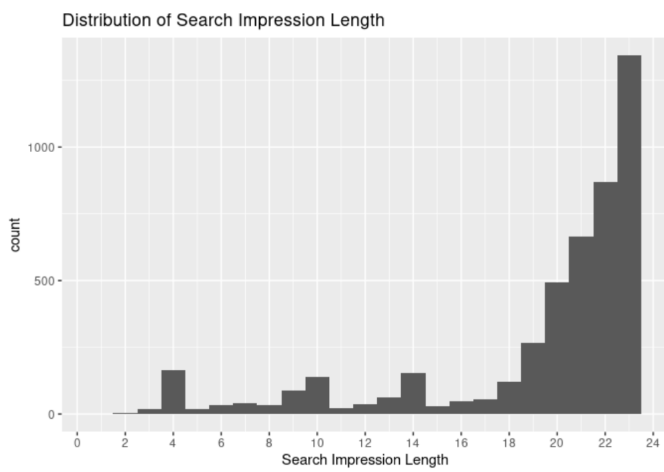

Figure 3: *The Distribution of Search Impression Length After Data Cleaning.* This figure shows the distribution of the length of search impressions over all users and listings in the data provided by Expedia after the author's data cleaning. The data contains a sampling of search impressions of OTA users from 2012 and 2013. The uncleaned data is available to the public on *Kaggle.com*. The author generated this figure using R.

[23]The companion data set is from the Wharton Customer Analytics Initiative and contains several statistics on *Expedia.com* users.

6.2 Demand Estimation Strategy

I fit a model of product choice where, for product j in consumer i's consideration set at time t,

$$u_{ijt} = \alpha(y - p_{ijt}) + \beta' X_j + \eta_{ijt}$$
$$u_{i0t} = \alpha y + \eta_{i0t}$$

Here, η_{ijt} is a standard Type I extreme value distributed random variable that is independent over j and t, given i's consideration set. The vector X_j contains the product characteristics of good j. Income y is not observed but is not needed since it differences out of product decisions. I assumed that each search impression was a unique user and that consideration sets were exactly the products listed in the search impression.[24] For simplicity of analysis, I assumed that the prices and product characteristics are independent of η_{ijt}.[25] In order to ensure my demand parameters could be estimated, I dropped all of the properties that were chosen fewer than 50 times and all search impressions where these infrequently booked properties were chosen. As there were a small number of observations where the recorded prices were much higher than the actual prices consumers observed (as discussed in Ursu (2018)), I also drop all bookings with a listed nightly price above \$1500USD and all search ids that choose a booking with a recorded price over this amount; there are only four of the latter and a handful of the former. In the end, 90,474 rows containing 4,694 search impressions remained; 1,775 impressions are from the random listing rule and the remaining 2,919 impressions are from the proprietary listing rule. Consideration sets, not including the outside product, have an average length of 19.27 listings and a median length of 21 listings; see Figure 3. The outside product was chosen 38.6% of the time. The average price of a room is \$129.29 per night.

I chose components of X_j with the aid of previous studies that looked into the covariates of hotel booking choice in this data set. Following the results of Liu et al. (2013), I included property star ratings, property branding information, a property-location score[26] and an

[24]This is justified in the previous subsection.

[25]Alternatively, this assumption was justified in Ursu (2018) by including product fixed-effects and check-in-time fixed effects in the regression and appealing to standard pricing practices in the OTA industry. When I run a more highly-specified regression, with check-in-week-fixed effects and product-fixed effects (including the latter necessitates the removal of star rating due to collinearity), I obtain a price coefficient of -.012. This is very similar to the price coefficient from my simpler regression. For ease of interpretation, I proceed with the simpler regression for the rest of the paper's results. Full results under the alternative fixed-effects-rich model are computed in the R code accompanying this paper.

[26]This is calculated by Expedia, using information about the user's query and the property's locations.

indicator variable for promotions. I ran the regression in R (R Core Team 2017), using the `mlogit` package (Croissant 2019). The results are shown in Table 1. I find that all of the included coefficients are highly significant. As expected for demand, the coefficient on price is negative and highly significant.[27] A picture of the demand for the most popular booking under a variety of deviations in the price of the second most popular booking is shown in Figure 4. Figure 4 shows that there is little difference between the demand for the top booking when the price of the second most popular booking is $166 above its market price and the demand for the top booking when the price of the second most popular booking is infinity.

Table 1: *Demand Parameter Estimates and Standard Errors for Online Searches and Bookings of Hotel Rooms in Expedia data set (2012-2013).* This table includes demand parameter estimates for the regression outlined in Section 6.2. Standard errors are listed in parenthesis, below coefficient estimates. All coefficient estimates are, statistically, highly significant.

	Dependent variable:
	Hotel Booked
property star rating	0.513
	(0.048)
property brand boolean	0.418
	(0.053)
property location score 1	−0.922
	(0.043)
price in USD	−0.010
	(0.001)
promotion flag	0.266
	(0.047)
Observations	4,694

[27] I performed an alternative regression, with product fixed effects, the promotion flag, check-in-week fixed effects, price and price interacted with property star ratings. I find the price and property star rating interaction term does not have a highly significant coefficient, which supports the assumption that utility is linear in money. Details of this regression and its output are included in the R code accompanying this paper.

Figure 4: *Demand for the Top Booking Under Variation in the Price of the Second Most Popular Booking.* These figures show the demand for the most popular booking in the data. Each line is drawn given a price deviation of the second most popular booking. The standard deviation of the second most popular booking, denoted in the figures by σ, is 33.14. (More summary statistics are available in Table 2.) Deviations in prices are from their observed market prices. The y-axis for both figures denotes the price increase in the most popular product above it's market price.

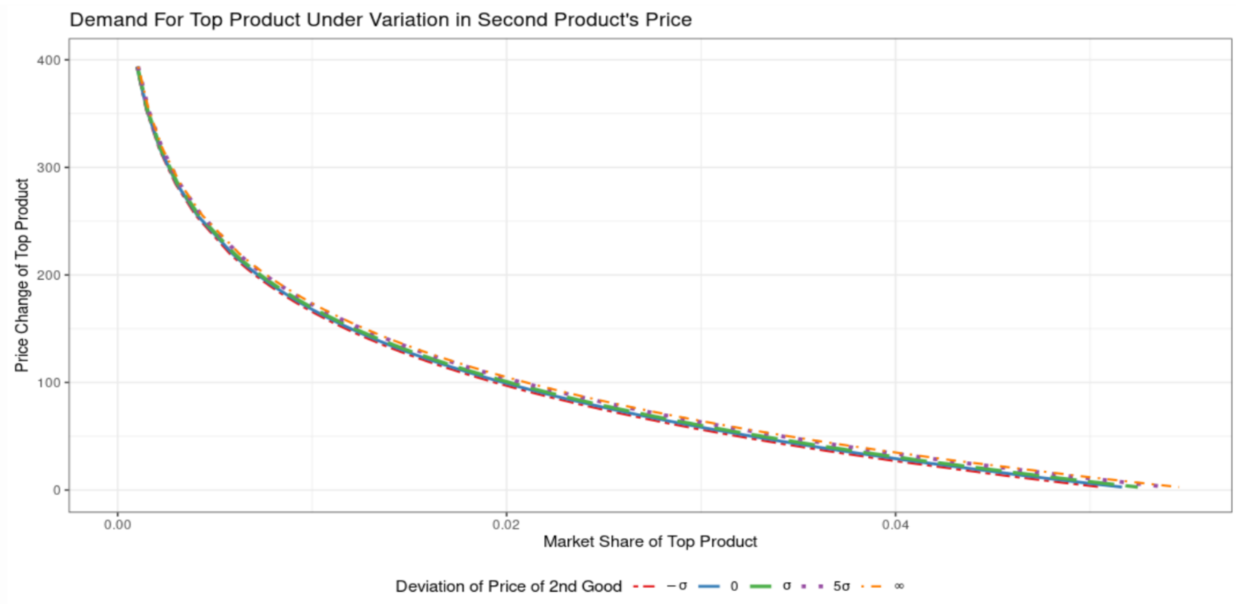

(a) *Demand for the Top Booking Under Variation in the Price of Second Most Popular Booking.*

(b) *Zoomed View of the Demand for the Top Booking Under Variation in the Price of the Second Most Popular Booking.* This figure zooms in on the figure in panel a.

6.3 Welfare Change from Random Rankings to Purchase Rankings

As discussed in the data overview, the data contains information over two different listing rules. The first listing rule is a "random" listing rule. Under the random listing rule, the list order is filled in using an *almost* random ordering of products matching consumers' first page selections. It is almost random because some of these listing positions are reserved for sponsored search impressions.[28] The second listing rule ranks products using a proprietary algorithm known to Expedia. The rule ranks products (at least in part) by their relevance (or probability of purchase). Indeed, the goal of the data challenge was to produce an algorithm that could learn to rank the products in order of their purchase-and-click likelihood.[29]

In order to estimate the change in welfare between the two listing rules, I used demand estimates from the previous section. Given the search impressions observed over the two listing rules, I predicted average demand for each good under each listing rule and over all prices. Using Equation (11) and Theorem 2, I found the total welfare under the random-listing rule to be \$96.77 and the total welfare under the proprietary-ranking rule to be \$104.88.[30] Thus, I conclude that welfare was improved by an average of \$8.11 per person when the listings were ordered by the proprietary-ranking rule. Intuitively, this is appealing. While there are valid reasons for concern about the welfare harm that could come from manipulating search-result lists, it is also true that a well-ordered list can improve welfare over a randomly ordered one.

In order to estimate a confidence interval for the welfare change, I performed 500 bootstrap resamples of the data. Each bootstrap resampling drew 4,694 impressions—the same number as total in the data—randomly and with replacement. For each bootstrap resampling, I recalculated demand and then re-estimated the total change in welfare. I then use the .025 and .975 quantile estimates of the bootstrapped estimates as the bounds of a 95% confidence interval on the welfare benefits of the proprietary listing rule. These bounds are from \$6.11 to \$9.74; \$8.11 is a statistically significant estimate.

Given my assumption that consumers only look at the first page of search results, price independence will hold reasonably well if, as the price of a good increases, the OTA's listing rule does not remove this good from the first page of search results before most consumers

[28] For this paper's study, it is not essential the ranking be perfectly random. The key is that this (imperfectly) random listing rule has welfare consequences that are measurably different from the alternative listing rule.

[29] These algorithms are called LeToR algorithms and the data science community has a literature around them.

[30] Note that Theorem 2 holds when there is a coefficient on money. This is clear as it is a simple change of variables in the integrals.

no longer prefer to purchase it. Since the dataset includes ranking information, a simple regression of product ranking on booking price and the other covariates from the earlier demand regression is informative of this relationship.[31] I ran this regression on the bookings that were ordered by the random-ranking rule and then again, separately, on the bookings that were ordered by the proprietary-ranking rule. As expected, the price coefficient in the former case is a weak predictor of ranking. In the latter case, while the price coefficient is statistically significant large changes in price are predicted to have small changes in ranking. For example, in the latter case and for each of the five most popular bookings, I estimate that a three-standard-deviation increase in a good's price increases a product's ranking by less than three positions.[32] This is a small amount, given the first page of search results has space for more than 30 positions. At the same time, the previous section's demand regression predicts that a three-standard-deviation increase in price lowers the quantity demanded for each of the top five products by more than 50%. (See Table 2 for details.) Thus, as the price of a good increases, it is likely to become undesirable before it is removed from the first page of search results, and price independence approximately holds in this environment.[33]

Product	Average Price (Standard Deviation)	Estimated Ranking Change	Estimated Demand Change
A	\$61.63 (\$36.27)	.735 positions	-65.7%
B	\$53.90 (\$33.14)	.672	-62.6%
C	\$89.55 (\$46.92)	.951	-75.4%
D	\$167.37 (\$82.55)	1.67	-91.7%
E	\$107.31 (\$52.04)	1.05	-79.2%

Table 2: *Comparing the Effects of a Price Increase on Predicted Ranking and Predicted Market Share.* This table shows my estimates of the effects of a good's price increase on its own rankings. The values are estimated for a three-standard-deviation price increase from their observed levels. Products A through E above are the same A through E in Table 3 below. That is, these are the five most popular bookings in the Expedia data set. This table also includes my estimates of the effects of a good's three-standard-deviation price increase on its quantity demanded. I estimate that this sharp price increase will only mildly increase its ranking but will decrease its quantity demanded sharply. Thus, price independence is tenable. Only data from the search impressions that were ranked by the proprietary listing rule were included in this analysis. The data was from 2012 to 2013 and was provided by Expedia.

[31] That is, for each search impression, I regressed the products ranking on price, property star rating, property brand boolean, property location score 1 and the promotion flag.

[32] If a hotel's ranking increases, it is found farther down the page. A ranking of 1 is the closest to the top of the page and is the first among the search results to be seen by the consumer.

[33] Repeating the regression with a richer-specification of product-fixed effects and check-in-week-fixed effects further reduces the average predicted change in ranking to less than 2 positions for 3 standard deviations of price change. The code for analysis under this richer-specification is also included with the paper.

6.4 Welfare Changes from Removing the Top 5 Products

In this section, I provide estimates of the welfare loss that results from a new listing rule that hides the top five products from the consideration sets. This analysis allowed me to simulate the welfare harm a search platform could cause by suddenly removing certain third-party listings from its search-result lists.[34] The top five products are the products with the largest estimated market share in the sample data. These market share estimates are listed in Table 3. Together, these firms account for about 20% of the observed bookings. The market is not dominated by any one hotel: even the hotel with the largest market share accounted for less than 6% of total sales in the observed data.

Given my demand estimates from earlier in this section, I calculated welfare using Corollary 2. This calculation amounted to removing product A, then B, then C, et cetera. The results are shown in Table 3. The average welfare lost from the combined removal of all five products is $23.87 per person. The marginal removal of each additional product reduced consumer welfare by around $5. Bootstrap confidence intervals were calculated with 500 bootstrap resamples of the 4,694 search impressions. With each resample, I re-estimated demand and then re-estimated the welfare losses. The .025 and .975 quantiles of the bootstrapped estimates are the boundaries of the 95% confidence intervals.

Product	Estimated Market Share	Estimated Marginal Welfare Loss	Bootstrap 95% Confidence Interval
A	.0530	$5.35	[$4.75, $5.96]
B	.0454	$4.88	[$4.36,$5.39]
C	.0432	$4.92	[$4.59, $5.28]
D	.0382	$4.49	[$3.97, $5.09]
E	.0344	$4.23	[$3.86,$4.62]
Total	.214	$23.87	[$21.85, $26.06]

Table 3: *Counterfactual Marginal and Cumulative Welfare Losses from Removing the Top Five Products from the Search-Result Lists.* This table shows my estimates of the marginal welfare loss from removing product A, then B, then C, etc. until E is removed last. Products A through E are the products estimated to be most frequently purchased in the data set. I performed 500 bootstrap resamples of the data to calculate the 95% confidence intervals. The table also indicates each product's estimated market share in the data provided by Expedia.

[34] This was a concern that drove the EU Antitrust authorities to fine Alphabet $2.7 billion in 2017 and there is evidence that Amazon also does this.

7 Conclusion

I have presented several formulas for measuring the changes in consumer welfare that result from an online shopping platform changing the way it lists its search results. Under monotonicity, compensating variation and equivalent variation can be bound with straightforward integrals of aggregate demand. Under quasi-linearity, the exact compensating variation and equivalent variation can be recovered. I have also provided formulas for estimating the counterfactual welfare changes that occur under certain simple listing rule changes. Applications to data that features an OTA's search-result experiment show that ordered listings improve welfare over random listings by an average of \$8.11 per user. They also show that the removal of the five most popular products from search-result lists would lower welfare by an average of \$23.87 per user.

Appendix A Welfare When Consideration Sets Depend on Prices

In this section, I show equivalent variation need not equal the difference between initial and final utility when price independence fails, despite utility being linear in money. In particular, suppose consideration sets depend on the difference between income and good prices. That is, a consumer (y, η, ζ) has consideration set under listing rule α denoted $\mathcal{C}(\{ y - p_k \}_{k\in\mathcal{J}}, \eta, \zeta, \alpha)$. Equivalent variation S^{EV} can be adapted to this setting in the case of a welfare-improving change in listing rule as follows:[35]

$$S^{EV} = \sup\{S \in \mathbb{R} : \max_{j\in\mathcal{C}(\{ y-p_k-S \}_{k\in\mathcal{J}},\eta,\zeta,A)} u_j(y - p_j - S, \eta) \geq \max_{j\in\mathcal{C}(\{ y-p_k \}_{k\in\mathcal{J}},\eta,\zeta,B)} u_j(y - p_j, \eta)\} \tag{14}$$

For simplicity, this market has a single consumer (y, η, ζ). This market has two platforms, Affordable Store and Big Mart. This market has two goods for sale; $\mathcal{J} = \{ 0, 1, 2 \}$. Affordable Store sells good 1 on its platform and Big Mart sells good 2 on its platform. The consumer forms her consideration set using the following rule: she considers goods exclusively from Affordable Store if good 1 is available on Affordable Store and has price no lower than \$20. Otherwise, she considers all goods on each platform.[36] The consumer's utility for each good is as follows:

$$\begin{cases} u_0(y, \eta) = y \\ u_1(y - p_1, \eta) = y - p_1 + \gamma \\ u_2(y - p_2, \eta) = y - p_2 + \gamma + \delta \end{cases}$$

[35] In the case of a welfare-decreasing change in listing rule, the supremum would be replaced by an infimum.

[36] To motivate this, suppose a price lower than \$20 may make her believe both stores are having a sale and that she could get an even better deal deal at Big Mart.

where $\delta > 3$ and $\gamma > 20$. Suppose that under listing rule A, good 1 is available for \$22 on Affordable Mart and good 2 is available for the same price, \$22, at Big Mart. Thus, under listing rule A, the consumer's consideration set is $\{0, 1\}$, she purchases good 1 and her utility is $y - \$22 + \gamma$.

Next, suppose that under listing rule B, Affordable Store removes good 1 from its platform.[37] Thus, the consumer expands her consideration set to include good 2 and receives utility $y - 22 + \gamma + \delta$.

The consumer's utility differs by $-\delta$ from case A to case B. However, inspection of Equation (14) shows that when $S < -2$, the consumer searches all goods and the left-hand side of the inequality in Equation (14) is equal to $y - 22 + |S| + \gamma + \delta$ and the inequality holds. Thus, $S^{EV} \geq -2 > -3 > -\delta$ and therefore equivalent variation is closer to 0 than the change in utility. Intuitively, a small increase in the price of good 1 encourages her to expand her search and ultimately benefits her just as much as the change in listing rule from A to B. In this case, the difference in utility overestimates the price decrease (or income increase) that would compensate her for the change in listing rule.

It is straightforward to devise a similar example showing utility differences need not equal compensating variation, either. Thus, welfare measures in search settings require some scrutiny. Simple differences in utility should be interpreted based on the entire search model, and need not coincide with the classical welfare measures of compensating variation or equivalent variation, even when utility is linear in money, without additional modeling assumptions.

Appendix B Proofs

B.1 Proof of Theorem 1

To prove Equation (9) in Theorem 4, I first show that the individual inequalities in Equation (15) below hold for an arbitrary consumer. That is,

$$0 \geq S^{CV} \geq \sum_{j \in \mathcal{R}} \lim_{p_k \to \infty \forall k \in \mathcal{R} \setminus \{j\}} \int_{p_j^m}^{\infty} \Delta q_j(y, p_j, (p_{\mathcal{R} \setminus \{j\}}, p^m_{-(R,j)}), \eta, \zeta) dp_j, \tag{15}$$

where

$$\Delta q_j(y, p_j, (p_{\mathcal{R} \setminus \{j\}}, p^m_{-(R,j)}), \eta, \zeta) \equiv \\ q_j(y, p_j, (p_{\mathcal{R} \setminus \{j\}}, p^m_{-(R,j)}), A, \eta, \zeta) - q_j(y, p_j, (p_{\mathcal{R} \setminus \{j\}}, p^m_{-(R,j)}), B, \eta, \zeta). \tag{16}$$

I start by proving the left-hand inequality in Equation (15). To that end, fix an arbitrary consumer (y, η, ζ) and fix consumer-listing rules A and B. For ease of notation, I define the

[37] To motivate this, Affordable Store could be removing good 1 to try to push a 3rd good onto the consumer that gives Affordable Store better margins. However, the consumer values this third good at its offered price less than the outside good and still expands her consideration set to include Big Mart's goods.

following terms. Let

$$j^\star(\mathcal{C}) := \arg\max_{j\in\mathcal{C}} u_j(y - p_j^m, \eta),$$

for an aribitrary consideration set $\mathcal{C} \supseteq \{0\}$. For further ease, I abbreviate $j^\star(\mathcal{C}_A)$ as $j_A^\star$ and $j^\star(\mathcal{C}_B)$ as $j_B^\star$. Next, note that by monotonicity for any consideration set $\mathcal{C} \supseteq \{0\}$ there is some real number $\bar{p}_{j_B^\star}(\mathcal{C})$ defined such that

$$\max_{j\in\mathcal{C}} u_j(y - p_j^m, \eta) = u_{j_B^\star}(y - \bar{p}_{j_B^\star}(\mathcal{C}), \eta).$$

Note that if $\mathcal{C} \subseteq \mathcal{D}$ then $\bar{p}_{j_B^\star}(\mathcal{C}) \geq \bar{p}_{j_B^\star}(\mathcal{D})$ by monotonicity and since,

$$\max_{j\in\mathcal{C}} u_j(y - p_j^m, \eta) = u_{j_B^\star}(y - \bar{p}_{j_B^\star}(\mathcal{C}), \eta) \leq \max_{j\in\mathcal{D}} u_j(y - p_j^m, \eta) = u_{j_B^\star}(y - \bar{p}_{j_B^\star}(\mathcal{D}), \eta).$$

Moreover, note that when $\mathcal{C} = \mathcal{C}_A$,

$$u_{j_A^\star}(y - p_{j_A^\star}^m, \eta) = u_{j_B^\star}(y - p_{j_B^\star}^m + S^{CV}, \eta) = u_{j_B^\star}(y - \bar{p}_{j_B^\star}(\mathcal{C}_A), \eta).$$

Therefore, we can conclude that $\bar{p}_{j_B^\star}(\mathcal{C}_A) = p_{j_B^\star}^m - S^{CV}$. Thus,

$$\begin{cases} S^{CV} = 0 \text{ if } j_B^\star = j_A^\star \\ S^{CV} \leq 0 \text{ if } j_B^\star \neq j_A^\star \end{cases}, \tag{17}$$

which proves the left-hand inequality of Equation (15). The second part of Equation (17) follows from the assumption in *Theorem* 1 that $\mathcal{C}_A \subseteq \mathcal{C}_B$.

To see the right-hand inequality in Equation (15), I prove two claims for our arbitrary consumer: (1) that

$$\lim_{p_k\to\infty \forall k\in\mathcal{R}\setminus\{j\}} \int_{p_j^m}^{\infty} \Delta q_j(y, p_j, (p_{\mathcal{R}\setminus\{j\}}, p_{-(R,j)}^m), \eta, \zeta) dp_j \leq 0,$$

for every $j \in \mathcal{R}$; and (2) there is some $j \in \mathcal{R}$ such that

$$S^{CV} \geq \lim_{p_k\to\infty \forall k\in\mathcal{R}\setminus\{j\}} \int_{p_j^m}^{\infty} \Delta q_j(y, p_j, (p_{\mathcal{R}\setminus\{j\}}, p_{-(R,j)}^m), \eta, \zeta) dp_j.$$

It is clear that these two claims together are sufficient for the right-hand inequality to hold.

For the first claim, fix $j \in \mathcal{R}$ arbitrarily. It sufficies to show that

$$\lim_{p_k\to\infty \forall k\in\mathcal{R}\setminus\{j\}} q_j(y, p_j, (p_{\mathcal{R}\setminus\{j\}}, p_{-(R,j)}^m), A, \eta, \zeta) \leq \lim_{p_k\to\infty \forall k\in\mathcal{R}\setminus\{j\}} q_j(y, p_j, (p_{\mathcal{R}\setminus\{j\}}, p_{-(R,j)}^m), B, \eta, \zeta),$$

for all $p_j \in [p_j^m, \infty)$. Note that for a fixed $p_j \in [p_j^m, \infty)$,

$$\lim_{p_k\to\infty \forall k\in\mathcal{R}\setminus\{j\}} q_j(y, p_j, (p_{\mathcal{R}\setminus\{j\}}, p_{-(R,j)}^m), A, \eta, \zeta) = 1 \text{ only if } j = j^\star(\mathcal{C}_A \setminus (\mathcal{R} \setminus j)).$$

But if $j = j^\star(\mathcal{C}_A \setminus (\mathcal{R} \setminus j))$ then it must be the case that $j = j^\star(\mathcal{C}_B \setminus (\mathcal{R} \setminus j))$ and moreover $\mathcal{C}_A \setminus (\mathcal{R} \setminus j) = \mathcal{C}_B \setminus (\mathcal{R} \setminus j)$ since $\mathcal{C}_A$ and $\mathcal{C}_B$ can differ by at most $\mathcal{R}$. Thus, it must be that for a fixed $p_j \in [p_j^m, \infty)$,

$$\lim_{p_k \to \infty \forall k \in \mathcal{R} \setminus \{j\}} q_j(y, p_j, (p_{\mathcal{R} \setminus \{j\}}, p^m_{-(R,j)}), A, \eta, \zeta) = 1$$
$$\text{only if} \lim_{p_k \to \infty \forall k \in \mathcal{R} \setminus \{j\}} q_j(y, p_j, (p_{\mathcal{R} \setminus \{j\}}, p^m_{-(R,j)}), B, \eta, \zeta) = 1,$$

and 0 otherwise. Since $\lim_{p_k \to \infty \forall k \in \mathcal{R} \setminus \{j\}} q_j(y, p_j, (p_{\mathcal{R} \setminus \{j\}}, p^m_{-(R,j)}), B, \eta, \zeta) \geq 0$, this completes the proof of the first claim.

For the second claim, note that if $S^{CV} = 0$, then the second claim holds true for all $j \in \mathcal{R}$ trivially as a result of the first claim. Thus, it just remains to show that the second claim holds in the case that $S^{CV} < 0$. In this case, it must be that $j^\star_B \neq j^\star_A$ and that $j^\star_B \in \mathcal{C}_B \setminus \mathcal{C}_A \subseteq \mathcal{R}$. Thus, fix $j = j^\star_B$ and note,

$$\begin{aligned}
&\lim_{p_k \to \infty \forall k \in \mathcal{R} \setminus \{j\}} \int_{p_j^m}^{\infty} \Delta q_j(y, p_j, (p_{\mathcal{R} \setminus \{j\}}, p^m_{-(R,j)}), \eta, \zeta) dp_j \\
&= - \lim_{p_k \to \infty \forall k \in \mathcal{R} \setminus \{j\}} \int_{p_j^m}^{\infty} q_j(y, p_j, (p_{\mathcal{R} \setminus \{j\}}, p^m_{-(R,j)}), B, \eta, \zeta) dp_j && (j \notin \mathcal{C}_A) \\
&= - \int_{p_j^m}^{\infty} \lim_{p_k \to \infty \forall k \in \mathcal{R} \setminus \{j\}} q_j(y, p_j, (p_{\mathcal{R} \setminus \{j\}}, p^m_{-(R,j)}), B, \eta, \zeta) dp_j \\
& && \text{(Monotone Convergence Theorem)} \\
&= p^m_{j^\star_B} - \bar{p}_{j^\star_B}(\mathcal{C}_B \setminus \mathcal{R}) \\
&= p^m_{j^\star_B} - \bar{p}_{j^\star_B}(\mathcal{C}_A \setminus \mathcal{R}) && (\mathcal{C}_B \setminus \mathcal{R} = \mathcal{C}_A \setminus \mathcal{R}) \\
&\leq p^m_{j^\star_B} - \bar{p}_{j^\star_B}(\mathcal{C}_A) && (\mathcal{C}_A \setminus \mathcal{R} \subseteq \mathcal{C}_A) \\
&= S^{CV}.
\end{aligned}$$

This completes the proof of the second claim and the proof that Equation (15) holds for an arbitrary consumer.

Finally, to get Equation (9) from Equation (15), all we need to do is integrate each part over η and ζ, pass the limit through the integral by the monotone convergence theorem (which is possible by monotonicity) and finally switch the integral order by Tonnelli's theorem (which is possible since we showed the integrand is non-positive).

B.2 Welfare Theorems Under Monotonicity in the Case of the Removal of Listings from Search Results

The following theorem bounds the equivalent variation in the case of the removal of a collection of goods from the search-result list.

Theorem 4 (Probable Removal of Goods)**.** *Suppose the listing rule changes exogenously from A to B in a way that makes each consumer unable to consider goods in the collection*

$\mathcal{R}$.[38] *Then, the average equivalent variation of this change in listing rule is such that*

$$0 \leq \mu^{EV} \leq \sum_{j \in \mathcal{R}} \lim_{p_k \to \infty \forall k \in R \setminus \{j\}} \int_{p_j^m}^{\infty} \Delta Q_j(y, p_j, (p_{\mathcal{R}\setminus\{j\}}, p_{-\mathcal{R}}^m)) dp_j. \tag{18}$$

The proof of this result is analogous to the proof of Theorem 1 above.

B.3 Proof of Corollary 1 and Corollary 3

In the case of the addition of a single product that was previously unconsidered or the complete removal of a single product from the search-result lists, the non-zero bounds in Theorem 1 and Theorem 4, respectively collapse into an equality. The results of Corollary 1 and Corollary 3 build on this result.

First, I state Corollary 3 and then I prove both.

Corollary 3 (Probable Removal of a Single Good)**.** *Suppose the listing rule changes exogenously from A to B in a way that makes each consumer less likely to consider product* 1.[39] *Then, the average equivalent variation of this change in listing rule is*

$$\mu^{EV} = \int_{p_1^m}^{\infty} \Delta Q_1^m(y, p_1, p_{-1}^m) dp_1. \tag{19}$$

B.3.1 Proof of Corollary 3

Fix a consumer (y, η, ζ). There are two cases. In case 1, the consumer makes the same good choice under A as under B. In case 2, her choice changes. In case 1, S^{EV} must be zero, and we see

$$\int_{p_1^m}^{\infty} q_1^m(y, p, \eta, \zeta, A) dp - \int_{p_1^m}^{\infty} q_1^m(y, p, \eta, \zeta, B) dp = 0$$

since if 1 is her choice then it must be that $1 \in \mathcal{C}_A, \mathcal{C}_B$ and if 1 is not her choice under A, then it mustn't be under B as well and both integrals are 0. This finishes the proof for case 1.

In case 2, our fixed consumer must have purchased good 1 at time $t = 0$. We know from Theorem 4 that

$$S^{EV} = \int_{p_1^m}^{\infty} q_1^m(y, p, \eta, \zeta, A) dp.$$

Moreover, for this consumer, since $1 \notin \mathcal{C}_B$ (otherwise she would have purchased it),

$$0 = \int_{p_1^m}^{\infty} q_1^m(y, p, \eta, \zeta, B) dp.$$

[38]Specifically, $\mathcal{C}_B \subseteq \mathcal{C}_A \subseteq \mathcal{C}_B \cup \mathcal{R}$.

[39]Specifically, I mean for every (y, η, ζ), $\mathcal{C}_B$ is either equal to $\mathcal{C}_A$ or $\mathcal{C}_A \setminus \{1\}$.

Thus, for this case as well,

$$S^{EV} = \int_{p_1^m}^{\infty} q_1^m(y, p, \eta, \zeta, A)dp - \int_{p_1^m}^{\infty} q_1^m(y, p, \eta, \zeta, B)dp.$$

Aggregating over consumers and cases and then switching the integration order by Tonelli's theorem yields the desired result.

B.3.2 Proof of Corollary 1

This proof is very similar to the proof for Corollary 3 above and is therefore omitted.

B.4 Proof of Theorem 2

Fix an arbitrary consumer (y, η, ζ) and let ω_α be defined by

$$\omega_\alpha := \lim_{p_2^m,\ldots,p_J^m \to \infty} \int_{p_1^m}^{\infty} q_1(y, p, p_{-1}^m, \eta, \zeta, \alpha)dp + \lim_{p_3^m,\ldots,p_J^m \to \infty} \int_{p_2^m}^{\infty} q_2(y, p, p_{-2}^m, \eta, \zeta, \alpha)dp + \cdots + \int_{p_J^m}^{\infty} q_J(y, p, p_{-J}^m, \eta, \zeta, \alpha)dp. \tag{20}$$

By Quasi-linearity and Tonelli's Theorem, it suffices to prove that

$$\omega_\alpha = \max_{j \in \mathcal{C}_\alpha} \left[y - p_j^m + \tilde{U}_j(\eta) \right] - u_0(y, \eta) \tag{21}$$

because then

$$\omega_A - \omega_B = \max_{j \in \mathcal{C}_A} \left[y - p_j^m + \tilde{U}_j(\eta) \right] - \max_{j \in \mathcal{C}_B} \left[y - p_j^m + \tilde{U}_j(\eta) \right] = S^W$$

and

$$\Omega_A - \Omega_B = \int \omega_A dF - \int \omega_B dF = \mu^W.$$

Thus, to show Equation (21), first note that

$$\begin{aligned}
&\lim_{p_2^m,\ldots,p_J^m \to \infty} \int_{p_1^m}^{\infty} q_1(y, p, p_{-1}^m, \eta, \zeta, \alpha)dp \\
&= \lim_{p_2^m,\ldots,p_J^m \to \infty} \int_{p_1^m}^{\infty} \mathbb{1}(1 \in \mathcal{C}_\alpha) \cdot \mathbb{1}\left[-p_1 + \tilde{U}_1(\eta) > \max_{j \neq 1, j \in \mathcal{C}_\alpha} -p_j + \tilde{U}_j(\eta) \right] dp_1 \\
&= \mathbb{1}(1 \in \mathcal{C}_\alpha) \cdot \int_{p_1^m}^{\infty} \mathbb{1}\left[-p_1 + \tilde{U}_1(\eta) > \tilde{U}_0(\eta) \right] dp_1 \qquad \text{(Monotone Convergence Thm)} \\
&= \begin{cases} -p_1^m + \tilde{U}_1(\eta) - \tilde{U}_0(\eta) & \text{if } \tilde{U}_1(\eta) - p_1^m > \tilde{U}_0(\eta) \text{ and } 1 \in \mathcal{C}_\alpha \\ 0 & \text{otherwise.} \end{cases}
\end{aligned}$$

Similarly,

$$\lim_{p_3^m,\ldots,p_J^m\to\infty}\int_{p_2^m}^{\infty} q_2(y,p,p_{-2}^m,\eta,\zeta,\alpha)dp$$
$$= \begin{cases} -p_2^m+\tilde{U}_2(\eta)-(\max_{j\in\mathcal{C}_\alpha\cap\{\,0,1\,\}}-p_j^m+\tilde{U}_j(\eta)) \\ \qquad \text{if } -p_2^m+\tilde{U}_2(\eta)>\max_{j\in\mathcal{C}_\alpha\cap\{\,0,1\,\}}-p_j^m+\tilde{U}_j(\eta) \text{ and } 2\in\mathcal{C}_\alpha \\ 0 \text{ otherwise.}\end{cases}$$

Continuing this pattern and putting this all together, we see

$$\begin{aligned}\omega_\alpha &= \lim_{p_2^m,\ldots,p_J^m\to\infty}\int_{p_1^m}^{\infty} q_1(y,p,p_{-1}^m,\eta,\zeta,\alpha)dp+\lim_{p_3^m,\ldots,p_J^m\to\infty}\int_{p_2^m}^{\infty} q_2(y,p,p_{-2}^m,\eta,\zeta,\alpha)dp \\ &\qquad +\cdots+\int_{p_J^m}^{\infty} q_J(y,p,p_{-J}^m,\eta,\zeta,\alpha)dp \\ &= \mathbb{1}\left(1=\arg\max_{\mathcal{C}_\alpha\cap\{\,0,1\,\}}\tilde{U}_j(\eta)-p_j^m\right)\times \\ &\qquad \left[\max_{j\in\mathcal{C}_\alpha\cap\{\,0,1\,\}}\tilde{U}_j(\eta)-p_j^m-\left(\max_{j\in\mathcal{C}_\alpha\cap\{\,0\,\}}\tilde{U}_j(\eta)-p_j^m\right)\right] \\ &\quad +\mathbb{1}\left(2=\arg\max_{\mathcal{C}_\alpha\cap\{\,0,1,2\,\}}\tilde{U}_j(\eta)-p_j^m\right)\times \\ &\qquad \left[\max_{j\in\mathcal{C}_\alpha\cap\{\,0,1,2\,\}}\tilde{U}_j(\eta)-p_j^m-\left(\max_{j\in\mathcal{C}_\alpha\cap\{\,0,1\,\}}\tilde{U}_j(\eta)-p_j^m\right)\right] \\ &\quad +\cdots+\mathbb{1}\left(J=\arg\max_{\mathcal{C}_\alpha}\tilde{U}_j(\eta)-p_j^m\right)\times \\ &\qquad \left[\max_{j\in\mathcal{C}_\alpha}\tilde{U}_j(\eta)-p_j^m-\left(\max_{j\in\mathcal{C}_\alpha\setminus\{\,J\,\}}\tilde{U}_j(\eta)-p_j^m\right)\right].\end{aligned} \tag{22}$$

I now conclude the proof of Equation (21) using Equation (22) with induction on $J\in\mathbb{N}$.

Note that for $J=1$, $\mathcal{C}_\alpha\subseteq\{\,0,1\,\}$ and

$$\begin{aligned}\omega_\alpha &= \mathbb{1}(1=\arg\max_{\mathcal{C}_\alpha})\left[\tilde{U}_1(\eta)-p_1^m-\tilde{U}_0(\eta)\right] \\ &= \max_{j\in\mathcal{C}_\alpha}\left[U_j(\eta)-p_j^m\right]-\tilde{U}_0(\eta),\end{aligned}$$

which proves the base case. Now suppose this holds for the collection of goods $\{\,0,1,\ldots,K\,\}$.

Then, for $J = K + 1$,

$$\begin{aligned}
\omega_\alpha &= \left[\max_{j\in\mathcal{C}_\alpha\setminus\{J\}} \tilde{U}_j(\eta) - p_j^m\right] - \tilde{U}_0(\eta) && \text{(by inductive hypothesis)}\\
&\quad + \mathbb{1}\left(J = \arg\max_{\mathcal{C}_\alpha} \tilde{U}_j(\eta) - p_j^m\right)\left[\max_{j\in\mathcal{C}_\alpha} \tilde{U}_j(\eta) - p_j^m - \left(\max_{j\in\mathcal{C}_\alpha\setminus\{J\}} \tilde{U}_j(\eta) - p_j^m\right)\right]\\
&= \max_{j\in\mathcal{C}_\alpha}\left[\tilde{U}_j(\eta) - p_j^m\right] - \tilde{U}_0(\eta),
\end{aligned}$$

which concludes the proof.

B.5 Proof of Theorem 3

Fix a consumer (y, η, ζ). I start by showing

$$S^W = \int_{p_1^m}^{\infty} q_1^m(y, p, \eta, \zeta, A)dp - \int_{p_M^m}^{\infty} q_M^m(y, p, \eta, \zeta, B)dp$$

In the cases where only good 1 exits this consumer's consideration set, or only good M enters this consumer's consideration set, or her purchase behavior does not change, the result is clear from Corollary 1 and Corollary 3. This leaves only the case where the consumer purchases good 1 under A and purchases good M under B. In this case, $\mathcal{C}_A\setminus\{1\} = \mathcal{C}_B\setminus\{M\}$

Thus,

$$\begin{aligned}
&\int_{p_1^m}^{\infty} q_1(y, p, p_{-1}^m, \eta, \zeta, A)dp - \int_{p_M^0}^{\infty} q_M^m(y, p, p_{-M}^m, \eta, \zeta, B)dp\\
&= u_1(y - p_1^m, \eta) - \max_{j\in\mathcal{C}_A\setminus\{1\}} u_j(y - p_1^m, \eta)\\
&\quad - \left[u_M(y - p_M^m, \eta) - \max_{j\in\mathcal{C}_B\setminus\{M\}} u_j(y - p_j, \eta)\right]\\
&&\text{(Corollary 1, Corollary 3 and Quasi-linearity)}\\
&= u_1(y - p_1^m, \eta) - u_M(y - p_M^m, \eta)\\
&= S^W.
\end{aligned}$$

Extending the results from S^W to μ^W proceeds exactly as in the rest of the proofs.

B.6 Proof of Corollary 2 and Analogous Result for Multiple Product Removal from Search-Result Lists

B.6.1 Proof of Corollary 2

This is a corollary of Theorem 2. As discussed in Section B.4, the ordering of the goods does not matter. So, let the goods $1, \ldots, J = 1, 2 \ldots, r_1, \ldots, r_R$. That is, good $J = r_R$, good

$J-1=r_{R-1}, \ldots, J-R+1=r_1$. Then

$$\Omega_\alpha = \lim_{p_2,\ldots,p_J\to\infty} \int_{p_1^m}^{\infty} Q_1(y,p,p_{-1},\alpha)dp + \cdots + \lim_{p_{r_1},\ldots,p_{r_R}\to\infty} \int_{p_{r_1-1}^m}^{\infty} Q_{r_1-1}(y,p,(p_1^m,\ldots,p_{r_1-2}^m,p_{r_1},\ldots,p_{r_R}),\alpha)dp + \Gamma_\alpha(\mathcal{R}).$$

Since

$$\Lambda_\alpha := \lim_{p_2,\ldots,p_J\to\infty} \int_{p_1^m}^{\infty} Q_1(y,p,p_{-1},\alpha)dp + \cdots + \lim_{p_{r_1},\ldots,p_{r_R}\to\infty} \int_{p_{r_1-1}^m}^{\infty} Q_{r_1-1}(y,p,(p_1^m,\ldots,p_{r_1-2}^m,p_{r_1},\ldots,p_{r_R}),\alpha)dp$$

is invariant to the inclusion of products $\{r_1,\ldots,r_R\}$ in consideration sets or not—the price limits make Λ_α independent of them—we see $\Lambda_A=\Lambda_B$ and, thus

$$\Omega_A-\Omega_B=\Gamma_A(\mathcal{R})-\Gamma_B(\mathcal{R}),$$

which completes the proof.

B.6.2 Multiple Product Removal

Suppose now that the collection of products $\mathcal{R}$ is removed from search-result listings. Under Quasi-linearity, this simply requires reversing the sign on the results in Corollary 2. The only difference is that the roles of A and B are switched.[40] For completeness, I state the results below.

Corollary 4. *Let $\mathcal{R}=\{r_1,\ldots,r_R\}$. Suppose that a change in the listing rule from A to B increases the probability that each of the products in $\mathcal{R}$ is considered. Then, under Quasi-linearity*

$$\mu^W=\Gamma_A(\mathcal{R})-\Gamma_B(\mathcal{R}).$$

[40]Indeed, renaming A and B above for their reversed roles gives the proof for Corollary 4.

Appendix C Additional Figures

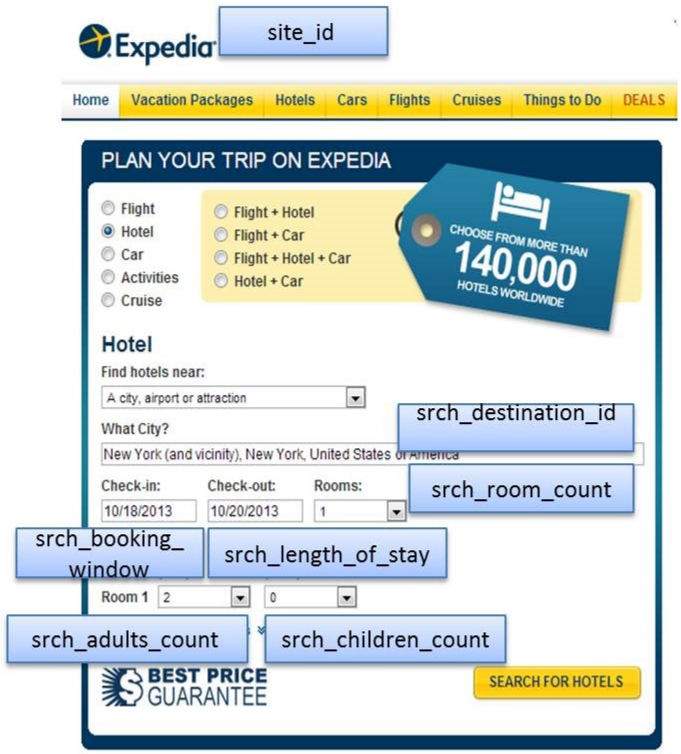

Figure 5: *2012 and 2013 Home Page of Expedia.com.* This is the first page encountered by users of *Expedia.com*. Users select their travel destination, the number of rooms they wish to book, the number of days they wish to spend at the destination, and the number of adults and children who will be staying in the room. All of this information, highlighted by the blue boxes, was collected by Expedia and included in the data set they provided. This figure was provided by Expedia.

References

Abaluck, Jason and Abi Adams (2018). "What Do Consumers Consider Before They Choose? Identification from Asymmetric Demand Responses." In: *Working Paper*. May 5, 2018 version. URL: https://som.yale.edu/sites/default/files/DiscreteChoiceInattention_master.pdf.

Figure 6: *A Search Listing and the Final Booking Page on Expedia.com in 2012 and 2013.*

(a) *A Typical Search Listing on Expedia.com in 2012 and 2013.* This picture shows a typical listing *Expedia.com* users would have observed during the period in which the data was collected. Each search impression contains between 1 and 34 of these listings. This figure was provided by Expedia.

(b) *A Typical Booking Page on Expedia.com in 2012 and 2013.* This is an example of the page a consumer would encounter when finalizing her booking on *Expedia.com* between 2012 and 2013. The consumer would have encountered this page after she clicked on a listing, such as the one shown above in Figure 6a. This figure was provided by Expedia.

Athey, Susan and Glenn Ellison (2011). "Position Auctions with Consumer Search." In: *The Quarterly Journal of Economics* 126.3, pp. 1213–1270. DOI: 10.1093/qje/qjr028. URL: http://dx.doi.org/10.1093/qje/qjr028.

Barseghyan, Levon, Maura Coughlin, et al. (2019). *Heterogeneous Choice Sets and Preferences*. arXiv: 1907.02337 [econ.EM].

Barseghyan, Levon, Francesca Molinari, and Matthew Thirkettle (2019). *Discrete Choice under Risk with Limited Consideration*. arXiv: 1902.06629 [econ.EM].

Bhattacharya, Debopam (2015). "Nonparametric Welfare Analysis for Discrete Choice." In: *Econometrica* 83.2, pp. 617–649. DOI: 10.3982/ECTA12574.

— (2018). "Empirical welfare analysis for discrete choice: Some general results." In: *Quantitative Economics* 9.2, pp. 571–615. DOI: https://doi.org/10.3982/QE931. URL: https://onlinelibrary.wiley.com/doi/abs/10.3982/QE931.

Brynjolfsson, Erik, Avinash Collis, et al. (Mar. 2019). *GDP-B: Accounting for the Value of New and Free Goods in the Digital Economy*. Working Paper 25695. National Bureau of Economic Research. DOI: 10.3386/w25695. URL: http://www.nber.org/papers/w25695.

Brynjolfsson, Erik and JooHee Oh (2012). "The Attention Economy: Measuring the Value of Free Digital Services on the Internet." In: *ICIS*.

Croissant, Yves (2019). *Estimation of multinomial logit models in R: The mlogit Packages*. Version 0.4-2.

Diewert, W. Erwin and Robert Feenstra (Dec. 2017). *Estimating the Benefits and Costs of New and Disappearing Products*. Microeconomics.ca working papers. Vancouver School of Economics. URL: https://ideas.repec.org/p/ubc/pmicro/tina_marandola-2017-12.html.

Diewert, W. Erwin, Kevin Fox, and Paul Schreyer (2018). *The Digital Economy, New Products and Consumer Welfare*. Economic Statistics Centre of Excellence (ESCoE) Discussion Papers ESCoE DP-2018-16. Economic Statistics Centre of Excellence (ESCoE). URL: https://EconPapers.repec.org/RePEc:nsr:escoed:escoe-dp-2018-16.

Dinerstein, Michael et al. (July 2018). "Consumer Price Search and Platform Design in Internet Commerce." In: *American Economic Review* 108.7, pp. 1820–59. DOI: 10.1257/aer.20171218. URL: http://www.aeaweb.org/articles?id=10.1257/aer.20171218.

Ershov, Daniel (2016). "The Effect of Consumer Search Costs on Entry and Quality in the Mobile App Market." In: *University of Toronto Job Market Paper*. URL: http://www.law.northwestern.edu/research-faculty/searlecenter/events/internet/documents/Ershov_Consumer_Search.pdf.

Feldstein, Martin (May 2017). "Underestimating the Real Growth of GDP, Personal Income, and Productivity." In: *Journal of Economic Perspectives* 31.2, pp. 145–64. DOI: 10.1257/jep.31.2.145. URL: http://www.aeaweb.org/articles?id=10.1257/jep.31.2.145.

Fradkin, Andrey (2018). "Search, Matching, and the Role of Digital Marketplace Design in Enabling Trade: Evidence from Airbnb." In: *Working Paper*. URL: https://andreyfradkin.com/assets/SearchMatchingEfficiency.pdf.

Gentzkow, Matthew (2007). "Valuing New Goods in a Model with Complementarity: Online Newspapers." In: *American Economic Review* 97.3, pp. 713–744. URL: https://EconPapers.repec.org/RePEc:aea:aecrev:v:97:y:2007:i:3:p:713-744.

Goeree, Michelle Sovinsky (2008). "Limited Information and Advertising in the U.S. Personal Computer Industry." In: *Econometrica* 76.5, pp. 1017–1074. DOI: 10.3982/ECTA4158. URL: https://onlinelibrary.wiley.com/doi/abs/10.3982/ECTA4158.

Greenstein, Shane and Ryan McDevitt (2011). "The broadband bonus: Estimating broadband Internet's economic value." In: *Telecommunications Policy* 35.7, pp. 617–632. URL: https://EconPapers.repec.org/RePEc:eee:telpol:v:35:y:2011:i:7:p:617-632.

Groshen, Erica L. et al. (May 2017). "How Government Statistics Adjust for Potential Biases from Quality Change and New Goods in an Age of Digital Technologies: A View from the Trenches." In: *Journal of Economic Perspectives* 31.2, pp. 187–210. DOI: 10.1257/jep.31.2.187. URL: http://www.aeaweb.org/articles?id=10.1257/jep.31.2.187.

Hausman, Jerry (1981). "Exact Consumer's Surplus and Deadweight Loss." In: *American Economic Review* 71.4, pp. 662–76. URL: https://EconPapers.repec.org/RePEc:aea:aecrev:v:71:y:1981:i:4:p:662-76.

— (1996). "Valuation of New Goods under Perfect and Imperfect Competition." In: *The Economics of New Goods*. National Bureau of Economic Research, Inc, pp. 207–248. URL: https://EconPapers.repec.org/RePEc:nbr:nberch:6068.

Honka, Elisabeth (2014). "Quantifying Search and Switching Costs in the US auto insurance industry." In: *RAND Journal of Economics* 45.4, pp. 847–884.

Honka, Elisabeth, Ali Hortaçsu, and Maria Ana Vitorino (2017). "Advertising, consumer Awareness and Choice: Evidence from the U.S. Banking Industry." In: *RAND Journal of Economics* 48.3, pp. 611–646.

Iaria, Alessandro, Gregory Crawford, and Rachel Griffith (Jan. 2020). "A Survey of Preference Estimation with Unobserved Choice Set Heterogeneity." In: *Journal of Econometrics*. ISSN: 0304-4076.

Khan, Lina M. (Jan. 2017). "Amazon's Antitrust Paradox." In: *Yale Law Yournal* 126.3, pp. 564–907. URL: http://www.yalelawjournal.org/article/amazons-antitrust-paradox.

Koulayev, Sergei (2014). "Search for Differentiated Products: Identification and Estimation." In: *RAND* 45.3, pp. 553–575.

Lewis, Gregory and Albert Wang (2013). "Who benefits from improved search in platform markets." In: *SSRN*. URL: http://ssrn.com/abstract=2249816.

Liu, Xudong et al. (2013). "Combination of Diverse Ranking Models for Personalized Expedia Hotel Searches." In: *CoRR* abs/1311.7679. URL: http://dblp.uni-trier.de/db/journals/corr/corr1311.html#LiuXZYPLSW13.

Manski, Charles (1975). "Maximum score estimation of the stochastic utility model of choice." In: *Journal of Econometrics* 3.3, pp. 205–228. URL: https://EconPapers.repec.org/RePEc:eee:econom:v:3:y:1975:i:3:p:205-228.

Matzkin, Rosa L. (2007). "Nonparametric Identification." In: *Handbook of Econometrics*. Ed. by J Heckman and E. Leamer. Vol. 6B. Amsterdam: Elsevier.

— (2013). "Nonparametric Identification in Structural Economic Models." In: *Annual Review of Economics* 5.1, pp. 457–486. URL: http://www.annualreviews.org/doi/abs/10.1146/annurev-economics-082912-110231.

Moraga-Gonzalez, José Luis, Zsolt Sándor, and Matthijs R. Wildenbeest (2017). "Nonsequential Search Equilibrium with Search Cost Heterogeneity." In: *International Journal of Industrial Organization* 50, pp. 392–414.

Nevo, Aviv (May 2003). "New products, quality changes, and welfare measures computed from estimated demand systems." English (US). In: *Review of Economics and Statistics* 85.2, pp. 266–275. ISSN: 0034-6535. DOI: 10.1162/003465303765299792.

Petrin, Amil (2002). “Quantifying the Benefits of New Products: The Case of the Minivan.” In: *Journal of Political Economy* 110.4, pp. 105–130.

Quan, Thomas W. and Kevin R. Williams (2018). “Product variety, across-market demand heterogeneity, and the value of online retail.” In: *The RAND Journal of Economics* 49.4, pp. 877–913. DOI: 10.1111/1756-2171.12255. URL: https://onlinelibrary.wiley.com/doi/abs/10.1111/1756-2171.12255.

R Core Team (2017). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing. Vienna, Austria. URL: http://www.R-project.org/.

Seiler, Stephan and Song Yao (2017). “The Impact of Advertising along hte Conversion Funnel.” In: *Quantitative Marketing and Economics* 15.3, pp. 241–278.

Small, Kenneth A. and Harvey S. Rosen (1981). “Applied Welfare Economics with Discrete Choice Models.” In: *Econometrica* 49.1, pp. 105–130. URL: http://www.jstor.org/stable/1911129.

Syverson, Chad (May 2017). “Challenges to Mismeasurement Explanations for the US Productivity Slowdown.” In: *Journal of Economic Perspectives* 31.2, pp. 165–86. DOI: 10.1257/jep.31.2.165. URL: http://www.aeaweb.org/articles?id=10.1257/jep.31.2.165.

Ursu, Raluca M. (2018). “The Power of Rankings: Quantifying the Effect of Rankings on Online Consumer Search and Purchase Decisions.” In: *Marketing Science* 37.4, pp. 530–552. DOI: 10.1287/mksc.2017.1072. URL: https://doi.org/10.1287/mksc.2017.1072.