

On Bayesian Analysis and Computation for Functions with Monotonicity and Curvature Restrictions

William J. McCausland*

Département de sciences économiques, Université de Montréal,
CIREQ[†] and CIRANO[‡]
C.P. 6128, succursale Centre-ville
Montréal QC H3C 3J7
Canada

Email: william.j.mccausland@umontreal.ca

Phone: (514) 343-7281

Fax: (514) 343-7221

April 26, 2007

Abstract

Following many in the literature on inference for functions with monotonicity and curvature restrictions, we use a functional form consisting of finite linear combinations of basis functions. We consider Bayesian inference on the number of basis functions and their coefficients. Prior elicitation is difficult because of the irregular shape and infinite dimension of the parameter space. We show how to elicit priors

*We appreciate comments by John Geweke, Narayana Kocherlakota, Dale Poirier, John Stevens, Walter Thurman and three referees. We thank William Harbaugh for providing experimental data. Software for computations reported in this paper is available from the author.

[†]Centre Interuniversitaire de Recherche en Économie Quantitative

[‡]Centre Interuniversitaire de Recherche en Analyse des Organisations

that are *flexible, theoretically consistent*, and proper. We demonstrate that uniform priors over coefficients imply priors over economically relevant features of the function that many may consider too informative. We give an example of a non-uniform prior that addresses this issue. We introduce methods for prior and posterior simulation that meet challenges posed by the irregular shape of the parameter space. In an example, we analyze data from a consumer demand experiment.

Key words: Bayesian methods; Flexible functional forms; Shape restrictions

1 Introduction

There are many examples in economics where theoretically consistent choice behavior is described by multivariate functions subject to monotonicity and curvature restrictions. These functions include utility, expenditure, indirect utility, production, cost, and profit functions.

Much empirical analysis in economics involves learning about these functions using data on the choices of consumers and firms. There is a large literature on such inference. See Deaton and Muellbauer (1980), Diewert and Wales (1987), Lau (1986), Matzkin (1994) and Terrell (1996).

Analysis typically begins with two choices: a parametric class of functions, and constraints on the parameter vector. The constraints define a restricted parameter set.

The literature identifies two important objectives governing these choices, theoretical consistency and flexibility. To a large extent, they are competing. Theoretical consistency refers to the extent to which the functions indexed by elements of the restricted parameter set satisfy the applicable monotonicity and curvature restrictions over their domain. If the functions satisfy the restrictions throughout the domain, we have *global theoretical consistency*. If they satisfy them at a point or in a region, we have *local theoretical consistency* or *regional theoretical consistency*, respectively. Flexibility refers to the variety of functions indexed by elements of the restricted parameter set, and it too may be more or less global, depending on how large is the subset of the domain where the relevant flexibility properties hold.

Consider, for example, two commonly used classes of indirect utility functions: the Constant Elasticity of Substitution (CES) class, and the trans-log class. According to standard theory, indirect utility functions are non-increasing and quasi-convex in income-normalized prices. The CES class of indirect utility functions, with non-negativity restrictions on its parameters, is globally theoretically consistent. However, it is quite inflexible, in the sense that elasticities of substitution cannot vary with prices and income. The trans-log class is locally flexible in the sense that with appropriate choices of the parameters one can attain arbitrary elasticities at a given point. However, it is not globally theoretically consistent: there are values of the parameters for which the function is not everywhere on its domain non-increasing and quasi-convex. We cannot rule out these values without renouncing local flexibility.

There are at least three distinct classes of functions whose flexibility al-

allows the simultaneous approximation of a continuous function, and any continuous derivatives it may have, on a compact subset \bar{X} of its theoretical domain X . We call \bar{X} the *restricted domain* and note that it can be chosen to include the empirically relevant region. The three classes consist of linear combinations of basis functions.

The simultaneous approximation of a function and its derivatives is important for two reasons. First, it is desirable to approximate the behavior that a function represents, and theoretically consistent choices are often given in terms of the function's derivatives. Roy's identity, for example, give choices as functions of derivatives of the indirect utility function. The proximity of two functions in, for example, the sup norm does not guarantee the proximity of their derivatives: the difference of the two functions may have low amplitude but high frequency ripples. A second reason is that by simultaneously approximating derivatives, we can guarantee that the approximating function satisfies the applicable monotonicity and curvature restrictions, which we can express in terms of derivatives.

Gallant (1981) launches this literature with his Fourier flexible form. Basis functions are sinusoidal, and any continuous function on \bar{X} can be approximated arbitrarily closely in sup norm by a linear combination of a finite number of these basis functions. If the function has bounded derivatives up to some order, we can simultaneously approximate the function and these derivatives in sup norm.

Unfortunately, sinusoidal functions do not satisfy typical monotonicity and curvature restrictions and so it can take many terms to build up an approximation. In the context of approximating indirect utility functions, Gallant (1981) proposes adding a negative definite quadratic term to the expansion.

Barnett and Jonas (1983) use a multivariate Müntz-Szasz expansion to approximate a firm's unit cost function, a function of the prices p_1, \dots, p_n of n input factors. The set of basis functions is

$$\left\{ \prod_{i=1}^n p_i^{\lambda(\iota_i)} : \iota \in \mathbb{N}_0^n \right\},$$

where $\mathbb{N}_0 \equiv 0, 1, 2, \dots$ and the sequence $\lambda(k)$, $k = 1, 2, \dots$, satisfies $\sum_{k=1}^{\infty} (1/\lambda(k)) = \infty$. Barnett and Jonas (1983) take $\lambda(k) = 2^{-k}$. Barnett and Yue (1988) give conditions for various modes of convergence of the function and its derivatives. An advantage of this approach is that all basis functions satisfy the

appropriate monotonicity and curvature restrictions for unit cost functions: they are non-decreasing and concave.

In unpublished work, Geweke and Petrella also approximate a firm’s unit cost as a function of input prices p_1, \dots, p_n , but use the following set of basis functions:

$$\left\{ \prod_{i=1}^n p_i^{b\iota_i} : \iota_i \in \mathbb{N}_0^n \right\}, \quad (1)$$

where $b > 0$. The functions $\prod_{i=1}^n p_i^{b\iota_i}$ satisfying $\iota_i < b^{-1}$, $i = 1, \dots, n$, are themselves non-decreasing and concave, which is convenient for constructing approximations of non-decreasing and concave unit cost functions using a small number of terms. If the function to approximate has continuous derivatives, we can simultaneously approximate the function and these derivatives. The approximation result is based on a mathematical result on the simultaneous approximation of a function and its derivatives using polynomials due to Evard and Jafari (1994).

For the remainder of the paper, we restrict attention to a particular set of monotonicity and curvature restrictions: we define a function as *regular* if it is non-decreasing and concave. In Section 2, we will discuss the applicability of our results to other cases.

Once we have an approximation of regular functions as linear combinations of basis functions, we can do inference for functions by doing inference for the number of basis functions to include and their coefficients. We adopt a Bayesian approach, which has many advantages in this context. Inequality restrictions, which figure prominently, are more easily handled using a Bayesian, rather than frequentist analysis. Inference on parameter values, numbers of terms and even competing functional forms¹ is done in a unified approach using methods with known and desirable properties.

Bayesian predictive inference automatically takes into account *a posteriori* uncertainty about the regular function. Within a set of functions satisfying the appropriate monotonicity and curvature restrictions and compatible with observed data, individual functions may differ considerably on regions outside the sample. Averaging over these functions by taking the posterior expectation may lead to better out-of-sample prediction than using only one function, however great its “likelihood”.

However, prior elicitation is difficult. The parameter space is infinite-

¹See Gordon, (1996) on Bayesian comparison of competing functional forms.

dimensional and irregularly shaped. Parameters have no obvious direct economic interpretation. It is difficult to think about distributions over functions.

The main contributions of the paper are the following. We describe classes of basis functions that allow one to construct a wide variety of regular functions with a small number of terms. We prove a result, similar to that of Geweke and Petrella, on the simultaneous approximation of a function and its derivatives. We illustrate using simulations that for a fixed finite number of basis functions, a uniform prior on the set of coefficients associated with regular functions implies quite informative prior distributions for certain economically relevant quantities. The unsuspecting practitioner may be assigning very low probability to regions of the parameter space that are quite plausible. We give an example of a proper, non-uniform prior over coefficients that implies a more diffuse prior distribution for these economically relevant quantities. We propose methods for prior and posterior simulation that meet the challenge posed by the shape of the parameter space and made more difficult when we use the kind of non-uniform prior that we suggest in this paper.

Section 2 concerns the approximation of functions by linear combinations of a finite number of basis functions. We introduce a wide variety of sets of basis functions nesting the set of basis functions used by Geweke and Petrella (unpublished). Following them, we apply an approximation result by Evard and Jafari (1994). We show that using any of our sets of basis functions, any twice continuously differentiable regular function can be arbitrarily well approximated on a compact subset \bar{X} of its theoretical domain X by a function that is regular on \bar{X} . The approximation is a simultaneous approximation (in sup norm) of the function, its gradient and its Hessian.

In Section 3, we discuss prior distributions over regular functions. We define notions of flexibility and theoretical consistency for these prior distributions, and give sufficient conditions for both. We describe an approach to prior elicitation that emphasizes economically relevant quantities rather than the parameters themselves, which are difficult to interpret. We discuss the non-trivial issue of ensuring that the prior is proper, which is essential for model comparison using Bayes factors.

We simulate from a uniform prior distribution to illustrate what this prior implies about the distribution of economically relevant features of the regular function. Interpreting the regular function as a utility function, we consider the marginal rate of substitution between two goods at a typical

point in the consumption set. Interpreting the function as a negative indirect utility function, we consider ratios of quantities demanded at a typical point in the space of income-normalized prices. We show that the implied prior distribution of these features of the regular function is quite informative, and argue that the uniform prior assigns low probability to sets of functions that many would find plausible. We give an example of a non-uniform prior and show that it implies more diffuse prior distributions over marginal rates of substitution and ratios of quantities demanded.

In Section 4, we discuss the problem of prior and posterior simulation of the unknown coefficients. The shape of the parameter space presents special problems, and these are exacerbated when we adopt the kind of non-uniform prior we propose in Section 3. To overcome these problems we propose Markov chain Monte Carlo (MCMC) simulation methods based on three Metropolis-Hastings proposal distributions adapted to the shape of the parameter space.

In Section 5, we present an empirical application of our econometric methods. It is common in empirical applications to adopt a measurement error approach to construct a data density, whereby an error distribution accounts for discrepancies between observed choices and the optimal choices given by the regular function. Instead, we use a stochastic model for observed choices where distributions over choices are determined by the regular function itself. Theil (1974) and McCausland (2004) give very different theoretical underpinnings for the model. We analyze individual choice data from a consumer experiment described in Harbaugh et al. (2001), and find that out-of-sample predictions are better on average than any model based only on numbers of observed violations of the Generalized Axiom of Revealed Preference.

We conclude in Section 6.

2 Parametric Classes of Functions

Our objective is Bayesian learning, from choice data, about a regular function representing choice behavior. For Bayesian analysis, we must be able to express prior and posterior uncertainty about these functions in terms of probability distributions. We make use of a parametric class of functions, which permits us to express this uncertainty in terms of distributions over parameters. Functions in our parametric class are linear combinations of basis functions. Unknown parameters are the coefficients of the basis functions

and their number. The regularity of functions defines a restricted parameter set.

We defined, for definiteness, a particular kind of regularity, but we will see that this does not restrict the applicability of our results and methods as much as it may first appear.

We do not know of any results that allow us to approximate any regular function on an unbounded set such as the classical consumption set $X = \mathbb{R}_+^n$. We settle for approximation on a compact restricted domain $\bar{X} \subseteq X$, which we can always choose to include the empirically relevant region for any given application.

We also restrict functions to be twice continuously differentiable. This ensures that we can simultaneously approximate the function, its gradient and its Hessian.

We want to be able to generate a reasonable variety of regular functions using a small number of terms. To this end, we introduce a transformation ϕ on X . Rather than directly approximate a function u on X , we will approximate $u \circ \phi^{-1}$ on $\phi(\bar{X})$. Choosing a ϕ that is regular itself makes it easier to approximate regular functions, since first order monomials of ϕ are then regular.

2.1 Regularity

Most of the results of this paper pertain to non-decreasing and concave functions, which we call regular. With simple and obvious modifications, we can replace non-decreasing with non-increasing or concave with convex.

Replacing concavity with quasi-concavity (for application to utility functions) or quasi-convexity (indirect utility functions) is a little more problematic. This is because we take advantage of the fact that the set of concave functions is closed under convex combination. Quasi-concave and quasi-convex functions do not have this property.

However, we argue that imposing concavity on utility functions is innocuous for our purposes. We grant that it is restrictive: while all complete, continuous, monotone and convex preferences have quasi-concave utility representations, not all have concave utility representations: see de Finetti (1949) or the document “A Pedagogical Example of Non-concavifiable Preferences” by James Schummer, available at the author’s website. Kannai (1974) shows, however, that on any non-empty compact set, a complete, continuous, monotone and convex preference can be arbitrarily well approximated (in a sense

made precise in Lemma 2 of that article) by a complete, continuous, monotone and convex preference that has a concave utility representation. Given that we are already restricting the domain of utility functions to a compact set \bar{X} to obtain an approximation result, concavity is not a severe additional restriction.

Interpreting the indirect utility function as the representation of preferences over income-normalized prices, we can make a similar argument about imposing convexity on indirect utility functions.

2.2 Restricted Domains

The choice of a restricted domain $\bar{X} \subseteq X$ is part of the prior specification, so it should not depend on observed choices. If the decision maker is a price-taker, then prices and income are exogenous and \bar{X} can depend on observed prices and income. We choose \bar{X} to include the empirically relevant region. One natural choice for \bar{X} is a hyper-rectangle of the form $[0, \bar{x}_1] \times \dots \times [0, \bar{x}_n]$, for some $\bar{x} \in \mathbb{R}_{++}^n$. Another is a simplex of the form $\{x \in X : w'x \leq m\}$, for some $w \in \mathbb{R}_{++}^n$ and $m \in \mathbb{R}_{++}$.

Since we impose regularity on the approximating function only on \bar{X} , enlarging \bar{X} serves the goal of theoretical consistency. However, one trades off flexibility: as \bar{X} grows, the set of parameters for which the function is regular on \bar{X} shrinks. Our approximation result holds for any compact \bar{X} , however large, but we lose flexibility for any fixed number of terms.

2.3 A Parametric Functional Form

We first introduce some definitions that help simplify notation for monomials and polynomials on \mathbb{R}^n . A *multi-index* of length n is an $\iota \in \mathbb{N}_0^n$, where \mathbb{N}_0 is the set of non-negative integers. For vectors $x \in \mathbb{R}^n$ and multi-indices $\iota \in \mathbb{N}_0^n$, we define the notation $x^\iota \equiv \prod_{i=1}^n x_i^{\iota_i}$. We will call a finite subset $I \subset \mathbb{N}_0^n$ a *constellation* of multi-indices. Thus for any multi-index ι , x^ι is a monomial and for any constellation I and vector $(\lambda_\iota)_{\iota \in I}$ of coefficients, $\sum_{\iota \in I} \lambda_\iota x^\iota$ is a polynomial.

We want to be able to construct a wide variety of regular functions using a small number of terms. To do this, we introduce a transformation $\phi: X \rightarrow \mathbb{R}^n$ of the domain and choose as our basis functions the set of monomials $\{[\phi(x)]^\iota : \iota \in \mathbb{N}_0^n\}$. In effect, we replace the problem of approximating the regular function u on \bar{X} with that of approximating $u \circ \phi^{-1}$ on $\phi(\bar{X})$.

In the simulations and empirical example of this paper, we use the following choice of ϕ :

$$\phi(x) = \left(\log \left(\frac{x_1 + \xi_1}{x_1^* + \xi_1} \right), \dots, \log \left(\frac{x_n + \xi_n}{x_n^* + \xi_n} \right) \right) \quad \forall x \in X, \quad (2)$$

where $\xi \in \mathbb{R}_{++}^n$ and $x^* \in \bar{X}$ are fixed constants that the econometrician chooses in advance for computational convenience.

The vector ξ must be strictly positive, to ensure that the region $\phi(\bar{X})$ on which we approximate $u \circ \phi^{-1}$ is compact. By choosing small elements of ξ , we obtain flexibility at small scales, but we also enlarge $\phi(\bar{X})$, which adversely affects flexibility at other scales. The constant x^* establishes the origin of the transformed choice space $\phi(\bar{X})$, since $\phi(x^*) = 0$. We find in practice that we obtain better flexibility for a fixed number of terms by choosing \bar{X} , ξ and x^* such that $\phi(x^*)$ is near the centroid of $\phi(\bar{X})$.

For the particular choice of $\phi(x)$ in equation (2), the set of monomials in $\phi(x)$ has some useful properties. It includes the regular functions $\log((x_i + \xi_i)/(x_i^* + \xi_i))$, $i = 1, \dots, n$. For small values of ξ_i , linear combinations $\sum_{i=1}^n \lambda_i \log((x_i + \xi_i)/(x_i^* + \xi_i))$ of these regular functions approximate Cobb-Douglas utility functions with arbitrary expenditure shares $\lambda_i / \sum_{j=1}^n \lambda_j$. When we add higher order monomials in $\phi(x)$ to the utility function, we allow the shares to change with prices and income.

Geweke and Petrella's basis functions, given in equation (1), can be interpreted as monomials in ϕ for the following alternative choice of ϕ :

$$\phi(x) = (x_1^b, \dots, x_n^b) \quad \forall x \in X.$$

We want to construct a sequence of families of polynomials in $\phi(x)$ with growing flexibility. We choose a sequence $\{I_k\}_{k=0}^\infty$ of constellations increasing² towards $\mathbb{N}_0^n \setminus \{0\}$. The set of polynomials $\sum_{\iota \in I_k} \lambda_\iota [\phi(x)]^\iota$ becomes more flexible as k increases. Excluding the multi-index 0 means excluding the constant term of $u(\cdot; \lambda)$. We do this to simplify prior elicitation and simulation. The exclusion means that we can only approximate a regular function u up to an additive constant, but this is innocuous.

One natural choice of I_k is the sequence of hyper-rectangular lattices $\{\iota \in \mathbb{N}_0^n : \iota_i \leq k - 1 \text{ for all } i\}$, which we use in this paper. Another is the sequence of simplicial lattices $\{\iota \in \mathbb{N}_0^n : \sum_{i=1}^n \iota_i \leq k - 1\}$.

²A sequence of sets $\{I_k\}_{k=0}^\infty$ increases towards set I , denoted $I_k \uparrow I$, if $I_k \subseteq I_{k+1}$ for all $k \in \mathbb{N}_0$ and $\cup_{k=0}^\infty I_k = I$

For given constellation index $k \in \mathbb{N}_0$, coefficient vector $\lambda \equiv (\lambda_\iota)_{\iota \in I_k}$, and transformation $\phi: \bar{X} \rightarrow \mathbb{R}$, we define the function $u(\cdot, \lambda): X \rightarrow \mathbb{R}$ by

$$u(x; \lambda) \equiv \sum_{\iota \in I_k} \lambda_\iota [\phi(x)]^\iota \quad \forall x \in X.$$

The unknown parameters of our approximation are the constellation index k and the vector $\lambda \equiv (\lambda_\iota)_{\iota \in I_k}$ of monomial coefficients. For each $k \in \mathbb{N}_0$, we define $\Lambda_{\bar{X}}^k$ as the set of vectors λ of length³ $|I_k|$ associated with functions that are regular on \bar{X} :

$$\Lambda_{\bar{X}}^k \equiv \{\lambda \in \mathbb{R}^{|I_k|} : u(\cdot; \lambda) \text{ is regular on } \bar{X}\}, \quad (3)$$

We also define $\Lambda_{\bar{X}} \equiv \cup_{k=1}^{\infty} \Lambda_{\bar{X}}^k$, the complete set of vectors λ associated with functions that are regular on \bar{X} .

2.4 Results

We now present two results on the parametric functional form. We will see in Section 3 that the following result is relevant for prior elicitation. In Section 4, we show that it is important for prior and posterior simulation.

Result 2.1 *For every $k \in \{1, 2, \dots\}$, $\Lambda_{\bar{X}}^k$ is a convex cone.*

Proof. The set of functions regular on \bar{X} is closed under addition and positive scalar multiplication. Therefore $\Lambda_{\bar{X}}^k$ is a convex cone. \square

The following alternate proof yields some important insights that are relevant for later sections. We can express the regularity conditions as follows. For all $x \in \bar{X}$,

$$\frac{\partial u(x; \lambda)}{\partial x} = \sum_{\iota \in I_k} \lambda_\iota \frac{\partial [\phi(x)]^\iota}{\partial x} \geq 0, \quad (4)$$

and for all $x \in X$ and $v \in \mathbb{R}^n$,

$$v' \frac{\partial^2 u(x; \lambda)}{\partial x \partial x'} v = \sum_{\iota \in I_k} \lambda_\iota v' \frac{\partial^2 [\phi(x)]^\iota}{\partial x \partial x'} v \leq 0. \quad (5)$$

For every choice of $x \in \bar{X}$, equation (4) gives n linear inequalities in λ , one for each component of the gradient. For every choice of $x \in \bar{X}$ and

³For any set A , we denote by $|A|$ the cardinality of A .

$v \in \mathbb{R}^n$, equation (5) gives another linear inequality in λ . We see, therefore, that the parameter space $\Lambda_{\bar{X}}^k$ is the intersection of half spaces defined by the inequalities above. The half spaces are convex, so their intersection is as well.

The following approximation result tells us that we can simultaneously approximate any twice continuously differentiable function u , regular on \bar{X} , together with its gradient and Hessian, arbitrarily closely on \bar{X} , up to an additive constant. The approximation of the gradient is important because we want to approximate the behavior that the function represents, which is often given in terms of derivatives. The approximation of the gradient and Hessian is important for guaranteeing that the approximating function is regular on \bar{X} . In practice, it is much easier to check that the approximating function is regular on \bar{X} than to verify its proximity to some function regular on \bar{X} . For notational convenience, we define the following norm for twice continuously differentiable functions:

$$\|f\| \equiv \max \left[\sup_{x \in \bar{X}} |f(x)|, \sup_{x \in \bar{X}, i \in \{1, \dots, n\}} \left| \frac{\partial f(x)}{\partial x_i} \right|, \sup_{x \in \bar{X}, i, j \in \{1, \dots, n\}} \left| \frac{\partial f(x)}{\partial x_i \partial x_j} \right| \right].$$

The result is a generalization of an unpublished result by Geweke and Petrella. It is they who recognized the significance of a result by Evard and Jafari (1994) on the simultaneous approximation of a function and its derivatives by polynomials. The modest contributions in this section include the recognition that different transformations ϕ can be used to generate sets of basis functions and a complete proof of the following result.

Result 2.2 (Approximation) *Suppose the transformation $\phi: \bar{X} \rightarrow \mathbb{R}$ is such that $\phi(\bar{X})$ is compact and that the inverse ϕ^{-1} on $\phi(\bar{X})$ exists and is twice continuously differentiable. Then for every twice continuously differentiable $u: X \rightarrow \mathbb{R}$ regular on \bar{X} , and every $\epsilon > 0$, there exists a $\lambda \in \Lambda_{\bar{X}}$ and a constant c such that*

$$\|c + u(\cdot; \lambda) - u(\cdot)\| < \epsilon. \tag{6}$$

Proof. Let u be a twice continuously differentiable regular function and let $\epsilon > 0$.

Rather than approximating u directly, which may be on the boundary of the regular region, we will approximate a nearby function $\hat{u}: \bar{X} \rightarrow \mathbb{R}$ in the

interior. We choose \hat{u} close enough to u that the approximation of \hat{u} is a sufficiently close approximation of u itself. The function \hat{u} is defined by

$$\hat{u}(x) \equiv u(x) + \frac{\epsilon}{2} \prod_{i=1}^n (1 - e^{\bar{x}_i - x_i}) \quad \forall x \in \bar{X},$$

where $\bar{x}_i = \max_{x \in \bar{X}} x_i$, $i = 1, \dots, n$. Since u is non-decreasing, concave and twice continuously differentiable on \bar{X} , \hat{u} is increasing, strictly concave and twice continuously differentiable on \bar{X} . Also,

$$\|\hat{u}(\cdot) - u(\cdot)\| \leq \frac{\epsilon}{2}. \quad (7)$$

A direct corollary of Corollary 3 of Evard and Jafari (1994) is that for every twice continuously differentiable function $f: \bar{X} \rightarrow \mathbb{R}$, and every $\epsilon' > 0$, there exists a polynomial $p: \bar{X} \rightarrow \mathbb{R}$ such that for all $i, j \in \{1, \dots, n\}$ and all $x \in \bar{X}$,

$$|f(x) - p(x)| < \epsilon', \quad \left| \frac{\partial f}{\partial x_i} - \frac{\partial p}{\partial x_i} \right| < \epsilon', \quad \text{and} \quad \left| \frac{\partial^2 f}{\partial x_i \partial x_j} - \frac{\partial^2 p}{\partial x_i \partial x_j} \right| < \epsilon'.$$

Since ϕ has an inverse ϕ^{-1} that is twice continuously differentiable on \bar{X} , $\hat{u} \circ \phi^{-1}$ is also twice continuously differentiable on $\phi(\bar{X})$. Furthermore, $\phi(\bar{X})$ is compact. Therefore the corollary implies that for all $\epsilon' > 0$, there exists a $\lambda \in \bigcup_{k=1}^{\infty} \mathbb{R}^{|I_k|}$ and a c such that for all $i, j \in \{1, \dots, n\}$ and all $z \in \phi(\bar{X})$,

$$\left| c + \sum_{\iota \in I_k} \lambda_{\iota} z^{\iota} - (\hat{u} \circ \phi^{-1})(z) \right| < \epsilon',$$

$$\left| \frac{\partial}{\partial z_i} \sum_{\iota \in I_k} \lambda_{\iota} z^{\iota} - \frac{\partial}{\partial z_i} (\hat{u} \circ \phi^{-1})(z) \right| < \epsilon'$$

and

$$\left| \frac{\partial^2}{\partial z_i \partial z_j} \sum_{\iota \in I_k} \lambda_{\iota} z^{\iota} - \frac{\partial^2}{\partial z_i \partial z_j} (\hat{u} \circ \phi^{-1})(z) \right| < \epsilon'.$$

The function ϕ maps \bar{X} to $\phi(\bar{X})$, and therefore for all $i, j \in \{1, \dots, n\}$ and $x \in \bar{X}$,

$$\left| c + \sum_{\iota \in I_k} \lambda_{\iota} [\phi(x)]^{\iota} - \hat{u}(x) \right| = |c + u(x; \lambda) - \hat{u}(x)| < \epsilon',$$

$$\left| \frac{\partial}{\partial x_i} u(x; \lambda) - \frac{\partial}{\partial x_i} \hat{u}(x) \right| < \epsilon' M_1$$

and

$$\left| \frac{\partial^2}{\partial x_i \partial x_j} u(x; \lambda) - \frac{\partial^2}{\partial x_i \partial x_j} \hat{u}(x) \right| < \epsilon' M_2,$$

where M_1 and M_2 , derived from uniform bounds on the derivatives of ϕ on \bar{X} , do not depend on x .

We can choose ϵ' such that

$$\|c + u(x; \lambda) - \hat{u}(x)\| < \frac{\epsilon}{2}, \quad (8)$$

and for all $x \in \bar{X}$ and all $i, j \in \{1, \dots, n\}$,

$$\frac{\partial u(x; \lambda)}{\partial x} \geq 0, \quad (9)$$

and

$$v' \frac{\partial^2 u(x; \lambda)}{\partial x \partial x'} v \leq 0 \quad \forall v \in \mathbb{R}^n. \quad (10)$$

Equations (7), (8), and the triangle inequality guarantee that (6) holds. Equations (9) and (10) guarantee that $\lambda \in \Lambda_{\bar{X}}$. \square

3 Priors

We have just described a parametric class of regular functions. In this section, we discuss the problem of expressing prior uncertainty about regular functions by specifying a prior probability distribution for the constellation index k and the coefficient vector $\lambda = (\lambda_\iota)_{\iota \in I_k}$. The ideas expressed in this section apply to various choices of the transformation ϕ and even to different systems of basis functions, such as those of the Fourier flexible form and the Müntz-Szasz expansion. For definiteness, we take $\bar{X} = [0, \bar{x}_1] \times \dots \times [0, \bar{x}_n]$ and ϕ defined in (2).

We first define notions of flexibility and theoretical consistency for these prior distributions and give sufficient conditions for both. We then focus on the problem of eliciting prior distributions over $\Lambda_{\bar{X}}^k$ for fixed k . We address the issue of ensuring that these are proper, which is important for inference on k and model comparison using Bayes factors.

We show that if the regular function is a utility function, a uniform prior implies prior distributions over marginal rates of substitution that many would consider too informative. When the regular function is a negative indirect utility function, the prior distributions of ratios of quantities demanded is very informative. We give an example of a non-uniform prior that implies more diffuse prior distributions over marginal rates of substitution or ratios of quantities.

3.1 Flexibility and Theoretical Consistency

We will say that a prior is *flexible* if for any twice continuously differentiable function u that is regular on \bar{X} , there is a constant c such that the prior assigns positive probability to any $\|\cdot\|$ -neighborhood of $u - c$. We will say that it is *theoretically consistent* if it assigns zero probability to the set of functions that are not regular on \bar{X} . The following result gives sufficient conditions on the prior for flexibility and theoretical consistency.

Result 3.1 *If $\sum_{k=K}^{\infty} \pi_k > 0$ for all $K \in \mathbb{N}$, and the conditional distributions $\lambda|k$ have support $\Lambda_{\bar{X}}^k$, then the prior is flexible and theoretically consistent.*

Proof. Let $u: X \rightarrow \mathbb{R}$ be twice continuously differentiable and regular on \bar{X} , and let $\epsilon > 0$. By Result 2.2, we can find a $k \in \mathbb{N}$ and a $\lambda^* \in \Lambda_{\bar{X}}^k$ such that $\|c + u(x; \lambda^*) - u(x)\| < \epsilon/2$. The conditional distribution $\lambda|k$, whose support is $\Lambda_{\bar{X}}^k$, assigns positive probability to the set

$$\Lambda^* \equiv \left\{ \lambda \in \Lambda_{\bar{X}}^k : |\lambda_{\iota}^* - \lambda_{\iota}| < \frac{\epsilon}{2|I_k| \|\phi(x)^{\iota}\|} \quad \forall \iota \in I_k \right\}.$$

For all $\lambda \in \Lambda^*$, $\|u(\cdot; \lambda) - u(\cdot; \lambda^*)\| < \frac{\epsilon}{2}$, and therefore $\|c + u(\cdot; \lambda) - u(\cdot)\| < \epsilon$. Since $\sum_{k=K}^{\infty} \pi_k > 0$ for all $K \in \mathbb{N}$, the prior assigns positive probability to Λ^* , and therefore positive probability to the $\|\cdot\|$ -neighborhood of u . The fact that the prior assigns zero probability to the set of functions that are not regular on \bar{X} follows trivially from the fact that the support of the prior is $\Lambda_{\bar{X}}$. \square

3.2 Conditional Prior Distributions on $\Lambda_{\bar{X}}^k$

The conditional priors $\lambda|k$ are distributions on the $\Lambda_{\bar{X}}^k$. Except for being convex cones, the $\Lambda_{\bar{X}}^k$ are irregularly shaped, and this makes elicitation difficult.

Utility functions that differ only up to the multiplication of a positive constant are observationally equivalent. The same holds for negative indirect utility functions. If the regular function is a utility function or a negative indirect utility functions, then without loss of generality, we can normalize it by imposing the condition $u(\bar{x}; \lambda) - u(0; \lambda) = 1$. This is a linear equality in λ , so it defines a hyperplane in $\mathbb{R}^{|I_k|}$. By Result A.1, the intersection of this hyperplane with $\Lambda_{\bar{X}}^k$ is bounded. This implies that a uniform prior on this intersection is proper.

This is not the only possible normalization. For illustrative purposes we will consider the following normalization:

$$\sum_{i=1}^n \beta_i \frac{\Delta u_i(\lambda)}{\bar{x}_i} = 1,$$

where $\Delta u_i(\lambda)$ is the increase in u as x_i goes from the lower extreme of 0 to the upper extreme of \bar{x}_i , with all other x_j set to the reference value x_j^* . That is,

$$\Delta u_i(\lambda) \equiv u(x^* + (\bar{x}_i - x_i^*)e_i; \lambda) - u(x^* + (0 - x_i^*)e_i; \lambda),$$

where e_i is the unit vector on the i 'th coordinate axis in \mathbb{R}^n . We can think of $\Delta u_i/\bar{x}_i$ as the average of the gradient component $\partial u/\partial x_i$ on the line segment from $x^* + (0 - x_i^*)e_i$ to $x^* + (\bar{x}_i - x_i^*)e_i$.

This normalization is also a linear equality in λ , and using Result A.1, we can show that the hyperplane it defines has a bounded intersection with $\Lambda_{\bar{X}}^k$. The uniform distribution on this intersection is therefore proper.

Using the methods described in Section 4, we simulate from this uniform distribution. We take $n = 2$, $\xi = (0.1, 0.1)$, $x^* = (1.0, 1.0)$, $\bar{x} = (12.0, 12.0)$, and $\beta = (0.5, 0.5)$.

The first panel of Figure 1 shows a histogram of the prior distribution of $\log(\partial u(x^*; \lambda)/\partial x_1) - \log(\partial u(x^*; \lambda)/\partial x_2)$. If we interpret u as a utility function, this quantity is the log marginal rate of substitution between goods 1 and 2 at the commodity bundle x^* . If u is a negative indirect utility function, it is the log ratio of the quantities demanded of goods 1 and 2 at income-normalized prices x^* .

We see that despite being uniform on the set

$$\{\lambda \in \Lambda_{\bar{X}}^k : \sum_{i=1}^n \beta_i \frac{\Delta u_i(\lambda)}{\bar{x}_i} = 1\},$$

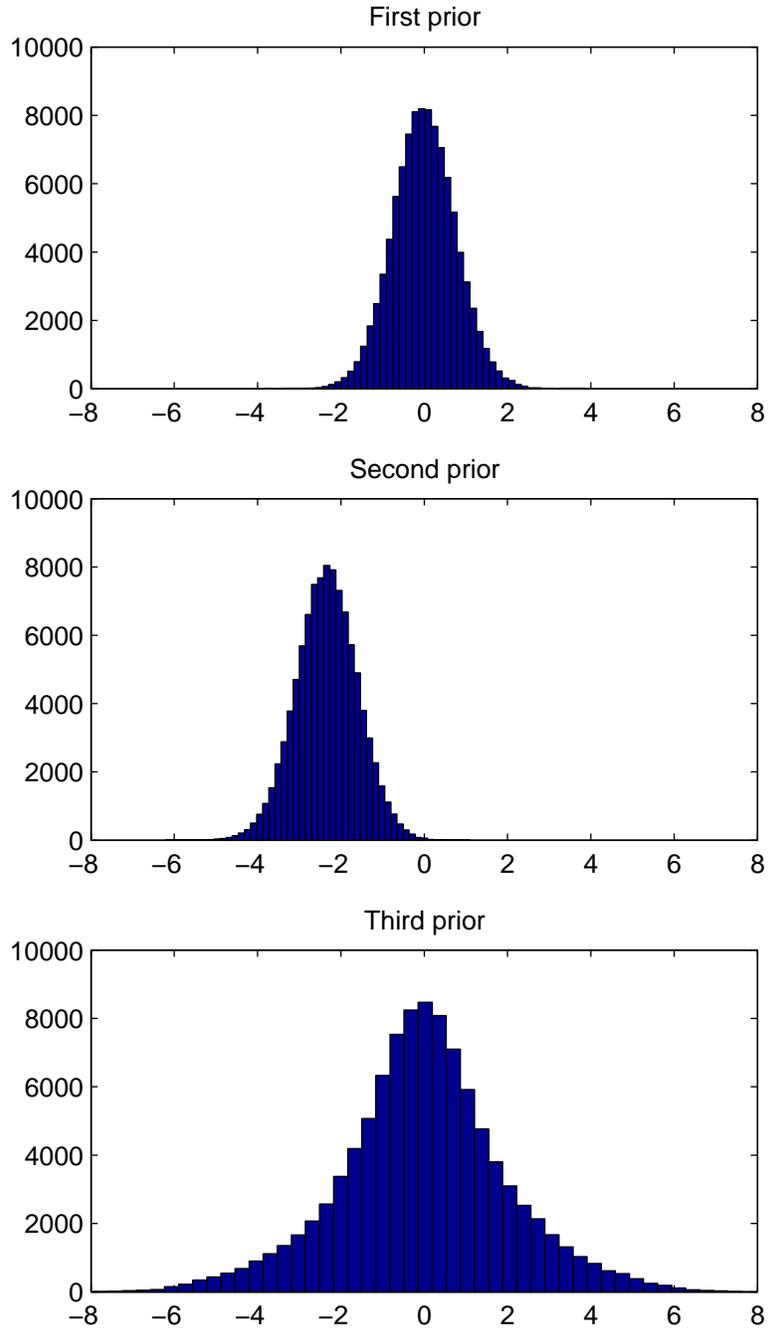


Figure 1: Histograms of $\log(\partial u(x^*; \lambda)/\partial x_1) - \log(\partial u(x^*; \lambda)/\partial x_2)$ for three different priors.

the prior is quite informative about the value of $\log(\partial u(x^*; \lambda)/\partial x_1) - \log(\partial u(x^*; \lambda)/\partial x_2)$. Specifically, it puts very low probability on the regions where $\partial u(x^*; \lambda)/\partial x_1$ is much smaller or much greater than $\partial u(x^*; \lambda)/\partial x_2$. Many will find these regions quite plausible for certain applications and will thus find the prior inappropriate.

We would like to make clear that the point we are making concerns scale, not location. We can always change the units of measurement of commodity quantities (and adjust prices accordingly) to change the location of this distribution.

To persuade the skeptical reader that the prior is not diffuse enough for some applications, we consider the following prior elicitation exercise, concerning individual consumption of two goods: gasoline and a composite good representing everything else. Hypothetical individual A earns \$30,000 after tax, takes the subway to work and buys 5 gallons of gasoline per month. B earns the same amount, drives to work 10 miles from home and spends 50 gallons of gasoline a month. The two individuals are hardly unusual, but the ratio of quantities differs by a factor of 10 between them, even though they have the same income and face the same prices. A factor of ten corresponds to slightly more than three standard deviations. It is quite plausible that if the price of gasoline doubled, A would not reduce consumption to less than 15 gallons a month. It is also plausible that if the price of gasoline dropped by a factor of 5, B would not increase consumption to more than 15 gallons a month. This would imply that their marginal rates of substitution at 15 gallons per month differ by at least a factor of ten.

Given the uniformity of the prior, the assignment of low probability to the regions where one gradient component is much smaller than the other must be because these regions have very low volume. Since these regions are also near the boundary of the set, we will refer to them (loosely) as the “tight corners” of $\Lambda_{\bar{X}}^k$. Intuitively, as one utility gradient component approaches zero it becomes more and more difficult to choose λ to satisfy monotonicity and concavity.

We illustrate using an example that the distribution of $\log(\partial u(x; \lambda)/\partial x_1) - \log(\partial u(x; \lambda)/\partial x_2)$ is sensitive to the choice of normalization. We tilt the hyperplane defined by the normalization $\beta_1 \Delta u_1 / \bar{x}_i + \beta_2 \Delta u_2 / \bar{x}_i = 1$ by replacing $\beta = (1/2, 1/2)$ with $\beta = (1/6, 5/6)$. The second panel of Figure 1 shows the prior histogram for $\log(\partial u(x^*; \lambda)/\partial x_1) - \log(\partial u(x^*; \lambda)/\partial x_2)$ implied by a uniform prior on the intersection of $\Lambda_{\bar{X}}^k$ with the new hyperplane. We see that the location of the distribution changes, but little else. In particular,

the prior is still very informative.

We conjecture that the tight distribution of $\log(\partial u(x^*; \lambda)/\partial x_1) - \log(\partial u(x^*; \lambda)/\partial x_2)$ arises not because of the particular choice of normalization but because the prior is uniform on a truncated hyperplane: such a prior puts low probability in the vicinity of any given bounding hyper-plane of the form $v'\lambda = 0$, such as $\partial u(x^*; \lambda)/\partial x_i = 0$. Simulations we have done suggest that as k increases and $u(\cdot; \lambda)$ becomes more flexible, the prior puts less and less probability in the vicinity of the bounding hyper-plane defined by $\partial u(x^*; \lambda)/\partial x_i = 0$. Intuitively, as k increases and the dimension of $\Lambda_{\bar{X}}^k$ increases with it, the curvature of its boundary increases as well, reducing the relative volume close to the hyper-plane.

If we want to make the prior distribution of $\log(\partial u(x^*; \lambda)/\partial x_1) - \log(\partial u(x^*; \lambda)/\partial x_2)$ more diffuse, we can do so by putting more probability mass near the boundary of $\Lambda_{\bar{X}}^k$. As an example, we multiply the prior by the factor

$$\prod_{i=1}^n \left(\beta_i \frac{\Delta u_i}{\bar{x}_i} + \delta \right)^{-p},$$

where δ and p are positive constants. The factor is bounded on the truncated hyper-plane $\{\lambda \in \Lambda_{\bar{X}}^k : \beta_1 \Delta u_1(\lambda)/\bar{x}_1 + \beta_2 \Delta u_2(\lambda)/\bar{x}_2 = 1\}$, so the prior remains proper.

The third panel of Figure 1 shows the prior histogram of $\log(\partial u(x^*; \lambda)/\partial x_1) - \log(\partial u(x^*; \lambda)/\partial x_2)$ for the modified prior, with $\delta = 0.001$ and $p = -4$. We see that it more diffuse.

The three examples we have just seen suggest that we can center the prior over whatever value of $\log(\partial u(x^*; \lambda)/\partial x_1) - \log(\partial u(x^*; \lambda)/\partial x_2)$ is appropriate for a given application and also independently choose the degree of dispersion.

The priors in these three examples have support equal to the intersection of $\Lambda_{\bar{X}}^k$ with a hyperplane defined by $v'\lambda$, where v is a vector defining a normalization of the regular function u . While utility and indirect utility functions can be normalized without affecting optimal choices, expenditure, cost and profit functions cannot. For flexibility in these cases, we would like to elicit a prior with full support $\Lambda_{\bar{X}}^k$. Rather than normalizing the regular function, we can choose a density f with the property that for all $a > 0$, $f(a\lambda) = a^{-|I_k|-1} f(\lambda)$. This makes the scale $v'\lambda$ and the normalized regular function $\frac{1}{v'\lambda} u(\cdot; \lambda)$ statistically independent.

If we want to elicit a prior whose contours are not all parallel to the hyperplane $v'\lambda = 1$, but has the property that the scale and the normalized

regular function are independent, we can choose a prior of the form

$$f(\lambda) \propto g(v'\lambda)(v'\lambda)^{-(|I_k|-1)}h(\lambda), \quad (11)$$

where h is a homogeneous degree 0 function.

For practical reasons, it turns out to be useful to elicit a prior with full support $\Lambda_{\bar{X}}^k$ even when the regular function is a utility or negative indirect utility function. And we can always ignore the unidentified scale. This means, for example, that we can simulate the uniform distribution on the intersection of $\Lambda_{\bar{X}}^k$ and the hyperplane defined by $v'\lambda = 1$ by simulating the density

$$f(\lambda) \propto g(v'\lambda)(v'\lambda)^{-(|I_k|-1)},$$

where g is a proper density function, and normalizing draws by dividing by $v'\lambda$. The distribution of the scale $v'\lambda$ will have density given by g .

To simulate the third, non-uniform, prior we can simulate the density in equation (11) with g arbitrary and

$$h(\lambda) \equiv \prod_{i=1}^n \left(\frac{\beta_i \Delta u_i / \bar{x}_i}{\sum_{j=1}^n \beta_j \Delta u_j / \bar{x}_i} + \delta \right)^{-p}, \quad (12)$$

and normalize by dividing by $v'\lambda$.

4 Prior and Posterior Simulation

In the previous section, we describe conditional prior densities $f(\lambda|k)$ on the convex cones $\Lambda_{\bar{X}}^k$ whose union is the regular parameter space. In this section we describe methods for Markov chain Monte Carlo (MCMC) simulation of distributions on $\Lambda_{\bar{X}}^k$.

We saw in the previous section that the regions near the boundary of the parameter space $\Lambda_{\bar{X}}^k$ where one or more of the gradient components is much smaller than the rest have very low volume. Thus the problems of simulating a distribution on this space are similar to those of simulating a distribution on a cigar-shaped object with long pointy ends. The problem is exacerbated when we want to simulate a distribution, such as the third prior of the previous section, that concentrates probability density in the tight corners of the parameter space. An efficient chain for simulating the target distribution must be able to spend enough time in the tight corners, but also move quickly in and out of them.

Random walk Metropolis chains with constant proposal variance matrices work poorly. If the eigenvalues of the variance are large enough to generate steps big enough for efficiently exploring the central region of $\Lambda_{\bar{X}}^k$, then the acceptance probability in the tight corners is intolerably low. We experimented with Metropolis-Hastings random walks where the proposal variance varied according to how close the current state of the chain was to various bounding planes. This is quite difficult to do, and the chain's numerical efficiency is fairly low. The problem is that it takes too many iterations to move in and out of the tight corners.

The Metropolis-Hastings proposal distributions we present here do not depend on the target distribution. So we make no distinction between prior and posterior simulation and leave open the question of how the data density is specified. For the simulations reported in this paper, we obtain satisfactory numerical efficiency for both prior and posterior simulation. We believe there is scope for improving these proposals by taking advantage of features of the prior and likelihood, but we do not explore this here.

We describe three Metropolis-Hastings updates that can be used in combination (either a mixture or a sweep) to simulate prior and posterior distributions efficiently on the parameter subspace $\Lambda_{\bar{X}}^k$.

The first update generates a line passing through the current state λ in a random direction, then draws a random proposal λ^* on a segment of this line containing its intersection with $\Lambda_{\bar{X}}^k$. This means that large jumps across $\Lambda_{\bar{X}}^k$ have reasonably high probability.

The second and third updates draw a random proposal λ^* on a ray emanating from a point near the boundary of $\Lambda_{\bar{X}}^k$ and passing through the current state λ . It is these two updates which ensure that the chain can quickly move in and out of the tight corners of $\Lambda_{\bar{X}}^k$.

4.1 Definitions

We introduce a few preliminary definitions that will be important in the next sections. We define, for each good i , the univariate function $u_i: [0, \bar{x}_i] \rightarrow \mathbb{R}$ as the restriction of $u(x; \lambda)$ to the line segment $\{x \in \bar{X} : x_j = x_j^*, j \neq i\}$, parallel to the i 'th coordinate axis but shifted so that the $x_j, j \neq i$, are fixed at x_j^* rather than zero. That is,

$$u_i(x_i) = u(x^* + (x_i - x_i^*)e_i), \quad x_i \in [0, \bar{x}_i],$$

where e_i is the unit length n -vector on the i 'th co-ordinate axis.

Note that

$$\phi(x^* + (x_i - x_i^*)e_i) = \log \left(\frac{x_i + \xi_i}{x_i^* + \xi_i} \right) e_i.$$

The only non-zero element of this vector is the i 'th, which means that the monomial $[\phi(x^* + (x_i - x_i^*)e_i)]^\iota$ is non-zero only for $\iota = \kappa e_i$, $\kappa = 1, \dots, k-1$. Thus the univariate u_i depends only the elements λ_ι such that $\iota = \kappa e_i$, $\kappa = 1, \dots, k-1$.

This leads to the following partition of the vector λ :

$$\lambda = (\lambda^{(1)}, \dots, \lambda^{(n)}, \lambda^\circ),$$

where $\lambda^{(i)}$ consists of the elements λ_ι for which $\iota_j = 0$ for all $j \neq i$. In other terms,

$$\lambda^{(i)} \equiv (\lambda_{e_i}, \lambda_{2e_i}, \dots, \lambda_{(k-1)e_i}), \quad i = 1, \dots, n.$$

The subvector λ° consists of all remaining elements of λ .

Since $\phi(x^* + (x_i - x_i^*)e_i)$ is non-zero only for $\iota = \kappa e_i$, $\kappa = 1, \dots, k-1$, u_i depends only on the sub-vector $\lambda^{(i)}$.

4.2 A Convex Cone Enclosing $\Lambda_{\bar{X}}^k$

We draw all three random proposals within a convex cone $\{\lambda \in \mathbb{R}^{|I_k|} : V\lambda \geq 0\}$, where V is an $N_v \times |I_k|$ matrix such that the cone contains $\Lambda_{\bar{X}}^k$. The tighter the fit of $V\lambda \geq 0$ to $\Lambda_{\bar{X}}^k$, the less often we draw proposals that are not in the regular set $\Lambda_{\bar{X}}^k$.

It is easy to construct matrices V such that $V\lambda \geq 0$ contains $\Lambda_{\bar{X}}^k$. Equations (4) and (5), which give various necessary conditions for regularity, are inequalities of the form $v\lambda \geq 0$, where v is a row vector of length $|I_k|$. Constructing V involves vertically stacking N_v row vectors v_j satisfying $v_j\lambda \geq 0$.

For the prior and posterior simulations reported in this paper, we use $N_v = n \cdot J^n$ vectors, where $J = 20$. Each vector is indexed by a pair $(i, \iota) \in \{1, \dots, n\} \times I_J$. For each i and ι , we generate a row of V using the necessary conditions

$$\begin{cases} \frac{\partial u(x^{(\iota)}; \lambda)}{\partial x_i} - \frac{\partial u(x^{(\iota+e_i)}; \lambda)}{\partial x_i} \geq 0 & \iota_i \leq J \\ \frac{\partial u(x^{(\iota)}; \lambda)}{\partial x_i} \geq 0 & \iota_i = J, \end{cases}$$

where $\{x^{(\iota)} : \iota \in I_J\}$ is a grid of points. Each $x^{(\iota)}$ is defined by

$$\phi(x^{(\iota)}) = \left(z_1^{lo} + \frac{\iota_1}{J-1}(z_1^{hi} - z_1^{lo}), \dots, z_n^{lo} + \frac{\iota_n}{J-1}(z_n^{hi} - z_n^{lo}) \right),$$

where $z^{lo} \equiv \phi(0)$ and $z^{hi} \equiv \phi(\bar{x})$.

4.3 A First Proposal

We start at the current state λ and generate a random proposal λ^* .

The proposal consists of three steps. First, we draw a random direction vector $w \in \mathbb{R}^{|I_k|}$ from a discrete uniform distribution over a set of precomputed direction vectors. Then we compute $\pi_- \leq 0$ and $\pi_+ \geq 0$ such that $\lambda + (\pi_-)w$ and $\lambda + (\pi_+)w$ are on the boundary of the cone $V\lambda \geq 0$ enclosing $\Lambda_{\bar{X}}^k$. Finally, we draw π from the uniform distribution on $[\pi_-, \pi_+]$ and construct the proposal $\lambda^* = \lambda + \pi w$. We accept λ^* with probability

$$\min \left(1, \frac{f(\lambda^*)}{f(\lambda)} \right),$$

where f is the unnormalized target density. Evaluating $f(\lambda^*)$ includes determining whether $\lambda^* \in \Lambda_{\bar{X}}^k$. We now describe some of these steps in more detail.

4.3.1 Computing π_- and π_+

We compute π_- as the smallest positive value of π for which $V(\lambda + \pi w) \geq 0$. That is, we choose

$$\pi_- = \max_{j \in \{1, \dots, N_v\}} \left\{ \frac{(V\lambda)_j}{(Vw)_j} : \frac{(V\lambda)_j}{(Vw)_j} \leq 0 \right\}.$$

Similarly,

$$\pi_+ = \min_{j \in \{1, \dots, N_v\}} \left\{ \frac{(V\lambda)_j}{(Vw)_j} : \frac{(V\lambda)_j}{(Vw)_j} \geq 0 \right\}.$$

4.3.2 Drawing Random Directions

For the simulations reported in this paper, we use the following set of directions. There is a direction $w^{(\iota)} \in \mathbb{R}^{|I_k|}$ for each multi-index $\iota \in I_k \setminus \{e_1, \dots, e_n\}$. For multi-indices ι and κ , $w_{\kappa}^{(\iota)}$ is element κ of the direction vector $w^{(\iota)}$ and is given by the coefficient of z^{κ} in the n -variate polynomial $\prod_{i=1}^n P_{\iota_i}(a_i z_i + b_i)$, where P_j is the j 'th Legendre polynomial, and $a_i z_i + b_i$ is the linear transformation of $z_i \equiv \phi_i(x)$ mapping $[z_i^{lo}, z_i^{hi}]$ to $[-1, 1]$. As before, $z^{lo} \equiv \phi(0)$ and

$z^{hi} \equiv \phi(\bar{x})$. In other terms,

$$w_{\kappa}^{(\iota)} = \prod_{i=1}^n \sum_{j=\kappa_i}^{\iota_i} L_{\iota_i, j} \binom{j}{\kappa_i} a_i^{\kappa_i} b_i^{j-\kappa_i},$$

where $a_i = 2/(z_i^{hi} - z_i^{lo})$, $b_i = -(z_i^{hi} + z_i^{lo})/(z_i^{hi} - z_i^{lo})$, and L_{ij} is the coefficient of the j 'th order monomial in the i 'th Legendre polynomial.

This set of directions has two useful properties. First, the n -variate polynomials $\sum_{\kappa \in I_k} w_{\kappa}^{(\iota)} z^{\kappa}$ are orthogonal on $\phi(\bar{X})$, inheriting the orthogonality of the (univariate) Legendre polynomials on $[-1, 1]$. This minimizes redundancy among the directions. Second, because the coordinate vectors e_i are excluded, $Vw_{\kappa}^{(\iota)}$ always has both positive and negative elements, which ensures that the values π_- and π_+ exist and are always finite.

4.3.3 Checking Regularity

Evaluating the target density, if its support is $\Lambda_{\bar{X}}^k$, typically involves verifying $\lambda^* \in \Lambda_{\bar{X}}^k$ or, equivalently, verifying that $u(\cdot; \lambda^*)$ is regular on \bar{X} . This is a difficult problem, and we do not know of any tractable algorithm that does this without error.

Various imperfect checks of regularity appear in the literature. Gallant (1981) proposes a simple algorithm for checking convexity or concavity, but the conditions it verifies are sufficient but not necessary. Gallant and Golub (1984) check quasi-concavity by searching for a minimum of a function whose non-negativity is necessary and sufficient for quasi-concavity, but there is no guarantee that the minimization procedure finds a global minimum.

All the basis functions of the Müntz Szasz expansion are regular, so there is an easy to verify sufficient condition for regularity: non-negativity of all coefficients. However, the condition is not necessary. In fact, it imposes substitutability on all inputs. Terrell (1994) uses simulations to show that even when inputs are substitutes, much flexibility is lost. See also Koop, Osiewalski and Steel (1994) on this point.

To check the regularity of $u(\cdot; \lambda^*)$, we run a battery of tests verifying necessary conditions. So although we will never attribute irregularity to a regular function, we cannot guarantee that we will detect the irregularity of an irregular function.

We note that any proposal λ^* satisfies $V\lambda^* \geq 0$ by construction, so it already survives all of the regularity testing implied by this condition.

For all $i \in \{1, \dots, n\}$, the expressions

$$g(z_i) \equiv (x_i + \xi_i) \frac{\partial u_i}{\partial x_i}(x_i; \lambda^{(i)}) \quad \text{and} \quad h(z_i) \equiv (x_i + \xi_i)^2 \frac{\partial^2 u_i}{\partial x_i^2}(x_i; \lambda^{(i)})$$

are polynomials of order $k-1$ in $z_i \equiv \phi_i(x)$. We find, numerically if necessary, all the roots of h and verify that none of these roots are in $[z_i^{lo}, z_i^{hi}]$, that $h(0) < 0$, and that $g(z_i^{hi}) > 0$.

We then verify that the gradient of u at \bar{x} is positive. Finally, we search for the maximum value of the largest eigenvalue of the Hessian of u over \bar{X} , using a simplex algorithm, and verify that its value is negative.

4.4 A Second Proposal

The second proposal consists of three steps. First we draw a random good i^* from the discrete uniform distribution on $\{1, \dots, n\}$. Then we find π_{\min} , the value of the multiplier π such that $(\lambda^{(1)}, \dots, \lambda^{(i^*-1)}, \pi \lambda^{(i^*)}, \lambda^{(i^*+1)}, \dots, \lambda^{(n)}, \lambda^\circ)$ is on the boundary of the cone $V\lambda \geq 0$. Finally, we draw π^* from a log-normal distribution truncated to $[\pi_{\min}, \infty)$ and construct the proposal

$$\lambda^* = (\lambda^{(1)}, \dots, \lambda^{(i^*-1)}, \pi^* \lambda^{(i^*)}, \lambda^{(i^*+1)}, \dots, \lambda^{(n)}, \lambda^\circ). \quad (13)$$

The advantage of the second proposal is that the chain can move quickly out of the tight corners of $\Lambda_{\bar{X}}^k$.

We compute π_{\min} as the smallest positive value of π for which $V(\lambda + (\pi - 1)w) \geq 0$, where $w = (0, \dots, 0, \lambda^{(i)}, 0, \dots, 0, 0)$. That is, we choose

$$\pi_{\min} = \max_{j \in \{1, \dots, N_v\}} \left\{ 1 - \frac{(V\lambda)_j}{(Vw)_j} \right\}.$$

We point out that there is no π_{\max} . That is, there is no $\pi > 1$ such that $V(\lambda + (\pi - 1)w) \not\geq 0$. This is because w and λ are both in $\Lambda_{\bar{X}}^k$ and $\Lambda_{\bar{X}}^k \subseteq \{\lambda : V\lambda \geq 0\}$.

Once we have π_{\min} , we draw π^* from a log-normal distribution truncated to $[\pi_{\min}, \infty)$. For the simulations described in this paper, we draw π^* such that $\log \pi^*$ has mean zero and standard deviation $\sigma \equiv 0.25$.

We then construct λ^* as in equation (13). We accept with probability

$$\min \left(1, \frac{f(\lambda^*)}{f(\lambda)} \cdot (\pi^*)^{k-2} \cdot \pi^* \cdot \frac{1 - \Phi((\log \pi_{\min})/\sigma)}{1 - \Phi((\log \pi_{\min} - \log \pi^*)/\sigma)} \right).$$

The factor $(\pi^*)^{k-2}$ is due to the fact that this is a radial draw in a $(k-1)$ -dimensional subspace of $\Lambda_{\bar{X}}^k$: the volume of the differential element increases as $(\pi^*)^{k-2}$. For more rigor on this point, and for more information on radial draws in radial co-ordinate systems, see Bauwens et al. (2004). The factor π^* comes from the Jacobian of the exponential transformation of the Gaussian draw $\log \pi^*$.

4.5 A Third Proposal

The third proposal is similar to the second proposal, except that we multiply both $\lambda^{(i^*)}$ and λ° by the same random multiplier. First we draw a random good i^* from the discrete uniform distribution on $\{1, \dots, n\}$. Then we find π_{\max} , the value of the multiplier π such that

$$(\lambda^{(1)}, \dots, \lambda^{(i^*-1)}, \pi \lambda^{(i^*)}, \lambda^{(i^*+1)}, \dots, \lambda^{(n)}, \pi \lambda^\circ)$$

is on the boundary of the cone $V\lambda \geq 0$. Finally, we draw π^* from a log-normal distribution truncated to $[0, \pi_{\max}]$ and construct the proposal

$$\lambda^* = (\lambda^{(1)}, \dots, \lambda^{(i^*-1)}, \pi^* \lambda^{(i^*)}, \lambda^{(i^*+1)}, \dots, \lambda^{(n)}, \pi^* \lambda^\circ). \quad (14)$$

The advantage of the third proposal is that the chain can move quickly into the tight corners of $\Lambda_{\bar{X}}^k$.

We compute π_{\max} as the largest value of π for which $V(\lambda + (\pi - 1)w) \geq 0$, where $w = (0, \dots, 0, \lambda^{(i)}, 0, \dots, 0, \lambda^\circ)$. That is, we choose

$$\pi_{\max} = \min_{j \in \{1, \dots, N_v\}} \left\{ 1 - \frac{(V\lambda)_j}{(Vw)_j} \right\}.$$

Regularity of the u_i , $i \neq i^*$ guarantees regularity of $u(\cdot; \lambda - w)$. Since, in addition, $\Lambda_{\bar{X}}^k$ is a cone, there is no $\pi \in (0, 1)$ such that $V(\lambda + (\pi - 1)w) \not\geq 0$.

Once we have π_{\max} , we draw π^* from a log-normal distribution truncated to $[0, \pi_{\max}]$. For the simulations described in this paper, $\log \pi^*$ is normal with mean 0 and standard deviation $\sigma \equiv 0.25$.

We then construct λ^* as in equation (14) and accept with probability

$$\min \left(1, \frac{f(\lambda^*)}{f(\lambda)} \cdot (\pi^*)^{k^n - (n-1)(k-1) - 2} \cdot \pi^* \cdot \frac{\Phi((\log \pi_{\max})/\sigma)}{\Phi((\log \pi_{\max} - \log \pi^*)/\sigma)} \right),$$

where the exponent $k^n - (n-1)(k-1) - 2$ is one less than the dimension of $(\lambda^{(i^*)}, \lambda^\circ)$.

5 An Empirical Application

We present a consumer demand application to illustrate our methods. In theory, utility or indirect utility functions represent choices exactly. Following common practice, we include a random component, or disturbance, to choices in order to obtain a data density and do likelihood-based statistical inference. However, we take an unconventional approach to the specification of the random disturbance.

Usually a measurement error approach is taken, whereby an error distribution, unrelated to preference, accounts for discrepancies between observed choices and choices which maximize utility. In contrast, we use a model for observed choices where distributions over choices are given by the utility function itself. Specifically, if $u: X \rightarrow \mathbb{R}$ is the utility function of the universe X of choices, then the distribution of observed choices is proportional to $\exp(u(x))$ on the frontier of the set of feasible choices. Theil (1974) and McCausland (2004) give very different theoretical underpinnings. There are several advantages of this approach. First, it is theoretically grounded. In usual practice, distributions of disturbances are given without theoretical justification. Second, the specification is parsimonious: a single function describes not only how choices broadly respond to changes in prices and income, but also the distribution of demand on any given budget. Third, the fit of an observed choice is measured by the relative desirability of the choice and its feasible alternatives, rather than by some measure on the choice set. Varian (1990), in a paper on goodness-of-fit measures, argues for preferring the former to the latter. Finally, the theories of Theil (1974) and McCausland (2004) do not rule out violations of the axioms of revealed preference. In practice, such violations are sometimes observed. The theories are more forgiving than standard consumer theory, without being undisciplined.

We analyze data from the Harbaugh et al. (2001) “GARP for Kids” experiment, undertaken in a study of the development of rational behavior. Subjects are 31 second grade students, 42 sixth grade students and 55 undergraduates. There are two goods, chips and juice, in indivisible packages. There are no prices and income as such: subjects are offered a budget of choices directly, and the budgets do not include off-frontier bundles. Figure 2 illustrates the eleven different budgets.

The experiment has the following features.

1. Choices are individual, rather than aggregate, so consumer theory (and

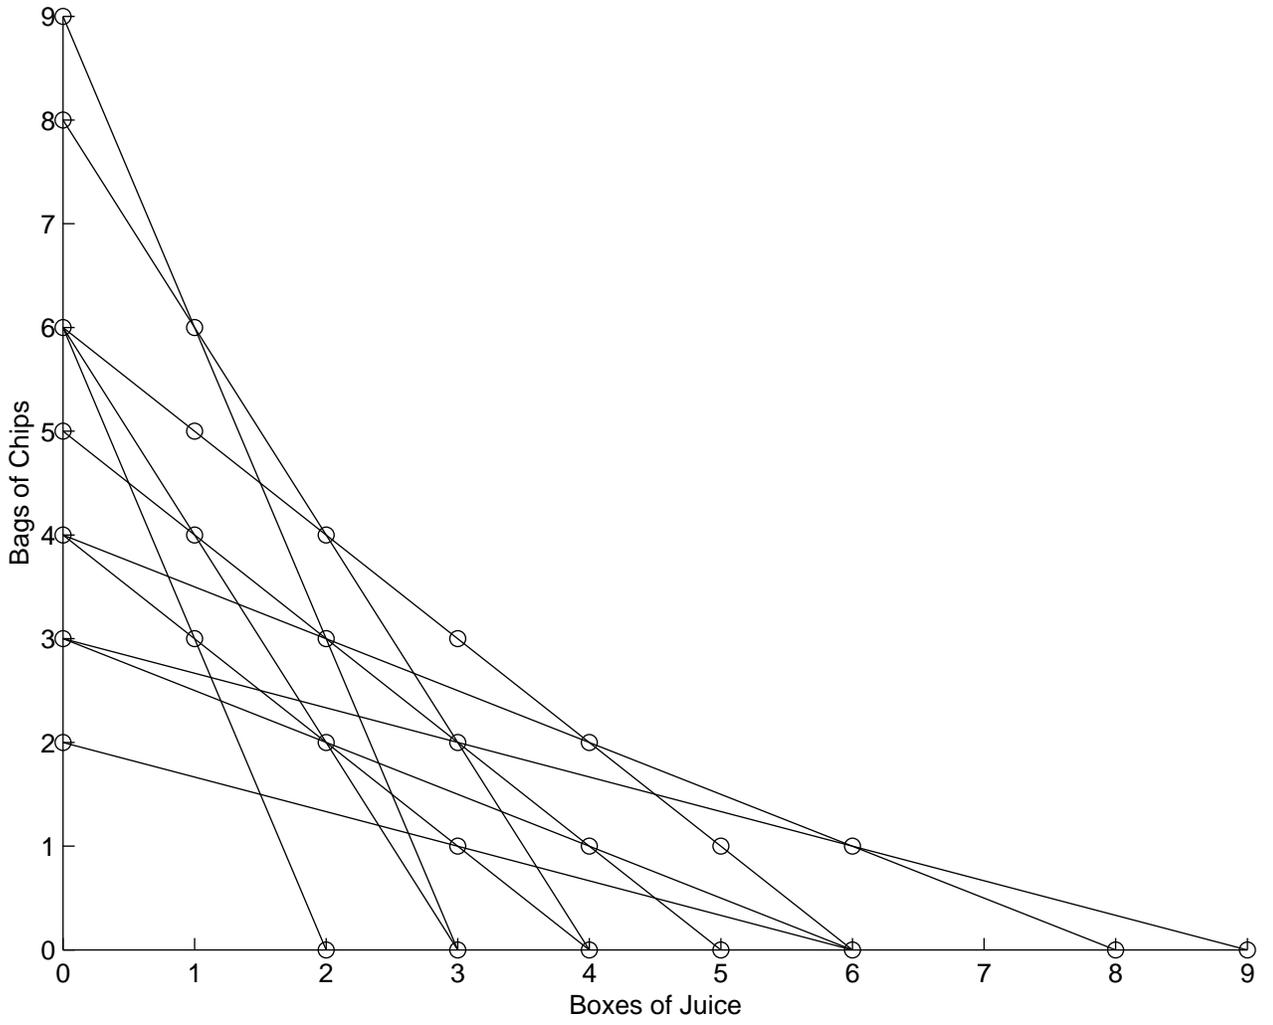


Figure 2: Budgets for the “GARP for Kids” experiment.

in particular the theory of random consumer demand in McCausland (2004)) applies.

2. Consumers select bundles from several different budgets, in the knowledge that after all decisions are made, exactly one of the budgets will be selected at random, and the consumer will be given their choice from only that budget. We can thus plausibly consider choices as being simultaneous or static, rather than dynamic.
3. Consumers have the opportunity to go back and change earlier choices, before a budget is selected at random. This mitigates the problem of learning during the experiment.
4. Choices are recorded in a laboratory. We can be fairly confident that measurement error is not a problem.
5. The number of goods, the coarseness of the indivisibilities, and prices and income are such that the number of possible choices is small. The likelihood function is therefore easily computed.

Using the theory of Theil (1974) or that of McCausland (2004), we obtain the following distribution for the t 'th choice x_t .

$$\Pr(x_t = x) = \begin{cases} \frac{\exp u(x)}{\sum_{y \in B_t} \exp u(y)} & x_t \in B_t \\ 0 & \text{otherwise,} \end{cases} \quad (15)$$

where B_t is the set of choices on the frontier of the t 'th choice set, which in this experiment is the choice set itself. We assume that the choices x_t are independent and appeal to points 2 and 3 above to justify this assumption.

A few comments on the applicability of the theory by Theil (1974) are in order, as its relevance might not be immediately clear to the reader consulting it. Theil assumes that the choice set is \mathbb{R}^n for some n , and uses a quadratic approximation for u that leads to multivariate normal choice distributions. We point out that the main result of Theil (1974) also holds for finite choice sets, and that we have no need for the approximation: if we use u directly, we obtain choice distributions given by (15).

5.1 Prior Specification

The following choices define the prior distribution over the constellation index k and the vector λ of coefficients. The restricted consumption set is $\bar{X} =$

$[0, 12]^2$, which is compact and contains all eleven budgets. The constants defining the transformation ϕ are $x^* = (1.0, 1.0)$ and $\xi = (0.1, 0.1)$. Thus ϕ is given by

$$\begin{aligned}\phi(x_1, x_2) &= \left(\log \left(\frac{x_1 + \xi_1}{x_1^* + \xi_1} \right), \log \left(\frac{x_2 + \xi_2}{x_2^* + \xi_2} \right) \right) \\ &= \left(\log \left(\frac{x_1 + 0.1}{1.1} \right), \log \left(\frac{x_2 + 0.1}{1.1} \right) \right).\end{aligned}$$

The multi-index constellations are rectangular lattices given by

$$I_k \equiv \{\iota \in \mathbb{N}_0^n : \iota \neq 0 \text{ and } 0 \leq \iota_i \leq k - 1 \text{ for } i = 1, \dots, n\}.$$

The multi-indices in I_3 , for example, are $(0, 1)$, $(0, 2)$, $(1, 0)$, $(1, 1)$, $(1, 2)$, $(2, 0)$, $(2, 1)$ and $(2, 2)$, corresponding to the basis functions ϕ_2 , ϕ_2^2 , ϕ_1 , $\phi_1\phi_2$, $\phi_1\phi_2^2$, ϕ_1^2 , $\phi_1^2\phi_2$, and $\phi_1^2\phi_2^2$.

The prior on k , the constellation index, is given by $\pi_k = 2^{-(k-1)}$ for $k \geq 2$. Thus we have three terms with probability $1/2$, eight with probability $1/4$, 15 with probability $1/8$, and so on. The conditional prior density of λ given k has the form given in equation (11), where v is the vector such that $(\Delta u_1 + \Delta u_2)/2 = v'\lambda$, g is the gamma density with shape parameter 4 and scale parameter 50, and h is the factor defined in (12), with $\delta = 0.001$ and $p = k - 1$.

5.2 Results

We present results for the ‘‘GARP for Kids’’ experiment. The objective here is not to study the development of rational behavior in children, and so we report results only for the benchmark undergraduate subjects.

Table 1 shows conditional posterior probabilities for k given $k \in \{2, 3, 4, 5\}$, for subjects 1074 through 1103, the first 30 of 55. We use the method of Newton and Raftery (1994) to compute $f(x_1, \dots, x_{11}|k)$ for each value of k .

Table 2 lists the log marginal likelihoods for all 55 subjects, given $k \in \{2, 3, 4, 5\}$. They are in row major order, so the first row gives the marginal likelihoods for subjects 1074 through Here, the marginal likelihood is the marginal probability the model, including prior, assigns to the sequence of observed choices that a subject makes. Standard errors for the numerical approximation of the log marginal likelihoods reported in Table 2 are all less than 0.15. The average log marginal likelihood is -12.49 . Subject 1088 (row

3, column 5) always spent all his income on the more *expensive* good. This exuberant irrationality earned him a log marginal likelihood of -30.36, by far the lowest.

Subject 1105 (row 7, column 2) always spent his income on the cheaper good, and chose equal quantities of the goods when their prices were equal. Although rational, these choices require very specific abrupt changes in expenditure shares in response to small changes in relative prices. Only a small set of λ account for this behavior, and the marginal likelihood is a low -21.43. However, we can see in Figure 3 that the posterior distribution puts high probability on these abrupt changes. Both panels of Figure 3 show 20 curves. Each curve gives the mode of the expenditure share of good one as a function of the ratio w_1/w_2 of income-normalized prices w_1 and w_2 , adding to one. The first panel shows curves for 20 draws of the utility function from its conditional prior distribution given $k = 4$. The second panel does the same for the conditional posterior distribution given $k = 4$ and the observed choices of Subject 1105. The curves with the most abrupt changes tend to be associated with high amplitude utility functions, for which the expenditure share distribution is quite tight. The smoother curves tend to correspond to more diffuse expenditure share distributions.

To put the log marginal likelihoods in perspective, we consider the average log marginal likelihood arising from various models. The model assigning equal probability to all possible sequences of eleven choices implies a log marginal likelihood of -16.31 for every subject. A model which correctly and with certainty predicts the behavior of all subjects on all budgets implies a log marginal likelihood of zero for every subject. Any model that assigns probability zero to every sequence featuring at least one violation of the Generalized Axiom of Revealed Preference (GARP) gives a log marginal likelihood of negative infinity to the sequences of the 19 out of 55 subjects who violated the GARP, and therefore an average log marginal likelihood of negative infinity.

We use the data in Table 3 to derive a maximum log marginal likelihood of -13.36 over all models assigning equal probabilities to all sequences with the same number of GARP violations. The second column gives, for the number of GARP violations in the first column, the number of subjects having that number of violations. The third column gives the total number of distinct sequences of eleven choices having that number of violations.

Table 1: Posterior probabilities of k for the first 30 Subjects in the “GARP for Kids” experiment

| Subject | $k = 2$ | $k = 3$ | $k = 4$ | $k = 5$ |
|---------|----------|----------|----------|----------|
| 1074 | 0.268031 | 0.407503 | 0.219064 | 0.105402 |
| 1075 | 0.345172 | 0.324044 | 0.184888 | 0.145895 |
| 1076 | 0.258968 | 0.462476 | 0.203739 | 0.074816 |
| 1077 | 0.967457 | 0.032059 | 0.000479 | 0.000005 |
| 1078 | 0.968034 | 0.031557 | 0.000405 | 0.000004 |
| 1079 | 0.198167 | 0.339118 | 0.300826 | 0.161889 |
| 1080 | 0.025543 | 0.821864 | 0.147249 | 0.005344 |
| 1081 | 0.233941 | 0.318601 | 0.266747 | 0.180711 |
| 1082 | 0.359153 | 0.390829 | 0.172782 | 0.077236 |
| 1083 | 0.309870 | 0.333014 | 0.223047 | 0.134069 |
| 1084 | 0.363238 | 0.184935 | 0.419443 | 0.032384 |
| 1085 | 0.966052 | 0.033497 | 0.000446 | 0.000005 |
| 1086 | 0.125813 | 0.624758 | 0.222555 | 0.026874 |
| 1087 | 0.817362 | 0.100623 | 0.075123 | 0.006892 |
| 1088 | 0.414538 | 0.250175 | 0.197775 | 0.137512 |
| 1089 | 0.415736 | 0.303844 | 0.176036 | 0.104384 |
| 1090 | 0.072000 | 0.025861 | 0.900104 | 0.002035 |
| 1091 | 0.632163 | 0.222689 | 0.105932 | 0.039216 |
| 1092 | 0.968132 | 0.031479 | 0.000385 | 0.000004 |
| 1093 | 0.257544 | 0.405650 | 0.222257 | 0.114550 |
| 1094 | 0.269312 | 0.462051 | 0.203084 | 0.065554 |
| 1095 | 0.967018 | 0.032522 | 0.000457 | 0.000003 |
| 1096 | 0.101630 | 0.505159 | 0.335330 | 0.057882 |
| 1097 | 0.265254 | 0.415457 | 0.226562 | 0.092727 |
| 1098 | 0.371295 | 0.358028 | 0.181920 | 0.088757 |
| 1099 | 0.894937 | 0.088172 | 0.016743 | 0.000148 |
| 1100 | 0.118235 | 0.306723 | 0.319137 | 0.255905 |
| 1101 | 0.970699 | 0.028890 | 0.000407 | 0.000004 |
| 1102 | 0.967782 | 0.031750 | 0.000465 | 0.000004 |
| 1103 | 0.427316 | 0.360936 | 0.129459 | 0.082289 |

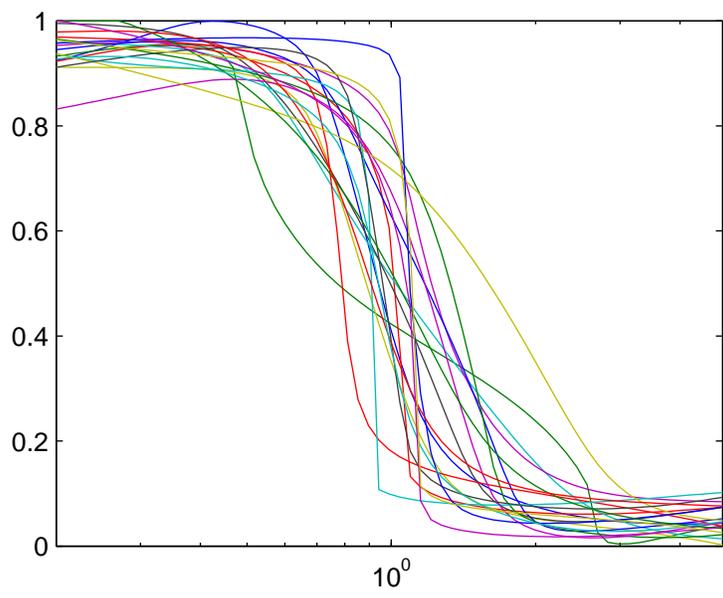
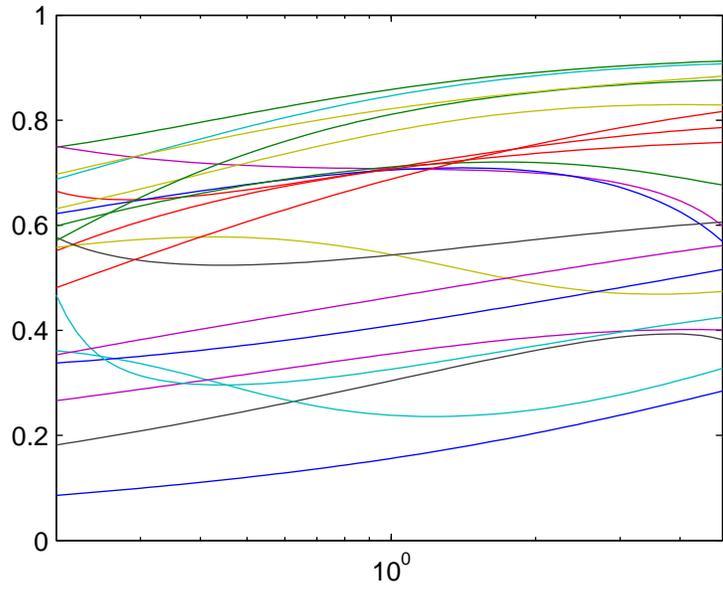


Figure 3: Prior and posterior scattergraph, for subject 1105, of the modal expenditure share of good one versus w_1/w_2 , where w_1 and w_2 are income-normalized prices adding to 1. 33

Table 2: Log marginal likelihoods for all subjects in the “GARP for Kids” experiment, in row major order

| | | | | |
|------------|------------|------------|------------|------------|
| -7.351133 | -4.608758 | -9.940656 | -4.197768 | -4.183579 |
| -14.935639 | -6.301195 | -15.442465 | -8.161736 | -16.676400 |
| -16.713183 | -4.233398 | -17.822427 | -10.557720 | -30.356065 |
| -22.978277 | -24.813786 | -23.485596 | -4.202415 | -7.330709 |
| -9.970015 | -4.236868 | -16.950571 | -7.352623 | -8.396829 |
| -10.407859 | -12.575744 | -4.160754 | -4.224852 | -8.894845 |
| -8.242575 | -21.439450 | -11.572254 | -12.535711 | -18.998176 |
| -7.359645 | -7.946763 | -14.665421 | -18.625966 | -4.218446 |
| -9.468318 | -9.972209 | -21.650330 | -8.078603 | -13.919956 |
| -22.521666 | -15.678075 | -18.931139 | -7.500323 | -21.632655 |
| -20.188184 | -4.211171 | -16.273644 | -7.366309 | -22.680512 |

6 Conclusions

We have pointed out that instead of approximating a function u on a compact restricted domain \bar{X} , we can approximate $u \circ \phi^{-1}$ on $\phi(\bar{X})$, where ϕ is a function that can be chosen to be non-decreasing and concave, facilitating the approximation of non-decreasing concave functions. With modest restrictions on ϕ , we can simultaneously approximate the function, its gradient and Hessian. We apply this idea using basis functions that are polynomials, but we note that we could easily do the same with the sinusoidal basis functions of the Fourier flexible form or the basis functions of the Müntz-Szasz expansion.

We have put considerable emphasis on prior elicitation, to which we believe the literature has not paid sufficient attention. We have described an approach to prior elicitation that puts more emphasis on economically relevant quantities than on the parameters themselves. We have shown that there are tight corners of the parameter space which, despite their low volume, may be quite plausible in many applications. We have shown that by concentrating prior probability in these corners, we can obtain more diffuse prior distributions over economically relevant quantities such as marginal rates of substitution and ratios of quantities demanded.

We believe this is relevant for applied economic research. Practitioners should be aware that flat priors on the parameter space can imply quite

Table 3: Counts of numbers of GARP violations

| Number of violations | Experimental subjects | All sequences |
|----------------------|-----------------------|---------------|
| 0 | 36 | 108,846 |
| 1 | 0 | 0 |
| 2 | 1 | 140,788 |
| 3 | 5 | 171,718 |
| 4 | 0 | 272,978 |
| 5 | 7 | 438,074 |
| 6 | 1 | 646,288 |
| 7 | 0 | 928,790 |
| 8 | 0 | 1,567,246 |
| 9 | 2 | 2,081,452 |
| 10 | 1 | 2,555,030 |
| 11 | 2 | 3,184,790 |
| Total | 55 | 12,096,000 |

informative priors about economically relevant quantities. We recommend prior simulation to ensure that these quantities have reasonable prior distributions. We have provided an example of how one can modify a prior if it assigns too informative a distribution for these quantities.

We have noted that the tight corners of the parameter space also pose a problem for prior and posterior simulation and that this problem is aggravated by the concentration of prior probability within them. We have introduced Metropolis-Hastings chains that are able to move quickly in and out of these tight corners, and thereby efficiently sample the prior and posterior distributions.

We have demonstrated the use of these prior distributions and simulation methods for the analysis of data from a consumer experiment. We find that out-of-sample predictions are better on average than any model based only on numbers of observed violations of the Generalized Axiom of Revealed Preference. This result provides some support for the theory of random consumer demand in McCausland (2004), which we used to construct the likelihood functions.

A On the Boundedness of Cone Truncations

This appendix is on the boundedness of a truncation of the cone $\Lambda_{\bar{X}}^k$ defined in equation (3). This boundedness is relevant for the construction of proper prior distributions on $\Lambda_{\bar{X}}^k$ as described in Section 3.2.

Result A.1 *The set $\{\lambda \in \Lambda_{\bar{X}}^k : u(\bar{x}; \lambda) \leq 1\}$ is bounded.*

Proof. We first establish bounds on the values of $u(x; \lambda)$ on the subset $[x_1^*, \bar{x}_1] \times \dots \times [x_n^*, \bar{x}_n]$ of the restricted domain \bar{X} .

Claim A.1 *For every $x \in [x_1^*, \bar{x}_1] \times \dots \times [x_n^*, \bar{x}_n]$ and every $\lambda \in \{\lambda \in \Lambda_{\bar{X}}^k : u(\bar{x}; \lambda) \leq 1\}$*

$$0 \leq u(x; \lambda) \leq 1.$$

Proof. The claim follows directly from the monotonicity of $u(\cdot; \lambda)$, the fact that $u(x^*; \lambda) = 0$, and the fact that $u(\bar{x}; \lambda) \leq 1$. \square

We now bound $\{\lambda \in \Lambda_{\bar{x}}^k : u(\bar{x}; \lambda) \leq 1\}$ by enclosing it in a hyperparallelogram defined by $\{\lambda \in \mathbb{R}^{|I_k|} : 0 \leq C\lambda \leq (1, 1, \dots, 1)\}$, where C is a non-singular matrix. The non-singularity of C guarantees that the hyperparallelogram is bounded, and thus that $\{\lambda \in \Lambda_{\bar{x}}^k : u(\bar{x}; \lambda) \leq 1\}$ is bounded.

Let $J = |I_k|$, and order the multi-indices $\iota \in I_k$ as $\iota^{(1)}, \dots, \iota^{(J)}$. Choose vector $q = (q_1, \dots, q_n)$ such that for every $i \in \{1, \dots, n\}$,

1. There exist positive integers m_N and m_D such that $q_i = p_{2i}^{m_N} / p_{2i-1}^{m_D}$, where p_i is the i 'th prime number, and
2. $[\phi_i(\frac{x^* + \bar{x}}{2}) / \phi_i(\bar{x})]^{1/J} \leq q_i \leq 1$.

A simple modification of the proof in Rudin (1976) of the denseness of the rational numbers in the reals shows that we can do this. The inequalities $x^* < (x^* + \bar{x})/2 < \bar{x}$ ensure that we are taking the J 'th root of a positive real number strictly less than one.

Now define, for all $j \in \{1, \dots, J\}$,

$$z_j \equiv (q_1^j \phi_1(\bar{x}), \dots, q_n^j \phi_n(\bar{x})), \quad x_j \equiv \phi^{-1}(z_j),$$

and

$$C \equiv \begin{bmatrix} z_1^{\iota^{(1)}} & \dots & z_1^{\iota^{(J)}} \\ \vdots & \ddots & \vdots \\ z_J^{\iota^{(1)}} & \dots & z_J^{\iota^{(J)}} \end{bmatrix}.$$

For all $\lambda \in \Lambda_{\bar{x}}^k$,

$$C\lambda = [u(x_1; \lambda), \dots, u(x_J; \lambda)]' = [u(\phi^{-1}(z_1); \lambda), \dots, u(\phi^{-1}(z_J); \lambda)]',$$

and for all $j \in \{1, \dots, J\}$, $x_j = \phi^{-1}(z_j) \in [x_1^*, \bar{x}_1] \times \dots \times [x_n^*, \bar{x}_n]$. Claim A.1 gives us $(0, \dots, 0)' \leq C\lambda \leq (1, \dots, 1)'$.

We now show that C is non-singular.

Claim A.2 C is non-singular.

Proof. C can be written as

$$\begin{aligned} & \begin{bmatrix} (q^{\iota^{(1)}})^1 [\phi(\bar{x})]^{\iota^{(1)}} & \cdots & (q^{\iota^{(J)}})^1 [\phi(\bar{x})]^{\iota^{(J)}} \\ \vdots & \ddots & \vdots \\ (q^{\iota^{(1)}})^J [\phi(\bar{x})]^{\iota^{(1)}} & \cdots & (q^{\iota^{(J)}})^J [\phi(\bar{x})]^{\iota^{(J)}} \end{bmatrix} \\ &= \begin{bmatrix} (q^{\iota^{(1)}})^1 & \cdots & (q^{\iota^{(J)}})^1 \\ \vdots & \ddots & \vdots \\ (q^{\iota^{(1)}})^J & \cdots & (q^{\iota^{(J)}})^J \end{bmatrix} \cdot \text{diag} \left([\phi(\bar{x})]^{\iota^{(1)}}, \dots, [\phi(\bar{x})]^{\iota^{(J)}} \right) \end{aligned}$$

We will show that both these factors are non-singular, which will then imply that C is non-singular. The first factor is a Vandermonde matrix, and to establish its non-singularity, it suffices to show that for all $j, l \in \{1, \dots, J\}$, $j \neq l \Rightarrow q^{\iota^{(j)}} \neq q^{\iota^{(l)}}$. This follows from the fact that there is a unique representation of any rational number as the ratio of two integers with no common factors, and unique prime factorizations of the two integers. The second factor is a diagonal matrix whose elements are non-zero, and so it is also non-singular. Since the two factors are non-singular, so is C . \square

References

- [1] William A. Barnett, John F. Geweke, and Michael Wolfe. Semionparametric Bayesian estimation of the asymptotically ideal production model. *Journal of Econometrics*, 49:5–50, 1991.
- [2] William A. Barnett, John F. Geweke, and Piyu Yue. Semionparametric Bayesian estimation of the Asymptotically Ideal Model: The AIM consumer demand system. In William A. Barnett, James Powell, and George E. Tauchen, editors, *Nonparametric and semiparametric methods in econometrics and statistics: Proceedings of the Fifth International Symposium in Economic Theory and Econometrics*, chapter 6, pages 127–173. Cambridge University Press, Cambridge, New York and Melbourne, 1991.
- [3] Luc Bauwens, Charles S. Bos, Herman K. van Dijk, and Rutger D. van Oest. Adaptive radial-based direction sampling: some flexible and robust Monte Carlo integration methods. *Journal of Econometrics*, 123:201–225, 2004.

- [4] B. de Finetti. Sulle stratificazioni convesse. *Ann. de Mat. Pura Appl.*, 30:173–183, 1949.
- [5] Angus Deaton and John Muellbauer. *Economics and Consumer Behavior*. Cambridge University Press, Cambridge, 1980.
- [6] W. E. Diewert and T. J. Wales. Flexible functional forms and global curvature conditions. *Econometrica*, 55(1):43–68, 1987.
- [7] J. C. Evard and F. Jafari. Direct computation of the simultaneous Stone-Weierstrass approximation of a function and its partial derivatives in banach spaces, and combination with hermite interpolation. *Journal of Approximation Theory*, 78:351–363, 1994.
- [8] A. Ronald Gallant. On the bias in flexible functional forms and an essentially unbiased form: The Fourier flexible form. *Journal of Econometrics*, 15:211–245, 1981.
- [9] A. Ronald Gallant and Gene H. Golub. Imposing curvature restrictions on flexible functional forms. *Journal of Econometrics*, 26:295–321, 1984.
- [10] John Geweke. Monte Carlo simulation and numerical integration. In Hans M. Amman, David A. Kendrick, and John Rust, editors, *Handbook of Computational Economics: Volume 1*, chapter 15, pages 731–800. Elsevier Science Publishers, North–Holland, 1996.
- [11] John Geweke. Using simulation methods for Bayesian econometric models: Inference, development, and communication. *Econometric Reviews*, 18:1–126, 1999.
- [12] John Geweke and Lea Petrella. Inference for regular functions using Müntz approximations. Unpublished Manuscript, 2000.
- [13] Stephen Gordon. Using mixtures of flexible functional forms to estimate factor demand elasticities. *Canadian Journal of Economics*, 29(3):717–736, 1996.
- [14] William Harbaugh, K. Krause, and T. Berry. GARP for kids: On the development of rational choice behavior. *American Economic Review*, 91:1539–1545, 2001.

- [15] W. K. Hastings. Monte Carlo sampling methods using Markov chains and their applications. *Biometrika*, 57:97–109, 1970.
- [16] Y. Kannai. Approximation of convex preferences. *Journal of Mathematical Economics*, 1:101–106, 1974.
- [17] Gary Koop, Jacek Osiewalski, and Mark F. J. Steel. Bayesian efficiency analysis with a flexible form: the AIM cost function. *Journal of Business and Economic Statistics*, 12(3):339–346, 1994.
- [18] Lawrence Lau. Functional forms in econometric model building. In Zvi Gliliches and M. D. Intriligator, editors, *Handbook of Econometrics: Volume 3*, chapter 26, pages 1515–1566. Elsevier Science Publishers, North–Holland, 1986.
- [19] Rosa L. Matzkin. Restrictions of economic theory in nonparametric methods. In R. F. Engle and D. L. McFadden, editors, *Handbook of Econometrics: Volume 4*, chapter 42, pages 2523–2558. Elsevier Science Publishers, North–Holland, 1994.
- [20] William J. McCausland. A theory of random consumer demand. Cahiers de recherche du Département de sciences économiques, Université de Montréal, no. 2004-04, 2004.
- [21] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller. Equation of state calculations by fast computing machines. *J. Chem. Phys.*, 21:1087–1092, 1953.
- [22] M. A. Newton and A. E. Raftery. Approximate Bayesian inference by the weighted likelihood bootstrap. *Journal of the Royal Statistical Society Series B*, 3:3–48, 1994.
- [23] Walter Rudin. *Principles of Mathematical Analysis*. McGraw-Hill, New York, 1976.
- [24] Dek Terrell. Incorporating monotonicity and concavity conditions in flexible functional forms. *Journal of Applied Econometrics*, 11:179–194, 1996.
- [25] Henri Theil. A theory of rational random behavior. *Journal of the American Statistical Association*, 69:310–314, 1974.

- [26] Hal R. Varian. Goodness-of-fit in optimizing models. *Journal of Econometrics*, 46:125–140, 1990.