Journal of Econometrics 186 (2015) 129–141

Contents lists available at ScienceDirect

Journal of Econometrics

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

Revealed preference tests for weak separability: An integer programming approach[☆]

Laurens Cherchye [a], Thomas Demuynck [b], Bram De Rock [c,*], Per Hjertstrand [d]

[a] Center for Economic Studies, University of Leuven, E. Sabbelaan 53, B-8500 Kortrijk, Belgium
[b] Maastricht University, Tongersestraat 53, 6370 Maastricht, Netherlands
[c] ECARES-ECORE, Université Libre de Bruxelles, Avenue F. D. Roosevelt 50, CP 114, B-1050 Brussels, Belgium
[d] Research Institute of Industrial Economics (IFN), P.O. Box 55665, SE-102 15 Stockholm, Sweden

ARTICLE INFO

Article history:
Received 20 July 2012
Received in revised form
21 May 2014
Accepted 7 July 2014
Available online 21 July 2014

JEL classification:
C14
C60
D01
D10

Keywords:
Weak separability
Revealed preference
Integer programming

ABSTRACT

We present the revealed preference conditions that characterize the data sets that are consistent with the maximization of a weakly separable utility function. We show that verifying these revealed preference conditions is NP-hard. We also present an integer programming approach, which is particularly attractive in view of empirical analysis. We demonstrate the versatility of this integer programming approach by showing that it allows for testing homothetic separability and weak separability of the indirect utility function. We illustrate the practical usefulness of the approach by an empirical application to Spanish household consumption data.

> *The universe cannot be dealt with in one stroke and so a bit has to be broken off and treated as if the rest did not matter.*
> [Afriat, 1969]

1. Introduction

We focus on the revealed preference conditions for consistency of a finite data set with the maximization of a weakly separable utility function. Our main contribution is twofold.[1] First, we show that verification of these revealed preference conditions is a difficult problem. In particular, the problem is NP-complete, which essentially means that it cannot be solved in polynomial time. As we will discuss below, this actually motivates our second contribution. Specifically, we show that the revealed preference conditions can be verified by means of elementary integer programming procedures, which are easily implemented in practice. We demonstrate the versatility of this integer programming approach by showing that it can also assess homothetic separability and weak separability of the indirect utility function. We illustrate our approach by applying it to a Spanish panel data set.

Demand analysis is a powerful tool for the measurement of the behavioral and distributional effects of counterfactual price and income changes. For example tax changes alter the relative prices faced by the consumer. How the consumer reacts to this, by choosing an alternative consumption bundle, is the subject of the analysis of demand behavior. Typically the researcher estimates the unknown parameters of a parametric demand system and uses these estimates to calculate pre- and post-reform demands and its corresponding welfare implications (see, for example, Banks et al. (1997)).

In this respect, weak separability of the utility function is a frequently used assumption in both theoretical and applied demand analyses. A group of goods is said to be weakly separable if the marginal rate of substitution between any two goods in the group is independent from the quantities consumed of any good outside

☆ We thank the Editor, the Associate Editor and two anonymous referees for useful comments, which substantially improved the paper. We are also grateful to Adrian Fleissig, James Swofford and Gerald Whitney for helpful comments. Ian Crawford generously provided us the data of the Encuesta Continua de Presupuestos Familiares (ECPF), which we use in our empirical study.
* Corresponding author.
E-mail addresses: laurens.cherchye@kuleuven-kortrijk.be (L. Cherchye), t.demuynck@maastrichtuniversity.nl (T. Demuynck), bderock@ulb.ac.be (B. De Rock), per.hjertstrand@ifn.se (P. Hjertstrand).
[1] This paper is a merger of two working papers, Cherchye et al. (2011d) and Hjertstrand (2011).

http://dx.doi.org/10.1016/j.jeconom.2014.07.001
0304-4076/

this group (Leontief, 1947; Sono, 1961). Weak separability has several convenient implications.[2] First of all, it allows for representing consumption in terms of two stage budgeting. This means that, in order to determine the demanded quantities of the goods in the separable group, it suffices to know the prices of the goods in this group and the total within-group expenditure. Further, weak separability is a crucial condition for the construction of group price and quantity indices. Such aggregates can be useful, for example, to compute group cost of living indices to be used in the welfare analysis. Finally, from an empirical point of view, weak separability significantly reduces the number of parameters of the demand system to be estimated in practical applications.[3]

Considering these advantages for both theoretical and empirical work, an important issue concerns empirically testing the validity of the separability assumption (prior to effectively imposing it). In the literature, there are two approaches to test for weak separability. One approach uses econometric techniques to verify certain parameter restrictions given a specific demand model. Although this approach is fairly flexible in terms of the demand model that is used, it also poses a number of problems.

First, the separability restriction is often tested using the Wald or likelihood ratio test procedures, which require estimating the full (unrestricted) demand model. Consequently, these tests may suffer from a degrees of freedom problem in the sense that too many parameters must be estimated given the amount of data.[4] Next, if the hypothesis of weak separability is rejected, it is impossible to verify whether this implies a rejection of weak separability as such or, instead, a rejection of the specific functional form imposed on demand a priori. In other words, if the null hypothesis of weak separability is rejected, this may well be due to the use of a wrong functional form rather than a non-separable utility structure per se.[5] Finally, most econometric tests for separability are based on separability of the indirect utility function (i.e. separability in prices), which does not imply separability of the direct utility function (i.e. separability in quantities).[6]

An alternative approach to test for weak separability is based on revealed preference theory. In several seminal contributions to the literature, Afriat (1969), Varian (1983) and Diewert and Parkan (1985) developed revealed preference conditions that characterize the collection of data sets that are rationalizable by a (weakly) separable utility function.[7] The revealed preference approach remedies different problems associated with the econometric approach. First, the revealed preference conditions can meaningfully be applied to data sets with as few as two observations, which avoid the degrees of freedom problem discussed above. Further, the revealed preference approach abstains from imposing a specific functional form on the utility functions. As such, the tests are insensitive to model misspecification. Finally, the revealed preference approach does not require additional assumptions like homotheticity of the subutility function or separability of the indirect utility function (although such additional assumptions can be imposed and tested; see below).

Unfortunately, the revealed preference conditions have the drawback that they take the form of a set of nonlinear, quadratic inequalities, which are very hard to verify. In order to avoid this problem, several heuristics have been proposed that provide separate sufficient and necessary conditions for data consistency with weak separability (see Section 2 for an overview). The lack of an efficient algorithm to verify the revealed preference conditions raises the question whether such an algorithm exists at all. In this study, we show that the answer is no. In particular, we prove that the verification of the revealed preference conditions for weak separability is an NP-complete problem. This NP-completeness result implies that it is impossible to find a polynomial time algorithm that verifies whether a data set is consistent with the maximization of a weakly separable utility function (unless one can prove P = NP). This indicates that we should better look for a widely applied and (for moderately sized problems) reasonably quick non-polynomial time algorithm to verify the revealed preference conditions. Given this, we present an easy-to-implement (non-polynomial time) integer programming procedure to verify the revealed preference conditions. Our approach exploits the equivalence of the generalized axiom of revealed preference (GARP) and a set of mixed integer inequalities. Such an integer programming approach has proven very useful in the literature that applies revealed preference theory to collective consumption models, which studies the behavior of multi-person households, and in the literature that investigates the testable implications of general equilibrium models.[8] We extend the insights from this literature to the model of utility maximization with a weakly separable utility function. This provides a further demonstration of the usefulness of integer programming techniques to deal with computational issues related to the practical implementation of revealed preference conditions.

From a theoretical point of view, the core motivation for adopting an integer programming approach is that this is a widely accepted and a well known approach to handle NP-complete problems. Besides this, we also have a number of other motivations. First of all, our approach can be applied to data sets with any number of observations. Second, any mixed integer program provides an exact solution in finite time. Hence, our approach provides a way to verify the necessary and sufficient conditions for a given data set to be consistent with maximization of a weakly separable utility function. Until now, all existing heuristics either provide a necessary or a sufficient condition. Our algorithm is the first to provide an exact solution. A third important argument pro

2 See also Deaton and Muellbauer (1980) for a more thorough discussion.

3 In this respect, it is important to point out the importance of employing a correct separability structure in empirical demand modeling. On the one hand, using a too narrow structure (i.e. omitting goods that should be included in the separable grouping) leads to an omitted variables problem, which consequently produces inconsistent parameter estimates in the estimated demand model. On the other hand, including redundant goods in the separability structure may inflate the variances of the parameters, which causes inefficient parameter estimates.

4 The degrees of freedom problem could in theory be circumvented instead by using Lagrange multiplier tests. However, similar to Wald tests, Lagrange multiplier tests require a consistent estimate of the covariance matrix. Although it is relatively easy to obtain such estimates, these are often biased in small samples, implying that the Lagrange multiplier test may suffer from a small sample bias.

5 Imposing separability conditions on a particular functional form might lead to additional difficulties. In particular, Blackorby et al. (1978) showed that testing for separability using several econometric specifications based on local approximations of the true model (i.e. flexible functional forms) is actually equivalent to testing a much stronger condition. For example, it turns out impossible to test separability for the translog model without imposing the much more stringent assumption of additive separability.Barnett and Choi (1989) confirmed this result by means of Monte Carlo simulations.

6 A well known sufficient condition to obtain that the direct and indirect separabilities coincide is that the subutility function is homothetic. We refer to Blackorby and Russell (1994) for more discussion.

7 The revealed preference conditions for weak separability have been used in many different types of applications. See, for example, Swofford and Whitney (1987, 1988, 1994), Barnhart and Whitney (1988), Patterson (1991), Belongia and Chrystal (1991), Choi and Sosin (1992), Jones and Mazzi (1996), Cox (1997), Fisher and Fleissig (1997), Rickertsen (1998), Spencer (2002), Fleissig and Whitney (2003), Fleissig and Whitney (2008), Swofford (2005), Serletis and Rangel-Ruiz (2005), Jones et al. (2005), Jha and Longjam (2006), Blundell et al. (2007), Hjertstrand (2007, 2009), Elger et al. (2008), Elger and Jones (2008), and Drake and Fleissig (2008).

8 See Cherchye et al. (2007, 2008, 2009, 2011a,b), for integer programming characterizations of household consumption models, Cherchye et al. (2011c) for integer programming characterizations of general equilibrium models and Cherchye et al. (2013) for an integer programming characterization for the transferable utility model.