Journal of Econometrics 185 (2015) 259–282

Contents lists available at ScienceDirect

Journal of Econometrics

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

Specification tests for partially identified models defined by moment inequalities☆

Federico A. Bugni [a,*], Ivan A. Canay [b], Xiaoxia Shi [c]

[a] Department of Economics, Duke University, 213 Social Sciences Road, Box 90097, Durham, NC, 27708, USA
[b] Department of Economics, Northwestern University, 2001 Sheridan Road, Evanston, IL, 60208, USA
[c] Department of Economics, University of Wisconsin, Madison, 1180 Observatory Drive, Madison, WI, 53706, USA

ARTICLE INFO

Article history:
Received 25 July 2013
Received in revised form
4 September 2014
Accepted 6 October 2014
Available online 18 November 2014

JEL classification:
C01
C12
C15

Keywords:
Partial identification
Moment inequalities
Specification tests
Hypothesis testing

ABSTRACT

This paper studies the problem of specification testing in partially identified models defined by moment (in)equalities. This problem has not been directly addressed in the literature, although several papers have suggested a test based on checking whether confidence sets for the parameters of interest are empty or not, referred to as Test BP. We propose two new specification tests, denoted Test RS and Test RC, that achieve uniform asymptotic size control and dominate Test BP in terms of power in any finite sample and in the asymptotic limit.

1. Introduction

This paper studies the problem of specification testing in partially identified models defined by a finite number of moment equalities and inequalities (henceforth, referred to as (in)equalities). The model can be written as follows. For a parameter vector (θ, F), where $\theta \in \Theta$ is a finite dimensional parameter of interest and F denotes the distribution of the observed data, the model states that

$$E_F[m_j(W_i, \theta)] \geq 0 \quad \text{for } j = 1, \ldots, p,$$
$$E_F[m_j(W_i, \theta)] = 0 \quad \text{for } j = p+1, \ldots, k, \tag{1.1}$$

where $\{W_i\}_{i=1}^n$ is an i.i.d. sequence of random variables with distribution F and $m : \mathbb{R}^d \times \Theta \to \mathbb{R}^k$ is a known measurable function. This model is *partially identified* because the sampling process and the maintained assumptions (that is, Eq. (1.1) together with regularity conditions) restrict the value of the parameter of interest θ to a set, called the *identified set*, which is smaller than Θ but potentially larger than a single point.

The model is said to be *correctly specified* (or statistically adequate) when the moment (in)equalities hold for at least one parameter value, i.e., when the identified set is non-empty.[1] A specification test takes correct specification of the model as the null hypothesis and rejects if the data seem to be inconsistent

☆ We thank Han Hong and two anonymous referees whose valuable suggestions helped greatly improve the paper. We also thank Adam Rosen, the participants at the 2011 Triangle Econometrics Conference, the 2012 Junior Festival on New Developments in Microeconometrics at Northwestern University, the 2012 Summer Meetings of the Econometrics Society, the Workshop on Inference in Nonstandard Problems at Princeton University, the Workshop on Mathematical Statistics of Partially Identified at Oberwolfach, and seminar participants at IUPUI, UC Berkeley, UC Davis, Yale, Columbia, Rochester, Ohio State, UT Austin, Wisconsin, UC San Diego, and Stanford for helpful comments. Bugni and Canay thank the National Science Foundation for research support via grants SES-1123771 and SES-1123586, respectively. Takuya Ura provided excellent research assistance. Any and all errors are our own.

* Corresponding author.
E-mail addresses: federico.bugni@duke.edu (F.A. Bugni), iacanay@northwestern.edu (I.A. Canay), xshi@ssc.wisc.edu (X. Shi).

http://dx.doi.org/10.1016/j.jeconom.2014.10.013
0304-4076/

[1] The concept of statistical adequacy was introduced by Koopmans (1937) and referred to as the Fisher's axiom of correct specification. The discussion of the importance of a correct specification for inference purposes dates back to Haavelmo (1944).

with it. Specification tests for partially identified models have been studied by a small number of authors (reviewed below), but the only existent test applicable to the general specification of Eq. (1.1) is the one based on checking whether a confidence set for θ is empty or not. We refer to this procedure as "Test BP", to emphasize that it is a *by-product* of confidence sets for θ, and describe it formally in the next section.

In this paper, we propose two new specification tests for the model above and show that they have the following properties. First, our tests achieve uniform size control, just like Test BP. Second, our tests dominate Test BP in terms of power in any finite sample and in the asymptotic limit. Specifically, our tests have more or equal power than Test BP in all finite samples, and there are sequences of local alternative hypotheses for which our tests have strictly higher asymptotic power.

Both of our tests use the same "infimum" test statistic $\inf_{\theta\in\Theta} Q_n(\theta)$, where $Q_n(\theta)$ is the criterion function typically used to construct confidence sets for θ, much in the spirit of the popular J-test in (point-identified) GMM models (see Remark 4.1). The difference between them lies in the critical value used to implement the test. Computing one of these critical values requires little additional work beyond the computation involved in the confidence set construction, just like in Test BP. We therefore always recommend the use of this test, as it attains better power at almost no additional cost. On the other hand, our second test has even better power, but it requires a separate resampling procedure to implement. For this reason, we recommend its use when one has serious interest in the statistical adequacy of the model.[2]

From a methodological point of view, there are two aspects of our paper worth highlighting. First, we derive the limiting distribution of the "infimum" test statistic under drifting sequences of data distributions and provide two methods to approximate its quantiles. To the best of our knowledge, we are the first ones to obtain these kinds of results in partially identified moment (in)equality models. These methodological contributions are relevant in problems that go well beyond specification testing. For example, Bugni et al. (2014) show that hypothesis tests based on the "infimum" test statistic can be adapted to address a large class of interesting new problems, which includes inference on a particular coordinate of a multivariate parameter θ. Second, the asymptotic framework we use is one where the tuning parameter κ_n that determines if a moment inequality is binding, diverges to infinity at an appropriate rate, c.f. Andrews and Soares (2010). In this framework, the arguably best possible implementation of Test BP is the one we use, see Definition 2.4. Recent contributions to the literature have used an alternative asymptotic framework where this tuning parameter κ_n converges to a constant $\kappa < \infty$ that affects the limiting distribution, see Andrews and Barwick (2012), Romano et al. (2014), McCloskey (2014). One could potentially use these methods to define another version of Test BP, and then study the behavior of our tests using fixed-κ asymptotics. We do not pursue this strategy as it involves technical tools that are well beyond those developed here.[3]

The motivation behind our interest in misspecified models stems from the view that most econometric models are only approximations to the underlying phenomenon of interest. This is also the case for partially identified models, where strong and usually unrealistic assumptions are replaced by weaker and more credible ones (see, e.g., Manski, 1989, 2003). In other words, the partial identification approach to inference allows the researcher to conduct inference on the parameter of interest without imposing assumptions on certain fundamental aspects of the model, typically related to the behavior of economic agents. Still, for computational or analytical convenience, the researcher has to impose certain other assumptions, that are typically related to functional forms or distributional assumptions.[4] If these assumptions are not supported by the data, and so the model is misspecified, the resulting statistical inferences are usually invalid (see, e.g., Ponomareva and Tamer, 2011; Bugni et al., 2012).

Specification tests for partially identified models have been studied in Guggenberger et al. (2008), Romano and Shaikh (2008), Andrews and Guggenberger (2009), Andrews and Soares (2010), Santos (2012). Guggenberger et al. (2008) propose to transform a linear moment (in)equality model into a dual form that does not involve parameters and, in this way, eliminate the partial-identification problem. Innovative as it is, their approach only applies to linear models and is not practical when the dimension of the parameter is large because the dimension of the dual form grows exponentially with the dimension of the parameters. Santos (2012) defines specification tests in a partially identified nonparametric instrumental variable model and, thus, his results are not directly applicable to the model in Eq. (1.1). To the best of our knowledge, the only valid specification test for the model in Eq. (1.1) that has been described in the literature is Test BP. This specification test has been proposed by Romano and Shaikh (2008, Remark 3.7), Andrews and Guggenberger (2009, Section 7), and Andrews and Soares (2010, Section 5).[5]

It is worth mentioning that the specification tests we propose in this paper are a type of omnibus tests, in the sense that the specific structure of certain nonparametric alternatives is unknown. However, a partially identified model is typically the result of removing undesirable restrictions in a certain point identified model. As a consequence, refuting the partially identified model leaves the researcher with a reduced set of assumptions that could potentially be wrong. In addition, in some cases testing the specification of a partially identified model can be analogous to directly testing an interesting economic behavior. For example, Kitamura and Stoye (2012) recently proposed a specification test for the Axiom of Revealed Stochastic Preference that shares similarities to our specification tests. In their case, rejecting the specification of the model through their non-parametric test directly means rejection of the Axiom of Revealed Stochastic Preferences. We note, however, that there are substantial differences between our approach and that in Kitamura and Stoye (2012) in terms of the nature of the model, the construction of the test statistic, and the range of applications in which each of these tests can be applied.

The rest of the paper is organized as follows. Section 2 introduces the basic notation we use in our formal analysis and describes the aforementioned Test BP. The tests proposed in this paper compare a test statistic with a critical value. Section 3 introduces our test statistic. The description of our tests is then completed by introducing appropriate critical values that are presented in the succeeding sections. Section 4 describes a critical value based on the asymptotic approximation or bootstrap approximation of the limiting distribution of the test statistic. We call this test the *re-sampling* test or "Test RS". Section 5 describes a critical value that is based on recycling critical values that have already been

[2] It is worth pointing out that a version of our second test has been used in Gandhi et al. (2013), with $p = 401$ and a parameter θ with more than 20 coordinates, which illustrates the feasibility of this test in real scale applications.

[3] For example, all tests would suffer from asymptotic size distortion and size correction would be needed.

[4] See Manski (2003) and Tamer (2003) for a discussion on the role of different assumptions and partial identification.

[5] It is important to clarify that Test BP was conceived by papers whose main objective was the construction of confidence sets and not the design of a specification test. In addition, Test BP has some robustness properties, see Remark 6.7.