Journal of Econometrics 191 (2016) 129–144

Contents lists available at ScienceDirect

Journal of Econometrics

journal homepage: www.elsevier.com/locate/jeconom

Testing multivariate economic restrictions using quantiles: The example of Slutsky negative semidefiniteness

Holger Dette [a], Stefan Hoderlein [b,*], Natalie Neumeyer [c]

[a] Ruhr-Universität Bochum, Fakultät für Mathematik, 44780 Bochum, Germany

[b] Department of Economics, Boston College, 140 Commonwealth Avenue, Chestnut Hill, MA 02467, USA

[c] University of Hamburg, Department of Mathematics, Bundesstrasse 55, 20146 Hamburg, Germany

ARTICLE INFO

Article history:
Received 21 April 2011
Received in revised form
21 February 2015
Accepted 28 July 2015
Available online 17 October 2015

Keywords:
Nonparametric testing
Heterogeneity
Integrability
Nonseparable models
Consumer demand
Quantile regression

ABSTRACT

This paper is concerned with testing a core economic restriction, negative semidefiniteness of the Slutsky matrix. We consider a system of nonseparable structural equations with infinite dimensional unobservables, and employ quantile regression methods because they allow us to utilize the entire distribution of the data. Difficulties arise because the restriction involves several equations, while the quantile is a univariate concept. We establish that we may use quantiles of linear combinations of the dependent variable, develop a new empirical process based test that applies kernel quantile estimators, and investigate its finite and large sample behavior. Finally, we apply all concepts to Canadian microdata.

1. Introduction

Economic theory yields strong implications for the actual behavior of individuals. In the standard utility maximization model for instance, economic theory places strong restrictions on individual responses to changes in prices and wealth, the so-called integrability constraints. These restrictions are inherently restrictions on individual level: They have to hold for every preference ordering and every single individual, at any price wealth combination. Other than obeying these restrictions, the individuals' idiosyncratic preference orderings may exhibit a lot of differences. Indeed, standard parametric cross section mean regression methods applied to consumer demand data often exhibit R^2 between 0.1 and 0.2. Today, the consensus is that the majority of the unexplained variation is precisely due to unobserved preference heterogeneity. For this reason, the literature has become increasingly interested in exploiting all the information about unobserved heterogeneity contained in the data, in particular using the quantiles of the dependent variable.

To lay out our model, let y denote the $L-1$ vector of quantities demanded. At this stage, we have already imposed the adding up constraint (i.e., out of L goods we have deleted the last). Let p denote the L vector of prices, and x denote income (total expenditure).[1] For every individual, define the cost function $C(p,u)$ to give the minimum cost to attain utility level u facing the L-vector of prices p, and given income (more precisely, total outlay) x. The Slutsky negative semidefiniteness restriction arises from the fact that the cost function is concave, and hence the matrix of second derivatives is negative semidefinite (nsd, henceforth). For brevity, we will sometimes equate negative semidefiniteness with "rationality", even though it is only one facet of rationality in this setup with linear budget constraint.[2]

Obviously, this hypothesis has to hold for any preference ordering u. However we do not observe the individual's preference ordering u, and only observe a K dimensional vector of household covariates (denoted q). Specifically, we assume to have n *iid* observations on individuals from an underlying heterogeneous

* Corresponding author. Tel.: +1 617 552 6042.
E-mail addresses: holger.dette@rub.de (H. Dette), stefan_hoderlein@yahoo.com (S. Hoderlein), neumeyer@math.uni-hamburg.de (N. Neumeyer).

http://dx.doi.org/10.1016/j.jeconom.2015.07.004
0304-4076/

1 This is the income concept commonly used in consumer demand. It is motivated by the assumption of separability of preferences over time and from other decisions (e.g., the labor supply decision). We use the phrases "total expenditure", "income" and "wealth" interchangeably throughout this paper.

2 Since negative semidefiniteness is arguably the core rationality restriction, we feel that this shortcut is justified, but we would like to alert the reader to this. See also the discussion in Section 2 on the relation to the weak axiom, and after Theorem 1.

population characterized by random variables U, Y, X, P, Q which have a nondegenerate joint distribution $F_{U,Y,X,P,Q}$.

The question of interest is now as follows: What can we learn from the observable part of this distribution, i.e. $F_{Y,X,P,Q}$, about whether the Slutsky matrix is negative semidefinite across a heterogeneous population, for all values of (p, x, u). In Hoderlein (2011), we consider testing negative semidefiniteness in such a setting with mean and second moment regressions only. However, these lower order moment regressions have the disadvantage that they use only one feature of $F_{Y,X,P,Q}$, and not the entire distribution. Therefore, in this paper we propose to exploit the distributional information by using all the α-quantiles of the conditional distribution of observables, which (with varying α) employ all the information that may be obtained from the data about the economic hypothesis of interest.

There are two immediate difficulties now, and solving them is the major innovation this paper introduces. The first is how to relate a specific economic property in the (unobservable) world of nonseparable functions to observable regression quantiles. The second one is how to use quantiles in systems of equations. The solution for the second difficulty is to consider linear combinations of the dependent variable, i.e. $Y(b) = b'Y$ for all vectors b of unit length and consider the respective conditional α-quantiles of this quantity. This can be thought of as an analogue to the Cramer–Wold device, and is a strategy that is feasible more generally, e.g., when testing omission of variables. As b and α vary, we exploit the entire distribution of observables.

The solution to the first of these two difficulties involves obviously identifying assumptions. To this end, since we are dealing with nonseparable models we require full conditional independence, i.e., we require that $U \perp (P, X) | Q$, or versions of this assumption that control for endogeneity. These assumptions are versions of the "selection on observables" assumptions in the treatment effect literature. Essentially they require that, in every subpopulation defined by $Q = q$, preferences as well as prices and income be independently distributed. Although endogeneity is not relevant for our application, our treatment covers the control function approach to handle endogeneity in nonseparable models discussed in Altonji and Matzkin (2005), Imbens and Newey (2009) or Hoderlein (2011), by simply adding endogeneity controls V to the set of household control variables Q. From now on, we denote by W the set of all observable right hand side variables, i.e., (P', X, Q'), and potentially in addition V, if we are controlling for endogeneity.[3]

Under this assumption and some regularity conditions, our first main contribution is as follows: We establish that the rationality hypothesis in the underlying population has a testable implication on the distribution of the data, specifically, on the conditional quantiles of linear combinations of the dependent variables. Consequently, we can test a null hypothesis in the underlying (unobservable) heterogeneous population model in the sense that a rejection of the testable implication leads also to a rejection of the original null hypothesis. While this procedure controls size, though likely conservative, it may suffer from low power: If we do not reject the implication, we cannot conclude that the data could not have been generated by some other, nonrational mechanism. This is the price we pay for being completely general, as the only material assumption that we require to relate the observable object and the underlying heterogeneous population is the conditional independence assumption $U \perp (P, X) | Q$, and no other material assumption on the functional form of demand or their distribution enters the model. In particular, we have not assumed any monotonicity or triangularity assumption; there can be infinitely many unobservables, and they can enter in arbitrarily complicated form; in a sense, every individual can have its own nonparametric utility function, and hence this framework is closer to a random functions setup. Assessing and testing the validity of the related weak axiom of revealed preferences in a similar random utility function setup has been investigated in Hoderlein and Stoye (2014), and Kitamura and Stoye (2014).

Our second main contribution is proposing a quantile regression based nonparametric test statistic. Specifically, we apply the sample counterparts principle to obtain a nonparametric test statistic of the testable implication, and derive its' large sample properties. We show weak convergence of a corresponding standardized stochastic process to a Gaussian process and obtain an asymptotically valid hypothesis test. Moreover, we propose a bootstrap version of our test statistic. To avoid the generation of bootstrap observations under the null we adapt the well known idea of residual bootstrap for our specific model and use a centered version of the stochastic process. Nonparametric tests involving quantiles are surprisingly scant, and we list the closest references in the following paragraph. Specifically, in a system of equations setup we are the first to propose a quantile based test of an economic hypothesis, and to implement such a test using real world data.

Our test is a pointwise test, meaning that it holds locally for $W = w_0$. The main reason for this is that we aim at a more detailed picture of possible rejections, providing a better description of the rationality of the population (e.g., one outcome is that negative semidefiniteness is rejected for 20% of the population (=representative positions at which the test is evaluated)).[4]

Literature Testing the key integrability constraints that arise out of utility maximization dates back at least to the early work of Stone (1954), and has spurned the extensive research on (parametric) flexible functional form demand systems (e.g., the Translog, cf. Jorgenson et al. (1982), and the Almost Ideal, cf. Deaton and Muellbauer (1980)). Nonparametric analysis of some derivative constraints was performed by Stoker (1989) and Härdle et al. (1991), but none of these has its focus on modeling unobserved heterogeneity. More closely related to our approach is Lewbel (2001) who analyzes integrability constraints in a purely exogenous setting, but does not use distributional information nor suggests or implements an actual test. An alternative method for checking some integrability constraints is revealed preference analysis, see Blundell et al. (2003), and references therein. An approach that combines revealed preference arguments with a demand function structure is Blundell et al. (2011). As a side result, this paper develops a test of the weak axiom of revealed preferences, but in contrast to our paper, this paper assumes a scalar unobservable that enters monotonically in a single equation setup.

While our approach extends earlier work on demand systems, it is very much a blueprint for testing all kinds of economic hypothesis in systems of equations. Due to the nonseparable framework we employ, our approach extends the recent work on nonseparable models—in particular Hoderlein (2011), Hoderlein and Mammen (2007), Imbens and Newey (2009), Matzkin (2003). When it comes to dealing with unobserved heterogeneity, there are two strands in this literature: The first assumes triangularity and monotonicity in the unobservables (Chesher, 2003; Matzkin, 2003). The triangularity and monotonicity assumptions are,

[3] In classical consumer demand, it is typically income that is considered endogenous, see the discussion in Section 2.

[4] As an alternative, we could consider the behavior of the test at a fixed grid of values $w_1, \ldots, w_B$. Since the pointwise estimators at the different grid points are asymptotically independent, the present theory extends straightforwardly, and merely results in a more cumbersome notation. We could also consider the sup over these values, see the discussion in Remark 2 below.