ELSEVIER

Journal of Econometrics 109 (2002) 167–193

JOURNAL OF Econometrics

www.elsevier.com/locate/econbase

Bootstrap J tests of nonnested linear regression models

Russell Davidson[a,b], James G. MacKinnon[b,*]

[a] GREQAM, Centre de la Vieille Charité, 2 rue de la Charité, 13002 Marseille, France
[b] Department of Economics, Queen's University, Dunning Hall, Kingston, ON, Canada K7L 3N6

Received 6 December 2000; received in revised form 20 August 2001; accepted 13 November 2001

Abstract

The J test for nonnested regression models often overrejects very severely as an asymptotic test. We provide a theoretical analysis which explains why and when it performs badly. This analysis implies that, except in certain extreme cases, the J test will perform very well when bootstrapped. Using several methods to speed up the simulations, we obtain extremely accurate Monte Carlo results on the finite-sample performance of the bootstrapped J test. These results fully support the predictions of our theoretical analysis, even in contexts where the analysis is not strictly applicable. © 2002 Elsevier Science B.V. All rights reserved.

JEL classification: C12; C15

Keywords: Bootstrap test; Nonnested hypotheses; Simulation; Linear regression

1. Introduction

Numerous procedures for testing nonnested regression models have been developed, directly or indirectly, from the pathbreaking work of Cox (1961, 1962). The most widely used, because of its simplicity, is the J test proposed in Davidson and MacKinnon (1981); see McAleer (1995) for evidence on this point. Like almost all nonnested hypothesis tests, the J test is not exact in finite samples. Indeed, as many Monte Carlo experiments have shown, its finite-sample distribution can be very far from the N(0, 1) distribution that it follows asymptotically.

Several ways have been proposed to improve the finite-sample properties of the J test. Fisher and McAleer (1981) proposed a variant, called the J_A test, which is exact in finite samples under the usual conditions for t tests in linear regression models to be

* Corresponding author. Tel.: +1-613-533-2293; fax: +1-613-533-6668.
E-mail addresses: russell@ehess.cnrs-mrs.fr (R. Davidson), jgm@qed.econ.queensu.ca (J.G. MacKinnon).

PII: S0304-4076(01)00146-4

exact; see Godfrey (1983). Unfortunately, the J_A test is often very much less powerful than other nonnested tests; see, among others, Davidson and MacKinnon (1982) and Godfrey and Pesaran (1983). The latter paper suggested a different approach, applied not to the J test but to variants of the Cox test based on the work of Pesaran (1974). This approach first corrects the bias in the numerator of the test statistic, then estimates the variance of the corrected numerator, and finally calculates a t-like statistic. It does not yield exact tests, but it does yield tests that perform considerably better than the J test under the null and have good power.

More recently, Fan and Li (1995) and Godfrey (1998) have suggested bootstrapping the J test and other nonnested hypothesis tests. Because the J test is cheap and easy to compute, this is very easy to do. The Monte Carlo results in these papers suggest that bootstrapping the J test often works very well. However, neither paper provides any theoretical explanation of why it does so.

In this paper, we develop a theoretical approach that enables us to show precisely what determines the finite-sample distribution of the J test. We explain why it often works very badly without bootstrapping and why it almost always works very well indeed when bootstrapped. The theory allows us to identify situations in which the tests can be expected to achieve their worst behavior, and our Monte Carlo experiments focus on these. Since the tests perform very well even in such situations, the experiments need to be very accurate. Fortunately, our theory provides a low-cost way to perform experiments that use extremely large numbers of replications.

The assumptions needed for our theoretical analysis are fairly restrictive: The errors are assumed to be normally distributed, and the regressors are assumed to be exogenous. However, additional Monte Carlo experiments strongly suggest that these assumptions are not crucial. Even when both of them are violated, the bootstrap J test performs in almost exactly the same way as it does when they are satisfied.

In the next section, we briefly describe the J test. In Section 3, we derive a theoretical expression for the test statistic and use it to obtain a number of interesting results. In Section 4, we use a combination of theory and simulation to study the finite-sample properties of the asymptotic J test. In Section 5, we study the finite-sample properties of the bootstrap J test. In Section 6, we relax the restrictive assumptions made up to this point and show that the bootstrap J test works extraordinarily well in almost every case in which a nonnested test is worth doing. Finally, in Section 7, we briefly discuss the effect of bootstrapping on the power of the J test.

2. The *J* test

Although the J test can be applied to both linear and nonlinear regression models, we restrict our attention to the linear case, since it would be extremely difficult to obtain general results about the finite-sample properties of the J test in the nonlinear case. Consider two nonnested, linear regression models with IID normal errors:

$$\mathrm{H}_1\colon \boldsymbol{y} = \boldsymbol{X\beta} + \boldsymbol{u}, \quad \boldsymbol{u} \sim \mathrm{N}(\mathbf{0}, \sigma^2\mathbf{I})$$