

NBER WORKING PAPER SERIES

A METHOD TO ESTIMATE DISCRETE CHOICE MODELS THAT IS ROBUST
TO CONSUMER SEARCH

Jason Abaluck
Giovanni Compiani

Working Paper 26849
<http://www.nber.org/papers/w26849>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
March 2020

Thanks to Abi Adams, Judy Chevalier, Magne Mogstad, Barry Nalebuff, Aniko Oery, Nicholas Ryan, Fiona Scott Morton, Raluca Ursu, Miguel Villas-Boas, and seminar participants at Yale, SICS 2019, Marketing Science 2019, Caltech, QME 2019, University of Bologna, Northwestern, and UT Austin for helpful comments and suggestions. Tianyu Han and Jaewon Lee provided excellent research assistance. Jason blames any remaining errors on the widening partisan divide in this country (the US) and the lackluster proofreading efforts of the economics Twitter community. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2020 by Jason Abaluck and Giovanni Compiani. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

A Method to Estimate Discrete Choice Models that is Robust to Consumer Search
Jason Abaluck and Giovanni Compiani
NBER Working Paper No. 26849
March 2020
JEL No. C5,C8,C9,D0,D6,D8

ABSTRACT

We state a sufficient condition under which choice data alone suffices to identify consumer preferences when choices are not fully informed. Suppose that: (i) the data generating process is a search model in which the attribute hidden to consumers is observed by the econometrician; (ii) if a consumer searches good j , she also searches goods which are better than j in terms of the non-hidden component of utility; and (iii) consumers choose the good that maximizes overall utility among searched goods. Canonical models will be biased: the value of the hidden attribute will be understated because consumers will be unresponsive to variation in the attribute for goods that they do not search. Under the conditions above and additional mild restrictions, an alternative method of recovering preferences using cross derivatives of choice probabilities succeeds regardless of the search protocol and is thus robust to whether consumers are informed. The approach nests several standard models, including full information. Our methods suggest natural tests for full information and can be used to forecast how consumers will respond to additional information. We verify in a lab experiment that our approach succeeds in recovering preferences when consumers engage in costly search.

Jason Abaluck
Yale School of Management
Box 208200
New Haven, CT 06520-8200
and NBER
jason.abaluck@yale.edu

Giovanni Compiani
University of California at Berkeley
2220 Piedmont Ave
Berkeley, CA 94720
gcompiani@berkeley.edu

1 Introduction

When consumers purchase cars, houses, food, insurance, schooling and much else, they are often imperfectly informed about the attributes of relevant products in ways that substantially alter their choices (Allcott and Knittel 2019; Woodward and Hall 2012; Abaluck and Gruber 2011; Allcott, Lockwood, and Taubinsky 2019; Hastings and Weinstein 2008). Given this, models which assume full information may generate wrong conclusions about welfare and cannot be used to assess how choices would respond to more information. However, despite the emergence of behavioral economics as a major subfield of economic analysis, most work in applied economics continues to assume that choices are fully informed. We count 350 articles published in the AER, QJE, JPE, ECTA or ReStud since 2015 that estimate discrete choice models. Of these 350, 315 (90%) assume that consumers are fully informed.¹

We believe this occurs for three reasons. First, for some positive purposes, it is irrelevant whether choices are informed since all that is required is to estimate how demand responds to price (Berry and Haile 2014). For instance, price elasticities are sufficient to predict equilibrium prices after a counterfactual merger between two firms.² Second, the data necessary to directly measure consumers' beliefs is often unavailable, and even when it is available, survey data is viewed with suspicion (Gul and Pesendorfer 2008). Full information is viewed as a parsimonious assumption in the absence of evidence to the contrary. Third, choice data alone does not suffice to separately identify preferences and beliefs without further assumptions (Manski 2002). Structural search models in which consumer beliefs can be identified (e.g. Ursu (2018)) require assumptions regarding whether consumers take into account option value, whether they solve an optimal stopping problem or “satisfice”, distributional assumptions about prior beliefs and search costs, and whether choices are simultaneous or sequential, among others. The empirical literature suggests that canonical assumptions in all of these cases are often rejected by the data (respectively, Gabaix et al. (2006), Schwartz et al. (2002), Jindal and Aribarg (2018), Honka and Chintagunta (2016)).

In this paper, we state what we believe is a more plausible condition under which choice data alone suffices to recover preferences whether consumers are fully or only partially informed. Importantly, this condition does not require specifying the full structural search model. The approach relies on what we call *visible utility*, the component of utility visible prior to search. The main restriction we impose on search is that consumers never search items with lower visible utility and then leave unsearched items with higher visible utility (a condition we make precise below). We show that if this condition is satisfied, there is a function of choice probabilities which recovers preferences whether consumers are fully or partially informed. Given preferences recovered by our approach, we show that it is possible to identify other primitives of interests. For example, in a sequential search model à la Weitzman (1979), one can identify the distribution of search costs and thus simulate counterfactuals in which search costs are altered by information provision or other means. However, for many counterfactuals, such as simulating full information choices

¹The list of papers and their classification is available upon request from the authors.

²Even among the 126 articles in our survey that conduct welfare analyses and thus must take a stand on whether consumers are informed, 109 (86.5%) assume full information without testing this assumption.

and evaluating the benefits of full information via better choices, search costs need not be recovered at all. Preferences are identified—and thus can be estimated—without imposing a structural search model beyond the visible utility assumption. Specifically, no additional assumptions about option value, optimization vs. satisficing, simultaneous or sequential search, or distributional assumptions about beliefs and preferences are necessary for identification. Our results also suggest straightforward tests of full information, which we prove are powered within the large class of models covered by the visible utility assumption.

One can think of our approach as a data-driven method of isolating consumers who maximize utility. Consider the example of consumers purchasing items in a grocery store: nutritional information is accessible, but at some cost. Consumers may fail to maximize utility if they do not pay the cost to examine labels. In this case, visible utility represents utility from all non-nutrient sources, e.g. a combination of prices and perceived taste. Our assumption states that if you bother to check the nutrition label for good j , you will also check the label for a good j' that you would otherwise prefer were it equally nutritious. This assumption implies that consumers who search the most nutritious good *always choose the good that maximizes utility among all options* (which is not necessarily the most nutritious good). To see this, note that if some other good has higher utility than the most nutritious good, it has higher visible utility and thus is searched and then chosen by the consumer. Further, only consumers who search the most nutritious good are sensitive to nutrient content for that good. Therefore, by looking at the sensitivity of choices to the nutrient content of the most nutritious good we are able to isolate consumers that behave *as if* they were fully informed; standard arguments then recover their preferences.

To spell things out in more detail, consider first a J -good model with linear utility $U_{ij} = x_j\alpha + z_j\beta + \epsilon_{ij}$ where $\alpha > 0$ and $\beta > 0$.³ In the text, we extend this result to allow vector-valued x_j and z_j , as well as random coefficients and nonlinear utility. Suppose that consumer i observes x_j and ϵ_{ij} for all goods, but needs to engage in search to observe z_j . On the other hand, the researcher observes x_j , z_j , and choice probabilities s_j , but not ϵ_{ij} . With full information, we have $s_j = P(U_{ij} \geq U_{ij'} \forall j' \neq j)$ and we could estimate marginal rates of substitution using $\frac{\partial s_j}{\partial z_j} / \frac{\partial s_j}{\partial x_j} = \beta/\alpha$; in other words, β/α is identified by whether the choice probability for good j is more sensitive to z_j or x_j . If the underlying model is a search model in which consumers are informed about z_j only for some alternatives, then the standard approach will suffer from attenuation bias: $\left| \frac{\partial s_j}{\partial z_j} / \frac{\partial s_j}{\partial x_j} \right| < |\beta/\alpha|$.⁴ Some consumers will be insensitive to z_j variation not because they don't value it, but because they are not aware of it; thus, the observed sensitivity of choices to z_j will understate consumers' valuation of z_j .

For each individual i and good j , we define *visible utility* as $VU_{ij} \equiv x_j\alpha + \epsilon_{ij}$. We call this quantity “visible utility” because it defines the utility that i receives from good j based only on x_j and ϵ_{ij} , the attributes of goods that consumers can observe without engaging in search. Visible utility has the property that if $VU_{ij} \geq VU_{ij'}$ and $z_j \geq z_{j'}$, then $U_{ij} \geq U_{ij'}$. Given this definition, suppose consumer search is characterized by the following assumptions:

³We will show that the signs of α and β are identified, so assuming they are positive is without loss.

⁴We prove this in Section 2.1.

1. If i searches a good with visible utility VU_{ij} , she also searches all goods j' with $VU_{ij'} \geq VU_{ij}$.
2. Conditional on having utility \bar{u} in hand, consumer i searches j if and only if $g_{ij}(x_j, \bar{u}) \geq 0$ where g_{ij} is decreasing in \bar{u} .
3. Consumers choose the good which maximizes utility among searched goods.

The first assumption states that if you search a good with a given level of visible utility, you always search all goods with higher visible utility. This assumption would be satisfied if consumers searched in order of visible utility but it is weaker (e.g. consumers might search the second ranked good in terms of visible utility and then the first, provided they eventually search both). We relax this assumption in Section 3 by allowing for variables (e.g. rank on a webpage) to affect search but not utility and by considering the case where consumers form expectations about z based on x . The second assumption states that a consumer's decision to search good j depends on j 's visible utility and on the utility \bar{u} the consumer already has in hand. If consumers search sequentially, then \bar{u} is the highest utility uncovered so far in the search process. Then, the assumption says that the goods searched so far only matter for the decision to continue searching through the utility in hand. This rules out, for example, the case where consumers stop searching when they uncover a high enough z . Note that our framework also allows for search protocols where consumers decide which goods to inspect simultaneously. In this case, the function g_{ij} will only depend on j 's visible utility (see Example 5). The third assumption simply says that you maximize utility conditional on the goods you have searched.

These assumptions are consistent with a broad class of search models. For example, in a Weitzman (1979) search model where the priors and search costs are the same across goods (but the latter vary across consumers), it is optimal to search the good with the highest visible utility and decide whether to search the next good by comparing the expected benefits with search costs. These assumptions are also consistent with many behavioral models: consumers may myopically decide whether to continue searching by comparing utility in hand with expected utility of the next good (the “directed cognition” model of Gabaix, Laibson, Moloche, and Weinberg (2006)), consumers might engage in satisficing, i.e. searching in order of visible utility and stopping whenever utility in hand is good enough, or they might simultaneously search all goods with visible utility above a certain threshold and then give up. One important special case is when consumers search all goods, yielding a conventional full-information discrete choice model.

Our main result (for the case with linear utility and no random coefficients) is that, if the above assumptions are met, then $\frac{\partial^2 s_1}{\partial z_1 \partial z_j} / \frac{\partial^2 s_1}{\partial z_1 \partial x_j} = \beta / \alpha$ for $j \neq 1$, where good 1 is defined as the good with the largest value of the hidden attribute z (which, again, is known to the econometrician but not necessarily to the consumer). This expression holds for any models where consumers search according to our assumptions above, including the full information case. The intuition is as follows: under our assumptions, consumers who search good 1 will always choose the good that maximizes utility among all options (although they might not be fully informed).⁵ Therefore, z_1 only impacts choices for consumers who maximize utility. By

⁵Again, if some other good had higher utility than good 1, that good must have higher visible utility since good 1 has the best z .

examining how $\frac{\partial s_1}{\partial z_1}$ changes with attributes of rival goods, we are effectively learning about the impact of z_j and x_j on choices for consumers who act *as if* they were fully informed. For these consumers, both of these attributes impact choices only via U_j , and they do so in proportion to β/α . Since α is straightforwardly identified by examining choices where $z_j = z$ for all j ,⁶ this means that preferences are identified. The quantity $\frac{\partial^2 s_1}{\partial z_1 \partial z_j} / \frac{\partial^2 s_1}{\partial z_1 \partial x_j}$ thus provides robust estimates of marginal rates of substitution in the sense that it works regardless of the search process that consumers engage in provided the assumptions above are met.

How general is this result? Using additional derivatives of the share function, we can recover nonlinear utility functions $U_{ij} = v(x_j, z_j) + \epsilon_{ij}$. Additionally, the approach extends to random coefficients on product characteristics. Specifically, letting $U_{ij} = x_j \alpha_i + z_j \beta_i + \epsilon_{ij}$ we can recover the distribution of random coefficients (α_i, β_i) over a known grid. With a sufficiently long panel and time-invariant preferences, $U_{ijt} = v_i(x_j, z_j) + \epsilon_{ijt}$, we can recover individual-specific, possibly nonlinear utility functions $v_i(x_j, z_j)$. Thus, we can allow for a similar degree of unobserved heterogeneity as other constructive results on discrete choice demand with full information.⁷

Our identification proof lends itself naturally to estimation and testing. If one can nonparametrically estimate choice probabilities as a function of product attributes, then our results can be used to directly recover preferences. For example, when the number of goods or attributes is not too large, one could approximate the share functions via the method of sieves, which makes it possible to conveniently impose natural restrictions such as monotonicity and exchangeability across goods. When the number of goods is too large for nonparametric estimation, we suggest using “flexible logits,” which allow attributes of rival goods to directly enter utility for each alternative and thus flexibly parameterize the cross-derivatives that identify preferences in our model. Given estimates of choice probabilities, one can use our result to test for full information by checking whether our “search-robust” estimates of preferences are equal to the conventional estimates based on first derivatives. Additionally, the visible utility assumption implies bounds on observed choice probabilities and our model is overidentified; in the case of homogeneous, linear preferences, for example, $\frac{\partial^2 s_1}{\partial z_1 \partial z_j} / \frac{\partial^2 s_1}{\partial z_1 \partial x_j}$ will be equal for all goods $j \neq 1$.

To validate our approach, we attempt to recover preferences in a lab experiment where individuals engage in costly search. Individuals choose from sets of three books with visible titles, authors, genre, star ratings and prices, but hidden discounts that can only be observed at some cost. We place no restrictions on which goods consumers search given these constraints. In particular, we do not constrain search to satisfy the visible utility assumption. For each individual, we also observe treatments where consumers

⁶We show this formally in Section 2.

⁷Fox and Gandhi (2016) provide identification results for more general models allowing for both nonlinearity and flexible heterogeneity but these results are non-constructive and assume utility maximization; we use their result to identify parameters in corner cases where consumers maximize utility, but they otherwise are difficult to adapt to the more general case where choice probabilities need not maximize utility. This is in contrast to the constructive methods in Fox, Kim, Ryan, and Bajari (2012), who recover distributions satisfying the “Carleman condition,” which implies that the distribution of preferences is uniquely characterized by its moments. Alternatively, we recover weights for distributions supported on a known and fixed grid, in line with the approach of Fox, Kim, Ryan, and Bajari (2011). Berry and Haile (2009) and Berry and Haile (2014) also provide related results on nonparametric demand estimation. Their focus is on recovery of the conditional distribution of utilities rather than the structural parameters of utility; the latter are essential for our task of assessing whether consumers are informed about relevant attributes.

choose given full information. As expected, conventional logit estimates using the costly search data give attenuated coefficients on the discount variable relative to the full-information case. By contrast, our search-robust estimates successfully recover full-information preferences. Estimated choice probabilities also satisfy bounds implied by the visible utility assumption.

Our result relates to several existing literatures. A large theoretical and empirical literature investigates the formation of “consideration sets” (e.g. Roberts and Lattin (1991), Conlon and Mortimer (2013), Goeree (2008), and Gaynor, Propper, and Seiler (2016)). These papers attempt to estimate preferences when consumers may only consider some alternatives. Manzini and Mariotti (2014) and Abaluck and Adams (2017) are particularly closely related in attempting to characterize when we can recover consideration probabilities for alternative goods using choice data alone. This paper considers the complementary problem of imperfect information at the level of attributes rather than goods. The recent theoretical literature on this question includes Branco, Sun, and Villas-Boas (2012) and Ke, Shen, and Villas-Boas (2016). A few papers, such as Mehta, Rajiv, and Srinivasan (2003), Honka and Chintagunta (2016) and Ursu (2018) consider estimating utility by specifying full search models. We are, as far as we know, the first to provide formal identification results for preferences for a class of models without the need to commit to a specific structural search model.⁸

An important assumption in our model is that the searched attributes are observable to the econometrician, even if they are not known without searching to consumers. Ericson, Kircher, Spinnewijn, and Starc (2015) consider the related problem of inferring risk preferences separately from risk types using insurance choices. Their model differs from ours in that, in the special case they consider, the covariate “risk type” is not observed by the econometrician either.

A second related literature attempts to analyze whether consumers make informed choices by comparing the choices of regular consumers to that of a more informed subgroup. Bronnenberg, Dubé, Gentzkow, and Shapiro (2015) ask whether pharmacists make similar prescription drug choices to consumers, Handel and Kolstad (2015) ask whether better informed consumers make different health insurance choices, and Johnson and Rehavi (2016) study whether physicians treat differently when their patients are other physicians. Our paper can be thought of as a data driven way of isolating a subgroup of consumers who maximize utility (those who search the good with the highest value of the hidden attribute). Importantly, the subgroup we identify is *not necessarily informed* in the sense of searching all available goods. What we know is that they maximize utility among all options, meaning that their choices can be used to estimate preferences even though they do not themselves know that they do so, as they cannot verify that they have searched the good with the highest z_j without searching all goods. This insight—that we can recover preferences by

⁸There is one special case where the problem of imperfect information about attributes has been addressed in the existing literature. This is the case in which all attributes can be expressed in dollar terms. For example, consumers should not care whether a health insurance plan saves them \$100 in premiums or out of pocket costs (see Abaluck and Gruber (2011)), or whether a light bulb saves them money in upfront costs or shelf life (as in Allcott and Taubinsky (2015)). If one dollar-equivalent attribute is assumed to be visible to consumers, it can provide a benchmark for how consumers should respond to a hidden dollar-equivalent attribute. However, in many cases, attributes cannot easily be translated into dollars without first estimating consumer preferences. In these cases, our results still allow one to recover preferences given imperfectly informed consumers.

isolating consumers who maximize utility rather than those who are informed—is what makes possible the surprising recovery of preferences from choices in cross-sectional data without auxiliary information.

Section 2 lays out our formal framework and proves our identification results, Section 3 considers several empirically important extensions such as endogenous attributes, Section 4 provides details of estimation and simulation results, Section 5 reports results from our experiment, and Section 6 concludes.

2 Model

There are $J \geq 2$ goods indexed by $j = 1, \dots, J$ with attributes x_j observed by consumers and the econometrician and attribute z_j observed by the econometrician but not necessarily by consumers.^{9,10} For simplicity, we assume that x_j is scalar for all j , but the identification argument immediately extends to the case of vector-valued x_j 's. Without loss of generality, we let good 1 be the good with the largest value of z_j . Again, for simplicity, we focus on the case where z_j is scalar-valued. If there are multiple hidden attributes for each good j , call them z_{kj} , the same argument applies with good 1 defined as the good satisfying $z_{k1} \geq z_{kj}$ for all j and k .¹¹ Let individual i 's utility from alternative j be denoted by $U_{ij}(x_j, z_j)$. In what follows, we often omit the dependence of U_{ij} on (x_j, z_j) unless it is necessary to avoid confusion.

We can always write: $U_{ij} = a_{ij}(x_j) + b_{ij}(x_j, z_j)$ where $b_{ij}(x_j, 0) = 0$ (to see this, define $b_{ij}(x_j, z_j) = U_{ij}(x_j, z_j) - U_{ij}(x_j, 0)$). Since in our setting $a_{ij}(x_j)$ is the component of utility that is known to the consumer before engaging in search, we label it “visible utility,” VU_{ij} .

We make the following assumptions on the utility function.

Assumption 1. (i) For all i and j , U_{ij} is strictly monotonic in z_j .

(ii) For all i , the function $b_{ij}(x_j, z_j)$ is not alternative-specific, i.e. $b_{ij}(x_j, z_j) = b_i(x_j, z_j)$ for all j , and continuous in its first argument.

The class of utility functions satisfying Assumption 1 is broad and subsumes most specifications commonly used in empirical work as special cases, including logit with possibly nonlinear-in-characteristics utilities¹² and mixed-logit. For instance, in a mixed-logit model, one may specify $U_{ij} = \alpha_i x_j + \beta_i z_j + \epsilon_{ij}$. To map this specification into our notation, let $a_{ij}(x_j) = \alpha_i x_j + \epsilon_{ij}$, and $b_i(x_j, z_j) = \beta_i z_j$. As another example, consider the logit specification $U_{ij} = \alpha x_j + \beta z_j + \gamma x_j z_j + \epsilon_{ij}$. This is subsumed in our notation by letting $a_{ij}(x_j) = \alpha x_j + \epsilon_{ij}$, and $b_i(x_j, z_j) = \beta z_j + \gamma x_j z_j$.

Next, we state the assumptions that characterize the class of search models we consider.

⁹Our model also permits the more general case where attributes are potentially both good and individual-specific, but we write x_j and z_j rather than x_{ij} and z_{ij} for notational simplicity.

¹⁰Since our model only requires variation in x and z for two goods, any of the remaining $J - 2$ goods may be taken to be the outside option.

¹¹Note that, while in the scalar z_j case a good satisfying our definition of good 1 always exists (we can always rank options based on their z_j 's), this is not true of the vector-valued z_j case. In this sense, the argument for the vector-valued z_j case is more demanding of the data. The multiple attributes case also requires that search reveals the value of all attributes concurrently.

¹²We allow for nonlinearities subject to Assumption 1(i) being satisfied.

Assumption 2. (i) If consumer i searches a good with visible utility VU_{ij} , she also searches all goods j' with $VU_{ij'} \geq VU_{ij}$.

(ii) Conditional on having utility \bar{u} in hand, consumer i searches j if and only if $g_{ij}(x_j, \bar{u}) \geq 0$ where g_{ij} is decreasing in \bar{u} .

(iii) Consumers choose the good which maximizes utility among searched goods.

(iv) Only the value of z_j is unknown to consumers prior to search.

We discuss these conditions at length in Section 2.5. To briefly clarify, Assumption 2(i) states that if you search a good j , you always search all goods with higher visible utility than j . This would be satisfied if consumers searched in order of visible utility but it is weaker (e.g. consumers might search the second ranked good in terms of visible utility and then the first, provided they eventually search both). Assumption 2(ii) states that consumers decide whether or not to search a good based on their utility in hand and the visible utility of the good they are considering searching. This rules out, for example, a search protocol whereby one stops searching after discovering a large z_j irrespective of utility in hand. We subscript the function g by i to emphasize that the function may depend on any individual (unobserved) heterogeneity in utility or search. For example, in a Weitzman search model, the stopping rule would depend on consumer i 's reservation value, which in turn depends on i 's search cost. In addition, we subscript the g function by j' to indicate that g is allowed to vary with the identity of the good that is being considered as a candidate for search. In particular, g could vary with any component of the visible utility of good j' (but crucially and intuitively not with $z_{j'}$). Further, in a model where consumers pay a cost to search goods, the search costs may vary across goods (at least provided the visible utility assumption 2(i) still holds); this is again captured by the j' subscript. Assumption 2(iii) simply states that consumers must search a good before choosing it. Assumption 2(iv) (implicit in the model already) states that the econometrician observes all the information which is revealed by search.

We pause here to highlight that Assumption 2 accommodates several commonly used models of search.

Example 1 (Sequential Search). *Suppose that utility takes the form $U_{ij} = x_j\alpha_i + z_j\beta_i + \epsilon_{ij}$, consumers search sequentially and consumer i must pay a cost c_i every time she uncovers the z attribute for a good. Further, assume that the consumer has the same prior F_z for all goods. Then, following Weitzman (1979), the consumer will rank goods according to their reservation value rv'_{ij} defined implicitly by*

$$c_i = \int_{rv'_{ij}}^{\infty} (u - rv'_{ij}) dF_{U_{ij}}(u) = \int_{rv_i}^{\infty} \beta_i (t - rv_i) dF_z(t) \quad (1)$$

where $rv_i \equiv \frac{rv'_{ij} - \alpha_i x_j - \epsilon_{ij}}{\beta_i}$ and the last steps follows from a change of variable. We can interpret rv_i as the reservation value in units of z . To see this, note that consumer i ranks goods according to the visible utility $x_j\alpha_i + \epsilon_{ij}$ and for each good j' she chooses to uncover $z_{j'}$ if and only if the maximum utility secured so far is lower than $x_{j'}\alpha_i + rv_i\beta_i + \epsilon_{ij'}$. Once she stops searching, she maximizes utility among the searched goods. Thus, Assumption 2 is satisfied with $g_{ij}(x_j, \bar{u}) = x_j\alpha_i + rv_i\beta_i + \epsilon_{ij} - \bar{u}$.

Example 2 (Directed Cognition Model). Suppose that utility takes the form $U_{ij} = x_j\alpha_i + z_j\beta_i + \epsilon_{ij}$. Further, as in the model of Gabaix, Laibson, Moloche, and Weinberg (2006), consumers rank goods in terms of expected utility¹³ and myopically check whether searching the next good is worth the cost. The directed cognition model has the same g_{ij} function as the Weitzman model,¹⁴ but the order of search (and which goods are ultimately searched) may differ.

Example 3 (Satisficing). Suppose that consumer i searches in order of visible utility and stops whenever utility in hand is above a threshold τ_i . Then Assumption 2 is satisfied with $g_{ij}(x_j, \bar{u}) = \tau_i - \bar{u}$.

Example 4 (Full Information). The full information model is subsumed within the previous example by letting $\tau_i = \infty$ for all i .

Example 5 (Simultaneous Search). Suppose that utility takes the form $U_{ij} = x_j\alpha_i + z_j\beta_i + \epsilon_{ij}$ and that consumer i simultaneously searches all goods that have visible utility above a threshold $\tilde{\tau}_i$. Then Assumption 2 is satisfied with $g_{ij}(x_j, \bar{u}) = \alpha x_j + \epsilon_{ij} - \tilde{\tau}_i$.

Our results will not require the researcher to take a stand on the specific model of search that consumers follow (provided that our assumptions are met). Therefore, as illustrated by the examples above, the approach will be agnostic as to whether consumers search sequentially or simultaneously, are forward-looking or myopic and have biased or unbiased beliefs, among other things.

Throughout the rest of the paper, we assume that $\frac{\partial b_i}{\partial z_j} > 0$, i.e. we treat z_j as an attribute that customers value in good j .¹⁵ We are now ready to state and prove a lemma that is at the core of our results.

Lemma 1. Let Assumptions 1 and 2 hold and let $x_j \in [\bar{x} - \eta, \bar{x} + \eta]$ for all j , for some $\eta > 0$ sufficiently small. If consumer i searches good 1 (i.e. the good with the highest value of z), then i chooses the utility-maximizing good.

Proof. If good 1 is searched but utility is not maximized, then for some unsearched j , $U_{ij} > U_{i1}$. Since $z_1 > z_j$, by monotonicity, $b_i(\bar{x}, z_1) > b_i(\bar{x}, z_j)$. By continuity of b_i in its first argument, this implies that for η sufficiently small, $b_i(x_1, z_1) \geq b_i(x_j, z_j)$.¹⁶ Given this, $U_{ij} > U_{i1}$ implies $VU_{ij} > VU_{i1}$. But by Assumption 2(i), this implies that good j is searched, which is a contradiction. \square

¹³Note that we may assume without loss that $E(z_j) = 0$ for all j since the mean value of the hidden attribute (known by rational consumers before search) is subsumed by visible utility.

¹⁴The result that consumers in the fully rational Weitzman model decide whether to continue searching “as if” they were myopic is one of the main insights of Weitzman (1979).

¹⁵Conditional on Assumption 1(i), this is without loss, since Assumption 2 implies that an increase in U_{ij} can only induce consumer i to switch from not choosing j to choosing j , but never vice versa. Thus, by the chain rule, the sign of $\frac{\partial b_i}{\partial z_j}$ is identified by the sign of $\frac{\partial s_j}{\partial z_j}$, where s_j is the choice probability function for good j from the data.

¹⁶More formally, by continuity, for all $\delta > 0$ there exists $\eta > 0$ such that if $|x_1 - x_j| < 2\eta$, then $b_i(x_j, z_j) - b_i(x_1, z_j) \leq \delta$. Therefore, we have:

$$\begin{aligned} b_i(x_j, z_j) &= b_i(x_1, z_j) + b_i(x_j, z_j) - b_i(x_1, z_j) \\ &\leq b_i(x_1, z_j) + \delta \\ &\leq b_i(x_1, z_1) \end{aligned}$$

where the last inequality follows by choosing $\delta \equiv \frac{b_i(x_1, z_1) - b_i(x_1, z_j)}{2}$.

Note that Lemma 1 does *not* imply that good 1 always maximizes utility if it is searched. Rather, it implies that if good 1 is searched, the utility-maximizing good will also be searched (whether it is good 1 or not) and thus the consumer will choose that good. Note that the lemma also does *not* mean that consumers searching good 1 are fully informed (in a search model they typically will not be), but just that those consumers act *as if* they were fully informed. When utility is linear, the same result holds under weaker conditions on the variation in x across products.

Lemma 2. *Let $U_{ij} = x_j\alpha_i + z_j\beta_i + \epsilon_{ij}$ and let Assumption 2 hold. If consumer i searches good 1 (i.e. the good with the highest value of z), then i chooses the utility-maximizing good.*

Proof. If good 1 is searched but utility is not maximized, then for some unsearched j , $U_{ij} > U_{i1}$. Since $z_1 > z_k$ for all $k \neq 1$, it must be that $x_j\alpha_i + \epsilon_{ij} > x_1\alpha_i + \epsilon_{i1}$. But by Assumption 2(i), this implies that good j is searched, which is a contradiction. \square

Lemmas 1 and 2 will have far-reaching implications. To understand them, it will be convenient to define the choice probability for good j as:

$$s_j \equiv P \left(\left\{ U_{ij} = \max_k U_{ik} \text{ for } k \in \mathcal{G}_i \right\} \cap \{j \in \mathcal{G}_i\} \right) \quad (2)$$

where \mathcal{G}_i denotes the set of searched goods for individual i . Note that this probability is computed by integrating over any individual-specific unobserved heterogeneity in utility or search. Therefore, s_j is a function of $\mathbf{x} \equiv [x_1, \dots, x_J]$ and $\mathbf{z} \equiv [z_1, \dots, z_J]$, but we will often omit the dependence from the notation. Throughout the paper, the sources of unobserved heterogeneity will vary with the specific models we consider, so the symbol P will denote integrals over different distributions depending on the context.

Now, Lemmas 1 and 2 imply that z_1 only impacts choice probabilities for individuals who maximize utility. Therefore, looking at $\frac{\partial s_1}{\partial z_1}$ will isolate individuals who maximize utility and allow us to recover preferences using standard arguments. To formalize this, note that Lemmas 1 and 2 imply we can write:

$$s_1 = P(U_{i1} \geq U_{ik} \forall k) - P(\{U_{i1} \geq U_{ik} \forall k\} \cap \{\text{for some } j \neq 1, VU_{ij} \geq VU_{i1} \text{ and } g_{i1}(x_1, U_{ij}) \leq 0\}) \quad (3)$$

In other words, the probability that good 1 is chosen is the probability that good 1 is utility-maximizing minus the probability that good 1 is not searched *even though it is utility-maximizing*.¹⁷ Failing to search good 1 requires that there exists some other good j with $VU_{ij} \geq VU_{i1}$ and utility high enough that $g_{i1}(x_1, U_{ij}) \leq 0$.

Our proof will use the fact that certain derivatives of the choice probability functions are linear in the preference parameters we hope to recover with known (or recoverable) weights. We consider identification for three specifications of utility that satisfy Assumption 1:

1. Cross-sectional data where $U_{ij} = v(x_j, z_j) + \epsilon_{ij}$

¹⁷This follows because Lemmas 1 and 2 imply that if good 1 is chosen (and thus searched), good 1 must maximize utility.

2. Panel data where $U_{ijt} = v_i(x_{jt}, z_{jt}) + \epsilon_{ijt}$
3. Cross-sectional data where $U_{ij} = x_j\alpha_i + z_j\beta_i + \epsilon_{ij}$

These cases are comparable in generality to existing constructive identification results for preferences in full information discrete choice models, such as Fox, Kim, Ryan, and Bajari (2012). In each of these cases, we assume that the unobservables are independent of the product characteristics (\mathbf{x}, \mathbf{z}) . We will extend the approach to deal with endogenous attributes in Section 3.3.

2.1 Case 1: Cross-sectional data with $U_{ij} = v(x_j, z_j) + \epsilon_{ij}$

We start from the case where utility takes the form $U_{ij} = v(x_j, z_j) + \epsilon_{ij}$ for an unknown function v . In what follows, we use x and z to denote generic arguments of v and let $\epsilon_i \equiv [\epsilon_{i1}, \dots, \epsilon_{iJ}]$.

Theorem 1. *Let Assumption 2 hold and utility be given by $U_{ij} = v(x_j, z_j) + \epsilon_{ij}$ with v increasing in both arguments and infinitely differentiable. Further, assume that $\frac{\partial^2 s_1}{\partial z_1 \partial x_j}(\mathbf{x}^*, \mathbf{z}^*) \neq 0$ for some $(\mathbf{x}^*, \mathbf{z}^*)$ and $j^* \neq 1$, s_1 is infinitely differentiable and $\epsilon_i \perp (\mathbf{x}, \mathbf{z})$. Then, v is identified up to an additive constant.*

Proof. See Appendix A.1. □

This theorem applies to a broad class of utility functions. The cost of this level of generality is that it requires the share function s_1 to be infinitely differentiable. However, the marginal rates of substitution are recovered under much weaker differentiability requirements.

Corollary 1. *Let Assumption 2 hold and utility be given by $U_{ij} = v(x_j, z_j) + \epsilon_{ij}$ with v increasing and differentiable in both arguments. Further, assume that s_1 is twice differentiable and $\epsilon_i \perp (\mathbf{x}, \mathbf{z})$. Then, the marginal rates of substitution, $\frac{\partial v}{\partial z}(x, z) / \frac{\partial v}{\partial x}(x, z)$, can be recovered using:*

$$\frac{\frac{\partial^2 s_1}{\partial z_1 \partial z_j}(\mathbf{x}, \mathbf{z})}{\frac{\partial^2 s_1}{\partial z_1 \partial x_j}(\mathbf{x}, \mathbf{z})} = \frac{\frac{\partial v}{\partial z}(x, z)}{\frac{\partial v}{\partial x}(x, z)} \quad (4)$$

for all $j \neq 1$ such that $\frac{\partial^2 s_1}{\partial z_1 \partial x_j}(\mathbf{x}, \mathbf{z}) \neq 0$.

We postpone the proof of Theorem 1 to Appendix A.1. Here, we focus on the special case with linear utility and $J = 2$ goods to simplify the proof and facilitate intuition. In this case, the preference parameters can be fully recovered by looking at second derivatives.

Lemma 3. *Let utility be given by $U_{ij} = x_j\alpha + z_j\beta + \epsilon_{ij}$ and let Assumption 2 hold. Further, assume that $\frac{\partial^2 s_1}{\partial z_1 \partial x_2}(\mathbf{x}^*, \mathbf{z}^*) \neq 0$ for some $(\mathbf{x}^*, \mathbf{z}^*)$, s_1 is twice differentiable and $\epsilon_i \perp (\mathbf{x}, \mathbf{z})$. Then,*

$$\frac{\frac{\partial^2 s_1}{\partial z_1 \partial z_2}(\mathbf{x}^*, \mathbf{z}^*)}{\frac{\partial^2 s_1}{\partial z_1 \partial x_2}(\mathbf{x}^*, \mathbf{z}^*)} = \frac{\beta}{\alpha} \quad (5)$$

In addition, α is identified by focusing on markets with $z_j = z$ for all j and thus β is also identified.

Proof. First, we prove equation (5). In order to ease notation, we often suppress the subscript i in what follows. As above, good 1 is defined as the good with the highest value of z_j . Further, we let $\beta > 0$ without loss.¹⁸ Using Lemma 2, the probability of choosing good 1 can be written as:

$$s_1 = P(U_1 > U_2) - P(\{U_1 > U_2\} \cap \{\mathcal{G} = 2\}) \quad (6)$$

where, as above, \mathcal{G} denotes the set of searched goods. This follows because (i) if good 1 is utility-maximizing, you will always choose it unless you search only good 2; and (ii) you only choose good 1 if it is utility-maximizing, since otherwise, good 2 must have higher visible utility, meaning it must be searched (and chosen) if good 1 is searched.

Let $\tilde{u}_j \equiv x_j\alpha + z_j\beta$, so that $U_{ij} = \tilde{u}_j + \epsilon_{ij}$, and let $(\mathbf{x}, \mathbf{z}) = (\mathbf{x}^*, \mathbf{z}^*)$. Our goal will be to show that both z_2 and x_2 only impact $\frac{\partial s_1}{\partial z_1}$ via \tilde{u}_2 . This, in turn, implies that $\frac{\partial^2 s_1}{\partial z_1 \partial z_2} = \frac{\partial^2 s_1}{\partial z_1 \partial \tilde{u}_2} \frac{\partial \tilde{u}_2}{\partial z_2}$ and $\frac{\partial^2 s_1}{\partial z_1 \partial x_2} = \frac{\partial^2 s_1}{\partial z_1 \partial \tilde{u}_2} \frac{\partial \tilde{u}_2}{\partial x_2}$, and the result in equation (5) follows. To establish this, note that we can write:

$$\begin{aligned} P(\{U_1 > U_2\} \cap \{\mathcal{G} = 2\}) &= P(\{U_1 > U_2\} \cap \{VU_2 > VU_1\} \cap \{g_1(x_1, U_2) \leq 0\}) \\ &= P(\{U_1 > U_2\} \cap \{g_1(x_1, U_2) \leq 0\}) - P(\{VU_1 > VU_2\} \cap \{g_1(x_1, U_2) \leq 0\}) \end{aligned} \quad (7)$$

where the second line follows since $VU_1 > VU_2$ implies $U_1 > U_2$ and thus $P(\{VU_1 > VU_2\} \cap \{g_1(x_1, U_2) \leq 0\}) = P(\{U_1 > U_2\} \cap \{VU_1 > VU_2\} \cap \{g_1(x_1, U_2) \leq 0\})$. The second term in equation (7) is not a function of z_1 . The first term is only a function of x_2 and z_2 via \tilde{u}_2 . This, together with equation (6), is sufficient to show that both z_2 and x_2 only impact $\frac{\partial s_1}{\partial z_1} = \frac{\partial P(U_1 > U_2)}{\partial z_1} - \frac{\partial P(\{U_1 > U_2\} \cap \{\mathcal{G} = 2\})}{\partial z_1}$ via \tilde{u}_2 , thus proving equation (5).

Finally, we show that we can identify α using standard techniques by looking at choice sets where $z_j = z$ for all j . To see this, note that when $z_j = z$ for all j then consumers maximize utility if and only if they maximize visible utility. Since by assumption they always search the good with the highest visible utility, it follows that they maximize utility. Thus, one can pin down α by looking at how the choice probabilities vary with \mathbf{x} conditional on $z_j = z$ for all j , just like in the full information case.¹⁹ Given (5) and α , identification of β follows immediately. \square

Finally, we note that in many models of interest the conventional way of identifying preferences based on the ratio of *first* derivatives leads to understating consumers' taste for z . For this result, we further assume that the function $g_{ij}(x_j, \bar{u})$ is weakly increasing in x_j .²⁰ For simplicity, consider the model with

¹⁸This is without loss, since the sign of β is immediately identified from the data (footnote 15).

¹⁹There is one subtle exception to this argument. Suppose there is an outside option with utility normalized to 0, and we wish to identify a fixed effect which gives the utility of all inside goods relative to the outside good. In this case, consumers do not necessarily maximize utility when $z_j = z$ for all goods because consumers may decide to search *none* of the inside goods, and they may do so even when the outside good has lower utility than some of the inside goods if search costs are sufficiently high (in other words, an outside option may violate our assumption that consumers must search a good before they choose it). When consumers search none of the inside goods, it is never possible to separately identify whether consumers do not value the inside goods or have high search costs to examine any of the inside goods. It is possible to say something about the utility of consumers who are induced to search at least one of the inside goods when price is low enough (for example), but parametric assumptions are needed to make claims about the utility of consumers who never search any of the inside goods.

²⁰This condition is satisfied in all the search model considered above (Examples 1–5) when the coefficient on x in utility is

linear utility $U_{ij} = x_j\alpha + z_j\beta + \epsilon_{ij} \equiv \tilde{u}_j + \epsilon_{ij}$ and $\alpha > 0$. Let $\tilde{\mathbf{u}} \equiv (\tilde{u}_1, \tilde{u}_2)$ and

$$\begin{aligned} P_{j,2}^*(\tilde{\mathbf{u}}, \mathbf{x}) &\equiv P(\{U_j > U_{-j}\} \cap \{VU_{-j} > VU_j\} \cap \{g_{ij}(x_j, U_{-j}) \leq 0\}) \\ P_{j,3}^*(\tilde{\mathbf{u}}, \mathbf{x}) &\equiv P(\{U_{-j} > U_j\} \cap \{VU_j > VU_{-j}\} \cap \{g_{i-j}(x_{-j}, U_j) \leq 0\}) \end{aligned}$$

For good 1, $P_{1,3}^*(\tilde{\mathbf{u}}, \mathbf{x}) = 0$, but this is not necessarily true for other goods. More generally, $s_j = P(U_j > U_{-j}) - P_{j,2}^*(\tilde{\mathbf{u}}, \mathbf{x}) + P_{j,3}^*(\tilde{\mathbf{u}}, \mathbf{x})$.

Differentiating, we obtain:

$$\begin{aligned} \frac{\partial s_j}{\partial z_j} &= \beta \left[\frac{\partial P(U_j > U_{-j})}{\partial \tilde{u}_j} - \frac{\partial P_{j,2}^*}{\partial \tilde{u}_j}(\tilde{\mathbf{u}}, \mathbf{x}) + \frac{\partial P_{j,3}^*}{\partial \tilde{u}_j}(\tilde{\mathbf{u}}, \mathbf{x}) \right] \\ \frac{\partial s_j}{\partial x_j} &= \alpha \left[\frac{\partial P(U_j > U_{-j})}{\partial \tilde{u}_j} - \frac{\partial P_{j,2}^*}{\partial \tilde{u}_j}(\tilde{\mathbf{u}}, \mathbf{x}) + \frac{\partial P_{j,3}^*}{\partial \tilde{u}_j}(\tilde{\mathbf{u}}, \mathbf{x}) - \frac{1}{\alpha} \frac{\partial P_{j,2}^*}{\partial x_j}(\tilde{\mathbf{u}}, \mathbf{x}) + \frac{1}{\alpha} \frac{\partial P_{j,3}^*}{\partial x_j}(\tilde{\mathbf{u}}, \mathbf{x}) \right] \end{aligned}$$

Note that $\frac{\partial P(U_j > U_{-j})}{\partial \tilde{u}_j} - \frac{\partial P_{j,2}^*}{\partial \tilde{u}_j}(\tilde{\mathbf{u}}, \mathbf{x}) + \frac{\partial P_{j,3}^*}{\partial \tilde{u}_j}(\tilde{\mathbf{u}}, \mathbf{x}) = \frac{\partial s_j}{\partial \tilde{u}_j} \geq 0$.²¹ Further, $\frac{1}{\alpha} \frac{\partial P_{j,2}^*}{\partial x_j}(\tilde{\mathbf{u}}, \mathbf{x}) \leq 0$ and $\frac{1}{\alpha} \frac{\partial P_{j,3}^*}{\partial x_j}(\tilde{\mathbf{u}}, \mathbf{x}) \geq 0$ due to our assumptions about the function g .²² Therefore,

$$\frac{\frac{\partial s_j}{\partial z_j}}{\frac{\partial s_j}{\partial x_j}} \leq \frac{\beta}{\alpha} \quad (8)$$

2.2 Case 2: Panel data where $U_{ijt} = v_i(x_{jt}, z_{jt}) + \epsilon_{ijt}$

This case closely parallels the proof in the previous section. Now, rather than observing only $s_j(\mathbf{x}, \mathbf{z})$, the choice probabilities for each alternative as a function of the attributes, panel data allows us to observe $s_{ij}(\mathbf{x}, \mathbf{z})$, the choice probabilities for each individual as (\mathbf{x}, \mathbf{z}) vary over a long period of time. Given these, the following result holds:

Theorem 2. *Let Assumption 2 hold and utility be given by $U_{ijt} = v_i(x_{jt}, z_{jt}) + \epsilon_{ijt}$ with v_i increasing in both arguments and infinitely differentiable. Further, assume that $\frac{\partial^2 s_{i1}}{\partial z_1 \partial x_{j*}}(\mathbf{x}^*, \mathbf{z}^*) \neq 0$ for some $(\mathbf{x}^*, \mathbf{z}^*)$ and $j^* \neq 1$, s_{i1} is infinitely differentiable and $\epsilon_i \perp (\mathbf{x}, \mathbf{z})$. Then, v_i is identified up to an additive constant.*

Corollary 2. *Let Assumption 2 hold and utility be given by $U_{ijt} = v_i(x_{jt}, z_{jt}) + \epsilon_{ijt}$ with v_i increasing and differentiable in both arguments. Further, assume that s_{i1} is twice differentiable and $\epsilon_i \perp (\mathbf{x}, \mathbf{z})$. Then, the*

positive and corresponds to the mild requirement that consumers are (weakly) more prone to searching a good the higher the value of x for that good.

²¹Increasing u_j can only switch consumers from not choosing good j to choosing j but never the reverse. To see this, note first that conditional on searching any given set of goods, increasing u_j increases the probability good j is chosen. Second, changing u_j doesn't change the probability that good j is searched, which depends on $g_{ij}(x_j, U_{-j})$ for each alternative good which is searched. Third, changing u_j never makes other goods more likely to be searched. Specifically, an alternative good k is searched if and only if $g_{ik}(x_k, U_{k'}) \geq 0$ for all goods k' currently searched. This quantity is unchanged for $k' \neq j$ and weakly decreasing for $k = j$, so no good can become more likely to be searched. Therefore, $\frac{\partial s_j}{\partial \tilde{u}_j} \geq 0$.

²²Recall also that we assume $\alpha > 0$ without loss.

marginal rates of substitution, $\frac{\partial v_i}{\partial z}(x, z) / \frac{\partial v_i}{\partial x}(x, z)$, can be recovered using:

$$\frac{\partial^2 s_{i1}}{\partial z_1 \partial z_j}(\mathbf{x}, \mathbf{z}) / \frac{\partial^2 s_{i1}}{\partial z_1 \partial x_j}(\mathbf{x}, \mathbf{z}) = \frac{\partial v_i}{\partial z}(x, z) / \frac{\partial v_i}{\partial x}(x, z)$$

for all $j \neq 1$ such that $\frac{\partial^2 s_{i1}}{\partial z_1 \partial x_j}(\mathbf{x}, \mathbf{z}) \neq 0$.

The proofs of these two results exactly parallel the arguments in the previous section.

2.3 Case 3: Cross-sectional data where $U_{ij} = x_j \alpha_i + z_j \beta_i + \epsilon_{ij}$

The cases in the previous two sections assume either that we have panel data or that all individual heterogeneity is additively separable. Due to the difficulty of separately identifying preferences and search as well as more practical difficulties with estimation, most empirical structural search models that we are aware of do not allow for non-separable unobserved heterogeneity (see, e.g., Ursu (2018), Honka, Hortaçsu, and Vitorino (2017)).

Of course, we would like to understand both from a theoretical perspective whether the assumption of separable heterogeneity is required for identification and from a practical perspective whether our results are applicable in such cases. The canonical case of non-separable heterogeneity that has been studied in the literature and for which constructive identification results exist is that of the linear random coefficients model. We maintain linearity and impose two additional assumptions.

Assumption 3. (i) Utility is given by $U_{ij} = x_j \alpha_i + z_j \beta_i + \epsilon_{ij}$.

(ii) The coefficients α_i and β_i take values on a known finite support, i.e. $\alpha_i \in \{\alpha_1, \dots, \alpha_{K_\alpha}\}$ and $\beta_i \in \{\beta_1, \dots, \beta_{K_\beta}\}$ with probabilities given by $\tilde{\pi}_{k_\alpha, k_\beta} \equiv P(\{\alpha_i = \alpha_{k_\alpha}\} \cap \{\beta_i = \beta_{k_\beta}\})$. Further, the elements of $\{\beta_1, \dots, \beta_{K_\beta}\}$ all have the same sign and, without loss, we assume that they are positive.

(iii) The distribution of ϵ_i is known (or independently identified) and the three random vectors $\epsilon_i, (\alpha, \beta)$ and (\mathbf{x}, \mathbf{z}) are mutually independent.

Assumption 3(ii) follows Fox, Kim, Ryan, and Bajari (2011) in assuming that the random coefficients are supported on a finite and known grid of points. Given the restriction that $\{\beta_1, \dots, \beta_{K_\beta}\}$ all have the same sign, assuming that they are positive is without loss (see footnote 15). Assumption 3(iii) maintains knowledge of the distribution of all unobservables other than the random coefficients, consistent with recent papers on identification and estimation of demand (e.g. Fox, Kim, Ryan, and Bajari (2012), Fox, Kim, and Yang (2016)).

Theorem 3. *Let Assumptions 2 and 3 hold. If the market share of good 1 is $K_\alpha K_\beta$ -time differentiable, then the probability weights $\tilde{\pi}_{k_\alpha, k_\beta}$ for $k_\alpha = 1, \dots, K_\alpha$, $k_\beta = 1, \dots, K_\beta$ are identified.*

Proof. See Appendix A.2. □

As it is clear from the statement of Theorem 3, allowing for heterogeneity across consumers in preferences for attributes typically requires taking derivatives of order higher than two. Thus, identifying heterogeneous

preferences is more demanding of the data. When the sample size does not allow for direct application of our result, a natural approach is to impose more structure by specifying a structural search model. Theorem 3 may then be used to establish nonparametric identification of preferences within the specified model of search.²³

2.4 Alternative Approaches and Support Assumptions

So far we have not focused on the support assumptions required for identification. These are nonetheless essential to understand our contribution. Alternative approaches to identification exist which differ principally in requiring much stronger support assumptions.

For instance, one could assume that the data exhibits “at-infinity” variation to effectively go back to a setting that is analogous to full information. As the visible utility for a subset of goods grows to infinity (minus infinity), the probability of searching those goods goes to one (zero) under reasonable assumptions on the search process. Using this, one could identify preferences using conventional arguments. However, in practice, it is often implausible that any goods are searched with probability close to 1, so this strategy would require substantial parametric extrapolation.

In contrast, our proof requires much more plausible support assumptions. There is always a good which maximizes z_j when the latter is a scalar. To recover preferences in the homogeneous linear case, we only need sufficient variation to estimate second derivatives of s_1 at a single point. Of course, flexibly recovering a nonparametric function v or nonparametric random coefficients requires substantially more variation, requiring that we can estimate many higher order derivatives of choice probabilities. We discuss these challenges further in Section 4.

2.5 Discussion of Search Model Assumptions

To reiterate, we consider search models satisfying the following assumptions:

1. If consumer i searches a good with visible utility VU_{ij} , she also searches all goods j' with $VU_{ij'} \geq VU_{ij}$.
2. Conditional on having utility \bar{u} in hand, consumer i searches j if and only if $g_{ij}(x_j, \bar{u}) \geq 0$ where g_{ij} is decreasing in \bar{u} .
3. Consumers choose the good which maximizes utility among searched goods.
4. Only the value of z_j is unknown prior to search.

As discussed above, a sufficient condition for the first assumption is that consumers search in order of visible utility. In the Weitzman (1979) search model, consumers search goods in order of reservation utility, which is a function of the visible attributes of those goods, the distribution of the hidden attribute z_j , and search costs. If z_j is i.i.d. across goods and consumers have the same search cost for all goods, then it

²³See Section 4.3 for more on this point.

follows that consumers will search in order of visible utility.²⁴ There are three reasons this might fail in the Weitzman (1979) model: first, there may be more uncertainty about the hidden attribute for some goods than others, and this might lead individuals to search such goods first. Second, unobservables might be correlated across goods, and learning good news about good 1 might cause one to positively update about good 2 and choose to search it before good 3 even if $VU_3 > VU_2$. Third, search costs might vary across goods, meaning that consumers prefer to search first goods with lower search costs even if the payoff is potentially lower.

While the restriction that priors be i.i.d. and search costs be constant across goods is sufficient for Assumption 2(i) (the first assumption above), this is not necessary. Priors may be heterogeneous but consumers may be unsophisticated and fail to take into account option value, as in the directed cognition model studied in Gabaix, Laibson, Moloche, and Weinberg (2006). Consumers searching for a laptop online may enter some attributes into a search function and look at the items which rank highly according to those attributes without regard for whether a lower item is worth searching first because its value is more uncertain despite its lower average utility. Such examples also raise the natural concern that in many settings, factors like the order in which items appear in search may impact search costs separately from visible utility. Applications in the marketing literature often allow search costs to vary with observable attributes, such as the position of a good in search Ursu (2018). In Section 3.1, we extend our main result to these cases where some observable attributes impact search but not utility (thus violating the visible utility assumption stated above). We can also relax the i.i.d. priors assumption by allowing consumers to form beliefs about the hidden attribute as a function of observed attributes. In Section 3.2, we extend our approach to the case where beliefs about z_j are a linear function of observables.

Our second assumption on search is that consumers search good j if and only if $g_{ij}(x_j, \bar{u}) \geq 0$ where \bar{u} is utility in hand; we also impose the natural restriction that one is less likely to search as \bar{u} increases. This assumption is satisfied in most search models we are aware of in the literature, including Weitzman search, satisficing, simultaneously searching all goods with visible utility above a threshold, random search, and directed cognition. One exception is a model in which consumers commit to searching the top k goods in terms of visible utility prior to engaging in search. This model would violate the assumption because the function which determines whether you search good j cannot be written only as a function of VU_j since it will depend on the visible utility of all goods. We investigate the robustness of our approach to this violation of our assumptions in the simulations of Section 4.²⁵

Our third assumption, that consumers choose the good which maximizes utility among searched goods, embeds two separate ideas: the first is that consumers do not choose a good they have not searched, and the second is that they maximize utility given the information available. To the degree that search can plausibly be modeled as binary, the first assumption is relatively innocuous. In cases where the hidden attribute is

²⁴See Example 1.

²⁵Contrast this with a model where the decision to search depends only on the highest-utility good encountered so far. Since $g_{ij'}(x_{j'}, U_{ij}) \geq 0$ for the highest-utility good implies $g_{ij'}(x_{j'}, U_{ij}) \geq 0$ for all goods searched so far, such models can be accommodated by a $g_{ij'}$ function which depends separately on the utility of each good.

not easily available upon search, binary search may be implausible. For example, if the hidden attribute is “school-value added,” a consumer who searches more may learn about test scores and graduation rates, but these are signals of the underlying variable. Our model is more suitable for cases where the covariates included in the model are ones that consumers can observe directly given search. The assumption that consumers maximize utility given the information available can also be relaxed. One could specify a positive utility function that allows for consumer errors; as long as consumers maximize that positive utility function, the weight that they would attach to the hidden attribute given full information will be revealed. It is then up to the researcher whether to take this weight as the normative benchmark or whether to use some external standard.

The fourth assumption, that only the value of z_j is unknown prior to search, is more demanding. A consumer who clicks through to the product information page of an Amazon product might learn information about the attributes of a good (“the battery is compatible with USB-c”), but they also might learn information not observable to the econometrician (“one reviewer said the battery exploded into flames”). Our model does not accommodate cases where search reveals information not available to the econometrician. To the degree that this is a concern, one should try to code all information that may be revealed to the consumer upon search as attributes in the model. The viability of this strategy will of course vary from case to case.

2.6 Testing Search Model Assumptions with and without Observable Search

Our proof so far has proceeded as if search were not observed; that is, we observe final choices as a function of \mathbf{x} and \mathbf{z} but we do not observe which specific goods were searched. Datasets increasingly contain some information on what is searched: for example, in online clickstream data, one observes not only which product was purchased, but also which products were clicked on en route to purchase (e.g. Ursu (2018)). In many settings, it is plausible to assume that such clicks reveal which products were searched.

Can preferences be identified without resorting to our approach or an explicit search model in these cases? One might naively assume that our identification results would be unnecessary in such cases; given data on which products were searched, perhaps preferences can be estimated conditional on search without resorting to the assumptions we require here.

This is not generally the case because the unobservable component of utility may also drive the search decision. One example would be if search depends on ϵ . In such cases, goods with undesirable observables that are searched likely have an especially high realization of ϵ . Thus, it will appear from conditional choice probabilities as though the observable attributes are not so bad when in practice, individuals dislike those attributes but this dislike is offset by a large ϵ . A second reason unobservable components of utility might impact search is if preferences are unobservably heterogeneous (random coefficients). Even if search does not depend on ϵ , preferences cannot generally be recovered using only conditional choices unless IIA is satisfied.²⁶

²⁶To see why heterogeneous preferences create a problem, imagine products have quality ratings from 1-5. There are two types of consumers, one type that cares about quality and one type that does not. The type that cares about quality is

Thus, with heterogeneous preferences, the existing literature requires specifying a search model in order to estimate preferences even when search is observed. Our approach avoids the need to do this under the assumptions we have outlined. Once our approach is used to identify preferences, clickstream data can be used to conduct additional overidentifying tests if we assume that the distribution of ϵ_{ij} is known. In the linear case, visible utility is given by $VU_{ij} = x_j\alpha_i + \epsilon_{ij}$. As shown in Lemma 3, examining choices with equal values of the hidden attribute is sufficient to identify the distribution of α_i . Given α_i , the known distribution of ϵ_{ij} , and the number of goods searched $|\mathcal{G}_i|$, we can thus compute:

$$P(j \in \mathcal{G}_i | \mathbf{x}, \mathbf{z}) = \sum_k P(|\mathcal{G}_i| = k | \mathbf{x}, \mathbf{z}) P(j \in \mathcal{G}_i | |\mathcal{G}_i| = k, \mathbf{x}, \mathbf{z}) \quad (9)$$

since the first probability on the RHS is observed and the second is pinned down by the model assumptions (specifically, the fact that with k goods searched, those k goods must be the k goods with the highest visible utility). Checking (9) against the observed search probabilities provides a test of the model.

Even when we do not observe auxiliary information on which goods are searched, the assumptions in our model can be jointly tested by checking whether the observed choice probabilities are consistent with bounds implied by the estimated preferences and assumed search rule.

To construct an upper-bound on choice probabilities, note that a good j cannot be chosen if there is an alternative good with higher visible utility and higher utility. Thus, we have:

$$s_j(\mathbf{x}, \mathbf{z}) \leq 1 - P(U_{ik} \geq U_{ij} \text{ and } VU_{ik} \geq VU_{ij} \text{ for some } k) \quad (10)$$

The latter probability can be directly computed from knowledge of preferences and the distribution of ϵ .

To construct a lower-bound, note that the probability of choosing good j is at least as large as the probability that good j maximizes both utility and visible utility. That is:

$$s_j(\mathbf{x}, \mathbf{z}) \geq P(U_{ij} \geq U_{ik} \text{ and } VU_{ij} \geq VU_{ik} \text{ for all } k) \quad (11)$$

Once again, this probability can be computed given knowledge of preferences and the distribution of ϵ . We can then check whether our estimated choice probabilities are consistent with these bounds.

Finally, our model is overidentified. For example, in the case of linear utility and homogeneous preferences, we have shown that $\frac{\partial^2 s_1}{\partial z_1 \partial z_j} / \frac{\partial^2 s_1}{\partial z_1 \partial x_j} = \beta / \alpha$ for all alternative goods $j \neq 1$ and values of (\mathbf{x}, \mathbf{z}) . This provides a number of overidentifying restrictions which could be used to further test the model.

indifferent about quality over the 4-5 range, but values quality over the 1-4 range sufficiently that quality differences outweigh any other differences observable to consumers. Suppose that quality is observable to consumers (x) but price is only observed conditional on search (z). Quality conscious consumers only search goods with quality of at least 4. Other consumers will search all goods. If we estimate preferences conditional on search, we will wrongly conclude that no one cares about quality: quality conscious consumers don't care about quality given the goods they have searched (quality ranging from 4-5) and non-quality conscious consumers don't care about quality at all. To estimate preferences correctly, we would have to jointly model the decision of which goods to search and preferences conditional on searching.

2.7 Testing for Full Information

In the homogeneous preferences case of Section 2.1, our results suggest a natural test for full information. Under the null hypothesis of full information, $s_j = P(U_j \geq U_k \forall k)$ and therefore:

$$\frac{\partial^2 s_1}{\partial z_1 \partial z_j} / \frac{\partial^2 s_1}{\partial z_1 \partial x_j} = \frac{\partial s_k}{\partial z_{k'}} / \frac{\partial s_k}{\partial x_{k'}} = \frac{\partial v}{\partial z} / \frac{\partial v}{\partial x} \quad (12)$$

for all $j \neq 1$ and all k, k' . On the contrary, when consumers are unaware of z_j for some goods, then the ratios of first derivatives need not be equal to the ratios of the second derivatives. For example, equation (8) showed that $\frac{\partial s_1}{\partial z_1} / \frac{\partial s_1}{\partial x_1} \leq \frac{\partial^2 s_1}{\partial z_1 \partial z_j} / \frac{\partial^2 s_1}{\partial z_1 \partial x_j}$ in the class of search models we consider. Since both the ratios of first derivatives and the ratios of second derivatives in (12) are estimable from the data, this immediately leads to a test based on the discrepancy between the two sets of ratios. More specifically, given estimators of the share functions, one can compute a Wald test-statistic based on the discrepancy between the two sets of ratios and reject the null hypothesis of full information if the statistic exceeds a critical value.

Note that this test is valid even if our assumptions on the search process fail to hold since with full information the two sets of ratios will be equal regardless. When our assumptions on the search process do hold, we expect the test to have power, since the first derivative ratio will be attenuated relative to the true preferences, which are recovered by the cross-derivative ratio.

3 Extensions

In this section, we consider three extensions to the baseline model: the first two relax the visible utility assumption, while the third deals with the case of endogenous attributes.

3.1 Allowing for variables affecting search but not utility

One important case in which the visible utility assumption 2(i) is likely to fail is when factors exist which impact search costs but not utility. An example might be search position for online purchases. Arguably, search position impacts the order in which people search but has no direct impact on utility conditional on searching (Ursu 2018). In this case, consumers might first search items with higher search position even if they do not have higher visible utility. For example, if we randomly assign search order, this is likely to impact choices even though we are not changing the utility of each item conditional on search.

Our model from Section 2.1 can be extended to deal with cases where the factors which impact search but not utility are observable and the sign of their impact on search probabilities is known (such as position in search). Denote the variable which perturbs search but not utility by r_j , suppose that r_j is observed and that higher values of r_j make a good weakly more likely to be searched. Now, rather than assuming that goods are searched based on VU_{ij} alone, we assume that goods are searched based on $m(VU_{ij}, r_j)$ where m is strictly increasing in both VU_{ij} and r_j .

We formalize this, together with two additional restrictions, in the following assumption.

Assumption 4. (i) If consumer i searches j , then i also searches all j' s.t. $m(VU_{ij'}, r_{j'}) \geq m(VU_{ij}, r_j)$, where m is strictly increasing in both arguments;

(ii) There is at least one good $j \neq 1$ such that $r_j > r_1$;

(iii) The support of $(\mathbf{x}, \mathbf{z}) \mid (r_1, \dots, r_J)$ has positive Lebesgue measure for all (r_1, \dots, r_J) .

Assumption 4(iii) is substantive: for identification purposes, we consider variation in product characteristics holding fixed product search position. In practice, search position is likely to vary as a function of observables (e.g. products are sorted in order of price). However, because of the discrete nature of search position, we are likely to see variation conditional on search position and this is the variation we will use to identify our model.

Violations of the visible utility assumption due to search position will cause Lemma 1 to no longer hold as stated: the good with the highest value of z_j can be searched, another good j' may have higher utility (and thus higher visible utility), but good j' may not be searched because it has lower search position. However, an extension of Lemma 1 will still hold in this case, which then allows us to prove identification of preferences.

Lemma 4. *Let Assumptions 1, 2(ii)-2(iv), and 4 hold and let $x_j \in [\bar{x} - \eta, \bar{x} + \eta]$ for all j , for some $\eta > 0$ sufficiently small. Then, if consumer i searches good 1 (i.e. the good with the highest value of z), then i chooses the good which maximizes utility among all goods with $r_j \geq r_1$.*

Proof. Suppose there was a good j with $r_j \geq r_1$ and $U_{ij} > U_{i1}$ that consumer i does not search. We can follow the proof of Lemma 1 to show that $VU_{ij} > VU_{i1}$. By Assumption 4(i), this implies that good j is searched, which is a contradiction. \square

In other words, if higher search position only makes a good more likely to be searched, then goods with higher visible utility *and* higher search position will always be searched if good 1 is searched. Given this Lemma, we can apply a modification of the identification argument in Theorem 1 after conditioning on the subset of goods with higher search position than good 1 (defined as usual as the good with the largest value of z_j). We state the result here and prove it in the Appendix.

Theorem 4. *Let the assumptions of Lemma 4 hold and let utility be given by $U_{ij} = v(x_j, z_j) + \epsilon_{ij}$ with v increasing in both arguments and infinitely differentiable. Further, assume that $\frac{\partial^2 s_1}{\partial z_1 \partial x_{j^*}}(\mathbf{x}^*, \mathbf{z}^*) \neq 0$ for some $(\mathbf{x}^*, \mathbf{z}^*)$ and $j^* \neq 1$, s_1 is infinitely differentiable and $\epsilon_i \perp (\mathbf{x}, \mathbf{z})$. Then, v is identified up to an additive constant.*

Proof. See Appendix A.3. \square

3.2 Allowing for consumers' expectations on z to depend on x

Another reason why the visible utility assumption 2(i) might fail is that consumers could form expectations about z based on x . For instance, if x is price and z is quality, consumers might infer that more expensive products tend to be higher quality. As a consequence, if they value quality to a sufficient degree relative to

price, they may search a high-priced product and not search a low-priced product even if the former has a lower visible utility than the latter.

In our proofs so far, we have assumed that consumers do not update about z_j given x_j . We now show that we can identify preferences given consumer beliefs about z_j given x_j in a linear model. Further, under additional assumptions, we will show that we can identify β/α , the relative value of the hidden attribute, even when beliefs are unknown. In other words, we can do so without taking a stand on whether consumers have rational expectations and form beliefs based on the empirical relationship between z_j and x_j or naively update.

Consider the linear model $U_{ij} = x_j\alpha + z_j\beta + \epsilon_{ij}$ and re-write it as

$$\begin{aligned} U_{ij} &= x_j\alpha + (z_j - E(z_j|x_j))\beta + E(z_j|x_j)\beta + \epsilon_{ij} \\ &= \beta\gamma_0 + x_j(\alpha + \beta\gamma_1) + \tilde{z}_j\beta + \epsilon_{ij} \end{aligned}$$

where the second equality assumes that consumers use the linear projection $E(z_j|x_j) = \gamma_0 + \gamma_1 x_j$ and we let $\tilde{z}_j \equiv z_j - E(z_j|x_j)$. Visible utility is then given by $\beta\gamma_0 + x_j(\alpha + \beta\gamma_1) + \epsilon_{ij}$ and consumers learn the deviation from their expectation on z_j , \tilde{z}_j , upon searching. For the moment, we assume that γ_1 is known to the researcher. Note that γ_0 is not identified, but also does not generally impact choices since it enters utility as an additive constant.²⁷ Given γ_1 , we can identify the ranking of goods in terms of \tilde{z} and we label good 1 as the good with the largest value of \tilde{z} . Then, an argument analogous to that in Lemma 3 yields identification of $\frac{\beta}{\alpha + \beta\gamma_1}$.²⁸ We can also recover $\alpha + \beta\gamma_1$ in a manner that parallels our usual identification of α (Lemma 3). When $\tilde{z}_j = \tilde{z}$ for all j , consumers who search based on our visible utility assumption always maximize utility, and thus we can directly estimate $\alpha + \beta\gamma_1$ as the coefficient on x_j for those consumers (we provide a formal proof in Appendix A.4.1). Therefore, this gives separate identification of β and α given γ_1 .

When γ_1 is unknown, we can identify β/α if we know its sign and make a further support assumption. Suppose that the sign of γ_1 is known (e.g. higher priced goods have weakly higher quality). Without loss, we assume $\gamma_1 > 0$. In addition, suppose that there exist choice sets in which a good has both the highest value of z and the lowest value of x . Even when γ_1 is unknown, this good is known to maximize \tilde{z} ; thus, we can label it by 1. Note that we cannot differentiate with respect to \tilde{z} as in the case above since γ_1 and thus \tilde{z} is unknown. However, with good 1 defined appropriately, Corollary 1 shows that cross-derivatives with respect to z_1, z_j, x_j for $j \neq 1$ identify β/α (specifically, consumers who search the good with the highest value of \tilde{z} will always maximize utility, and so their sensitivity to x_j and z_j identifies their true preferences). The quantity β/α is not sufficient to simulate choices with full information, since we cannot

²⁷There is one exception to the above claim, which is the case when there is an outside option for which the x and z attributes are not defined, so that a systematic bias in beliefs about the distribution of z_j given x_j would change the relative value of all the inside goods relative to that outside option. This might mean that the relative utility of the outside option cannot be separately identified from γ_0 ; the model could still be estimated, but the normative interpretation of fixed effects might change.

²⁸In Lemma 3, we showed identification of $\frac{\beta}{\alpha}$ by taking derivatives of s_1 w.r.t. z_1, z_2, x_2 . Similarly, here we obtain identification of $\frac{\beta}{\alpha + \beta\gamma_1}$ by taking derivatives of s_1 wrt $\tilde{z}_1, \tilde{z}_2, x_2$.

tell how responsive consumers would be to x_j were choices fully-informed. However, it is sufficient to identify the relative value placed on the hidden attribute as well as to conduct tests for full information as in Section 2.7.

3.3 Endogenous attributes

So far, we have assumed that the observed product attributes are independent of all unobservables. This is restrictive, especially in settings in which product attributes — notably price — are chosen by firms who might know more about consumer preferences relative to what is captured by the observed data. As highlighted by a large literature (e.g. Berry, Levinsohn, and Pakes (1995)), this typically leads to correlation between the attributes chosen by firms and product-level unobservables.

Here we consider an extension of our model that allows for endogenous product attributes. We specify the utility that consumer i gets from good j as

$$U_{ij} = \alpha x_j + \beta_i z_j + \lambda_i p_j + \xi_j + \epsilon_{ij} \quad (13)$$

where p_j denotes the endogenous characteristic and ξ_j is a product-specific characteristic that is known by consumers before search, but is not observed by the researcher.²⁹ If firms also know ξ_j when choosing p_j , then the two will typically be correlated, thus leading to endogeneity of p_j . We consider both the case where p_j is part of visible utility and that in which consumers need to search good j to uncover p_j (as well as possibly other non-endogenous attributes z_j). If p_j is price, the first scenario corresponds to settings such as e-commerce where typically price is visible on the results page and does not require any further clicking by the user. On the other hand, the second scenario covers cases in which price is itself the object of consumer search (there is a large literature on this, particularly in relation to the often observed price dispersion for relatively homogeneous goods; see, e.g., Stahl (1989), Hong and Shum (2006) and Hortaçsu and Syverson (2004)). In what follows, we will show identification of preferences for each of these two cases.

To this end, we introduce two mutually exclusive variants of assumption 2(ii). Let $\delta_j = \alpha x_j + \xi_j$ for all j .

Assumption 5. (i) The attribute p_j is part of the visible utility of good j . Conditional on having utility \bar{u} in hand, consumer i searches j if and only if $g_{ij}(\delta_j, p_j, \bar{u}) \geq 0$ where g_{ij} is decreasing in \bar{u} .

(ii) The attribute p_j is uncovered by consumers only upon searching good j . Conditional on having utility \bar{u} in hand, consumer i searches j if and only if $g_{ij}(\delta_j, \bar{u}) \geq 0$ where g_{ij} is decreasing in \bar{u} .

Like Assumption 2(ii), Assumption 5 states that consumers decide whether to search good j based on

²⁹Note that the utility specification in (13) allows for random coefficients on both z_j and p_j , but not on x_j . This is stronger than needed, since the identification argument below only requires that x_j and ξ_j enter the demand functions via a linear index. Thus, another possible specification is

$$U_{ij} = \tilde{\alpha}_i (\alpha x_j + \xi_j) + \beta_i z_j + \lambda_i p_j + \epsilon_{ij}$$

The latter is weaker, but also less common in the discrete choice literature, so we focus on model (13) in what follows.

utility in hand and the visible utility of j . Letting $\delta = (\delta_1, \dots, \delta_J)$, we may write the share of good j as

$$s_j = \sigma_j(\delta, \mathbf{z}, \mathbf{p}) \quad (14)$$

for some function σ_j . Repeating this for all j and stacking the equations, we obtain a demand system of the form

$$\mathbf{s} = \sigma(\delta, \mathbf{z}, \mathbf{p}) \quad (15)$$

where $\mathbf{s} = (s_1, \dots, s_J)$. We also define the share of the outside option as $s_0 \equiv 1 - \sum_{j=1}^J s_j$, with associated function $\sigma_0(\delta, \mathbf{z}, \mathbf{p})$. We establish nonparametric identification of this demand system by invoking results from Berry and Haile (2014) (henceforth, BH).³⁰ Specifically, the results in BH yield identification of $(\xi_j)_{j=1}^J$ for every unit (individual or market) in the population. This means that all the arguments of σ are known, which immediately implies (nonparametric) identification of σ itself. Once σ is identified, one may apply our results in Section 2.3 to identify the distribution of the preference parameters α , β_i and λ_i . Note that, while knowledge of σ is sufficient for several counterfactuals of interest (e.g. computing equilibrium prices after a potential merger or tax), the preference parameters are required to predict how choices and welfare would change if consumers were given full information, among other things. In this sense, our approach complements the identification results in BH within the class of search models we consider.

To prove identification of σ , we first note that model (14) satisfies the index restriction in BH's Assumption 1. Second, we assume that we have excluded instruments \mathbf{w} which, together with the exogenous attributes, satisfy the following mean-independence restriction

$$\mathbb{E}(\xi_j | \mathbf{x}, \mathbf{z}, \mathbf{w}) = 0 \quad \text{for all } j \quad (16)$$

almost surely (Assumption 3 in BH) and assume that the instruments shift the endogenous variables (market shares and endogenous attributes \mathbf{p}) to a sufficient degree (as in BH's Assumption 4). Finally, we verify that the demand system satisfies the “connected substitutes” restriction defined in BH's Assumption 2. To this end, we prove the following result.

Lemma 5. *Let utility be given by (13) with ϵ_i supported on \mathbb{R}^J and let Assumptions 2(i), 2(iii), 2(iv), and either 5(i) or 5(ii) hold. Then, for all $j, k = 1, \dots, J$ with $j \neq k$, σ_j is (i) strictly increasing in δ_j and (ii) strictly decreasing in δ_k .*

Proof. See Appendix A.5 for the proof. □

Lemma 5 implies that the goods are connected substitutes in δ (see Definition 1 in BH), which in turn allows us to prove identification of σ by invoking Theorem 1 in BH.³¹ Since Lemma 5 holds under either

³⁰See also Berry, Gandhi, and Haile (2013).

³¹Note that the proof of Theorem 1 in BH only uses the fact that goods are connected substitutes in δ , not in $-\mathbf{p}$.

Assumption 5(i) or 5(ii), we obtain identification of preferences both in the case where p_j is part of the visible utility of good j and in the case where p_j is only uncovered upon searching j . Moreover, Theorem 1 of BH implies that one can invert the demand system σ for the indices δ and write

$$\alpha x_j + \xi_j = \sigma_j^{-1}(\mathbf{s}, \mathbf{z}, \mathbf{p}) \quad (17)$$

for all j . Equations (17) and (16) naturally lead to a nonparametric instrumental variable approach to estimate σ_j^{-1} (and thus σ_j).³²

4 Estimation

Our identification result shows that preferences can be recovered given knowledge of the choice probability function for good 1, denoted by $s_1(\mathbf{x}, \mathbf{z})$. To this end, note that the model implies the following conditional moment restrictions

$$E(y_j - s_j(\mathbf{x}, \mathbf{z}) | \mathbf{x}, \mathbf{z}) = 0 \quad \forall j \quad (18)$$

where y_j is a dummy variable equal to 1 if a consumer chooses good j .³³

Thus, methods designed to estimate conditional moment restriction models can be used. Of course, the performance of an estimator will depend on how flexibly it captures the derivatives that identify preferences in our approach.

Here, we consider two approaches to estimating $s_1(\mathbf{x}, \mathbf{z})$: (i) an approximation via Bernstein polynomials which is viable when the number of goods and attributes is small; and (ii) a “flexible logit” model which is more ad hoc, but scales better as the number of goods increases.

Throughout this section, we focus on the linear homogeneous case of $U_{ij} = x_j\alpha + z_j\beta + \epsilon_{ij}$. Our result in Section 2.1 shows that β/α can be recovered from $\frac{\partial^2 s_1}{\partial z_1 \partial z_j} / \frac{\partial^2 s_1}{\partial z_1 \partial x_j}$ for $j \neq 1$. As discussed in Section 2.3, the model with linear, homogeneous preferences is the current standard in the empirical literature on search.

In more general non-linear or random coefficients models, our identification arguments require recovery of higher-order derivatives and thus might not directly translate into viable estimation strategies in small to medium sample sizes or with a large number of goods. In these cases, the best way forward might be to parametrically specify a full structural search model and estimate it via standard methods, e.g. MLE. We would then view our identification results as providing reassurance that preferences are indeed identified, something that had not been previously established in the literature (see Section 4.3 and Appendix A.7 for more on this). Additionally, given the estimated structural search model, if it satisfies the visible utility assumption, the tests in Section 2.6 could be used to check consistency with the data.

³²Compiani (2019) proposes to approximate σ_j^{-1} using Bernstein polynomials. We use a similar approach in Section 4 to estimate the demand function for the case without endogeneity.

³³Here, we focus on the case where data on individual-level choices are available, as in the experiment of Section 5. However, our identification approach could also be applied to aggregate (i.e. market share) data as long as one can consistently estimate the share functions s_j .

4.1 Approximation via Bernstein polynomials

Following Compiani (2019), one can approximate the demand function via Bernstein polynomials. This allows the researcher to impose natural restrictions on the demand functions via linear (and thus easy-to-enforce) constraints on the coefficients to be estimated. Specifically, we impose monotonicity in \mathbf{x} and \mathbf{z} (s_j increasing in x_j and z_j and decreasing in \mathbf{x}_{-j} and \mathbf{z}_{-j}) as well as exchangeability across goods.³⁴ The purpose of these restrictions is twofold. First, they discipline the estimation routine in the sense that they help obtain reasonable estimates of quantities of interest (e.g. negative price elasticities). Second, they help partially alleviate the curse of dimensionality that arises as the number of goods increases. We estimate the Bernstein coefficients by minimizing a quadratic form based on the restrictions in (18) subject to the constraints.³⁵

Simulation Results To test the performance of our approach, we consider several simulations. In all simulations, we generate $N = 20,000$ choices with utility given by:

$$U_{ij} = \alpha x_{ij} + \beta z_{ij} + \epsilon_{ij} \quad (19)$$

with $\alpha = \beta = 1$, $x_{ij} \sim_{i.i.d} N(0, 1)$, $z_{ij} \sim_{i.i.d} N(0, 1)$, and ϵ_{ij} i.i.d. Type 1 extreme value.

We simulate data from four data generating processes, three of which satisfy the assumptions of our theorem and one of which does not. These are:

1. Weitzman search, with search costs $c \sim \text{LogNormal}(-2, 2.25)$
2. Satisficing, searching in order of visible utility until utility-in-hand is at least $T \sim \text{LogNormal}(-0.35, 2.25)$
3. Search all goods with visible utility above a threshold given by $c \sim N(-1, 16)$ (if no goods are above the threshold, search and choose the good with the highest visible utility)
4. Randomly search $k \in \{1, \dots, J\}$ goods, where the searched goods are the k highest in terms of visible utility

DGPs 1-3 satisfy our assumptions. By contrast, DGP 4 violates Assumption 2(ii) because the decision of whether to search a good does not just depend on that good's visible utility, but on the visible utilities of all goods.

Table 1 reports results from the Bernstein approximation of the cross-derivative ratio which identifies β/α . For comparison, we also report estimates of $\frac{\partial s_j / \partial z_j}{\partial s_j / \partial x_j}$, which would recover β/α with full information. In all cases, the estimates based on first-derivatives are attenuated relative to the true values. This occurs for the reason discussed in Section 2.1: consumer insensitivity to variation in z for goods that are not searched biases the coefficients towards zero. In contrast, the confidence intervals from Bernstein estimation of the

³⁴See Compiani (2019) for a formal definition of exchangeability. In words, a demand system is exchangeable if it only depends on the attributes of the goods (in our case x_j and z_j), but not their identity.

³⁵Compiani (2019) provides more details on the implementation of the estimator.

cross-derivative ratio include the true values in DGPs 1-3, and are fairly precise for the $J = 3$ case. For DGP 4, where the assumptions of our model do not hold, the coefficient is attenuated for $J = 3$, although the point estimates remain much closer to the true values the first-derivative estimates.

Table 1: Bernstein Approximation

DGP	Number of Goods			
	2		3	
	First-Derivatives	Cross-Derivatives	First-Derivatives	Cross-Derivatives
1	0.610 (0.024)	0.977 (0.304)	0.403 (0.012)	0.997 (0.076)
2	0.691 (0.024)	1.280 (0.538)	0.361 (0.014)	0.935 (0.068)
3	0.527 (0.021)	0.870 (0.190)	0.330 (0.010)	0.872 (0.071)
4	0.444 (0.018)	0.801 (0.301)	0.206 (0.010)	0.626 (0.075)

Note: Across all rows, the data the sample size is $N = 20,000$ and the data in each row is generated by the corresponding DGP described in the main text. In all cases, the true value is 1 and standard errors are in parentheses.

4.2 “Flexible Logit”

With a large number of goods, nonparametric methods face a curse of dimensionality, and thus it becomes necessary to place some parametric structure on the problem. In this section, we develop one such parametric approximation which performs well in simulations for a relatively large number of goods.

To motivate our parametric approach to estimating $s_1(\mathbf{x}, \mathbf{z})$, note that full-information logit models typically impose strong restrictions on the structure of the derivatives of choice probabilities. Specifically, if $u_{ij} = v_j^* + \epsilon_{ij}$ and ϵ_{ij} is i.i.d. extreme value where v_j^* is a differentiable function of x_j and z_j , then for $q_j \in \{x_j, z_j\}$:

$$\begin{aligned}
\frac{\partial s_j}{\partial q_j} &= \frac{\partial s_j}{\partial v_j^*} \frac{\partial v_j^*}{\partial q_j} = \frac{\partial v_j^*}{\partial q_j} s_j (1 - s_j) \\
\frac{\partial s_j}{\partial q_{j'}} &= \frac{\partial s_j}{\partial v_{j'}^*} \frac{\partial v_{j'}^*}{\partial q_{j'}} = -\frac{\partial v_{j'}^*}{\partial q_{j'}} s_j s_{j'} \\
\frac{\partial^2 s_j}{\partial z_j \partial q_{j'}} &= -\frac{\partial v_j^*}{\partial q_{j'}} \frac{\partial v_{j'}^*}{\partial z_j} s_j s_{j'} (1 - 2s_j)
\end{aligned} \tag{20}$$

for $j' \neq j$. Thus, in a conventional logit model, $\frac{\partial^2 s_1}{\partial z_1 \partial z_{j'}} / \frac{\partial^2 s_1}{\partial z_1 \partial x_{j'}} = \frac{\partial s_1}{\partial z_{j'}} / \frac{\partial s_1}{\partial x_{j'}} = \frac{\partial v_{j'}^*}{\partial z_{j'}} / \frac{\partial v_{j'}^*}{\partial x_{j'}}$ for all $j' \neq 1$, and this further equals $\frac{\partial s_1}{\partial z_1} / \frac{\partial s_1}{\partial x_1}$ when $\frac{\partial v_j^*}{\partial z_j} = \frac{\partial v_{j'}^*}{\partial x_{j'}}$ for all j, j' . We would like to estimate a model of s_1 which is sufficiently flexible that ratios of first-derivatives differ from ratios of second cross-derivatives, as will generally occur if consumers engage in search. To allow for this additional flexibility, we let the utility for

good 1 depend directly on attributes of rival goods as follows:

$$v_1 = \tilde{v}(x_1, z_1) + b_1 z_1 + \sum_{k \neq 1} (\gamma_k w_{z1k} z_k + \gamma_{2k} w_{x1k} x_k + w_{z2k} \delta_k z_k z_1 + w_{x2k} \delta_{2k} x_k z_1) \quad (21)$$

where $\tilde{v}(x, z)$ is a differentiable function of x and z , w_{z1k} , w_{x1k} , w_{z2k} and w_{x2k} are known weights, and b_1 , γ_k , γ_{2k} , δ_k and δ_{2k} are coefficients to be estimated. Further, we let $v_k = \tilde{v}(x_k, z_k)$ for $k \neq 1$. In Appendix A.6, we describe one way of choosing the weights which we find works well in simulations, and for which the ratio of second derivatives (which recovers β/α) is a particularly convenient function of model parameters. We note that the parameters in (21) do not have the usual interpretation (i.e. we are not positing that the actual utility of good 1 depends on the attributes of good k for $k \neq 1$). Instead, (21) is simply a flexible function of (\mathbf{x}, \mathbf{z}) that captures the second derivatives of s_1 well.

Simulation Results For each of the DGPs described in Section 4.1, we consider simulations with $J \in \{2, 3, 5, 10\}$. Additionally, we report estimates from the flexible logit model as well as the naive logit model. We bootstrap the standard errors using 250 repetitions.

Results from these simulations are reported in Table 2. The table shows estimates of β/α from a conditional logit model with no adjustment for imperfect information, as well as the cross-derivative ratio estimates from the flexible logit model. In the naive logit model, the coefficient is attenuated, typically biased towards zero by 30-50%. The flexible logit model performs substantially better, with 95% confidence intervals including the true estimates in DGPs 1-3. Perhaps surprisingly, the flexible logit model also performs well for DGP 4; the confidence intervals include the true values for 2 and 5 goods, and have less bias than the naive logit model for 5 and 10 goods.

Table 2: Estimator based on cross-derivatives ratio (flexible logit) vs naive logit

DGP	Number of Goods							
	2		3		5		10	
	Naive	Flexible	Naive	Flexible	Naive	Flexible	Naive	Flexible
1	0.6590 (0.0158)	1.0214 (0.1208)	0.6330 (0.0122)	1.0671 (0.1259)	0.6050 (0.0095)	0.9633 (0.1254)	0.5770 (0.0089)	0.8986 (0.1053)
2	0.7403 (0.0162)	0.9976 (0.1034)	0.6194 (0.0135)	1.0854 (0.1300)	0.4587 (0.0102)	1.0407 (0.1578)	0.2909 (0.0083)	1.0004 (0.2603)
3	0.5424 (0.0149)	1.1177 (0.1716)	0.5945 (0.0117)	1.0286 (0.1469)	0.6543 (0.0099)	0.9017 (0.1071)	0.7246 (0.0106)	0.8822 (0.0733)
4	0.4543 (0.0140)	1.1358 (0.1906)	0.5568 (0.0118)	0.9614 (0.1659)	0.6691 (0.0105)	0.8015 (0.1012)	0.7887 (0.0104)	0.8151 (0.0679)

Note: Across all rows, the data the sample size is $N = 20,000$ and the data in each row is generated by the corresponding DGP described in the main text. “Naive” refers to estimates of β/α from a conventional logit model, and “Flexible” refers to estimates from the flexible logit model. In all cases, the true value is 1 and standard errors are in parentheses.

4.3 When Should You Estimate a Search Model?

Our results can be used to identify preferences under the assumptions we outline without estimating a full search model. In this section, we discuss for which counterfactuals preference estimation is sufficient. Additionally, we discuss how our results can be used to aid in estimation of search costs given a fully specified search model or for specification testing after such a search model is estimated.

One important class of counterfactuals asks: how would consumers choose if search costs were reduced? For intermediate values of search costs, this question would require an explicit search model. On the other hand, in many cases, we want to know how consumers would choose with *zero* search costs. This question can be answered given knowledge of preferences and available choice sets without a search model. Given the implied choice probabilities, we can also evaluate the benefits of perfect information via better choices by comparing full information choices with status quo choices.

A full normative evaluation of an information intervention would also directly include search costs: information may benefit consumers both by helping them make better choices and by helping them make choices more easily, and search costs quantify the value of making choices more easily. Structural search models can aid in identification of these search costs. Of course, this exercise is intrinsically challenging because such structural estimates are often sensitive to the specific assumptions made about the search process. Back of the envelope estimates of search costs based on survey data or other information on the time consumers spend choosing may be more credible than structural estimates.

Nonetheless, if search costs are of interest either to conduct counterfactuals with intermediate levels of search costs or for normative purposes, our model can be used to identify search costs given preferences and an underlying structural search model. In Appendix A.7, we give an explicit example of how search costs can be recovered in a Weitzman model once preferences are known. Intuitively, when preferences are known, we know how consumers would respond to the hidden attribute with zero search costs, and thus we can trace out the distribution of search costs from the observed responsiveness of choice probabilities to the hidden attribute.

A final reason to estimate a full structural search model is to impose plausible parametric restrictions on the data necessary for estimation in finite samples. Our identification proof shows that, in principle, these parametric restrictions are unnecessary for identification. Our simulations suggest that in realistic samples, our results can be used to directly recover preferences in models with linear utility and homogeneous preferences over observables. However, estimation of the higher-order derivatives of choice probabilities necessary to nonparametrically identify the distribution of random coefficients may not be possible given the data available. In such cases, a natural approach is to specify a structural search model with random coefficients in order to place some parametric structure on these higher-order derivatives. This requires taking an explicit stand on the underlying search model. Nonetheless, once the model has been estimated and preferences recovered, the results in Section 2.6 can be used to conduct specification tests for the internal consistency of our estimates. If these tests reject, an alternative search model may fit the data better.

5 Experimental Validation

Our identification proof and simulation results show that preferences can be estimated regardless of whether consumers are fully informed, provided the underlying model is a search model satisfying the stated assumptions. Of course, the theorem does not tell us whether those assumptions are likely to be satisfied in practice.

In this section, we test in a lab experiment whether we can recover preferences in a setting where consumers engage in costly search. Unlike in our simulations, the search protocol is unknown to us and not restricted to satisfy the assumptions of our model. We nonetheless show that we are able to correctly recover preferences using our “search-robust” estimation technique.

5.1 Set-up

We selected 1,000 books for sale on Amazon Kindle chosen from a wide variety of genres. For each book, we observe its average rating on the site “Goodreads.com” as well as the average rating from Amazon.com, the number of reviews on Goodreads, and the price of the book for Amazon Kindle.

In our experiment, conducted via Mechanical Turk, each participant made 40 choices from sets of 3 randomly selected books. For all books, participants could see a photo of the cover, the title, author and genre, as well as the Goodreads rating and the number of ratings. Prices were randomized to integers from \$11-\$15 (equally likely). All books were then further discounted by an integer amount from \$0-\$10 (equally likely). All users were given a \$25.00 bank at the start of each choice, from which any costs incurred were deducted. There were a total of 93 participants, yielding 3,720 choices.


The discount is our key variable of interest. For 10 of the 40 choices, users could see all discounts and thus could see the net price of all options at no cost. For 30 of the 40 choices, discounts were hidden and users had to pay a cost to see the discount for any given book.³⁶ The cost per click was constant for each user across the 30 choices, and randomly chosen from {\$0.10, \$0.25, \$0.35, \$0.50}. For the 30 choices with hidden information, users could only choose books after they clicked to reveal the discount and had to choose at least one book. One of the 40 choices made by each user was randomly chosen to be realized, and users received the chosen book as well as any money left over from the original \$25.00.

Figure 1 shows a sample product selection screen from a choice where discounts were hidden. In this case, the user clicked to reveal the discount of the second book and could either choose that book or continue by revealing the discounts for additional books. Note that the user could search books in any order she wished.

The 10 choices where all information is revealed are our benchmark for the “truth”. The goal is then to test whether the relative weight on discounts and prices that we estimate in the cases where discounts are costly to observe matches the relative weight we see when discounts are visible to everyone. Further, because both discounts and prices are in dollar terms, and because they are randomized (and so not signals

³⁶The full information and costly information choice situations were randomly ordered, so that the 10 “full information” choices were intermixed with the costly information choices.

Figure 1: Lab Experiment: Sample Product Selection Screen



of quality), there is a second benchmark: if consumers are rational, the weight on discounts and prices should be equal.

In other words, we will model choices using:

$$U_{ij} = price \cdot \alpha - discount \cdot \beta + rating \cdot \alpha_2 + \epsilon_{ij} \quad (22)$$

Fully informed and rational consumers should have $\alpha = \beta$. Our goal will be to show that we can recover these fully informed preferences using the choices of beneficiaries for whom revealing discounts is costly.

5.2 Estimation Results

Columns 1 and 2 of Table 3 show results from estimating a conditional logit model on consumer choices for the 10 choice situations (per consumer) where all information is revealed (Full Info) and the 30 choice situations where consumers must pay to reveal information (Costly Info). With full information, consumers place equal weight on prices and (negative) discounts, so they pass our test of rationality. In other words, they care only about the final price of the product. By contrast, when discounts are costly to reveal, the coefficient on the discount variable is attenuated. This is because consumers are insensitive to variation in discounts for books whose prices they do not reveal. The ratio of the two coefficients is 0.974 in the full information treatment and 0.682 in the costly information treatment.

Following Section 4.1, we estimate the demand function $s_1(\mathbf{x}, \mathbf{z})$ via Bernstein polynomials. The main result of this procedure is an estimate of $E_j[\frac{\partial^2 s_{i1}}{\partial z_1 \partial z_j} / \frac{\partial^2 s_{i1}}{\partial z_1 \partial x_j}] = \beta / \alpha$ where the expectation averages over all

Table 3: Conditional Logit and Cross-Derivative Estimation Results

Variable	Full Info	Costly Info	Cross-Derivative Ratio
Price	-0.386*** (0.038)	-0.302*** (0.018)	-
Discount (-)	-0.376*** (0.020)	-0.206*** (0.009)	-
Rating	0.591*** (0.190)	0.421*** (0.099)	-
Discount (-) / Price	0.986*** (0.093)	0.683*** (0.044)	1.032*** (0.102)
N	930	2790	2790

Note: tables show estimation results from a conditional logit model estimated on the full information and costly information treatments in columns 1 and 2, and estimation of the cross-derivative ratio based on Bernstein polynomials in column 3. The minus sign indicates that discount multiplied by -1 so that the coefficient on discount should equal that of price with full information. Standard errors on the ratio of the discount and price coefficients are computed using 250 bootstrap draws. *** denotes significance at the 1% level, ** at 5% level, and * at 10%.

goods other than good 1. This estimate is given by 1.032 (0.102). In other words, using data only on choices when information is costly, we successfully recover informed preferences. The confidence interval is sufficiently tight to exclude the logit estimates in the costly information treatment.


As in most real-world settings, visible utility is not observable to the econometrician: while we can see attributes of the goods in question, we do not know how individuals will weigh these attributes, nor do we know their preferences for specific genres or book titles and images. The assumption that consumers search according to the visible utility assumption is substantive and could be violated in numerous ways: users might always reveal discounts for the lowest priced book first or they might search in the order in which books are displayed. Nonetheless, our “robust” estimation approach succeeds in recovering the preferences consumers reveal with full information.

As discussed in Section 2.6, while the visible utility assumption cannot be verified directly, it can be tested indirectly. One such test is to compute bounds on the choice probabilities for each set of observables consistent with the visible utility assumption. We operationalize this test as follows: first, we estimate α , the coefficient on price, using a conditional logit model, but using only choices from beneficiaries who face little variation in discounts (we use the first quartile of the variance of the discount variable within each choice set). Second, given α , we use our estimates of β/α based on the Bernstein approximation to recover β . Third, we assume that ϵ_{ij} are i.i.d. type I extreme value. Given these assumptions, we can compute the upper and lower bounds described in Section 2.6 for each individual via simulation. We sort the data by the lower bound, bin the data into 100 quantiles, and graph in each quantile the mean of the upper and

lower bounds, as well as the choice probabilities estimated via Bernstein polynomials.

Figure 2 shows the results of this exercise. We can see that the bounds in the experimental data have some bite: the range between the lower bound and the upper bound ranges from 15 to 30 percentage points. The estimated choice probabilities in nearly all cases lie within this range. These probabilities thus appear broadly consistent with the visible utility assumption.

Figure 2: Choice Probabilities, Upper and Lower Bounds from Visible Utility Assumption



6 Conclusion

We prove that it is possible to estimate preferences using only data on attributes and choices in cross-sectional data even when consumers must search to acquire information about product attributes. This result holds in a broad class of search models: those in which consumers search according to the “visible utility” assumption and in which the searched attribute is visible to the econometrician even if unknown to consumers.

The functions of choice probabilities which identify preferences in our model are “robust” in the sense that they work in both full information and search models. For example, in the case of homogeneous preferences, if the ratio of cross-derivatives does not agree with the ratio of first derivatives, one can reject the full information model as misspecified even if the assumptions we make about search do not hold. If those further conditions hold, the ratio of cross-derivatives will recover preferences while the ratio of first derivatives will not. The comparison between the two gives a (powered) test of full-information.

Because our conditions allow preferences to be recovered when consumers are imperfectly informed, our results allow a wide range of inquiries that are impossible using conventional methods. Prior to conducting

an information intervention, choice data can be used to estimate counterfactually how consumers would choose were they fully informed. If preferences are not informed, the preferences consumers would have if they were informed can be used to conduct a more defensible welfare analyses.

Preferences can (sometimes) be identified in structural search models, but such models require making many explicit assumptions about how consumers search. Do consumers consider option value or are they myopic? Do they solve an optimal stopping problem or search until they find a good enough option? Is search sequential or simultaneous? If search costs vary across consumers, what is their statistical distribution? Our approach attempts to avoid these complexities by instead relying on a sufficient condition satisfied in a broad class of search models that can be falsified by the available data via bounds on choice probabilities and overidentification tests.

Of course, our assumptions are not guaranteed to hold and must be assessed on a case by case basis. However, it is important to keep in mind the alternatives. If the choice is between the visible utility assumption and the assumption that consumers are fully informed, the former may be closer to correct even if neither is exactly satisfied. Likewise, if the true search process is unknown and the choice is between the visible utility assumption and many untested structural assumptions about search, visible utility may be more parsimonious and leads to clear, testable predictions.

Our assumptions are sufficient for identification but not necessary. This raises several questions for future research: can our methods be extended to recover preferences when consumers do not search in order of visible utility? Likewise, can preferences be recovered if there is an “ ϵ in the box”: in other words, if search reveals both attributes observable to the econometrician and observable only to consumers?

Increasingly, empirical analyses relax the assumption that consumers make informed choices. Typically, behavioral welfare analysis is done using auxiliary data, restrictions on preferences, or by testing whether consumers choose differently when provided with information. Despite this, absent data to the contrary, the default assumption in most economic analysis remains that consumers make informed choices. Our result suggests this need not be the case. Even with no auxiliary data, researchers can use observed choices both to test whether choices are informed and to recover what preferences would be were consumers more informed. This removes the (often compelling) excuse that while consumers may not be informed, assuming informed choices is the only constructive way to proceed given the data available.

References

- Abaluck, J. and A. Adams (2017). What do consumers consider before they choose? identification from asymmetric demand responses. Technical report, National Bureau of Economic Research.
- Abaluck, J. and J. Gruber (2011). Choice inconsistencies among the elderly: evidence from plan choice in the medicare part d program. *The American economic review* 101(4), 1180–1210.
- Allcott, H. and C. Knittel (2019). Are consumers poorly informed about fuel economy? evidence from two experiments. *American Economic Journal: Economic Policy* 11(1), 1–37.
- Allcott, H., B. Lockwood, and D. Taubinsky (2019). Regressive sin taxes, with an application to the optimal soda tax. Technical report, National Bureau of Economic Research.
- Allcott, H. and D. Taubinsky (2015). Evaluating behaviorally motivated policy: Experimental evidence from the lightbulb market. *American Economic Review* 105(8), 2501–38.
- Armstrong, M. (2017). Ordered consumer search. *Journal of the European Economic Association* 15(5), 989–1024.
- Berry, S., A. Gandhi, and P. Haile (2013). Connected substitutes and invertibility of demand. *Econometrica* 81, 2087–2111.
- Berry, S., J. Levinsohn, and A. Pakes (1995). Automobile prices in market equilibrium. *Econometrica: Journal of the Econometric Society*, 841–890.
- Berry, S. T. and P. A. Haile (2009). Nonparametric identification of multinomial choice demand models with heterogeneous consumers. Technical report, National Bureau of Economic Research.
- Berry, S. T. and P. A. Haile (2014). Identification in differentiated products markets using market level data. *Econometrica* 82(5), 1749–1797.
- Branco, F., M. Sun, and J. M. Villas-Boas (2012). Optimal search for product information. *Management Science* 58(11), 2037–2056.
- Bronnenberg, B. J., J.-P. Dubé, M. Gentzkow, and J. M. Shapiro (2015). Do pharmacists buy bayer? informed shoppers and the brand premium. *The Quarterly Journal of Economics* 130(4), 1669–1726.
- Choi, M., A. Y. Dai, and K. Kim (2018). Consumer search and price competition. *Econometrica* 86(4), 1257–1281.
- Compiani, G. (2019). Market counterfactuals and the specification of multi-product demand: a nonparametric approach. *Working Paper*.
- Conlon, C. T. and J. H. Mortimer (2013). Demand estimation under incomplete product availability. *American Economic Journal: Microeconomics* 5(4), 1–30.
- Ericson, K. M., P. Kircher, J. Spinnewijn, and A. Starc (2015). Inferring risk perceptions and preferences using choice from insurance menus: theory and evidence. Technical report, National Bureau of Economic Research.

- Fox, J. T. and A. Gandhi (2016). Nonparametric identification and estimation of random coefficients in multinomial choice models. *The RAND Journal of Economics* 47(1), 118–139.
- Fox, J. T., K. i. Kim, S. P. Ryan, and P. Bajari (2011). A simple estimator for the distribution of random coefficients. *Quantitative Economics* 2, 381–418.
- Fox, J. T., K. i. Kim, S. P. Ryan, and P. Bajari (2012). The random coefficients logit model is identified. *Journal of Econometrics* 166(2), 204–212.
- Fox, J. T., K. i. Kim, and C. Yang (2016). A simple nonparametric approach to estimating the distribution of random coefficients in structural models. *Journal of Econometrics* 195(2), 236–254.
- Gabaix, X., D. Laibson, G. Moloche, and S. Weinberg (2006). Costly information acquisition: Experimental analysis of a boundedly rational model. *American Economic Review* 96(4), 1043–1068.
- Gaynor, M., C. Propper, and S. Seiler (2016). Free to choose? reform, choice, and consideration sets in the english national health service. *American Economic Review*.
- Goeree, M. S. (2008). Limited information and advertising in the us personal computer industry. *Econometrica* 76(5), 1017–1074.
- Gul, F. and W. Pesendorfer (2008). The case for mindless economics. *The foundations of positive and normative economics: A handbook* 1, 3–42.
- Handel, B. R. and J. T. Kolstad (2015). Health insurance for “humans”: Information frictions, plan choice, and consumer welfare. *The American Economic Review* 105(8), 2449–2500.
- Hastings, J. S. and J. M. Weinstein (2008). Information, school choice, and academic achievement: Evidence from two experiments. *The Quarterly journal of economics* 123(4), 1373–1414.
- Hong, H. and M. Shum (2006). Using price distributions to estimate search costs. *RAND Journal of Economics* 37(2), 257–275.
- Honka, E. and P. Chintagunta (2016). Simultaneous or sequential? search strategies in the us auto insurance industry. *Marketing Science* 36(1), 21–42.
- Honka, E., A. Hortaçsu, and M. A. Vitorino (2017). Advertising, consumer awareness, and choice: Evidence from the us banking industry. *RAND Journal of Economics* 48(3), 611–646.
- Hortaçsu, A. and C. Syverson (2004). Product differentiation, search costs, and competition in the mutual fund industry: a case study of s&p 500 index funds. *Quarterly Journal of Economics* 119(2), 403–456.
- Jindal, P. and A. Aribarg (2018). The value of price beliefs in consumer search.
- Johnson, E. M. and M. M. Rehavi (2016). Physicians treating physicians: Information and incentives in childbirth. *American Economic Journal: Economic Policy* 8(1), 115–41.
- Ke, T. T., Z.-J. M. Shen, and J. M. Villas-Boas (2016). Search for information on multiple products. *Management Science* 62(12), 3576–3603.

- Manski, C. F. (2002). Identification of decision rules in experiments on simple games of proposal and response. *European Economic Review* 46(4-5), 880–891.
- Manzini, P. and M. Mariotti (2014). Stochastic choice and consideration sets. *Econometrica* 82(3), 1153–1176.
- Mehta, N., S. Rajiv, and K. Srinivasan (2003). Price uncertainty and consumer search: A structural model of consideration set formation. *Marketing science* 22(1), 58–84.
- Roberts, J. H. and J. M. Lattin (1991). Development and testing of a model of consideration set composition. *Journal of Marketing Research*, 429–440.
- Schwartz, B., A. Ward, J. Monterosso, S. Lyubomirsky, K. White, and D. R. Lehman (2002). Maximizing versus satisficing: Happiness is a matter of choice. *Journal of personality and social psychology* 83(5), 1178.
- Stahl, D. (1989). Oligopolistic pricing with sequential consumer search. *American Economic Review* 79(4), 700–712.
- Ursu, R. M. (2018). The power of rankings: Quantifying the effect of rankings on online consumer search and purchase decisions. *Marketing Science* 37(4), 530–552.
- Weitzman, M. L. (1979). Optimal search for the best alternative. *Econometrica: Journal of the Econometric Society*, 641–654.
- Woodward, S. E. and R. E. Hall (2012). Diagnosing consumer confusion and sub-optimal shopping effort: Theory and mortgage-market evidence. *American Economic Review* 102(7), 3249–76.

Appendix A: Additional Proofs

In this appendix, we collect the proofs not included in the main text. Throughout, we let $\mathcal{J} \equiv \{1, \dots, J\}$ and use the notational convention $\frac{\partial f}{\partial x^{\bar{n}}}(x) = f(x) \forall x$ for any function f .

A.1 Proof of Theorem 1

We let $2 = j^*$ and $(\mathbf{0}, \mathbf{0}) = (\mathbf{x}^*, \mathbf{z}^*)$ for notational convenience (the proof is unchanged for other values of j^* and $(\mathbf{x}^*, \mathbf{z}^*)$).

We show that the derivatives of the share functions identify the derivatives of v at a point and thus the entire function up to a constant. Fix any (x, z) in the domain of v and consider a Taylor expansion of v around the point $(0, 0)$:

$$v(x, z) = v_0 + \frac{\partial v_0}{\partial x}x + \frac{\partial v_0}{\partial z}z + \dots + \frac{1}{n!} \frac{\partial^n v_0}{\partial z^{\bar{n}} \partial x^{n-\bar{n}}} z^{\bar{n}} x^{n-\bar{n}} + \dots \quad (23)$$

where $v_0 \equiv v(0, 0)$ and $\frac{\partial^n v_0}{\partial z^{\bar{n}} \partial x^{n-\bar{n}}} \equiv \frac{\partial^n v}{\partial z^{\bar{n}} \partial x^{n-\bar{n}}}(0, 0)$ for all $n \geq 1, \bar{n} \leq n$. To recover $v(x, z)$, it is sufficient to recover all derivatives $\frac{\partial^n v_0}{\partial z^{\bar{n}} \partial x^{n-\bar{n}}}$. First, we map this into the notation of Assumption 1 as follows:

$$U_{ij} = \underbrace{v(x_j, 0) + \epsilon_{ij}}_{a_i(x_j)} + \underbrace{v(x_j, z_j) - v(x_j, 0)}_{b(x_j, z_j)}$$

Note that we can directly recover $v(x, 0)$ by the following argument. When $z = 0$, $b(x, 0) = 0$ which means that consumers who search the highest visible utility good (which is guaranteed by Assumption 2(i)) maximize utility and we can identify $v(x, 0)$, a function of x only, using standard techniques. By differentiating $v(x, 0)$ and evaluating at $x = 0$, we can recover all of the terms that contain no z 's, i.e. $\frac{\partial^n v}{\partial x^n}(0, 0)$.

To recover derivatives of v with respect to z , we will use Lemma 1. Specifically, we will take $x_j \in [\bar{x} - \eta, \bar{x} + \eta]$ for all j , where \bar{x} and η are defined in Lemma 1, and use the fact that $\frac{\partial s_1}{\partial z_1}$ can be written as a function of terms which only depend on x_2 and z_2 via U_2 . To formalize this, we let $\mathcal{J}_1 \equiv \{2, \dots, J\}$, $v_j \equiv v(x_j, z_j)$ for all j , and $\mathbf{v} \equiv (v_1, \dots, v_J)$. Similarly, we let $v_j^0 = v(x_j, 0)$ and $\mathbf{v}^0 = (v_1^0, \dots, v_J^0)$. Then, by (3) we can write for all (\mathbf{x}, \mathbf{z}) with $z_1 \geq z_j$ for all j :

$$\begin{aligned} s_1 &= P(U_1 \geq U_k \forall k) - \sum_{\mathcal{S} \subset \mathcal{J}_1, \mathcal{S} \neq \emptyset} P(\{U_1 \geq U_k \forall k\} \cap \{VU_j \geq VU_1 \text{ for at least one } j \in \mathcal{S}\} \\ &\quad \cap \{g_1(x_1, U_j) \leq 0 \text{ for all } j \in \mathcal{S}\} \cap \{g_1(x_1, U_j) \geq 0 \text{ for all } j \in \mathcal{J}_1 \setminus \mathcal{S}\}) \\ &\equiv P_4(\mathbf{v}) - \sum_{\mathcal{S} \subset \mathcal{J}_1, \mathcal{S} \neq \emptyset} P_5^{\mathcal{S}}(\mathbf{v}, \mathbf{v}^0, x_1) \end{aligned} \quad (24)$$

Further, for every $\mathcal{S} \subset \mathcal{J}_1, \mathcal{S} \neq \emptyset$, we have

$$\begin{aligned}
P_5^{\mathcal{S}} &= P(\{U_1 \geq U_k \forall k\} \cap \{g_1(x_1, U_j) \leq 0 \text{ for all } j \in \mathcal{S}\} \cap \{g_1(x_1, U_j) \geq 0 \text{ for all } j \in \mathcal{J}_1 \setminus \mathcal{S}\}) - \\
&\quad P(\{U_1 \geq U_k \forall k\} \cap \{VU_1 \geq VU_j \text{ for all } j \in \mathcal{S}\} \\
&\quad \cap \{g_1(x_1, U_j) \leq 0 \text{ for all } j \in \mathcal{S}\} \cap \{g_1(x_1, U_j) \geq 0 \text{ for all } j \in \mathcal{J}_1 \setminus \mathcal{S}\}) \\
&= P(\{U_1 \geq U_k \forall k\} \cap \{g_1(x_1, U_j) \leq 0 \text{ for all } j \in \mathcal{S}\} \cap \{g_1(x_1, U_j) \geq 0 \text{ for all } j \in \mathcal{J}_1 \setminus \mathcal{S}\}) - \\
&\quad P(\{VU_1 \geq VU_j \text{ for all } j \in \mathcal{S}\} \cap \{g_1(x_1, U_j) \leq 0 \text{ for all } j \in \mathcal{S}\} \cap \{g_1(x_1, U_j) \geq 0 \text{ for all } j \in \mathcal{J}_1 \setminus \mathcal{S}\}) \\
&\equiv P_{5,1}^{\mathcal{S}}(\mathbf{v}, x_1) - P_{5,2}^{\mathcal{S}}(\mathbf{v}_{-1}, \mathbf{v}^0, x_1)
\end{aligned}$$

where $\mathbf{v}_{-1} \equiv (v_2, \dots, v_J)$. The first equality follows from basic set algebra while the second follows from the fact that for all $j \in \mathcal{S}$ and all $k \in \mathcal{J}_1 \setminus \mathcal{S}$, (i) $VU_1 \geq VU_j$ implies $U_1 \geq U_j$ since $z_1 \geq z_j$ for all $j \in \mathcal{J}_1$; and (ii) $g_1(x_1, U_k) \geq 0 \geq g_1(x_1, U_j)$ implies $U_k \leq U_j$, which (together with the implication in (i)) implies $U_1 \geq U_k$. Thus the event $U_1 \geq U_k \forall k \in \mathcal{J}_1$ is implied by the other events inside the probability and can be dropped.

Note that $P_{5,2}^{\mathcal{S}}$ does not depend on z_1 . Thus, omitting the function arguments, we have

$$\frac{\partial s_1}{\partial z_1} = \frac{\partial P_4}{\partial v_1} \frac{\partial v_1}{\partial z_1} - \sum_{\mathcal{S} \subset \mathcal{J}_1, \mathcal{S} \neq \emptyset} \frac{\partial P_{5,1}^{\mathcal{S}}}{\partial v_1} \frac{\partial v_1}{\partial z_1} \quad (25)$$

Differentiating again with respect to z_2 gives:

$$\frac{\partial^2 s_1}{\partial z_1 \partial z_2} = \frac{\partial^2 P_4}{\partial v_1 \partial v_2} \frac{\partial v_1}{\partial z_1} \frac{\partial v_2}{\partial z_2} - \sum_{\mathcal{S} \subset \mathcal{J}_1, \mathcal{S} \neq \emptyset} \frac{\partial^2 P_{5,1}^{\mathcal{S}}}{\partial v_1 \partial v_2} \frac{\partial v_1}{\partial z_1} \frac{\partial v_2}{\partial z_2} \quad (26)$$

Differentiating equation (25) with respect to x_2 gives:

$$\frac{\partial^2 s_1}{\partial z_1 \partial x_2} = \frac{\partial^2 P_4}{\partial v_1 \partial v_2} \frac{\partial v_1}{\partial z_1} \frac{\partial v_2}{\partial x_2} - \sum_{\mathcal{S} \subset \mathcal{J}_1, \mathcal{S} \neq \emptyset} \frac{\partial^2 P_{5,1}^{\mathcal{S}}}{\partial v_1 \partial v_2} \frac{\partial v_1}{\partial z_1} \frac{\partial v_2}{\partial x_2} \quad (27)$$

Combining (26) and (27), we obtain

$$\frac{\partial^2 s_1}{\partial z_1 \partial z_2} / \frac{\partial^2 s_1}{\partial z_1 \partial x_2} = \frac{\frac{\partial v_2}{\partial z_2}}{\frac{\partial v_2}{\partial x_2}} \quad (28)$$

Since this equation holds for all (\mathbf{x}, \mathbf{z}) such that $\frac{\partial^2 s_1}{\partial z_1 \partial x_2} \neq 0$ and we already showed that we can recover $\frac{\partial v}{\partial x}(0, 0)$, we can also recover $\frac{\partial v}{\partial z}(0, 0)$.

Next, note that, fixing $z_k = 0$ for all $k = 1, \dots, J$ and $x_j = 0$ for all $j \neq 2$ in (25), we can write

$$\frac{\partial s_1}{\partial z_1} = k(l(x_2)) \quad (29)$$

where

$$k(v_2) : v_2 \mapsto \frac{\partial v}{\partial z_1}(\mathbf{0}) \left[\frac{\partial P_4}{\partial v_1}(v(\mathbf{0}), v_2, v(\mathbf{0}), \dots, v(\mathbf{0})) - \sum_{S \subset \mathcal{I}_1, S \neq \emptyset} \frac{\partial P_{5,1}^S}{\partial v_1}(v(\mathbf{0}), v_2, v(\mathbf{0}), \dots, v(\mathbf{0}), 0) \right]$$

and $l(x_2) : x_2 \mapsto v(x_2, 0)$. So by the chain rule we have that, for $n > 1$, $\frac{\partial^n s_1}{\partial z_1 \partial x_2^{n-1}}$ is a linear function of the $(n-1)$ -th derivative of k with slope depending on the first derivative of l and intercept depending on derivatives of l and derivatives of k of order strictly less than $n-1$. Further, by the above, all derivatives of l are known. Thus, we have a system of equations that can be uniquely solved for the derivatives of k by recursion.³⁷

Next, we differentiate $\frac{\partial s_1}{\partial z_1}$ once with respect to z_2 and $n-2$ times with respect to x_2 . Similar to the above, we can write

$$\frac{\partial s_1}{\partial z_1} = k(v(x_2, z_2)) \quad (30)$$

where now note that z_2 is no longer fixed at 0. Again by the chain rule we have that, for $n \geq 3$, $\frac{\partial^n s_1}{\partial z_1 \partial z_2 \partial x_2^{n-2}}$ evaluated at $(\mathbf{0}, \mathbf{0})$ is a linear function of $\frac{\partial^{n-1} v}{\partial z_2 \partial x_2^{n-2}}(0, 0)$ with slope coefficient depending on $k'(v(0, 0))$ and intercept depending on lower-order derivatives of v as well as derivatives of k .³⁸ Because all derivatives of k are known by the argument above, we can iteratively solve for $\frac{\partial^{n-1} v}{\partial z_2 \partial x_2^{n-2}}(0, 0)$ for all $n \geq 3$.

The remaining terms in the Taylor expansion can be recovered by an analogous argument. Specifically, for any $n \geq 3, m \geq 2$, by differentiating (30) m times wrt z_2 and again $n-m-1$ times wrt x_2 , one can write $\frac{\partial^n s_1}{\partial z_1 \partial z_2^m \partial x_2^{n-m-1}}$ as a linear function of $\frac{\partial^{n-1} v}{\partial z_2^m \partial x_2^{n-m-1}}(0, 0)$ with known, nonzero slope and known intercept. This system can then be solved iteratively for $\frac{\partial^{n-1} v}{\partial z_2^m \partial x_2^{n-m-1}}(0, 0)$ for all $n > m \geq 2$.

Therefore, we know all the coefficients in the Taylor-expansion of $v(x, z)$ except the constant $v(0, 0)$, i.e. we can recover $v(x, z)$ up to a constant.

A.2 Proof of Theorem 3

Note that the analog of equation (24) in the random coefficients model holds for any given value of α and $\beta > 0$, so that we can write:

$$s_1 = \sum_{k_\alpha=1}^{K_\alpha} \sum_{k_\beta=1}^{K_\beta} \left[P_6(\mathbf{v}(k_\alpha, k_\beta)) - \sum_{S \subset \mathcal{I}_1, S \neq \emptyset} P_7^S(\mathbf{v}(k_\alpha, k_\beta), \mathbf{v}^0(k_\alpha), x_1) \right] \tilde{\pi}_{k_\alpha, k_\beta} \quad (31)$$

³⁷Here, we use the fact that, by assumption, the first derivative of l is nonzero.

³⁸Note that $\frac{\partial^2 s_1}{\partial z_1 \partial x_2}(\mathbf{0}, \mathbf{0}) = k'(v(0, 0))l'(0)$, so we have $k'(v(0, 0)) \neq 0$ by assumption.

where $v_j(k_\alpha, k_\beta) \equiv x_j \alpha_{k_\alpha} + z_j \beta_{k_\beta}$, $\mathbf{v}(k_\alpha, k_\beta) \equiv (v_1(k_\alpha, k_\beta), \dots, v_J(k_\alpha, k_\beta))$, and similarly $v_j^0(k_\alpha) \equiv x_j \alpha_{k_\alpha}$, $\mathbf{v}^0(k_\alpha) \equiv (v_1^0(k_\alpha), \dots, v_J^0(k_\alpha))$. Differentiating (31), we have, for all integers $n \geq \tilde{n} \geq 0$:

$$\frac{\partial^{1+n} s_1}{\partial z_1 \partial z_2^{\tilde{n}} \partial x_2^{n-\tilde{n}}} = \sum_{k_\alpha=1}^{K_\alpha} \sum_{k_\beta=1}^{K_\beta} \frac{\partial^{1+n} \left[P_6(\mathbf{v}(k_\alpha, k_\beta)) - \sum_{S \subset \mathcal{J}_1, S \neq \emptyset} P_7^S(\mathbf{v}(k_\alpha, k_\beta), \mathbf{v}^0(k_\alpha), x_1; \alpha_{k_\alpha}, \beta_{k_\beta}) \right]}{\partial z_1 \partial z_2^{\tilde{n}} \partial x_2^{n-\tilde{n}}} \tilde{\pi}_{k_\alpha, k_\beta} \quad (32)$$

Next, we evaluate (32) at a value of (\mathbf{x}, \mathbf{z}) such that $x_j = \bar{x}$ and $z_j = \bar{z}_j$ for all j . Note that at such (\mathbf{x}, \mathbf{z}) , P_6 no longer depends on k_α, k_β . Using this and letting $\bar{x} = \bar{z} = 0$ for notational convenience (the proof is unchanged for any other values), we may re-write (32) as

$$\begin{aligned} \frac{\partial^{1+n} s_1}{\partial z_1 \partial z_2^{\tilde{n}} \partial x_2^{n-\tilde{n}}} &= \frac{\partial^{1+n} P_6}{\partial v_1 \partial v_2^n}(\mathbf{0}) \sum_{k_\alpha=1}^{K_\alpha} \sum_{k_\beta=1}^{K_\beta} \alpha_{k_\alpha}^{n-\tilde{n}} \beta_{k_\beta}^{\tilde{n}+1} \tilde{\pi}_{k_\alpha, k_\beta} - \sum_{k_\alpha=1}^{K_\alpha} \sum_{k_\beta=1}^{K_\beta} \alpha_{k_\alpha}^{n-\tilde{n}} \beta_{k_\beta}^{\tilde{n}+1} \tilde{\pi}_{k_\alpha, k_\beta} \sum_{S \subset \mathcal{J}_1, S \neq \emptyset} \frac{\partial^{1+n} P_7^S(\mathbf{0}, \mathbf{0}, 0; \alpha_{k_\alpha}, \beta_{k_\beta})}{\partial v_1 \partial v_2^n} \\ &\equiv \sum_{k=1}^K [a_{k,n,\tilde{n}} + b_{k,n,\tilde{n}} f_{k,n}] \pi_k \end{aligned} \quad (33)$$

where we let $K \equiv K_\alpha K_\beta$ and let k represent the double index (k_α, k_β) , $a_{k,n,\tilde{n}} \equiv \frac{\partial^{1+n} P_6}{\partial v_1 \partial v_2^n}(\mathbf{0}) \alpha_{k_\alpha}^{n-\tilde{n}} \beta_{k_\beta}^{\tilde{n}+1}$, $b_{k,n,\tilde{n}} \equiv \alpha_{k_\alpha}^{n-\tilde{n}} \beta_{k_\beta}^{\tilde{n}+1}$ are known scalars, and $f_{k,n} \equiv \frac{\partial^{1+n} P_7^S(\mathbf{0}, \mathbf{0}, 0; \alpha_{k_\alpha}, \beta_{k_\beta})}{\partial v_1 \partial v_2^n}$ and $\pi_k \equiv \tilde{\pi}_{k_\alpha, k_\beta}$ are unknown scalars.

Setting $n = K - 1$ and stacking the equations corresponding to $\tilde{n} = 0, \dots, K - 1$, we get

$$q = A\pi + B(f * \pi)$$

where q is a known column K -vector, A, B are known K -by- K matrices, and $f * \pi$ denotes the column vector given by the element-by-element product of $f = (f_{1,K-1}, \dots, f_{K,K-1})'$ and $\pi \equiv (\pi_1, \dots, \pi_K)'$. We re-write this system of equations in a way that highlights which objects depend on $\mathbf{z} \equiv (z_1, \dots, z_J)$ as follows

$$q(\mathbf{z} = \mathbf{0}) = A\pi + B(f(\mathbf{z} = \mathbf{0}) * \pi)$$

Note that A depends on z only through $z_1 - z_J$ (i.e. it exhibits a lack of nominal illusion property) and we leave that dependence implicit. Now consider increasing z_j by Δz for all j relative to the baseline $\mathbf{z} = \mathbf{0}$. Then we can write

$$q(\mathbf{z} = \Delta \mathbf{z}) = A\pi + B(f(\mathbf{z} = \Delta \mathbf{z}) * \pi)$$

Combining the last two systems, we get

$$q(\mathbf{z} = \Delta \mathbf{z}) - q(\mathbf{z} = \mathbf{0}) = B[(f(\mathbf{z} = \Delta \mathbf{z}) - f(\mathbf{z} = \mathbf{0})) * \pi]$$

If B is full rank,³⁹ we obtain identification of $(f(\mathbf{z} = \Delta \mathbf{z}) - f(\mathbf{z} = \mathbf{0})) * \pi$. Also, note that, for all k ,

³⁹Note that this condition is immediately verifiable since the points in the support of α and β are chosen by the researcher.

$\lim_{\Delta z \rightarrow 0} \frac{f_{k,K-1}(\mathbf{z}=\Delta\mathbf{z}) - f_{k,K-1}(\mathbf{z}=\mathbf{0})}{\Delta z}$ is the directional derivative of $f_{k,K-1}$ in the direction $\mathbf{1} = (1, \dots, 1)$ and thus is equal to $\sum_{j=1}^J \frac{\partial f_{k,K-1}}{\partial z_j}(\mathbf{z}=\mathbf{0})$ if $f_{k,K-1}$ is differentiable. Therefore, we can write

$$\lim_{\Delta z \rightarrow 0} \frac{q(\mathbf{z}=\Delta\mathbf{z}) - q(\mathbf{z}=\mathbf{0})}{\Delta z} = B \left[\left(\sum_{j=1}^J \frac{\partial f}{\partial z_j}(\mathbf{z}=\mathbf{0}) \right) * \pi \right]$$

Because the lhs is identified, this shows that we can identify $\left(\sum_{j=1}^J \frac{\partial f}{\partial z_j}(\mathbf{z}=\mathbf{0}) \right) * \pi$.

Next, for $j \in \mathcal{J}$, we can take another derivative wrt z_j in (33) and write

$$q_{(j)}(\mathbf{z}=\mathbf{0}) = A_{(j)}\pi + B_{(j)} \left(\frac{\partial f}{\partial z_j}(\mathbf{z}=\mathbf{0}) * \pi \right) \quad (34)$$

for known K -by- K matrices $A_{(j)}, B_{(j)}$ and a known column K -vector $q_{(j)}(\mathbf{z}=\mathbf{0})$. Note that $B_{(j)} = B$ for all $j \in \mathcal{J}$ and so we can write

$$\sum_{j=1}^J q_{(j)}(\mathbf{z}=\mathbf{0}) = \left(\sum_{j=1}^J A_{(j)} \right) \pi + B \left[\left(\sum_{j=1}^J \frac{\partial f}{\partial z_j}(\mathbf{z}=\mathbf{0}) \right) * \pi \right] \quad (35)$$

From above, $\left(\sum_{j=1}^J \frac{\partial f}{\partial z_j}(\mathbf{z}=\mathbf{0}) \right) * \pi$ is identified. This implies that π is identified if the matrix $\sum_{j=1}^J A_{(j)}$ is invertible.⁴⁰

A.3 Proof of Theorem 4 (when Observables Impact Search but not Utility)

Note that if good 1 doesn't maximize utility in R , good 1 will never be chosen. If some other good in R has higher utility, it has higher visible utility and will be searched before good 1. Then, by Lemma 4, we can write for all (\mathbf{x}, \mathbf{z}) with $z_1 \geq z_j$ and all j :

$$\begin{aligned} s_1 &= P(U_1 \geq U_k \forall k \in R) - \sum_{\mathcal{S} \subset \mathcal{J}_1, \mathcal{S} \neq \emptyset} P(\{U_1 \geq U_k \forall k \in R\} \cap \{m(VU_j, r_j) \geq m(VU_1, r_1) \text{ for at least one } j \in \mathcal{S}\} \\ &\quad \cap \{g_1(x_1, U_j) \leq 0 \text{ for all } j \in \mathcal{S}\} \cap \{g_1(x_1, U_j) \geq 0 \text{ for all } j \in \mathcal{J}_1 \setminus \mathcal{S}\}) \\ &\equiv P_{4new}(\mathbf{v}) - \sum_{\mathcal{S} \subset \mathcal{J}_1, \mathcal{S} \neq \emptyset} P_{5new}^{\mathcal{S}}(\mathbf{v}, \mathbf{v}^0, x_1, \mathbf{r}) \end{aligned} \quad (36)$$

⁴⁰Again, this condition is immediately verifiable given the support points for α and β chosen by the researcher.

where \mathbf{v} and \mathbf{v}^0 are defined as in the proof of Theorem 1 (Appendix A.1) and $\mathbf{r} \equiv (r_1, \dots, r_J)$. Further, for every $\mathcal{S} \subset \mathcal{J}_1, \mathcal{S} \neq \emptyset$, we have

$$\begin{aligned}
P_{5new}^{\mathcal{S}} &= P(\{U_1 \geq U_k \ \forall k \in R\} \cap \{g_1(x_1, U_j) \leq 0 \text{ for all } j \in \mathcal{S}\} \cap \{g_1(x_1, U_j) \geq 0 \text{ for all } j \in \mathcal{J}_1 \setminus \mathcal{S}\}) - \\
&\quad P(\{U_1 \geq U_k \ \forall k \in R\} \cap \{m(VU_1, r_1) \geq m(VU_j, r_j) \text{ for all } j \in \mathcal{S}\} \\
&\quad \cap \{g_1(x_1, U_j) \leq 0 \text{ for all } j \in \mathcal{S}\} \cap \{g_1(x_1, U_j) \geq 0 \text{ for all } j \in \mathcal{J}_1 \setminus \mathcal{S}\}) \\
&= P(\{U_1 \geq U_k \ \forall k \in R\} \cap \{g_1(x_1, U_j) \leq 0 \text{ for all } j \in \mathcal{S}\} \cap \{g_1(x_1, U_j) \geq 0 \text{ for all } j \in \mathcal{J}_1 \setminus \mathcal{S}\}) - \\
&\quad P(\{m(VU_1, r_1) \geq m(VU_j, r_j) \text{ for all } j \in \mathcal{S}\} \cap \{g_1(x_1, U_j) \leq 0 \text{ for all } j \in \mathcal{S}\} \cap \{g_1(x_1, U_j) \geq 0 \text{ for all } j \in \mathcal{J}_1 \setminus \mathcal{S}\}) \\
&\equiv P_{5new,1}^{\mathcal{S}}(\mathbf{v}, x_1, \mathbf{r}) - P_{5new,2}^{\mathcal{S}}(\mathbf{v}_{-1}, \mathbf{v}^0, x_1, \mathbf{r})
\end{aligned}$$

where $\mathbf{v}_{-1} \equiv (v_2, \dots, v_J)$. This argument exactly parallels the argument in Appendix A.1, except now we have additionally used the fact that $U_1 \geq U_j$ for all $j \in R$, since (i) if $j \in \mathcal{S}$, then $m(VU_1, r_1) \geq m(VU_j, r_j)$ implies $VU_1 \geq VU_j$, which in turn implies $U_1 \geq U_j$; (ii) if $j \notin \mathcal{S}$, then $g_1(x_1, U_j) \geq 0 \geq g_1(x_1, U_k)$ for all $k \in \mathcal{S}$ implies $U_j \leq U_k$. Note that $P_{5new,2}^{\mathcal{S}}$ does not depend on z_1 and $P_{5new,1}^{\mathcal{S}}(\mathbf{v}, x_1, \mathbf{r})$ only depends on x_j and z_j via v_j for $j \neq 1$, so the remainder of the argument in Appendix A.1 applies.

A.4 Identification of a model where consumers form expectations on z_j based on x_j

A.4.1 Identification of $\alpha + \beta\gamma_1$

Note that if $\tilde{z}_j = 0$ for all j , then consumers always maximize utility. Thus, seeing how choice probabilities change with x conditional on $\tilde{z}_j = 0$ for all j should help identify $\alpha + \beta\gamma_1$. Because the event $\tilde{z}_j = 0$ involves x_j , we need to differentiate choice probabilities with respect to x_j on the envelope satisfying the condition $\tilde{z}_j = 0$ for all x_j . Formally, fix any $j \in \mathcal{J}$ and choose (x_k, z_k) so that $z_k = \gamma_0 + \gamma_1 x_k$ (which implies $\tilde{z}_k = 0$) for all $k \neq j$. For every $\delta > 0$, let $\epsilon(\delta) \equiv \gamma_0 + (x_j + \delta)\gamma_1 - z_j$, so that $z_j + \epsilon(\delta) - \mathbb{E}(z_j | x_j + \delta) = 0$. Note that $\epsilon(\delta)$ is known to the econometrician. Denoting by $\mathbf{x}_{-j} = (x_k)_{k \neq j}$ and similarly for \mathbf{z}_{-j} , we have

$$\frac{s_j(x_j + \delta, \mathbf{x}_{-j}, z_j + \epsilon(\delta), \mathbf{z}_{-j}) - s(\mathbf{x}, \mathbf{z})}{\delta} \tag{37}$$

$$\begin{aligned}
&= \frac{P((x_j + \delta)(\alpha + \beta\gamma_1) + \epsilon_{ij} \geq x_k(\alpha + \beta\gamma_1) + \epsilon_{ik} \ \forall k) - P(x_j(\alpha + \beta\gamma_1) + \epsilon_{ij} \geq x_k(\alpha + \beta\gamma_1) + \epsilon_{ik} \ \forall k)}{\delta} \\
&\xrightarrow{\delta \rightarrow 0} \frac{\partial}{\partial x_j} P(x_j(\alpha + \beta\gamma_1) + \epsilon_{ij} \geq x_k(\alpha + \beta\gamma_1) + \epsilon_{ik} \ \forall k) \tag{38}
\end{aligned}$$

where the first equality follows from the fact that all consumers always maximize utility at the chosen values of (\mathbf{x}, \mathbf{z}) . Now note that (i) the expression in (37) is known for all $\delta > 0$; and (ii) at $\mathbf{x} = \mathbf{0}$, the term in (38) factors into $(\alpha + \beta\gamma_1)$ and a term that only depends the distribution of ϵ_i . Thus, evaluating the last display at $\mathbf{x} = \mathbf{0}$ yields identification of $(\alpha + \beta\gamma_1)$ under a parametric assumption on ϵ_i .

A.5 Proof of Lemma 5 which shows identification with endogeneity

First, assume that p_j is part of the visible utility of good j and Fix (δ_j, p_j, z_j) for all j . To prove claim (i), we show that an increase in δ_j can only induce a consumer to switch from not choosing j to choosing j but

never vice versa, and that a positive mass of consumers will switch to choosing j . To see this, consider the case where consumer i initially searches j , which happens if and only if $g_{ij}(\delta_j, p_j, U_{ik}) \geq 0$ for all k such that $VU_{ik} \geq VU_{ij}$. Let $\Delta \geq 0$ be the change in δ_j . Since g_{ij} is increasing in its first argument, we have $g_{ij}(\delta_j + \Delta, p_j, U_{ik}) \geq 0$ for all k such that $VU_{ik} \geq VU_{ij} + \Delta$ and thus i will still search j . Moreover, since g_{ij} is decreasing in its last argument, if $g_{ik}(\delta_k, p_k, U_{ij}) \leq 0$ for some k such that $VU_{ik} \leq VU_{ij}$ (i.e. if k is initially not searched), then $g_{ik}(\delta_k, p_k, U_{ij} + \Delta) \leq 0$ (i.e. k is also not searched after the change in δ_j), which means that the set of goods searched by i never becomes larger. Next, note that if $U_{ij} \geq U_{ik}$ for all k in the set of searched goods \mathcal{G}_i , then $U_{ij} + \Delta \geq U_{ik}$ for all $k \in \mathcal{G}_i$. Further, since ϵ_i is supported on all of \mathbb{R}^J , there is a positive mass of consumers for which $U_{ik} \geq U_{ij}$ for some $k \in \mathcal{G}_i$, but $U_{ij} + \Delta \geq U_{ik}$ for all $k \in \mathcal{G}_i$. An analogous argument proves claim (ii).

Since the argument above does not rely on the fact that p_j is part of the visible utility of good j , the conclusion also holds for the case in which p_j is only uncovered upon searching good j .

A.6 Derivation of Flexible Logit Weights and Choice Probabilities

In this section, we derive the relevant derivatives of choice probabilities for the flexible logit model described in the text. In this model:

$$v_1 = \tilde{v}(x_1, z_1) + b_1 z_1 + \sum_{k \neq 1} (\gamma_k w_{z1k} z_k + \gamma_{2k} w_{x1k} x_k + w_{z2k} \delta_k z_k z_1 + w_{x2k} \delta_{2k} x_k z_1) \quad (39)$$

and $v_k = \tilde{v}(x_k, z_k)$ for $k \neq 1$ where b_1 , γ_k , γ_{2k} , δ_k and δ_{2k} are coefficients to be estimated which allow greater flexibility in how derivatives with respect to z_1 vary with attributes of rival goods. The weights w_{x1k} , w_{z1k} , w_{x2k} and w_{z2k} are chosen so that, given the logit functional form, $\frac{\partial^2 s_1}{\partial z_1 \partial z_j} / \frac{\partial^2 s_1}{\partial z_1 \partial x_j}$ can be constant across goods as our structural model implies when these weights are regarded as constant in derivatives. With these weights, we have the following derivatives (where we use the notation \tilde{v}_j to refer to the function

\tilde{v} evaluated at (x_j, z_j) :

$$\begin{aligned}
\frac{\partial v_1}{\partial z_1} &= \frac{\partial \tilde{v}_1}{\partial z} + b_1 + \sum_{k \neq 1} (w_{z2k} \delta_k z_k + w_{x2k} \delta_{2k} x_k) \\
\frac{\partial s_1}{\partial x_1} &= \frac{\partial s_1}{\partial v_1} \frac{\partial v_1}{\partial x_1} = \frac{\partial \tilde{v}_1}{\partial x} s_1 (1 - s_1) \\
\frac{\partial s_1}{\partial z_1} &= \frac{\partial s_1}{\partial v_1} \frac{\partial v_1}{\partial z_1} = \frac{\partial v_1}{\partial z_1} s_1 (1 - s_1) \\
\frac{\partial s_1}{\partial x_{j'}} &= \frac{\partial s_1}{\partial v_{ij'}} \frac{\partial v_{ij'}}{\partial x_{j'}} + \frac{\partial s_1}{\partial v_{i1}} \frac{\partial v_{i1}}{\partial x_{j'}} = -\frac{\partial \tilde{v}_{j'}}{\partial x} s_1 s_{j'} + [w_{x1j'} \gamma_{2j'} + w_{x2j'} \delta_{2j'} z_1] s_1 (1 - s_1) \\
\frac{\partial s_1}{\partial z_{j'}} &= \frac{\partial s_1}{\partial v_{ij'}} \frac{\partial v_{ij'}}{\partial z_{j'}} + \frac{\partial s_1}{\partial v_{i1}} \frac{\partial v_{i1}}{\partial z_{j'}} = -\frac{\partial \tilde{v}_{j'}}{\partial z} s_1 s_{j'} + [w_{z1j'} \gamma_{j'} + w_{z2j'} \delta_{j'} z_1] s_1 (1 - s_1) \\
\frac{\partial^2 s_1}{\partial z_1 \partial x_{j'}} &= \frac{\partial^2 s_1}{\partial v_1 \partial x_{j'}} \frac{\partial v_1}{\partial z_1} + \frac{\partial s_1}{\partial v_1} \frac{\partial^2 v_1}{\partial z_1 \partial x_{j'}} \\
&= \frac{\partial v_1}{\partial z_1} (1 - 2s_1) \frac{\partial s_1}{\partial x_{j'}} + s_1 (1 - s_1) w_{x2j'} \delta_{2j'} \\
\frac{\partial^2 s_1}{\partial z_1 \partial z_{j'}} &= \frac{\partial^2 s_1}{\partial v_1 \partial z_{j'}} \frac{\partial v_1}{\partial z_1} + \frac{\partial s_1}{\partial v_1} \frac{\partial^2 v_1}{\partial z_1 \partial z_{j'}} \\
&= \frac{\partial v_1}{\partial z_1} (1 - 2s_1) \frac{\partial s_1}{\partial z_{j'}} + s_1 (1 - s_1) w_{z2j'} \delta_{j'}
\end{aligned} \tag{40}$$

And also:

$$\begin{aligned}
\frac{\partial^2 s_1}{\partial z_1 \partial z_{j'}} / \frac{\partial^2 s_1}{\partial z_1 \partial x_{j'}} &= \frac{\frac{\partial v_1}{\partial z_1} (1 - 2s_1) \frac{\partial s_1}{\partial z_{j'}} + s_1 (1 - s_1) w_{z2j'} \delta_{j'}}{\frac{\partial v_1}{\partial z_1} (1 - 2s_1) \frac{\partial s_1}{\partial x_{j'}} + s_1 (1 - s_1) w_{x2j'} \delta_{2j'}} \\
&= \frac{\frac{\partial v_1}{\partial z_1} (1 - 2s_1) \left(-\frac{\partial \tilde{v}_{j'}}{\partial z} s_1 s_{j'} + [w_{z1j'} \gamma_{j'} + w_{z2j'} \delta_{j'} z_1] s_1 (1 - s_1) \right) + s_1 (1 - s_1) w_{z2j'} \delta_{j'}}{\frac{\partial v_1}{\partial z_1} (1 - 2s_1) \left(-\frac{\partial \tilde{v}_{j'}}{\partial x} s_1 s_{j'} + [w_{x1j'} \gamma_{2j'} + w_{x2j'} \delta_{2j'} z_1] s_1 (1 - s_1) \right) + s_1 (1 - s_1) w_{x2j'} \delta_{2j'}} \tag{41}
\end{aligned}$$

If we define the weights: $w_{x1j'} = w_{z1j'} = \frac{s_{j'}}{1-s_1}$ and $w_{x2j'} = w_{z2j'} = \left[\frac{z_1(1-s_1)}{s_{j'}} + \frac{(1-s_1)}{(\partial v_1/\partial z_1)(1-2s_1)s_{j'}} \right]^{-1} = \frac{(1-2s_1)s_{j'}}{1-s_1} \left(\frac{1}{\partial v_1/\partial z_1} + (1-2s_1)z_1 \right)^{-1} = \frac{(\partial v_1/\partial z_1)(1-2s_1)s_{j'}}{1-s_1} (1 + (1-2s_1)z_1(\partial v_1/\partial z_1))^{-1}$, then:

$$\begin{aligned}
\frac{\partial^2 s_1}{\partial z_1 \partial z_{j'}} / \frac{\partial^2 s_1}{\partial z_1 \partial x_{j'}} &= \frac{\frac{\partial v_1}{\partial z_1} (1 - 2s_1) s_1 s_{j'} \left(-\frac{\partial \tilde{v}_{j'}}{\partial z} + \gamma_{j'} w_{z1j'} \frac{(1-s_1)}{s_{j'}} + \delta_{j'} w_{z2j'} \left[\frac{z_1(1-s_1)}{s_{j'}} + \frac{(1-s_1)}{(\partial v_1/\partial z_1)(1-2s_1)s_{j'}} \right] \right)}{\frac{\partial v_1}{\partial z_1} (1 - 2s_1) s_1 s_{j'} \left(-\frac{\partial \tilde{v}_{j'}}{\partial x} + \gamma_{j'} w_{x1j'} \frac{(1-s_1)}{s_{j'}} + \delta_{j'} w_{x2j'} \left[\frac{z_1(1-s_1)}{s_{j'}} + \frac{(1-s_1)}{(\partial v_1/\partial z_1)(1-2s_1)s_{j'}} \right] \right)} \\
&= \frac{\frac{\partial v_1}{\partial z_1} (1 - 2s_1) s_1 s_{j'} \left(-\frac{\partial \tilde{v}_{j'}}{\partial z} + \gamma_{j'} + \delta_{j'} \right)}{\frac{\partial v_1}{\partial z_1} (1 - 2s_1) s_1 s_{j'} \left(-\frac{\partial \tilde{v}_{j'}}{\partial x} + \gamma_{2j'} + \delta_{2j'} \right)} \\
&= \frac{-\frac{\partial \tilde{v}_{j'}}{\partial z} + \gamma_{j'} + \delta_{j'}}{-\frac{\partial \tilde{v}_{j'}}{\partial x} + \gamma_{2j'} + \delta_{2j'}} \tag{42}
\end{aligned}$$

Thus, we have:

$$\frac{\partial^2 s_1}{\partial z_1 \partial z_{j'}} / \frac{\partial^2 s_1}{\partial z_1 \partial x_{j'}} = \frac{-\frac{\partial \tilde{v}_{j'}}{\partial z} + \gamma_{j'} + \delta_{j'}}{-\frac{\partial \tilde{v}_{j'}}{\partial x} + \gamma_{2j'} + \delta_{2j'}} \quad (43)$$

where $w_{z1j'} = w_{x1j'} = \frac{s_{j'}}{1-s_1}$ and $w_{z2j'} = w_{x2j'} = (\partial v_1 / \partial z_1) \frac{(1-2s_1)s_{j'}}{1-s_1} (1 + (\partial v_1 / \partial z_1)(1-2s_1)z_1)^{-1}$. Given a linear specification of \tilde{v} , $\tilde{v}(x_j, z_j) = x_j a_1 + z_j a_2$, this implies that the above ratio is a constant for each j' .

Estimation of the model with these weights is infeasible since the levels of the choice probabilities s_1 and s_k , as well as the derivatives $\partial v_1 / \partial z_1$ are unknown ex ante and thus we do not know the weights. We estimate the model via a two-step process where s_1 and s_k are estimated using a naive logit model (where utility for each good is a linear function of x_j and z_j), these estimates are used to construct weights, and then the model in equation (21) is estimated treating these weights as constants.⁴¹

To recover estimates of β/α from the flexible logit model, we use the ratio in equation (43). With the linear specification of \tilde{v} , this ratio is given by $\frac{\beta}{\alpha} = \frac{-a_2 + \gamma_{j'} + \delta_{j'}}{-a_1 + \gamma_{2j'} + \delta_{2j'}}$. In cases where the identity of goods is not meaningful (e.g. “good 2” does not refer to the same good across different choice sets and there are no alternative-specific fixed effects), we can further impose $\gamma_k = \gamma$, $\gamma_{2k} = \gamma_2$, $\delta_k = \delta$ and $\delta_{2k} = \delta_2$, which gives a single estimate of $\frac{\beta}{\alpha}$.

A.7 Recovery of Search Costs Given Preferences in the Weitzman Model

Suppose that utility is given by $U_{ij} = x_j \alpha + z_j \beta + \epsilon_{ij}$ and that consumers search sequentially according to the model of Weitzman (1979).

As shown in Armstrong (2017),⁴² the optimal search strategy is for consumers to behave as if they were choosing among options in a static model with utilities given by $\tilde{U}_{ij} = x_j \alpha + \min\{z_j, rv_i\} \beta + \epsilon_{ij}$, where rv_i denotes i 's reservation value in units of z (see Example 1). Thus, dropping i subscripts, ordering goods so that $z_1 \geq z_2 \geq \dots \geq z_J$, and letting

$$E_t \equiv \{\epsilon : \epsilon_k - \epsilon_1 \leq (x_1 - x_k) \alpha, k = 2, \dots, J - t - 1\} \cap \{\epsilon : \epsilon_h - \epsilon_1 \leq (x_1 - x_h) \alpha + (rv - z_h) \beta, h = J - t, \dots, J\}$$

we can write

$$\begin{aligned} s_1 &= P(x_1 \alpha + \min\{z_1, rv\} \beta + \epsilon_1 \geq x_k \alpha + \min\{z_k, rv\} \beta + \epsilon_k \ \forall k) \\ &= P(\epsilon_k - \epsilon_1 \leq (x_1 - x_k) \alpha \ \forall k) P(rv \leq z_J) \\ &\quad + \sum_{t=0}^{J-2} \int P(\{\epsilon \in E_t\} \cap \{z_{J-t} \leq rv \leq z_{J-t-1}\}) dF_{rv}(rv) \\ &\quad + P(\epsilon_k - \epsilon_1 \leq (x_1 - x_k) \alpha + (z_1 - z_k) \beta \ \forall k) P(rv \geq z_1) \end{aligned}$$

⁴¹Since $\partial v_1 / \partial z_1$ is estimated imprecisely from the naive logit, when $1 + (\partial v_1 / \partial z_1)(1 - 2s_1)z_1$ is close to 0 (leading to very large weights), we set $\partial v_1 / \partial z_1 = 0$ when the former term falls below 1 in absolute value.

⁴²See also Choi, Dai, and Kim (2018).

where F_{rv} denotes the cdf of rv and the second equality assumes that search costs (and thus rv) are independent of ϵ . Therefore, we have

$$\frac{\partial s_1}{\partial z_1} = \left[\frac{\partial}{\partial z_1} P(\epsilon_k - \epsilon_1 \leq (x_1 - x_k)\alpha + (z_1 - z_k)\beta \forall k) \right] P(rv \geq z_1) \quad (44)$$

Given identification of (α, β) by the argument in Section 2.1, the first term on the rhs of (44) is identified given parametric assumptions on the distribution of ϵ . Thus, $P\{rv \geq z_1\}$ is identified. Repeating the argument for all z_1 , one can trace out the entire distribution of rv . Since c , the search cost for consumer i , is a known transformation of rv ,⁴³ the distribution of c is also identified.

Equation (44) also lends itself to a different argument that does not require making a parametric assumption on the distribution of ϵ , but instead relies on “at-infinity” variation. Note that the first term on the rhs of (44) is invariant to increasing all z_j ’s by the same amount. Thus, we can write

$$\frac{\frac{\partial s_1}{\partial z_1}(\mathbf{z} + \mathbf{\Delta})}{\frac{\partial s_1}{\partial z_1}(\mathbf{z})} = \frac{P(rv \geq z_1 + \Delta)}{P(rv \geq z_1)} \quad (45)$$

where $\mathbf{\Delta}$ is a J -vector with all elements equal to some Δ . Letting $\Delta \rightarrow -\infty$, the numerator on the rhs of (45) goes to 1, which yields identification of $P(rv \geq z_1)$. Repeating the argument for all z_1 , one can trace out the entire distribution of rv and recover the distribution of c as above.

⁴³This assumes that the prior F_z used by consumers in forming expectations are known to the researcher, as in the case where consumers have rational expectations and F_z coincides with the observed distribution of z across goods and/or markets.