Available online at www.sciencedirect.com
ScienceDirect
Journal of Econometrics 142 (2008) 484–507
JOURNAL OF Econometrics
www.elsevier.com/locate/jeconom

On Bayesian analysis and computation for functions with monotonicity and curvature restrictions

William J. McCausland[a,b,c,*]

[a]*Département de sciences économiques, Université de Montréal, C.P. 6128, Succursale Centre-ville, Montréal, QC, Canada H3C 3J7*
[b]*Centre Interuniversitaire de Recherche en Économie Quantitative, Canada*
[c]*Centre Interuniversitaire de Recherche en Analyse des Organisations, Canada*

Available online 14 September 2007

Abstract

Our goal is inference for shape-restricted functions. Our functional form consists of finite linear combinations of basis functions. Prior elicitation is difficult due to the irregular shape of the parameter space. We show how to elicit priors that are *flexible*, *theoretically consistent*, and proper. We demonstrate that uniform priors over coefficients imply priors over economically relevant quantities that are quite informative and give an example of a non-uniform prior that addresses this issue. We introduce simulation methods that meet challenges posed by the shape of the parameter space. We analyze data from a consumer demand experiment.

JEL classification: C11; C14; C15; D12; D21

Keywords: Bayesian methods; Flexible functional forms; Shape restrictions; Consumer demand

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

*Département de sciences économiques, Université de Montréal, C.P. 6128, Succursale Centre-ville, Montréal, QC, Canada H3C 3J7. Tel.: +1 514 343 7281; fax: +1 514 343 7221.

E-mail address: william.j.mccausland@umontreal.ca

0304-4076/$ - see front matter
doi:10.1016/j.jeconom.2007.08.018

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 (see Christensen et al. (1975)) 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 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, gives 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 linear and quadratic terms.

Barnett and Jonas (1983) use a multivariate Müntz–Szasz expansion to approximate a firm's unit cost function, a function of the prices $p_1, \ldots, p_n$ 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 $\mathbb{N}_0 \equiv 0, 1, 2, \ldots$ and the sequence $\lambda(k)$, $k = 1, 2, \ldots$, satisfies $\sum_{k=1}^{\infty} (1/\lambda(k)) = \infty$. Barnett and Jonas (1983) and Barnett et al. (1991a, b) 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 (2000) also approximate a firm's unit cost as a function of input prices $p_1, \ldots, p_n$, but use the following set of basis functions:

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

where $b > 0$. The functions $\prod_{i=1}^{n} p_i^{b\iota_i}$ satisfying $\iota_i < b^{-1}$, $i = 1, \ldots, n$, are themselves non-decreasing and concave, which is convenient for constructing approximations of non-decreasing concave unit cost functions using a