

Aristodemou, Eleni; Rosen, Adam M.

Working Paper

A discrete choice model for partially ordered alternatives

cemmap working paper, No. CWP35/21

Provided in Cooperation with:

Institute for Fiscal Studies (IFS), London

Suggested Citation: Aristodemou, Eleni; Rosen, Adam M. (2021) : A discrete choice model for partially ordered alternatives, cemmap working paper, No. CWP35/21, Centre for Microdata Methods and Practice (cemmap), London,
<https://doi.org/10.47004/wp.cem.2021.3521>

This Version is available at:

<https://hdl.handle.net/10419/246803>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

A discrete choice model for partially ordered alternatives

Eleni Aristodemou
Adam M. Rosen

The Institute for Fiscal Studies
Department of Economics, UCL

cemmap working paper CWP35/21



A Discrete Choice Model for Partially Ordered Alternatives.*

Eleni Aristodemou[†]
University of Cyprus

Adam M. Rosen[‡]
Duke University, CeMMAP and IFS

July 31, 2021

Abstract

In this paper we analyze a discrete choice model for partially ordered alternatives. The alternatives are differentiated along two dimensions, the first an unordered “horizontal” dimension, and the second an ordered “vertical” dimension. The model can be used in circumstances in which individuals choose amongst products of different brands, wherein each brand offers an ordered choice menu, for example by offering products of varying quality. The unordered-ordered nature of the discrete choice problem is used to characterize the identified set of model parameters. Following an initial nonparametric analysis that relies on shape restrictions inherent in the ordered dimension of the problem, we then provide a specialized analysis for parametric specifications that generalize common ordered choice models. We characterize conditional choice probabilities as a function of model primitives with particular analysis focusing on cases in

*We thank Tim Christensen, Francesca Molinari, Lars Nesheim and several conference and seminar audiences for helpful discussions and comments. Xinyue Bei and Khuong (Lucas) Do provided excellent research assistance. We gratefully acknowledge financial support from the UK Economic and Social Research Council through a grant (RES-589-28-0001) to the ESRC Centre for Microdata Methods and Practice (CeMMAP) and through the funding of the “Programme Evaluation for Policy Analysis” node of the UK National Centre for Research Methods, as well as from the European Research Council (ERC) under grants ERC-2009-StG-240910-ROMETA and ERC-2009-AdG, grant agreement number 249529. Eleni Aristodemou gratefully acknowledges financial support from ESRC and UCL. Data supplied by TNS UK Limited. The use of TNS UK Ltd. data in this work does not imply the endorsement of TNS UK Ltd. in relation to the interpretation or analysis of the data. All errors and omissions remain the responsibility of the authors.

[†]Address: Department of Economics, Faculty of Economics and Management, University of Cyprus, P.O.Box 20537, CY1678 Nicosia, Cyprus. Email: aristodemou.d.eleni@ucy.ac.cy.

[‡]Address: Department of Economics, Duke University, 213 Social Sciences, Box 90097, Durham, NC, United States. Email: adam.rosen@duke.edu.

which unobservable taste for quality of each brand offering is multivariate normally distributed. We provide explicit formulae used for estimation and inference via maximum likelihood, and we consider inference based on Wald and quasi-likelihood ratio statistics, the latter of which can be robust to a possible lack of point identification. An empirical illustration is conducted using consumer purchase data in the UK to study consumers' choice of razor blades in which each brand has product offerings vertically differentiated by quality.

JEL classification: C01, C31, C35.

1 Introduction

In this paper we study a discrete choice model in which alternatives are distinguished by two dimensions. The alternatives are first horizontally differentiated according to one of a number of unordered categories. In the context of a consumer choice problem the alternatives could be products differentiated by brands $b = 1, \dots, \bar{b}$. Within each such category, alternatives are vertically differentiated by quality $y = 1, 2, \dots, \bar{y}_b$. Individuals are assumed to have ordered preferences over the vertical quality dimension, *within* each horizontally differentiated category, but preferences *across* horizontal categories are unordered.

There are several real-world examples of product markets that feature multiple firms competing to sell vertically differentiated alternatives to consumers. For example, airlines sell seat tickets for routes that are vertically differentiated by travel class. Cable and streaming television packages offered by different providers are often vertically-differentiated, with more expensive offerings providing more viewing options or live channels. Cellular phone providers offer vertically differentiated data and cellular plans. Competing ride-share companies offer regular and premium transportation services. Automobile dealers sell quality differentiated cars, for example luxury versus economy, offered by horizontally differentiated manufacturers. These all lead to the type of vertically and horizontally differentiated choice menus that our model aims to capture.

As initially set out by McFadden (1974), and as is now standard in the discrete choice literature, we assume that each consumer chooses the brand-quality combination that maximizes her latent utility. Our model differs from standard models of discrete choice by explicitly incorporating both the horizontal and vertical dimensions of differentiation. Models that consider choice amongst unordered discrete alternatives, such as those of McFadden (1974) and Hausman and Wise (1978), allow for horizontal differentiation by brand but do

not incorporate vertical differentiation. Models for choice amongst totally ordered alternatives can be used to estimate demand for vertically differentiated products, as in Bresnahan (1987). Nested choice models such as that of Goldberg (1995) allow flexible substitution patterns across nests, but feature unordered (e.g. logit) differentiation within nests.

Our goal here is to combine features of models for ordered and unordered choice in order to incorporate both horizontal and vertical aspects of differentiation. Relative to existing methods, this approach allows the model to respect the unordered-ordered nature of the choice problem when both kinds of differentiation are present. This may be useful for accurately estimating important features of substitution patterns in such scenarios.

A related line of research, and an important area of potential application, is the modeling of consumer choice in oligopoly markets in which competing firms each offer vertically differentiated products. Some empirical work in this area includes Davies, Waddams, and Wilson (2009) and Song (2015). Davies, Waddams, and Wilson (2009) focus on two-part tariffs and bundling in the British gas and electricity markets, and use linear panel data regression and instrumental variables to investigate whether the market operates in accord with economic theory. Song (2015) develops an explicit model of consumer demand for vertically and horizontally differentiated products, but our model and Song’s model are quite distinct and suited for different contexts. Song’s (2015) model is a hybrid of those of Berry, Levinsohn, and Pakes (1995) and Berry and Pakes (2007) and is well-suited to settings where products span multiple markets. Moreover, Song (2015) models demand for attributes in characteristics space, and is thus capable of handling a large product space. Our model is instead focused at the consumer level, requiring individual-specific choice data, and is best suited to competition among relatively few brands, or firms, with vertically differentiated product offerings, where demand depends on individual characteristics as well as on product attributes.

In our model, if attention is restricted to any single brand b , the quality of the utility-maximizing option offered by that brand for a given consumer is determined by a standard ordered choice structure. That is, the shape of the latent utility function results in an ordered choice model, e.g. ordered probit or logit, when consumers’ choices are restricted to brand b . From a modeling standpoint, this can be used to recover an indirect utility function for each brand b . The solution to the problem of choosing the best brand-quality offering from among all products can then be recovered as the brand that maximizes the indirect utility function, and the quality level that maximizes the corresponding brand-specific utility.

The structure of the problem is thus analogous to that of the mixed discrete-continuous

choice model of Dubin and McFadden (1984). However, due to the discrete nature of both dimensions of choice, one cannot use differential arguments and in particular Roy’s Identity to characterize the optimal choice of either dimension. Nonetheless, the model is complete in that conditional on any value of exogenous variables, there is a unique solution to the consumer choice problem with probability one. This is because the model is for a single-agent decision problem, rather than a simultaneous move game with strategic interactions and multiple equilibria, as encountered for instance in simultaneous equations model for ordered actions considered by Aradillas-Lopez and Rosen (2021).

The paper proceeds as follows. In Section 2 we layout our econometric model for partially ordered response in its most general form, providing shape restrictions on individuals’ underlying utility functions that deliver the within-brand ordered choice structure. In Section 3 we characterize conditional choice probabilities and provide identification analysis. The most general formulation of these restrictions offers minimal requirements for this structure but will in general yield set identification, where the identified set of model primitives can be characterized as the set of maximizers of the expected log-likelihood. In Section 4 we consider additional parametric restrictions that preserve the unordered-ordered nature of the underlying choice problem. These lend tractability to the log-likelihood and narrow the size of the identified set, and may result in point identification. Section 5 provides details for computation of the log-likelihoods and the construction of confidence intervals, and Monte Carlo analysis illustrating the relative performance of inference employing both the Wald and QLR tests, the latter having been shown to be robust to partial identification by Chen, Christensen, and Tamer (2018) under certain conditions not formally verified here.¹ Section 6 presents an application to the market for women’s razor blades using consumer data from the United Kingdom in the early 2000s. Section 7 concludes and discusses directions of continuing research. Proofs of propositions and theorems are provided in the Appendix.

2 The Model

Each individual in the population is characterized by observables (Y, B, X) and an unobservable vector V . It is assumed that each individual chooses either an ordered alternative $Y \in \mathcal{Y}_b \equiv \{1, \dots, \bar{y}_b\}$ of some type $B \in \mathcal{B} \equiv \{1, \dots, \bar{b}\}$, or an outside alternative denoted

¹Further results on the distribution of likelihood ratio statistics when point identification fails include those of Liu and Shao (2003) for parametric likelihood models and Chen, Tamer, and Torgovitsky (2011) for semi-parametric models.

by $(B, Y) = (0, 0)$. The set $\mathcal{M}_{BY} \equiv \{(b, y) : B \in \mathcal{B}, Y \in \mathcal{Y}_b\} \cup \{(0, 0)\}$ denotes the joint support of (B, Y) . The set \mathcal{X} denotes the support of observable covariates X , such as individual characteristics. The vector $V \in \mathbb{R}^{\bar{b}}$ represents unobserved heterogeneity that affects individuals' preferences, and the distribution of V is denoted $G(\cdot)$ so that for any set $\mathcal{V} \subseteq \mathbb{R}^{\bar{b}}$, $G(\mathcal{V}) \equiv \Pr[V \in \mathcal{V}]$.

The utility obtained by an individual with covariates x and unobservable v from any choice $(b, y) \in \mathcal{M}_{BY}$ is given by

$$U_{by} \equiv u(b, y, x, v_b) \text{ if } (b, y) \neq (0, 0), \quad U_{00} \equiv 0, \quad (2.1)$$

where $u(b, y, x, v_b)$ is strictly increasing in v_b for each (b, y, x) and the utility from the outside alternative is normalized to zero.

We assume that each individual chooses the alternative that maximizes her utility.² For any $b \in \mathcal{B}$ define $\bar{\mathcal{Y}}_b \equiv \mathcal{Y}_b \cup \{0\}$ and $U_{b0} \equiv 0$, and let

$$U_b^* \equiv \max_{y \in \bar{\mathcal{Y}}_b} U_{by}, \quad Y_b^* \equiv \operatorname{argmax}_{y \in \bar{\mathcal{Y}}_b} U_{by},$$

denote the indirect utility and optimal choice of Y , respectively, if the individual's alternatives were limited to only those of type b or the outside alternative. The structure of the model will be such that for any fixed b , the choice of the ordered outcome Y produces a standard model of ordered response, in the sense that this choice is weakly increasing in V_b . For example, if V_b is normally distributed, independent of X , and the consumer may only purchase from brand b , then we have an ordered probit model. A consumer who has the option to choose any quality level from any brand then chooses

$$B = 1 \left[\max_{b \in \mathcal{B}} U_b^* > 0 \right] \cdot \operatorname{argmax}_{b \in \mathcal{B}} U_b^*, \quad Y = 1 \left[\max_{b \in \mathcal{B}} U_b^* > 0 \right] \cdot Y_B^*. \quad (2.2)$$

Restriction A1: (Probability space) (B, Y, X, V) are defined on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$, where \mathcal{F} contains the Borel sets. The support of (X, V) is $\mathcal{X} \times \mathbb{R}^{\bar{b}}$, and the support of (B, Y) is $\mathcal{M}_{BY} \equiv \{(b, y) : B \in \mathcal{B}, Y \in \mathcal{Y}_b\} \cup \{(0, 0)\}$.

²Under Restriction A3 below ties in the utility obtained from different alternatives occur with zero probability conditional on any realization of x . How ties are handled is therefore of no consequence in the determination of conditional choice probabilities, but to simplify notation we adopt the convention that if alternatives (b, y) and (b, y') , $y < y'$, achieve the same utility, then (b, y) is chosen, and if $(b, y) \neq (b', y')$, $b < b'$ achieve the same utility, then (b', y') is chosen.

Restriction A2: (Identification of $f_x^0(b, y)$) For each value $x \in \mathcal{X}$ there is a proper conditional distribution of (B, Y) given $X = x$ and $f_x^0(b, y) \equiv \mathbb{P}[B = b \wedge Y = y | X = x]$ is point identified over the support of (B, Y) for almost every $x \in \mathcal{X}$.

Restriction A3: (Distribution of unobserved heterogeneity) The conditional distribution of V given $X = x$ is absolutely continuous with respect to Lebesgue measure with everywhere positive density on $\mathbb{R}^{\bar{b}}$.

Restriction A4: (Independence) X and V are stochastically independent.

Restriction A5: (Admissible structures) Structure $S \equiv (u, G)$ belongs to a known collection \mathcal{S} of pairs of utility functions and distributions of unobserved heterogeneity, (u, G) .

Restriction A6: (Utility maximization) Given (X, V) , (B, Y) are chosen to maximize $u(B, Y, X, V_b)$, where u belongs to a known class of functions \mathcal{U} satisfying (i) $u(0, 0, x, v_b) = 0$ for all (x, v_b) , (ii) for all $(b, y) \neq (0, 0)$ and all x , $u(b, y, x, v_b)$ is strictly increasing and continuous in v_b , and (iii) for each $(b, x) \in \mathcal{B} \times \mathcal{X}$, $\{u(b, y, x, v_b) : v_b \in \mathbb{R}\}$ satisfies the single-crossing property in (y, v_b) , namely that if $v'_b > v_b$ and $y' > y$, then

$$u(b, y', x, v_b) - u(b, y, x, v_b) \geq (>) 0 \Rightarrow u(b, y', x, v'_b) - u(b, y, x, v'_b) \geq (>) 0.$$

Restrictions A1-A3 are standard. Restriction A1 defines the underlying probability space and notation for the support of random variables (B, Y, X, V) . Restriction A2 stipulates that the conditional distribution of (B, Y) given covariates x is point identified for almost every $x \in \mathcal{X}$, as is the case for example under random sampling. Restriction A3 requires that unobserved heterogeneity V is absolutely continuously distributed with full support in Euclidean space.

Restriction A4 imposes independence of X and V . This is an important restriction. If X includes prices, then it requires that prices are exogenous, ruling out the possibility that unobserved components of individual utility are correlated with prices. This could be violated if different sellers offer different prices for the products being sold and if some individuals choose where to shop based on these prices. This assumption may still be appropriate however if the price of the product makes up only a small fraction of total expenditure – as is the case in our application to razor blade purchases – such that individuals are unlikely to choose where to shop for a basket of many goods (e.g. groceries and household products) based on this one product’s price. In applications with individual level variation in which prices are thought to be endogenous, one could instead use an instrumental variable approach

applicable to models of discrete choice such as those of Chesher, Rosen, and Smolinski (2013) and Chesher and Rosen (2017). This precludes the likelihood approach to inference employed here, resulting in moment inequality characterizations of identified sets, such as those used in applications to IV models for ordered response and models of interdependent binary outcomes in Chesher, Rosen, and Siddique (2019) and Chesher and Rosen (2020), respectively.

Restriction A5 defines a structure S as a utility function and distribution of unobserved heterogeneity, assumed to belong to some class of admissible pairs \mathcal{S} . Note that any given structure S gives rise to a collection of conditional distributions $f_x^0(b, y)$ for almost every $x \in \mathcal{X}$. The identification problem is to determine the set of structures that can generate the observed distributions $f_x^0(b, y)$. The set of structures \mathcal{S} admitted by the model can be restricted to a parametric, semiparametric, or nonparametric class.

The underlying structure S maps to conditional distributions $f_x^0(b, y)$ through the specification of the individual choice problem. Restriction A6 specifies that individuals choose (B, Y) to maximize utility $u(B, Y, X, V_b)$, on which we impose some conditions. First, the specification (2.1) requires that there is a single, separate component of unobserved heterogeneity for each brand b , and through Restriction A6(ii) that utility from each product of this brand is weakly increasing in the associated unobservable. The components of V may however be jointly dependent, allowing for potential correlation across brand preferences, and quality tastes across brands. With Restriction A6(i) we normalize the utility from the outside option to zero. Restriction A6(iii) requires that the utility function satisfies the single-crossing property in (y, v_b) . By Milgrom and Shannon (1994) Theorem 4 this guarantees that for all consumers and all $b \in \mathcal{B}$, the optimal choice within brand b , Y_b^* , is nondecreasing in v_b , so that quality-choice within any brand b assumes the structure of an ordered choice problem. This combined with the within brand monotonicity given by Restriction A6(ii) allows for characterizations of regions of unobservables that give rise to conditional choice probabilities for each brand-quality combination. This plays a key role in the identification of structure S .

3 Identification

We begin this section by characterizing in Section 3.1 the form of the multivariate integral delivering conditional choice probabilities as a function of the underlying structure S . In Section 3.2 we then provide a general characterization of the identified set of structures compatible with Restrictions A1-A6, and we show that if the model is correctly specified,

the identified set can be written as the maximizers of the expected log-likelihood. In Section 4 we then provide specialized identification results for a class of parametric models.

3.1 Conditional Choice Probabilities

The utility maximization hypothesis together with the shape restrictions in Restriction A6 enable concise characterization of the conditional choice probabilities

$$\wp_{by}(x; S) \equiv G(\mathcal{V}_{by}(x; u))$$

for brand-quality pair (b, y) given $X = x$, considered as a function of any structure $S = (u, G)$. Without parametric restrictions on u , the monotonicity and single-crossing conditions suffice to establish the representation of each choice probability $\wp_{by}(x; S)$ as a particular form of a \bar{b} -variate integral. Thus, given a specific (u, G) , $\wp_{by}(x; S)$ can be computed by either numerical integration or simulation. The formal result follows.

Theorem 1 *Let Restriction A6 hold. Then for each $(b, y, x) \in \mathcal{M}_{BY} \times \mathcal{X}$, the region $\mathcal{V}_{by}(x; u)$ is a convex polytope in $\mathbb{R}^{\bar{b}}$ defined by the inequalities*

$$V_b \in (g_b(y), g_b(y+1)], \quad (3.1)$$

$$\forall k < b, V_k \leq h_{b,k}(y), \text{ and} \quad (3.2)$$

$$\forall k > b, V_k < h_{b,k}(y), \quad (3.3)$$

and the conditional choice probability $\wp_{by}(x; S)$ takes the form

$$\wp_{by}(x; S) = \int_{g_b(y)}^{g_b(y+1)} \left(\int_{H(b,y)} dG_{V_{-b}|V_b}(v_{-b}|v_b) \right) dG_{V_b}(v_b), \quad (3.4)$$

where

$$H(b, y) \equiv \{r \in \mathbb{R}^{\bar{b}-1} : r_1 \leq h_{b,1}(y), \dots, r_{b-1} \leq h_{b,b-1}(y), r_b \leq h_{b,b+1}(y), \dots, r_{\bar{b}-1} \leq h_{b,\bar{b}}(y)\}$$

with $\{h_{b,k}(y) : k \neq b\}$ cross-brand threshold functions, and $\{g_b(y) : y = (0, \dots, \bar{y}_b + 1)\}$ within-brand threshold functions satisfying $g_b(0) \equiv -\infty$ and $g_b(\bar{y}_b + 1) \equiv \infty$. Here $G_{V_{-b}}(\cdot|v_b)$ and $G_{V_b}(\cdot)$ denote the conditional distribution of V_{-b} given $V_b = v_b$ and the marginal distribu-

tion of V_b , respectively. The threshold function $g_b(\cdot)$ may depend on x and each function $h_{b,k}(\cdot) : k \neq b$, may depend on both v_b and x .

3.2 General Characterization of the Identified Set

Before adding further restrictions we first characterize the identified set of structures \mathcal{S} under Restrictions A1-A6, denoted \mathcal{S}_0 and defined as

$$\mathcal{S}_0 \equiv \{(u, G) \in \mathcal{S} : \forall (b, y) \in \mathcal{M}_{BY}, G(\mathcal{V}_{by}(x; u)) = f_x^0(b, y) \text{ a.e. } x \in \mathcal{X}\}, \quad (3.5)$$

where $\mathcal{V}_{by}(x; u)$ denotes the set of values of V defined by (3.1)–(3.3) in Theorem 1. In words, \mathcal{S}_0 is the set of admissible structures (u, G) that generate identified conditional choice probabilities $f_x^0(b, y)$ for all (b, y) and almost every $x \in \mathcal{X}$. Given the absolute continuity of the distribution of V and the continuity of utility in unobserved heterogeneity, sets $\mathcal{V}_{by}(x; u)$ and $\mathcal{V}_{\tilde{b}\tilde{y}}(x; u)$, $(\tilde{b}, \tilde{y}) \neq (b, y)$, overlap at most on a set of Lebesgue measure zero, so that there is a unique utility maximizing pair (b, y) with probability one given any $x \in \mathcal{X}$. Hence $G(\mathcal{V}_{by}(x; u))$ is the conditional probability of observing (b, y) given $X = x$ when the utility function is u and $V \sim G$. Structures (u, G) that do not belong in the identified set \mathcal{S}_0 in (3.5) are those such that the set

$$\mathcal{X}^*(u, G) \equiv \{x \in \mathcal{X} : \exists (b, y) \in \mathcal{M}_{BY} \text{ s.t. } G(\mathcal{V}_{by}(x; u)) \neq f_x^0(b, y)\}, \quad (3.6)$$

has positive measure \mathbb{P}_X .

Given the representation of the identified set through the equalities $G(\mathcal{V}_{by}(x; u)) = f_x^0(b, y)$ we can equivalently characterize the identified set as those structures that maximize the log-likelihood. For this we require that the model is correctly specified, formalized with the following additional assumption.

Restriction A7: (Correct Specification) $\exists S^* \in \mathcal{S}$, $S^* \equiv (u^*, G^*)$ such that $\forall (b, y) \in \mathcal{M}_{BY}$ $G^*(\mathcal{V}_{by}(x; u^*)) = f_x^0(b, y)$ a.e. $x \in \mathcal{X}$.

Consider the expected log-likelihood function

$$Q(u, G) \equiv E[\ln G(\mathcal{V}_{BY}(X; u))],$$

where the expectation is taken with respect to population measure \mathbb{P} . It follows by arguments identical to those with singleton $\mathcal{S}_0(\mathcal{X})$ that $Q(u, G)$ attains its maximum at all $(u, G) \in \mathcal{S}_0$,

since by definition these all produce the same probabilities $G(\mathcal{V}_{by}(x; u))$ for almost every x . The general observation that when point identification is lacking the set of maximizers of the expected log-likelihood are precisely those observationally equivalent to the population data generating structure has been made previously, see e.g. Bowden (1973) and Redner (1981). The formal statement in the present setting, a proof of which is included in the appendix for completeness, is made in the following Proposition.

Proposition 1 *Let restrictions A1-A7 hold. Then $\mathcal{S}_0 = \underset{(u, G) \in \mathcal{S}}{\operatorname{argmax}} Q(u, G)$.*

Unless sufficiently strong parametric restrictions on \mathcal{S} are imposed, \mathcal{S}_0 may not be a singleton, so that there may not be point identification. When sufficiently strong restrictions for point identification do hold, estimation and inference can proceed under the classical maximum likelihood paradigm. When these restrictions do not hold, the classical results do not apply. But the characterization of \mathcal{S}_0 as the (set of) maximizers of the expected log-likelihood enables us to apply inference techniques for maximum likelihood estimators when point identification is lacking. The subsequent characterization of choice probabilities $G(\mathcal{V}_{by}(x; u))$ under parametric restrictions enables derivation of some sufficient conditions for point identification, as well as computation of set estimates and inferential statistics when point identification fails.

4 Identification under Parametric Restrictions

In this section we consider a parametric restriction on the underlying utility functions with \bar{b} firms $b \in \mathcal{B} \equiv \{1, 2, \dots, \bar{b}\}$, each selling a low-quality product offering ($Y = 1$) and a high-quality product offering ($Y = 2$), so that for all $b \in \mathcal{B}$, $\mathcal{Y}_b \equiv \{1, 2\}$.³ Specifically, we specify

$$u(b, y, x, v_b) \equiv y \times (x_b \beta_b + v_b) - \alpha_{by}, \quad \alpha_{00} \equiv 0, \quad (4.1)$$

for each $(b, y) \in \mathcal{M}_{BY}$, $x \in \mathcal{X}$, and $v_b \in \mathbb{R}$.

The parameters α_{by} capture that component of utility from choice (b, y) not restricted to scale linearly with y . These in turn determine thresholds λ_{by} for individuals' within-brand ordered choice preferences as described in (4.2) below. We consider variants of our

³Extension of the analysis to the case where firms sell more than two product offerings is conceptually straightforward. We focus on the case where each brand has two ordered product offerings to economize on notation and also because this is the setting encountered in our application.

model in which α_{by} are fixed parameters for each (b, y) , as well as cases where they may be parametrically specified functions of observable covariates, such as prices. This will lead to generalizations of ordered choice models in which threshold parameters may or may not depend on observable exogenous variables.

This model generalizes a three-choice ordered response model, such as ordered probit or logit, in that for any fixed $b \in \mathcal{B}$ we have

$$\begin{aligned} Y_b^* &= 0 \Leftrightarrow \lambda_{b0} < V_b \leq \lambda_{b1} - X_b\beta_b, \\ Y_b^* &= 1 \Leftrightarrow \lambda_{b1} - X_b\beta_b < V_b \leq \lambda_{b2} - X_b\beta_b, \\ Y_b^* &= 2 \Leftrightarrow \lambda_{b2} - X_b\beta_b < V_b, \end{aligned} \tag{4.2}$$

where

$$\lambda_{b0} \equiv -\infty, \quad \lambda_{b1} \equiv \min \left\{ \alpha_{b1}, \frac{\alpha_{b2}}{2} \right\}, \quad \lambda_{b2} \equiv \max \left\{ \alpha_{b2} - \alpha_{b1}, \frac{\alpha_{b2}}{2} \right\}. \tag{4.3}$$

denote threshold parameters. Some algebra reveals that

$$\alpha_{b2} > 2\alpha_{b1} \Rightarrow \lambda_{b1} = \alpha_{b1}, \text{ and } \lambda_{b2} = \alpha_{b2} - \alpha_{b1}, \tag{4.4}$$

while

$$\alpha_{b2} \leq 2\alpha_{b1} \Rightarrow \lambda_{b1} = \lambda_{b2} = \frac{\alpha_{b2}}{2}. \tag{4.5}$$

The inequality on the left hand side of (4.4) ensures that for each b ,

$$\mathbb{P} [\alpha_{b1} - X_b\beta_b \leq V_b \leq \alpha_{b2} - \alpha_{b1} - X_b\beta_b | X = x] > 0,$$

or equivalently that some randomly chosen individuals prefer $y = 1$ to both the other alternative of type b and the outside alternative. When instead the inequality on the left hand side of (4.5) holds, then the probability of this event is zero. In this case, if one were to imagine taking a randomly selected individual and increasing their unobservable V_b continuously from $-\infty$ to ∞ , that individual would choose the outside alternative for values of V_b up to $\frac{\alpha_{b2}}{2} - X_b\beta_b$, and then switch to $Y = 2$ for all $V_b > \frac{\alpha_{b2}}{2} - X_b\beta_b$, respecting the ordered nature of the quality dimension y , but skipping over the lower quality alternative $y = 1$.

Given the parametric specification (4.1) for u , the resulting regions of unobserved variables \mathcal{V}_{by} defined by (3.1)–(3.1) in Theorem 1 take the form of convex polytopes in $\mathbb{R}^{\bar{b}}$. Figure 1 gives an example illustrating these regions for a case in which $\bar{b} = 2$ with a particular parameter vector θ and a given value of the conditioning variables x in which the inequality

$\alpha_{b2} > 2\alpha_{b1}$ on the left hand side of (4.4) holds.

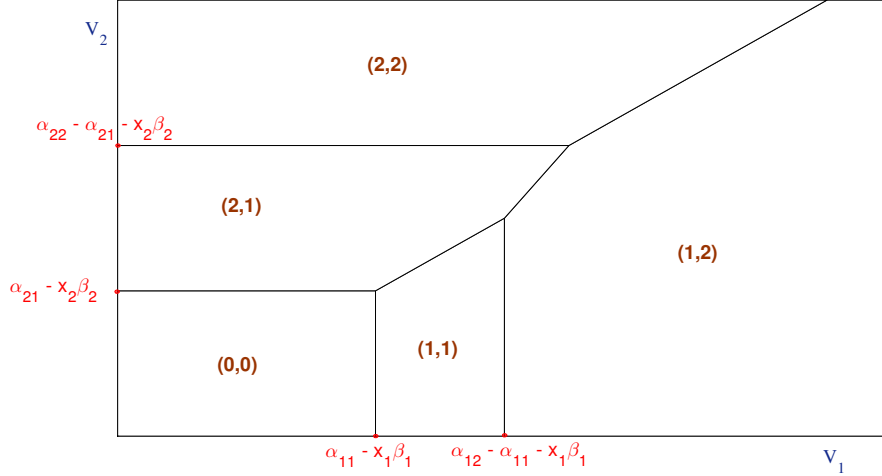


Figure 1: Regions of unobservables V resulting in each choice of $(b, y) \in \mathcal{M}_{BY}$ with utility as specified in (4.1).

The specification considered so far has implicitly treated $(\alpha_{11}, \alpha_{12}, \dots, \alpha_{\bar{b}1}, \alpha_{\bar{b}2})$ as fixed parameters to be estimated. Fixed threshold specifications for ordered response models are common, but it is straightforward to allow thresholds to be functions of observable variables. This is important in our application, where observed prices may affect the utility of purchasing each product.

Here and in our application for each $b \in \mathcal{B}$ and $y \in \{1, 2\}$, α_{by} is specified by

$$\alpha_{by} = \delta_b + k(p_{by}, \gamma_b), \quad k(p_{by}, \gamma_b) = \gamma_b p_{by}, \quad (4.6)$$

where the observed variable $p_{by} \in \mathbb{R}$ denotes the price of alternative (b, y) , and δ_b and γ_b denote parameters on the real line. More generally, p_{by} could be used to denote any vector of observable variables thought to influence α_{by} . As before the utility of the outside alternative is normalized to zero. The optimal choice Y_b^* for fixed brand b is still determined by (4.2) and (4.3), but with thresholds determined by (4.6).⁴

With price now entering the utility function, the possibility of values of exogenous variables that imply that consumers never choose quality offering $Y = 1$ for a given brand

⁴We focus here on a linear specification for $k(p_{by}, \gamma_b)$, but alternative specifications are possible. For instance, one could use the CRRA or isoelastic utility specification $k(p_{by}, \gamma_b) = (1 - \gamma_b)^{-1} p_{by}^{1-\gamma_b}$ with $\gamma_b > 0$, or the exponential utility specification $k(p_{by}, \gamma_b) = 1 - \exp(-\gamma_b p_{by})$ for some $\gamma_b \in \mathbb{R}$.

$B = b$ becomes particularly relevant. This is because unlike the brand-specific covariates X_b , prices vary across the vertical dimension y within brand. For ease of exposition, let $Z_b \equiv (X_b, P_{b1}, P_{b2})$ and $Z \equiv (Z_1, \dots, Z_{\bar{b}})$, and let Z replace X as observable utility shifters in all that follows, with support \mathcal{Z} . In contrast to the fixed threshold specification, it is possible now that there are values of the conditioning variables $Z = z$ such that the conditional choice probability $\mathbb{P}[(B, Y) = (b, 1) | Z = z]$ equals zero for either b , while conditional on other values $Z = \tilde{z}$, $\mathbb{P}[(B, Y) = (b, 1) | Z = \tilde{z}] > 0$. This is practically relevant because there may be consumers who face prices such that the higher quality product offering will always be more desirable than the lower product quality offering no matter their realization of unobservables V , as could happen when a firm introduces a sale for the high quality offering in order to induce consumers to try it. Thus both cases (4.4) and (4.5) are allowed in all that follows, depending on the value of conditioning variables Z .

We now specialize the characterization of the set \mathcal{S}_0 from Proposition 1 and the form of the conditional choice probabilities given in Theorem 1 to models in which utilities satisfy (4.1) such that formally $\mathcal{S} \equiv \mathcal{U} \times \mathcal{G}$, where

$$\mathcal{U} \equiv \left\{ \begin{array}{l} u : \mathcal{B} \times \mathcal{Y} \times \mathcal{Z} \times \mathbb{R}^{\bar{b}} \rightarrow \mathbb{R} : u(b, y, z, v; \theta) \equiv y(x_b \beta_b + v_b) - \delta_b - \gamma_b p_{by} \\ \text{for some } \theta \equiv \{(\delta_b, \gamma_b, \beta_b) : b \in \mathcal{B}\} \in \Theta. \end{array} \right\}, \quad (4.7)$$

where Θ is a compact subset of Euclidean space and \mathcal{G} denotes a collection of distribution functions for V indexed by parameter Σ , belonging to compact parameter space Λ . The utility function and distribution pair (u, G) is completely specified given ζ , so we define $\zeta \equiv (\theta, \Sigma)$ and use the full parameter vector ζ to denote the corresponding structure $(u, G) \in \mathcal{S}$, henceforth writing $G(\cdot; \Sigma)$ to denote the distribution G indicated by parameter Σ . Define

$$\mathcal{V}_{by}(z; \theta) \equiv \left\{ v \in \mathbb{R}^{\bar{b}} : \forall (b', y') \neq (b, y), u(b, y, z, v; \theta) \geq u(b', y', z, v; \theta) \right\}$$

to be the set of values of v for which choice (b, y) is optimal for parameter value θ and

$$\wp_{by}(z; \zeta) \equiv G(\mathcal{V}_{by}(z; \theta); \Sigma),$$

to be the conditional probability that $B = b$ and $Y = y$ given $Z = z$ implied by ζ .⁵ Then

⁵The set $\mathcal{V}_{by}(z; \theta)$ is equivalent to the set $\mathcal{V}_{by}(x; u)$ with the additional parametric restrictions imposed in this section, up to a set of Lebesgue measure zero.

by definition the identified set for ζ is

$$\mathcal{S}_0 \equiv \left\{ \zeta \in \Theta \times \Lambda : \forall (b, y) \in \mathcal{M}_{BY}, \varphi_{by}(z; \zeta) = f_z^0(b, y) \text{ a.e. } z \in \mathcal{Z} \right\},$$

where $f_z^0(b, y)$ is the population probability that $B = b$ and $Y = y$.

Application of Proposition 1 gives the likelihood characterization of the identified set

$$\mathcal{S}_0 \equiv \operatorname{argmax}_{\zeta \in \Theta \times \Lambda} E [\ln \varphi_{BY}(Z; \zeta)] = E_z E [\ln \varphi_{BY}(Z; \zeta) | Z = z].$$

Using the parametric structure set out above together with Theorem 1 we have

$$\varphi_{by}(z; \zeta) = \int_{g_b(y; z, \theta)}^{g_b(y+1; z, \theta)} \left(\int_{H(y, z, v_b, \theta)} dG_{V_{-b}|V_b=v_b}(v_{-b}; \Sigma) \right) dG_{V_b}(v_b; \Sigma),$$

where

$$g_b(y; z, \theta) \equiv \lambda_{by} - X_b \beta_b, \quad (4.8)$$

with each λ_{by} as defined by (4.3) and (4.6), and where the region of integration $H(b, y, v_b, \theta)$ in the inner integral is the set of $r \in \mathbb{R}^{\bar{b}-1}$ such that

$$r_1 \leq h_{b,1}(y, z, v_b, \theta), \dots, r_{b-1} \leq h_{b,b-1}(y, z, v_b, \theta), r_b \leq h_{b,b+1}(y, z, v_b, \theta), \dots, r_{\bar{b}-1} \leq h_{b,\bar{b}}(y, z, v_b, \theta),$$

where for all $d \neq b$,

$$h_{b,d}(y, z, v, \theta) \equiv \min_{\bar{y} \in \{1, \dots, \bar{y}_d\}} \frac{1}{\bar{y}} [y(x_b \beta_b + v_b) - (\alpha_{by} - \alpha_{d\bar{y}})] - x_d \beta_d, \quad (4.9)$$

with α_{by} as specified in (4.6). Thus each $\varphi_{by}(z; \zeta)$ takes the form of an integral over a region defined by inequalities that are linear in θ . Written in this form it is straightforward to verify that $\varphi_{by}(z; \zeta) = \varphi_{by}(z; (\theta, \Sigma))$ is log-concave in θ for any fixed Σ and each value (b, y, z) if G is a known log-concave distribution with density $f_V(\cdot, \Sigma)$. This in turn implies that the maximizers of $\mathcal{L}(\zeta)$ for any fixed Σ comprise a convex set.

Theorem 2 *Suppose that Restrictions A1-A7 hold, that $u \in \mathcal{U}$ defined in (4.7), and that $G(\cdot) = G(\cdot; \Sigma)$ with known Σ such that $G(\cdot; \Sigma)$ has log-concave density $f_V(\cdot, \Sigma)$. Then the identified set for θ is*

$$\Theta^*(\Sigma) \equiv \operatorname{argmax}_{\theta \in \Theta} \mathcal{L}(\theta, \Sigma),$$

with the expected log-likelihood

$$\mathcal{L}(\theta, \Sigma) \equiv \sum_{(b,y) \in \mathcal{M}_{BY}} f_z^0(b, y) \ln \int_{\mathbb{R}^{\bar{b}}} f_V(v, \Sigma) 1[v \in \mathcal{V}_{by}(z; \theta)] dv,$$

concave in θ .

Many commonly used distributions are log-concave, with the multivariate normal distribution a leading example. If the distribution G is not known, but the elements of the admissible set of distributions \mathcal{G} are all log-concave, for example if all such distributions are multivariate normal but with different variances, then it follows that the identified set for θ is contained in a union of convex sets, namely the union of sets delivered by Theorem 2 for each $G \in \mathcal{G}$.

Under some additional conditions on the variation in observable variables Z , a known G can deliver point identification, as stated in Theorem 3 below. The Theorem is a generalization of a result in Theorem 2 of Aradillas-Lopez and Rosen (2021), up to minor changes in notation, allowing for $\bar{b} > 2$ and also focusing on regions of the parameter space in which for all b $\alpha_{b2} > 2\alpha_{b1}$, equivalently $\wp_{b1}(z; \zeta) > 0$. Accordingly, define this region as

$$\mathcal{Z}^* \equiv \{z \in \mathcal{Z} : \forall b \in \mathcal{B}, \wp_{b1}(z; \zeta) > 0\}.$$

The reason the result from Aradillas-Lopez and Rosen (2021) applies with $\bar{b} = 2$ is the equivalence of the conditional probability of consumer choosing not to purchase, i.e. $\wp_0(z; \zeta)$ in the present model, to the conditional probability that $(0, 0)$ is an equilibrium in the ordered outcome simultaneous equations model studied by Aradillas-Lopez and Rosen (2021). While both models feature the same conditional probabilities for these outcomes, the rest of their observable implications differ. The simultaneous equations model of Aradillas-Lopez and Rosen (2021) produces *inequalities* on the conditional probabilities of other outcomes, due to the presence of strategic interactions and multiple equilibria. They then combine the conditional moment equality from the probability of outcome $(0, 0)$ with conditional moment inequalities to produce a test statistic for inference. In the single agent decision problem studied here, the model delivers *equalities* for the conditional probabilities of *all* outcomes. The formal result is now provided, wherein Θ_b is used to denote the projection of the parameter space for θ onto the space of admissible $(\delta_b, \gamma_b, \beta_b)$.

Theorem 3 *Suppose that Restrictions A1-A7 hold, that $\mathbb{P}\{Z \in \mathcal{Z}^*\} > 0$, and that we have the parametric structure $\mathcal{S} = \mathcal{U} \times \mathcal{G}$ given in (4.7) with singleton \mathcal{G} so that Σ is known. Then if (i) conditional on $Z \in \mathcal{Z}^*$ for each $b \in \mathcal{B}$ there exists no proper linear subspace of the support of $\tilde{Z}_b \equiv (1, P_{b1}, -X_b)$ that contains \tilde{Z}_b with probability one, and (ii) for all conformable column vectors $c_1, \dots, c_{\bar{b}}$ satisfying $c_b \in \{\theta_b - \tilde{\theta}_b : \theta_b \in \Theta_b, \tilde{\theta}_b \in \Theta_b\}$ for each b , with $c_b \neq 0$ for some b , we have that at least one of*

$$\mathbb{P}\left\{\tilde{Z}_1 c_1 \leq 0, \dots, \tilde{Z}_{\bar{b}} c_{\bar{b}} \leq 0 \mid \tilde{Z}_b c_b < 0, Z \in \mathcal{Z}^*\right\} > 0 \text{ with } \mathbb{P}\left\{\tilde{Z}_b c_b < 0, Z \in \mathcal{Z}^*\right\} > 0, \quad (4.10)$$

or

$$\mathbb{P}\left\{\tilde{Z}_1 c_1 \geq 0, \dots, \tilde{Z}_{\bar{b}} c_{\bar{b}} \geq 0 \mid \tilde{Z}_b c_b > 0, Z \in \mathcal{Z}^*\right\} > 0 \text{ with } \mathbb{P}\left\{\tilde{Z}_b c_b > 0, Z \in \mathcal{Z}^*\right\} > 0 \quad (4.11)$$

holds then θ is point identified.

The theorem above shows that under conditions that guarantee sufficient variation in exogenous variables, θ is point identified. The first requirement is that $\mathbb{P}\{Z \in \mathcal{Z}^*\} > 0$, i.e. that there is positive probability of values of exogenous variables such that $\mathbb{P}[B = b, Y = 1 \mid Z = z] > 0$ for all b . The subsequent statements (i) and (ii) are then made conditional on $Z \in \mathcal{Z}^*$. Condition (i), is standard. Note that this requires that each X_b contains no constant components. Condition (ii) implies that conditional on $\tilde{Z}_b c_b$ negative (positive), each $\tilde{Z}_d c_d$, $d \neq b$ takes nonpositive (nonnegative) values with nonzero probability. This condition helps to achieve identification because it ensures that for any $\tilde{\theta} \neq \theta$ there exist values of z such that as a function of $\tilde{\theta}$ the indices that define the cutoffs for the outside option are all either above or below the corresponding indices at $\tilde{\theta}$, with the comparison being strict for at least one index. This implies that the implied conditional probabilities of choosing the outside alternative for the two values $\tilde{\theta}$ and θ differ for such values of Z .

The support requirement of condition (ii) will in general depend on the specification of the parameter space Θ . If the parameter space is an arbitrary compact subset of Euclidean space with no restrictions imposed on $\theta_b = (\delta_b, \gamma_b, \beta_b)$ across b it will hold under standard large support assumptions requiring each Z_b to have a component taking arbitrarily large values with positive probability conditional on $Z \in \mathcal{Z}^*$. In this case it is difficult to come up with general weaker support conditions that will imply (ii). However, in models in which there are restrictions involving both θ_b and $\theta_{b'}$, $b \neq b'$, weaker conditions can suffice. For example, if the parameters θ_b relevant to the $\mathcal{V}_{00}(z; \theta)$ region are restricted to be the same

for all b , as is the case in the application of Aradillas-Lopez and Rosen (2021) in which $\bar{b} = 2$, then this condition can be satisfied even if each \tilde{Z}_b has a discrete distribution. In this case $\theta_b = \theta$ for all b , so we can also write $c_b = c$ for each b . Condition (4.10) then becomes $\mathbb{P} \left\{ \tilde{Z}_1 c \leq 0, \dots, \tilde{Z}_{\bar{b}} c \leq 0 \mid \tilde{Z}_b c < 0, Z \in \mathcal{Z}^* \right\} > 0$, which is easier to satisfy. It is trivial for instance if $\tilde{Z}_b = \tilde{Z}$ for all b .

Notably, the theorem hinges only on implied differences in the conditional probability of $Y = 0$ given Z , and exploits no additional information from the conditional probabilities of other brand-quality combinations. The conditions are shown to be sufficient, but are not shown to be necessary, so point identification may hold under weaker conditions.

Although the conditions of Theorem 3 are not directly applicable to specifications in which G is restricted to belong to a family of admissible distributions \mathcal{G} , it still carries meaningful implications in such contexts, and it is useful for understanding what type of variation is helpful for identification. Under the other conditions of the Theorem, associated with each possible $G' \in \mathcal{G}$ there can be only a singleton identified set for θ , say $\theta(G')$. Thus the identified set can only consist of parameter values for θ that are $\theta(G')$ for some $G' \in \mathcal{G}$. This implies that with a parametric specification in which \mathcal{G} is parameterized by Σ , each element of the identified set for θ must maximize $\mathcal{L}(\zeta) = \mathcal{L}(\theta, \Sigma)$ for some Σ .

A leading example of a model for which this Theorem is applicable is a partially ordered logit model in which unobservable variables $V_1, \dots, V_{\bar{b}}$ are assumed to be i.i.d. logit variates. This restriction would simplify computation, but at the cost of imposing that unobservable taste for the vertical dimension is independent of unobservable taste in the horizontal dimension. While this may be reasonable in some settings, in our application this would require unobserved preference for quality to be independent of brand preferences. Thus, we employ a specification in which the unobservables are restricted to be multivariate normal, with variance governed by parameters to be estimated. When allowance for potential correlation among unobservables is desired other parametric specifications for their joint distribution could also be used, such as the Farlie-Gumbel-Morgenstern copula in conjunction with logit marginals employed by Aradillas-Lopez and Rosen (2021).⁶

⁶This would retain the convenience of logit marginals while allowing unobservables to be correlated, but only to a limited extent relative to the bivariate normal specification employed in our application.

5 CCPs for the Partially Ordered Probit Model

In this section ζ will be used to denote the full vector of model parameters of a bivariate probit model with $\bar{b} = \bar{y} = 2$ and utility function as specified by (4.7), and $G(\cdot; \Sigma)$ corresponding to a bivariate normal distribution with variance matrix Σ with Σ_{11} normalized to one, $\Sigma_{22} = \sigma$ and $\Sigma_{12} = \Sigma_{21} = \rho\sigma$, where ρ, σ are unknown parameters. Thus $\zeta \equiv (\gamma_1, \gamma_2, \delta_1, \delta_2, \beta_1, \beta_2, \rho, \sigma)$, where each β_b is a vector of coefficients on variables X_b that affect utility from alternatives from brand b . This is the specification used in our application in Section 6. The parameter space for ζ is denoted Υ , the parameter space for ζ_k is denoted Υ_k , and Υ coincides with the product of Υ_k across $k = 1, \dots, \dim(\zeta)$. We focus attention on the partially ordered probit model, but other specifications for the distribution of unobservable heterogeneity could be used as alternatives, such as the partially ordered logit specification discussed at the end of the previous section.

The choice probabilities $\varphi_{by}(z, \zeta)$ implied by the partially ordered probit model must be computed in order to compute the likelihood. These are of the form set out in (3.4), which with $\bar{b} = 2$ take the form of a bivariate integral and can thus be computed using numerical integration or simulation for any given (ζ, z_i) . Although we experimented with implementing both approaches, maximization of the log-likelihood was found to perform relatively slowly using these methods. With (ρ, σ) unknown the log-likelihood is generally not concave in parameters.

To compute the choice probabilities (and therefore the log-likelihood) more quickly, we used results from Owen (1980) that allow us to show equivalence of the choice probabilities to a closed form expression that does not involve integration. This alternative formulation involves univariate and bivariate normal CDFs evaluated at functions of parameters and observable variables. Software was used that vectorizes function evaluation to compute the CDF at each value of z_i in one function call.⁷ This enabled computing likelihood contributions significantly more quickly than performing numerical integration or computing simulated probabilities separately for each observation. The details of how the conditional choice probabilities were manipulated to bypass the need for explicitly computing or simulating integrals are now set out. Section 6.2 then explains how inference was implemented. The finite sample performance of inference methods using both Wald and quasi-likelihood ratio (QLR) statistics is investigated in Monte Carlo experiments reported in Appendix D.

In the partially ordered probit model in which $G(\cdot; \Sigma)$ is the bivariate normal distribution,

⁷We used the pbivnorm R package Kenkel (2015), which is based on Azzalini and Genz (2016).

application of (3.4) gives the following representation for the conditional choice probabilities:

$$\forall (b, y) \in \mathcal{M}_{BY}, \quad \wp_{by}(z, \zeta) = \frac{1}{\sigma_b} \int_{g_b(y; z, \theta)}^{g_b(y+1; z, \theta)} \Phi \left(\frac{h_b(y, z, v, \theta) - \rho \frac{\sigma_d}{\sigma_b} v}{\sigma_d \sqrt{1 - \rho^2}} \right) \phi \left(\frac{v}{\sigma_b} \right) dv, \quad (5.1)$$

$$h_b(y, z, v, \theta) \equiv \min_{\bar{y} \in \{1, \dots, \bar{y}_d\}} \frac{1}{\bar{y}} [y(x_b \beta_b + v_b) - (\alpha_{by} - \alpha_{d\bar{y}})] - x_d \beta_d, \quad (5.2)$$

where $d = 3 - b$ denotes the brand other than b .⁸

To remove the need to simulate or numerically approximate the above integral, conditional choice probabilities $\wp_{by}(z, \zeta)$ can be further simplified using formulas for integrals of normal densities and distribution functions collected in Owen (1980). The representation so obtained is given in the following Proposition.

Proposition 2 *Let Restrictions A1-A7 hold with $\bar{b} = 2$, $\bar{y}_b = 2$ for each b , and*

$$u(b, y, z, v_b) \equiv y \times (x_b \beta_b + v_b) - \alpha_{by},$$

as in (4.1) with $\alpha_{00} \equiv 0$ and $V = (V_1, V_2)$ normally distributed with mean zero and variance matrix $\Sigma = \begin{pmatrix} 1 & \rho\sigma \\ \rho\sigma & \sigma^2 \end{pmatrix}$ with unknown parameters $\rho \in [-1, 1]$ and $\sigma > 0$. Then the conditional choice probabilities for each $b = 1, 2$ and $y = 1, 2$ can be expressed as

$$\wp_{by}(z, \zeta) = \begin{pmatrix} 1 [z_{by}^* < g_b(y+1; z, \theta)] \Delta(\sigma_b^{-1} \max\{z_{by}^*, g_b(y; z, \theta)\}, \sigma_b^{-1} g_b(y+1; z, \theta), m_1^+, m_2^+) \\ +1 [z_{by}^* > g_b(y; z, \theta)] \Delta(\sigma_b^{-1} g_b(y; z, \theta), \sigma_b^{-1} \min\{z_{by}^*, g_b(y+1; z, \theta)\}, m_1^-, m_2^-) \end{pmatrix}, \quad (5.3)$$

where $g_b(\cdot; z, \theta)$ is as defined in (4.8) and for any h, k, m_1, m_2 ,

$$\Delta(h, k, m_1, m_2) \equiv \Phi_2(k, m_1; m_2) - \Phi_2(h, m_1; m_2), \quad (5.4)$$

where $\Phi_2(\cdot, \cdot, \rho)$ denotes the CDF of a bivariate normal random vector Z with mean zero and unit variance components with correlation ρ , and for $d = 3 - b$,

$$z_{by}^* \equiv \frac{\alpha_{d2} + \alpha_{by} - 2\alpha_{d1}}{y} - x_b \beta_b,$$

⁸Here $h_b(y, z, v, \theta)$ is equal to $h_{b,d}(y, z, v, \theta)$ defined in (4.9) with $d = 3 - b$ since there are only two firms.

and

$$m_1^+ \equiv \frac{yx_b\beta_b + \alpha_{d2} - \alpha_{by} - 2x_d\beta_d}{\sqrt{\sigma_b^2 y^2 - 4\rho\sigma_b\sigma_d y + 4\sigma_d^2}}, \quad m_2^+ \equiv \frac{2\rho\sigma_d - \sigma_b y}{\sqrt{\sigma_b^2 y^2 - 4\rho\sigma_b\sigma_d y + 4\sigma_d^2}}, \quad (5.5)$$

$$m_1^- \equiv \frac{yx_b\beta_b + \alpha_{d1} - \alpha_{by} - x_d\beta_d}{\sqrt{\sigma_b^2 y^2 - 2\rho\sigma_b\sigma_d y + \sigma_d^2}}, \quad m_2^- \equiv \frac{\rho\sigma_d - \sigma_b y}{\sqrt{\sigma_b^2 y^2 - 2\rho\sigma_b\sigma_d y + \sigma_d^2}}. \quad (5.6)$$

6 Application to Razor Blade Purchases

This section presents an application of the parametric bivariate probit model in Section 5 to the market for women’s razor blades using consumer data from the United Kingdom in the early 2000s. We discuss the data, computational details, and empirical results in turn.

6.1 Data

We use purchase data for a rolling panel of households from the Kantar Worldpanel. The data comprise a representative sample of UK households observed making repeated purchases, obtained by a handheld scanner used to record all household grocery purchases at the UPC level. Data on razor blade purchases is used for the years 2004 – 2005.⁹ In particular we focus on consumers’ decisions to buy a double or triple blade cartridge from one of the two leading razor blade brands in the UK, Gillette and Wilkinson Sword.

In our application, we consider households observed purchasing either razor blade cartridges for a reusable non-electric women’s razor (which we refer to as “system blades”) or disposable women’s razors, where the main shopper of the household is a female between the age of 18 and 50 years old. The outside alternative is the purchase of a disposable razor. The total sample size consists of 4842 observations. Table 1 shows the observed market shares of Gillette and Wilkinson Sword system blades and disposable razors. Out of the 4842 observations, 1973 observations correspond to system blade purchases and Table 2 shows the observed market shares conditional on buying either double or triple blade cartridges from either Gillette or Wilkinson Sword.

The covariates used for each household are indicators for age of the main shopper being 31-40 and 41-50, indicators for the main shopper’s marital and employment status, and an indicator specifying the presence of more than one female in the household. Table 3 provides descriptive statistics. Further details are provided in Appendix B.

⁹Razor blades can be purchased in three forms: as disposable razors, as reusable razors sold with razor blade cartridges, or as razor blade cartridges for use with a previously purchased handle.

Razor Blade Type	Market Share
Gillette Blades	29.82%
Wilkinson Sword Blades	10.93%
Disposable Razors	59.25%

Table 1: 2004–2005 market shares.

Trading Company	Blade Type		Total
	Double Blade	Triple Blade	
Gillette	17.74%	55.45%	73.19%
Wilkinson Sword	9.22%	17.59%	26.81%
Total	26.96%	73.04%	100.00%

Table 2: Market shares conditional on purchasing double or triple blade cartridges from Gillette or Wilkinson Sword in 2004-2005.

For each observation in the sample the brand (b)-quality (y) combination of cartridges purchased is observed, in addition to the individual characteristics. For each purchase we observe the amount spent (w), the quantity purchased (q) and the specific pack size (v), as well as the month (m) and store (s) of purchase. The average per cartridge price p_{bymsv} , defined as the total amount spent divided by the total quantity purchased in a specific month and store, for each specific brand-quality combination and each pack size was computed as

$$p_{bymsv} = \frac{\sum_{i=1}^{N^{bymsv}} w_{bymsv,i}}{\sum_{i=1}^{N^{bymsv}} q_{bymsv,i}}. \quad (6.1)$$

where $i = 1, \dots, N^{bymsv}$ and N^{bymsv} is the number of observations for each $bymsv$ combination. Regarding the unit prices faced by the consumers of the cartridges not bought, since these are not observed, we estimated them using the best linear predictor (BLP) under two different specifications,

$$p_{bymsv} = \beta_0 + 1[b = 2]\beta_1 + 1[y = 2]\beta_2 + \sum_{\tilde{m}=2}^{24} 1[m = \tilde{m}]\kappa_m + \sum_{\tilde{s}=2}^{15} 1[s = \tilde{s}]\tau_s + \sum_{\tilde{v}=2}^3 1[v = \tilde{v}]\mu_v + \varepsilon, \quad (6.2)$$

Age group		Marital status		Employment status		Number of females	
18-30	34.30%	Married	61.83%	Employed	70.76%	One female	43.04%
31-40	40.75%	Other	38.17%	Unemployed	29.24%	More than one	56.96%
41-50	24.95%						

Table 3: Main shopper characteristics for 2004-2005.

$$p_{bymsv} = \beta_0 + 1[b = 2]\beta_1 + 1[y = 2]\beta_2 + \sum_{\tilde{m}=2}^{24} 1[m = \tilde{m}]\kappa_m + \sum_{\tilde{s}=2}^{15} 1[s = \tilde{s}]\tau_s + \varepsilon, \quad (6.3)$$

where ε denotes the residual of the BLP regression, $b = 2$ corresponds to Wilkinson Sword system blades, and $y = 2$ corresponds to triple blade cartridges. The intercept, β_0 , corresponds to the price of a Gillette double blade cartridge sold in January 2004 in the store group category “All other” and for specification (6.2) in the small pack size. All of the other Greek letters denote BLP coefficients. These were estimated using all 1973 observations of cartridge purchases in our data for the 24 months spanning 2004 – 2005, 15 stores, and three different pack sizes. To impute counterfactual prices for the brand-quality combination of system blades not purchased, the best linear predictors were matched to each observation according to the actual month, store and (in the case of specification (6.2)) pack size purchased. The difference in specification (6.2) and specification (6.3) is that in the former we also control for pack size. We chose to differentiate between the two specifications as not all the blade-types and/or brands offer all (original) pack sizes. For example, double blade razors were only offered in five cartridge packs. In order to deal with this, pack size was categorized according to small, medium and large. See Appendix B for further details.

Using this data our application employs the partially ordered probit specification with $\bar{b} = 2$ and $\bar{y}_b = 2$ for each b considered in Section 5 with

$$U_{by} \equiv Y(X\beta_b + V_b) - \delta_b - \gamma_b P_{by} \quad \text{if } (b, y) \neq (0, 0), \quad U_{00} \equiv 0,$$

where X are the individual specific indicators for the household’s main shopper’s age categories, marital status, employment, and the presence of multiple females in the household as previously described, and which here do not vary across brand. Prices P_{by} on the other hand vary at both the product and individual level. The index $X\beta_b + V_b$ captures the marginal effect on utility from an increase in quality Y when choosing brand b . All else equal, a higher value of this index increases both the utility of choosing the higher quality option ($Y = 2$) relative to the lower quality option ($Y = 1$) of brand b . It also increases the relative utility of brand b ’s offerings relative to those of its rival. The distribution of $V \equiv (V_1, V_2)$ is multivariate normal with parameters as indicated in Proposition 2. Thus, a positive value of the correlation parameter ρ indicates positive correlation in the taste for higher quality cartridges across the different brands, as would be expected if individuals who have higher unobserved taste for Gillette’s triple blade cartridges also have higher unobserved taste for

Wilkinson Sword’s triple blade cartridges.

6.2 Empirical Implementation

Estimation and inference were carried out by maximum likelihood. Specifically, we report point estimates obtained by maximizing the likelihood

$$L_n(\zeta) \equiv \sum_{i=1}^n [\ln \wp_{b_i y_i}(z_i; \zeta)] \quad (6.4)$$

with respect to ζ , and 95% confidence intervals for each parameter component ζ_k obtained using both Wald and QLR statistics. Wald confidence intervals are computed as those values that come within 1.96 estimated standard deviations of the point estimate, where the Hessian form of the information matrix was used to estimate the asymptotic variance. The QLR confidence intervals are more time consuming to compute, but can be robust to a possible lack of point identification under certain conditions as shown by Chen, Christensen, and Tamer (2018), which are not verified here.¹⁰ With point identification, both approaches are valid under standard conditions. In Appendix D we report the results of several Monte Carlo experiments comparing the two approaches.

The steps taken to compute the QLR confidence intervals $\mathcal{C}_{\alpha,k}^{LR}$ were as follows, for each $k = 1, \dots, \dim(\zeta)$ with $\alpha = 0.95$. First the quasi-likelihood ratio statistic

$$Q_{k,n}(\mu) \equiv \inf_{\zeta \in \Upsilon: \zeta_k = \mu} 2n [L_n^* - L_n(\zeta)], \quad (6.5)$$

¹⁰See Chen, Christensen, and Tamer (2018) for sufficient conditions for asymptotic validity of profile QLR confidence intervals for the *identified set* of parameter components ζ_k , each k , when ζ_k is not point identified. Note that a confidence interval that guarantees coverage of the entire identified set for ζ_k with pre-specified probability asymptotically is also guaranteed to contain the parameter ζ_k itself with at least the same probability asymptotically. This ensures that the confidence sets are asymptotically valid for ζ_k , although for coverage of the true parameter value only, it may in principle be possible to establish either weaker sufficient conditions or strictly smaller confidence intervals. Furthermore, while Chen, Christensen, and Tamer (2018) provide sufficient conditions for profile QLR confidence intervals to achieve at least the desired asymptotic coverage for identified sets of univariate *components* of the structural parameter, one could consider using the likelihood ratio statistic for inference on other objects, such as partial effects or various counterfactuals, although we are unaware of sufficient conditions for this to achieve valid inference under partial identification. These are questions we leave open to future research.

was computed at values of μ over a grid of values \mathcal{M} .¹¹ The values

$$\begin{aligned}\underline{\mu}_0 &\equiv \min \{ \mu \in \mathcal{M} : Q_{k,n}(\mu) \leq \chi_{1,\alpha}^2 \}, & \underline{\mu}_0 &\equiv \max \{ \mu \in \mathcal{M} : \mu < \overline{\mu}_0 \}, \\ \underline{\mu}_1 &\equiv \max \{ \mu \in \mathcal{M} : Q_{k,n}(\mu) \leq \chi_{1,\alpha}^2 \}, & \overline{\mu}_1 &\equiv \min \{ \mu \in \mathcal{M} : \mu > \underline{\mu}_1 \},\end{aligned}$$

were recorded, with $\chi_{1,\alpha}^2$ the α quantile of the chi-square distribution with one degree of freedom. Here $\overline{\mu}_0$ and $\underline{\mu}_1$ are the lowest and greatest values of μ on the grid \mathcal{M} that pass the criterion $Q_{k,n}(\mu) \leq \chi_{1,\alpha}^2$ required for $\mu \in \mathcal{C}_{\alpha,k}^{LR}$. The value $\underline{\mu}_0$ is the next lowest value to $\overline{\mu}_0$ on \mathcal{M} while $\overline{\mu}_1$ is the next highest value to $\underline{\mu}_1$ on the grid. Then a minimal tolerance $\varepsilon > 0$ was set for the desired precision within which to compute each endpoint of $\mathcal{C}_{\alpha,k}^{LR}$ and the following steps were iterated.

1. Set $\tilde{\mu} \equiv (\underline{\mu}_0 + \overline{\mu}_0) / 2$ the halfway point between $\underline{\mu}_0$ and $\overline{\mu}_0$. Compute $Q_{k,n}(\tilde{\mu})$.
2. If $Q_{k,n}(\tilde{\mu}) \leq \chi_{1,\alpha}^2$ then set $\overline{\mu}_0 \equiv \tilde{\mu}$. Otherwise set $\underline{\mu}_0 \equiv \tilde{\mu}$.
3. If $|\overline{\mu}_0 - \underline{\mu}_0| > \varepsilon$ then return to step 1 and continue. Otherwise set the terminal value $\mu_0 \equiv \underline{\mu}_0$ and stop iterating.

Then the same steps were carried out for the upper bound of $\mathcal{C}_{\alpha,k}^{LR}$ by setting $\tilde{\mu} \equiv (\underline{\mu}_1 + \overline{\mu}_1) / 2$ and replacing $\overline{\mu}_0$ with $\underline{\mu}_1$ and $\underline{\mu}_0$ with $\overline{\mu}_1$ in the subsequent step. Here we let the terminal value be denoted μ_1 . When the procedure is done, μ_0 and μ_1 serve as lower and upper bounds for $\mathcal{C}_{\alpha,k}^{LR}$.

6.3 Empirical Results

Tables 4 and 5 report maximum likelihood point estimates and confidence intervals constructed as described above using specifications (6.2) and (6.3), respectively, for counterfactual prices. The results in Table 4 lead to several observations. The coefficients γ_1 and γ_2 on price are positive, so that utility is measured to be decreasing in price. The coefficient γ_1 on price for Gillette is considerably smaller than the coefficient γ_2 for Wilkinson Sword, even after scaling by the estimate of σ_2 . The coefficient on the dummy variables for both age groups 31-40 and 41-50 are negative, as are their associated confidence intervals, with the exception of the coefficient on the 31-40 age group for Wilkinson Sword. This indicates a lower utility of system blade purchases of these age groups relative to the 18-30 age group.

¹¹The likelihood $L_n(\zeta)$ is as defined in (6.4) with $\wp_{by}(z_i; \zeta)$ defined in (5.3).

Likewise, the estimated coefficient on the more females indicator for both brands is negative. Coefficient estimates for employment and married dummy variables are negative and statistically indistinguishable from zero for both brands.

The estimated correlation coefficient between brand-specific unobservables is effectively one, indicating perfect correlation in unobserved preference for quality as reflected by blades per cartridge across the two different brands.¹² This indicates that consumers who have higher unobservable taste for the high quality product of one brand have higher unobservable taste for the higher quality product of the other brand too. This reflects choice patterns delivered by our fitted conditional choice probabilities with respect to variations in price, and is consistent with a setting in which consumers purchasing the higher quality offering of one brand would likely choose the higher quality offering of the rival brand if their chosen brand's product was not available or was prohibitively high-priced. While a high correlation in unobserved taste in the quality dimension thus seems reasonable in this application, a value of $\hat{\rho} \approx 1$ indicates roughly perfect correlation in unobserved taste for quality across brands. It also suggests that the population parameter ρ may be on the boundary of the parameter space, which can be problematic for inference. Indeed, the Hessian of the log-likelihood computed at the maximizing parameter vector was singular. Consequently, Wald confidence intervals are not reported for this specification. The QLR statistic can still be computed and the confidence intervals reported here use the χ_1^2 critical value which is valid if the parameter is on the interior of the parameter space.¹³ The QLR confidence interval obtained for ρ is very tightly concentrated around 1.

Table 5 reports results obtained using specification (6.3) for counterfactual prices. The estimate of the correlation coefficient ρ between brand-specific unobservables is again very close to one, again indicating near perfect correlation in preference for quality (blades per cartridge) across the two different brands. However, the likelihood-maximizing value of ρ was slightly lower than was found using specification (6.2) for counterfactual prices. The Hessian was nonsingular and both Wald and QLR confidence intervals for each parameter are reported.¹⁴ The point estimate for γ_1 , the price coefficient for Gillette, is negative, indicating that utility is increasing in price, although its magnitude is small and the confidence intervals indicate it is nearly statistically indistinguishable from zero. The coefficient estimate on price

¹²The maximum likelihood estimate was 0.99999992, effectively indistinguishable from 1.

¹³If in fact the population value of ρ is one, alternative critical values are needed, see for example Self and Liang (1987) and Andrews (1999).

¹⁴Small perturbations of ρ near the maximizing parameter vector were investigated and confirmed to result in a decrease in the log-likelihood.

for Wilkinson Sword is positive, and statistically significantly different from zero, indicating that utility from purchasing these products is decreasing in price. For the most part, the signs of coefficient estimates and confidence intervals on other variables accord qualitatively with those of the prior specification. Two slight exceptions are that although $\beta_{2Age31-40}$ and $\beta_{2married}$ are again estimated to be negative, their associated confidence intervals now lie fully below zero. The estimate of σ_2 is slightly larger than it was using the previous specification, but of similar magnitude.

Parameter	ML Point Estimate	QLR 95% CI
γ_1	0.0617	[-0.0380, 0.1737]
γ_2	4.2788	[3.4006, 5.6788]
δ_1	-0.202	[-0.2485, -0.2020]
δ_2	-4.0525	[-5.3889, -3.3035]
$\beta_{1Age31-40}$	-0.1566	[-0.2358, -0.0772]
$\beta_{1Age41-50}$	-0.2264	[-0.3202, -0.1323]
$\beta_{1Married}$	-0.0444	[-0.1092, 0.0202]
$\beta_{1Employed}$	-0.0382	[-0.1072, 0.0320]
$\beta_{1Females}$	-0.2236	[-0.2920, -0.1559]
$\beta_{2Age31-40}$	-0.2392	[-0.4837, 0.0042]
$\beta_{2Age41-50}$	-0.3644	[-0.6418, -0.0868]
$\beta_{2Married}$	-0.1864	[-0.3953, 0.0128]
$\beta_{2Employed}$	-0.1157	[-0.3316, 0.0915]
$\beta_{2Females}$	-0.6098	[-0.8768, -0.4080]
ρ	1.0000	[0.9994, 1.0000]
σ_2	2.6941	[2.3474, 3.3302]

Table 4: 95% confidence intervals with prices as specified in (6.2).

Table 6 reports features of the estimated own- and cross-price elasticities

$$\eta_{bykl}(z_i; \zeta) \equiv \frac{\partial \wp_{by}(z_i; \zeta)}{\partial p_{kl}} \cdot \frac{p_{kl}}{\wp_{by}(z_i; \zeta)}$$

for each (b, y) product combination and for the outside good implied by the parameter estimates reported in Table 4. For each household i the corresponding elasticity of the choice probability with respect to each price was computed. Table 6 displays the resulting means and 0.2, 0.5, and 0.8 quantiles of these elasticities.¹⁵ Relative to the means, the

¹⁵Fitted shares less than 10^{-8} were rounded to zero, and corresponding elasticity estimates were likewise set to zero to avoid approximation error. Code for computing these elasticities is included with the code used to compute the likelihood and produce the Monte Carlo results in Appendix D, which can be downloaded at <https://sites.google.com/site/amr331/home/por-code>. An additional link on this page provides a pdf web supplement that derives the expressions that were used to compute the elasticities.

Parameter	ML Point Estimate	Wald 95% CI	QLR 95% CI
γ_1	-0.0643	[-0.1274, -0.0013]	[-0.1157, -0.0036]
γ_2	2.6258	[2.3479, 2.9037]	[2.3449, 2.9131]
δ_1	-0.1650	[-0.1862, -0.1438]	[-0.1850, -0.1467]
δ_2	-2.5891	[-2.7832, -2.3951]	[-2.7805, -2.3872]
$\beta_{1Age31-40}$	-0.1595	[-0.2378, -0.0813]	[-0.2369, -0.0822]
$\beta_{1Age41-50}$	-0.2162	[-0.3079, -0.1246]	[-0.3087, -0.1240]
$\beta_{1Married}$	-0.0814	[-0.1435, -0.0194]	[-0.1435, -0.01923]
$\beta_{1Employed}$	-0.0585	[-0.1321, 0.0152]	[-0.1266, 0.0098]
$\beta_{1Females}$	-0.2378	[-0.3085, -0.1670]	[-0.3045, -0.1710]
$\beta_{2Age31-40}$	-0.3057	[-0.4779, -0.1335]	[-0.4763, -0.1361]
$\beta_{2Age41-50}$	-0.4181	[-0.6207, -0.2154]	[-0.6209, -0.2171]
$\beta_{2Married}$	-0.1937	[-0.3342, -0.0532]	[-0.3319, -0.0553]
$\beta_{2Employed}$	-0.1126	[-0.2741, 0.0488]	[-0.262, 0.0369]
$\beta_{2Females}$	-0.5249	[-0.6796, -0.3702]	[-0.6719, -0.3800]
ρ	0.9998	[0.9996, 1.0000]	[0.9996, 1.0000]
σ_2	2.1580	[2.0505, 2.2655]	[2.0334, 2.2734]

Table 5: 95% confidence intervals with prices as specified in (6.3).

quantiles illustrate considerable heterogeneity in household substitution patterns. Focusing first on mean own-price elasticities, Gillette products ($B = 1$) appear to be less sensitive to changes in own-price than Wilkinson products ($B = 2$). However, shifting focus to the quantiles of elasticities shows a more nuanced picture. For example, the estimated mean own-price elasticity for alternative (2, 1) is roughly -45, which is extremely high, but the estimated median elasticity is zero and the estimated 0.2 quantile is -2.143. In addition, the 0.1 and 0.05 quantiles (not reported in the table) are roughly -32 and -66, respectively. This reflects that the majority of households are not considering this product at the margin. For these households $\varphi_{21}(z_i; \zeta)$ is unaffected by a local change in price. However there is a small subset of households that in fact are quite sensitive to a change in the price of this product. Indeed this is consistent with this product offering having by far the smallest market share in our data. A small price increase that leads to even a handful of households no longer purchasing can produce a large percentage change in the share purchasing.

A qualitatively similar observation holds for η_{2122} ; the purchase probability φ_{21} is insensitive to a change in p_{22} for households not locally considering product (2, 1), but purchase probabilities for the segment of households that are purchasing or are locally close to purchasing this product respond considerably to changes in p_{22} . For these households, product (2, 2) appears to be the closest substitute. The large estimated mean own-price elasticity of -3.885 for φ_{22} presents a more uniform level of sensitivity to price changes, as the quantiles

are much less disperse. This product has nearly double the market share of product (2, 1) and is generally the most expensive product available.

Cross-price elasticity estimates reflect various additional substitution patterns. For instance, the elasticity estimates for Gillette’s double blade offering (1, 1) with respect to Wilkinson blade cartridges ($B = 2$) are zero for most consumers, up to at least the 0.8 quantile. However, the substantially higher mean estimate with respect to p_{21} illustrates that there is a segment of consumers who are locally sensitive to price changes in Wilkinson’s double blade razors; in this sense product (2, 1) is a closer substitute for (1, 1) than is product (2, 2). On the other hand, for Gillette’s triple blade razors, product (1, 2), there are more consumers sensitive to price changes in Wilkinson’s triple blade cartridges (2, 2) than Wilkinson’s double blade cartridges (2, 1). Given the ordered nature of the product offerings, the probability of purchasing the outside good, disposable razors, is somewhat sensitive to changes in the price of double blade razors ($Y = 1$), but is insensitive to changes in the price of triple blade razors ($Y = 2$).

Before concluding, it should be noted that these observations come with at least two important caveats. First of all, our model is static. We do not model consumers’ purchase or prior possession of handle for any brand/quality offering. Ownership of a handle for product (b, y) that is not compatible with other products may be an important factor in consumer demand and observed substitution patterns. Addressing this would require a dynamic model and data on handle purchases that may have occurred well in the past, which we do not have. Second, we do not actually observe the menu of products each consumer faces when making their purchase. Even if we are able to accurately impute counterfactual prices based on prices paid by other consumers in a given store and month, it is possible that products a consumer did not purchase were actually not available in the store at the time the consumer was shopping. This could be problematic if some stores regularly did not carry or sold out of certain products, rendering imputed prices invalid.

Despite these caveats, the application illustrates the ability of the partially ordered response model to produce consumer substitution patterns that accord with the dual horizontal and vertical dimensions of product differentiation. To illustrate this point, for the sake of comparison we report in Table 15 in Appendix C elasticities instead obtained using the alternative specific conditional logit model of McFadden (1974), implemented as the `asclogit` command in Stata, using price specification (6.2). Relative to our model, this model ignores the vertical dimension of differentiation, thus treating all products as horizontally differentiated. In comparison to the average elasticities in Table 6, the magnitude of the average

own-price elasticities for Gillette products ($B = 1$) in Table 15 is larger, indicating a greater average substitution effect following an increase in own price, while the reverse is true for Wilkson products ($B = 2$). Comparing the quantiles the picture is also quite different. An implication of the logit model is that no matter what the value of z_i for a household, there is a strictly positive probability that a marginal change in the price of the product it purchases will induce it to switch to any one of the other alternatives. Consequently, at each quantile of price elasticities, there is a non-zero proportion of households considering each product. An even more striking difference is with respect to the cross-price elasticities. While in our model the cross-price elasticities are different, the cross-price elasticities for product (b, y) with respect to $p_{k\ell}$ implied by the alternative specific conditional logit model are the same regardless of (b, y) . For example, the mean cross-price elasticity of product (b, y) induced by a change in the price of the Gillette double blade cartridge (product $(1, 1)$) reported in Table 15 is estimated as 0.0874, across *all* $(b, y) \neq (1, 1)$. This is restrictive, as we would expect individuals to respond to price changes of a given product (e.g. $(1, 1)$) differently depending on which product they are currently purchasing.¹⁶ Our model allows this, as the first column of Table 6 indicates that the share of $(1, 2)$ and the outside alternative are both locally sensitive to changes in p_{11} , although the corresponding cross-price elasticities differ. Furthermore, Table 6 indicates consumers choosing the shares of $(2, 1)$ and $(2, 2)$ are locally insensitive to changes in p_{11} . More generally it is easy to see from the columns of Table 6 that cross-price elasticities reported with respect to any given product price may vary.

The estimated elasticities further indicate interesting substitution patterns within and across brands. Looking at the rows for φ_{11} and φ_{12} , we generally see non-zero cross-price elasticities for consumers choosing Gillette's alternative product offering, but zero quantiles of many of the cross-price elasticities with respect to Wilkinson's products. This suggests most consumers who purchase Gillette cartridges are more responsive to local changes in the price of Gillette's other product than they are to local changes in Wilkinson's prices. Considering instead the rows for φ_{21} and φ_{22} we instead see that consumers purchasing Wilkinson's products are comparatively insensitive to changes in the prices of other products, with exceptions at some quantiles of estimated cross-price elasticities.

¹⁶This issue as well as other implications of the the logit elasticities have been previously discussed in the literature, for example in Nevo (2000).

Price Elasticities												
$\frac{\partial \log \varphi_{by}}{\partial \log p_{k\ell}}$	p_{11}			p_{12}			p_{21}			p_{22}		
<hr/>												
φ_{11}												
mean		-0.589			0.400			2.891			0.001	
quantiles	-0.600	-0.546	-0.425 0.364	0.424	0.457	0.000	0.000	0.000	0.000	0.000	0.000	0.000
<hr/>												
φ_{12}												
mean		0.159			-0.508			14.825			1.141	
quantiles	0.062	0.073	0.081 -0.153	-0.135	-0.117	0.000	0.000	0.000	0.000	1.332	1.629	
<hr/>												
φ_{21}												
mean		0.000			0.563			-45.376			13.450	
quantiles	0.000	0.000	0.000 0.000	0.000	0.000	-2.143	0.000	0.000	0.000	0.000	1.016	
<hr/>												
φ_{22}												
mean		0.000			0.033			0.566			-3.885	
quantiles	0.000	0.000	0.000 0.000	0.040	0.046	0.000	0.000	2.243	-4.241	-3.845	-3.445	
<hr/>												
φ_0												
mean		0.033			0.000			0.025			0.000	
quantiles	0.028	0.033	0.040 0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000

Table 6: Estimated means and 0.2, 0.5, and 0.8 quantiles of household elasticities.

7 Conclusion

In this paper we proposed a new discrete choice model for partially ordered alternatives, applicable when discrete choices are differentiated along both vertical and horizontal dimensions. We provided general characterizations of the identified set of structures admitted by the model and conditional choice probabilities under mild shape restrictions that induce an ordered (i.e. vertically differentiated) choice problem holding the horizontal dimension of choice fixed. We considered specialized results for a partially ordered probit model in which brand-specific unobservables are restricted to be multivariate normal with parameterized covariance matrix. Closed form expressions for choice probabilities were obtained in a duopoly (two-brand) setting, useful for computing the log-likelihood. An empirical illustration was provided using data on razor blade cartridge purchases in the United Kingdom, a setting that features two dominant competing firms with vertically differentiated products. The application demonstrated the model’s ability to capture interesting substitution

patterns commensurate with the two levels of differentiation, and which reflect underlying heterogeneity in preferences.

Estimating such patterns as accurately as possible can be important for measuring welfare, predicting reactions to the introduction of new products, and for modeling firm competition. Indeed, future analysis combining demand modeled by way of a vertically and horizontally ordered choice model with a model of equilibrium firm behavior could be useful for estimating such quantities. In this way, the choice model developed here could potentially be applied to study equilibrium pricing by multiproduct firms engaging in second degree price discrimination. This could in turn be used to empirically measure the direction and magnitude of welfare effects of competition in markets with non-linear pricing, for which the incorporation of both horizontal and vertical differentiation is important.

References

- ANDREWS, D. W. K. (1999): “Estimation When a Parameter is on a Boundary,” *Econometrica*, 67(6), 1341–1383.
- ARADILLAS-LOPEZ, A., AND A. M. ROSEN (2021): “Inference in Ordered Response Games with Complete Information,” Working Paper, available at https://www.personal.psu.edu/aza12/ordered_game.pdf.
- AZZALINI, A., AND A. GENZ (2016): *The R package mnormt: The multivariate normal and t distributions (version 1.5-4)*.
- BERRY, S., J. LEVINSOHN, AND A. PAKES (1995): “Automobile Prices in Market Equilibrium,” *Econometrica*, 63(4), 841–890.
- BERRY, S., AND A. PAKES (2007): “The Pure Characteristics Demand Model,” *International Economic Review*, 48, 1193–1225.
- BOWDEN, R. (1973): “Theory of Parametric Identification,” *Econometrica*, 41(6), 1069–1074.
- BRESNAHAN, T. F. (1987): “Competition and Collusion in the American Automobile Industry: the 1955 Price War,” *Journal of Industrial Economics*, 35(4), 457–482.
- CAMERON, A. C., AND P. K. TRIVEDI (2009): *Microeconometrics using stata*. College Station (TX), Stata Press.

- CHEN, X., T. M. CHRISTENSEN, AND E. TAMER (2018): “Monte Carlo Confidence Sets for Identified Sets,” *Econometrica*, 86(6), 1965–2018.
- CHEN, X., E. TAMER, AND A. TORGOVITSKY (2011): “Sensitivity Analysis in Semiparametric Likelihood Models,” Cowles Foundation working paper 1836.
- CHESHER, A., AND A. ROSEN (2017): “Generalized Instrumental Variable Models,” *Econometrica*, 85(3), 959–989.
- CHESHER, A., AND A. M. ROSEN (2020): “Econometric Modeling of Interdependent Discrete Choice with Applications to Market Structure,” CeMMAP working paper CWP25/20.
- CHESHER, A., A. M. ROSEN, AND Z. SIDDIQUE (2019): “Estimating Endogenous Effects on Ordinal Outcomes,” CeMMAP working paper CWP66/19.
- CHESHER, A., A. M. ROSEN, AND K. SMOLINSKI (2013): “An Instrumental Variable Model of Multiple Discrete Choice,” *Quantitative Economics*, 4(2), 157–196.
- DATTA, Y. (2019): “The U.S. Womens Razor-Blade Market: A Competitive Profile,” *Journal of Economics and Public Finance*, 5, 491–508.
- DAVIES, S., C. WADDAMS, AND C. WILSON (2009): “How Far Does Economic Theory Explain Competitive Nonlinear Pricing in Practice?,” Centre for Competition Policy Working Paper 09-7.
- DUBIN, J., AND D. MCFADDEN (1984): “An Econometric Analysis of Residential Electric Appliance Holdings and Consumption,” *Econometrica*, 52, 345–362.
- GHALANOS, A., AND T. STEFAN (2015): *Rsolnp: General Non-linear Optimization Using Augmented Lagrange Multiplier Method (version 1.16)*.
- GOLDBERG, P. K. (1995): “Product Differentiation and Oligopoly in International Markets: The Case of the U.S. Automobile Industry,” *Econometrica*, 63(4), 891–951.
- HAUSMAN, J., AND D. WISE (1978): “A Conditional Probit Model for Qualitative Choice: Discrete Decisions Recognizing Interdependence and Heterogeneous Preferences,” *Econometrica*, 46(2), 403–426.

- KENKEL, B. (2015): *The R package pbivnorm: Vectorized Bivariate Normal CDF (version 0.60)*. Based on the `mnormt` package. Fortran code by Alan Genz.
- LIU, X., AND Y. SHAO (2003): “Asymptotics for Likelihood Ratio Tests under Loss of Identifiability,” *Annals of Statistics*, 31(3), 807–832.
- MCFADDEN, D. (1974): “Conditional Logit Analysis of Qualitative Choice Behavior,” in *Frontiers in Econometrics*, ed. by P. Zarembka. New York: Academic Press.
- MILGROM, P., AND C. SHANNON (1994): “Monotone Comparative Statics,” *Econometrica*, 62(1), 157–180.
- NEVO, A. (2000): “A practitioner’s guide to estimation of random-coefficients logit models of demand,” *Journal of economics & management strategy*, 9(4), 513–548.
- OWEN, D. B. (1980): “A Table of Normal Integrals,” *Communications in Statistics - Simulation and Computation*, 4(4), 389–419.
- PREKOPA, A. (1973): “On logarithmic concave measures and functions,” *Acta Sci. Math. (Szeged)*, 34, 335–343.
- REDNER, R. (1981): “Note on the Consistency of the Maximum Likelihood Estimate for Nonidentifiable Distributions,” *Annals of Statistics*, 9(1), 225–228.
- SELF, S. G., AND K.-Y. LIANG (1987): “Asymptotic Properties of Maximum Likelihood Estimators and Likelihood Ratio Tests Under Nonstandard Conditions,” *Journal of the American Statistical Association*, 82(398), 605–610.
- SONG, M. (2015): “A Hybrid Discrete Choice Model of Differentiated Product Demand with an Application to Personal Computers,” *International Economic Review*, 56(1), 265–301.

A Proofs

Proof of Proposition 1. First consider any $(u_0, G_0) \in \mathcal{S}_0$. By the same argument as when there is point identification we have for almost every $x \in \mathcal{X}$,

$$E[\ln G_0(\mathcal{V}_{BY}(x; u_0)) | x] \geq E[\ln G(\mathcal{V}_{BY}(x; u)) | x] \quad (\text{A.1})$$

for all $(u, G) \in \mathcal{S}$. Thus \mathcal{S}_0 is contained in the set of maximizers of $Q(u, G)$. Consider now $(\tilde{u}, \tilde{G}) \notin \mathcal{S}_0$. Then for some $(b, y) \in \mathcal{M}_{BY}$ there exists a positive measure set $\mathcal{X}^*(\tilde{u}, \tilde{G})$ as defined in (3.6) on which $\tilde{G}(\mathcal{V}_{by}(x; \tilde{u})) \neq G_0(\mathcal{V}_{BY}(x; u_0)) = f_x^0(b, y)$ for at least one (b, y) pair. We therefore have

$$\forall x \in \mathcal{X}^*(\tilde{u}, \tilde{G}), E[\ln G_0(\mathcal{V}_{BY}(x; u_0)) | x] > E[\ln \tilde{G}(\mathcal{V}_{BY}(x; \tilde{u})) | x].$$

Combining this with (A.1) it follows that $Q(u_0, G_0) > Q(\tilde{u}, \tilde{G})$, completing the proof. ■

Proof of Theorem 1. From the utility maximization hypothesis, (b, y) is chosen if and only if it maximizes $u(b, y, x, v_b)$. This is so if and only (i) (b, y) provides higher utility than that delivered by all within brand options $\{u(b, \tilde{y}, x, v_b) : \tilde{y} \neq y\}$, and (ii) (b, y) provides higher utility than that delivered by all alternative brand options $\{u(\tilde{b}, \tilde{y}, x, v_{\tilde{b}}) : (\tilde{b}, \tilde{y}) \neq (b, y)\}$.

Condition (i) requires that y maximizes $u(b, \cdot, x, v_b)$ for the stated brand b , that is $Y_b^* = y$. Given the single-crossing property of Restriction A6(iii) we can apply Theorem 4 of Milgrom and Shannon (1994), implying that Y_b^* is nondecreasing in v_b . It follows that for each $y \in \{0, \dots, \bar{y}_b + 1\}$, there is a nondecreasing sequence of thresholds $\{g_b(y) : y = (0, \dots, \bar{y}_b + 1)\}$ such that $Y_b^* = y$ if and only if $v_b \in (g_b(y), g_b(y + 1)]$, where possibly $g_b(y) = g_b(y + 1)$ if alternative (b, y) is never chosen. That $g_b(0) \equiv -\infty$ and $g_b(\bar{y}_b + 1) \equiv \infty$ follows from $y = 0$ and \bar{y}_b being the lowest and highest feasible values of y .

Condition (ii) stems from Restriction A6(ii), strict monotonicity of $u(b, y, x, v_b)$ in v_b for each b . The consumer will choose brand b if and only if for any other brand d , the utility from choosing (b, Y_b^*) exceeds that from choosing (d, Y_d^*) , that is if

$$u(b, Y_b^*, x, v_b) > \max_{y \in \mathcal{Y}_d} u(d, y, x, v_d), \text{ if } b < d, \quad (\text{A.2})$$

$$u(b, Y_b^*, x, v_b) \geq \max_{y \in \mathcal{Y}_d} u(d, y, x, v_d), \text{ if } b > d. \quad (\text{A.3})$$

By A6(ii) it follows that

$$u_d^*(x, v_d) \equiv \max_{y \in \mathcal{Y}_d} u(d, y, x, v_d)$$

is strictly monotone and thus invertible in v_d . Inequalities (A.2) and (A.3) thus simplify to

$$q_d \{u(b, Y_b^*, x, v_b); x\} > v_d, \text{ if } b < d,$$

$$q_d \{u(b, Y_b^*, x, v_b); x\} \geq v_d, \text{ if } b > d,$$

where $q_d(\cdot; x)$ denotes the inverse of $u_d^*(x, v_d)$ with respect to v_d , i.e. for any (x, v_d) ,

$$q_d(u_d^*(x, v_d); x) = v_d.$$

Then we have the inequalities (3.2) and (3.3) with

$$h_{b,d}(y) \equiv q_d\{u(b, Y_b^*, x, v_b); x\},$$

for each pair $b \neq d$. The integral (3.4) for the conditional choice probabilities then follows immediately from their definition $\wp_{by}(x; \mathcal{S}) \equiv G(\mathcal{V}_{by}(x; u))$. \blacksquare

Proof of Theorem 2. It is straightforward to verify that the function

$$h(v, \theta) \equiv g(v) 1[v \in \mathcal{V}_{by}(z; \theta)]$$

is log-concave in (v, θ) . This follows from log-concavity of $g(v)$ and log-concavity of $1[v \in \mathcal{V}_{by}(z; \theta)]$ in (v, θ) , which is easy to establish given $\mathcal{V}_{by}(z; \theta)$ comprises a system of linear inequalities in (v, θ) . By Theorem 6 of Prekopa (1973) it then follows that

$$\int h(v, \theta) dv$$

is log-concave in θ and concavity of $\mathcal{L}(\theta, G)$ follows. \blacksquare

Proof of Theorem 3. Let $\tilde{\theta} \neq \theta$ and for each $b \in \mathcal{B}$ let $\theta_b \equiv (\delta_b, \gamma_b, \beta_b)'$ and $\tilde{\theta}_b \equiv (\tilde{\delta}_b, \tilde{\gamma}_b, \tilde{\beta}_b)'$, $\theta_{-b} \equiv (\theta'_1, \dots, \theta'_{b-1}, \theta'_{b+1}, \dots, \theta'_b)'$, and likewise for $\tilde{\theta}_{-b}$. Let k_b denote the number of components of θ_b .

Let

$$\mathbf{z} \equiv \begin{pmatrix} \tilde{z}_{1,1} & \cdots & \tilde{z}_{1,k_1} & 0 & \cdots & \cdots & \cdots & \cdots & \cdots & 0 \\ 0 & \cdots & 0 & \tilde{z}_{2,1} & \cdots & \tilde{z}_{2,k_2} & 0 & \cdots & 0 & \vdots \\ 0 & \cdots & \cdots & \cdots & \cdots & 0 & \ddots & 0 & \cdots & 0 \\ 0 & \cdots & \cdots & \cdots & \cdots & \cdots & 0 & \tilde{z}_{\bar{b},1} & \cdots & \tilde{z}_{\bar{b},k_{\bar{b}}} \end{pmatrix}$$

and let \mathbf{z}_{-b} denote the matrix \mathbf{z} with the b_{th} row removed. Let \mathbf{z}_b denote the b_{th} row of \mathbf{z} excluding zero entries. As in the main text, z (not in bold font) is used to denote a realization of all exogenous variables excluding repetitions.

Define the sets

$$S_b^+ \equiv \left\{ z \in \mathcal{Z}^* : \mathbf{z}_b (\tilde{\theta}_b - \theta_b) > 0 \wedge \mathbf{z}_{-b} (\tilde{\theta}_{-b} - \theta_{-b}) \geq 0 \right\},$$

$$S_b^- \equiv \left\{ z \in \mathcal{Z}^* : \mathbf{z}_b \left(\tilde{\theta}_b - \theta_b \right) < 0 \wedge \mathbf{z}_{-b} \left(\tilde{\theta}_{-b} - \theta_{-b} \right) \leq 0 \right\}.$$

Let $\Phi_{\bar{b}}(\mathbf{z}\theta; \Sigma)$ denote the cumulative distribution function for the \bar{b} -variate distribution $G(\cdot, \Sigma)$. For any $z \in S_b^+$ we have that

$$\Phi_{\bar{b}}(\mathbf{z}\tilde{\theta}; \Sigma) > \Phi_{\bar{b}}(\mathbf{z}\theta; \Sigma) = f_x^0(0), \quad (\text{A.4})$$

and likewise for any $z \in S_b^-$,

$$\Phi_{\bar{b}}(\mathbf{z}\tilde{\theta}; \Sigma) < \Phi_{\bar{b}}(\mathbf{z}\theta; \Sigma) = f_x^0(0), \quad (\text{A.5})$$

where $f_x^0(0) = \mathbb{P}\{Y = 0 | X = x\}$. The probability that $Z \in S_b \equiv S_b^+ \cup S_b^-$ is

$$\begin{aligned} \mathbb{P}\{Z \in S_b\} &= \mathbb{P}\{Z \in S_b^+\} + \mathbb{P}\{Z \in S_b^-\} \\ &= \left(\begin{array}{l} \mathbb{P}\left\{ \mathbf{Z}_{-b} \left(\tilde{\theta}_{-b} - \theta_{-b} \right) \geq 0 \mid \mathbf{Z}_b \left(\tilde{\theta}_b - \theta_b \right) > 0, Z \in \mathcal{Z}^* \right\} \mathbb{P}\left\{ \mathbf{Z}_b \left(\tilde{\theta}_b - \theta_b \right) > 0, Z \in \mathcal{Z}^* \right\} \\ + \mathbb{P}\left\{ \mathbf{Z}_{-b} \left(\tilde{\theta}_{-b} - \theta_{-b} \right) \leq 0 \mid \mathbf{Z}_b \left(\tilde{\theta}_b - \theta_b \right) < 0, Z \in \mathcal{Z}^* \right\} \mathbb{P}\left\{ \mathbf{Z}_b \left(\tilde{\theta}_b - \theta_b \right) < 0, Z \in \mathcal{Z}^* \right\} \end{array} \right) \end{aligned}$$

At least one of $\mathbb{P}\left\{ \mathbf{Z}_b \left(\tilde{\theta}_b - \theta_b \right) > 0, Z \in \mathcal{Z}^* \right\}$ and $\mathbb{P}\left\{ \mathbf{Z}_b \left(\tilde{\theta}_b - \theta_b \right) < 0, Z \in \mathcal{Z}^* \right\}$ are strictly positive because of condition (i), and $\mathbb{P}\{Z \in \mathcal{Z}^*\} > 0$, and it follows that at least one of

$$\mathbb{P}\left\{ \mathbf{Z}_{-b} \left(\tilde{\theta}_{-b} - \theta_{-b} \right) \geq 0 \mid \mathbf{Z}_b \left(\tilde{\theta}_b - \theta_b \right) > 0, Z \in \mathcal{Z}^* \right\} \mathbb{P}\left\{ \mathbf{Z}_b \left(\tilde{\theta}_b - \theta_b \right) > 0, Z \in \mathcal{Z}^* \right\},$$

$$\mathbb{P}\left\{ \mathbf{Z}_{-b} \left(\tilde{\theta}_{-b} - \theta_{-b} \right) \leq 0 \mid \mathbf{Z}_b \left(\tilde{\theta}_b - \theta_b \right) < 0, Z \in \mathcal{Z}^* \right\} \mathbb{P}\left\{ \mathbf{Z}_b \left(\tilde{\theta}_b - \theta_b \right) < 0, Z \in \mathcal{Z}^* \right\},$$

or both must be strictly positive by condition (ii). Therefore $\mathbb{P}\{Z \in S_b\} > 0$, implying that $\tilde{\theta}$ is observationally distinct from θ since for each $z \in S_b$, $f_x^0(0) \neq \Phi_{\bar{b}}(\mathbf{z}\tilde{\theta}; \Sigma)$. \blacksquare

Before proving Proposition 2, the following Lemma is first proven.

Lemma 1 *When $\bar{y}_b = 2$ for each b , then (5.2) can be simplified to*

$$h_b(y, z, v, \theta) = 1 [v_b < z_{by}^*] m_{by}^-(z, v_b, \theta) + 1 [v_b \geq z_{by}^*] m_{by}^+(z, v_b, \theta), \quad (\text{A.6})$$

where

$$m_{by}^-(z, v_b, \theta) \equiv y(x_b \beta_b + v_b) + \alpha_{d1} - \alpha_{by} - x_d \beta_d, \quad (\text{A.7})$$

$$m_{by}^+(z, v_b, \theta) \equiv \frac{1}{2} [y(x_b \beta_b + v_b) + \alpha_{d2} - \alpha_{by}] - x_d \beta_d, \quad (\text{A.8})$$

Proof. Since $\bar{y}_b = 2$, (5.2) simplifies to

$$h_b(y, z, v, \theta) = \min_{\bar{y} \in \{1, \dots, \bar{y}_d\}} \frac{1}{\bar{y}} [y(x_b \beta_b + v_b) - (\alpha_{by} - \alpha_{d\bar{y}})] - x_d \beta_d \quad (\text{A.9})$$

$$= \min \{m_{by}^-(z, v_b, \theta), m_{by}^+(z, v_b, \theta)\}. \quad (\text{A.10})$$

Both $m_{by}^-(z, v_b, \theta)$ and $m_{by}^+(z, v_b, \theta)$ are linear and strictly increasing in v_b . Setting $m_{by}^-(z, v_b, \theta) = m_{by}^+(z, v_b, \theta)$ reveals that the two functions are equal at

$$v_b = z_{by}^* \equiv \frac{\alpha_{d2} + \alpha_{by} - 2\alpha_{d1}}{y} - x_b \beta_b,$$

and since $m_{by}^-(z, v_b, \theta)$ has a larger slope with respect to v_b , it follows that for all $v_b < z_{by}^*$, $m_{by}^-(z, v_b, \theta) < m_{by}^+(z, v_b, \theta)$, while for all $v_b > z_{by}^*$, $m_{by}^-(z, v_b, \theta) > m_{by}^+(z, v_b, \theta)$. Thus (A.10) simplifies to (A.6), completing the proof. \blacksquare

Proof of Proposition 2. The starting point is (5.1):

$$\wp_{by}(z, \zeta) = \frac{1}{\sigma_b} \int_{g_b(y; z, \theta)}^{g_b(y+1; z, \theta)} \Phi \left(\frac{h_b(y, z, v, \theta) - \rho \frac{\sigma_d}{\sigma_b} v}{\sigma_d \sqrt{1 - \rho^2}} \right) \phi \left(\frac{v}{\sigma_b} \right) dv,$$

which is broken into three cases, depending on whether z_{by}^* lies below, inside, or above the interval $[g_b(y; z, \theta), g_b(y+1; z, \theta)]$ on which the integral is to be evaluated.

1. $g_b(y; z, \theta) < z_{by}^* < g_b(y+1; z, \theta)$.

$$\begin{aligned} \wp_{by}(z, \zeta) &= \sigma_b^{-1} \int_{g_b(y; z, \theta)}^{z_{by}^*} \Phi \left(\frac{m_{by}^-(z, v, \theta) - \rho \frac{\sigma_d}{\sigma_b} v}{\sigma_d \sqrt{1 - \rho^2}} \right) \phi \left(\frac{v}{\sigma_b} \right) dv \\ &\quad + \sigma_b^{-1} \int_{z_{by}^*}^{g_b(y+1; z, \theta)} \Phi \left(\frac{m_{by}^+(z, v, \theta) - \rho \frac{\sigma_d}{\sigma_b} v}{\sigma_d \sqrt{1 - \rho^2}} \right) \phi \left(\frac{v}{\sigma_b} \right) dv \quad (\text{A.11}) \end{aligned}$$

2. $g_b(y; z, \theta) \leq g_b(y+1; z, \theta) \leq z_{by}^*$.

$$\wp_{by}(z, \zeta) = \sigma_b^{-1} \int_{g_b(y; z, \theta)}^{g_b(y+1; z, \theta)} \Phi \left(\frac{m_{by}^-(z, v, \theta) - \rho \frac{\sigma_d}{\sigma_b} v}{\sigma_d \sqrt{1 - \rho^2}} \right) \phi \left(\frac{v}{\sigma_b} \right) dv \quad (\text{A.12})$$

3. $z_{by}^* \leq g_b(y; z, \theta) \leq g_b(y+1; z, \theta)$.

$$\rho_{by}(z, \zeta) = \sigma_b^{-1} \int_{g_b(y; z, \theta)}^{g_b(y+1; z, \theta)} \Phi \left(\frac{m_{by}^+(z, v, \theta) - \rho \frac{\sigma_d}{\sigma_b} v}{\sigma_d \sqrt{1 - \rho^2}} \right) \phi \left(\frac{v}{\sigma_b} \right) dv \quad (\text{A.13})$$

The expressions in each case simplify as follows, using (A.7) and (A.8) and a change of variables substitution for $\frac{v}{\sigma_b}$.

$$\begin{aligned} \sigma_b^{-1} \int_{g_b(y; z, \theta)}^{z_{by}^*} \Phi \left(\frac{m_{by}^-(z, v, \theta) - \rho \frac{\sigma_d}{\sigma_b} v}{\sigma_d \sqrt{1 - \rho^2}} \right) \phi \left(\frac{v}{\sigma_b} \right) dv \\ = \int_{\sigma_b^{-1} g_b(y; z, \theta)}^{\sigma_b^{-1} z_{by}^*} \Phi \left(\frac{yx_b \beta_b + \alpha_{d1} - \alpha_{by} - x_d \beta_d + (\sigma_b y - \rho \sigma_d) v}{\sigma_d \sqrt{1 - \rho^2}} \right) \phi(v) dv \end{aligned} \quad (\text{A.14})$$

$$\begin{aligned} \sigma_b^{-1} \int_{z_{by}^*}^{g_b(y+1; z, \theta)} \Phi \left(\frac{m_{by}^+(z, v, \theta) - \rho \frac{\sigma_d}{\sigma_b} v}{\sigma_d \sqrt{1 - \rho^2}} \right) \phi \left(\frac{v}{\sigma_b} \right) dv \\ = \int_{\sigma_b^{-1} z_{by}^*}^{\sigma_b^{-1} g_b(y+1; z, \theta)} \Phi \left(\frac{\frac{1}{2} [yx_b \beta_b + \alpha_{d2} - \alpha_{by}] - x_d \beta_d + (\frac{1}{2} \sigma_b y - \rho \sigma_d) v}{\sigma_d \sqrt{1 - \rho^2}} \right) \phi(v) dv. \end{aligned} \quad (\text{A.15})$$

$$\begin{aligned} \sigma_b^{-1} \int_{g_b(y; z, \theta)}^{g_b(y+1; z, \theta)} \Phi \left(\frac{m_{by}^-(z, v, \theta) - \rho \frac{\sigma_d}{\sigma_b} v}{\sigma_d \sqrt{1 - \rho^2}} \right) \phi \left(\frac{v}{\sigma_b} \right) dv \\ = \int_{\sigma_b^{-1} g_b(y; z, \theta)}^{\sigma_b^{-1} g_b(y+1; z, \theta)} \Phi \left(\frac{yx_b \beta_b + \alpha_{d1} - \alpha_{by} - x_d \beta_d + (\sigma_b y - \rho \sigma_d) v}{\sigma_d \sqrt{1 - \rho^2}} \right) \phi(v) dv \end{aligned} \quad (\text{A.16})$$

$$\begin{aligned}
& \sigma_b^{-1} \int_{g_b(y;z,\theta)}^{g_b(y+1;z,\theta)} \Phi \left(\frac{m_{by}^+(z, v, \theta) - \rho \frac{\sigma_d}{\sigma_b} v}{\sigma_d \sqrt{1 - \rho^2}} \right) \phi \left(\frac{v}{\sigma_b} \right) dv \\
&= \int_{\sigma_b^{-1} g_b(y;z,\theta)}^{\sigma_b^{-1} g_b(y+1;z,\theta)} \Phi \left(\frac{\frac{1}{2} [yx\beta_b + \alpha_{d2} - \alpha_{by}] - x_d \beta_d + (\frac{1}{2} \sigma_b y - \rho \sigma_d) v}{\sigma_d \sqrt{1 - \rho^2}} \right) \phi(v) dv \quad (\text{A.17})
\end{aligned}$$

Page 403 of Owen (1980) gives us formula 10,010.4:

$$\int_h^k \Phi(c_1 + c_2 z) \phi(z) dz = \Lambda(k, c_1, c_2) - \Lambda(h, c_1, c_2) \quad (\text{A.18})$$

where the function $\Lambda(\cdot, \cdot, \cdot)$ is given by

$$\Lambda(k, c_1, c_2) = \int_{-\infty}^{\frac{c_1}{\sqrt{c_2^2 + 1}}} \phi(z) \Phi \left(k \sqrt{c_2^2 + 1} + c_2 z \right) dz. \quad (\text{A.19})$$

Formula 10,010.1 on page 402 of Owen (1980) is

$$\int_{-\infty}^y \phi(z) \Phi(a + bz) dz = \Phi_2 \left(\frac{a}{\sqrt{1 + b^2}}, y; \frac{-b}{\sqrt{1 + b^2}} \right). \quad (\text{A.20})$$

Applying this formula to (A.19) with

$$a = k \sqrt{c_2^2 + 1}, \quad b = c_2$$

gives

$$\Lambda(k, c_1, c_2) = \Phi_2 \left(k, \frac{c_1}{\sqrt{c_2^2 + 1}}; \frac{-c_2}{\sqrt{1 + c_2^2}} \right).$$

Define now

$$\begin{aligned}
c_1^- &\equiv \frac{yx_b \beta_b + \alpha_{d1} - \alpha_{by} - x_d \beta_d}{\sigma_d \sqrt{1 - \rho^2}}, & c_2^- &\equiv \frac{\sigma_b y - \rho \sigma_d}{\sigma_d \sqrt{1 - \rho^2}}, \\
c_1^+ &\equiv \frac{yx_b \beta_b + \alpha_{d2} - \alpha_{by} - 2x_d \beta_d}{2\sigma_d \sqrt{1 - \rho^2}}, & c_2^+ &\equiv \frac{\sigma_b y - 2\rho \sigma_d}{2\sigma_d \sqrt{1 - \rho^2}}.
\end{aligned}$$

as well as

$$\tilde{\Delta}(h, k, c_1, c_2) \equiv \Lambda(k, c_1, c_2) - \Lambda(h, c_1, c_2).$$

Referring back to (A.18), substitution of c_2 with those coefficients multiplying v and substitution of c_1 with those terms not multiplying v in the integrands on the right hand side of (A.14) - (A.17) combined with (A.11) - (A.13) gives the following expression for conditional choice probabilities according to where z_{by}^* lies with respect to the interval $[g_b(y; z, \theta), g_b(y+1; z, \theta)]$.

$$1. \ g_b(y; z, \theta) < z_{by}^* < g_b(y+1; z, \theta).$$

$$\begin{aligned} \wp_{by}(z, \zeta) &= \sigma_b^{-1} \left\{ \begin{array}{l} \int_{g_b(y; z, \theta)}^{z_{by}^*} \Phi\left(\frac{m_{by}^-(z, v, \theta) - \rho \frac{\sigma_d}{\sigma_b} v}{\sigma_d \sqrt{1 - \rho^2}}\right) \phi\left(\frac{v}{\sigma_b}\right) dv \\ + \int_{z_{by}^*}^{g_b(y+1; z, \theta)} \Phi\left(\frac{m_{by}^+(z, v, \theta) - \rho \frac{\sigma_d}{\sigma_b} v}{\sigma_d \sqrt{1 - \rho^2}}\right) \phi\left(\frac{v}{\sigma_b}\right) dv \end{array} \right\} \\ &= \tilde{\Delta}(\sigma_b^{-1} g_b(y; z, \theta), \sigma_b^{-1} z_{by}^*, c_1^-, c_2^-) + \tilde{\Delta}(\sigma_b^{-1} z_{by}^*, \sigma_b^{-1} g_b(y+1; z, \theta), c_1^+, c_2^+). \end{aligned}$$

$$2. \ g_b(y; z, \theta) \leq g_b(y+1; z, \theta) \leq z_{by}^*.$$

$$\begin{aligned} \wp_{by}(z, \zeta) &= \sigma_b^{-1} \int_{g_b(y; z, \theta)}^{g_b(y+1; z, \theta)} \Phi\left(\frac{m_{by}^-(z, v, \theta) - \rho \frac{\sigma_d}{\sigma_b} v}{\sigma_d \sqrt{1 - \rho^2}}\right) \phi\left(\frac{v}{\sigma_b}\right) dv \\ &= \tilde{\Delta}(\sigma_b^{-1} g_b(y; z, \theta), \sigma_b^{-1} g_b(y+1; z, \theta), c_1^-, c_2^-). \end{aligned}$$

$$3. \ z_{by}^* \leq g_b(y; z, \theta) \leq g_b(y+1; z, \theta).$$

$$\begin{aligned} \wp_{by}(z, \zeta) &= \sigma_b^{-1} \int_{g_b(y; z, \theta)}^{g_b(y+1; z, \theta)} \Phi\left(\frac{m_{by}^+(z, v, \theta) - \rho \frac{\sigma_d}{\sigma_b} v}{\sigma_d \sqrt{1 - \rho^2}}\right) \phi\left(\frac{v}{\sigma_b}\right) dv \\ &= \tilde{\Delta}(\sigma_b^{-1} g_b(y; z, \theta), \sigma_b^{-1} g_b(y+1; z, \theta), c_1^+, c_2^+). \end{aligned}$$

Using indicators for whether $z_{by}^* < g_b(y+1; z, \theta)$ and $z_{by}^* > g_b(y; z, \theta)$ to cover each of

these cases gives

$$\wp_{by}(z, \zeta) = \begin{pmatrix} 1 [z_{by}^* < g_b(y+1; z, \theta)] \tilde{\Delta}(\sigma_b^{-1} \max\{z_{by}^*, g_b(y; z, \theta)\}, \sigma_b^{-1} g_b(y+1; z, \theta), c_1^+, c_2^+) \\ +1 [z_{by}^* > g_b(y; z, \theta)] \tilde{\Delta}(\sigma_b^{-1} g_b(y; z, \theta), \sigma_b^{-1} \min\{z_{by}^*, g_b(y+1; z, \theta)\}, c_1^-, c_2^-) \end{pmatrix}. \quad (\text{A.21})$$

This produces (5.3) by noting that the variables defined in (5.5) and (5.6) satisfy

$$\begin{aligned} m_1^+ &= \frac{c_1^+}{\sqrt{(c_2^+)^2 + 1}}, & m_2^+ &= -\frac{c_2^+}{\sqrt{(c_2^+)^2 + 1}}, \\ m_1^- &= \frac{c_1^-}{\sqrt{(c_2^-)^2 + 1}}, & m_2^- &= -\frac{c_2^-}{\sqrt{(c_2^-)^2 + 1}}, \end{aligned}$$

from which it follows that $\wp_{by}(z, \zeta)$ in (A.21) is equal to (5.3) in the statement of the Proposition for each b and $y \in \{1, 2\}$. ■

B Data

In the application in Section 6 we used data on purchases of women’s razor blades for the years 2004-2005 in the UK. The razor blade market is divided into three different sectors: cartridges bought with a razor (referred to as “system razors”), cartridges bought alone (referred to as “system blades”), and disposable razors. The original data consists of 7234 observations in which the main shopper was a female. Table 7 shows the observed market share of the three sectors in 2004-2005.

Sector	Total
System razors	16.80%
System blades	33.67%
Disposable razors	49.53%

Table 7: Market shares of the sectors System razors, System blades and Disposable razors in 2004-2005

For the application we concentrate on the market for system blades, where consumers buy a set of cartridges to use with a handle they already own. We define the outside option as the purchase of a disposable razor. On average in our sample, using equation (6.1) for the per-cartridge price of system places, a double blade cartridge costs £0.79 and a triple blade

cartridge costs £1.45. Table 8 provides the average unit price per cartridge for each (b, y) combination of cartridges in our sample as estimated using equation (6.1).

As discussed in Section 2, Restriction A4 requires the explanatory variables and the unobservables to be stochastically independent. In our application prices are included as exogenous variable, so the independence assumption requires price to be independent of unobservable heterogeneity. Given the relatively small cost of razor blades as indicated in Table 8 it seems reasonable to assume that this cost makes up only a small fraction of total expenditure, such that consumers are unlikely to choose where to shop for their groceries and personal care items on the basis of razor blade prices. Thus we think the assumption of price exogeneity is reasonable in this context.

Trading Company	Blade Type	
	Double Blade	Triple Blade
Gillette	£0.75 (0.0786)	£1.32 (0.1028)
Wilkinson Sword	£0.85 (0.0421)	£1.86 (0.2398)

Table 8: Average unit prices of system blade cartridges in 2004-2005. Standard deviations in parentheses.

The women’s market for reusable razors for the years 2004 – 2005 is dominated by two firms, Gillette and Wilkinson Sword, each offering razors and cartridges with two or three blades.¹⁷ Gillette’s twin-blade reusable razor, Sensor for Women, was launched in 1992 and the three-blade reusable razor, Venus, was introduced in 2001. Wilkinson Sword introduced the double blade reusable razor, Lady Protector, in 1994, and the triple blade reusable razor, Intuition, in 2003.¹⁸

We use households in which the main shopper of the household is a female between 18-50 years old who is active in the labor force. This includes women who work full time, work part time, are unemployed or not working, or in full time education.¹⁹ In the analysis we also

¹⁷Some stores offer own-label reusable razors but these were dropped from the sample as they accounted for only 4.41% of the market share for system blades used in our analysis.

¹⁸Datta (2019) and Women’s Razors, Shavers & Shaving Products UK | Wilkinson Sword (source: www.wilkinsonsword.com/en-gb/womens/). In 1992 Schick-Wilkinson Sword was formed and the Schick name was use in North America and elsewhere, while the Wilkinson Sword name was used in Europe (History of Wilkinson Sword Ltd. FundingUniverse, source: www.fundinguniverse.com/company-histories/wilkinson-sword-ltd-history/.)

¹⁹Retired individuals were excluded from the sample. In our analysis “Employed” corresponds to working some hours while “Unemployed” corresponds to either unemployed/not working or full time education.

include the marital status of the main shopper,²⁰ and a variable indicating whether there is more than one female in the household. The sample used in our analysis consists of 4842 observations. Table 9 gives summary statistics of the main shopper characteristics.

Employment status		Marital status		No of females	
Works more than 30 hours	39.20%	Married	61.83%	One female	43.04%
Works 8-29 hours	29.18%	Single	31.50%	More than one	56.96%
Works less than 8 hours	2.38%	Divorced/Widowed/Separated	6.67%		
Unemployed/not working	28.13%				
Full time education	1.12%				

Table 9: Main shopper characteristics for 2004-2005.

For each observation in the sample we observe whether they purchased cartridges for reusable razors or disposable razors and the type of blade they bought, as well as the total amount they spent, the pack size of the product they bought, the month they made the purchase, and the store in which the purchase was made. As shown in Table 10, cartridges of system blades were offered in pack sizes of 3-8 cartridges, with double blade cartridges only offered in a pack size of 5. In the calculation of the average price in equation (6.1) and of the counterfactual prices in equation (6.2) the pack sizes were redefined as small (S) if they contained 3 or 4 cartridges, medium (M) if they contained 5 or 6 cartridges, and large (L) if they contained 8 or more cartridges, as shown in Table 11.

Pack size	System Blades				Disposables
	Double Blade	Triple Blade	Gillette	Wilkinson Sword	
3		22.00%		59.92%	0.10%
4		72.80%	72.65%		23.74%
5	100.00%		24.24%	34.40%	43.67%
6		3.05%	0.97%	5.67%	3.62%
8		2.15%	2.15%		13.21%
more than 8					15.65%

Table 10: Pack sizes offered.

The counterfactual prices in equations (6.2) and (6.3) were calculated by conditioning and not conditioning on the pack size of the product purchased, respectively. As is evident from Table 10 not all blade types and not all brands offer all pack sizes. Table 13 gives the estimates of regressions (6.2) and (6.3). For the calculation of the average price in equation

²⁰The not married category in our analysis corresponds to either single or divorced/widowed/separated.

Pack size	System Blades	Total
S	69.23%	42.34%
M	29.19%	39.92%
L	1.57%	17.74%

Table 11: Pack size grouping.

(6.1) and of the counterfactual prices in equations (6.2) and (6.3), the individual shops were grouped using the company groups in Table 12.

Company group	Total
Asda	22.24%
Boots	8.57%
Co-op	0.48%
Kwiksave	0.56%
Morrisons	7.95%
Safeway	1.47%
Sainsbury's	7.97%
Savacentre	0.64%
Somerfield	0.91%
Superdrug	3.90%
Tesco	25.20%
Waitrose	0.33%
Wilkinsons	13.78%
Default ²¹	3.61%
All other	2.40%

Table 12: Company groups of stores observed.

C Alternative Specific Conditional Logit Model

In this section the results in Section 6.3 are compared to the standard alternative specific conditional logit (asclogit) model. Following McFadden (1974) choice probabilities of the asclogit model are specified as:

$$\wp_{by}(z; \zeta) = \frac{\exp(p_{by}\gamma_1 + 1[y = 2]\gamma_2 + 1[b = 2]\gamma_3 + x\beta_{by})}{\sum_{(b',y') \in \mathcal{M}_{BY}} \exp(p_{b'y'}\gamma_1 + 1[y' = 2]\gamma_2 + 1[b' = 2]\gamma_3 + x\beta_{b'y'})}$$

²¹This category contains stores such as Savers and Bodycare/Grahams, among others.

	Specification (6.2)	Specification (6.3)
Wilkinson Sword	0.4048*** (0.008)	0.3961*** (0.0081)
Triple blade	0.5043*** (0.024)	0.6948*** (0.0082)
Medium pack size	-0.1988*** (0.0236)	
Large pack size	-0.0205 (0.0277)	
Constant	0.7385*** (0.0429)	0.5456*** (0.037)

Notes: Monthly and store dummies are suppressed. Standard errors in parentheses. *** denotes significance at 1%.

Table 13: OLS estimates of (6.2) and (6.3) conditional on system blade purchases

where x denotes the same individual characteristics as in our model and p_{by} is the price for each of the four system blades alternatives each individual would have faced at the time of purchase. The specification includes both brand and high-quality (triple blade) product-specific dummy variables $1(b = 2)$ and $1(y = 2)$, respectively.²² The mean value of the utility of the outside good, disposable razors, is imposed to be zero by setting the value of each of the alternative specific covariates, i.e. the price and the quality and brand dummies, for option (0, 0) to zero. The estimated coefficients using only specification (6.2) are given in Table 14. These results were then used to calculate the predicted own and cross price elasticities, which following e.g. Cameron and Trivedi (2009) take the form:

$$\eta_{ibykl} = \begin{cases} \gamma_1 [1 - \rho_{by}(z_i; \beta, \gamma)] p_{iby} & \text{if } (b, y) = (k, l), \\ -\gamma_1 \rho_{kl}(z_i; \beta, \gamma) p_{ikl} & \text{if } (b, y) \neq (k, l). \end{cases}$$

Table 15 gives the estimates of the average and the 0.2, 0.5 and 0.8 quantiles of household own- and cross-price elasticities. For example, the mean own price elasticity of a Gillette double blade cartridge, (1, 1), is -0.8872 , while the mean cross-price elasticity of the probability of choosing any other alternative with respect to the price of the Gillette double blade is 0.0874.

²²Constant terms are excluded from the alternative-specific utilities as these would not be identified with the inclusion of the brand and quality dummy variables.

	Specification (6.2)				
	γ	β_{11}	β_{12}	β_{21}	β_{22}
Price	-1.2552*** (0.125)				
Triple Blade	0.5211*** (0.1582)				
Wilkinson Sword	-1.3581*** (0.1651)				
Age31-40		-0.2304** (0.1172)	-0.1439* (0.0854)	-0.3202* (0.1831)	0.2973** (0.1443)
Age41-50		-0.6835*** (0.1567)	-0.43*** (0.1014)	-0.1935 (0.2013)	0.1617 (0.164)
Married		-0.1051 (0.1066)	0.168** (0.0769)	0.1166 (0.1589)	0.2867** (0.1259)
Employed		-0.3357*** (0.1034)	0.3436*** (0.0812)	0.1396 (0.1558)	0.8254*** (0.1416)
Females		-0.7163*** (0.1077)	-0.0819 (0.0754)	0.2128 (0.1572)	-0.2498** (0.1199)

Notes: Estimates of Alternative Specific Conditional Logit Regression. Standard errors in parentheses. *significant at 10%; ** significant at 5%; *** significant at 1%

Table 14: Alternative specific conditional logit regression estimates 2004-2005

	Own Price Elasticities			Cross Price Elasticities		
	Specification (6.2)					
(1, 1)						
mean		-0.8872			0.0874	
quantiles	-1.0055	-0.9075	-0.7555	0.0514	0.0797	0.1207
(1, 2)						
Mean		-1.2302			0.3584	
Median	-1.3571	-1.2364	-1.0943	0.3131	0.3563	0.405
(2, 1)						
mean		-1.4075			0.0562	
quantiles	-1.5374	-1.4533	-1.2661	0.0489	0.0546	0.0685
(2, 2)						
mean		-1.9678			0.1478	
quantiles	-2.1031	-1.9762	-1.8039	0.094	0.1551	0.1953

Notes: Estimates of Own and Cross Price Elasticities as the average and the 0.2, 0.5 and 0.8 quantiles across all observations

Table 15: Estimated means and 0.2, 0.5 and 0.8 quantiles of household elasticities 2004-2005

D Monte Carlo Experiments

In this Appendix we report the results of Monte Carlo experiments to compare the finite sample performance of Wald and profile QLR confidence intervals for inference with our model.²³ For these experiments we generated data from the partially ordered probit model, with the number of parameters matching those employed in the subsequent application. There were five individual-specific dummy variables with corresponding coefficients $\beta_{b1}, \dots, \beta_{b5}$ for each $b = 1, 2$. Each product offering had a price p_{by} generated differently in each of the three data generation processes (DGPs) – referred to as DGP1, DGP2, and DGP3 – as described below. The linear-in-price specification described by (4.6) was used.

To simulate data population parameter values were set as follows.

$$\begin{aligned} \gamma_1 = 1, \quad \gamma_2 = 0.8, \quad \delta_1 = -1.5, \quad \delta_2 = -1.2, \quad \rho = 0.5, \quad \sigma = 1, \\ \beta_1 = (1.3, 0.3, -0.1, -0.3, 0.7)', \quad \beta_2 = (1.0, 0.3, -0.1, -0.3, 0.7)'. \end{aligned} \tag{D.1}$$

In our application the first two components of X , X_1 and X_2 , are dummy variables indicating whether age of a female shopper is from 31-40, or 41-50, with 18-30 denoting the base category. These variables were drawn such that $\Pr[X_1 = 1] = 0.426$, $\Pr[X_2 = 1] = 0.234$, and $X_1 + X_2 \leq 1$. The remaining components of X are dummy variables for marriage, employment, and a variable “more_females” indicating the presence of more than one female in the household. These were generated from the Bernoulli distribution with parameters 0.4, 0.85, and 0.554, respectively. The age variable from which X_1 and X_2 were generated and the remaining dummy variables were drawn independently of one another.

Prices $(p_{11}, p_{12}, p_{21}, p_{22})$ were generated independently of X , as follows. First, for each DGP and for each observation a vector ε was drawn from the bivariate normal distribution with each component having mean zero and variance one, with correlation 0.25. In DGP1 prices p_{11} and p_{21} were generated independently, and uniformly on the intervals $[1, 4]$ and $[1.35, 2.15]$, respectively. Prices p_{12} and p_{22} were then set to $p_{12} = p_{11} + \varepsilon_1$ and $p_{22} = p_{21} + \varepsilon_2$. In this DGP, prices p_{12} and p_{22} thus both have positive density on all of \mathbb{R} conditional on all other variables. This implies that there is positive probability that the price of the higher quality product for either brand b undercuts the price of the lower quality product, i.e. $p_{b2} < p_{b1}$, as could happen under a promotion for the higher quality product. In such cases the conditional probability of choosing the lower quality product for the brand will be

²³Code for the Monte Carlo experiments is available at <https://sites.google.com/site/amr331/home/por-code>.

zero. Moreover, the large support for both p_{12} and p_{22} imply that this happens with positive probability for both brands, in which case the choice problem reduces to a simple multinomial choice setting between each brand’s higher quality product and the outside option. Thus, the large support of these variables, artificial though it may be, demonstrates a setting in which point identification may be achieved. This is borne out in the Monte Carlo simulations below.

In practice prices will not have support on the entire real line, and neither DGP2 nor DGP3 have this feature. In DGP2 p_{11} and p_{21} were generated independently from the uniform distribution on $[1, 2]$ and $[1.35, 2.15]$, respectively, and each p_{b2} was set to $p_{b1} + \max\{1, \min\{|\varepsilon_b|, 2\}\}$. Thus the higher quality product for each brand always has a higher price than the lower quality product of that brand. Moreover, all prices have continuous, but bounded support. In DGP3 p_{11} and p_{21} were generated the same way, but the term added on to p_{b1} to determine p_{b2} was instead rounded to the nearest integer (which was either one or two) before adding. In this design prices again have bounded continuous support, but for each b the conditional support of p_{b2} given p_{b1} is discrete.

With variables X and prices $P = (p_{11}, p_{12}, p_{21}, p_{22})$ generated as described above, and unobservables $V = (V_1, V_2)$ drawn from the bivariate normal distribution with parameters ρ and σ , data (b_i, y_i, x_i, p_i) were generated with each (b_i, y_i) solving the individual choice problem with the corresponding (x_i, p_i, v_i) and utility parameters as in (D.1). The expression (5.3) obtained for choice probabilities in Proposition 2 was used in the log-likelihood function based on n observations in each experiment, with $n \in \{200, 500, 1000, 2000\}$. In preliminary investigation, choice probabilities computed using (5.3) conditional on several values of observable variables were compared to those obtained using the integral formula (5.1) and those obtained by simulation, and these were all found to be in close agreement up to negligible computation difference.

In order to compare the empirical coverage frequencies of the Wald and QLR confidence intervals in each repetition of our Monte Carlo simulations, we carried out the following steps. First, the R package Ghalanos and Stefan (2015) was used to minimize $-L_n(\zeta)$ with respect to the full parameter vector ζ , producing an optimizing vector $\hat{\zeta}_{ML}$ and an optimal value L_n^* . To ensure accuracy 500 randomly generated starting values were employed using the function `gosolnp`.²⁴ The optimization routine also returned a numerical approximation

²⁴In Monte Carlo simulations the population parameter value was also used as an additional starting value. The number of randomly generated starting values was chosen based on experimentation; increasing it further was not found to be beneficial.

to the Hessian at the optimal value, and this was used to compute asymptotic variance estimators $\widehat{var}(\hat{\zeta}_k)$ for each component of the maximum likelihood estimator $\hat{\zeta}_{ML}$. There is no guarantee that ζ is point identified. If it is point identified confidence intervals for each parameter component based on the usual asymptotic normal approximation should be expected to perform well, but if it is not point identified the classical theory will be invalid. Nonetheless QLR confidence intervals can sometimes remain valid under partial identification as shown by Chen, Christensen, and Tamer (2018).²⁵

In Monte Carlo experiments where the true population parameter is known, the same routine was also used to compute the maximum likelihood estimator taking the values of ρ and σ fixed at their population values. In our application ρ and σ are not known, so this approach is infeasible. However, with these parameters known, the rest of the parameters are point identified under mild conditions on the variation in observable payoff shifters. Thus, these confidence intervals should be expected to perform well, and in our Monte Carlo experiments this was indeed the case. In our reported results, we refer to confidence intervals for each ζ_k based on the maximum likelihood estimator $\hat{\zeta}_k$ plus or minus $1.96 \times \sqrt{\widehat{var}(\hat{\zeta}_k)}$ as Wald confidence intervals, considering both cases where maximum likelihood was carried out with ρ and σ fixed at their population values, as well as with ρ and σ as additional parameters to estimate. Only the latter approach is feasible if the population values of ρ and σ are unknown. In order to compute Monte Carlo coverage frequencies of $\mathcal{C}_{\alpha,k}^{LR}$ for ζ_k we additionally computed $Q_{k,n}(\zeta_k)$ as defined in (6.5) at the population value of ζ_k and checked whether $Q_{k,n}(\mu) \leq \chi_{1,\alpha}^2$ in each simulation.²⁶

The empirical coverage frequency of the three different procedures for DGPs 1-3 out of 1000 Monte Carlo repetitions for each sample size are reported in Tables 16, 17, and 18. The target coverage level in each case was 0.95. All procedures performed reasonably well, although coverage probabilities for the distributional parameters are substantially below the nominal level at smaller sample sizes. This is particularly so for the Wald procedure in which ρ and σ are treated as unknown parameters, while under-coverage from the QLR procedure is less severe at smaller sample sizes.

²⁵All reported coverage frequencies are for individual parameters ζ_k . Chen, Christensen, and Tamer (2018) provide sufficient conditions for the QLR confidence interval to achieve asymptotic coverage for the identified set for ζ_k , which also guarantee at least as high a level of asymptotic coverage for ζ_k .

²⁶Our implementation employed the `solnp` function from the package Ghalanos and Stefan (2015) to compute $\sup_{\zeta \in \Upsilon: \zeta_k = \mu} L_n(\zeta)$. In the constrained optimizations conducted with `solnp`, the population value of ζ_{-k} was used as a starting value to speed up computations. In terms of coverage frequency for ζ_k , this was found to produce the same results in a subset of the Monte Carlo iterations attempted when as many of 500 random starting values were used.

DGP 1	Realized Coverage Probability, $n = 200, 500, 1000, 2000$											
Parameter	Wald (ρ, σ known)				Wald (ρ, σ unknown)				QLR			
γ_1	.935	.939	.954	.950	.952	.951	.950	.944	.929	.941	.947	.935
γ_2	.944	.954	.950	.961	.934	.956	.951	.937	.938	.949	.947	.934
δ_1	.927	.952	.951	.958	.937	.960	.963	.954	.943	.955	.953	.950
δ_2	.940	.949	.954	.957	.907	.945	.947	.939	.932	.948	.941	.937
β_{11}	.957	.952	.951	.949	.951	.956	.946	.950	.942	.948	.939	.942
β_{12}	.947	.961	.944	.949	.950	.963	.955	.945	.933	.954	.949	.933
β_{13}	.959	.953	.951	.949	.961	.961	.953	.950	.947	.956	.944	.943
β_{14}	.954	.948	.959	.952	.962	.950	.960	.953	.942	.941	.959	.946
β_{15}	.954	.956	.957	.952	.949	.968	.961	.956	.945	.962	.956	.948
β_{21}	.960	.963	.958	.957	.922	.943	.938	.942	.938	.946	.946	.941
β_{22}	.942	.956	.947	.953	.938	.961	.961	.952	.931	.950	.947	.941
β_{23}	.953	.949	.961	.959	.969	.955	.964	.963	.951	.946	.951	.951
β_{24}	.964	.946	.941	.962	.930	.940	.926	.951	.945	.937	.931	.942
β_{25}	.957	.957	.956	.961	.924	.952	.957	.941	.955	.955	.953	.938
ρ	–				.899	.928	.934	.944	.933	.940	.940	.941
σ	–				.908	.941	.956	.927	.942	.947	.957	.923

Table 16: Monte Carlo coverage frequencies out of 1000 simulations for sample sizes $n = 200, 500, 1000, 2000$ for DGP1, as described in the text.

DGP 2	Realized Coverage Probability, $n = 200, 500, 1000, 2000$											
Parameter	Wald (ρ, σ known)				Wald (ρ, σ unknown)				QLR			
γ_1	.956	.956	.957	.944	.936	.961	.958	.942	.936	.951	.955	.939
γ_2	.955	.947	.954	.948	.836	.905	.926	.943	.899	.935	.942	.948
δ_1	.962	.960	.955	.957	.923	.957	.958	.963	.939	.963	.956	.959
δ_2	.961	.945	.951	.950	.812	.897	.923	.940	.892	.921	.941	.942
β_{11}	.942	.962	.953	.955	.928	.947	.942	.952	.939	.958	.946	.959
β_{12}	.942	.950	.953	.944	.951	.953	.951	.942	.943	.953	.953	.942
β_{13}	.953	.950	.951	.956	.951	.952	.952	.950	.948	.949	.949	.949
β_{14}	.949	.950	.946	.963	.961	.956	.948	.957	.949	.955	.947	.956
β_{15}	.940	.958	.938	.957	.927	.938	.934	.949	.934	.952	.938	.948
β_{21}	.942	.956	.958	.941	.852	.905	.926	.934	.912	.939	.945	.937
β_{22}	.948	.957	.947	.941	.900	.937	.940	.946	.927	.947	.945	.943
β_{23}	.931	.960	.947	.935	.949	.966	.959	.952	.927	.954	.951	.943
β_{24}	.946	.939	.947	.962	.894	.933	.932	.962	.925	.947	.940	.953
β_{25}	.940	.948	.941	.941	.866	.915	.932	.938	.925	.933	.942	.936
ρ	–				.710	.840	.885	.904	.899	.929	.938	.942
σ	–				.843	.910	.928	.955	.896	.929	.940	.950

Table 17: Monte Carlo coverage frequencies out of 1000 simulations for sample sizes $n = 200, 500, 1000, 2000$ for DGP2, as described in the text.

DGP 3	Realized Coverage Probability, $n = 200, 500, 1000, 2000$											
Parameter	Wald (ρ, σ known)				Wald (ρ, σ unknown)				QLR			
γ_1	.948	.958	.939	.952	.954	.964	.948	.950	.947	.958	.947	.947
γ_2	.949	.955	.945	.943	.911	.930	.951	.950	.929	.946	.944	.952
δ_1	.954	.950	.950	.956	.955	.959	.956	.958	.951	.963	.956	.960
δ_2	.960	.944	.948	.948	.907	.932	.953	.953	.923	.931	.957	.941
β_{11}	.936	.945	.948	.952	.937	.953	.958	.954	.938	.954	.958	.951
β_{12}	.949	.950	.956	.948	.956	.954	.959	.940	.952	.948	.953	.937
β_{13}	.948	.936	.954	.951	.955	.944	.950	.950	.951	.941	.947	.947
β_{14}	.939	.953	.954	.956	.940	.950	.954	.960	.935	.950	.951	.955
β_{15}	.941	.952	.943	.946	.941	.957	.954	.959	.944	.953	.952	.956
β_{21}	.938	.945	.953	.954	.917	.948	.953	.943	.934	.953	.950	.940
β_{22}	.942	.956	.943	.948	.947	.961	.952	.947	.946	.956	.950	.951
β_{23}	.945	.959	.942	.944	.968	.965	.959	.946	.931	.953	.942	.944
β_{24}	.948	.945	.950	.963	.923	.950	.943	.960	.936	.948	.944	.961
β_{25}	.938	.945	.946	.948	.931	.933	.949	.947	.950	.941	.954	.946
ρ	–				.795	.897	.925	.923	.916	.943	.947	.937
σ	–				.918	.939	.957	.949	.927	.932	.950	.943

Table 18: Monte Carlo coverage frequencies out of 1000 simulations for sample sizes $n = 200, 500, 1000, 2000$ for DGP3, as described in the text.