

NBER WORKING PAPER SERIES

FOUNDATIONS OF WELFARE ECONOMICS AND PRODUCT MARKET APPLICATIONS

Daniel McFadden

Working Paper 23535

<http://www.nber.org/papers/w23535>

NATIONAL BUREAU OF ECONOMIC RESEARCH

1050 Massachusetts Avenue

Cambridge, MA 02138

June 2017

I am indebted to Kenneth Train, Professor of Economics, University of California, Berkeley, who made major contributions to the contents of this paper, including the welfare calculus formulas given in Sections 5 and 7, the application given in Section 8, and Appendix C. I also thank Moshe Ben-Akiva, Andrew Daly, Mogens Fosgerau, Garrett Glasgow, Stephane Hess, Armando Levy, Douglas MacNair, Charles Manski, Rosa Matzkin, Kevin Murphy, Frank Pinter, Joan Walker, and Ken Wise for useful suggestions and comments. The views expressed herein are those of the author and do not necessarily reflect the views of the National Bureau of Economic Research.

The author has disclosed a financial relationship of potential relevance for this research. Further information is available online at <http://www.nber.org/papers/w23535.ack>

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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2017 by Daniel McFadden. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Foundations of Welfare Economics and Product Market Applications
Daniel McFadden
NBER Working Paper No. 23535
June 2017
JEL No. D11,D12,D60,D61,K13,L51

ABSTRACT

A common problem in applied economics is to determine the impact on consumers of changes in prices and attributes of marketed products as a consequence of policy changes. Examples are prospective regulation of product safety and reliability, or retrospective compensation for harm from defective products or misrepresentation of product features. This paper reexamines the foundations of welfare analysis for these applications. We consider discrete product choice, and develop practical formulas that apply when discrete product demands are characterized by mixed multinomial logit models and policy changes affect hedonic attributes of products in addition to price. We show that for applications that are retrospective, or are prospective but compensating transfers are hypothetical rather than fulfilled, a Market Compensating Equivalent measure that updates Marshallian consumer surplus is more appropriate than Hicksian compensating or equivalent variations. We identify the welfare questions that can be answered in the presence of partial observability on the preferences of individual consumers. We examine the welfare calculus when the experienced-utility of consumers differs from the decision-utility that determines market demands, as the result of resolution of contingencies regarding attributes of products and interactions with consumer needs, or as the result of inconsistencies in tastes and incomplete optimizing behavior. We conclude with an illustrative application that calculates the welfare impacts of unauthorized sharing of consumer information by video streaming services.

Daniel McFadden
University of California, Berkeley
Department of Economics
508-1 Evans Hall #3880
Berkeley, CA 94720-3880
and NBER
mcfadden@econ.berkeley.edu

1. INTRODUCTION

A common problem in applied economics is assessment of the welfare consequences for consumers of policies/scenarios that regulate markets for products, or correct for past product defects or misrepresentations. Examples are (1) *prospective* regulation of information provided on coverage and costs in insurance contracts and other financial instruments such as mortgages, and *retrospective* redress of harm from failures to properly disclose information; (2) harm from environmental damage to recreation facilities such as ocean beaches; (3) safety regulation of consumer products such as automobile air bags, mobile phones, and privacy protection in video streaming services, or redress of harm from safety defects; and (4) evaluation of overall market performance; e.g., the prospective benefit of blocking a merger of dominant suppliers, or retrospective harm from collusion or restraints on entry. This paper reexamines the foundations of welfare analysis for these applications, and provides a practical framework for analysis that rests on these foundations.

Measuring changes in consumer well-being from policies that affect the availability, prices, and/or attributes of goods and services has been a central concern of economics from its earliest days. Adam Smith (1776) observed that “haggling and bargaining in the market” would achieve “rough equality” between value in use and value in exchange. Working at the fringes of mainstream economics, Jules Dupuit (1844) and Hermann Gossen (1854) deduced further that consumers exhibiting diminishing marginal utility would achieve maximum utility by equalizing marginal utilities per unit of expenditure across goods. Dupuit was remarkably prescient, recognizing that if the marginal utility of income is constant, then the demand curve for a commodity (illustrated in Figure 1) is a marginal utility curve, so that the area to the left of this demand curve between the prices established by scenarios labeled *a* and *b* gives a measure (in monetary units) of “relative utility”. Dupuit’s measure was later named *Marshallian Consumer Surplus* (MCS); see Alfred Marshall (1890, III.IV.2-8).

To rephrase these propositions in microeconomic language, suppose the consumer maximizes a utility function $U(q_0, q_1)$ of two goods subject to a budget constraint $I = p_0 q_0 + p_1 q_1$, where I is income and p_0 and p_1 are the goods prices. Let $q_0 = D_0(I, p_0, p_1)$ and $q_1 = D_1(I, p_0, p_1) \equiv (I - p_0 D_0(I, p_0, p_1))/p_1$ denote the demands that come out of this maximization, and let $V(I, p_0, p_1) = U(D_0(I, p_0, p_1), D_1(I, p_0, p_1)) \equiv \max_{q_0} U(q_0, (I - p_0 q_0)/p_1)$ denote the resulting *indirect* (or maximized) utility. Then $\partial V/\partial I = (1/p_1)\partial U/\partial q_1 + [\partial U/\partial q_0 - (p_0/p_1)\partial U/\partial q_1](\partial q_0/\partial I) \equiv (1/p_1)\partial U/\partial q_1$ since the first-order condition for maximization makes the term in square brackets zero. Rearranging this expression for $\partial V/\partial I$ gives Smith’s proposition: “value in exchange” $\equiv p_1 = (\partial U/\partial q_1)/(\partial V/\partial I) \equiv$ “value in use” or marginal utility per unit of good 1 measured in money units. Further, $\partial V/\partial p_1 = - (D_1(I, p_0, p_1)/p_1)\partial U/\partial q_1 \equiv - D_1(I, p_0, p_1) \cdot (\partial V/\partial I)$, giving Roy’s identity when both sides are divided by $(\partial V/\partial I)$, or giving Dupuit’s link between demand and changes

in utility when $(\partial V/\partial I)$ is constant. A measure in monetary units of the consumer's change in well-being is the difference in indirect utilities in the two scenarios, scaled by the marginal utility $\mu = \partial V(I, p_0, p_{1a})/\partial I$ of income at scenario a . We term this the *market compensating equivalent* (MCE),

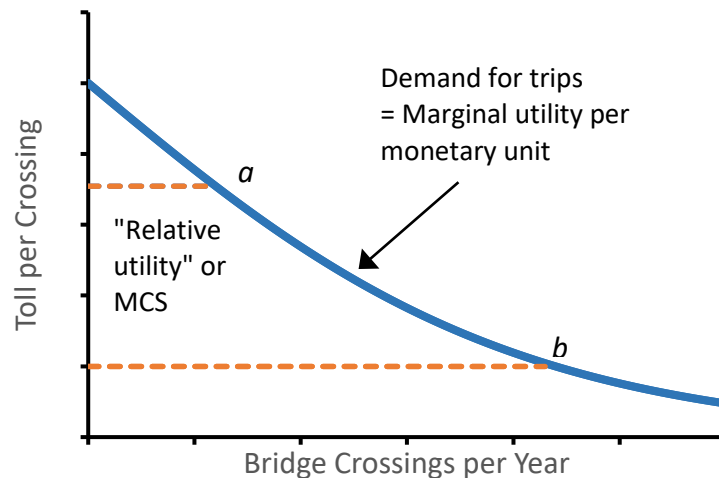
$$(1) \quad MCE = [V(I, p_0, p_{1b}) - V(I, p_0, p_{1a})]/\mu.$$

Given p_0 , the marginal utility of income is independent of I and p_1 if and only if the indirect utility function has an additively separable form $V(I, p_0, p_1) = \mu I/p_0 - G(p_1/p_0)$ for some function G and constant μ , in which case Roy's identity establishes that the demand for good 1, $D_1(I, p_0, p_1) = G'(p_1/p_0)/\mu$, is independent of income. This is the case where the "relative utility" or "consumer surplus" of Dupuit and Marshall coincide with MCE, and these measures are given by integrating Roy's identity between prices p_{1b} and p_{1a} .

$$(2) \quad MCS = \int_{p_{1b}}^{p_{1a}} D_1(I, p_0, p_1) dp_1 = MCE$$

Even if the marginal utility of income is not necessarily constant, applying the first mean value theorem for integrals to Roy's identity establishes that $\int_{p_{1b}}^{p_{1a}} D_1(I, p_0, p_1) dp_1 = [V(I, p_0, p_{1b}) - V(I, p_0, p_{1a})]/\mu^*$ for some intermediate value μ^* of the marginal utility of income, so that MCE equals MCS scaled by a factor μ^*/μ . Obviously, when the marginal utility of income is constant, $\mu^*/\mu = 1$.

Figure 1. Dupuit's Calculation of Relative Utility



Dupuit's idea of solving the *inverse problem*, recovering utility from demand, was brought into mainstream economics at the end of the 19th century by William Stanley Jevons (1871), Francis Edgeworth (1881), Alfred Marshall (1890), Vilfredo Pareto (1906), and Eugen Slutsky (1915). MCS became the accepted measure of the change in consumer well-being. However, John Hicks (1939) observed that when the marginal utility of income μ

$= \partial V(I, p_0, p_1) / \partial I$ is not constant, MCS is a measure of relative utility whose scaling is specific to the consumer and circumstance. Hicks introduced two closely related alternative measures to remedy this defect: *Hicksian Contingent Valuation* (HCV), the net decrease in scenario *b* income that equates utility in the two scenarios, and *Hicksian Equivalent Variation* (HEV), the net increase in scenario *a* income that equates utility in the two scenarios.

Importantly, the MCE, HCV, and HEV measures correspond to different consumer choice assumptions: The HCV measure assumes that the transfer is fulfilled in scenario *b* before the consumer makes a choice in that scenario, while the HEV measure assumes the transfer is fulfilled in scenario *a* before the scenario *a* choice. The MCE measure assumes that choices are made under actual market and income conditions, without compensation, in each scenario, and that the post-choice transfer is determined after this as a remedy for the utility gain or loss caused by the change in scenario. Then, HCV is appropriate for prospective welfare analysis when the transfer is made before choice in scenario *b*, and HEV when the transfer is made before choice in scenario *a*. However, for retrospective welfare analysis where the objective is to redress past harm, or for prospective analysis where the transfers are hypothetical and not fulfilled, MCE is an appropriate measure of what it takes to “make the consumer whole” relative to a benchmark scenario subsequent to the choices the consumer did make or would have made in the uncompensated “as is” and “but for” scenarios. MCE is also appropriate for assessment of residual gains and losses subsequent to prospective analysis where an inexact compensation scheme is fulfilled.

HCV and HEV are often defined as areas to the left of income-compensated demand curves (i.e., demands with income adjusted as price changes to keep utility fixed at the scenario *a* or scenario *b* levels, respectively), as in Figure 2.² However, their definition in terms of the indirect utility function, solutions to $V(I - \text{HCV}, p_0, p_{1b}) = V(I, p_0, p_{1a})$ and $V(I, p_0, p_{1b}) = V(I + \text{HEV}, p_0, p_{1a})$, are revealing. Applying the mean value theorem, they satisfy

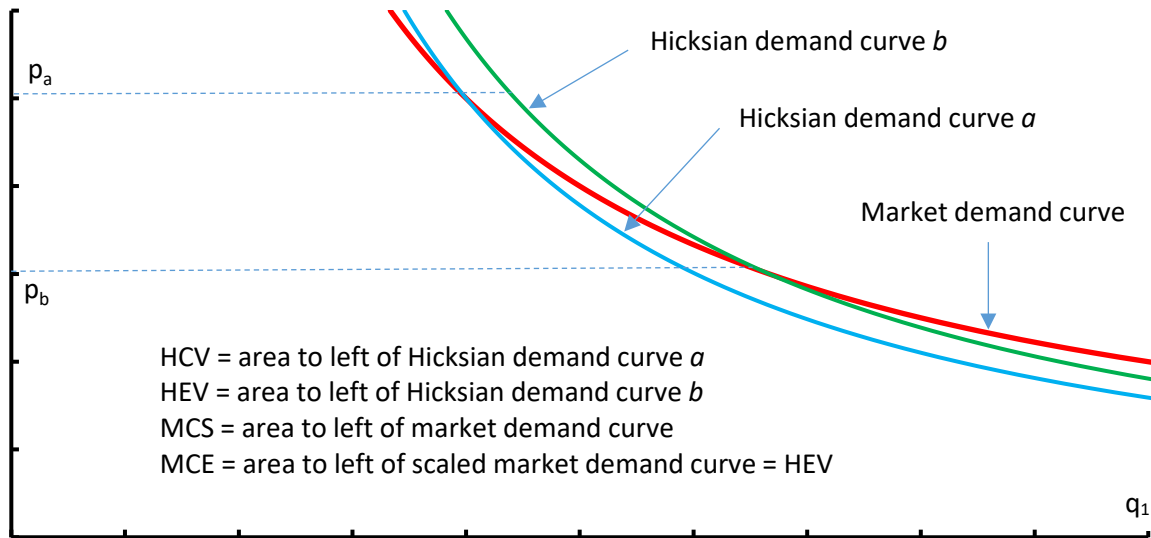
$$(3) \quad \begin{aligned} \text{HCV} &= [V(I, p_0, p_{1b}) - V(I, p_0, p_{1a})] / \mu' \\ \text{HEV} &= [V(I, p_0, p_{1b}) - V(I, p_0, p_{1a})] / \mu'' \end{aligned}$$

where μ' and μ'' are some intermediate values of the marginal utility of income. Then, these two measures, the MCE measure from (1), and MCS are all proportional to the difference in utilities of the two scenarios, and differ only in scalings by marginal utilities of income valued at different points. Obviously, if the marginal utility of income is constant, then MCE, HCV, HEV, and MCS are identical, and in applications where the marginal utility of

² The figure comes from a log-linear utility function $u = 0.8 \log q_0 + 0.2 \log q_1$. With a budget $50 = I = p_0 q_0 + p_1 q_1$, $p_0 = 1$, $p_{1a} = 5$, and $p_{1b} = 3$, the corresponding money-metric indirect utility based on scenario *a* is $u = I \cdot (p_{1a}/p_1)^{0.2}$, and the welfare measures are $\text{MCE} = 5.38 = \text{HEV}$, $\text{HCV} = 4.86$, and $\text{MCS} = 5.11$.

income varies little, they will be close approximations. MCE has a closed form when the indirect utility function is known, a computational advantage over HCV and HEV.

Figure 2. Consumer Surplus Measures for a Continuously Supplied Good



The Hicksian welfare apparatus was refined by Kaldor (1939) and Scitovsky (1942), and the Kaldor-Hicks criterion for public policy decisions states that if a policy change from a to b has a positive aggregate HCV, then this aggregate can be distributed so that the change is a Pareto improvement; the same argument applies to HEV and MCE in appropriate circumstances. The MEC, HDV, and HEV measures are mute on the issue of income inequality and socially desirable redistribution, but can be adapted for this purpose by applying social weights that depend on socially weighted individual marginal utilities of income, estimated from observed consumer socioeconomic characteristics. Samuelson (1947) and Hurwicz and Uzawa (1971) updated the Hicksian analysis using modern consumer theory, and their approach has been adapted to consumers making discrete choices by Diamond and McFadden (1974), Small and Rosen (1981), McFadden (1981, 1994, 1999, 2004, 2012, 2014), Yatchew (1985), and Zhou et al (2012). For the most part, this literature assumes that consumers are strictly neoclassical utility maximizers, with self-interest defined narrowly to include only personally purchased and consumed goods. Mostly, social motives are ignored and no allowance is made for ambiguities and uncertainties regarding tastes, budgets, hedonic attributes of goods and services, the reliability of transactions, or the consistency and completeness of preference maximization, and there is no distinction between the decision-utility postulated to determine market behavior and the experienced-utility of outcomes. Public and environmental

goods are incorporated only if they have active margins that allow them to be valued from market behavior. The market demand functions of individual consumers are assumed to be completely observed, and consumers fully informed about policy regimes, so that utility can be recovered from the demand behavior it produces and the compensating transfers can be calculated and fulfilled exactly each consumer. The primary focus of welfare theory has been prospective, assuming that compensating transfers are fulfilled before consumer choices are made. The analysis has been fundamentally static, with the consumer pictured as making a once-and-for-all utility-maximizing choice for contingent deliveries of market goods, even if resolution of uncertainties and fulfillment of contracts extend over time; as in Debreu (1959). Analysis typically starts from prespecified scenarios, although in retrospective applications there are often substantive questions regarding the nature of the “but for” scenario, particularly when the “as is” scenario leads to experienced utility different from decision utility. Two further assumptions are tacit in most practical welfare calculations: First, policy scenario differences are limited in scope and magnitude, so that after accounting for a few major margins, general equilibrium effects can be neglected. Second, if compensating transfers are incomplete within a class of consumers, conducted say using a simple formula such as uniform transfers rather than an exact consumer-by-consumer calculation, the loss in social welfare from this imperfect redistribution can be neglected relative to the aggregate welfare change for the class.

We review these assumptions. Section 2 gives a foundation in consumer theory for the welfare calculus, with explicit treatment of discrete alternatives and their hedonic attributes. Section 3 restates the welfare measures in Section 1 for general applications, using the consumer theory of Section 2. Section 4 distinguishes retrospective and prospective policy applications of the welfare calculus. Section 5 discusses partial observability of individual consumer preferences, and its implications for welfare measurement and aggregation. Section 6 distinguishes decision-utility and experienced-utility foundations for calculation of well-being. Section 7 gives computational formulas for common policy problems. Section 8 contains an illustrative empirical application. Appendices collect relevant mathematical results on approximation, give properties of extreme-value distributed random variables, and give R-code for discrete welfare calculations.

2. CONSUMER FOUNDATIONS

A common starting assumption for welfare analysis is that consumers have “nice” demand functions that allow recovery of indirect utility. For example, Hurwicz and Uzawa (1971) give local and global sufficient conditions for recovery of money-metric indirect utility³ when the market demand function is single-valued and smooth; see also Katzner (1970) and Border (2014). Another approach, originating in the revealed preference analysis of Samuelson (1948), Houthakker (1950), and Richter (1966), gives necessary and sufficient conditions for recovery of a preference order whose maximization yields the market demand function; Afriat (1967) and Varian (2006) provide constructive methods for recovery of utility under some conditions. Technical difficulties arise because quite strong smoothness and curvature conditions on utility are needed to assure smoothness properties on market demand, while preferences recovered from upper hemicontinuous demand functions are not necessarily continuous; see Peleg (1970), Rader (1973), Conniffe (2007). This section gives a restatement of the consumer theory behind welfare measurement, with extensions that include a “no local cliffs” Lipschitz continuity axiom on the preference map that avoids the Peleg-Rader problem and guarantees representation of preferences by utility, expenditure, indirect utility, and demand functions that satisfy bi-Lipschitz⁴ conditions in economic variables. These results facilitate practical welfare measurement, and are of independent interest. Readers may find it useful to refer to Table 1 for notation, and consult as needed the technical material in the remainder of this section.

Suppose consumers face scenarios $m = a, b$, and a universe of possible discrete alternatives indexed by a finite set $J \equiv \{0, \dots, J\}$. Let $J_m \subseteq J$ denote the set of alternatives available in the market under policy m , with $|J_m|$ elements, and characterize then by *real* prices p_{jm} in a compact interval $P = [0, p^U]$ with $p^U > 0$, and *observed* hedonic attributes z_{jm} in a compact finite-dimensional set Z . Let \mathbf{p}_m and \mathbf{z}_m denote the vectors of p_{jm} and z_{jm} for $j \in J_m$. Assume that there are “baseline” or “no purchase” alternatives that are always available and are unaffected by policy; e.g., a “no purchase” j is assigned $p_{jm} = 0$ and $z_{jm} = 0$ by convention. Market goods supplied in continuous quantities are described by commodity vectors $\mathbf{q} \subseteq Q = [0, \mathbf{q}^U]$ with $0 \ll \mathbf{q}^U$, a bounded rectangle in n -dimensional space, with *real* market prices $\mathbf{r} \in R = [\mathbf{r}^L, \mathbf{r}^U]$, a commensurate bounded rectangle with $0 \ll \mathbf{r}^L \ll \mathbf{r}^U$. We assume that Z is a finite union of disjoint rectangles; this avoids technical complications and covers applications where measured attributes either vary continuously in some interval or take on a finite number of discrete levels.

³ Indirect utility is money-metric if the marginal utility of (real) income in a baseline scenario remains one as income changes.

⁴ An increasing function is bi-Lipschitz if its left and right derivatives are bounded away from 0 and $+\infty$.

Table 1. Notation

$m = a, b$	"As-Is"/baseline policy/scenario a and "But-For"/counterfactual policy/scenario b
$s \in S$	Finite-dimensional vector in a compact set S describing observed demographics and history of the decision-maker
$I_m \in [I^L, I^U]$	Consumer real income, in an interval $[I^L, I^U]$, with $0 < I^L < I^U < +\infty$, in scenario m
$j \in J_m \subseteq J = \{0, \dots, J\}$	Mutually exclusive discrete choices (e.g., "products"), including "benchmark" or "no-purchase" alternatives that are not affected by policy change
$z_{jm} \in Z$	Vector \mathbf{z}_m of observed hedonic attributes z_{jm} for alternatives $j \in J_m$ in scenario m , in a compact finite-dimensional set Z
$\mathbf{q} \in Q' \subseteq Q$	Vector of the goods and services that are supplied in continuous quantities, in a finite rectangle $Q \equiv [0, \mathbf{q}^U]$ in n -dimensional Euclidean space, or in a subrectangle $Q' = [0, \mathbf{q}^A]$, where \mathbf{q}^A is an upper bound on vectors that are affordable, $0 \ll \mathbf{q}^A \ll \mathbf{q}^U$
$\mathbf{w}_{jm} = (\mathbf{q}, z_{jm}) \in W$	Consumption vector given discrete choice j in scenario m , in $W = Q \times Z$ or in $W' = Q' \times Z$
$p_{jm} \in P$	Real price, in a compact interval $P = [0, p^U]$ with $p^U > 0$, of discrete product j in scenario m ; \mathbf{p}_m is the vector with components p_{jm} for $j \in J_m$
$\mathbf{r}_m \in R$	Finite-dimensional vector in a rectangle $R = [\mathbf{r}^L, \mathbf{r}^U]$, with $0 \ll \mathbf{r}^L \ll \mathbf{r}^U \ll +\infty$, of real prices of the goods and services that are available in continuous quantities; benchmark \mathbf{r}_a
$\mathbf{r}_m \cdot \mathbf{q}_m + p_{jm} \leq I_m$	Budget constraint given discrete alternative j in scenario m
$\succsim \in H$	A field H of complete transitive reflexive preference preorders \succsim on $Q \times Z$, represented by sets $G(\succsim) \subseteq W \times W$ with $(w', w'') \in G(\succsim) \Leftrightarrow w' \succsim w''$
$U(\mathbf{q}, z, \succsim)$	A direct utility function conditioned on choice j in scenario m , defined on $Q' \times Z \times H$ as the minimum over $\mathbf{q}' \in Q'$ of $\mathbf{r}_a \cdot \mathbf{q}'$ such that $(\mathbf{q}', z_{0a}) \succsim (\mathbf{q}, z_{jm}) \equiv \mathbf{w}_{jm}$
$M(u, \mathbf{r}, z, \succsim)$	An expenditure function, the minimum over $\mathbf{q} \in Q$ of $\mathbf{r} \cdot \mathbf{q}$ such that $U(\mathbf{q}, z, \succsim) \geq u$
$\tilde{V}(I, \mathbf{r}, z, \succsim)$	$\tilde{V} \equiv I + \tilde{v}(I, \mathbf{r}, z, \succsim)$, a money-metric indirect utility function, the maximum of $U(\mathbf{q}, z, \succsim)$ subject to the budget constraint $\mathbf{r} \cdot \mathbf{q} \leq I$
$\mathcal{V}(I_m, \mathbf{p}_m, \mathbf{r}_m, \mathbf{z}_m, \succsim)$	$\mathcal{V} = \max_{j \in J_m} \tilde{V}(I_m - p_{jm}, \mathbf{r}_m, z_{jm}, \succsim)$ unconditional maximum utility in scenario m
$\tilde{P}_k(I, \mathbf{p}_m, \mathbf{r}_m, \mathbf{z}_m, S)$	The probability that choice k in scenario m attains maximum utility \mathcal{V}
$x_{jm} = X(I - p_{jm}, \mathbf{r}_m, z_{jm})$	A finite-dimensional vector of predetermined functions
$v(I_m - p_{jm}, \mathbf{r}_m, z_{jm}, \succsim)$	$v = x_{jm} \beta$ with $I = I_m$, parameters $\beta = \beta(\succsim)$, an approximation to $\tilde{v}(I_m - p_{jm}, \mathbf{r}_m, z_{jm}, \succsim)$
$V(I_m - p_{jm}, \mathbf{r}_m, z_{jm}, \succsim)$	$V = I_m - p_{jm} + x_{jm} \beta + \sigma \varepsilon_{jm}$ approximation to \tilde{V} with additive EV1 "noise", $\sigma = \sigma(\succsim)$
$P_k(I, \mathbf{p}_m, \mathbf{r}_m, \mathbf{z}_m, S)$	$P_{km} = \mathbf{E}_{\beta s} \exp(\frac{x_{km} \beta - p_{km}}{\sigma}) / \sum_{j=0}^{J_m} \exp(\frac{x_{jm} \beta - p_{jm}}{\sigma})$ MMNL approximation to $\tilde{P}_k(I, \mathbf{p}_m, \mathbf{r}_m, \mathbf{z}_m, S)$

Assume that consumers are characterized by a vector s of observed demographics and history, and by *real* income I_m in a bounded interval $[I^L, I^U]$. In many applications, $I_a = I_b$, but if changing from scenario a to scenario b entails a net production cost or fulfilled transfer, assessed as a lump sum net tax, then I_a and I_b will differ by the amount of this tax. A consumer's market opportunities under policy m are summarized in prices $\mathbf{r}_m \in \mathbb{R}$, and for $j \in \mathbf{J}_m$, attributes $z_{jm} \in \mathbb{Z}$ and prices $p_{jm} \in \mathbb{P}$, giving the budget constraint $\mathbf{r}_m \cdot \mathbf{q} \leq I_m - p_{jm}$ for vectors $\mathbf{q} \in \mathbb{Q}$ when a product j from \mathbf{J}_m is chosen. Let $\mathbf{q}^A \in \mathbb{Q}$ denote a vector that bounds all affordable vectors (i.e., $\mathbf{r}^L \cdot \mathbf{q} \leq I^U$ implies $\mathbf{q} \ll \mathbf{q}^A$) and define $\mathbb{Q}' = [0, \mathbf{q}^A]$. Let \mathbf{z} denote the vector of attributes and \mathbf{p} the vector of prices for the discrete alternatives in \mathbf{J} , and let \mathbf{z}_m and \mathbf{p}_m denote their subvectors for the available alternatives in \mathbf{J}_m .

We adopt a description of consumer preferences and optimizing behavior that is sufficiently flexible to encompass neoclassical preference maximization and some behavioral deviations, but is also empirically tractable. Assume that consumers have complete transitive reflexive preference preorders \succsim over vectors $(\mathbf{q}, \mathbf{z}) \in \mathbb{Q} \times \mathbb{Z} \equiv \mathbb{W}$. A preference preorder \succsim is characterized by the non-empty set $G(\succsim)$ of pairs $((\mathbf{q}', \mathbf{z}'), (\mathbf{q}'', \mathbf{z}'')) \in \mathbb{W} \times \mathbb{W}$ that satisfy $(\mathbf{q}', \mathbf{z}') \succsim (\mathbf{q}'', \mathbf{z}'')$. Let \mathbb{H} denote the field of preference preorders of consumers in the population. We will assume that preferences for continuous goods are monotonic (i.e., $\mathbf{q}' \geq \mathbf{q}'' \implies \mathbf{q}' \succsim \mathbf{q}''$), and that \mathbf{q}^U is sufficiently large and continuous goods are sufficiently desirable so that they can substitute for any affordable (\mathbf{q}, \mathbf{z}) ; i.e., $(\mathbf{q}^U, z_{0a}) \succsim (\mathbf{q}^A, \mathbf{z})$ for all $\mathbf{z} \in \mathbb{Z}$ and $\succsim \in \mathbb{H}$. We use the notation " $\succsim \in \mathbb{H}$ " and " $G(\succsim) \in \mathbb{H}$ " interchangeably, the notation " $>$ " for strict preference, and the notation " \sim " for indifference. We use the Euclidean norm on \mathbb{Q} , \mathbb{R} , and \mathbb{Z} ; e.g., $\|\mathbf{q}\| = \sqrt{\mathbf{q} \cdot \mathbf{q}}$ for $\mathbf{q} \in \mathbb{Q}$.

For non-empty subsets A, B of the metric space $\mathbb{W} \times \mathbb{W}$, define the Hausdorff distance $h(A, B)$ to be the greatest lower bound of positive scalars η such that each set is contained in an η -neighborhood of the other; i.e., if $N_\eta(A)$ denotes the union of the open balls of radius η centered at the points in A , then $h(A, B)$ is the greatest lower bound of η satisfying $B \subseteq N_\eta(A)$ and $A \subseteq N_\eta(B)$. The set $\mathbb{W} \times \mathbb{W}$ is compact, so h is bounded, and if $A, B \in \mathbb{W} \times \mathbb{W}$ are closed, then $h(A, B) = 0$ if and only if $A = B$. If the sets in \mathbb{H} are all closed, then h is a metric on \mathbb{H} termed the Hausdorff set metric, and the closure of \mathbb{H} is compact in its metric topology. We make a series of assumptions on preferences and budgets, beginning with a basic assumption on continuity of preferences:


A1. If a sequence of preorders $\succsim^i \in \mathbb{H}_+$ and sequences of consumption vectors $(\mathbf{w}^i, \mathbf{w}''^i) \in G(\succsim^i)$ satisfy $h(G(\succsim^i), G(\succsim^0)) \rightarrow 0$, $\mathbf{w}^i \rightarrow \mathbf{w}''^0$, and $\mathbf{w}''^i \rightarrow \mathbf{w}''^0$, then $G(\succsim^0) \in \mathbb{H}$ and $(\mathbf{w}''^0, \mathbf{w}''^0) \in G(\succsim^0)$.

Since our attention is primarily on discrete choice, we will make strong and simple assumptions on continuous good preferences. Fix baseline values $(\mathbf{r}_a, z_a) \in \mathbb{R} \times \mathbb{Z}$. For $(\mathbf{q}, \mathbf{z}) \in \mathbb{Q} \times \mathbb{Z}$ and $\succsim \in \mathbb{H}$, define $A(\mathbf{q}, \mathbf{z}, \succsim) = \{\mathbf{q}' \in \mathbb{Q} \mid (\mathbf{q}', z_a) \succsim (\mathbf{q}, \mathbf{z})\}$, the set of continuous commodity vectors $\mathbf{q}' \in \mathbb{Q}$ that combined with "benchmark" attributes z_a are *at least*

as good as (\mathbf{q}, z) . We will assume for $\mathbf{q} \in Q'$ that $\mathbf{q}^u \in A(\mathbf{q}, z, \succsim)$, so this set is non-empty. Assumption A1 implies that $A(\mathbf{q}, z, \succsim)$ is compact, and if $\mathbf{q}', \mathbf{q}'' \in Q', z', z'' \in Z$, and $(\mathbf{q}', z') \succ (\mathbf{q}'', z'')$, then $A(\mathbf{q}', z', \succsim)$ is contained in the interior of $A(\mathbf{q}'', z'', \succsim)$. The next assumption strengthens our monotonicity requirement for continuous goods and imposes Lipschitz continuity conditions on preferences. Let $h_\alpha(A', A'')$ denote the Hausdorff distance between non-empty subsets $A', A'' \subseteq Q$. If $A' \subseteq A''$, then $h_\alpha(A', A'') \equiv \inf\{\eta > 0 \mid A'' \subseteq N_\eta(A')\}$.

A2. For $\mathbf{q}', \mathbf{q}'' \in Q, z', z'' \in Z$, and $\succsim', \succsim'' \in H$, $\mathbf{q}'' \ll \mathbf{q}'$ implies $(\mathbf{q}', z') \succ' (\mathbf{q}'', z'')$. If $\mathbf{q} \in Q'$, then $(\mathbf{q}^u, z_a) \succ' (\mathbf{q}, z')$. There exist scalars $\alpha, \delta > 0$ such that (i) $(\mathbf{q}', z') \sim' (\mathbf{q}'', z'')$ implies $h_\alpha(A(\mathbf{q}', z', \succsim'), A(\mathbf{q}'', z'', \succsim'')) \leq \alpha \cdot h(\succsim', \succsim'')$, (ii) $h_\alpha(A(\mathbf{q}', z', \succsim'), A(\mathbf{q}'', z'', \succsim'')) \leq \alpha \cdot (|\mathbf{q}' - \mathbf{q}''| + |z' - z''| + h(\succsim', \succsim''))$, and (iii) $(\mathbf{q}', z') \succ' (\mathbf{q}'', z'')$ implies $\inf_{\mathbf{q} \in A(\mathbf{q}', z', \succsim')} \sup\{\eta > 0 \mid N_\eta(\mathbf{q}) \subseteq A(\mathbf{q}'', z'', \succsim'')\} \geq \delta h_\alpha(A(\mathbf{q}', z', \succsim'), A(\mathbf{q}'', z'', \succsim''))$.

Figure 3. Lowest and Highest Slopes between Contours



For each \succsim , the $A(\mathbf{q}, z, \succsim)$ define a contour map on Q , as illustrated in Figure 3. Assumptions A1 and A2 do not impose any convexity condition on preferences, but do require that the open quadrant to the northeast of any point in $A(\mathbf{q}, z, \succsim)$ is contained in the interior of this set. The Lipschitz continuity conditions (i) and (ii) rule out precipitous changes in $A(\mathbf{q}, z, \succsim)$ when $(\mathbf{q}, z, \succsim)$ changes. The line between boundary points in the upper contour sets $A' = A(\mathbf{q}', z', \succsim')$ and $A'' = A(\mathbf{q}'', z'', \succsim'')$ that achieves their Hausdorff distance gives the lowest slope between them, while the distance between boundary points with the highest slope is some fraction of their Hausdorff

distance. Condition (iii) bounds this fraction from below by $\delta > 0$. Then condition (iii) rules out “local cliffs” by bounding the ratio of the highest slope to the lowest slope between two contours as they converge.⁵

A last assumption confirms that income is in a range where consumer budgets are limiting but nevertheless allow choice of any available discrete alternative, and all affordable \mathbf{q} are in Q' :

A3. The consumption vector $(0,z)$ is affordable for each $z \in Z$ (i.e., $l^z > p^z$), and $\mathbf{q} \in Q$ and $\mathbf{r}^l \cdot \mathbf{q} \leq l^z \implies \mathbf{q} \in Q'$.

We next establish that A1-A3 are sufficient to guarantee well-behaved representations of preferences by utility, expenditure, and indirect utility functions.

Lemma 2.1. If A1, then each $G(\succsim) \in H$ is a compact set, and H is a compact metric space with metric h .

Proof: For fixed \succsim , A1 establishes that $G(\succsim)$ is a non-empty closed subset of the compact space $W \times W$. The properties of h and H are given in Aliprantis and Border (2006, Sections 3.16-3.18), particularly Definition 3.70, Theorem 3.85, and Corollary 3.95. ■

Lemma 2.2. Suppose A1 and A2, $(\mathbf{q}', z'), (\mathbf{q}'', z'') \in Q' \times Z$, $\succsim \in H$, and $(\mathbf{q}', z') \succ (\mathbf{q}'', z'')$. Let δ denote the bound from A2, $\mathbf{1}$ denote a vector of ones, and $\gamma = h_Q(A(\mathbf{q}', z', \succsim), A(\mathbf{q}'', z'', \succsim))$. If \mathbf{q}''^* is in the boundary of $A(\mathbf{q}'', z'', \succsim)$, \mathbf{q}''^* is in the boundary of $A(\mathbf{q}', z', \succsim)$, and $\mathbf{r} \in R$, then $\mathbf{q}''^* - \gamma \delta \mathbf{r} / \|\mathbf{r}\| \in A(\mathbf{q}'', z'', \succsim)$ and $\mathbf{q}''^* + \gamma \mathbf{1} \in A(\mathbf{q}', z', \succsim)$.

Proof: By A2, all points in a neighborhood of \mathbf{q}''^* with radius $\gamma \delta$ are contained in $A(\mathbf{q}'', z'', \succsim)$, giving the first result. If $\mathbf{q}''^* + \gamma \mathbf{1} \notin A(\mathbf{q}', z', \succsim)$, then no point within radius γ of \mathbf{q}''^* is in $A(\mathbf{q}', z', \succsim)$, contradicting the definition of γ as the Hausdorff distance. ■

The next result uses the expenditure at baseline continuous good prices needed to achieve the level of satisfaction of a given vector of goods to define a well-behaved (i.e., bi-Lipschitz) utility function.

Theorem 2.3. Suppose A1 and A2. For each $(\mathbf{q}, z) \in Q' \times Z$ and $\succsim \in H$, define

$$(5) \quad U(\mathbf{q}, z, \succsim) = \min\{\mathbf{r}_a \cdot \mathbf{q}' \mid \mathbf{q}' \in A(\mathbf{q}, z, \succsim)\}.$$

Then, U is a continuous *direct utility function* on $Q' \times Z \times H$; i.e., U is continuous in its arguments and $U(\mathbf{q}', z', \succsim) \geq U(\mathbf{q}'', z'', \succsim) \iff (\mathbf{q}', z') \succsim (\mathbf{q}'', z'')$. Further, U is locally non-satiated with a range $U(Q', z, \succsim)$ contained in the bounded

⁵ Preferences on a compact set that are representable by absolutely continuous utility functions with gradients that are bounded positive and finite satisfy the Lipschitz continuity condition. The utility function in \mathbb{R}^2 satisfying $u = q_1 - \sqrt{1 - q_2}$ for $0 \leq q_2 \leq 1$ and $u = q_1 + \sqrt{q_2 - 1}$ for $q_2 > 1$ is an example that has a local cliff and fails to satisfy the condition at $q_2 = 1$.

interval $[0, r_a \cdot q^U]$, is Lipschitz in $(\mathbf{q}, z, \succcurlyeq)$, and is *bi-Lipschitz* in \mathbf{q} ; i.e., there exist scalars $\alpha_U, \lambda_U > 0$ such that $\mathbf{q}', \mathbf{q}'' \in Q'$ and $z', z'' \in Z$ imply

$$(6) \quad \begin{aligned} h_Q(A(\mathbf{q}', z', \succcurlyeq), A(\mathbf{q}'', z'', \succcurlyeq)) \cdot \lambda_U &\geq |U(\mathbf{q}', z', \succcurlyeq) - U(\mathbf{q}'', z'', \succcurlyeq)| \geq h_Q(A(\mathbf{q}', z', \succcurlyeq), A(\mathbf{q}'', z'', \succcurlyeq)) / \lambda_U \\ |U(\mathbf{q}', z', \succcurlyeq) - U(\mathbf{q}'', z'', \succcurlyeq)| &\leq \alpha_U \cdot (|\mathbf{q}' - \mathbf{q}''| + |z' - z''| + h(\succcurlyeq', \succcurlyeq'')) \end{aligned}$$

Proof: A1 and A2 imply that $A(\mathbf{q}, z, \succcurlyeq)$ is non-empty and closed, hence compact, so (5) is well-defined. Suppose $(\mathbf{q}', z') \succcurlyeq (\mathbf{q}'', z'')$. If $(\mathbf{q}, z_a) \in A(\mathbf{q}', z', \succcurlyeq)$, then $(\mathbf{q}, z_a) \succcurlyeq (\mathbf{q}', z')$ and transitivity implies $(\mathbf{q}, z_a) \succcurlyeq (\mathbf{q}'', z'')$, and hence $(\mathbf{q}, z_a) \in A(\mathbf{q}'', z'', \succcurlyeq)$. Therefore, $r_a \cdot \mathbf{q} \geq U(\mathbf{q}'', z'', \succcurlyeq)$, and hence $U(\mathbf{q}', z', \succcurlyeq) \geq U(\mathbf{q}'', z'', \succcurlyeq)$. A1 implies that U is continuous in its arguments. In particular, U is sequentially continuous in \succcurlyeq in the Hausdorff metric topology on H , or equivalently the topology of closed convergence.

If $\mathbf{q}^{*'} \in \operatorname{argmin}\{r_a \cdot \mathbf{q} \mid \mathbf{q} \in A(\mathbf{q}', z', \succcurlyeq)\}$ and $\mathbf{q}^{*''} \in \operatorname{argmin}\{r_a \cdot \mathbf{q} \mid \mathbf{q} \in A(\mathbf{q}'', z'', \succcurlyeq)\}$, then $U(\mathbf{q}', z', \succcurlyeq) - U(\mathbf{q}'', z'', \succcurlyeq) = r_a \cdot (\mathbf{q}^{*'} - \mathbf{q}^{*''})$. Let $\gamma \equiv h_Q(A(\mathbf{q}', z', \succcurlyeq), A(\mathbf{q}'', z'', \succcurlyeq))$. Lemma 2.2 with $\mathbf{r} = r_a$ implies $\tilde{\mathbf{q}}'' = \mathbf{q}^{*'} - \delta \gamma r_a / \|r_a\| \in A(\mathbf{q}'', z'', \succcurlyeq)$ and $\tilde{\mathbf{q}}' = \mathbf{q}^{*''} + \gamma \mathbf{1} \in A(\mathbf{q}', z', \succcurlyeq)$. Then $r_a \cdot (\mathbf{q}^{*'} - \tilde{\mathbf{q}}'') = \delta \gamma \|r_a\| \leq r_a \cdot (\mathbf{q}^{*'} - \mathbf{q}^{*''})$ and $r_a \cdot (\tilde{\mathbf{q}}' - \mathbf{q}^{*''}) = \gamma r_a \cdot \mathbf{1} \geq r_a \cdot (\mathbf{q}^{*'} - \mathbf{q}^{*''})$. Defining $\lambda_U = \max(r_a \cdot \mathbf{1}, 1 / \delta \|r_a\|)$, this gives the first row of (6). Then, the first Lipschitz condition in A2 gives the second inequality in (6),

$$|U(\mathbf{q}', z', \succcurlyeq) - U(\mathbf{q}'', z'', \succcurlyeq)| \leq \lambda_U h_Q(A(\mathbf{q}', z', \succcurlyeq), A(\mathbf{q}'', z'', \succcurlyeq)) \leq \lambda_U \alpha \cdot (|\mathbf{q}' - \mathbf{q}''| + |z' - z''| + h(\succcurlyeq', \succcurlyeq'')). \quad \blacksquare$$

Theorem 2.4. Suppose A1 and A2, and $U: Q' \times Z \times H \rightarrow [0, r_a \cdot q^U]$ from (5). For each $(\mathbf{r}, z) \in R \times Z$, $\succcurlyeq \in H$, and $u \in U(Q', z, \succcurlyeq)$, define

$$(7) \quad M(u, \mathbf{r}, z, \succcurlyeq) = \min_{\mathbf{q}' \in Q'} \{r \cdot \mathbf{q}' \mid U(\mathbf{q}', z, \succcurlyeq) \geq u\}.$$

Then M is an *expenditure function* that is continuous in its arguments, and concave, linear homogeneous, and non-decreasing in \mathbf{r} . Further, $M(u, r_a, z_a, \succcurlyeq) \equiv u$, is Lipschitz in $(\mathbf{r}, z, \succcurlyeq)$, and is *bi-Lipschitz and increasing* in u ; i.e., there exists scalars $\alpha_M, \lambda_M > 0$ such that $(\mathbf{r}', z', \succcurlyeq'), (\mathbf{r}'', z'', \succcurlyeq'') \in R \times Z \times H$ and $u', u'' \in U(Q', z, \succcurlyeq)$ with $u' \geq u''$ imply

$$(8) \quad \begin{aligned} (u' - u'') \lambda_M &\geq M(u', \mathbf{r}', z', \succcurlyeq') - M(u'', \mathbf{r}', z', \succcurlyeq') \geq (u' - u'') / \lambda_M \\ |M(u', \mathbf{r}', z', \succcurlyeq') - M(u', \mathbf{r}'', z'', \succcurlyeq'')| &\leq \alpha_M \cdot (|\mathbf{r}' - \mathbf{r}''| + |z' - z''| + h(\succcurlyeq', \succcurlyeq'')) \end{aligned}$$

Proof: Mas-Colell, Whinston, and Green (1995, Proposition 3E2) demonstrate that M is concave, linear homogeneous, and non-decreasing in \mathbf{r} . The continuity of M in its arguments follows from the Berge Maximum Theorem (Aliprantis and Border, 2006, Theorem 17.31). Suppose $(\mathbf{r}', z', \succcurlyeq'), (\mathbf{r}'', z'', \succcurlyeq'') \in R \times Z \times H$, and $u' \geq u''$.

Consider $\mathbf{q}^{*'} \in \operatorname{argmin} \{\mathbf{r}' \cdot \mathbf{q} \mid U(\mathbf{q}, z', \succsim') \geq u'\}$ and $\mathbf{q}^{*''} \in \operatorname{argmin} \{\mathbf{r}'' \cdot \mathbf{q} \mid U(\mathbf{q}, z'', \succsim'') \geq u''\}$. Let $\gamma = h_Q(A(\mathbf{q}^{*'}, z', \succsim'), A(\mathbf{q}^{*''}, z'', \succsim''))$. From Lemma 2.2, $\tilde{\mathbf{q}}'' = \mathbf{q}^{*'} - \delta\gamma\mathbf{r}/\|\mathbf{r}\| \in A(\mathbf{q}^{*''}, z'', \succsim'')$ and $\tilde{\mathbf{q}}' = \mathbf{q}^{*''} + \gamma\mathbf{1} \in A(\mathbf{q}^{*'}, z', \succsim')$, so that $\mathbf{r} \cdot (\mathbf{q}^{*'} - \tilde{\mathbf{q}}'') = \delta\gamma\|\mathbf{r}\| \leq \mathbf{r} \cdot (\mathbf{q}^{*'} - \mathbf{q}^{*''}) = M(u', \mathbf{r}', z', \succsim') - M(u'', \mathbf{r}', z', \succsim') \leq \mathbf{r} \cdot (\tilde{\mathbf{q}}' - \mathbf{q}^{*''}) = \gamma\mathbf{r} \cdot \mathbf{1}$. This establishes the first row inequality in (8) with $\lambda_M = \lambda_U \max(\mathbf{r}^U \cdot \mathbf{1}, 1/\delta\|\mathbf{r}^L\|)$.

Next consider $\mathbf{q}^{*'} \in \operatorname{argmin}_{U(\mathbf{q}, z', \succsim') \geq u'} \mathbf{r}' \cdot \mathbf{q}$ and $\mathbf{q}^{*''} \in \operatorname{argmin}_{U(\mathbf{q}, z'', \succsim'') \geq u''} \mathbf{r}'' \cdot \mathbf{q}$, and define $\gamma = h_Q(A(\mathbf{q}^{*''}, z'', \succsim''), A(\mathbf{q}^{*'}, z', \succsim'))$ $\leq \alpha(|z' - z''| + h(\succsim', \succsim''))$. Then, $A(\mathbf{q}^{*'}, z', \succsim') = A(\mathbf{q}^{*''}, z'', \succsim'')$. If $M(u', \mathbf{r}', z', \succsim') > M(u', \mathbf{r}'', z'', \succsim'')$, then $\mathbf{q}^{*''} + \gamma\mathbf{1} \in A(\mathbf{q}^{*'}, z', \succsim')$, implying $M(u', \mathbf{r}', z', \succsim') - M(u', \mathbf{r}'', z'', \succsim'') \leq \gamma\mathbf{r}' \cdot \mathbf{1} + |(\mathbf{r}'' - \mathbf{r}') \cdot \mathbf{q}^{*''}| \leq \|\mathbf{r}^U\| \alpha(|z' - z''| + h(\succsim', \succsim'')) \leq |\mathbf{r}'' - \mathbf{r}'| \cdot \|\mathbf{q}^U\| + \|\mathbf{r}^U\| \alpha(|z' - z''| + h(\succsim', \succsim''))$, proving the second row in (8) with $\alpha_M = \max(\|\mathbf{q}^U\|, \|\mathbf{r}^U\| \alpha)$. ■

Theorem 2.5. Suppose A1-A3 and $U: Q' \times Z \times H \rightarrow [0, r_a \cdot \mathbf{q}^U]$ from (5). For $I \in [I^L - p^U, I^U]$, $\mathbf{r} \in R$, $z \in Z$, and $\succsim \in H$, define a *money-metric indirect utility function*

$$(9) \quad V(I, \mathbf{r}, z, \succsim) \equiv \max_{\mathbf{q} \in Q'} \{U(\mathbf{q}, z, \succsim) \mid \mathbf{r} \cdot \mathbf{q} \leq I\}$$

satisfying $V(I, \mathbf{r}_a, z_a, \succsim) \equiv I$ that is continuous in its arguments with a range contained in the bounded interval $[0, r_a \cdot \mathbf{q}^U]$, is quasi-convex and homogeneous of degree zero in (I, \mathbf{r}) , is non-increasing in \mathbf{r} , is Lipschitz in $(\mathbf{r}, z, \succsim)$, and is bi-Lipschitz increasing in I ; i.e., there exist scalars $\alpha_V \geq r_a \cdot \mathbf{q}^U$ and $\lambda_V > 0$ such that $I' > I''$, $\mathbf{r}', \mathbf{r}'' \in R$, $z', z'' \in Z$, and $\succsim', \succsim'' \in H$ imply

$$(10) \quad \begin{aligned} \lambda_V(I' - I'') &\geq V(I', \mathbf{r}', z', \succsim') - V(I'', \mathbf{r}', z', \succsim') \geq (I' - I'')/\lambda_V \\ |V(I', \mathbf{r}', z', \succsim') - V(I'', \mathbf{r}'', z'', \succsim'')| &\leq \alpha_V(|I' - I''| + |\mathbf{r}' - \mathbf{r}''| + |z' - z''| + h(\succsim', \succsim'')) \end{aligned}$$

Further, $V(I, \mathbf{r}, z, \succsim)$ is twice continuously differentiable in (I, \mathbf{r}) except on a set of measure zero, and continuous good demands are almost everywhere in (I, \mathbf{r}) single-valued and continuously differentiable, and satisfy Roy's identity,

$$\mathbf{q} = D(I, \mathbf{r}, z, \succsim) \equiv - \frac{\partial V(I, \mathbf{r}, z, \succsim) / \partial \mathbf{r}}{\partial V(I, \mathbf{r}, z, \succsim) / \partial I}.$$

Proof: Theorem 2.3 implies that U is continuous in its arguments, and A3 assures that the budget set for \mathbf{q} is a non-empty subset of Q' for all discrete choices. Then, (9) is well-defined. Local non-satiation from A2 implies that V is the inverse with respect to u of $I = M(u, \mathbf{r}, z, \succsim)$; see Mas-Colell, Whinston, and Green (1995, Propositions 3D3 and 3E1). Result (8) then implies the first row of (10) with $\lambda_V = \lambda_U$. The continuity of M in its arguments from Lemma 2 and the bi-Lipschitz condition (8) imply V is continuous in its arguments and Lipschitz-continuous in $(\mathbf{r}, z, \succsim)$. It is immediate from the properties of M that V is homogeneous of degree zero in (I, \mathbf{r}) , and non-increasing

in r . Consider budgets (I^i, r^i) for $i = 0, 1$ and $(I^\theta, r^\theta) = \theta(I^0, r^0) + (1-\theta)(I^1, r^1)$ for $\theta \in (0, 1)$, and let q^* denote a maximand of (5) subject to the budget (I^θ, r^θ) . Then $V(I^\theta, r^\theta, z, \succcurlyeq) = U(q^*, z, \succcurlyeq)$. But $r^\theta \cdot q^* \leq I^\theta$ implies either or both $r^0 \cdot q^* \leq I^0$ or $r^1 \cdot q^* \leq I^1$, and therefore either $V(I^0, r^0, z, \succcurlyeq) = U(q^*, z, \succcurlyeq)$ or $V(I^1, r^1, z, \succcurlyeq) = U(q^*, z, \succcurlyeq)$, so that V is quasi-convex: $V(I^\theta, r^\theta, z, \succcurlyeq) \leq \max\{V(I^0, r^0, z, \succcurlyeq), V(I^1, r^1, z, \succcurlyeq)\}$. From the definition of quasi-convexity, there exists an increasing transformation ψ such that $v(I, r, z, \succcurlyeq) = \psi(V(I, r, z, \succcurlyeq))$ is a convex function of (I, r) . Results of Rademacher (1919) and Alexandrov (1939) establish that since $v(I, r, z, \succcurlyeq) \equiv \psi(I)$ is convex in I , ψ is bi-Lipschitz on $(I^L - p^U, I^U)$, continuously differentiable in I except possibly on a countable set, and almost everywhere twice continuously differentiable in I . Then, $\psi^{-1}(v)$ is also bi-Lipschitz and increasing, and hence continuously differentiable except on a countable set. This implies that $V(I, r, z, \succcurlyeq) = \psi^{-1}(v(I, r, z, \succcurlyeq))$ is increasing and bi-Lipschitz in I , and hence continuously differentiable in I except for a countable set, and almost everywhere twice continuously differentiable in (I, r) . Then the Roy (1942) identity applied to $v(I, r, z, \succcurlyeq)$, or equivalently to $V(I, r, z, \succcurlyeq)$, establishes that continuous good demands are almost surely single-valued and continuously differentiable in (I, r) . ■

Lemma 2.6. Suppose A1-A3, the direct utility function $U(q, z, \succcurlyeq)$ from (5), and its associated money metric indirect utility function $V(I, r, z, \succcurlyeq)$ from (9). For $q \in Q'$, define $U^*(q, z, \succcurlyeq) = \min_{r \in R} V(r \cdot q, r, z, \succcurlyeq)$. Then $U^*(q, z, \succcurlyeq)$ is quasi-concave and R-monotone⁶, $U^*(q, z, \succcurlyeq) \geq U(q, z, \succcurlyeq)$, with equality if for some $r \in R$, the conditions $q' \in Q'$ and $r \cdot q' \leq r \cdot q$ imply $U(q', z, \succcurlyeq) \leq U(q, z, \succcurlyeq)$. Then, U and U^* are observationally equivalent; i.e., $V(I, r, z, \succcurlyeq) \equiv \max_{q \in Q'} \{U^*(q, z, \succcurlyeq) \mid r \cdot q \leq I\}$ and the continuous good demands from maximization of U and U^* subject to the budget constraint $r \cdot q \leq I$ coincide except on a set of (I, r) of measure zero.

Proof: $V(r \cdot q, r, z, \succcurlyeq) = \max\{U(q', z, \succcurlyeq) \mid r \cdot q' \leq r \cdot q\} \geq U(q, z, \succcurlyeq)$ implies $U^*(q, z, \succcurlyeq) \geq U(q, z, \succcurlyeq)$. Suppose for some $r \in R$, the conditions $q' \in Q'$ and $r \cdot q' \leq r \cdot q$ imply $U(q', z, \succcurlyeq) \leq U(q, z, \succcurlyeq)$, so that q' is a maximand of U subject to this budget constraint. Then $U^*(q, z, \succcurlyeq) \leq V(r \cdot q', r, z, \succcurlyeq) \leq U(q', z, \succcurlyeq) \leq U(q, z, \succcurlyeq)$. Hence, U and U^* are observationally equivalent. Suppose $q^0, q^1 \in Q'$ and $q^\theta = \theta q^0 + (1-\theta)q^1$ for $\theta \in (0, 1)$, and let $r^\theta \in R$ be such that $U^*(q^\theta, z, \succcurlyeq) = V(r^\theta \cdot q^\theta, r^\theta, z, \succcurlyeq)$. Since $r^\theta \cdot q^\theta \geq \min\{r^\theta \cdot q^0, r^\theta \cdot q^1\}$, $U^*(q^\theta, z, \succcurlyeq) = V(r^\theta \cdot q^\theta, r^\theta, z, \succcurlyeq) \geq \min\{V(r^\theta \cdot q^0, r^\theta, z, \succcurlyeq), V(r^\theta \cdot q^1, r^\theta, z, \succcurlyeq)\} \geq \min\{U^*(q^0, z, \succcurlyeq), U^*(q^1, z, \succcurlyeq)\}$, so U^* is quasi-concave in q . Suppose $r \cdot q'' > r \cdot q'$ for all $r \in R$. Then there exists r'' satisfying $U^*(q'', z, \succcurlyeq) = V(r'' \cdot q'', r'', z, \succcurlyeq) > V(r'' \cdot q', r'', z, \succcurlyeq) \geq U^*(q', z, \succcurlyeq)$, so U^* is R-monotone. ■

⁶ A function $u(q)$ is quasi-concave if $0 < \theta < 1$ implies $u(\theta q' + (1-\theta)q'') \geq \min(u(q'), u(q''))$, and is R-monotone if $r \cdot q' \geq r \cdot q''$ for all $r \in R$ implies $u(q') \geq u(q'')$.

The indirect utility function (9) in Theorem 2.5 gives the utility obtainable conditioned on a discrete alternative j with attributes z_{jm} . The consumer is assumed to choose j to maximize utility subject to the budget constraint $\mathbf{r}_m \cdot \mathbf{q} + p_{jm} \leq I$ when the alternatives $j \in \mathbf{J}_m$ have attributes z_{jm} , achieving *unconditional indirect utility*

$$(11) \quad u = \mathcal{V}(I, \mathbf{p}_m, \mathbf{r}_m, \mathbf{z}_m, \succsim) \equiv \max_{j \in \mathbf{J}_m} V(I - p_{jm}, \mathbf{r}_m, z_{jm}, \succsim).$$

Associated with (11) is an *unconditional expenditure function* $I = \mathcal{M}(u, \mathbf{p}_m, \mathbf{r}_m, \mathbf{z}_m, \succsim)$ obtained as an implicit solution of (11), or equivalently as

$$(12) \quad I = \mathcal{M}(u, \mathbf{p}_m, \mathbf{r}_m, \mathbf{z}_m, \succsim) \equiv \min_{j \in \mathbf{J}_m} [M(u, \mathbf{r}_m, z_{jm}, \succsim) + p_{jm}].$$

Theorems 2.4 and 2.5 imply that \mathcal{M} is bi-Lipschitz increasing in u and \mathcal{V} is bi-Lipschitz increasing in I .

Next characterize the choices and demands coming from the indirect utility function (11). For $k \in \mathbf{J}_m$, and $I \in [I^L, I^U]$, $\mathbf{p}_m \in \mathbb{P}^{|\mathbf{J}_m|}$, $\mathbf{z}_m \in \mathbb{Z}^{|\mathbf{J}_m|}$, and $\mathbf{r}_m \in \mathbb{R}$, define the set of preferences that make alternative k uniquely optimal,

$$(13) \quad H^k(I, \mathbf{p}_m, \mathbf{r}_m, \mathbf{z}_m) = \{\succsim \in H \mid V(I - p_{km}, \mathbf{r}_m, z_{km}, \succsim) > V(I - p_{jm}, \mathbf{r}_m, z_{jm}, \succsim) \text{ for } j \in \mathbf{J}_m \setminus \{k\}\}$$

and let $H^\#(I, \mathbf{p}_m, \mathbf{r}_m, \mathbf{z}_m) = \bigcup_{k \in \mathbf{J}_m} H^k(I, \mathbf{p}_m, \mathbf{r}_m, \mathbf{z}_m)$ denote the set of all preferences that result in a unique utility-maximizing choice. For $\succsim \in H^\#(I, \mathbf{p}_m, \mathbf{r}_m, \mathbf{z}_m)$, choice is indicated by

$$(14) \quad \delta_k(I, \mathbf{p}_m, \mathbf{r}_m, \mathbf{z}_m, \succsim) = \mathbf{1}_{H^k(I, \mathbf{p}_m, \mathbf{r}_m, \mathbf{z}_m)}(\succsim) \equiv \begin{cases} 1 & \text{if } \succsim \in H^k(I, \mathbf{p}_m, \mathbf{r}_m, \mathbf{z}_m), \\ 0 & \text{otherwise} \end{cases},$$

and using Theorem 2.5 and Roy's identity, continuous good demands are given for almost all (I, \mathbf{r}) by

$$(15) \quad D(I, \mathbf{p}_m, \mathbf{r}_m, \mathbf{z}_m, \succsim) = - \sum_{k \in \mathbf{J}_m} \delta_k(I, \mathbf{p}_m, \mathbf{r}_m, \mathbf{z}_m, \succsim) \cdot \frac{\partial V(I - p_{km}, \mathbf{r}_m, z_{km}, \succsim) / \partial \mathbf{r}}{\partial V(I - p_{km}, \mathbf{r}_m, z_{km}, \succsim) / \partial I}.$$

In applications, the preferences \succsim of an individual consumer that appear as an argument in (9) are unobserved and heterogeneous across consumers. Limited observations on the market choices of a single consumer will provide only partial identification of \succsim , insufficient in general to determine (9) with precision. The approach we take is to treat the \succsim as "random effects", with a probability $F^H(\cdot | s)$ on the field of preferences H whose salient features we can hope to identify from market data. Poor identification of individual consumer preferences is worsened if there is preference randomness within the individual, arising from the limits of psychophysical discrimination as in Thurstone (1927). However, intra-consumer randomness does provide a mechanism for breaking ties in utilities so that $H^\#$ has probability one. We will not at this point distinguish intra-consumer and inter-consumer preference heterogeneity, but this can be done in applications where one has multiple

observations on market choices of individuals; see Ben-Akiva, McFadden, and Train (2016). We follow the dictates of consumer sovereignty in assuming that while preferences can depend on consumer history through s , they are predetermined in analysis of current policies and market circumstances. Then, (11) provides a policy-independent yardstick for welfare analysis.⁷ Our final preference assumption is that preference heterogeneity is almost surely sufficient to break ties.⁸

$$\mathbf{A4.} \quad F^H(H^\#(I, \mathbf{p}_m, \mathbf{r}_m, \mathbf{z}_m) | s) \equiv 1 \text{ for } (I, \mathbf{p}_m, \mathbf{r}_m, \mathbf{z}_m) \in [I^L, I^U] \times P^{J_m+1} \times R \times Z^{J_m+1} \text{ and } s \in S.$$

Given A4, the discrete alternatives $k \in J_m$, are chosen with probabilities

$$(16) \quad P_k(I, \mathbf{p}_m, \mathbf{r}_m, \mathbf{z}_m, s) = F^H(H^k(I, \mathbf{p}_m, \mathbf{r}_m, \mathbf{z}_m) | s).$$

The conditional probability of \succcurlyeq given choice k then satisfies $F^H(A | s, k) = F^H(A \cap H^k(I, \mathbf{p}_m, \mathbf{r}_m, \mathbf{z}_m) | s) / P_k(I, \mathbf{p}_m, \mathbf{r}_m, \mathbf{z}_m, s)$ for measurable $A \subseteq H$, and the conditional probability of continuous good demand in a measurable set $B \subseteq Q$, given choice k , is $F^H(\{\succcurlyeq \in H^k(I, \mathbf{p}_m, \mathbf{r}_m, \mathbf{z}_m) | s, D(I, \mathbf{p}_m, \mathbf{r}_m, \mathbf{z}_m, \succcurlyeq) \in B\}) / P_k(I, \mathbf{p}_m, \mathbf{r}_m, \mathbf{z}_m, s)$.

We next show that the field of indirect utility functions (9) with the properties given by Assumptions A1-A4 can be approximated uniformly by a practical finitely-parameterized family, with random parameters in the population that are also finitely parameterized. Then, this family can be estimated from observed choices in sufficiently rich arrays of market environments faced by samples of consumers, and the estimated family can be used to carry out welfare calculations with no essential loss of generality.

Theorem 2.7. Suppose A1-A4. Let $\tilde{V}(I, \mathbf{r}_m, \mathbf{z}_{jm}, \succcurlyeq)$ for $(I, \mathbf{r}_m, \mathbf{z}_{jm}, \succcurlyeq) \in [I^L - p^U, I^U] \times R \times Z \times H$ denote the true indirect utility function from Theorem 2.5, and define $\tilde{v}(I, p_{jm}, \mathbf{r}_m, \mathbf{z}_{jm}, \succcurlyeq) \equiv \tilde{V}(I - p_{jm}, \mathbf{r}_m, \mathbf{z}_{jm}, \succcurlyeq) - I$ on $[I^L, I^U] \times [0, p^U] \times R \times Z \times H$. Given a small scalar $\gamma \in (0, 1)$, there exists a bound $\eta = -\ln(\gamma/4 |J_m|)$; a vector of predetermined twice continuously differentiable functions $X: [I^L - p^U, I^U] \times R \times Z \rightarrow \mathbb{R}^N$ drawn from a Schauder basis⁹ for the space $\mathcal{C}([I^L - p^U, I^U] \times R \times Z)$; a

⁷ Our analysis lumps unobserved perceptions and attributes of alternatives together with unobserved preferences. To maintain taste sovereignty when these unobserved factors are influenced by policy, we would need to make these sources of randomness explicit and consider how to detect their presence and identify their influence on welfare.

⁸ A sufficient condition for A4 is that the indirect utility (9) have the form $V(I, \mathbf{r}_m, \mathbf{z}_{jm}, \succcurlyeq) = V^*(I, \mathbf{r}_m, \mathbf{z}_{jm}, \succcurlyeq) + \zeta_j(\succcurlyeq)$, where $\zeta_j(\succcurlyeq)$ is a component of a vector of absolutely continuous psychometric disturbances in $|J|$ -dimensional space, conditioned on the heterogeneity in V^* . More generality, A4 is satisfied by a probabilistic transversality condition that ensures that the distribution of the vector of indirect utilities for the various alternatives is absolutely continuous in $|J|$ -dimensional space; see Shannon (2006).

⁹ A Schauder basis may be polynomials, Fourier series, or other series of functions that span the space of continuous functions on a compact finite-dimensional space. The basis may be tailored to reduce the number of terms required to achieve a given tolerance.

commensurate vector of Lipschitz-continuous real functions β from a compact subset $\mathcal{B} \subseteq \mathbb{C}(H, \mathbb{R}^N)$, and a Lipschitz-continuous real function $\sigma: H \rightarrow [\sigma^L, \sigma^U]$ from a compact subset $\mathcal{S} \subseteq \mathbb{C}(H, [\sigma^L, \sigma^U])$ with $\sigma^L > 0$ and $\sigma^U < \gamma/2\eta$; and independent standard type I extreme value distributed random variables ε_j such that:

(i) There is an approximate indirect utility function¹⁰

$$(17) \quad V(I - p_{jm}, \mathbf{r}_m, z_{jm}, \beta, \sigma, \varepsilon_j) = I + v(I, p_{jm}, \mathbf{r}_m, z_{jm}, \beta) + \sigma \varepsilon_j$$

on $[I^L - p^U, I^U] \times \mathbb{R} \times \mathbb{Z} \times \mathcal{B} \times \mathcal{S} \times \mathbb{R}$ with $v(I, p_{jm}, \mathbf{r}_m, z_{jm}, \beta) \equiv X(I - p_{jm}, \mathbf{r}_m, z_{jm}) \cdot \beta - p_{jm}$ such that $|\tilde{v}(I, p_{jm}, \mathbf{r}_m, z_{jm}, \succ) - v(I, p_{jm}, \mathbf{r}_m, z_{jm}, \beta(\succ))| < \gamma$ uniformly. Further, in the event $C = \{\varepsilon \mid |\varepsilon_j| \leq \eta \text{ for } j \in \mathbf{J}\}$ that has $\text{Prob}(C) > 1 - \gamma/2$, $|\tilde{V}(I, \mathbf{r}_m, z_{jm}, \succ) - V(I, \mathbf{r}_m, z_{jm}, \beta(\succ), \sigma(\succ), \varepsilon_j)| < \gamma$ uniformly.

(ii) Suppose $\tilde{\delta}_{km}(I, \succ) \equiv \tilde{\delta}_k(I, \mathbf{p}_m, \mathbf{r}_m, \mathbf{z}_m, \succ)$ is the choice indicator given by (14) for \tilde{V} , and let $\delta_{km}(I, \beta, \sigma) \equiv \delta_k(I, \mathbf{p}_m, \mathbf{r}_m, \mathbf{z}_m, \beta, \sigma, \varepsilon)$ be an indicator for the discrete alternative that maximizes $V(I - p_{jm}, \mathbf{r}_m, z_{jm}, \beta, \sigma, \varepsilon_j)$ on \mathbf{J}_m . Then except for \succ and ε each in sets that have probability at most $\gamma/3$, $\delta_k(I, \mathbf{p}_m, \mathbf{r}_m, \mathbf{z}_m, \beta(\succ), \sigma(\succ), \varepsilon) = \tilde{\delta}_k(I, \mathbf{p}_m, \mathbf{r}_m, \mathbf{z}_m, \succ)$. Letting $\tilde{P}_k(I, \mathbf{p}_m, \mathbf{r}_m, \mathbf{z}_m, s)$ denote the true discrete choice probability, from (16), and $P_k(I, \mathbf{p}_m, \mathbf{r}_m, \mathbf{z}_m, s) = \mathbf{E} \delta_k(I, \mathbf{p}_m, \mathbf{r}_m, \mathbf{z}_m, \beta, \sigma, \varepsilon) \equiv \mathbf{E}_{\beta, \sigma | s} \frac{\exp(v_{km}(I, \beta)/\sigma)}{\sum_{j \in \mathbf{J}_m} \exp(v_{jm}(I, \beta)/\sigma)} \equiv \int_{\mathcal{B} \times \mathcal{S}} \frac{\exp(v_{km}(I, \beta)/\sigma)}{\sum_{j \in \mathbf{J}_m} \exp(v_{jm}(I, \beta)/\sigma)} F(d\beta, d\sigma | s)$, where $v_{km}(I, \beta) \equiv X(I, \mathbf{r}_m, z_{km}) \cdot \beta - p_{km}$, one then has, uniformly, $|\tilde{P}_k(I, \mathbf{p}_m, \mathbf{r}_m, \mathbf{z}_m, s) - P_k(I, \mathbf{p}_m, \mathbf{r}_m, \mathbf{z}_m, s)| < \gamma$.

(iii) Let $F(A|s) \equiv F^H(\{\succ \in H \mid (\beta(\succ), \sigma(\succ)) | s \in A\})$ for Borel sets $A \subseteq \mathcal{B} \times \mathcal{S}$ and $s \in \mathcal{S}$, and let $F_T(A|s)$ denote the empirical probability obtained from T independent draws from F . Let \mathcal{F}_1 denote the family of functions of the form (17) for $j \in \mathbf{J}_m$, \mathcal{F}_2 denote the family of functions formed as differences of the functions in \mathcal{F}_1 , and \mathcal{F} denote the family of functions of the form $\min(f_1, \dots, f_k)$ for $f_k \in \mathcal{F}_2$ and $1 \leq k \leq |\mathbf{J}|$, plus the function $f \equiv 1$. Let \mathcal{K} denote the family of functions $\exp(v(I, p_{jm}, \mathbf{r}_m, z_{jm}, \beta)/\sigma) / \sum_{i \in \mathbf{J}_m} \exp(v(I, p_{im}, \mathbf{r}_m, z_{im}, \beta)/\sigma)$ for v given in (17). Let \mathcal{J} denote the family of indicator functions $i = \mathbf{1}(f > 0)$ for $f \in \mathcal{F}$, and \mathcal{G} denote the family of functions of the form $i \cdot f$ for $i \in \mathcal{J}$ and $f \in \mathcal{F}$. Letting $\mathbf{E}_{\beta, \sigma}$ and $\mathbf{E}_{\beta, \sigma, T}$ denote expectation operators with respect to F and F_T respectively, there exists T such that $\text{Prob}(\sup_{T' \geq T} \sup_{f \in \mathcal{F} \cup \mathcal{K} \cup \mathcal{J} \cup \mathcal{G}} |(\mathbf{E}_{T'} - \mathbf{E})f| > \gamma/3) < \gamma/3$.

(iv) Let $\tilde{D}(I, \mathbf{p}_m, \mathbf{r}_m, \mathbf{z}_m, \succ)$ and $D(I, \mathbf{p}_m, \mathbf{r}_m, \mathbf{z}_m, \beta, \sigma, \varepsilon)$ denote the continuous good demands given by (15) for the indirect utility functions \tilde{V} and V respectively. If on a closed subset A of $[I^L - p^U, I^U] \times \mathbb{R} \times \mathbb{Z}$, V is continuously differentiable in (I, r) , then X can be selected with a sufficient number of terms so that on the set A and except for sets of \succ and ε that each have probability at most $\gamma/3$, $|\tilde{D}(I, \mathbf{p}_m, \mathbf{r}_m, \mathbf{z}_m, \succ) - D(I, \mathbf{p}_m, \mathbf{r}_m, \mathbf{z}_m, \beta(\succ), \sigma(\succ), \varepsilon_j)| < \gamma$ uniformly.

Proof: Let $H_\delta^\#(I, \mathbf{p}_m, \mathbf{r}_m, \mathbf{z}_m) = \cup_{k \in \mathbf{J}_m} \{\succ \in H \mid \tilde{V}(I - p_{km}, \mathbf{r}_m, z_{km}, \succ) > \tilde{V}(I - p_{jm}, \mathbf{r}_m, z_{jm}, \succ) + \delta \text{ for } j \in \mathbf{J}_m \& j \neq k\}$ for $0 < \delta \leq \gamma$. Then $H_\delta^\#(I, \mathbf{p}_m, \mathbf{r}_m, \mathbf{z}_m) \supseteq H^\#(I, \mathbf{p}_m, \mathbf{r}_m, \mathbf{z}_m)$, and A4 implies that there exists $\delta(I, \mathbf{p}_m, \mathbf{r}_m, \mathbf{z}_m) > 0$ such that

¹⁰ The approximation V is not guaranteed to satisfy the slope and curvature properties of \tilde{V} , but at each point where \tilde{V} is twice continuously differentiable with non-zero slopes and a non-singular (bordered) hessian, the approximation V for a sufficiently small tolerance γ will also have these properties and preserve signs, and hence locally have the same slope and curvature properties as \tilde{V} .

$F(H_{\delta(I, \mathbf{p}_m, \mathbf{r}_m, \mathbf{z}_m)}^\#(I, \mathbf{p}_m, \mathbf{r}_m, \mathbf{z}_m) | s) \geq 1 - \gamma/2$. Further, the continuity of \tilde{V} on $[I^L - \rho^U, I^U] \times \mathbb{R} \times \mathbb{Z} \times \mathbb{H}$ implies there exists an open neighborhood $N(I, \mathbf{p}_m, \mathbf{r}_m, \mathbf{z}_m)$ in $[I^L, I^U] \times P^{J_m} \times \mathbb{R} \times Z^{J_m}$ such that $\succcurlyeq \in H_{\delta(I, \mathbf{p}_m, \mathbf{r}_m, \mathbf{z}_m)}^\#(I, \mathbf{p}_m, \mathbf{r}_m, \mathbf{z}_m)$ and $(\tilde{I}, \tilde{\mathbf{p}}_m, \tilde{\mathbf{r}}, \tilde{\mathbf{z}}_m) \in N(I, \mathbf{p}_m, \mathbf{r}_m, \mathbf{z}_m)$ imply $\max_{k \in J_m} \{\tilde{V}(\tilde{I} - \tilde{p}_{km}, \tilde{\mathbf{r}}, \tilde{z}_{km}, \succcurlyeq) - \max_{j \neq k} \tilde{V}(\tilde{I} - \tilde{p}_{jm}, \tilde{\mathbf{r}}, \tilde{z}_{jm}, \succcurlyeq)\} > \delta(I, \mathbf{p}_m, \mathbf{r}_m, \mathbf{z}_m)/2$. One can then extract a finite family of these open neighborhoods that cover $[I^L, I^U] \times P^{J_m} \times \mathbb{R} \times Z^{J_m}$. Let $\delta_0 > 0$ denote the minimum of the $\delta(I, \mathbf{p}_m, \mathbf{r}_m, \mathbf{z}_m)$ for this finite family and define a constant $\sigma = \sigma(\succcurlyeq) \equiv \frac{\delta_0}{12\eta}$. Recall that Z is a finite union of disjoint rectangles. Combine each of these rectangles with the rectangular domains of income and prices, shift and scale these rectangles so they form a unit cube, and apply Appendix Theorem A.1 to establish the existence of a vector of multivariate polynomials $v(I, \mathbf{r}_m, z_{jm}, \succcurlyeq) \equiv X(I, \mathbf{r}_m, z_{jm}) \cdot \beta(\succcurlyeq)$ that satisfy $|\tilde{v}(I, \mathbf{r}_m, z_{jm}, \succcurlyeq) - v(I, \mathbf{r}_m, z_{jm}, \succcurlyeq)| < \frac{\delta_0}{12} \leq \gamma$. From the properties of EV1 variates, the event C has $\text{Prob}(C) > 1 - \gamma/2$, and if C , then $|\sigma \varepsilon_j| < \delta_0/12$. In the event C , $V(I, \mathbf{r}_m, z_{jm}, \succcurlyeq)$ given by (17) satisfies

$$|\tilde{V}(I, \mathbf{r}_m, z_{jm}, \succcurlyeq) - V(I, \mathbf{r}_m, z_{jm}, \succcurlyeq)| \leq |v(I, \mathbf{r}_m, z_{jm}, \succcurlyeq) - \tilde{v}(I, \mathbf{r}_m, z_{jm}, \succcurlyeq)| + |\sigma \varepsilon_j| < \delta_0/6.$$

This establishes (i).

For any point $(\tilde{I}, \tilde{\mathbf{p}}_m, \tilde{\mathbf{r}}, \tilde{\mathbf{z}}_m) \in [I^L, I^U] \times P^{J_m+1} \times \mathbb{R} \times Z^{J_m+1}$, let $(I, \mathbf{p}_m, \mathbf{r}_m, \mathbf{z}_m)$ be the center of a neighborhood in the open cover that includes $(\tilde{I}, \tilde{\mathbf{p}}_m, \tilde{\mathbf{r}}, \tilde{\mathbf{z}}_m)$. The probability of the event $C \cap H_{\alpha(I, \mathbf{p}_m, \mathbf{r}_m, \mathbf{z}_m)}^\#(I, \mathbf{p}_m, \mathbf{r}_m, \mathbf{z}_m)$ is at least $1 - \gamma$. In this case, $V(\tilde{I} - \tilde{p}_{km}, \tilde{\mathbf{r}}, \tilde{z}_{km}, \succcurlyeq) - \max_{j \neq k} V(\tilde{I} - \tilde{p}_{jm}, \tilde{\mathbf{r}}, \tilde{z}_{jm}, \succcurlyeq) > \frac{\delta(I, \mathbf{p}_m, \mathbf{r}_m, \mathbf{z}_m)}{2} > \delta_0/2$ for some k implies $\tilde{V}(\tilde{I} - \tilde{p}_{km}, \tilde{\mathbf{r}}, \tilde{z}_{km}, \succcurlyeq) - \max_{j \neq k} \tilde{V}(\tilde{I} - \tilde{p}_{jm}, \tilde{\mathbf{r}}, \tilde{z}_{jm}, \succcurlyeq) > \delta_0/6$. Then (ii) holds with probability at least $1 - \gamma$. The bound on the difference between the exact and approximate choice probabilities then follows.

The proof of (iii) utilizes results on convergence of empirical processes given in Appendix A. The functions in the family \mathcal{F}_1 , and hence in \mathcal{F}_2 , are linear in $(\beta, \gamma, \varepsilon_0, \dots, \varepsilon_{J|J})$. Then these families are contained in a finite-dimensional subspace defined by their intercepts and slope coefficients. The functions in \mathcal{F}_2 are Lipschitz in these intercepts and slope coefficients, implying that \mathcal{F} is Lipschitz in these parameters. By construction, the domain $[I^L - \rho^U, I^U] \times \mathbb{R} \times \mathbb{Z} \times \mathcal{B} \times \mathcal{S}$ and the domain $\mathcal{B} \times [\sigma^L, \sigma^U]$ of (β, σ) are compact, so that $v(I, \mathbf{p}_{jm}, \mathbf{r}_m, z_{jm}, \beta)$ is bounded on its domain by a constant M . Therefore, $f^* = M + \sigma^U |\varepsilon|$ is an envelope function for \mathcal{F}_1 and $2f^*$ is an envelope function for \mathcal{F}_2 , and hence for \mathcal{F} , that from Appendix B(a) satisfies $\mathbf{E} f^* \leq M + 1.219384 \sigma^U$. The family \mathcal{K} is Lipschitz in $(\beta, \sigma) \in \mathcal{B} \times [\sigma^L, \sigma^U]$ since $\sigma^L > 0$, with envelope function $f^* \equiv 1$. Apply Theorem A.3 to establish the result for \mathcal{F} and \mathcal{K} , and Theorem A.4 to establish the result for \mathcal{J} and \mathcal{G} .

From Theorem 2.5, \tilde{V} is almost everywhere continuously differentiable in (I, \mathbf{r}) , and where it is, continuous good demands are unique and are given by (15). Let A be a closed set on which this continuous differentiability holds. Then, Lemma 6 establishes that the derivatives of V approximate uniformly on A the corresponding derivatives of \tilde{V} . Combined with the bi-Lipschitz property of \tilde{V} in I , this establishes (v). ■

Result (iii) in this theorem shows that utility and the distribution of tastes can be approximated using the empirical distribution of the finite-dimensional taste parameters (β, σ) ; thus, using a distribution that has finite support, and is trivially a finite-parameter approximation. In general, it is then possible to achieve an approximation of the same precision using other finite-parameter families of distributions. The proof of this theorem assumes a constant for the scaling factor σ , and this is sufficient for the approximation results, but allowing heterogeneity in σ may give more parsimonious approximations with respect to the specification of β .

The direct utility function (5) given by Lemma 2 can be interpreted as a continuous mapping from the compact space of preferences H onto a compact subset \mathcal{U} of the normed linear space $\mathcal{C}(Q' \times Z)$ of continuous real-valued functions $u: Q' \times Z \rightarrow \mathbb{R}$. The probability F^H on H induces a probability F^U on \mathcal{U} that satisfies $F^U(A|s) = F^H(\{\succsim \in H \mid U(\cdot, \cdot, \succsim) \in A\} \mid s)$ for measurable subsets A of \mathcal{U} . Then, given A1-A4, the field of preferences can be characterized by (\mathcal{U}, F^U) rather than (H, F^H) . With this characterization, the money-metric utility function (9) in Lemma 4 is written $V: [I^L - p^U, I^U] \times R \times Z \times \mathcal{U} \rightarrow \mathbb{R}$, and correspondingly the choice indicator $\delta_k(I, \mathbf{p}_m, \mathbf{r}_m, \mathbf{z}_m, U)$ from (14) and continuous good demand $D(I, \mathbf{p}_m, \mathbf{r}_m, \mathbf{z}_m, U)$ from (15) are written as functions of $U \in \mathcal{U}$.

Consider the family (17) of the indirect utility functions $V(I_m - p_{jm}, \mathbf{r}_m, z_{jm}, \beta, \sigma, \varepsilon_j) \equiv I_m + v(I_m, p_{jm}, \mathbf{r}_m, z_{jm}, \beta) + \sigma \varepsilon_j$, where $v(I_m, p_{jm}, \mathbf{r}_m, z_{jm}, \beta) \equiv X(I_m - p_{jm}, \mathbf{r}_m, z_{jm}) \cdot \beta - p_{jm}$. A major simplification of the welfare calculus occurs when $v(I, p_{jm}, \mathbf{r}_m, z_{jm}, \beta)$ is independent of income. Make explicit the price index $\pi = \pi(\mathbf{r}_m)$ used to deflate income and prices to real terms, where $\pi(\mathbf{r}_a) = \pi(\mathbf{r}_b)$ by assumption, and rewrite V as

$$(18) \quad V(I_m - p_{jm}, \mathbf{r}_m, z_{jm}, \beta, \sigma, \varepsilon_j) = I_m / \pi(\mathbf{r}_m) + v(p_{jm} / \pi(\mathbf{r}_m), \mathbf{r}_m / \pi(\mathbf{r}_m), z_{jm}, \beta) + \sigma \varepsilon_j.$$

But this is a Gorman Polar Form, with the properties that choice among the products $j \in J_m$ is independent of income, and continuous good demands have the form

$$(19) \quad D(I, \mathbf{p}_m, \mathbf{r}_m, \mathbf{z}_m, \beta, \sigma, \varepsilon_m) = -\partial v(p_{jm} / \pi(\mathbf{r}_m), \mathbf{r}_m / \pi(\mathbf{r}_m), z_{jm}, \beta) / \partial \mathbf{r}_m + [I_m + v(p_{jm} / \pi(\mathbf{r}_m), \mathbf{r}_m / \pi(\mathbf{r}_m), z_{jm}, \beta)] \partial \ln \pi(\mathbf{r}_m) / \partial \mathbf{r}_m,$$

so that the only goods showing income effects are those whose prices influence the index π , and the Engle curves for these goods are affine linear. The Gorman polar preference field has been studied extensively in welfare economics, and has important aggregation properties for both continuous and discrete choice; see Chipman and

Moore (1980,1990), Small and Rosen (1981), and McFadden (2004,2014). The Gorman form (19) defines a *hedonic preference field* in which product attributes influence tastes only through an *effective price*, $\tilde{p}_{jm} = p_{jm} - X(\mathbf{r}_m, \mathbf{z}_{jm}) \cdot \beta$.

3. WELFARE ANALYSIS

We restate for product markets the neoclassical welfare calculus outlined in Section 1, utilizing the treatment of consumer theory given in Section 2.¹¹ There is a baseline, “as is,” or “incumbent” policy/scenario $m = a$ and a counterfactual, “but for,” or “replacement” policy/scenario $m = b$.¹² Consumers face menus of mutually exclusive products $j \in \mathbf{J}_m \subseteq \mathbf{J}$ with at least one “benchmark” or “no purchase” alternative whose attributes are unaffected by scenario changes.¹³ Our analysis will be carried out for a population of consumers who have preferences that satisfy assumptions A1-A4 in Section 2, and hence have indirect utility functions, conditioned on selection of product j and stated in terms of real income and prices, that from Lemma 7 are uniformly approximated by

$$(20) \quad V(I - p_{jm}, \mathbf{r}_m, \mathbf{z}_{jm}, \beta, \sigma, \varepsilon_j) \equiv I + v_{jm}(I, \beta) + \sigma \varepsilon_{jm},$$

where $v_{jm}(I, \beta)$ is shorthand for $v(I, p_{jm}, \mathbf{r}_m, \mathbf{z}_{jm}, \beta) \equiv X(I - p_{jm}, \mathbf{r}_m, \mathbf{z}_{jm})\beta - p_{jm}$. The vector β and positive scalar σ are randomly distributed in the population with a probability $F(\beta, \sigma | s)$ that is in a parametric family with parameter α , given observed socioeconomic history s , and the ε_{jm} are independent standard Extreme Value type I random variables. By construction, V is money-metric for a “no purchase” or “benchmark” alternative in scenario a (e.g., $v(I, p_{0a}, \mathbf{r}_a, \mathbf{z}_{0a}, \beta) \equiv 0$). In the event $H_m^k(I, \beta, \sigma) = \{\varepsilon_m \mid V(I - p_{km}, \mathbf{r}_m, \mathbf{z}_{km}, \beta, \sigma, \varepsilon_k) > V(I - p_{jm}, \mathbf{r}_m, \mathbf{z}_{jm}, \beta, \sigma, \varepsilon_j) \text{ for } j \in \mathbf{J}_m \setminus \{k\}\}$, the

¹¹ The basics of this theory can be found in Varian (1992, Chap. 7, 10), Mas-Colell, Whinston, and Green (1995, Chap. 3), and other graduate-level textbooks. See also McFadden and Winter (1966) and Border (2014).

¹² For convenience we will use the “baseline/as is/incumbent” and “counterfactual/but for/replacement” labels for both retrospective analysis of past policy and prospective analysis of policies not yet implemented, noting that these labels are arbitrary and interchangeable in many prospective applications. In retrospective applications, associating a with the historical scenario and b with the counterfactual leads to measures of welfare change often termed “Willingness to Pay” (WTP), while reversing these labels and making b the baseline leads to “Willingness to Accept” (WTA) welfare measures.

¹³ In many applications, the attainable products will be the same in both scenarios except for specific changes in prices or other attributes. However, the notation allows the possibility that products may appear or disappear when scenarios change. The case where psychometric noise is independent in the two scenarios, even for “brands” that are common to both scenarios, will be handled by indexing alternatives, including the “no purchase” alternative, differently so that $\mathbf{J}_a \cap \mathbf{J}_b = \emptyset$. If the products in an application are not mutually exclusive, or the consumer can buy more than one unit of a product, then \mathbf{J}_m indexes the mutually exclusive possible portfolios of product purchases. In general, \mathbf{J}_m may index locations or “addresses” in physical or hedonic space, and with added technical machinery need not be restricted to be finite.

consumer maximizes (21) at alternative $k \in J_m$, an event indicated by $\delta_{km}(I, \beta, \sigma, \epsilon_m) = 1$, with a probability that is uniformly approximated by a mixed multinomial logit,

$$(21) \quad P_{km} = P_k(I, \mathbf{p}_m, \mathbf{r}_m, \mathbf{z}_m, S, \alpha) \equiv \mathbf{E}_{\beta, \sigma | S} L_{km}(I, \beta, \sigma),$$

where $L_{km}(I, \beta, \sigma) = \frac{\exp(v_{km}(I, \beta)/\sigma)}{\sum_{j \in J_m} \exp(v_{jm}(I, \beta)/\sigma)}$ is the “flat” multinomial logit probability of the event $\{\epsilon | \delta_{km}(I, \beta, \sigma, \epsilon) = 1\}$.

The components of the vector of taste parameters β are termed “part-worth” or Willingness-to-Pay (WTP) coefficients for unit changes in the corresponding components of X . The parameter σ scales psychometric noise and is large when consumer acuity is low.¹⁴ Unconditional maximum indirect utility in scenario m when income is I satisfies

$$(22) \quad u_m = \mathcal{V}_k(I, \mathbf{p}_m, \mathbf{r}_m, \mathbf{z}_m, \beta, \sigma, \epsilon) \equiv \max_{j \in J_m} [I + v(I, p_{jm}, \mathbf{r}_m, z_{jm}, \beta) + \sigma \epsilon_j].$$

We assume consumer sovereignty – tastes are predetermined at the time of market choice. This implies that the distribution $F(\beta, \sigma | S, \alpha)$ is the same in each scenario m , and that the psychometric disturbances ϵ_{jm} are independent of income, and hence independent of income transfers. There is a substantive conceptual question about how psychometric noise should be treated in welfare calculations. Its presence may be attributable to whims, which enter perceived tastes temporarily and influence choice, but may be too transient to enter experienced utility. The welfare calculus should arguably minimize attributions of apparent gains and losses from transient whims, even when these whims capture genuine momentary sensations of well-being. A more critical argument for excluding whimsy from the welfare calculus is that while it can influence observed choices, psychometric noise is impossible to incorporate into fulfilled compensating transfers. If whimsy is to be excluded, the question is how to do this. Our approach is to form expectations of welfare measures after first computing these measures for each configuration of psychometric noise. An alternative is to form *expected* utility by averaging out the effects of psychometric noise before calculating welfare measures; this requires appropriate scaling of utility to be meaningful. In either case, there will remain a question of the consequences of inexact compensating transfers.

¹⁴ One motivation for introducing “psychometric noise” is that it is behaviorally realistic, but these disturbances also act as a perturbation of preferences, simplifying the mathematics of consumer demand analysis by smoothing “knife-edge” responses. Maskin and McFadden (2011) and Fosgerau and McFadden (2012) show that these “trembling payoff” perturbations can be used to strengthen conventional results in consumer theory and tighten some of the properties of Harsanyi-Selten “trembling hand” equilibria in game theory.

A more complicated question is the relationship of psychometric noise in the two scenarios. If there are different products, and no natural correspondence between the indexing of the products in scenario a and scenario b , or there are changes in a broad spectrum of product attributes that arguably encompass unobserved attributes even for “branded” products, or psychometric noise comes primarily from transient context-dependent “animal spirits” rather than from unobserved attributes of products, then it is plausible to assume that the relevant components of ϵ in the two scenarios are statistically independent. This is handled flexibly in our notation by either assuming that ϵ_{ja} and ϵ_{jb} are statistically independent for each j , or by assuming $\epsilon_{ja} = \epsilon_{jb}$ but $J_a \cap J_b = \emptyset$. However, if the products in scenarios a and b are identifiably the same “brands”, either unchanged or with changes only in prices or selected attributes, it is plausible that unobserved attributes and whims do not change and the relevant components of ϵ are the same in both scenarios for $j \in J_a \cap J_b$. We consider both cases. Since psychometric noise in our setup is defined as part of the approximation of true preferences, and not established by psychological experiments as an element, even transitory, in actual preferences, there is good reason to avoid assumptions on ϵ that complicate the welfare calculus, but also concern if the assumptions made on ϵ have a substantial effect on welfare conclusions. Generally, if the scaling of psychometric noise that is assumed to be EV1 makes this noise substantial relative to scenario changes in observed product attributes and prices, the prerequisites of the approximation Lemma 7 will fail, and the welfare conclusions will be sensitive to assumptions on ϵ_{jb} versus ϵ_{ja} . The most effective remedy is then to expand X and tighten the specification (21) to make σ small.

There are three substantive issues whose resolution affects the form of the welfare calculus: Is the analysis *prospective*, comparing policies not yet put into place, or *retrospective*, comparing “as is” and “but for” past policies? Is information on the tastes of individual consumers *complete* or *partial*, and if partial what welfare measures are relevant to transfers that can actually be fulfilled? Should well-being be assessed in terms of the *decision-utility* postulated to determine economic demand behavior, or in terms of *experienced-utility* after taste ambiguities and uncertainties are resolved? These issues are discussed in Sections 4, 5, and 6 below. In the remainder of this section, we restate for our general model of discrete product choice and neoclassical assumptions the welfare measures introduced in Section 1. In this restatement, we will assume that there is a single vector of psychometric noise components that are invariant across the policy scenarios, and treat situations where the noise associated with a “branded” product is assumed to change across scenarios by giving it different indices. The case of statistically independent psychometric noise in the two scenarios is then handled by assuming $J_a \cap J_b = \emptyset$.

One standard Hicksian welfare measure for the net gain in well-being from scenario b relative to scenario a is *Hicksian Compensating Variation* (HCV), the net decrease in scenario b income that makes the two scenarios

indifferent. Let $HCV(s,k,\beta,\sigma,\epsilon)$ denote this measure for an observed history s and scenario a choice k , and a vector (β,σ,ϵ) of unobservables.¹⁵ In terms of the conditional indirect utility function (21), $HCV(s,k,\beta,\sigma,\epsilon)$ satisfies

$$(23) \quad \max_{j \in J_b} [I_b - HCV + v(I_b - HCV, p_{jb}, \mathbf{r}_b, z_{jb}, \beta) + \sigma \epsilon_j] = I_a + v(I_a, p_{ka}, \mathbf{r}_a, z_{ka}, \beta) + \sigma \epsilon_k.$$

Removing the conditioning on k when it is not observed gives the measure $HCV(s,\beta,\sigma,\epsilon) = \min_{k \in J_a} HCV(s, k, \beta, \sigma, \epsilon)$.

The *Hicksian Equivalent Variation* (HEV), denoted $HEV(s,\beta,\sigma,\epsilon)$, is the net increase in scenario a income that makes the two scenarios indifferent. Because the scenario a choice may change with the income transfer, HEV does not depend explicitly on the uncompensated scenario a choice k ; from (23), $HEV(s,\beta,\sigma,\epsilon)$ satisfies

$$(24) \quad \max_{j \in J_b} [I_b + v(I_b, p_{jb}, \mathbf{r}_b, z_{jb}, \beta) + \sigma \epsilon_j] = \max_{j \in J_a} [I_a + HEV + v(I_a + HEV, p_{ja}, \mathbf{r}_a, z_{ja}, \beta) + \sigma \epsilon_j].$$

Sometimes, HCV is termed *Willingness-to-Pay* (WTP), and HEV is termed *Willingness-to-Accept* (WTA); this terminology is related to the description of β as a vector of WTP coefficients, but only in special cases is there a simple mapping between β and HCV or HEV.

The definition of *Market Compensating Equivalent* (MCE) that generalizes (1) is the difference in the utilities (23) that the consumer would attain in scenarios a and b in the absence of compensation, scaled by a marginal utility of income in scenario a that converts the utility difference into monetary units. The conditional indirect utility $v(I - p_{jm}, \mathbf{r}_m, z_{jm}, \beta, \sigma, \epsilon_j) \equiv I + v(I, p_{km}, \mathbf{r}_m, z_{km}, \beta) + \sigma \epsilon_{km}$ is denominated in monetary units and is money-metric for a “no purchase” alternative. However, if there are neoclassical income effects for other alternatives k , one has $\partial v(I - p_{km}, \mathbf{r}_m, z_{km}, \beta, \sigma, \epsilon_k) / \partial I \equiv 1 + \partial v(I, p_{km}, \mathbf{r}_m, z_{km}, \beta) / \partial I \neq 1$. If alternative k is chosen in scenario a , then

$$(25) \quad MCE(s,k,\beta,\sigma,\epsilon) = \frac{\max_{j \in J_b} [I_b + v(I_b, p_{jb}, \mathbf{r}_b, z_{jb}, \beta) + \sigma \epsilon_j] - \max_{j \in J_a} [I_a + v(I_a, p_{ja}, \mathbf{r}_a, z_{ja}, \beta) + \sigma \epsilon_j]}{\mu_k(I_a, \beta)}.$$

where

$$(26) \quad \mu_k(I_a, \beta) = \partial v(I_a - p_{ka}, \mathbf{r}_a, z_{ka}, \beta) / \partial I_a \equiv 1 + [\partial X(I_a - p_{ka}, \mathbf{r}_a, z_{ka}) / \partial I_a] \beta,$$

gives a definition of $MCE(s,k,\beta,\sigma,\epsilon)$ that at least locally has the money-metric property in scenario a .¹⁶ Later, when we consider cases where choice k in scenario a is not observed, or one observes or uses only the information that

¹⁵ Note that when k is observed in scenario a , the distribution of ϵ is conditioned on the event $\{\epsilon | \delta_{ka}(I, \beta, \sigma, \epsilon) = 1\}$.

¹⁶ The scale factor $\mu_k(I_a, \beta)$ in the definition of MCE is natural for retrospective analysis where the consumer has experienced scenario a , or for prospective analysis when scenario a is a default that will occur unless there is a policy intervention, and

“as is” choice is from a set $\mathbf{D} \subseteq \mathbf{J}_m$, $\mu_k(I_a, \beta)$ will be replaced by a scale factor $\mu_{\mathbf{D}}(I_a, \beta)$. Note that $\mu_k(I_a, \beta)$ defined by (26) is independent of σ and $\boldsymbol{\varepsilon}$, and if X does not depend on income, then $\mu_k(I_a, \beta) \equiv 1$.

The measure (25) can be interpreted as a generalization to multiple products with varying attributes of the Marshallian consumer surplus (MCS) introduced in Section 1; this generalization is more useful for applications than multivariate extensions, path-dependent when there are income effects, of the integral form (2) for MCS. First-order Taylor’s expansions of (24) imply

$$(27) \quad \text{HCV}(s, k, \beta, \sigma, \boldsymbol{\varepsilon}) \cdot \mu' = \text{HEV}(s, k, \beta, \sigma, \boldsymbol{\varepsilon}) \cdot \mu'' = \text{MCE}(s, k, \beta, \sigma, \boldsymbol{\varepsilon}) \cdot \mu_k(I_a, \beta) ,$$

where $\mu' = 1 + \partial v(I' - p_{jb}, \mathbf{r}_b, \mathbf{z}_{jb}, \beta) / \partial I$ and $\mu'' = 1 + \partial v(I'' - p_{ka}, \mathbf{r}_a, \mathbf{z}_{kb}, \beta) / \partial I$ are marginal utilities of income evaluated at the chosen alternatives j in scenario b and k in scenario a when there is no compensation, evaluated at incomes I' and I'' intermediate between uncompensated and compensated levels. If the marginal utility of income is constant, then $\text{HCV} = \text{HEV} = \text{MCE}$. In general, (24) can be solved quickly for HCV or HEV by iteration starting from either zero or MCE.

With the apparatus above, practical welfare analysis of product markets can be carried out in three steps. First, observations on the market choices of surveyed consumers, augmented by extra-market data on stated preferences if necessary to identify tastes for relevant attributes, can be used to estimate the mixed MNL model (22) and recover the probability $F(\beta, \sigma | s)$. An obvious caution is that the vector of predetermined functions X in (21) has to be comprehensive enough to achieve the approximation accuracy promised by Lemma 7, so that estimation of (22) needs to include a careful econometric specification analysis. A “method of sieves” or “semi-non-parametric” approach to the specification of X provides practical guidelines for this specification search. With this caveat, this setup is both practical and sufficiently general to handle welfare analysis of policy changes that affect discrete choice without making unwarranted assumptions on preferences.

Second, construct a large synthetic population. Start from a random sample from the target population. For each sampled person, assign a history s , incomes I_a and I_b , choice sets \mathbf{J}_a and \mathbf{J}_b , and market environments $(\mathbf{p}_a, \mathbf{r}_a, \mathbf{z}_a)$ and $(\mathbf{p}_b, \mathbf{r}_b, \mathbf{z}_b)$, using available data for the sampled individual wherever possible in order to preserve ecological correlations in the target population. Make multiple draws of (β, σ) from the estimated probability $F(\beta, \sigma | s)$ and of $\boldsymbol{\varepsilon}_a$ and $\boldsymbol{\varepsilon}_b$ from the standard Extreme Value Type I distribution. If $\boldsymbol{\varepsilon}_a$ and $\boldsymbol{\varepsilon}_b$ are not statistically independent, the

is unambiguous when indirect utility has been transformed so that the marginal utility of income in scenario a remains constant when income changes. However, more generally, MCE will be affected by transformations of utility and the evaluation point for the marginal utility of income, and additional criteria may be needed to select among alternative versions of MCE.

analysis has to include a specification for their joint distribution and a mechanism for sampling from it. Assign utility-maximizing choices k in scenario a and j in scenario b . Each draw defines a synthetic consumer.

Third, calculate the measures $HCV(s, k, \beta, \sigma, \epsilon)$, $HEV(s, k, \beta, \sigma, \epsilon)$, and $MCE(s, k, \beta, \sigma, \epsilon)$ from (24) and (25) for each consumer in the synthetic population. These measures can be aggregated over this synthetic population or subpopulations to estimate hypothetical compensating transfers for relevant consumer classes. However, transfers that are actually fulfilled in the target population can depend only on observable history s and (if observed) the scenario a choice k . Define uniform transfers $UMCE(s, k) = \mathbf{E}_{\beta, \sigma, \epsilon | s, k} MCE(s, k, \beta, \sigma, \epsilon)$ and $UMCE(s) = \mathbf{E}_{k, \beta, \sigma, \epsilon | s} MCE(s, k, \beta, \sigma, \epsilon)$. Fulfillment of these transfers in the real population in retrospective welfare analysis will not in general make individual consumers “whole”, but will balance individual gains and losses in the sense that a MCE welfare measure taken subsequent to these uniform transfers aggregates to zero. In the same way, one can solve for uniform transfers $t_k = UHCV(s, k)$ and $t = UHCV(s)$ that if fulfilled in scenario b balance the gains and losses from the remaining unfulfilled Hicksian transfers, so that a subsequent MCE aggregates to zero:

$$(28) \quad 0 = \begin{cases} \mathbf{E}_{\beta, \sigma, \epsilon | s, k} \left\{ \max_{j \in J_b} [I_b - t_k + v(I_b - t_k, p_{jb}, \mathbf{r}_b, z_{jb}, \beta) + \sigma \epsilon_j] - [I_a + v(I_a, p_{ka}, \mathbf{r}_a, z_{ka}, \beta) + \sigma \epsilon_k] \right\} \\ \mathbf{E}_{k, \beta, \sigma, \epsilon | s} \left\{ \max_{j \in J_b} [I_b - t + v(I_b - t, p_{jb}, \mathbf{r}_b, z_{jb}, \beta) + \sigma \epsilon_j] - [I_a + v(I_a, p_{ka}, \mathbf{r}_a, z_{ka}, \beta) + \sigma \epsilon_k] \right\} \end{cases}.$$

Analogous definitions can be given for $UHEV(s, k)$ and $UHEV(s)$.¹⁷ Note that the measures considered in this paragraph are all determined by decision-utility, with no adjustment for possible differences in decision and experienced utility.

4. PROSPECTIVE VERSUS RETROSPECTIVE WELFARE ANALYSIS

Traditional welfare theory considers a prospective policy change in a static “what if” environment. An “incumbent” or “default” policy/scenario a is compared with a “replacement” policy/scenario b in a situation where neither has been implemented and both are on the table. The theory assumes that the policymaker has the information and authority to carry out net lump sum transfers in the event that policy b is adopted, adjusted for direct policy-induced effects on incomes, that make each consumer indifferent between the policies, and

¹⁷ Another approach to defining $UHCV(s, k)$ and $UHEV(s, k)$ is to first define “representative” utility for the class of consumers with history s , perhaps the expectation of (21) with respect to the unobservables, and then define $UHCV(s, k)$, or $UHEV(s, k)$ as analogs of HCV or HEV for “representative” utility. However, these definitions will not in general have the property that $UHCV(s, k) = \mathbf{E}_{\beta, \sigma, \epsilon | s, k} HCV(\beta, \sigma, \epsilon)$ or $UHEV(s, k) = \mathbf{E}_{\beta, \sigma, \epsilon | s, k} HEV(\beta, \sigma, \epsilon)$.

assumes that if policy b is adopted, these transfers are fulfilled before consumers maximize utility. Under these conditions, the Hicksian Contingent Variation (HCV) defined in (24) is the precise measure of each lump sum transfer required. If instead, a and b are reversed, so that transfers are fulfilled if a is adopted, then the Hicksian Equivalent Variation (HEV) is the precise measure of each lump sum transfer required. So long as population aggregate HCV or HEV, adjusted for policy-induced income changes, exceeds zero, a shift from policy a to policy b with the exact individual transfers fulfilled, plus any distribution of the residual surplus, is a Pareto improvement.

In practice many welfare calculations are *retrospective* rather than prospective. The welfare question is what transfers *after the fact* redress harm from a past “as is” or “baseline” scenario a in which some products were defective or improperly marketed, using as a benchmark a “but for” or “counterfactual” scenario b in which these flaws were absent.¹⁸ A key feature of these applications is that the transfer occurs after the decision-utility-maximizing choice would have been made in the “but for” scenario, and hence could not be a factor in “but for” choice. Put another way, the “but for” utility maximization that would have occurred at the consumer’s original income will not in general coincide with that assumed in the Hicksian compensating variation calculation in which the transfer would have been made prior to consumer choice and would have influenced that choice. Since at the time the corrective transfer is being considered, the consumer is in the “as is” situation, this transfer is denominated in “as is” monetary units. Then, the transfer that “makes whole the consumer with choice k in the baseline scenario” equals the difference in the utilities (21) that would have been attained in the “but for” and “as is” scenarios, scaled to “as is” monetary units, the MCE (25).

Suppose the purpose of a prospective policy analysis is not to actually fulfill the HCV or HEV transfers associated with a move from scenario a to scenario b , but simply to determine whether it is possible *in principle* to compensate consumers so that the move from scenario a to scenario b would be a Pareto improvement. Then, arguably, aggregate MCE rather than aggregate HCV or HEV is the appropriate welfare criterion. Further, MCE is easier to compute and aggregate than HCV or HEV, since it is obtained as an explicit solution (25) from the indirect utility functions (21) of individual consumers, and the distribution of these solutions in the target population. Equation (27) shows that HCV, HEV, and MCE differ only due to differences in the marginal utility of income at different arguments. Later, we show in examples that these differences are often but not always modest. Then,

¹⁸ Retrospective policy analysis is often conducted in conjunction with litigation, and statutes and legal rulings often control the definition of harm and the scope and magnitude of remedies. These legal standards are often rooted in economic arguments, but may nevertheless deviate from a purely economic analysis of harm and remedy. In this paper, we consider only the economic foundation of retrospective analysis, and do not take up legal considerations.

the distinction between prospective and retrospective welfare measures often will be empirically unimportant, but occasionally will be of practical as well as theoretical significance.

The distinction we have made between prospective and retrospective welfare analysis does not require explicit consumer dynamics, but a MCE transfer to redress past harm obviously occurs at some time later than the period of the harm, introducing issues such as discounting and pre-judgement interest, but more fundamentally the longer-run impacts of injury on consumer assets and opportunities. We leave this as a topic for future research, but note that in a fully dynamic model, the impact of policy on state variables justifies scaling MCE in monetary units that make the consumer whole in terms of lifetime well-being.¹⁹

5. PARTIAL OBSERVABILITY AND WELFARE AGGREGATES

Traditional welfare analysis assumes that the individual utility functions required to calculate measures of well-being can be recovered fully (with money-metric scaling) from observations on this consumer's market choices. This is unrealistic, first because the analyst typically has observations on a consumer's choices in only a small number of market environments, often only one, and because markets are observed only over a limited range of conditions. For example, variations in historical product prices are limited by production costs and competition between products, and the dimensionality of possible product attributes is high, with only a limited range of bundles of attributes appearing in historically available products. However, different consumers generally face somewhat different observed market environments, and if one can maintain the consumer sovereignty assumption that consumer tastes are predetermined at the time of market choice, and assume plausibly that given s there is no ecological correlation of market environments and tastes, then observations across consumers can be used to estimate the distribution of tastes in the population. Further, in many applications it is reasonable to assume that consumers value products using hedonic effective prices that adjust market price for the attributes of the product; then the analysis can recover distributions of hedonic weights. This

¹⁹ Technically, retrospective welfare analysis should be conducted with a multi-period consumer model, with redress in the second period from harm in the first period. If the consumer has intertemporally separable utility, then the ideal MCE measure satisfies $V_{1b}(I_1) + V_2(I_2 - \text{MCE}) = V_{1a}(I_1) + V_2(I_2)$, where V_1 and V_2 are indirect utilities for the respective periods, and non-income arguments in indirect utility are suppressed. Applying the first mean value theorem for integrals, $\text{MCE} = [V_{1b}(I_2) - V_{1a}(I_1)]/\mu_2$, where μ_2 is a marginal utility of income in the second period. But the consumer will allocate income between periods to equate marginal utilities of income (without accounting for MCE), so that μ_2 will to a first approximation equal $\mu(I_a, \beta, \sigma)$. Consequently, the MCE defined in (25) approximates the two-period ideal. Further analysis of intertemporal utility to sharpen the definition of MCE is left to the reader.

will often be sufficient to infer the distribution of consumer utilities for new or modified products even if their attribute bundles are novel.

A more challenging recovery problem arises when markets are incomplete, due to transaction costs, asymmetric information that causes market failure through adverse selection and moral hazard, or failure to establish ownership and control of the distribution of some goods and services. For example, consumers cannot insure against some kinds of events, cannot directly purchase environmental amenities such as clean air and unpolluted beaches, and lack market opportunities that show their tastes for “existence goods” such as protecting endangered species or reducing global warming. If there is sufficient market redundancy, or if there are active margins where unmarketed and marketed goods are complements or substitutes, then it may be possible to recover indirectly preferences for unmarketed goods. For example, consumer preferences for environmental amenities are reflected in their willingness to travel to unpolluted beaches or move to neighborhoods with cleaner air. However, when preferences for unmarketed goods and services leave no market trace, they obviously cannot be recovered from market data. Experimental methods for directly eliciting stated preferences for these goods in hypothetical markets are successful in some marketing contexts, but sensitivity to context and framing can make experimental data unreliable; see Ben Akiva et al (2016), McFadden (2017), Miller et al. (2011). For the remainder of this section, we assume that there is sufficient market information to recover distributions of preferences in the population, and study the construction of aggregate measures of welfare. These aggregates may be sufficient for policy decisions, or sufficient to determine transfers that are judged appropriate to remedy harm to a class of consumers even if the compensation is not exact for each individual.

When a welfare analysis seeks to fulfill the transfers HCV, HEV, or MCE that in retrospective or prospective applications leave a class of consumers indifferent to the policy change, an obvious limitation is that an actual transfer to a consumer can be a function only of observed characteristics. It is common in applied welfare analysis to estimate welfare effects by postulating a *representative consumer* whose demands are close to the per capita market demands of a consumer class, calculating the transfer that keeps “representative” utility constant, and assuming that this per capita transfer could in principle be redistributed to keep the utility of each consumer in the class constant. A necessary and sufficient condition for the existence of a representative consumer meeting these conditions exactly is that the utilities of individuals in the class be representable in Gorman Polar Form with possibly heterogeneous committed expenditures but a common price deflator; see Chipman and Moore (1990), McFadden (2004). In (21), this requires that the X functions be independent of income, so that discrete choices will exhibit no neoclassical income effect and HCV, HEV, and MCE coincide. In practical fulfillment of compensating transfers, the policymaker faces a decision-theory problem in which there will be social losses from under or over-

compensation of individuals, and some (Bayesian) criterion must be applied to determine a loss-minimizing transfer rule. For example, if the policymaker has a quadratic social loss function, and a diffuse Bayesian prior, the optimal transfer to an individual equals the expected compensating transfer given observed characteristics. This suggests two rules in the case of partial observability. First, if transfers are fulfilled, prospectively or retrospectively, then they should equal the expected value of the exact compensating transfer given available information on the individual. Second, the impact of a policy change on a class of consumers in either prospective or retrospective applications should equal the expected value of the exact aggregate compensating transfers, with the appropriate compensating transfers determined by whether or not the transfers are hypothetical or fulfilled, and if the latter, whether this occurs before or after preference-maximizing choices in each scenario.

Relevant aggregates defined in Section 3 are the expected values $UMCE(s,k) = \mathbf{E}_{\beta,\sigma,\epsilon|s,k} MCE(s, k, \beta, \sigma, \epsilon)$ and $UMCE(s) = \mathbf{E}_{k,\beta,\sigma,\epsilon|s} MCE(s, k, \beta, \sigma, \epsilon)$, or uniform Hicksian measures such as $UHCV(s, k)$ and $UHCV(s)$. Section 3 describes a computational approach to forming the relevant aggregates using a synthetic population; this approach can accommodate any assumptions the analyst chooses on the properties of ϵ and the observed histories on which the welfare measures are conditioned. However, in selected cases, it is possible to reduce computation by forming analytic expectations with respect to ϵ . In the remainder of this section, we do this for the case where the scenario a choice is not observed, and three cases where this choice is observed: (A) psychometric noise is independent across the two scenarios, even for “brands” whose prices and attributes do not change; (B) components of ϵ across scenarios are independent for “brands” with changing attributes and prices, but are the same for benchmark “brands” whose attributes and prices do not change; and (C) all “brands” are present in both scenarios a and b , and $\epsilon_{ja} = \epsilon_{jb}$ for $j \in J_a = J_b$.

Consider the unconditional indirect utility function (23). Appendix A(b) shows that its expectation with respect to (β,σ,ϵ) , given history s , income I , and scenario m is

$$(29) \quad \mathbf{E}_{\beta,\sigma|s} \mathbf{E}_{\epsilon} \max_{j \in J_m} [I + v(I, p_{jm}, r_m, z_{jm}, \beta) + \sigma \epsilon_j] = I + \mathbf{E}_{\beta,\sigma|s} \{ \sigma \cdot \ln \sum_{j \in J_m} \exp(v_{jm}(I, \beta)/\sigma) + \sigma \cdot \gamma_0 \},$$

where γ_0 denotes Euler’s constant and $v_{jm}(I, \beta) \equiv v(I, p_{jm}, r_m, z_{jm}, \beta)$. Scaling and differencing for $m = b, a$,

$$(30) \quad UMCE(s) \equiv \mathbf{E}_{k,\beta,\sigma,\epsilon|s} MCE(s, k, \beta, \sigma, \epsilon) = \mathbf{E}_{\beta,\sigma|s} [I_b - I_a + \sigma \cdot \ln \frac{\sum_{j \in J_b} \exp(v_{jb}(I_b, \beta)/\sigma)}{\sum_{j \in J_a} \exp(v_{ja}(I_a, \beta)/\sigma)}] / \mu(I_a, \beta, \sigma),$$

where the *mean scaling factor* $\mu(I_a, \beta, \sigma)$ is a weighted harmonic mean of the marginal utilities of income (26), with the MNL choice probabilities from (22) as weights, and depends on σ through these weights,

$$(31) \quad \frac{1}{\mu(I_a, \beta, \sigma)} \equiv \mathbf{E}_{k|I_a, \beta, \sigma} \frac{1}{\mu_k(I_a, \beta)} = \sum_{k \in J_a} \frac{L_{km}(I_a, \beta, \sigma)}{\mu_k(I_a, \beta)} .$$

A Hicksian analogue of (30) is obtained by forming the expectation of (24) and solving

$$(32) \quad 0 = \mathbf{E}_{\beta, \sigma | s} [I_b - \text{UHCV}(s) - I_a + \sigma \cdot \ln \frac{\sum_{j \in J_b} \exp(v_{jb}(I_b - \text{UHCV}(s), \beta) / \sigma)}{\sum_{j \in J_a} \exp(v_{ja}(I_a, \beta) / \sigma)}] .$$

A formula for UHEV(s) is more complicated. If scenario a optimal choices before and after the income transfer change due to income effects, then (29) is replaced by a complex expression from Appendix A(e). However, if the marginal utility of income $\mu_k(I_a, \beta)$ from (26) is independent of I_a and k remains the optimal choice after the transfer, then from (27), UHEV(s) = UMCE(s). Of course, if discrete choice exhibits no income effects, then the definitions above satisfy UMCE(s) = UHCV(s) = UHEV(s).

Next consider situations in which choice of alternative k in scenario a is observed, $\delta_k(I_a) = 1$. We consider the three cases (A), (B), and (C) defined above.

(A) Suppose psychometric noise is independent across the two scenarios for all alternatives, including alternatives whose prices and attributes do not change (e.g., the “no purchase” alternative). This supposition is appropriate if this noise is whimsy that is memoryless and transient, and arises anew every time a choice is made, so there is no carryover from one choice scenario to the next. From Appendix A(b), expected utility given the utility-maximizing choice k is the same as (29) when $m = a$ and $I = I_a$. Further, independence implies that (29) will also apply when $m = b$ and $I = I_b$. Then, (32) with conditioning on k added continues to hold in case (A), with a solution defining a uniform transfer UHCV(s, k) for each k . Further, (30) is altered only by substituting the choice- k specific scale factor (26), giving

$$(33) \quad \text{UMCE}(s, k) = \mathbf{E}_{\beta, \sigma | s, k} [I_b - I_a + \sigma \cdot \ln \frac{\sum_{j \in J_b} \exp(v_{jb}(I_b, \beta) / \sigma)}{\sum_{j \in J_a} \exp(v_{ja}(I_a, \beta) / \sigma)}] / \mu_k(I_a, \beta) .$$

Note that (30) is the expectation of (33) with respect to the MNL probability $L_{km}(I_a, \beta, \sigma)$. While the formula (33) depends on k only due to the scale factor, its expectation conditioned on choice k in a population with heterogeneous observed environments will in general vary substantially with k due to selection on the environments that yield this choice.

(B) Suppose the alternatives in $J_a \cup J_b$ can be partitioned into a set A of alternatives with indices that appear only in J_a ; a set B of alternatives that appear only in J_b ; and a set C of “benchmark” alternatives appearing in both

scenarios that satisfy $v_{ja}(I_a, \beta) = v_{jb}(I_b, \beta)$ for $j \in C$; this requires that their attributes and prices not change, and if these v_{jm} depend on income, that $I_a = I_b$. This setup is appropriate when the psychometric noise ε_j for product j is a durable unobserved attribute of this product, but the “brand” is reindexed and evaluated anew with a statistically independent draw of psychometric noise when the scenario change leads to changes in its attributes or price. As a result of this assumption, the $v_{jm}(I, \beta) \equiv v_j(I, \beta)$ for $j \in A \cup B \cup C$ do not depend on the scenario m .

Utilizing the conditional expectation formulas in Appendix A(c),

$$(34) \quad \text{UMCE}(s, k) = \mathbf{E}_{\beta, \sigma | s, k} \left[I_b - I_a + \max_{j \in B \cup C} (v_j(I_b, \beta) + \sigma \varepsilon_j) - (v_k(I_a, \beta) + \sigma \varepsilon_k) \right] / \mu_k(I_a, \beta)$$

$$= \mathbf{E}_{\beta, \sigma | s, k} \frac{1}{\mu_k(I_a, \beta)} \left[I_b - I_a + \sigma \cdot \ln \left(\frac{\sum_{j \in B \cup C} \exp(v_j(I_b, \beta) / \sigma)}{\sum_{j \in A \cup C} \exp(v_j(I_a, \beta) / \sigma)} \right) \right] + \begin{cases} \mathbf{E}_{\beta, \sigma | s, k} \frac{L(C|A, C)}{L(A|A, C)} \frac{\sigma \cdot \ln(1 - L(A|A, B, C))}{\mu_k(I_a, \beta)} & \text{if } k \in A \\ -\mathbf{E}_{\beta, \sigma | s, k} \frac{\sigma \cdot \ln(1 - L(A|A, B, C))}{\mu_k(I_a, \beta)} & \text{if } k \in C \end{cases},$$

$$\text{where } L(A|A, B, C) = \frac{\sum_{j \in A} e^{v_j(I_b, \beta) / \sigma}}{\sum_{j \in A \cup B \cup C} e^{v_j(I_a, \beta) / \sigma}} \text{ and } L(C|A, C) = \frac{\sum_{j \in C} e^{v_j(I_b, \beta) / \sigma}}{\sum_{j \in A \cup C} e^{v_j(I_a, \beta) / \sigma}}.$$

The left-hand expectation term in the last line of (34) coincides with the expression (33) for $\text{UMCE}(s, k)$ when the psychometric noise in scenario b is independent of the psychometric noise in scenario a . The right-hand expectation term is an adjustment for the effect of the conditioning event on the expected maximum utility from $B \cup C$, downward if $k \in A$ and upward if $k \in C$. This expectation incorporates the effects of selection, which can be powerful if σ is large: Many choices from A will come from favorable draws of psychometric noise even when observed attributes make these alternatives unattractive. Then, regression to the mean in draws of psychometric noise will tend to make alternatives in B less desirable than their analogues in A even if they are objectively better. In contrast, when the analogues in B of alternatives in A objectively improve, choices from C that result from a favorable draw will lead to an even better expected outcome in scenario b since alternatives with this draw remain available. If the scale factors $\mu_k(I_a, \beta)$ vary with k , then the interaction of selection and income effects no longer gives the result that (30) with scale factor (31) equals the expectation of (34) with respect to k . For example, if alternatives in A have $\mu_k(I_a, \beta) > 1$, and alternatives in C have $\mu_k(I_a, \beta) = 1$, then the expectation of (34) with respect to k exceeds $\text{MCE}(s)$ from (30). Equation (34) can also be adapted to calculate $\text{UHCV}(s, k)$ for this case: Reduce income I_b in scenario b by $\text{UHCV}(s, k)$, with this quantity adjusted to that (34) equals zero.

(C) Suppose $\mathbf{J}_a = \mathbf{J}_b = \mathbf{J} = \{0, \dots, J\}$, $\delta_k(I_a) = 1$, and $\varepsilon_a = \varepsilon_b = \varepsilon$, so that all alternatives are indexed the same and have the same psychometric noise in both scenarios. It is possible to obtain analytic formulas for $\text{MCE}(s, k)$ under

quite general conditions in which the differences $v_{jb}(I_b^*, \beta) - v_{ja}(I_a^*, \beta)$, evaluated at income levels that may differ from I_a or I_b respectively due to transfers, vary across multiple alternatives. Appendix A(e) provides formulas that can be assembled to program this calculation, but these are too complex to be useful for comparison to the previous cases. We instead consider the special circumstance in which $I_a = I_b = I$ and the scenario affects only product J , so $v_{jm}(I, \beta)$ is independent of m for $j < J$. For this case, Appendix A(d) implies the following results:

If $k = J$, then

$$(35) \quad \text{UMCE}(s, J) = \mathbf{E}_{\beta, \sigma | s, J} \frac{1}{\mu_J(I, \beta)} \begin{cases} v_{Jb}(I, \beta) - v_{Ja}(I, \beta) & \text{if } v_{Jb}(I, \beta) > v_{Ja}(I, \beta) \\ \frac{\sigma}{L_{Ja}} \cdot \ln \frac{\sum_{j \in J} e^{v_{jb}(I, \beta)/\sigma}}{\sum_{j \in J} e^{v_{ja}(I, \beta)/\sigma}} & \text{if } v_{Jb}(I, \beta) < v_{Ja}(I, \beta) \end{cases},$$

while if $k < J$,

$$(36) \quad \text{UMCE}(s, k) = \mathbf{E}_{\beta, \sigma | s, k} \frac{1}{\mu_k(I, \beta)} \begin{cases} -\frac{L_{Ja}}{L_{ka}} (v_{Jb}(I, \beta) - v_{Ja}(I, \beta)) + \frac{\sigma}{L_{ka}} \cdot \ln \frac{\sum_{j \in J} e^{v_{jb}(I, \beta)/\sigma}}{\sum_{j \in J} e^{v_{ja}(I, \beta)/\sigma}} & \text{if } v_{Jb}(I, \beta) > v_{Ja}(I, \beta) \\ 0 & \text{if } v_{Jb}(I, \beta) < v_{Ja}(I, \beta) \end{cases},$$

where $L_{ka} = e^{v_{ka}(I, \beta)/\sigma} / \sum_{j \in J} e^{v_{ja}(I, \beta)/\sigma}$. As in case (B), this formula can be adapted to solve for the transfer $\text{UHCV}(s, k)$ that when fulfilled makes a subsequent UMCE zero, while computation of $\text{UHEV}(s, k)$ is in general more complicated.

We consider an example where due to a fixing agreement the price of a single product, say a tablet computer, is higher in scenario a than in scenario b . For the alternative configurations of J_a , J_b , ϵ_a , and ϵ_b , we estimate in Table 2 the measures $\mathbf{E}_{\beta, \sigma, \epsilon | s, k} \text{MCE}(s, k, \beta, \sigma, \epsilon)$, $\mathbf{E}_{\beta, \sigma, \epsilon | s, k} \text{HCV}(s, k, \beta, \sigma, \epsilon)$, and $\mathbf{E}_{\beta, \sigma, \epsilon | s, k} \text{HEV}(s, k, \beta, \sigma, \epsilon)$ in a synthetic population, and the measure $\text{UMCE}(s, k)$. Suppose the product $J = 1$ has price p_{1m} in scenario m . Suppose the “no purchase” alternative has $p_{0m} = 0$. Suppose consumers have utilities of the form (21) with $v_{1m}(I, \beta) = \beta_1 I + \beta_2 - p_{1m}$ for alternatives where the product is purchased, and $v_{0m}(I, \beta) = 0$ for “no purchase” alternative, for scenarios $m = a, b$. Then $\mu_0(I_a, \beta) = 1$ and $\mu_1(I_a, \beta) = 1 + \beta_1$. The psychometric noise cases we consider are (A) independent noise across scenarios, represented by $J_a = \{0, 1\}$ and $J_b = \{2, 3\}$, with $j = 0, 2$ corresponding to “no purchase” and $j = 1, 3$ corresponding to “purchase”; (B) $J_a = \{0, 1\}$ and $J_b = \{0, 3\}$, with $j = 0$ corresponding to a common “no purchase” and $j = 1, 3$ corresponding to “purchase”; and (C) $J_a = J_b = \{0, 1\}$, with $j = 0$ corresponding to a common “no purchase” and $j = 1$ to a common purchase. Suppose that $\beta_1 = 0.002$ and $\sigma = 9$ are fixed parameters, and that β_2 is normal with mean zero and standard deviation 60. The choice probabilities are then mixed logit, with $P_{0m}(I) = \mathbf{E}_{\beta} \frac{1}{1 + \exp(v_{jm}(I, \beta)/\sigma)}$ for non-purchase of the product j in scenario m . Suppose the consumer faces $p_{1a} = \$110$

and $p_{1b} = \$90$, and the base income is $I = \$50,000$. The probabilities of buying the product in a synthetic population of 10,000 are $P_{1a}(50000) = 0.430$, $P_{1b}(50000) = 0.555$, and $P_{1a}(56000) = 0.505$. These probabilities imply an arc income elasticity of 1.45 and an arc price elasticity of -1.59 for the given market changes. The table shows first that for this example, $E_{\beta,\sigma,\epsilon|s,s,k}$, $MCE(s,k,\beta,\sigma,\epsilon)$, $E_{\beta,\sigma,\epsilon|s,s,k}$, $HCV(s,k,\beta,\sigma,\epsilon)$, and $E_{\beta,\sigma,\epsilon|s,k}$ $HEV(s,k,\beta,\sigma,\epsilon)$ estimated in the synthetic population are almost the same. This result is consistent with the conclusion of Willig (1976) that income effects are typically small. The value of UMCE using an analytic expectation with respect to ϵ differs modestly from the synthetic population estimate of $E_{\beta,\sigma,\epsilon|s,s,k}$, $MCE(s,k,\beta,\sigma,\epsilon)$, but the difference is well within sampling error. The Marshallian consumer surplus, estimated here using the trapezoid rule, is nearly identical to UMCE.

The example suggests that UMCE will be an adequate approximation when $E_{\beta,\sigma,\epsilon|s,k_a}$ $HCV(s,k,\beta,\sigma,\epsilon)$ is the ideal measure. However, the closeness of UMCE and expected HCV is sensitive to the magnitude of the change in price in the two scenarios, and larger changes can lead to a gap between these measures. In short, when UMCE is used as an approximation to expected HCV, it is desirable to use synthetic population methods with large samples to check the quality of the approximation.

Table 2. Comparisons of Welfare Measures (\$pp)

	Case A	Case B	Case C	$\sigma = 0$
Total Population				
Marshallian consumer surplus	9.819	9.799	9.781	9.854
UMCE (analytic E_ϵ)	9.848	9.858	9.802	9.840
MCE (synthetic population)	9.886	9.806	9.755	9.840
HCV (synthetic population)	9.883	9.803	9.753	9.837
HEV (synthetic population)	9.886	9.806	NC	9.840
Class of Product Purchasers				
UMCE (analytic E_ϵ)	18.568	18.283	19.960	19.960
MCE (synthetic population)	18.609	18.368	19.960	19.960
HCV (synthetic population)	18.610	18.368	19.960	19.960
HEV (synthetic population)	18.609	18.368	NC	19.960
Class of Non-Purchasers				
UMCE (analytic E_ϵ)	3.305	3.535	2.180	2.240
MCE (synthetic population)	3.340	3.382	2.097	2.240
HCV (synthetic population)	3.335	3.375	2.093	2.235
HEV (synthetic population)	3.340	3.382	NC	2.240

There is variation in the welfare measures when one moves from Case A with independent disturbances to Case C with common disturbances. In particular, Cases A and B attribute less welfare gain to purchasers and

more welfare gain to non-purchasers than does Case C and a $\sigma = 0$ case with no psychometric noise. Then, the assumptions made on the persistence of psychometric errors across scenarios makes some difference.

Table 3. Effect of Psychometric Noise on Distribution of Welfare Changes (\$pp)

UMCE at	Case A			Case B			Case C		
	Buyers	Non-buyers	Total	Buyers	Non-buyers	Total	Buyers	Non-buyers	Total
$\sigma = 0$	19.960	2.240	9.840	19.960	2.240	9.840	19.960	2.240	9.840
$\sigma = 9$	18.568	3.305	9.848	18.283	3.535	9.858	19.960	2.180	9.802
$\sigma = 36$	13.641	6.943	9.895	5.046	14.057	10.086	19.960	1.683	9.737
$\sigma = 64$	11.687	8.463	9.927	-9.086	25.560	10.370	19.960	1.224	9.734

Table 3 continues the example with different scale factors σ , and shows that at high levels of σ relative to the observed changes in the scenarios, the effects of selection on psychometric noise can drastically alter the distribution of welfare gains between purchasers and non-purchasers. We infer from this table that unless there is compelling evidence to support the case (B) assumptions, they should be rejected in favor of case (A) or case (C) assumptions that more closely approximate a model in which neoclassical tastes, heterogeneous across consumers but durable within each consumer, describe choice behavior without significant added psychometric noise. Finally, there is a substantial advantage in simplicity for the analytic expectations coming out of the case (A) compared to case (C), suggesting that case (C) be used only if there is persuasive evidence for durable psychometric noise.

6. DECISION-UTILITY VERSUS EXPERIENCED-UTILITY

Decision-utility is defined as the objective function that the neoclassical consumer optimizes in making her market choices, the function that can be recovered (with money-metric scaling) from sufficiently rich observations on these market choices. The foundations of welfare theory restated in Section 2 assume that decision-utility is a direct and complete measure of well-being. In reality, anticipated decision-utility and realized experienced-utility can differ. The most straightforward case, covered by neoclassical theory, is decision-making under uncertainty where the decision utility function equals the expectation over objective probabilities of a utility function of outcomes, and the experienced utility function equals this utility function evaluated at the realized outcome. For example, consumers may be uncertain about attributes of alternatives such as product durability,

so that buying a product is equivalent to buying a lottery ticket on its attributes. Under von Neumann-Morgenstern assumptions on utility, sufficiently rich market observations on choice among risky prospects will suffice to recover the utility function of outcomes.

Beyond neoclassical decision-making under uncertainty, there are a number of factors that can cause gaps between decision-utility and experienced-utility: (1) misperceptions of shrouded, ambiguous, or misleadingly promoted product attributes, (2) unrealistic personal probability judgements on uncertain events, (3) whims and psychometric noise that induce instabilities in tastes, (4) factors that influence the sensation of well-being but do not influence market choices, such as provision of pure public goods and services, (5) inconsistencies in preferences, such as time-inconsistent discounting and unanticipated habit-formation or addiction, and (6) flaws in the process of utility maximization, such as reference point bias and hypersensitivity to recent experience.

When there are gaps between decision-utility and experienced-utility, which should be used to measure well-being? Roughly, welfare measures based on decision-utility focus on equity in opportunity, while those based on experienced-utility focus on equity in outcomes. It may seem evident that consumer perceptions and decision-making in markets are simply instruments to achieve final outcomes, so experienced utility should be at the core of welfare assessment. However, there are complicating factors. First, when consumers are fully and accurately informed about the prospects and contingencies they face, and there are sufficient contingent markets so that they can insure against risks if they choose, then they have it in their own hands to make informed choices and live with the consequences of these choices. Further, interventions based on experienced utility can introduce “moral hazard” in which the anticipation of ex post remedies for bad outcomes leads consumers to take excess risks and be less diligent in their decisions, particularly by failing to take steps to avoid or mitigate harm. Then, for fully informed consumers facing complete contingent markets, policies should arguably be evaluated in terms of decision-utilities. However, when consumers are poorly informed or lack opportunities to manage risks, ex post equity is a social concern, and/or consumers are unable to look after their own interests, interventions by a benevolently paternalistic regulator based on experienced utility may be appropriate. In general, it will be important to know how perceptions, decision utility, and experienced utility are linked. Misperceptions of attributes, biased personal probabilities, and instabilities in tastes arising from psychometric noise, listed above as sources (1)-(3) of gaps between perception and experience, do not substantially alter the neoclassical preference structure, and can in principle be accounted for starting from decision utility and correcting these factors. For example, it should be straightforward in principle to correct consumer misperceptions arising from supplier misrepresentation of product attributes. In practice, identification and recovery of personal perceptions and probabilities may overburden market data and require extra-market experimental observations. Other

elements entering experienced utility such as sources (4)-(6) of gaps between perceptions and experience, particularly factors that leave no trace in market choices, can confound choice behavior so that there may be no identifiable decision-utility or linked experienced utility that capture consumer well-being. Then recovery of experienced-utility will often be beyond the capacity of economists using customary market-based techniques. Ben-Akiva, McFadden, and Train (2016) discuss experimental methods for direct elicitation of preference that can be used in principle to address these identification and recovery tasks. These are widely used in market research to forecast demand for new products and the value of extra-market resources, with varying degrees of reliability; see McFadden (2017). The focus of this paper is welfare analysis using market observations, and we will not here consider further experimental, cognitive, or neurological approaches to direct measurement of well-being.

An additional issue in welfare assessment arises when the impacts of real product defects are confounded by flaws in perception. For example, a homeowner in a neighborhood stigmatized by an environmental hazard coming from improper disposal of industrial waste is harmed by a reduced market value when refinancing, and this harm is compounded if misperceptions exaggerate the environmental hazard. Thus, a consumer may incur real harm from misperceptions of others, whether or not this consumer has accurate perceptions. What components of harm in this case are a compensable liability of the polluter is a question of the applicable law.

To facilitate analysis of the consequences of gaps between anticipation and experience, let superscript “d” denote decision utility and superscript “e” denote experienced utility. From (21), choice in scenario m at income I then maximizes $v_{jm}^d(I, \beta^d) + \sigma^d \varepsilon_{jm}^d$; let $\delta_{jm}(I, \beta^d, \sigma^d, \varepsilon^d)$ denote an indicator for this choice, and $P_{jm}(I, \beta^d, \sigma^d)$ the probability of this choice given β^d, σ^d . Let $v_{jm}^e(I, \beta^e) + \sigma^e \varepsilon_{jm}^e$ denote the experienced utility obtained from choice j . The application will determine the structure of $v_{jm}^e(I, \beta^e)$ and its linkage to $v_{jm}^d(I, \beta^d)$, and the mappings from (β^d, σ^d) to (β^e, σ^e) and from ε^d to ε^e .²⁰ The decision utility and experienced utility from a choice situation with income I are, in this notation,

$$(37) \quad u_m^d(I, \beta^d, \sigma^d) = I + \max_{j=0, \dots, J_m} [v_{jm}^d(I, \beta^d) + \sigma^d \varepsilon_{jm}^d]$$

and

²⁰ One case which is straightforward occurs when $v_{jm}^d(I)$ and $v_{jm}^e(I)$ differ only because of differences in observed anticipated and experienced product attributes, z_{jm}^d and z_{jm}^e , due to say false advertising of attributes, and $\beta^d = \beta^e$, $\sigma^d = \sigma^e$, and $\varepsilon_m^d = \varepsilon_m^e$. Cases that are more challenging for economic analysis occur when either anticipated or experienced attributes are unobserved, or v^e and v^d differ due to optimization errors and volatility in tastes. In such cases, the analyst will often have no recourse other than using extra-market observations such as experimental elicitation of stated preferences, with attendant questions of reliability.

$$(38) \quad u_m^e(I, \beta^e, \sigma^e, \beta^d, \sigma^d) = I + \sum_{j=0}^{J_m} \delta_{jm}(I, \beta^d, \sigma^d, \boldsymbol{\varepsilon}^d) [v_{jm}^e(I, \beta^e) + \sigma^e \varepsilon_{jm}^e] \\ \equiv u_m^d(I, \beta^d, \sigma^d) + \sum_{j=0}^{J_m} \delta_{jm}(I, \beta^d, \sigma^d, \boldsymbol{\varepsilon}^d) [v_{jm}^e(I, \beta^e) - v_{jm}^d(I, \beta^d) + \sigma^d \varepsilon_{jm}^d - \sigma^e \varepsilon_{jm}^e],$$

where experienced utility in the last expression equals anticipated utility plus a correction that comes from differences in anticipated and realized attributes and tastes plus a correction that comes from differences in anticipated and realized psychometric noise. Combined with (25) defining decision-utility MCE^d , (38) implies the experienced-utility welfare measure, given (β^d, σ^d) , (β^e, σ^e) , $\boldsymbol{\varepsilon}_m^d$ and $\boldsymbol{\varepsilon}_m^e$ for $m = a, b$, and $\delta_{ka}(I, \beta^d, \sigma^d, \boldsymbol{\varepsilon}^d) = 1$,

$$(39) \quad \mu_k(I_a, \beta) \cdot MCE^e(s, k, \beta^e, \sigma^e, \boldsymbol{\varepsilon}_a^e, \boldsymbol{\varepsilon}_b^e, \beta^d, \sigma^d, \boldsymbol{\varepsilon}_a^d, \boldsymbol{\varepsilon}_b^d) = \mu_k(I_a, \beta) \cdot MCE^d(s, k, \beta^d, \sigma^d, \boldsymbol{\varepsilon}_a^d, \boldsymbol{\varepsilon}_b^d) \\ + v_{kb}^e(I_b, \beta^e) - v_{kb}^d(I_b, \beta^d) + \sigma^e \varepsilon_{kb}^e - \sigma^d \varepsilon_{kb}^d - v_{ka}^e(I_a, \beta^e) + v_{ka}^d(I_a, \beta^d) - \sigma^e \varepsilon_{ka}^e + \sigma^d \varepsilon_{ka}^d,$$

and when $\sigma^d \varepsilon_{jm}^d = \sigma^e \varepsilon_{jm}^e$ for $m = a, b$,

$$(40) \quad \mu_k(I_a, \beta) \cdot MCE^e(s, k, \beta^e, \sigma^e, \boldsymbol{\varepsilon}_a^e, \boldsymbol{\varepsilon}_b^e, \beta^d, \sigma^d, \boldsymbol{\varepsilon}_a^d, \boldsymbol{\varepsilon}_b^d) = \mu_k(I_a, \beta) \cdot MCE^d(s, k, \beta^d, \sigma^d, \boldsymbol{\varepsilon}_a^d, \boldsymbol{\varepsilon}_b^d) \\ + v_{kb}^e(I_b, \beta^e) - v_{kb}^d(I_b, \beta^d) - v_{ka}^e(I_a, \beta^e) + v_{ka}^d(I_a, \beta^d).$$

In general, economists should be very cautious in applying the traditional welfare calculus when behavioral factors are needed in decision-utility to explain behavior; as transfers that restore decision utility levels can have unreliable and unintended effects on experienced well-being. If anticipated tastes are an unreliable guide to realized tastes, this is a challenge to the foundations of welfare economics; see Lowenstein and Ubel (2008), Thaler and Sunstein (2003, 2008), McFadden (2014), Train (2015). There is currently no accepted general welfare theory for non-neoclassical consumers who have shifts between anticipated and realized tastes, even though the random decision-utility setup itself can accommodate many non-neoclassical elements. However, there may be some special circumstances and assumptions that overcome this limitation. For example, differences in “as is” or “but for” (z_{jm}^d, p_{jm}^d) and (z_{jm}^e, p_{jm}^e) may be limited to identifiable misperceptions such as misinformation about product attributes, and the joint distribution of anticipated and realized tastes may by assumption be generated through limited differences such as personal misjudgments on the probabilities of contingent events or biases in risk preferences and time discounts used in making decisions. If it is plausible that such limited shifts in tastes can be fully described and modeled using specific external evidence, then welfare analysis based on (40) may be justified.

An example of consumer behavior that appears to be distorted by unrealistic personal probability judgements is consumer choice of health insurance policies. An argument, simplified from Heiss, McFadden, and Winter (2013) and McFadden and Zhou (2015), shows that misperceptions can be identified and corrected in some cases. Suppose consumers face stochastic medical expenses c , and have the subjective perception that these have

a distribution $K^d(c)$ with a mean μ_d and variance s_d^2 . Suppose they have a menu of insurance alternatives $j = 0, \dots, J$ with plan j characterized by a premium p_j and a copayment rate r_j , with $p_0 = 0$ and $r_0 = 1$. Suppose their decision-utility is a money-metric transformation of a constant-absolute-risk-aversion (CARA) expected utility function,

$$(41) \quad u_j = \frac{-1}{\beta} \ln \int_{c=0}^{+\infty} \exp\left(-\beta(I - p_j - r_j c)\right) K^d(dc) + \sigma \varepsilon_j \equiv I - p_j - \kappa^d(\beta r_j) / \beta + \sigma \varepsilon_j,$$

where I is income, κ^d is the cumulant generating function of K^d , β is a risk-aversion parameter with a probability distribution in the population, and the parameter σ scales psychometric noise ε_j . Replacing the cumulant generating function κ^d in (38) with a quadratic approximation gives a utility $u_j = I - p_j - \mu_d r_j - \frac{1}{2} s_d^2 \beta r_j^2 + \sigma \varepsilon_j$ of the form (21). Suppose $(\ln \beta, \ln \sigma)$ is distributed bivariate normal, and the psychometric noise is i.i.d. EV1. Then observations on consumer insurance choices in real or experimental markets allows estimation of the parameters of the bivariate normal distribution, and μ_d , and s_d^2 . Observations on objective probabilities $K^e(c)$ for health expenses allow estimation of μ_e and s_e^2 . Then specialization of (40) using (41) and the quadratic approximations to the cumulant generating functions κ^d and κ^e allow estimation of the money-metric loss in consumer utility arising from poor choices due to misperception of medical expense risk.

An open question about measurement of well-being of consumers who have behavioral elements in their decision-making and a gap between decision-utility and experienced-utility is whether experienced utility could be elicited directly in conjoint analysis experiments, either through experiments used to uncover the components of experienced utility, or through conjoint elicitation methods such as elicitation of stated personal probabilities. Established experimental designs for such elicitations are not available now, and there are major scientific challenges to their development, particularly known biases in personal probability judgements and the problem of verification, but there will be high payoffs to future scientific breakthroughs in these areas.

7. WELFARE CALCULUS FOR COMMON POLICY PROBLEMS

Suppose mixed MNL choice probabilities of the form (22), along with the associated parameter α of a population distribution of taste parameters $F(\beta, \sigma | \alpha)$ and a money-metric utility of the form (21), have been estimated from choice data collected in real or hypothetical markets. Using these estimates, *prospective* benefit-cost analysis using decision utility can be carried out by solving (24) or evaluating (25) for each consumer in a synthetic population defined by draws of s , parameters (β, σ) from $F(\beta, \sigma | s, \alpha)$, and psychometric noise ε_a and ε_b . The measures such as HCV, HEV, or MCE can then be averaged over the synthetic consumers falling into classes

defined by restrictions on s , with the law of large numbers operating to ensure reliable estimates of the net transfer to the class that when optimally distributed leaves its members indifferent to the policy change. Alternately, one can concentrate on estimating a UMCE measure (30) for this class. To simplify notation, suppress the “d” superscript for decision utility. Let \mathbf{C} denote the set of alternatives whose attributes are unchanged by a shift from policy a to policy b . By construction, \mathbf{C} is a proper subset of \mathbf{J}_a and \mathbf{J}_b which contain alternatives whose attributes do not change, and \mathbf{C} always contains at least $j = 0$. Then (30) can be rewritten as

$$(42) \quad \text{UMCE}(s) \equiv \mathbf{E}_{\beta, \sigma | s} \left[I_b - I_a + \sigma \cdot \ln \frac{L_{\mathbf{C}a}(I_a, \beta, \sigma)}{L_{\mathbf{C}b}(I_b, \beta, \sigma)} \right] / \mu(I_a, \beta, \sigma)$$

where $L_{\mathbf{C}m}$ is the logit probability at random parameters (β, σ) of choice from \mathbf{C} in scenario m . For example, introducing a set \mathbf{B} of new products with attributes included in $(\mathbf{x}_b, \mathbf{p}_b)$, keeping unchanged the attributes of existing products in \mathbf{C} , is $\text{UMCE}_b(s) = \mathbf{E}_{\beta, \sigma | s} [I_b - I_a - \sigma \cdot \ln L_{\mathbf{C}b}(I_b, \beta, \sigma)] / \mu(I_a, \beta, \sigma)$.

For small policy changes and $\mathbf{J}_a = \mathbf{J}_b = \mathbf{J}$, a Taylor’s expansion of the first form of (42) in variations $\Delta v_j \equiv v_{jb}(I_b, \beta) - v_{ja}(I_a, \beta) \equiv \Delta X_j \beta - \Delta p_j$, where $\Delta X_j \equiv X(I_b - p_{jb}, \mathbf{r}_b, z_{jb}) - X(I_a - p_{ja}, \mathbf{r}_a, z_{ja})$ and $\Delta p_j \equiv p_{jb} - p_{ja}$, gives the approximation,

$$(43) \quad \text{UMCE}(s) = \mathbf{E}_{\beta, \sigma | s} [I_b - I_a + \sum_{j \in \mathbf{J}} [L_{ja}(I_a, \beta, \sigma) \Delta v_j + O((\Delta v_j)^2 / \sigma)]] / \mu(I_a, \beta, \sigma).$$

Another useful approximation, due to Doug MacNair, applies the expansion $\ln(1 - y) = -y + O(y^2)$ to (42) with $L_{\mathbf{B}b} = 1 - L_{\mathbf{C}b}$ and $L_{\mathbf{A}a} = 1 - L_{\mathbf{C}a}$ the probabilities of choosing products whose attributes change, to obtain

$$(44) \quad \text{UMCE}_b(s) = \mathbf{E}_{\beta, \sigma | s} [I_b - I_a + \sigma \cdot [L_{\mathbf{B}b}(I_b, \beta, \sigma) - L_{\mathbf{A}a}(I_a, \beta, \sigma) + O((1 - L_{\mathbf{C}m})^2)]] / \mu(I_a, \beta, \sigma).$$

When σ and $\mu = \mu(I_a, \beta, \sigma)$ are homogeneous in the population, (44) has a particularly simple form,

$$(45) \quad \text{UMCE}_b(s) \approx \left[\frac{I_b - I_a}{\mu} + \frac{\sigma}{\mu} \mathbf{E}_{\beta | s} [L_{\mathbf{B}b}(I_b, \beta, \sigma) - L_{\mathbf{A}a}(I_a, \beta, \sigma)] \right],$$

the income difference scaled by μ plus the difference in the (σ/μ) -scaled full population market share of consumers choosing the products affected by the policy change. For example, if a set \mathbf{A} of products is “new” in scenario b , and income and the attributes of the remaining products in \mathbf{C} are unchanged, then $\text{UMCE}_b(s) \approx \frac{\sigma}{\mu} \mathbf{E}_{\beta | s} L_{\mathbf{A}b}(I_b, \beta, \sigma)$, a scaled market share of the new products.

Consider a policy that affects attributes and price of products, and let $X_{j\lambda} = X_{ja} + \lambda \Delta X_j$, $p_{j\lambda} = p_{ja} + \lambda \Delta p_j$, and $I_\lambda = I_a + \lambda(I_b - I_a)$ for $\lambda \in [0, 1]$ and $j = 1, \dots, J$ denote a linear path that achieves this change. Let UMCE_λ denote (42)

evaluated at point λ on this path. Let $v_{j\lambda} \equiv X_{j\lambda}\beta - p_{j\lambda}$ and $L_{j\lambda} = e^{v_{j\lambda}/\sigma} / \sum_{i \in J} e^{v_{i\lambda}/\sigma}$. Then $\partial v_{j\lambda} / \partial \lambda = \Delta X_j \beta - \Delta p_j$, and since the numerator of $L_{C\lambda}$ does not vary with λ ,

$$(46) \quad \frac{d \text{UMCE}_{\lambda}(s)}{d\lambda} = \mathbf{E}_{\beta, \sigma | s} [I_b - I_a + \sum_{j=1}^J L_{j\lambda} (\Delta X_j \beta - \Delta p_j)] / \mu(I_a, \beta, \sigma).$$

Then the incremental change in UMCE is a demand-weighted average of the changes $\Delta X_j \beta - \Delta p_j$ in the systematic components of utility. First, consider the common circumstance where ΔX_j and Δp_j do not depend on s ; this will be the case for example for a product offered in a national market where interactions of product attributes and individual characteristics are not needed to explain choice behavior. Then, (46) reduces to

$$(47) \quad \frac{d \text{UMCE}_{\lambda}(s)}{d\lambda} = \mathbf{E}_{\beta, \sigma | s} [I_b - I_a + \sum_{j=1}^J [\Delta X_j \bar{\beta}_{j\lambda} - \Delta p_j] L_{j\lambda}(I_{\lambda}, \beta, \sigma)] / \mu(I_a, \beta, \sigma),$$

where $\bar{\beta}_{j\lambda} = \frac{\mathbf{E}_{\beta, \sigma | s, \alpha} \beta L_{j\lambda}(I_{\lambda}, \beta, \sigma)}{\mathbf{E}_{\beta, \sigma | s, \alpha} P_j(L_{j\lambda}, \beta, \sigma)}$ denotes the mean of β among consumers who choose j when the alternatives are characterized by $(\mathbf{X}_{\lambda}, \mathbf{p}_{\lambda})$. In this case, $\bar{\beta}_{j\lambda}$ gives WTP for attribute changes that translate directly into incremental compensating variation. In the special sub-case of changes that are uniform in j , $\Delta x_j = \Delta x_1$ for $j \neq 0$, $\bar{\beta}_{j\lambda}$ is independent of j and is the mean of β among all buyers. In the special case that the relevant components of β are homogeneous, then $\bar{\beta}_{j\lambda} = \beta$ in the corresponding components, and these coefficients are unequivocal measures of “part-worths”. More generally, obtaining $\bar{\beta}_{j\lambda}$ is a calculation that requires estimates of both $F(\beta, \sigma | s, \alpha)$ and $L_{j\lambda}(I_{\lambda}, \beta, \sigma)$.

Second, when the relevant components of β are homogeneous, but Δx_j and Δp_j are heterogeneous over s , (46) reduces to

$$(48) \quad \frac{d \text{UMCE}_{\lambda}(s)}{d\lambda} = \mathbf{E}_{\beta, \sigma | s} [I_b - I_a + (\overline{\Delta X_j})_{\lambda} \cdot \beta - (\overline{\Delta p_j})_{\lambda}] / \mu(I_a, \beta, \sigma),$$

so the relevant components of β give WTP for mean attribute changes among consumers choosing j when product features are described by $(\mathbf{X}_{\lambda}, \mathbf{p}_{\lambda})$. Third, when β is heterogeneous and the ΔX_j are heterogeneous over the population (i.e., vary with s), the relationship between values of β and $\mathbf{E}_{\beta, \sigma | s, \alpha} \text{UMCE}_b(\beta, \sigma)$ is more complex; (42) requires a calculation that handles selection driven by both consumer history and taste heterogeneity.

The scaling parameter σ appears in (42) and (44) to have a prominent direct role in determining the level of $\text{UMCE}(s)$, but (43) indicates that this is offset elsewhere, so that the final impact of σ is only indirect, through its influence on the choice probabilities. To see this more generally, write (42) as

$$(49) \quad \text{UMCE}(s) \equiv \int_{\lambda=0}^1 \frac{d\text{UMCE}_{\lambda}(s)}{d\lambda} d\lambda = \mathbf{E}_{\beta, \sigma | s} [I_b - I_a + \sum_{j \in J} \int_{\lambda=0}^1 L_{j\lambda} (\Delta X_j \beta - \Delta p_j) d\lambda] / \mu(I_a, \beta, \sigma).$$

This is a line integral over the rectifiable path of the area behind the demand functions for the products in \mathbf{A} between the old and new vectors of quality-adjusted net values, which is the Marshallian consumer surplus associated with the change from policies a to b ; (49) depends on σ only through its influence on the acuity of consumer response to price changes. These price effects are usually bounded even when there is a positive probability of very small σ . Recall that the own price elasticity of a MNL probability $L_{j\lambda} = \frac{\exp(\frac{v_{j\lambda}}{\sigma})}{\sum_{i \in J} \exp(\frac{v_{i\lambda}}{\sigma})}$ equals $\frac{-p_{j\lambda}(1-L_{j\lambda})}{\sigma}$. Use the inequality $e^{-c/\sigma} \leq \sigma/c$ for $c > 0$. If $c_k = \max_{j=0, \dots, J} (v_{j\lambda} - v_{k\lambda}) > 0$, then $L_{k\lambda}$ is bounded above by σ/c_k , and if $c_{-k} \equiv \min_{i \neq k} c_i > 0$, then $1 - L_{k\lambda}$ is bounded by $J\sigma/c_{-k}$. Then the price elasticity is bounded in magnitude by $\max\{p_{k\lambda}/c_k, Jp_{k\lambda}/c_{-k}\}$ no matter how small σ . The limited sensitivity of (49) to σ is also seen by considering limiting cases. For constant $\sigma \rightarrow 0$, $\text{UMCE}(s) \rightarrow \mathbf{E}_{\beta, \sigma | s} \{ \min_{j \in J} v_{j\lambda} - \min_{j \in J} v_{j\lambda} \} / \mu(I_a, \beta, \sigma)$, and for $\sigma \rightarrow +\infty$, $\text{UMCE}(s) \rightarrow \mathbf{E}_{\beta, \sigma | s} \frac{1}{|J|} \sum_{j \in J} \{v_{j\lambda} - v_{j\lambda}\} / \mu(I_a, \beta, \sigma)$. The difference in these expressions comes only from the difference between least and average quality-adjusted net values, reflecting two extremes in the acuity of consumers in gravitating to alternatives with the greatest quality-adjusted net values.

Next consider *retrospective* welfare analysis that quantifies the harm to consumers from a past “as is” scenario a compared to a “but-for” scenario b in which product attributes are changed by altering attributes or seller conduct judged defective or improper. By its nature, retrospective analysis deals with loss of experienced utility, and with compensating transfers to make consumers whole after their “as is” choices have been made, so that experienced-utility MCE rather than HCV or HEV is the target of the analysis, even if choices are influenced by neoclassical income effects. The analysis in these applications is focused on objective changes in product attributes rather than shifts in consumer tastes, so it is reasonable to assume that the decision-utility and experienced-utility tastes are the same, that in most cases $I_a = I_b$, and that $\varepsilon_{jm} = \varepsilon_{jm}^d = \varepsilon_{jm}^e$, so that any gap between decision utility and experienced utility comes from differences in z_{jm}^d and z_{jm}^e . Then experienced utility MCE is given by (40), with MCE^d given by (25). The circumstances of the application will determine the configurations of $v_{jm}^d(I, \beta, \sigma)$ and $v_{jm}^e(I, \beta, \sigma)$ that prevail. A critical question is whether consumers are fully and accurately informed about the attributes of products in both the “as is” and “but for” scenarios, or whether the issue is misinformation or deception on product attributes in the “as is” scenario.

The first case we consider is one in which consumers have full information on the available products under both “as is” and “but for” conditions. One example is anti-trust litigation in which the question is the harm to

consumers caused by improper supplier conduct such as price collusion, market allocation, bundling, or artificial barriers to entry. Other examples are environmental litigation in which the question is the harm caused by improper disposal of hazardous wastes, and patent litigation in which the question is the value to consumers of infringing features. With full information, anticipations are realized, so that $v_{jm}^d(I, \beta) = v_{jm}^e(I, \beta)$ for $m = a, b$. In many applications, the class of consumers of interest is not the general population, but individuals meeting specific conditions, such as residence in a specified region. If the class is defined by consumer characteristics s in a set \mathbf{T} , and either the “as is” choice is unobserved, or it is observed but ϵ_a and ϵ_b are statistically independent, then the per capita transfer prescribed for this class is

$$(50) \quad \text{UMCE}(\mathbf{T}) = \mathbf{E}_{s|\mathbf{T}} \mathbf{E}_{\beta, \sigma | s} \sigma \cdot \ln \frac{L_{ca}(I_a, \beta, \sigma)}{L_{cb}(I_b, \beta, \sigma)} / \mu(I_a, \beta, \sigma).$$

Next consider cases where consumers are misinformed about products in the “as is” scenario 1, due to failure to deliver goods as promised, or to deceptive advertising, resulting in experienced utility that deviates from anticipated utility, an application studied by Chorus and Timmermans (2009), Alcott (2013), Schmeiser (2014), and Train (2015). In these cases, consumers are fully informed in scenario b . Then, $J = J_a = J_b$, (z_{ja}^d, p_{ja}^d) and (z_{ja}^e, p_{ja}^e) are different for j in a set of products \mathbf{D} where the misinformation occurs in scenario a , but (z_{ja}^d, p_{ja}^d) and (z_{ja}^e, p_{ja}^e) agree for $j \notin \mathbf{D}$ and (z_{jb}^d, p_{jb}^d) and (z_{jb}^e, p_{jb}^e) agree for all j . Continue to assume that anticipated and experienced taste parameters are the same. There are two leading possibilities for defining the “but for” scenario: the benchmark “but for” net values can match either the anticipated decision-utility net values when the anticipation is accurate, or match the realized utility net values when these net values are correctly anticipated in the “as is” scenario. The former benchmark applies to contract violations, where the violator is obligated to provide the promised product or equivalent compensation. The latter benchmark arguably applies to false advertising cases where the “but for” scenario correctly informs consumers of the actual product attributes, so that anticipations are realistic.

In the contract violation case, the “but for” net values are defined to match what consumers *anticipated* in the “as is” situation. Then $v_{jb}^e(I, \beta) = v_{jb}^d(I, \beta) = v_{ja}^d(I, \beta)$ for all j , but $v_{ja}^e(I, \beta) \neq v_{ja}^d(I, \beta)$ for $j \in \mathbf{D}$. The appropriate metric for comparing consumer welfare under the “as is” and the “but for” scenarios for consumers with the observed “as is” choice k is the experienced MCE, which becomes

$$(51) \quad \text{MCE}^e(s, k, \beta, \sigma, \epsilon) = \left\{ \max_{j \in J} [v_{jb}^d(I_b, \beta) + \sigma \epsilon_{jb}] - [v_{ka}^e(I_a, \beta) + \sigma \epsilon_{ka}] \right\} / \mu_k(I_a, \beta)$$

$$= \left\{ \max_{j \in J} [v_{jb}^d(I_b, \beta) + \sigma \epsilon_{kb}] - \max_{j \in J} [v_{ja}^d(I_a, \beta) + \sigma \epsilon_{ja}] + v_{ka}^d(I_a, \beta) - v_{ka}^e(I_a, \beta) \right\} / \mu_k(I_a, \beta).$$

If $\varepsilon_a = \varepsilon_b$, this expression reduces to $MCE^e(s, k, \beta, \sigma, \varepsilon) = \{v_{ka}^d(I_a, \beta) - v_{ka}^e(I_a, \beta)\} / \mu_k(I_a, \beta)$, the scaled difference in the anticipated and realized net value for the chosen alternative. Even without the last assumption,

$$(52) \quad UMCE^e(s, k) = \mathbf{E}_{\beta, \sigma | s, k} (v_{ka}^d(I_a, \beta) - v_{ka}^e(I_a, \beta)) / \mu_k(I_a, \beta).$$

Selection again enters the calculation of UMCE for a class of consumers. For consumers with $s \in \mathbf{T}$ and observed “as is” choices in a set \mathbf{D} ,

$$(53) \quad UMCE^e(s, k) = \frac{\mathbf{E}_{s | \mathbf{T}} \mathbf{E}_{\beta, \sigma | s} \sum_{k \in \mathbf{D}} L_{ka}(I_a, \beta, \sigma) \cdot [v_{ka}^d(I_a, \beta) - v_{ka}^e(I_a, \beta)] / \mu_k(I_a, \beta)}{\mathbf{E}_{s | \mathbf{T}} \mathbf{E}_{\beta, \sigma | s} L_{\mathbf{D}a}(I_a, \beta, \sigma)}.$$

These per capita transfers can be applied separately to disjoint \mathbf{D} sets, or combined into a weighted average of the form (52) to give a uniform transfer for all consumers in \mathbf{C} whose scenario a purchases are from \mathbf{D} . Since only consumers who choose an alternative in subset \mathbf{D} experience any difference between anticipated and realized net values, the numerator of (52) is the expected compensating variation per capita for all consumers with characteristics in \mathbf{T} , while the denominator is the share of the population with characteristics in \mathbf{T} and scenario a choices in \mathbf{D} . In (56), commonly $v_{ka}^d(I_a, \beta, \sigma) \geq v_{ka}^e(I_a, \beta, \sigma)$ for all tastes. However, it is possible that there are tastes appearing in reality, or in the utility model approximation to it, that lead to some “as is” winners with $v_{ka}^d(I_a, \beta, \sigma) < v_{ka}^e(I_a, \beta, \sigma)$. This raises two issues, first whether the transfers should be calculated including or excluding winners in the calculation of the aggregate needed to make losers whole. The argument hinges on whether the distribution fulfilling the aggregate transfer can in principle claw back gains from winners to compensate losers; if not, the calculation should exclude winners. A related issue is that it may be impossible to distinguish winners and losers in the class of consumers in \mathbf{C} and \mathbf{D} , in which case the per capita calculation excluding winners but applied to both losers and winners gives an unwarranted transfer to winners.

In the second case, with false advertising or other misinformation about alternatives’ actual attributes, the MCE is the difference between the realized utility obtained from (i) the alternative the person chose when misinformed and (ii) the alternative the person would have chosen if fully informed. If the chosen alternative is the same in the “but for” and “as is” scenarios, then $MCE^e(s, k, \beta, \sigma, \varepsilon) = 0$; i.e., there is no loss for consumers whose choice was unaffected by the misinformation. Since the “but for” anticipated net values are defined to match the net values that consumers *realized* in the “as is” situation, one has $v_{kb}^e(I_a, \beta, \sigma) = v_{kb}^d(I_a, \beta, \sigma) = v_{ka}^e(I_a, \beta, \sigma)$ for all k . Given $\varepsilon_a = \varepsilon_b$, the experienced-utility MCE has the form (51) specialized to this relation among the net values:

$$(54) \quad \text{UMCE}^e(s, k) \\ = \mathbf{E}_{\beta, \sigma | s, k} \left\{ \max \left(\frac{\sigma}{L_{ja}(I_a, \beta, \sigma)} \cdot \ln \frac{\sum_{k=0}^{J_b} \exp(v_{ka}^e(I_a, \beta, \sigma)/\sigma)}{\sum_{k=0}^{J_a} \exp(v_{ka}^d(I_a, \beta, \sigma)/\sigma)}, 0 \right) + v_{ja}^d(I_a, \beta, \sigma) - v_{ja}^e(I_a, \beta, \sigma) \right\} / \mu(I_a, \beta, \sigma).$$

Again, it is normal in false advertising situations (but not necessarily for all forms of misinformation) that $v_{ka}^e(I_a, \beta, \sigma) \leq v_{ka}^d(I_a, \beta, \sigma)$. Then (53) is less than (51); i.e., the transfer is lower when the “but for” scenario consists of providing the correct information that leads anticipated and realized utilities to agree than when the “but for” scenario consists of providing consumers with their anticipated utilities. When there are tastes such that $v_{ka}^e(I_a) > v_{ka}^d(I_a)$, so that these consumers win from the misrepresentation, there is again a question of whether they should be included or excluded in the calculation of the per capita transfer.

Analogously to (52), in the class of consumers with characteristics in \mathbf{T} who chose alternative J in scenario a ,

$$(55) \quad \mathbf{E}_{\varepsilon | \beta, \sigma, \delta_{Ja}(I_a)=1} \text{MCE}^e(\beta, \sigma, \varepsilon) \\ = \frac{\mathbf{E}_{s | \mathbf{T}} \mathbf{E}_{\zeta | s} P_{Ja}(I_a, \beta, \sigma) \cdot (v_{j1}^a(I_a, \beta) - v_{j1}^r(I_a, \beta)) + \mathbf{E}_{\varepsilon | \beta, \sigma, \delta_{Ja}(I_a)=1} \max \left(\sigma \cdot \ln \frac{\sum_{k=0}^{J_2} \exp(v_{ka}^e(I_a, \beta)/\sigma)}{\sum_{k=0}^{J_1} \exp(v_{ka}^d(I_a, \beta)/\sigma)}, 0 \right) / \mu(I_a, \beta, \sigma)}{\mathbf{E}_{\varepsilon | \beta, \sigma, \delta_{Ja}(I_a)=1} P_{Ja}(I_a, \beta, \sigma)}.$$

Retrospective welfare analysis for consumer durables whose attributes are affected by contract violations or deceptions can require a combination of the preceding calculations. For example, consider homeowners whose properties lose value due to groundwater contamination from an industrial site, or automobile owners whose vehicles fail to deliver promised performance after correction of defective emission controls, and lose resale value as a result. Then members of the class of owners of the affected durables at the time the defect is announced are harmed in the amount given by (51) if they are legally entitled to a non-defective product, as in the case of environmental injury, or given by (53) if they are legally entitled only to the opportunity to make a product choice with the correct information, as in the case of false advertising. Further, as long as there is no further contract violation or deception following the announcement, the harm is fully capitalized in the resale value of the durables and these calculations conclude the calculation of harm. Pre-announcement owners who choose to continue to hold their durables have willingly declined the opportunity to mitigate their losses by selling, and post-announcement buyers who find that the lower price offsets the reduced performance are not harmed.

8. AN ILLUSTRATIVE APPLICATION

An empirical example of applied welfare analysis using the methods of this paper, due to Kenneth Train (2015), examines the impact on consumers of video streaming services that share customers’ personal and usage information without their prior knowledge. This analysis is based on choice models estimated using data from a conjoint experiment designed and described by Butler and Glasgow (2015). Each choice experiment included four alternative video steaming services with specified price and the attributes listed in Table 4 plus a fifth alternative of not subscribing to any video streaming service.

Each of 260 respondents was presented with 11 choice experiments. The choice model was of the form (9) for money-metric utility, with $(\beta, \ln \sigma)$ having a multivariate normal distribution. Estimates obtained using maximum simulated likelihood are given in Table 5. The results indicate that people are willing to pay \$1.56 per month on average to avoid commercials. Fast availability is valued highly, with an average WTP of \$3.95 per month in order to see TV shows and movies soon after their original showing. On average, people prefer having a mix with more TV shows and fewer movies, but the mean is not significantly different from zero. Average willingness to pay for more content of both kinds is \$2.96 per month. Interestingly, people who want fast availability tend to be those who prefer more TV shows and fewer movies: the correlation between these two WTP’s is 0.51, while the correlation between WTP for fast availability and more content of both kinds is only 0.04. Apparently, the desire for fast availability mainly applies to TV shows.²¹

Table 4. Non-Price Attributes

Attribute	Levels
Commercials shown between content	Yes (“commercials’) No (baseline category)
Speed of content availability	TV episodes next day, movies in 3 months (“fast content”) TV episodes in 3 months, movies in 6 months (baseline category)
Catalogue	10,000 movies and 5,000 TV episodes (“more content”) 2,000 movies and 13,000 TV episodes (“more TV/fewer movies”) 5,000 movies and 2,500 TV episodes (the baseline category)
Data-sharing policies	Information is collected but not shared (baseline category) Usage information is share with third parties (“share usage”) ²² Usage and personal information are shared with third parties (“share usage and personal”)

²¹ The model was also estimated using an Allenby-Train hierarchical Bayes method, with similar results; the details of both estimation methods are given in Bhat (2001); Train (2000, 2009, 2015), and Ben Akiva, McFadden, and Train (2016).

²² Butler and Glasgow use the terms “non-personally identifiable information (NPPI)” and “personally identifiable information (PII)” for what we are labelling “share usage” and “share usage and personal”.

Table 5A. MSL Estimates of WTPs for Video Streaming Services

	Population Mean		Std Dev in Population	
	Estimate	Std Error	Estimate	Std Error
Ln(1/σ)	-2.002	0.0.945	1.0637	0.0755
WTP for:				
Commercials	-1.562	0.4214	3.940	0.5302
Fast Availability	3.945	0.4767	3.631	0.4138
More TV, fewer movies	-0.6988	0.4783	4.857	0.5541
More content	2.963	0.4708	2.524	0.4434
Share usage only	-0.6224	0.4040	2.494	0.4164
Share personal and usage	-2.705	0.5844	6.751	0.7166
No service	-27.26	2.662	19.42	2.333

Table 5B. Correlation Point Estimates

(* denotes significance at 5% level)

	Commer- cials	Fast Avail- ability	Mostly TV	Mostly movies	Share usage	Share personal and usage	No service
Ln(1/σ)	-0.5813*	-0.1371	0.0358	0.0256	0.0022	-0.1287	0.2801*
Commercials	1.0000	0.1172	-0.3473*	0.0109	-0.2562	-0.0079	-0.4108*
Fast Availability		1.0000	0.8042*	-0.4019*	-0.3542*	-0.4206*	0.2391*
Mostly TV			1.0000	-0.5890*	-0.1695	-0.3328*	0.4616*
Mostly movies				1.0000	0.5141*	0.5181*	-0.0147
Share usage					1.0000	0.9370*	-0.0563
Share personal and usage						1.0000	-0.0975
No service							1.0000

Consider how a video streaming service might share its subscribers’ personal and usage information with third parties who then use that information for targeted marketing to the subscribers. The Table 5 estimates imply that consumers have an average WTP of 62 cents per month to avoid having their usage data shared in aggregate form; however, the hypothesis of zero average WTP cannot be rejected. Consumers are much more concerned about their personal information being shared along with their usage information: The average WTP to avoid such sharing is \$2.71 per month. The correlation between WTP to avoid the two forms of sharing is a substantial 0.937. However, some people like having their data shared, because they value the targeted marketing that they receive as a result of the sharing. In the demand model, the WTP is normally distributed with a mean of -2.71 and standard deviation of 6.751, which implies that 34.4% of the population like to have their information shared.

For the welfare analysis, there are three providers, Netflix, Amazon Prime, and Hulu, and that customers can subscribe to any one of these services, any combination of them, or to no service. Table 6 gives the “as is” alternatives available to customers, and the shares of customers in the sample who chose each alternative. At

the time of the survey, Hulu had about 6 million subscribers, which, given the market shares above, imply that total market size is 31 million potential subscribers. This is less than the number of households in the US because the survey screened for people who either already subscribe, or were likely to subscribe, to a video-screening service if they did not currently have one. The market is then the US households who are open to the possibility of subscribing to a video streaming service.

Table 6: Market Shares of Video Steaming Service Portfolios

Alternative	Share
Netflix	0.2867
Amazon Prime	0.0467
Hulu	0.0400
Netflix + Amazon Prime	0.1167
Netflix + Hulu	0.0700
Amazon Prime + Hulu	0.0100
Netflix + Amazon Prime + Hulu	0.0733
No video streaming service	0.3567

In the “as is” scenario, customers think that none of the service providers shares their usage and personal information, but in fact one of them does. The analysis chooses Hulu as the one who shares, but the selection is arbitrary. How much are consumers hurt by the fact that Hulu shared its subscribers information without their knowing beforehand, and how much would Hulu be liable for under different theories of damages?

Assume for the welfare analysis that when people were choosing among services, they anticipated that these services would have the attributes given in Table 7. Note that none of the providers were thought to share their subscribers’ information.

Table 7: Anticipated Attributes for Decision Utility

	Netflix	Amazon Prime	Hulu
Price per month	7.99	6.58	7.99
Commercials	0	0	0
Fast Availability	0	0	1
More TV, fewer movies	0	1	0
More content	1	0	0
Share usage only	0	0	0
Share personal and usage	0	0	0

The attributes of the alternatives that represent multiple services are the sum of the attributes of the services within the packages. For example, the price of Netflix+Amazon Prime is \$14.67 per month and provides the “More

content” of Netflix and the “MoreTV, fewer movies” of Amazon Prime. Alternative specific constants were calibrated such that the predicted shares for the alternatives equal the observed shares in Table 7.

Now suppose that, in reality, Hulu shared its subscribers’ personal and usage information, and that this fact was revealed months after people began subscribing. The experienced utility is based on the attributes in Table 7 except that “Share personal and usage” receives a 1 for Hulu. What is the difference between the welfare that people expected to obtain when they made their choices compared to the welfare they actually obtained? Only Hulu subscribers obtained experienced utility that differed from decision utility. The aggregate difference is \$22.9 million per month, or \$3.81 on average for Hulu subscribers. Note that, for the population as a whole, the average WTP to avoid sharing is \$2.71, as stated above. The average WTP conditional on having subscribed to Hulu is \$3.81. That is, the average Hulu subscriber dislikes sharing their information more than the average person in the population does. How does this arise? Note in Table 5B that the correlation between the WTPs for between “Fast Availability” and Share personal and usage” is -0.42. Hulu is the only service that offered Fast Availability, and so people who valued this attributed tended to choose Hulu. However, the people who place a high value on Fast Availability also tend to dislike sharing their information more than other people. The difference between the conditional mean of \$3.80 and the unconditional mean of \$2.71 arises because of this correlation.

The damages that Hulu would need to pay in compensation for its sharing of its subscribers’ information depends critically on what was illegal: was it illegal for Hulu to share its customers’ information, or was it illegal for Hulu not to disclose that it was doing so. If it was illegal for Hulu to share its subscribers’ information, then the aggregate damage that Hulu is responsible for is \$22.9 million for each month that the sharing had been undisclosed. However, some customers like having their data shared, and this aggregate nets their gains from the losses that people who dislike sharing incurred. To obtain Pareto neutral compensation on a person-by-person basis, the \$22.80 would not be enough to compensate the people who were hurt by the sharing: the people who liked the sharing would need to contribute their gains too. We can calculate the welfare impact separately for people who like sharing and people who dislike sharing. Among the Hulu subscribers who have a negative WTP for sharing, the aggregate loss in welfare is \$30.4 million. Hulu subscribers who have a positive WTP for sharing obtained an aggregate gain of \$7.50 million. For Hulu to be able to compensate the people who were hurt from its sharing, Hulu would need to pay \$30.4, since it does not have the ability to claw back compensation from the people who gained.

Next suppose information sharing is legal, but nondisclosure is illegal. If Hulu is liable for nondisclosure, then the relevant comparison is between

- (i) the utility that consumers obtained in the “as is” situation, where they choose among the alternative under the concept that Hulu did not share but it in fact did; this is the realized utility for the alternative that the person chose based on decision utilities, and
- (ii) the utility that consumers would have obtained Hulu had disclosed its sharing practice before customers choose among the services; this is the realized utility that the customer would choose based on realized utilities.

Every Hulu subscriber who likes sharing would have chosen Hulu if they had known in advance that it shared information. And some of the Hulu subscribers who dislike sharing would still have chosen Hulu if they had known that Hulu shared their information. None of these subscribers were hurt by the nondisclosure. The only Hulu subscribers who were hurt by the nondisclosure are those who dislike sharing sufficiently that they would not have chosen Hulu if they had known the sharing practice. However, the welfare losses from non-disclosure are not borne only by Hulu subscribers. People who like sharing but didn’t know that Hulu shares and chose a different provider were potentially hurt because they were not able to take advantage of this undisclosed attribute of Hulu service. People who would have chosen Hulu if they had known that Hulu shares but didn’t obtained less welfare than they would have obtained under full disclosure. Table 8 gives the losses for each group of consumers from the non-disclosure of Hulu’s sharing practice.

Table 8: Damages Arising from Non-Disclosure

	Aggregate loss, in million \$ per month	Average loss per person in the market	Average loss for Hulu subscribers	Average loss for people who did not subscribe to Hulu
All people	16.5	0.53	2.16	0.14
People who dislike sharing	13.0	0.64	3.05	0.00
People who like sharing	3.5	0.33	0.00	0.39

The total loss is \$16.5 million per month, which consists of \$13.0 million loss to people who dislike sharing and 3.55 loss to people who like sharing. The \$13.0 million loss was incurred by Hulu subscribers who dislike sharing sufficiently to not choose Hulu if they had known its sharing practices. The \$3.5 million loss was incurred by people who did not subscribe to Hulu but like sharing sufficiently to have chosen Hulu if they had known its sharing practices. The average loss per person in the population is simply the aggregate loss divided by market size (31 million). The average loss for Hulu subscribers can best be explained by starting in the bottom row of Table 10. Hulu subscribers who like sharing their information incurred zero harm from the nondisclosure: they subscribed to Hulu and so obtained the benefits of the sharing even though they didn't realize beforehand that they would.

Importantly, they also did not gain from the nondisclosure. They obtained greater welfare from Hulu than they had expected when they chose Hulu. But they obtained the benefits of sharing even without prior disclosure, which would not have changed anything for them. Hulu subscribers who dislike sharing were hurt by \$3.05 on average. Not all Hulu subscribers who dislike sharing were hurt by the non-disclosure. Only those who would not have chosen Hulu if they had known of its sharing practices were hurt, and these people were hurt by more than \$3.05 on average (since the \$3.05 average include Hulu subscribers who dislike sharing but were not hurt from the nondisclosure since they still would have chosen Hulu.) The top row in Table 10 gives a loss per Hulu subscriber of \$2.16: it is the average of the \$3.05 in the second row and \$0.00 in the third row, weighted by the share of Hulu subscribers who dislike and like sharing. The losses for people who did not subscribe the Hulu are analogous. People who dislike sharing and did not subscribe to Hulu incurred no loss, since they would not have chosen Hulu if its sharing practices had been disclosed. Some people who did not subscribe to Hulu but like sharing would have chosen Hulu if they had known that Hulu shared their information. These people obtained less utility that they could have obtained under full disclosure.

In the “as is” situation, 19.3 percent of people in the market subscribed to Hulu. If everyone had been informed about Hulu’s sharing practice, then this share would have dropped to 16.0 percent, which is a 17 percent reduction in subscribers. However, as explained above, this change includes two different movements: the share drops because some Hulu subscribers would not have chosen Hulu if they had known that Hulu would share their information, and the share rises because some people who did not subscribe to Hulu would have subscribed if they had known. Table 9 gives the share of people in each group. 12.5% of people subscribed to Hulu and would still have also done so if the sharing practice had been disclosed. 6.8% subscribed to Hulu but would not have if they had known about its sharing practice. That is, about a third of Hulu’s subscribers would have not subscribed if they had been informed. 3.5% of people did not subscribe to Hulu but would have done so if they had known that Hulu shares their information.

Table 9: Choice Shares without and with Disclosure

	Would have subscribed to Hulu if its sharing practices had been disclosed	Would not have subscribed to Hulu if its sharing practices had been disclosed	Total
Subscribed to Hulu	0.125	0.068	0.193
Did not subscribe to Hulu	0.035	0.772	0.807
Total	0.160	0.840	

The share of people who subscribed to Hulu was 19.3%. If its sharing practices had been disclosed, then the share of subscribers would have been $0.193 - 0.068 + 0.035 = 0.16$, i.e., 16 % as stated above.

9. CONCLUSIONS

This paper provides a foundation for applied welfare analysis of product regulation or compensation for product defects. It gives a practical setup for money-metric indirect utility functions whose features can be estimated using data on choice in real or hypothetical markets, and shows that there is essentially no loss of generality in restricting analysis to this setup. It draws a distinction between prospective and retrospective policy applications, and between cases where compensating transfers are hypothetical or are actually fulfilled. It introduces a Market Compensating Equivalent (MCE) welfare measure, an updated version of Marshallian consumer surplus, and shows that when compensating transfers are not actually fulfilled, it is preferred to commonly prescribed Hicksian compensating or equivalent variations. Further, MCE is shown to have desirable computational and aggregation properties. The problem of carrying out welfare calculations when tastes of individual consumers are only partially observed is addressed, and computational formulas are given for calculation of expected compensating transfers. Decision-utility and experienced-utility are distinguished, and the issues of conducting welfare calculus in experienced utility are discussed. A number of common welfare calculus problems are treated, and formulas are given for their resolution. Finally, an application illustrates the use of these methods and the importance of the distinctions introduced in this paper.

REFERENCES

- Afriat, S. (1967) "The construction of utility functions from expenditure data," *International Economic Review*, 8, 67-77.
- Alexandrov, A. (1939) "Almost everywhere existence of the second differential of a convex function and surfaces connected with it," *Leningrad State University Annals, Mathematics Series* 6;3-35.
- Aliprantis, C.; K. Border (2006) *Infinite dimensional Analysis*, Springer: Berlin.
- Allcott, H., (2013) "The welfare effects of misperceived product costs: data and calibrations from the automobile market." *Am. Econ. J.: Econ. Policy* 5 (3), 30–66.
- Anas, A. and C. Feng (1988) "Invariance of Expected Utilities in Logit Models," *Economic Letters* 27:1, 41-45.
- Ben-Akiva, M.; D. McFadden; K. Train (2016) "Foundations of Stated Preference Elicitation: Consumer Behavior and Choice-Based Conjoint Analysis," University of California, Berkeley, working paper.
- Bentham, J. (1789) *An introduction to the principles of morals and legislation*, Oxford: The Clarendon Press, 1876.
- Border, K. (2014) "Monetary Welfare Measurement," Cal Tech lecture notes.
- Chipman, J.; J. Moore (1980) "Compensating Variation, Consumer's Surplus, and Welfare," *American Economic Review*. 70: 933-49
- Chipman, J.; J. Moore (1990) "Acceptable Indicators of Welfare Change, Consumer's Surplus Analysis, and the Gorman Polar Form," in D. McFadden, M. Richter (eds) *Preferences, uncertainty, and optimality: Essays in honor of Leonid Hurwicz*. Boulder and Oxford: Westview Press; 68-120.
- Chorus, C.G., H. Timmermans (2009) "Measuring user benefits of changes in the transport system when traveler awareness is limited," *Transportation Research Part A*, 43(5), 536-547.
- Conniffe, D. (2007) "A Note on Generating Globally Regular Indirect Utility Functions," *Journal of Theoretical Economics*, 7.1, 1-11.

- Dagsvik, J.; A. Karlstrom (2005) "Compensating Variation and Hicksian Choice Probabilities in Random Utility Models that are Nonlinear in Income," *Review of Economic Studies*, 72.1, 57-76.
- Debreu, G. (1959) *Theory of Value*, New Haven : Yale University Press.
- Diamond, P. and D. McFadden (1974), "Some uses of the expenditure function in public finance," *Journal of Public Economics* 3.1 3-21.
- Doha, E. H.; A. H. Bhrawy; M. A. Saker (2011) "On the Derivatives of Bernstein Polynomials," *Boundary Value Problems*, doi:10.1155/2011/829543 p. 1-16.
- Dubin, J. (1985) *Consumer Durable Choice and the Demand for Electricity*, Elsevier: New York.
- Dubin, J., D. McFadden (1984) "An Econometric Analysis of Residential Electric Appliance Holdings and Consumption," *Econometrica*, 52, 345-62
- Dudley, R. (2002) *Real Analysis and Probability*, Cambridge University Press, New York.
- Dunford, N.; J. Schwartz (1964) *Linear Operators*, Interscience, New York.
- Dupuit, J. (1844) "On the Measurement of the Utility of Public Works", *Annales des ponts et chaussées*. (English translation, *International Economic Review*, 1952).
- Edgeworth, Francis Y. (1881) *Mathematical Psychics; an essay on the application of mathematics to the moral sciences*, London, C. K. Paul & co.
- Fosgerau, M.; D. McFadden (2012) "A theory of the perturbed consumer with general budgets," working paper.
- Gorman, W. (1953) "Community Preference Fields," *Econometrica*, 21, 63-80.
- Gorman, W. (1961) "On a Class of Preference Fields," *Metroeconomica*, 13, 53-56.
- Gossen, Hermann (1854) *Die Entwicklung*, English translation: *The Laws of Human Relations*, Cambridge: MIT Press, 1983.
- Hammond, Peter (1994) "Money Metric Measures of Individual and Social Welfare Allowing for Environmental Externalities," in W. Eichhorn (ed) *Models and Measurement of Welfare and Inequality*, Springer-Verlag, 694-724.
- Heiss, F.; D. McFadden; J. Winter (2013) "Plan Selection in Medicare Part D: Evidence from Administrative Data," *Journal of Health Economics* 32.6, 1325-1344.
- Hicks, John (1939) *Value and Capital*, Oxford, Clarendon press.
- Houthakker, H. (1950) "Revealed preference and the utility function," *Economica*, N.S. 17, 159-174.
- Hurwicz, L.; H. Uzawa (1971) "On the Integrability of Demand Functions," in J. Chipman, L. Hurwicz, M. Richter, and H. Sonnenschein (eds) *Preferences, Utility, and Demand*, New York: Harcourt, 114-148.
- Jevons, W. (1871) *Theory of Political Economy*, reprinted by London, Macmillan, 1931.
- Johnson, N. and S. Kotz (1970, Ch. 21) *Continuous Univariate Distributions-1*, Houghton-Mifflin: New York.
- Kadison, R. and Z. Liu (2016) *Bernstein Polynomial and Approximation*, lecture notes.
- Kaldor, N. (1939) "Welfare Propositions of Economics and Interpersonal Comparisons of Utility," *Econ. Jour.*, XLIX, 549-52.
- Katzner, Donald (1970) *Static Demand Theory*, New York: Macmillan.
- Kosorok, M. (2008) *Introduction to Empirical Processes and Semiparametric Inference*, Springer: New York.
- Lorentz, G. (1937) "Zur theorie der polynome von S. Bernstein," *Matematicheskij Sbornik* 2, 543-556.
- Lowenstein, G.; Ubel (2008) "Hedonic adaptation and the role of decision and experience utility in public policy," *Journal of Public Economics*, 92, 1795-1810.
- Marshall, A. (1890) *Principles of Economics*, London: Macmillan.
- Mas-Colell, Andreu; M. Whinston, and J. Green (1995) *Microeconomic Theory*, Oxford: Oxford University Press.
- Matzkin, R. and D. McFadden (2011) "Trembling Payoff Market Games," working paper.
- McFadden, D. (1974) "The Measurement of Urban Travel Demand," *Journal of Public Economics*, 3, 303-328.
- McFadden, D. (1981) "Structural Discrete Probability Models Derived from Theories of Choice," in C. Manski and D. McFadden (eds) *Structural Analysis of Discrete Data and Econometric Applications*, MIT Press: Cambridge, 198-272.
- McFadden, D. (1986) "The Choice Theory Approach to Market Research," *Marketing Science*, 275-297.
- McFadden, D. (1994) "Contingent valuation and social choice," *American Journal of Agricultural Economics* 76, 689-708.
- McFadden, D. (1999) "Computing Willingness-to-Pay in Random Utility Models," in J. Moore, R. Riezman, and J. Melvin (eds.), *Trade, Theory, and Econometrics: Essays in Honour of John S. Chipman*, Routledge: London.
- McFadden, D. (2004) "Welfare Economics at the Extensive Margin: Giving Gorman Polar Consumers Some Latitude," University of California, Berkeley, working paper.
- McFadden, D. (2008) "Environmental Valuation of Environmental Projects," Univ. of California working paper.
- McFadden, D. (2012) "Economic Juries and Public Project Provision," *Journal of Econometrics*, 166, 116-126.

- McFadden, D. (2014) "The New Science of Pleasure: Consumer Behavior and the Measurement of Well-Being," in S. Hess and A. Daly, eds, *Handbook of Choice Modelling*, Elgar: Cheltenham, 7-48.
- McFadden, D. (2017) "Stated Preference Methods and their Applicability to Environmental Use and Non-Use Valuations," in D. McFadden and K. Train (eds) *Contingent Valuation of Environmental Goods: A Comprehensive Critique*, Elgar: Cheltenham, Chap. 6.
- McFadden, D., K. Train (2000) "Mixed MNL Models for Discrete Response," *Journal of Applied Econometrics*, 15, 447-470.
- McFadden, D.; B. Zhou (2015) "Measuring Lost Welfare from Poor Health Insurance Choices," Schaeffer Center, USC.
- Miller, K.; et al. (2011) "How Should Consumers' Willingness to Pay Be Measured? An Empirical Comparison of State-of-the-Art Approaches," *Journal of Marketing Research*, 48.1, 172-184.
- Pareto, Vilfredo (1906) *Manual of Political Economy*, English Translation, Augustus M. Kelley, NY, 1971.
- Peleg, B. (1970) "Utility functions for partially ordered topological spaces," *Econometrica*, 38, 93-96.
- Pollard, D. (1984) *Convergence of Stochastic Processes*, Springer, New York.
- Rademacher, H. (1919) "Über partielle und totale Differenzierbarkeit von Funktionen mehrerer Variablen und über die Transformation der Doppelintegrale," *Math. Ann.* 79, 340-359.
- Rader, T. (1973) "Nice demand functions," *Econometrica*, 41, 913-935.
- Resnic, S.; R. Roy (1990) "Leader and Maximum Independence for a Class of Discrete Choice Models," *Economic Letters*, 33.3, 259-263.
- Richter, M. (1966) "Revealed Preference Theory," *Econometrica*, 34, 635-645.
- Roy, Rene (1942) *Elements d'econometrie*. Presses universitaires de France, Paris.
- Samuelson, P. (1947) *Foundations of economic analysis*, Cambridge: Harvard University Press, 1983.
- Samuelson, P. (1948) "Consumption theory in terms of revealed preference," *Economica*, 15, 243-253.
- Schmeiser, S. (2014) "Consumer inference and the regulation of consumer information," *Int. J. Ind. Organ.* 37, 192-200.
- Scitovsky, T. (1951) "The State of Welfare Economics," *American Economic Review*, 41-3, 303-315.
- Shannon, C. (2006) "A Prevalent Transversality Theorem for Lipschitz Functions," *Proceedings of the American Mathematical Society*, 134.9, 2755-2755.
- Slutsky, E. (1915) "Sulla teoria del bilancio del consummatore", *Giornale degli Economisti*. English translation, "On the Theory of the Budget of the Consumer," in G. Stigler and K. Boulding, eds, *Readings in Price Theory*, Homewood: Irving.
- Small, K.; S. Rosen (1981), "Applied Welfare Economics with Discrete Choice Models," *Econometrica*, 49.1, 105-130.
- Smith, A. (1776) *An inquiry into the nature and causes of the wealth of nations*. London, W. Strahan and T. Cadell.
- Thaler, R.; C. Sunstein (2003) "Libertarian Paternalism," *American Economic Review*, 93.2, 175-179.
- Thaler, R.; C. Sunstein (2008) *Nudge: Improving Decisions about Health, Wealth, and Happiness*, Yale University Press: New Haven.
- Thurstone, L. (1927) "A Law of Comparative Judgment," *Psychological Review*, 34: 273-286.
- Train, K. (2015) "Welfare calculations in discrete choice models when anticipated and experienced attributes differ: A guide with examples," *Journal of Choice Modelling*, 16, 15-22.
- Van der Vaart, A.; J. Wellner (1996) *Weak Convergence and Empirical Processes*, Springer, New York.
- Varian, H. (1982) "The Nonparametric Approach to Demand Analysis," *Econometrica*, 50: 945-73.
- Varian, H. (2006) *Revealed Preference*. New York: Oxford University Press.
- Willig, Robert (1976) "Consumer's Surplus without Apology," *American-Economic-Review*. 66: 589-97.
- Yatchew, A. (1985) "Applied Welfare Analysis with Discrete Choice Models: Comment," *Economic Letters*, 18.1, 13-16.
- Zhao, Y.; K. Kockelman; A. Karlstrom (2012) "Welfare Calculations in Discrete Choice Settings: The Role of Error Term Correlation," *Transport Policy* 19.1, 76-84.

Appendix A: Approximation Theory for Functions and Probabilities

This appendix provides the mathematical basis for uniform parametric approximations to utility functions and probabilities. The first theorem adapts Bernstein-Weierstrauss approximation theory to the class of functions considered in this paper, and the second theorem utilizes Pollard's methods for establishing uniform weak convergence of empirical probabilities; see Lorentz (1937), Kadison-Liu (2016), Pollard (1984).

Let $b_{j,K}(p) = \binom{K}{j} p^j (1-p)^{K-j}$ denote the binomial probability of j successes in K draws, each with probability $p \in [0,1]$; and define $b_{j,K}(p) \equiv 0$ for $j < 0$ or $j > K$. Differentiating, $\frac{d}{dp} b_{j,K}(p) = K[b_{j-1,K-1}(p) - b_{j,K-1}(p)]$. Higher order derivatives can be defined recursively; see Doha et al (2011). Note that $\sum_{j=0}^K b_{j,K}(p) \equiv 1$ and $\sum_{j=0}^K \frac{d}{dp} b_{j,K}(p) \equiv 0$. The following result is a straightforward multivariate restatement of the Bernstein-Weierstrauss theorem on approximation of continuous functions by polynomials.

Theorem A.1. Let H denote a compact metric space with metric h . Consider $f \in \mathcal{C}([0,1]^n \times H)$. Let $\mathbf{K} = (K_1, \dots, K_n)$ denote a vector of positive integers, $\mathbf{j} = (j_1, \dots, j_n)$ a vector of integers satisfying $0 \leq j_i \leq K_i$ for $i = 1, \dots, n$, $\mathbf{p} = (p_1, \dots, p_n) \in [0,1]^n$, $\mathbf{j} \oslash \mathbf{K} = (j_1/K_1, \dots, j_n/K_n)$. Define the multivariate binomial probability $b_{\mathbf{j},\mathbf{K}}(\mathbf{p}) = \prod_{i=1}^n b_{j_i, K_i}(p_i)$, the vector $\beta_{\mathbf{K}}(h)$ of functions $\beta_{\mathbf{j},\mathbf{K}}(h) \equiv f(\mathbf{j} \oslash \mathbf{K}, h)$ on H for $\mathbf{0} \leq \mathbf{j} \leq \mathbf{K}$, and the multivariate polynomial $B_{\mathbf{K}}(\mathbf{p}, \beta_{\mathbf{K}}(h)) = \sum_{\mathbf{0} \leq \mathbf{j} \leq \mathbf{K}} b_{\mathbf{j},\mathbf{K}}(\mathbf{p}) \beta_{\mathbf{j},\mathbf{K}}(h)$. Let C denote the compact range of f . Then, $\beta_{\mathbf{j},\mathbf{K}} \in \mathcal{C}(H, C)$, and $B_{\mathbf{K}}(\mathbf{p}, \beta_{\mathbf{K}}(h))$ has the following approximation properties: (i) $\lim_{K \rightarrow \infty} \max_{[0,1]^n \times H} |B_{\mathbf{K}}(\mathbf{p}, \beta_{\mathbf{K}}(h)) - f(\mathbf{p}, h)| = 0$, and (ii) if $\partial f(\mathbf{p}, h)/\partial p_i$ exists and is continuous on closed set $A \subseteq [0,1]^n \times H$, then $\lim_{K \rightarrow \infty} \max_A \left| \frac{\partial B_{\mathbf{K}}(\mathbf{p}, \beta_{\mathbf{K}}(h))}{\partial p_i} - \frac{\partial f(\mathbf{p}, h)}{\partial p_i} \right| = 0$. If in addition, f is Lipschitz in its arguments, then $\beta_{\mathbf{K}}$ is Lipschitz on H .

Proof: The continuous function f is uniformly continuous on $[0,1]^n \times H$ and bounded by a constant M , so that given $\varepsilon > 0$, there exists $\delta \in (0,1)$ such that $\|\mathbf{p}' - \mathbf{p}\| \leq \delta$ and $h(h, h') \leq \delta$ imply $|f(\mathbf{p}', h') - f(\mathbf{p}, h)| < \varepsilon/6$. Define the set $\mathbf{J}_\delta = \{\mathbf{j} | \mathbf{0} \leq \mathbf{j} \leq \mathbf{K} \text{ and } |\mathbf{j} \oslash \mathbf{K} - \mathbf{p}| \leq \delta/2\}$. By Hoeffding's inequality, $\text{Prob}(\mathbf{J}_\delta^c) \leq 2 \sum_{i=1}^n \exp(-\delta^2 K_i/2)$. Select $\mathbf{K} \geq 192nM/\varepsilon\delta^3$ and $\mathbf{K} > 2/\delta$. In the inequality $|B_{\mathbf{K}}(\mathbf{p}, \beta_{\mathbf{K}}(h)) - f(\mathbf{p}, h)| \leq \left\{ \sum_{\mathbf{J}_\delta} + \sum_{\mathbf{J}_\delta^c} \right\} b_{\mathbf{j},\mathbf{K}}(\mathbf{p}) |f(\mathbf{j} \oslash \mathbf{K}, h) - f(\mathbf{p}, h)|$, the first sum is bounded by $\varepsilon/6$, while the second sum is bounded by $2M \cdot \text{Prob}(\mathbf{J}_\delta^c) \leq 4M \cdot \sum_{i=1}^n \exp(-\delta^2 K_i/2) \leq 4M \cdot \sum_{i=1}^n \frac{2}{\delta^2 K_i} \leq \varepsilon\delta/24 \leq \varepsilon/24$. This establishes $|B_{\mathbf{K}}(\mathbf{p}, \beta_{\mathbf{K}}(h)) - f(\mathbf{p}, h)| < \varepsilon/3$ for each $(\mathbf{p}, h) \in [0,1]^n \times H$.

Next suppose that on a compact set A , $\partial f(\mathbf{p}, h)/\partial p_1$ exists and is continuous. Then it is uniformly continuous and bounded on A ; let M be a bound. The δ above can be chosen so that $\left| \frac{\partial f(\mathbf{p}', h')}{\partial p_1} - \frac{\partial f(\mathbf{p}, h)}{\partial p_1} \right| \leq \frac{\varepsilon}{6}$ and $\frac{f(\mathbf{p} + \delta' \Delta_1, h) - f(\mathbf{p}, h)}{\delta'} - \frac{\partial f(\mathbf{p}, h)}{\partial p_1} \equiv \zeta(\delta', \mathbf{p}, h)$ with $|\zeta(\delta', \mathbf{p}, h)| \leq \frac{\varepsilon}{6}$ for $|\delta'| \leq \delta$ and $|\zeta(\delta', \mathbf{p}, h)| \leq M(1+2/\delta)$ for $|\delta'| > \delta$, where Δ_1 is a vector with a one in component 1, zeros elsewhere. Define $\mathbf{p}_{2+} = (p_2, \dots, p_n)$, $\mathbf{j}_{2+} = (j_2, \dots, j_n)$, $\mathbf{K}_{2+} = (K_2, \dots, K_n)$, and $b_{\mathbf{j}_{2+}, \mathbf{K}_{2+}}(\mathbf{p}_{2+}) = \prod_{i=2}^n b_{j_i, K_i}(p_i)$ on $[0,1]^{n-1}$. Then

$$\frac{\partial B_{\mathbf{K}}(\mathbf{p}, \beta_{\mathbf{K}}(h))}{\partial p_1} = K_1 \left\{ \sum_{\mathbf{J}_\delta} + \sum_{\mathbf{J}_\delta^c} \right\} b_{\mathbf{j}_{2+}, \mathbf{K}_{2+}}(\mathbf{p}_{2+}) \left(b_{j_1-1, K_1-1}(p_1) - b_{j_1, K_1-1}(p_1) \right) f(\mathbf{j}_{2+} \oslash \mathbf{K}_{2+}, h)$$

$$\begin{aligned}
&= \left\{ \sum_{\mathbf{J}_\delta} + \sum_{\mathbf{J}_\delta^c} \right\} b_{\mathbf{j}_{2+}, \mathbf{K}_{2+}}(\mathbf{p}_{2+}) b_{\mathbf{j}_1, \mathbf{K}_1-1}(p_1) \left[\frac{f((\mathbf{j}_1+1)/\mathbf{K}_1, \mathbf{j}_{2+} \circ \mathbf{K}_{2+}, h) - f(\mathbf{j}_1/\mathbf{K}_1, \mathbf{j}_{2+} \circ \mathbf{K}_{2+}, h)}{1/\mathbf{K}_1} \right] \\
&= \left\{ \sum_{\mathbf{J}_\delta} + \sum_{\mathbf{J}_\delta^c} \right\} b_{\mathbf{j}_{2+}, \mathbf{K}_{2+}}(\mathbf{p}_{2+}) b_{\mathbf{j}_1, \mathbf{K}_1-1}(p_1) \left[\frac{\partial f(\mathbf{j}_1/\mathbf{K}_1, \mathbf{j}_{2+} \circ \mathbf{K}_{2+}, h)}{\partial p_1} + \zeta(\delta', \mathbf{p}, h) \right] \\
&= \frac{\partial f(\mathbf{p}, h)}{\partial p_1} + \left\{ \sum_{\mathbf{J}_\delta} + \sum_{\mathbf{J}_\delta^c} \right\} b_{\mathbf{j}_{2+}, \mathbf{K}_{2+}}(\mathbf{p}_{2+}) b_{\mathbf{j}_1, \mathbf{K}_1-1}(p_1) \left[\frac{\partial f(\mathbf{j}_1/\mathbf{K}_1, \mathbf{j}_{2+} \circ \mathbf{K}_{2+}, h)}{\partial p_1} - \frac{\partial f(\mathbf{p}, h)}{\partial p_1} + \zeta(\delta', \mathbf{p}, h) \right].
\end{aligned}$$

On \mathbf{J}_δ , the term above in square brackets is bounded by $\frac{\varepsilon}{6}$, which then also bounds the first sum, and on \mathbf{J}_δ^c this term is bounded by $5M/\delta$. The probability of \mathbf{J}_δ^c is bounded by $2\sum_{i=1}^n \exp(-\delta^2 K_i/2)$, so the second sum is bounded by $\frac{10M}{\delta} \cdot \sum_{i=1}^n \exp(-\delta^2 K_i/2) \leq \frac{10M}{\delta} \cdot \sum_{i=1}^n \frac{1}{\delta^2 K_i/2}$. Then $\mathbf{K} \geq 192nM/\varepsilon\delta^3$ implies that the second sum is bounded by $20\varepsilon/192 < \varepsilon/6$. This establishes the approximation property $\left| \frac{\partial B_{\mathbf{K}}(\mathbf{p}, \beta_{\mathbf{K}}(h))}{\partial p_i} - \frac{\partial f(\mathbf{p}, h)}{\partial p_i} \right| < \varepsilon/3$ at each (\mathbf{p}, h) in A .

A final step to establish (i) and (ii) uniformly considers the open cover of neighborhoods where the results hold (with tolerance $\varepsilon/2$ rather than $\varepsilon/3$), extracts finite sub-coverings for the compact domains, and uses the minimum value of δ from these finite sub-coverings. By construction, $\beta_{\mathbf{K}}$ retains the properties of f with respect to h ; hence, in particular, if f is Lipschitz in H , then so is $\beta_{\mathbf{K}}$. ■

The next results will establish uniform convergence of empirical expectations for a family of functions that encompasses the applications in this paper. These results are obtained as specializations of the general theory of stochastic convergence treated in Dudley (2014), Kosorak (2008), Pollard (1984), and van der Vaart and Wellner (1996), referred to hereafter as VW. Let Y denote a closed subset of \mathbb{R}^n , \mathcal{Y} denote the Borel σ -field of subsets of Y , and F denote a probability on \mathcal{Y} . Define a family \mathcal{F} of functions $f: Y \rightarrow \mathbb{R}$ that is contained in the Banach space $\mathcal{L}_1(Y, \mathcal{Y}, F)$ and includes the constant function $f(y) \equiv 1$. We assume that the functions in \mathcal{F} are bounded by an envelope function $f^* \in \mathcal{L}_1(Y, F)$; i.e., $f^* \geq |f|$ for $f \in \mathcal{F}$. Let Θ denote a compact subset of \mathbb{R}^d , with a bound $\alpha > \max(1, \max_{\theta \in \Theta} \|\theta\|)$. Assume that the functions in \mathcal{F} are indexed by $\theta \in \Theta$ and are Lipschitz with respect to this index; specifically, $|f_\theta(y) - f_{\theta'}(y)| \leq \|\theta' - \theta\| \cdot f^*(y) \leq \alpha \cdot f^*(y)$. We will call \mathcal{F} with the properties above a *Lipschitz-parametric family*.

Let F_T denote the empirical probability defined by T independent draws $\{y_1, \dots, y_T\}$ from F ; i.e., for $A \in \mathcal{Y}$, $F_T(A) = \frac{1}{T} \sum_{t=1}^T \mathbf{1}(y_t \in A)$. For $f \in \mathcal{L}_1(Y, F)$ and a probability Q on \mathcal{Y} , define $\mathbf{E}_Q f \equiv \int_Y f(y) Q(dy)$ and $\mathbf{E}_T f \equiv \frac{1}{T} \sum_{t=1}^T f(y_t)$. Define $\|f\|_Q = \mathbf{E}_Q |f|$, and note that $\|f\|_F$ is the norm of $\mathcal{L}_1(Y, F)$. A strong law of large numbers establishes that $\mathbf{E}_T f \xrightarrow[as]{} \mathbf{E}_F f$ pointwise for each $f \in \mathcal{F}$ and for f^* . We give conditions under which this convergence is uniform on \mathcal{F} .

A measure of the “density” or “complexity” of \mathcal{F} is its *bracketing number* $N_{[]}(\gamma, \mathcal{F}, Q)$, defined for $\gamma > 0$ and a probability Q on \mathcal{Y} , the minimum cardinality of a family $\mathcal{F}_\gamma \subseteq \mathcal{L}_1(Y, \mathcal{Y}, F)$, not necessarily a subset of \mathcal{F} , such that for each $f \in \mathcal{F}$, there are $f', f'' \in \mathcal{F}_\gamma$ satisfying $f' \geq f \geq f''$ and $\mathbf{E}_Q(f' - f'') < \gamma$. A related measure of the complexity of \mathcal{F} is its *covering number* $N(\gamma, \mathcal{F}, Q)$, defined for $\gamma > 0$ and a probability Q on \mathcal{Y} as the minimum cardinality of a family $\mathcal{F}_\gamma \subseteq \mathcal{L}_1(Y, \mathcal{Y}, F)$, not necessarily a subset of \mathcal{F} , such that for each $f \in \mathcal{F}$, $\inf_{f' \in \mathcal{F}_\gamma} \mathbf{E}_Q |f' - f| < \gamma$. Obviously, $N(\gamma, \mathcal{F}, Q) \leq N_{[]}(\gamma, \mathcal{F}, Q)$. We will be interested in families of functions for which the bracketing or covering number is finite. The following result specializes VW Theorem 2.7.11:

Lemma A.2. Consider a Lipschitz-parametric family \mathcal{F} and a positive constant $M > 1$. For each probability Q on \mathcal{Y} such that $\mathbf{E}_Q f^* \leq M$ and each $\gamma > 0$, $N_{[]}(\gamma, \mathcal{F}, Q) \leq 2 + 2(8\alpha M/\gamma)^d$.

Proof: Let J be the largest integer no greater than $8\alpha M/\gamma$, and $\mathbf{j} = (j_1, \dots, j_d)$ a vector of indices with $1 \leq j_i \leq J$ for each i . Consider the family of open balls of radius $\gamma/2M$ centered at $\mathbf{b}_j = (-\alpha + j_1\gamma/4M, \dots, -\alpha + j_d\gamma/4M)$. This family covers $\Theta \subseteq [-\alpha, \alpha]^d$ and contains J^d elements. Discard the balls that do not intersect Θ . From each of the remainder, select a point $\theta_j \in \Theta$ and let \mathcal{F}_γ denote the family of functions $\min(f_{\theta_j} + \frac{\gamma f^*}{2M}, f^*)$ and $\max(f_{\theta_j} - \frac{\gamma f^*}{2M}, -f^*)$ plus f^* and $-f^*$. Then, \mathcal{F}_γ contains at most $2(1+J^d)$ functions. For θ in the ball containing θ_j , the Lipschitz condition gives $|f_{\theta_j}(\mathbf{y}) - f_\theta(\mathbf{y})| \leq \frac{\gamma}{2M} f^*(\mathbf{y})$, implying $f_{\theta_j}(\mathbf{y}) + \frac{\gamma}{2M} f^*(\mathbf{y}) - f(\mathbf{y}) \geq 0 \geq f_{\theta_j}(\mathbf{y}) - \frac{\gamma}{2M} f^*(\mathbf{y}) - f(\mathbf{y})$. Then \mathcal{F}_γ brackets \mathcal{F} , and $N_{[]}(\gamma, \mathcal{F}, Q) \leq 2 + 2(8\alpha M/\gamma)^d$. ■

Augment the Lipschitz-parametric family \mathcal{F} with the countable family $\mathcal{F}^0 \equiv \bigcup_{k=1}^{\infty} \mathcal{F}_{1/k}$ of the approximating functions in Lemma A.2 at tolerances $\gamma = 1/k$ for $k = 1, 2, \dots$; i.e., consider the family $\mathcal{F}^* \equiv \mathcal{F} \cup \mathcal{F}^0$. Then the bound on bracketing numbers that the lemma establishes for \mathcal{F} also holds for \mathcal{F}^* , and \mathcal{F}^0 is dense in \mathcal{F}^* . Then, \mathcal{F}^* is said to be *Q-measurable* for any probability Q on \mathcal{Y} such that $\mathbf{E}_Q f^* \leq M$; see VW, 2.2.3 and 2.2.4.

Theorem A.3. Consider a Lipschitz-parametric family $\mathcal{F} \subseteq \mathcal{L}_1(\mathcal{Y}, F)$ that has an envelope $f^* \in \mathcal{L}_1(\mathcal{Y}, F)$. For each tolerance $\gamma \in (0, 1)$, $\lim_{T \rightarrow \infty} \text{Prob}(\sup_{T' \geq T} \sup_{f \in \mathcal{F}} |(\mathbf{E}_{T'} - \mathbf{E})f| > \gamma) = 0$.

Proof: From the discussion following Lemma A.2, consider the augmented family \mathcal{F}^* that contains \mathcal{F} and also contains the countable dense subfamily \mathcal{F}^0 . Given $\gamma \in (0, 1)$, the condition $\mathbf{E}_f f^* < \infty$ implies there exists a constant $M > \mathbf{E}_f f^*$ such that $\mathbf{E}_f f^* \cdot \mathbf{1}(f^* > M) < \gamma/4$. Define $\mathcal{F}^M = \{\min(M, \max(f, -M)) \mid f \in \mathcal{F}^*\}$. From Lemma A.2, the bracketing number bound established on \mathcal{F} by the functions in \mathcal{F}_γ also holds for \mathcal{F}^M and the corresponding finite family $\mathcal{F}_\gamma^M = \{\min(M, \max(f, -M)) \mid f \in \mathcal{F}_\gamma\}$ for all probabilities Q on \mathcal{Y} , since $f_M^* = \min(M, \max(f^*, -M))$ is an envelope function for \mathcal{F}^M and $\mathbf{E}_Q f_M^* \leq M$. Then from Lemma A.2, $N(\gamma, \mathcal{F}^M, F_T) \leq 2 + 2(8\alpha M/\gamma)^d$. This bound is independent of T . Then, the result follows for \mathcal{F}^* , and hence for \mathcal{F} , from VW Theorem 2.4.3. ■

The following result is stated in a form sufficient for our needs; for more general results, see VW, 2.6.17.

Theorem A.4. Consider a finite-dimensional linear subspace \mathcal{K} of $\mathcal{L}_1(\mathcal{Y}, F)$. Without loss of generality, assume that \mathcal{K} includes the function $f(\mathbf{y}) \equiv 1$. For a fixed integer J , define \mathcal{F} to be a subset of the family of functions of the form $\min(f_1, \dots, f_j)$ for $f_j \in \mathcal{K}, 1 \leq j \leq J$. Let \mathcal{J} denote the family of indicator functions $i(\mathbf{y}) = \mathbf{1}(f(\mathbf{y}) > 0)$ for $f \in \mathcal{F}$, and \mathcal{G} denote the family of functions $g = f \cdot i$ for $f \in \mathcal{F}$ and $i \in \mathcal{J}$. Suppose \mathcal{F} has an envelope function $f^* \in \mathcal{L}_1(\mathcal{Y}, F)$. Then, for each tolerance $\gamma \in (0, 1)$,

$$\lim_{T \rightarrow \infty} \text{Prob}(\sup_{T' \geq T} \sup_{i \in \mathcal{J}} |(\mathbf{E}_{T'} - \mathbf{E})i| > \gamma) = 0 \text{ and } \lim_{T \rightarrow \infty} \text{Prob}(\sup_{T' \geq T} \sup_{g \in \mathcal{G}} |(\mathbf{E}_{T'} - \mathbf{E})g| > \gamma) = 0.$$

Proof: The proof utilizes a geometric measure of the complexity of a family of functions \mathcal{F} or a family of sets \mathcal{C} , the *Vapnik-Červonenkis (VC) index*, denoted $V(\mathcal{F})$ or $V(\mathcal{C})$; see VW 2.6.1, Dudley (2014, 2.6.1). Classes of functions or sets with a finite VC index are termed *VC-classes*. VW Lemma 2.6.15 establishes that \mathcal{K} is a VC-class. Then VW, Lemma 2.6.18(ii) establishes that \mathcal{F} is a VC-class with index $V(\mathcal{F})$, and the truncated class $\mathcal{F}^M = \{\min(M, \max(f, -M)) \mid f \in \mathcal{F}\}$ for $M > 0$ is a VC-class with index at most $V(\mathcal{F}) + 2$; see Dudley (2014, Theorem

4.41). Pollard (1984, Lemma 2.4.18) establishes that the family \mathcal{C} of sets $C = \{y \in Y \mid f(y) > 0\}$ for $f \in \mathcal{F}$ is a VC-class, implying that \mathcal{J} is a VC-class (see VW, p. 151, #9). VW Theorem 2.6.7 applied to \mathcal{F}^M with envelope $f^* \equiv M$ or to \mathcal{J} with envelope $i^* \equiv 1$ implies bounds $N(\gamma, \mathcal{F}^M, Q) \leq K(M/\gamma)^{V(\mathcal{F})+2}$ and $N(\gamma, \mathcal{J}, Q) \leq K(1/\gamma)^{V(\mathcal{F})+2}$ for $\gamma \in (0,1)$ and any probability Q on \mathcal{Y} , where K is a constant that does not depend on γ or M . Let $\mathcal{F}_{\gamma/2}^M$ and $\mathcal{J}_{\gamma/2M}$ denote the sets of centers of open balls of radius $\gamma/2$ and $\gamma/2M$ that cover \mathcal{F}^M and \mathcal{J} respectively, and satisfy $\text{card}(\mathcal{F}_{\gamma/2}^M) \leq 2K(2M/\gamma)^{V(\mathcal{F})+2}$ and $\text{card}(\mathcal{J}_{\gamma/2M}) \leq 2K(2M/\gamma)^{V(\mathcal{F})+2}$. Let $\mathcal{G}^M = \{i \cdot f \mid i \in \mathcal{J} \text{ and } f \in \mathcal{F}^M\}$. For $i \in \mathcal{J}$ and $f \in \mathcal{F}^M$, one has $\min_{i' \in \mathcal{J}_{\gamma/2M}} \min_{f' \in \mathcal{F}_{\gamma/2}} \mathbf{E}|i \cdot f - i' \cdot f'| \leq \max_{f \in \mathcal{F}^M} \min_{i' \in \mathcal{J}_{\gamma/2M}} \mathbf{E}|(i - i') \cdot f| + \max_{i' \in \mathcal{J}_{\gamma/2M}} \min_{f' \in \mathcal{F}_{\gamma/2}} \mathbf{E}|(f - f') \cdot i'| < \gamma$. Then the covering number $N(\gamma, \mathcal{G}^M, Q)$ for any probability Q on \mathcal{Y} is bounded by the number of functions in $\mathcal{G}_\gamma = \{i \cdot f \mid i \in \mathcal{J}_{\gamma/2M} \text{ and } f \in \mathcal{F}_{\gamma/2}\}$, which is in turn bounded by $4K^2(2M/\gamma)^{2V(\mathcal{F})+4}$. The countable families $\bigcup_{k=1}^{\infty} \mathcal{F}_{1/2k}^M$ and $\bigcup_{k=1}^{\infty} \mathcal{J}_{1/2kM}$ are dense in \mathcal{F}^M and \mathcal{J} respectively, so that these families are F-measurable. Then VW Theorem 2.4.3 applies to give the result. ■

Appendix B: Properties of Extreme Value Type 1 Random Variables

a. A standard Extreme Value Type 1 (EV1) random variable has CDF $F(\varepsilon) \equiv \exp(-e^{-\varepsilon})$, density $e^{-\varepsilon} \cdot \exp(-e^{-\varepsilon})$, and for $t < 1$ the moment generating function $\Gamma(1-t)$. Johnson and Kotz (1970, Ch. 21) show the linear transformation $\xi = v + \sigma\varepsilon$ with $\sigma > 0$ has CDF $\exp(-e^{-(\xi-v)/\sigma})$, mean $v + \sigma\gamma_0$, where $\gamma_0 = 0.5772156649\dots$ is Euler's constant, median $v - \sigma \ln \ln 2$, mode v , and variance $\sigma^2\pi^2/6$. For $0 < \rho < 0.08$, the tails of $F(\varepsilon)$ satisfy $F(2 \cdot \ln \rho) + 1 - F(-2 \cdot \ln \rho) < \rho$ and $\int_{|\varepsilon| > -2 \cdot \ln \rho} |\varepsilon| F(d\varepsilon) < \rho$. Also, $\mathbf{E}|\varepsilon| \leq 1.219384$ (i.e., integrating by parts, $\mathbf{E}|\varepsilon| = \int_0^{\infty} \exp(-e^{-\varepsilon}) d\varepsilon + \int_0^{\infty} [1 - \exp(-e^{-\varepsilon})] d\varepsilon \leq E_1(1) + \int_0^{\infty} \exp(-\varepsilon) d\varepsilon$, where $E_1(c) \equiv \int_c^{\infty} \frac{e^{-y}}{y} dy$ is the exponential integral with values given in Abramovitz and Stegun, 1964, Table 5.1). Finally, $\mathbf{E}\varepsilon^2 = \gamma_0^2 + \pi^2/6 = 1.978112\dots$

b. Consider $\mathbf{J} = \{0, \dots, J\}$, constants a_j and independent standard EV1 random variates ε_j for $j \in \mathbf{J}$, and a non-empty subset \mathbf{C} of \mathbf{J} . Define $q_{\mathbf{C}} = \ln \sum_{j \in \mathbf{C}} e^{a_j}$ and $\xi_{\mathbf{C}} = \max_{j \in \mathbf{C}} (a_j + \varepsilon_j) - q_{\mathbf{C}}$. Then $\xi_{\mathbf{C}}$ is again a standard EV1 random variable; i.e., $\text{Prob}(\xi_{\mathbf{C}} < c) = \text{Prob}(\varepsilon_j < c + q_{\mathbf{C}} - a_j \text{ for } j \in \mathbf{C}) = \prod_{j \in \mathbf{C}} \exp(-e^{-c - q_{\mathbf{C}} + a_j}) \equiv \exp(-e^{-c})$. Given $k \in \mathbf{C}$, the probability of the event $Y_{\mathbf{C}}(k) = \{\varepsilon \mid a_k + \varepsilon_k \geq a_j + \varepsilon_j \text{ for } j \in \mathbf{C}\}$ is multinomial logit,

$$L_{\mathbf{C}}(k) = \int_{\varepsilon_k = -\infty}^{+\infty} f(\varepsilon_k) \prod_{j \in \mathbf{C} \setminus \{k\}} F(\varepsilon_k + a_k - a_j) d\varepsilon_k = \int_{\varepsilon_k = -\infty}^{+\infty} e^{-\varepsilon_k} \exp(-e^{-\varepsilon_k} \sum_{j \in \mathbf{C}} e^{a_j - a_k}) d\varepsilon_k = \frac{e^{a_k}}{\sum_{j \in \mathbf{C}} e^{a_j}},$$

and for $\mathbf{A} \subseteq \mathbf{C}$, $L_{\mathbf{C}}(\mathbf{A}) = \frac{\sum_{j \in \mathbf{A}} e^{a_j}}{\sum_{j \in \mathbf{C}} e^{a_j}}$. The conditional CDF of ε_k , given $k \in \mathbf{C}$ and $Y_{\mathbf{C}}(k)$, is

$$\begin{aligned} \text{Prob}(\varepsilon_k < c \mid Y_{\mathbf{C}}(k)) &= \frac{1}{L_{\mathbf{C}}(k)} \int_{\varepsilon_k = -\infty}^c f(\varepsilon_k) \prod_{j \in \mathbf{C} \setminus \{k\}} F(\varepsilon_k + a_k - a_j) d\varepsilon_k \\ &= \frac{1}{L_{\mathbf{C}}(k)} \int_{\varepsilon_k = -\infty}^c e^{-\varepsilon_k} \exp(-e^{-\varepsilon_k} \sum_{j \in \mathbf{C}} e^{a_j - a_k}) d\varepsilon_k = \exp(-e^{-(c + a_k - q_{\mathbf{C}})}) \equiv F(c + a_k - q_{\mathbf{C}}). \end{aligned}$$

Then the payoff $a_k + \varepsilon_k$, conditioned on the event $Y_{\mathbf{C}}(k)$, has the same CDF as $\xi_{\mathbf{C}} + q_{\mathbf{C}}$, and is therefore the same for all k . Term this the *Optimizer Invariance Property* (OIP). An immediate implication of OIP is

$$\gamma_0 + q_C = E(\xi_C + q_C) \equiv E \max_{j \in C} (a_j + \varepsilon_j) \equiv E\{a_k + \varepsilon_k \mid Y_C(k)\}.$$

so these unconditional and conditional means are the same. This result is obtained in Dubin and McFadden (1984), Anas and Feng (1988), Resnick and Roy (1990), and Dubin (1985). A consequence is that when **B** and **C** are disjoint non-empty subsets of **J**, then the conditional (on $Y_C(k)$ for some $k \in C$) and unconditional expectations of utility differences are given by the same log sum difference:

$$E \left\{ \max_{j \in B} (a_j + \varepsilon_j) - \max_{j \in C} (a_j + \varepsilon_j) \mid Y_C(k) \right\} \equiv E \left\{ \max_{j \in B} (a_j + \varepsilon_j) - \max_{j \in C} (a_j + \varepsilon_j) \right\} \equiv \ln \frac{\sum_{j \in B} e^{a_j}}{\sum_{j \in C} e^{a_j}}.$$

If $k \notin C$, then the conditional CDF of ξ_C , given $a_k + \varepsilon_k > \xi_C + q_C$, is

$$\begin{aligned} \text{Prob}(\xi_C < c \mid a_k + \varepsilon_k > \xi_C + q_C) &= \frac{1}{L_{CU\{k\}}(k)} \int_{\xi_C = -\infty}^c f(\xi_C) [1 - F(\xi_C + q_C - a_k)] d\xi_C \\ &= \frac{F(c)}{L_{CU\{k\}}(k)} - \frac{\int_{\xi_C = -\infty}^c e^{-\xi_C} \exp(-e^{-\xi_C} [1 + e^{a_k - q_C}]) d\xi_C}{L_{CU\{k\}}(k)} = \frac{F(c) - L_{CU\{k\}}(C) F(c + q_C - \ln[e^{q_C} + e^{a_k}])}{L_{CU\{k\}}(k)}. \end{aligned}$$

Using the OIP, this result is unchanged if instead of a single alternative $k \notin C$, there is a set of alternatives **A** with $A \cap C = \emptyset$ and either k maximizes $a_j + \varepsilon_j$ for $j \in A$, with conditioning on the event $Y_A(k)$, or $q_A = \ln \sum_{j \in A} e^{a_j}$ replaces a_k , a standard EV1 variate ξ_A replaces ε_k , and **A** replaces $\{k\}$.

Next, given $k \notin C$, the conditional CDF of ξ_C , given $a_k + \varepsilon_k < \xi_C + q_C$, is

$$\begin{aligned} \text{Prob}(\xi_C < c \mid a_k + \varepsilon_k < \xi_C + q_C) &= \frac{1}{L_{CU\{k\}}(C)} \int_{\xi_C = -\infty}^c f(\xi_C) F(\xi_C + q_C - a_k) d\xi_C \\ &= \frac{1}{L_{CU\{k\}}(C)} \int_{\xi_C = -\infty}^c e^{-\xi_C} \exp(-e^{-\xi_C} [1 + e^{a_k - q_C}]) d\xi_C = F(c + q_C - \ln[e^{q_C} + e^{a_k}]). \end{aligned}$$

Again by the OIP, this result is unchanged if $k \in B$ with $B \cap C = \emptyset$, $q_B = \ln \sum_{j \in B} e^{a_j}$ replaces a_k , a standard EV1 variate ξ_B replaces ε_k , and **A** replaces $\{k\}$.

c. Let **A**, **B**, **C** denote disjoint non-empty subsets of **J**. Define $q_A = \ln \sum_{j \in A} e^{a_j}$, and define q_B and q_C analogously. Define $\xi_A = \max_{j \in A} (a_j + \varepsilon_j) - q_A$, with analogous definitions for ξ_B and ξ_C , and let “**ABC**” denote the event $\xi_A + q_A > \xi_B + q_B > \xi_C + q_C$, and so on. The possible events and outcomes are given below:

	ABC	ACB	BAC	BCA	CAB	CBA
Choice at a	A	A	A	C	C	C
Choice at b	B	C	B	B	C	C
Type	difference	difference	compound	compound	compound	compound

The probability of the event **ABC** is

$$\begin{aligned}
P(\mathbf{ABC}) &= \int_{\xi_B=-\infty}^{+\infty} f(\xi_B)F(\xi_B + q_B - q_C)[1 - F(\xi_B + q_B - q_A)]d\xi_B \\
&= \int_{\xi_B=-\infty}^{+\infty} e^{-\xi_B} \exp(-e^{-\xi_B} [1 + e^{q_C - q_B}]) d\xi_B - \int_{\xi_B=-\infty}^{+\infty} e^{-\xi_B} \exp(-e^{-\xi_B} [1 + e^{q_A - q_B} + e^{q_C - q_B}]) d\xi_B \\
&= P(\mathbf{B|B,C}) - P(\mathbf{B|A,B,C}) \equiv P(\mathbf{B|B,C}) \cdot P(\mathbf{A|A,B,C}),
\end{aligned}$$

where $P(\mathbf{A|A,B,C}) = e^{q_A} / (e^{q_A} + e^{q_B} + e^{q_C})$ and $P(\mathbf{B|B,C}) = e^{q_B} / (e^{q_B} + e^{q_C})$. This formula gives the probability of any other of the events by substituting the corresponding permutation of $\mathbf{A, B, C}$. Next,

$$\begin{aligned}
E\{(\xi_B + q_B) \cdot \mathbf{1}(\mathbf{ABC})\} &= \int_{\xi_B=-\infty}^{+\infty} (\xi_B + q_B) f(\xi_B) F(\xi_B + q_B - q_C) [1 - F(\xi_B + q_B - q_A)] d\xi_B \\
&= \int_{\xi_B=-\infty}^{+\infty} (\xi_B + q_B) e^{-\xi_B} \{ \exp(-e^{-\xi_B} [1 + e^{q_C - q_B}]) - \exp(-e^{-\xi_B} [1 + e^{q_A - q_B} + e^{q_C - q_B}]) \} d\xi_B \\
&= \{ P(\mathbf{B|B,C}) [\gamma_0 + \ln(e^{q_B} + e^{q_C})] - P(\mathbf{B|A,B,C}) [\gamma_0 + \ln(e^{q_A} + e^{q_B} + e^{q_C})] \},
\end{aligned}$$

The event \mathbf{ACB} also has an expectation satisfying this “difference type” formula with \mathbf{B} and \mathbf{C} interchanged.

The event \mathbf{BAC} has

$$\begin{aligned}
E\{(\xi_B + q_B) \cdot \mathbf{1}(\mathbf{BAC})\} &= \int_{\xi_B=-\infty}^{+\infty} (\xi_B + q_B) f(\xi_B) \int_{\xi_A=-\infty}^{\xi_B + q_B - q_A} f(\xi_A) F(\xi_A + q_A - q_C) d\xi_A d\xi_B \\
&= \int_{\xi_B=-\infty}^{+\infty} (\xi_B + q_B) f(\xi_B) \int_{\xi_A=-\infty}^{\xi_B + q_B - q_A} e^{-\xi_A} \exp(-e^{-\xi_A} [1 + e^{q_C - q_A}]) d\xi_A d\xi_B \\
&= P(\mathbf{A|A,C}) \int_{\xi_B=-\infty}^{+\infty} (\xi_B + q_B) e^{-\xi_B} \exp(-e^{-\xi_B}) \exp(-e^{-\xi_B + q_A - q_B} [1 + e^{q_C - q_A}]) d\xi_B \\
&= P(\mathbf{A|A,C}) \int_{\xi_B=-\infty}^{+\infty} (\xi_B + q_B) e^{-\xi_B} \exp(-e^{-\xi_B} [1 + e^{q_A - q_B} + e^{q_C - q_B}]) d\xi_B \\
&= P(\mathbf{A|A,C}) P(\mathbf{B|A,B,C}) \{ \gamma_0 + \ln(e^{q_A} + e^{q_B} + e^{q_C}) \}.
\end{aligned}$$

The events \mathbf{BCA} , \mathbf{CAB} , and \mathbf{CBA} also have expectations satisfying this “compound type” formula with the corresponding permutations of $\mathbf{A, B}$, and \mathbf{C} .

Consider the event \mathbf{AC} . From (b), $E \max(\xi_A + q_A, \xi_C + q_C) \cdot \mathbf{1}(\mathbf{AC}) = [\gamma_0 + \ln(e^{q_A} + e^{q_C})] \cdot P(\mathbf{A|A,C})$. Then,

$$\begin{aligned}
&E\{[\max(\xi_B + q_B, \xi_C + q_C) - \max(\xi_A + q_A, \xi_C + q_C)] \cdot \mathbf{1}(\mathbf{AC})\} \\
&= E\{[\xi_B + \frac{q_B}{\sigma} - (\xi_A + q_A)] \cdot \mathbf{1}(\mathbf{ABC}) + E\{[\xi_C + q_C - (\xi_A + q_A)] \cdot \mathbf{1}(\mathbf{ACB}) + E\{[\xi_B + q_B - (\xi_A + q_A)] \cdot \mathbf{1}(\mathbf{BAC})\} \\
&= \{ P(\mathbf{B|B,C}) [\gamma_0 + \ln(e^{q_B} + e^{q_C})] - P(\mathbf{B|A,B,C}) [\gamma_0 + \ln(e^{q_A} + e^{q_B} + e^{q_C})] \} \\
&\quad + \{ P(\mathbf{C|B,C}) [\gamma_0 + \ln(e^{q_B} + e^{q_C})] - P(\mathbf{C|A,B,C}) [\gamma_0 + \ln(e^{q_A} + e^{q_B} + e^{q_C})] \} \\
&\quad + P(\mathbf{A|A,C}) P(\mathbf{B|A,B,C}) \{ \gamma_0 + \ln(e^{q_A} + e^{q_B} + e^{q_C}) \} - [\gamma_0 + \ln(e^{q_A} + e^{q_C})] \cdot P(\mathbf{A|A,C})
\end{aligned}$$

$$\begin{aligned}
&= -P(\mathbf{C}|\mathbf{A},\mathbf{C}) \cdot \ln(e^{q_A} + e^{q_B} + e^{q_C}) + \ln(e^{q_B} + e^{q_C}) - P(\mathbf{A}|\mathbf{A},\mathbf{C}) \cdot \ln(e^{q_A} + e^{q_C}) \\
&= P(\mathbf{A}|\mathbf{A},\mathbf{C}) \cdot \ln \frac{e^{q_B+e^{q_C}}}{e^{q_A+e^{q_C}}} + P(\mathbf{C}|\mathbf{A},\mathbf{C}) \cdot \ln \frac{e^{q_B+e^{q_C}}}{e^{q_A+e^{q_B+e^{q_C}}}}.
\end{aligned}$$

Hence,

$$\mathbf{E}\{[\max(\xi_B + q_B, \xi_C + q_C) - \max(\xi_A + q_A, \xi_C + q_C)] \mid \mathbf{AC}\} = \ln \frac{e^{q_B+e^{q_C}}}{e^{q_A+e^{q_C}}} + \frac{P(\mathbf{C}|\mathbf{A},\mathbf{C})}{P(\mathbf{A}|\mathbf{A},\mathbf{C})} \cdot \ln \frac{e^{q_B+e^{q_C}}}{e^{q_A+e^{q_B+e^{q_C}}}}.$$

The first term in the last expression coincides with the unconditional expectation of the maximum, and the final term adjusts for the conditioning event. The adjustment is negative so that the information that the best in **A** is better than the best in **C** decreases the expected maximum utility over **B** and **C**. By application of the OIC as described at the end of (b), this result is the same no matter which event $Y_A(k)$ occurs. Next,

$$\begin{aligned}
&\mathbf{E}\{[\max(\xi_B + q_B, \xi_C + q_C) - \max(\xi_A + q_A, \xi_C + q_C)] \cdot \mathbf{1}(\mathbf{CA})\} \\
&= \mathbf{E}\{[\xi_B + q_B - (\xi_C + q_C)] \cdot \mathbf{1}(\mathbf{BCA})\} = P(\mathbf{C}|\mathbf{A},\mathbf{C})P(\mathbf{B}|\mathbf{A},\mathbf{B},\mathbf{C})[\gamma_0 + \ln(e^{q_A} + e^{q_B} + e^{q_C})] \\
&\quad - P(\mathbf{C}|\mathbf{A},\mathbf{C}) \cdot [\gamma_0 + \ln(e^{q_A} + e^{q_C})] + P(\mathbf{C}|\mathbf{A},\mathbf{B},\mathbf{C}) \cdot [\gamma_0 + \ln(e^{q_A} + e^{q_B} + e^{q_C})] \\
&= P(\mathbf{C}|\mathbf{A},\mathbf{C}) \cdot \left\{ \ln \frac{e^{q_B+e^{q_C}}}{e^{q_A+e^{q_C}}} - \ln \frac{e^{q_B+e^{q_C}}}{e^{q_A+e^{q_B+e^{q_C}}}} \right\},
\end{aligned}$$

$$\text{and hence } \mathbf{E}\{\max(\xi_B + q_B, \xi_C + q_C) - \max(\xi_A + q_A, \xi_C + q_C) \mid \mathbf{CA}\} = \ln \frac{e^{q_B+e^{q_C}}}{e^{q_A+e^{q_C}}} - \ln \frac{e^{q_B+e^{q_C}}}{e^{q_A+e^{q_B+e^{q_C}}}}.$$

As before, the first term in the last expression coincides with the unconditional expectation of the maximum, and the final term is a positive adjustment for the conditioning event, so that the information that the best in **C** is better than the best in **A** increases the expected maximum utility over **B** and **C**. Again, by application of the OIC, this result is the same no matter which event $Y_A(k)$ occurs.

d. Now consider $\mathbf{J} = \{0, \dots, J\}$ and $\mathbf{C} = \{0, \dots, J-1\}$. Assume that in a scenario change from $m = a$ to $m = b$, constants $a_{jm} \equiv a_j$ for $j \in \mathbf{C}$ do not change, but $a_{Ja} \neq a_{Jb}$. Assume ε_j for $j \in \mathbf{J}$ is the same in both scenarios. Define $q_C = \ln \sum_{j \in \mathbf{C}} e^{a_j}$ and $\xi_C = \max_{j \in \mathbf{C}} (a_j + \varepsilon_j) - q_C$. There is an alternative k that maximizes $a_j + \varepsilon_j$ over $j \in \mathbf{C}$, and from (b), the CDF of $a_k + \varepsilon_k$ given that k maximizes the payoff in **C** is the same as the CDF of $\xi_C + q_C$. Define $\omega = \xi_C - \varepsilon_J$ and $L(w) \equiv \text{Prob}(\omega \leq w) = 1/(1+e^{-w})$. The possible events are then:

Event	Case	Condition	Probability	Payoff
Y_{bak}	$a_{Ja} < a_{Jb}$	$\xi_C + q_C < a_{Ja} + \varepsilon_J < a_{Jb} + \varepsilon_J$	$L(a_{Ja} - q_C)$	$a_{Jb} - a_{Ja}$
Y_{bka}	$a_{Ja} < a_{Jb}$	$a_{Ja} + \varepsilon_J < \xi_C + q_C < a_{Jb} + \varepsilon_J$	$L(a_{Jb} - q_C) - L(a_{Ja} - q_C)$	$a_{Jb} - q_C - \omega$
Y_{kba}	$a_{Ja} < a_{Jb}$	$a_{Ja} + \varepsilon_J < a_{Jb} + \varepsilon_J < \xi_C + q_C$	$L(q_C - a_{Jb})$	0
Y_{abk}	$a_{Ja} > a_{Jb}$	$\xi_C + q_C < a_{Jb} + \varepsilon_J < a_{Ja} + \varepsilon_J$	$L(a_{Jb} - q_C)$	$a_{Jb} - a_{Ja}$
Y_{akb}	$a_{Ja} > a_{Jb}$	$a_{Jb} + \varepsilon_J < \xi_C + q_C < a_{Ja} + \varepsilon_J$	$L(a_{Ja} - q_C) - L(a_{Jb} - q_C)$	$q_C - a_{Ja} + \omega$
Y_{kab}	$a_{Ja} > a_{Jb}$	$a_{Jb} + \varepsilon_J < a_{Ja} + \varepsilon_J < \xi_C + q_C$	$L(q_C - a_{Ja})$	0

Note that $\int_s^t \omega dL(\omega) = \omega L(\omega)]_s^t - \int_s^t \frac{e^\omega}{1+e^\omega} d\omega = tL(t) - sL(s) - \ln \frac{1+e^t}{1+e^s}$. Then the expected payoff in the event \mathbf{Y}_{bka} is

$$\begin{aligned} E\{(a_{Jb} - q_c - \omega) \cdot \mathbf{1}(\omega \in \mathbf{Y}_{bka})\} &= (a_{Jb} - q_c)[L(a_{Jb} - q_c) - L(a_{Ja} - q_c)] - \int_{a_{Ja} - q_c}^{a_{Jb} - q_c} \omega dL(\omega) \\ &= (a_{Jb} - q_c)[L(a_{Jb} - q_c) - L(a_{Ja} - q_c)] - (a_{Jb} - q_c)L(a_{Jb} - q_c) + (a_{Ja} - q_c)L(a_{Ja} - q_c) + \ln \frac{1+e^{a_{Jb} - q_c}}{1+e^{a_{Ja} - q_c}} \\ &= (a_{Ja} - a_{Jb}) L(a_{Ja} - q_c) + \ln \frac{e^{q_c} + e^{a_{Jb}}}{e^{q_c} + e^{a_{Ja}}}, \end{aligned}$$

and the expected payoff in the event \mathbf{Y}_{akb} is

$$\begin{aligned} E\{(q_c - a_{Ja} + \omega) \cdot \mathbf{1}(\omega \in \mathbf{Y}_{akb})\} &= (q_c - a_{Ja})[L(a_{Ja} - q_c) - L(a_{Jb} - q_c)] + \int_{a_{Jb} - q_c}^{a_{Ja} - q_c} \omega dL(\omega) \\ &= (q_c - a_{Ja})[L(a_{Ja} - q_c) - L(a_{Jb} - q_c)] + (a_{Ja} - q_c)L(a_{Ja} - q_c) - (a_{Jb} - q_c)L(a_{Jb} - q_c) - \ln \frac{1+e^{a_{Ja} - q_c}}{1+e^{a_{Jb} - q_c}} \\ &= (a_{Ja} - a_{Jb}) L(a_{Jb} - q_c) + \ln \frac{e^{q_c} + e^{a_{Jb}}}{e^{q_c} + e^{a_{Ja}}}. \end{aligned}$$

Combining these results with other payoffs in the table gives

Scenario <i>a</i> Choice	Case	Expected Payoff Given Choice
J	$a_{Ja} < a_{Jb}$	$a_{Jb} - a_{Ja}$
J	$a_{Ja} > a_{Jb}$	$\frac{1}{L(a_{Ja} - q_c)} \ln \frac{e^{q_c} + e^{a_{Jb}}}{e^{q_c} + e^{a_{Ja}}}$
k	$a_{Ja} < a_{Jb}$	$-\frac{L(a_{Ja} - q_c)}{L(q_c - a_{Ja})} (a_{Jb} - a_{Ja}) + \frac{1}{L(q_c - a_{Ja})} \ln \frac{e^{q_c} + e^{a_{Jb}}}{e^{q_c} + e^{a_{Ja}}}$
k	$a_{Ja} > a_{Jb}$	0

e. Assume scenarios $m = a, b$, a set of alternatives $\mathbf{J}_a = \mathbf{J}_b = \mathbf{J} = \{0, \dots, J\}$, and noise $\boldsymbol{\varepsilon}$ that is the same in both scenarios. Let a_{jm} denote constants. Order the alternatives so that $\Delta_i \equiv a_{ib} - a_{ia}$ is non-decreasing in i . Define non-decreasing constants $c_i = \Delta_i + a_{ka} - a_{rb}$; then $c_k = a_{kb} - a_{rb}$, and $c_r = a_{ka} - a_{ra}$. Let \mathbf{A}_{jm} denote the event that alternative j is optimal in scenario m . Consider the event

$$\mathbf{B}_{kr} = \mathbf{A}_{ka} \cap \mathbf{A}_{rb} = \{\boldsymbol{\varepsilon} \mid \varepsilon_k + a_{ka} \geq \varepsilon_i + a_{ia} \text{ for } i \neq k \ \& \ \varepsilon_r + a_{rb} \geq \varepsilon_i + a_{ib} \text{ for } i \neq r\},$$

including both cases $k = r$ and $k \neq r$. The \mathbf{B}_{kr} are disjoint for different k or for different r except for sets of probability zero, and satisfy $\mathbf{A}_{ka} = \bigcup_{r=0}^J \mathbf{B}_{kr}$ and $\mathbf{A}_{rb} = \bigcup_{k=0}^J \mathbf{B}_{kr}$. The event \mathbf{B}_{kr} implies $(a_{rb} - a_{kb}) \geq \varepsilon_k - \varepsilon_r \geq (a_{ra} - a_{ka})$. Hence, \mathbf{B}_{kr} is non-empty if and only if $a_{rb} - a_{ra} \geq a_{kb} - a_{ka}$, or equivalently $c_k \leq c_r$, implying $r \geq k$.

Consider the case $r = k$. Then for $i \neq k$, $\varepsilon_i \leq \varepsilon_k + a_{ka} - \max(a_{ia}, a_{ib} + a_{ka} - a_{kb})$. Then

$$P(B_{kk}) = \int_{\varepsilon_k = -\infty}^{+\infty} e^{-\varepsilon_k} \exp\left(-e^{-\varepsilon_k} \sum_{i=0}^J e^{-a_{ka} + \max(a_{ia}, a_{ib} + a_{ka} - a_{kb})}\right) d\varepsilon_k = \frac{e^{a_{ka}}}{\sum_{i=0}^J e^{\max(a_{ia}, a_{ib} + a_{ka} - a_{kb})}}.$$

Then, the conditional probability $P_{kb|ka}$ that the optimal choice in scenario b is k , given that the optimal choice in scenario a is k , satisfies

$$P_{kb|ka} \equiv P(B_{kk})/P(A_{ka}) = \frac{\sum_{i=0}^J e^{a_{ia}}}{\sum_{i=0}^J e^{\max(a_{ia}, a_{ib} + a_{ka} - a_{kb})}}.$$

If $r < k$, B_{kr} is empty and $P_{rb|ka} = 0$. Finally, consider the case $r > k$. Then $\varepsilon \in B_{kr}$ requires $\varepsilon_r + (a_{ra} - a_{ka}) \equiv \varepsilon_r - c_r \leq \varepsilon_k \leq \varepsilon_r + (a_{rb} - a_{kb}) \equiv \varepsilon_r - c_k$. Let $B_{kri} = \{\varepsilon \in B_{kr} \mid \varepsilon_r - c_{i+1} \leq \varepsilon_k \leq \varepsilon_r - c_i\}$ for $i = k, \dots, r-1$ and consider the inequalities $\varepsilon_n \leq \min(\varepsilon_k + (a_{ka} - a_{na}), \varepsilon_r + (a_{rb} - a_{nb})) = (a_{ka} - a_{na}) + \min(\varepsilon_k, \varepsilon_r - c_n)$. If $n > i$, then $\varepsilon_r - c_n \leq \varepsilon_k$, implying $\varepsilon_n \leq (a_{ka} - a_{na} - c_n) + \varepsilon_r$; otherwise, $\varepsilon_n \leq (a_{ka} - a_{na}) + \varepsilon_k$. The probability of B_{kri} is then

$$\begin{aligned} P_{kri} &= \int_{\varepsilon_r = -\infty}^{+\infty} e^{-\varepsilon_r} \exp\left(-e^{-\varepsilon_r} \sum_{n=i+1}^J e^{(a_{na} - a_{ka} + c_n)}\right) \int_{\varepsilon_k = \varepsilon_r - c_{i+1}}^{\varepsilon_r - c_i} e^{-\varepsilon_k} \exp\left(-e^{-\varepsilon_k} \sum_{n=0}^i e^{(a_{na} - a_{ka})}\right) d\varepsilon_k d\varepsilon_r \\ &= \frac{e^{a_{ka}}}{\sum_{n=0}^i e^{a_{na}}} \left\{ \int_{\varepsilon_r = -\infty}^{+\infty} e^{-\varepsilon_r} \exp\left(-e^{-\varepsilon_r} \left[\sum_{n=i+1}^J e^{(a_{na} - a_{rb} + \Delta_n)} + \sum_{n=0}^i e^{a_{na} - a_{rb} + \Delta_i} \right]\right) d\varepsilon_r \right. \\ &\quad \left. - \int_{\varepsilon_r = -\infty}^{+\infty} e^{-\varepsilon_r} \exp\left(-e^{-\varepsilon_r} \left[\sum_{n=i+1}^J e^{a_{na} - a_{rb} + \Delta_n} + \sum_{n=0}^i e^{a_{na} - a_{rb} + \Delta_{i+1}} \right]\right) d\varepsilon_r \right\} \\ &= \frac{e^{a_{ka}}}{\sum_{n=0}^i e^{a_{na}}} \left\{ \int_{\varepsilon_r = -\infty}^{+\infty} e^{-\varepsilon_r} \exp\left(-e^{-\varepsilon_r} \left[\sum_{n=0}^J e^{(a_{na} - a_{rb} + \max(\Delta_i, \Delta_n))}\right]\right) d\varepsilon_r \right. \\ &\quad \left. - \int_{\varepsilon_r = -\infty}^{+\infty} e^{-\varepsilon_r} \exp\left(-e^{-\varepsilon_r} \left[\sum_{n=0}^J e^{(a_{na} - a_{rb} + \max(\Delta_{i+1}, \Delta_n))}\right]\right) d\varepsilon_r \right\} \\ &= \frac{e^{a_{ka}}}{\sum_{n=0}^i e^{a_{na}}} \left[\frac{e^{a_{rb}}}{\sum_{n=0}^J e^{a_{na} + \max(\Delta_i, \Delta_n)}} - \frac{e^{a_{rb}}}{\sum_{n=0}^J e^{a_{na} + \max(\Delta_{i+1}, \Delta_n)}} \right] = \frac{e^{a_{ka}} e^{a_{rb}}}{e^{a_{na} + \max(\Delta_i, \Delta_n)}} \cdot \frac{(e^{\Delta_{i+1}} - e^{\Delta_i})}{\sum_{n=0}^J e^{a_{na} + \max(\Delta_{i+1}, \Delta_n)}}. \end{aligned}$$

Then, the probability that the optimal choice in scenario b is $r > k$, given that the optimal choice in scenario a is k , satisfies $P_{rb|ka} \equiv \sum_{i=k}^{r-1} P_{kri}/P_{ka} = \sum_{i=k}^{r-1} \frac{e^{a_{rb}} \sum_{n=0}^i e^{a_{na}}}{\sum_{n=0}^J e^{a_{na} + \max(\Delta_i, \Delta_n)}} \cdot \frac{(e^{\Delta_{i+1}} - e^{\Delta_i})}{\sum_{n=0}^J e^{a_{na} + \max(\Delta_{i+1}, \Delta_n)}}$.

Next, consider the expectations of $(a_{kb} + \varepsilon_k) \cdot \mathbf{1}(\varepsilon \in B_{kk})$ for $r = k$ and $(a_{rb} + \varepsilon_r) \cdot \mathbf{1}(\varepsilon \in B_{kri})$ for $k \leq i < r$, given $\varepsilon \in A_{ka}$. First,

$$\begin{aligned} \mathbf{E}_{\varepsilon|A_{ka}}(a_{kb} + \varepsilon_k) \cdot \mathbf{1}(\varepsilon \in B_{kk}) &= \frac{1}{P_{ka}} \int_{\varepsilon_k = -\infty}^{+\infty} (a_{kb} + \varepsilon_k) e^{-\varepsilon_k} \exp\left(-e^{-\varepsilon_k} \sum_{i=0}^J e^{-a_{ka} + \max(a_{ia}, a_{ib} + a_{ka} - a_{kb})}\right) d\varepsilon_k \\ &= P_{kb|ka} [a_{kb} - a_{ka} + \ln(\sum_{i=0}^J e^{\max(a_{ia}, a_{ib} + a_{ka} - a_{kb})}) + \gamma_0]. \end{aligned}$$

Second, for $k \leq i < r$,

$$\begin{aligned}
\mathbf{E}_{\boldsymbol{\varepsilon}|A_{ka}}(a_{rb} + \varepsilon_r) \cdot \mathbf{1}(\boldsymbol{\varepsilon} \in B_{kri}) &= \frac{P_{kri}}{P_{ka}} a_{rb} \\
&+ \frac{\sum_{n=0}^J e^{a_{na}}}{\sum_{n=0}^i e^{a_{na}}} \left\{ \int_{\varepsilon_r=-\infty}^{+\infty} \varepsilon_r e^{-\varepsilon_r} \exp\left(-e^{-\varepsilon_r} \left[\sum_{n=0}^J e^{a_{na} - a_{rb} + \max(\Delta_i, \Delta_n)}\right]\right) d\varepsilon_r \right. \\
&\quad \left. - \int_{\varepsilon_r=-\infty}^{+\infty} \varepsilon_r e^{-\varepsilon_r} \exp\left(-e^{-\varepsilon_r} \left[\sum_{n=0}^J e^{(a_{na} - a_{rb} + \max(\Delta_{i+1}, \Delta_n))}\right]\right) d\varepsilon_r \right\} \\
&= \frac{P_{kri}}{P_{ka}} \gamma_0 + \frac{\sum_{n=0}^J e^{a_{na}} e^{a_{rb}}}{\sum_{n=0}^i e^{a_{na}}} \left[\frac{1}{\sum_{n=0}^J e^{a_{na} + \max(\Delta_i, \Delta_n)}} - \frac{1}{\sum_{n=0}^J e^{(a_{na} + \max(\Delta_{i+1}, \Delta_n))}} \right] \cdot \ln \sum_{n=0}^J e^{a_{na} + \max(\Delta_i, \Delta_n)} \\
&\quad + \frac{\sum_{n=0}^J e^{a_{na}} e^{a_{rb}}}{\sum_{n=0}^i e^{a_{na}}} \cdot \frac{1}{\sum_{n=0}^J e^{a_{na} + \max(\Delta_{i+1}, \Delta_n)}} \cdot \ln \frac{\sum_{n=0}^J e^{a_{na} + \max(\Delta_i, \Delta_n)}}{\sum_{n=0}^J e^{a_{na} + \max(\Delta_{i+1}, \Delta_n)}} \\
&= \frac{P_{kri}}{P_{ka}} [\gamma_0 + \ln \sum_{n=0}^J e^{a_{na} + \max(\Delta_i, \Delta_n)}] + \frac{e^{a_{rb}}}{\sum_{n=0}^i e^{a_{na}}} \cdot \frac{\sum_{n=0}^J e^{a_{na}}}{\sum_{n=0}^J e^{a_{na} + \max(\Delta_{i+1}, \Delta_n)}} \cdot \ln \frac{\sum_{n=0}^J e^{a_{na} + \max(\Delta_i, \Delta_n)}}{\sum_{n=0}^J e^{a_{na} + \max(\Delta_{i+1}, \Delta_n)}}
\end{aligned}$$

Hence, the conditional expectation of $(a_{rb} + \varepsilon_r) \cdot \mathbf{1}(\boldsymbol{\varepsilon} \in B_{kr})$ given $\boldsymbol{\varepsilon} \in A_{ka}$ is

$$P_{rb|ka} \gamma_0 + \sum_{i=k}^{r-1} \frac{P_{kri}}{P_{ka}} \cdot \ln \sum_{n=0}^J e^{a_{na} + \max(\Delta_i, \Delta_n)} + \sum_{i=k}^{r-1} \frac{1}{\sum_{n=0}^i e^{a_{na}}} \cdot \frac{e^{a_{rb}} \cdot \sum_{n=0}^J e^{a_{na}}}{\sum_{n=0}^J e^{a_{na} + \max(\Delta_{i+1}, \Delta_n)}} \cdot \ln \frac{\sum_{n=0}^J e^{a_{na} + \max(\Delta_i, \Delta_n)}}{\sum_{n=0}^J e^{a_{na} + \max(\Delta_{i+1}, \Delta_n)}}.$$

Summing this expression over $r > k$ gives

$$\begin{aligned}
&(1 - P_{kb|ka}) \gamma_0 + \sum_{r>k} \sum_{i=k}^{r-1} \frac{P_{kri}}{P_{ka}} \cdot \ln \sum_{n=0}^J e^{a_{na} + \max(\Delta_i, \Delta_n)} \\
&\quad + \sum_{r>k} \sum_{i=k}^{r-1} \frac{1}{\sum_{n=0}^i e^{a_{na}}} \cdot \frac{e^{a_{rb}} \cdot \sum_{n=0}^J e^{a_{na}}}{\sum_{n=0}^J e^{a_{na} + \max(\Delta_{i+1}, \Delta_n)}} \cdot \ln \frac{\sum_{n=0}^J e^{a_{na} + \max(\Delta_i, \Delta_n)}}{\sum_{n=0}^J e^{a_{na} + \max(\Delta_{i+1}, \Delta_n)}} \\
&= (1 - P_{kb|ka}) \gamma_0 + \sum_{i=k}^{J-1} \sum_{r=i+1}^J \frac{P_{kri}}{P_{ka}} \cdot \ln \sum_{n=0}^J e^{a_{na} + \max(\Delta_i, \Delta_n)} \\
&\quad + \sum_{i=k}^{J-1} \frac{1}{\sum_{n=0}^i e^{a_{na}}} \cdot \frac{\sum_{r=i+1}^J e^{a_{rb}} \cdot \sum_{n=0}^J e^{a_{na}}}{\sum_{n=0}^J e^{a_{na} + \max(\Delta_{i+1}, \Delta_n)}} \cdot \ln \frac{\sum_{n=0}^J e^{a_{na} + \max(\Delta_i, \Delta_n)}}{\sum_{n=0}^J e^{a_{na} + \max(\Delta_{i+1}, \Delta_n)}}.
\end{aligned}$$

Combining this expression with the earlier conditional expectation for $r = k$,

$$\begin{aligned}
\mathbf{E}_{\boldsymbol{\varepsilon}|A_{ka}} \max_{r \geq k}(a_{rb} + \varepsilon_r) &= \gamma_0 + \frac{\sum_{i=0}^J e^{a_{ia}}}{\sum_{i=0}^J e^{\max(a_{ia}, a_{ib} + a_{ka} - a_{kb})}} [a_{kb} - a_{ka} + \ln(\sum_{i=0}^J e^{\max(a_{ia}, a_{ib} + a_{ka} - a_{kb})})] \\
&\quad + \sum_{i=k}^{J-1} \frac{\sum_{r=i+1}^J e^{a_{rb}} \sum_{n=0}^J e^{a_{na}}}{\sum_{n=0}^J e^{(a_{na} + \max(\Delta_i, \Delta_n))/\sigma}} \cdot \frac{(e^{\Delta_{i+1}} - e^{\Delta_i})}{\sum_{n=0}^J e^{(a_{na} + \max(\Delta_{i+1}, \Delta_n))/\sigma}} \cdot \ln \sum_{n=0}^J e^{a_{na} + \max(\Delta_i, \Delta_n)} \\
&\quad + \sum_{i=k}^{J-1} \frac{1}{\sum_{n=0}^i e^{a_{na}}} \cdot \frac{\sum_{r=i+1}^J e^{a_{rb}} \cdot \sum_{n=0}^J e^{a_{na}}}{\sum_{n=0}^J e^{(a_{na} + \max(\Delta_{i+1}, \Delta_n))}} \cdot \ln \frac{\sum_{n=0}^J e^{a_{na} + \max(\Delta_i, \Delta_n)}}{\sum_{n=0}^J e^{a_{na} + \max(\Delta_{i+1}, \Delta_n)}}.
\end{aligned}$$

A consequence of this formula is

$$\mathbf{E} \max_{j \in J}(a_{jb} + \varepsilon_j) | A_{ka} - \mathbf{E} \max_{j \in J}(a_{ja} + \varepsilon_j) | A_{ka}$$

$$\begin{aligned}
&= \frac{\sum_{i=0}^J e^{a_{ia}}}{\sum_{i=0}^J e^{\max(a_{ia}, a_{ib} + a_{ka} - a_{kb})}} \left[a_{kb} - a_{ka} + \sigma \cdot \ln \left(\sum_{i=0}^J e^{\max(a_{ia}, a_{ib} + a_{ka} - a_{kb})} \right) \right] - \ln \sum_{n=0}^J e^{a_{na}} \\
&+ \sum_{i=k}^{J-1} \frac{\sum_{r=i+1}^J e^{a_{rb}} \sum_{n=0}^J e^{a_{na}}}{\sum_{n=0}^J e^{a_{na} + \max(\Delta_i, \Delta_n)}} \cdot \frac{(e^{\Delta_{i+1}} - e^{\Delta_i})}{\sum_{n=0}^J e^{a_{na} + \max(\Delta_{i+1}, \Delta_n)}} \cdot \ln \sum_{n=0}^J e^{a_{na} + \max(\Delta_i, \Delta_n)} \\
&+ \sum_{i=k}^{J-1} \frac{1}{\sum_{n=0}^J e^{a_{na}}} \cdot \frac{\sum_{r=i+1}^J e^{a_{rb}} \cdot \sum_{n=0}^J e^{a_{na}}}{\sum_{n=0}^J e^{a_{na} + \max(\Delta_{i+1}, \Delta_n)}} \cdot \sigma \cdot \ln \frac{\sum_{n=0}^J e^{a_{na} + \max(\Delta_i, \Delta_n)}}{\sum_{n=0}^J e^{a_{na} + \max(\Delta_{i+1}, \Delta_n)}}.
\end{aligned}$$

Appendix C. R-Code for the Discrete Welfare Calculus using a Synthetic Population

#Code to estimate losses from consumers not knowing their data were shared

```
sink('P:\\USER\\Kenneth.Train\\misperceptions\\simulation of privacy
results\\SimulationsForDan.txt')
```

```
#Attributes
```

```
# 1 price
# 2 Commercials shown between shows
# 3 Fast content availability
# 4 More TV shows
# 5 More movies
# 6 share usage data
# 7 share usage and personal data
# 8 No service
```

```
#Estimated parameters of distribution of WTP and scale
```

```
#alpha is scale parameter; 1/alpha is distributed log-normal
```

```
#WTPs are dsitributed normal
```

```
#Estimates are in Table 9 and 10 of Foundations
```

```
#Mean, stdev, and correlation matrix of underlying normals:
```

```
#Order: log(1/alpha), commercials, fast, mostly TV, mostly movies, share usage,
share personal and usage
```

```
normmmn <- c(-2.0002, -1.562, 3.945, -0.6988, 2.963, -0.6224, -2.705, -27.26)
```

```
normstd <- c(1.0637, 3.940, 3.631, 4.857, 2.524, 2.494, 6.751, 19.42)
```

```
r1 <- c(1, -0.5813, -0.1371, 0.0358, 0.0256, 0.0022, -0.1287, 0.2801)
```

```
r2 <- c(0, 1.0000, 0.1172, -0.3473, 0.0109, -0.2562, -0.0079, -0.4108)
```

```
r3 <- c(0, 0, 1.0000, 0.8042, -0.4019, -0.3542, -0.4206, 0.2391)
```

```
r4 <- c(0, 0, 0, 1.0000, -0.5890, -0.1695, -0.3328, 0.4616)
```

```
r5 <- c(0, 0, 0, 0, 1.0000, 0.5141, 0.5181, -0.0147)
```

```
r6 <- c(0, 0, 0, 0, 0, 1.0000, 0.9370, -0.0563)
```

```
r7 <- c(0, 0, 0, 0, 0, 0, 1.0000, -0.0975)
```

```
r8 <- c(0, 0, 0, 0, 0, 0, 0, 1.0000)
```

```
corrMat=rbind(r1,r2,r3,r4,r5,r6,r7,r8)
```

```
corrMat=corrMat+t(corrMat) - diag(1,8);
```

```
#Specification of services available and combinations of services.
```

```
# Look like netflix (N), Amazon Prime (A), huluplus (H), and combos eg NA
```

```
N <- c(7.99, 0, 0, 0, 1, 0, 0, 0)
```

```

A <- c(6.58, 0, 0, 1, 0, 0, 0, 0)
H <- c(7.99, 0, 1, 0, 0, 0, 0, 0)

nos <- c(0,0,0,0,0,0,0,1) #No service
#Create matrix of attributes of the 8 alternatives:
xmat <- rbind(N,A,H,N+A,N+H,A+H,N+A+H,nos) #8 alts x 8 attributes
#Indicator of which alternatives have Hulu:
hasH <- c(0,0,1,0,1,1,1,0)
ndraws <- 1000000

samplen <- c(86, 14, 12, 35, 21, 3, 22, 107) #Number of people in survey who chose
each of the 8 alternatives

mktshares <- samplen/sum(samplen)
market <- sum(samplen)*(6/58) #in million. We know Hulu has 6m customers and 58
people in the survey have Hulu

#Create draws of coefficients
set.seed(1234)
coef <- matrix(rnorm(8*ndraws),8,ndraws)
coef <- matrix(rep(normmn,times=ndraws),8,ndraws) + diag(normstd) %*%
(t(chol(corrMat)) %*% coef)
print("Check mean, std, and correlation matrix of draws against true")
print("Means: simulated and true")
print(cbind(rowMeans(coef),normmn)) #Check against normmn
print("Stds: simulated and true")
print(cbind( sqrt(diag(cov(t(coef)))), normstd)) #Check against normstd
print("Correlation matrix, simulated first, then true")
print(cor(t(coef))) #Check against corrMat
print(corrMat)
wtpsharing <- coef[7,]
pcoef <- exp(coef[1,]); #For lognormally distributed price
coef[1,]<- -pcoef; #First coef is for price
coef[2:8,] <- matrix(rep(pcoef, each=7),7,ndraws) * coef[2:8,] #Attribute coeffs are
wtp times price coef

#Calculate representative decision utility and choice probabilities
u <- xmat %*% coef
eu <- exp(u)
eu[is.infinite(eu)] <- 10^300
p <- eu / matrix(rep(colSums(eu),each=8),8,ndraws)
s <- rowMeans(p)

# Adjust constants to equal market shares

alpha <- matrix(0,8,1)
oldu <- u

for(count in 1:20){
  alpha <- alpha+log(mktshares / s)
  u <- oldu+matrix(rep(alpha,times=ndraws),8,ndraws)
  eu <- exp(u)
  eu[is.infinite(eu)] <- 10^300
  p <- eu / matrix(rep(colSums(eu),each=8),8,ndraws)
  s <- rowMeans(p)
}

```

```

print("ASCs")
print(alpha)
print("Predicted and actual market shares at ASCs")
print(cbind(s,mktshares))

#Calculate welfare impact of lack of knowlegde of sharing by hulu-like service
#Actual attributes; which includes sharing of usage and personal data by hulu
#Same as above but now 1 in column 7 for Hulu, to indicate that Hulu shares
personal and usage info:

H[7] <- 1

xmatnew <- rbind(N,A,H,N+A,N+H,A+H,N+A+H, nos) #8 alts x 8 attributes
unew <- xmatnew %*% coef
unew <- unew + matrix(rep(alpha,times=ndraws),8,ndraws)
eunew <- exp(unew);
eunew[is.infinite(eunew)] <- 10^300
pnew <- eunew / matrix(rep(colSums(eunew),each=8),8,ndraws)
newshares <- rowMeans(pnew)
diffu <- unew-u

hold <- log(colSums(eu)) / pcoef
lsdecision <- mean(hold) #expected log sum based on decision utility
hold <- log(colSums(eunew)) / pcoef
lsrealized <- mean(hold) #expected log sum based on realized utility
hold <- colSums(p * diffu) / pcoef
squareloss<- mean(hold) #expected difference between perceived and actual utility
in money metric

hulusubscribers= t(mktshares) %*% hasH

print("Difference between peoples realized utility and decision utility for chosen
alternative")
print("in money metric.")
print("Aggregate, and per-person who subscribed to Hulu")
print(cbind((squareloss * market), (squareloss / hulusubscribers)))

print("Note: Conditional mean WTP (second number above) differs from unconditional
mean of 2.70.")

print("Difference between peoples realized utility and the utility they would have
obtained if informed");
print("in money metric.");
print("Aggregate, and per-person who subscribed Hulu")
print(cbind(((lsdecision-lsrealized+squareloss) *market),((lsdecision-
lsrealized+squareloss) / hulusubscribers)))
print("Hulu share")
print("Actual choices, informed choices, percent difference")
print(cbind(hulusubscribers,(t(newshares) %*% hasH),((t(mktshares-newshares) %*%
hasH) / hulusubscribers)))

#Break down analysis further by conditioning on choice and whether person likes or
dislikes sharing info.

```

```

poswtp <- coef[7,]>=0 #These people like sharing their information; others dislike
it.

everrors <- matrix(runif(8*ndraws),8,ndraws)
everrors <- -log(-log(everrors))
util <- u+everrors
util <- util / matrix(rep(pcoef,each=8),8,ndraws) #So utils are back in money
metric
utilnew <- unew+everrors
utilnew <- utilnew / matrix(rep(pcoef,each=8),8,ndraws)

i=max.col(t(util))
inew=max.col(t(utilnew))
c=matrix(0,1,ndraws)
exper_util=matrix(0,1,ndraws)
cnew=matrix(0,1,ndraws)

for(n in 1:ndraws) {
  chosenalt <- i[n]
  c[1,n] <- util[chosenalt,n]
  exper_util[1,n] <- utilnew[chosenalt,n]
  newchosenalt <- inew[n]
  cnew[1,n] <- utilnew[newchosenalt,n]
}
Huluer <- i == 3 | i == 5 | i== 6 | i==7
NumHuluer <- sum(Huluer)
NumHuluerPosWTP <- sum(Huluer * poswtp)
NumHuluerNegWTP <- sum(Huluer *(1-poswtp))
print("Dollar difference in welfare relative what expected")
print("Aggregate, per person, per Hulu subscriber")
print("Everyone:")
xx <- exper_util-c
print(cbind(market * mean(xx), mean(xx), sum(xx)/NumHuluer))
print ("People whose dislike sharing:")
xx <-(exper_util-c) * (1-poswtp)
print(cbind(market * mean(xx), mean(xx), sum(xx)/NumHuluerNegWTP))
print("People whose like sharing:")
xx <-(exper_util-c) * poswtp
print(cbind(market * mean(xx), mean(xx), sum(xx)/NumHuluerPosWTP))

NumOther <- sum(1-Huluer)
NumOtherPosWTP <- sum(Huluer*poswtp)
NumOtherNegWTP <- sum(Huluer*(1-poswtp))
print("Dollar difference in welfare relative to being informed")
print("Aggregate, per person, average for Hulu subscribers, average for non-
subscribers")
print("Everyone:")
xx <- exper_util-cnew
print(cbind(market*mean(xx), mean(xx), mean(xx[Huluer==1]), mean(xx[Huluer==0]) ))
print("People whose dislike sharing:")
print(cbind(mean(1-poswtp)*market*mean(xx[poswtp==0]), mean(xx[poswtp==0]),
mean(xx[Huluer==1 & poswtp==0]), mean(xx[Huluer==0 & poswtp==0]) ))
print("People whose like sharing:")
print(cbind(mean(poswtp)*market*mean(xx[poswtp==1]), mean(xx[poswtp==1]),
mean(xx[Huluer==1 & poswtp==1]), mean(xx[Huluer==0 & poswtp==1]) ))

```

```

choicemat <- array(0,dim=c(8,8,2)) #Rows for actual choice, cols for choice if
informed, depth for wtp<>0
exper_diff <- array(0,dim=c(8,8,2))
welfare_diff <- array(0,dim=c(8,8,2))
for(rr in 1:8) {
  for(cc in 1:8) {
    k <- (i == rr) & (inew == cc) & (poswtp==0)
    choicemat[rr,cc,1] <- sum(k==1)
    exper_diff[rr,cc,1] <- sum(c[k==1]-exper_util[k==1])
    welfare_diff[rr,cc,1] <- sum(exper_util[k==1]-cnew[k==1])
    k <- (i == rr) & (inew == cc) & (poswtp==1)
    choicemat[rr,cc,2] <- sum(k==1)
    exper_diff[rr,cc,2] <- sum(c[k==1]-exper_util[k==1])
    welfare_diff[rr,cc,2] <- sum(exper_util[k==1]-cnew[k==1])
  }
}

print("Hulu subscriber or not")
math <- matrix(c( hasH,(1-hasH)),8,2)
subscriber1 <- t(math) %*% choicemat[, ,1] %*% math
subscriber2 <- t(math) %*% choicemat[, ,2] %*% math
print("Share of population who chose row and would have chosen col")
print((subscriber1+subscriber2)/ndraws )
print("Of people who dislike sharing, share who chose row and would have chosen
col")
print( subscriber1/sum(subscriber1) )
print("Of people who like sharing, share who chose row and would have chosen col")
print( subscriber2/sum(subscriber2) )

sink()

```