Journal of Econometrics 158 (2010) 130–141

Contents lists available at ScienceDirect

Journal of Econometrics

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

Speed of adjustment in cointegrated systems

Luca Fanelli[a], Paolo Paruolo[b,*]

[a] Department of Statistical Sciences, University of Bologna, Italy

[b] European Commission Joint Research Centre & Department of Economics, University of Insubria, Varese, Italy

ARTICLE INFO

Article history:
Available online 12 March 2010

JEL classification:
C32
C52
F31

Keywords:
Half-life
π-life
Speed of adjustment
Impact factors
Vector equilibrium correction
Interim and total multipliers

ABSTRACT

This paper discusses summary measures for the speed of adjustment in possibly cointegrated Vector Autoregressive Processes (VAR). In particular we propose long-run half-lives, based on interim and total multipliers. We discuss their relation with Granger-noncausality and other types of half-life, which are shown to convey different information, except in the univariate AR(1) case. We present likelihood-based inference on long-run half-lives, regarded as discrete functions of parameters in the VAR model. It is shown how asymptotic confidence regions can be defined. An empirical illustration concerning speed of adjustment to purchasing-power parity is provided.

1. Introduction

Economic theory often provides predictions on the stability of equilibrium relations, as well as on some degree of mean reversion or speed of adjustment. This is the case, for instance, in studies on aggregate consumption, where fast adjustment of consumption is interpreted by Morley (2007) as an indication of the validity of the permanent income hypothesis, whereas slow adjustment is associated with habit formation and precautionary savings. Similarly, in studies on purchasing power parity (PPP), validity of PPP is linked to a sufficient degree of mean reversion for the real exchange rate, see e.g. Kilian and Zha (2002), Murray and Papell (2002) and Rossi (2005).

Most empirical evidence on the speed of adjustment in these fields has been based on univariate time-series models and the associated measures of half-life. The half-life is usually defined as the number of periods required for the impulse response (IR) to a unit shock to dissipate by one half, see Kilian and Zha (2002) page 199 and references therein.

Multivariate half-lives have been recently discussed by van Dijk et al. (2007) in the context of nonlinear systems; other papers which apply the notion of half-life in VAR models are Crowder (2004), Cheung et al. (2004) and Malliaropulos et al. (2004). We follow the definition of half-life given in van Dijk et al. (2007), which is here generalized by allowing for different choices of band scale.

The present definition is general and applies to possibly nonlinear processes and to any indicator, not necessarily equal the IR. Examples include forecasting-ability indicators, such as the general predictability measures of Diebold and Kilian (2001), as well as variance-decomposition and impulse-response indicators, such as generalized impulse responses (GIR) and persistence profiles, see Pesaran and Shin (1996, 1998), or cumulative impulse responses (CIR), see Andrews and Chen (1994).

We focus on half-lives for possibly cointegrated VAR processes; confining our attention to VARs, we are able to discuss concrete examples of different notions of half-life. We compare the standard half-life with a long-run half-life, defined as the number of periods required for the CIR to converge to its terminal value, with an approximation of plus or minus half the terminal value. Values of the CIR at different horizons correspond to interim multipliers, while the terminal value of the CIR is the total multiplier.[1] The choice of band scale for long-run half-life is given by (the absolute value of) the terminal value of the CIR.

* Corresponding author. Fax: +39 0332 215509.
E-mail addresses: luca.fanelli@unibo.it (L. Fanelli), paolo.paruolo@uninsubria.it (P. Paruolo).

[1] Crowder (2004), Cheung et al. (2004) and Malliaropulos et al. (2004) calculate half-lives from IR. van Dijk et al. (2007) consider half-lives for GIR. All these notions correspond to a counterfactual experiment similar to the one of the IR, i.e. an impulse perturbation at time t. The counterfactual experiment for the CIR is given, instead, by step perturbation at time t, i.e. a unit perturbation from time t onwards. For this reason the CIR is called the step-response in engineering literature.

0304-4076/$ – see front matter © 2010 Elsevier B.V. All rights reserved.
doi:10.1016/j.jeconom.2010.03.020

The notion of half-life is related to the literature on mean and median lags in distributed lags models, see e.g. Hendry et al. (1984) p. 1051. However, while mean or median lags are defined only for non-negative distributed-lag coefficients, the concept of half-life discussed in this paper applies in general to any dynamic response.

The long-run half-life has a special interpretation for VAR systems integrated of order 1, I(1). The CIR in fact describes the effects of an impulse on the level of the variables. The terminal value of the CIR is called the impact factor by Omtzigt and Paruolo (2005) (henceforth OP), who provide a detailed discussion of its interpretation. In this prominent, albeit special case, the long-run half-life can be interpreted as the number of periods required for an impulse to exert half its total effect on the level of the variables. The present setup also accommodates the type of counterfactuals employed by Johansen (2005) for the interpretation of cointegration coefficients. Long-run half-lives can also be computed more in general for VARs integrated of any nonnegative order.

Long-run half-lives provide a measure of the speed of adjustment within a model that is consistent with all other model features. This allows one to analyze several economic-theory predictions simultaneously: one can test if long-run equilibria exist in the form of cointegrating relations, measure long-run effects and distinguish which variable is adjusting and how fast to which perturbation, all within the same model framework. We note that the same system can be characterized by different speeds of adjustment for different combinations of variables and perturbations; hence one should talk about speeds of convergence also within a single VAR system.

One area of application of half-lives is given by the evaluation of policy interventions. Finding, for instance, that adjustment is very slow implies that a substantial part of the effects of some economic policy will hardly materialize within the typical time-frame of a policy-maker. For example Giannini (1983) discussed the estimated consumption function of Davidson et al. (1978), finding that, after a shock to income, the time required before consumption reached 90% of its steady-state value was between 107 and 108 quarters, i.e. approximately 27 years.[2] The speed of adjustment hence plays an important role in determining the effectiveness of policy interventions within a given time-frame; the long-run measures of speed presented in this paper can be used to evaluate this type of adjustment.

The policy-evaluation application suggests generalizing the notion of half-life to the one of π-life, where π is any fraction $0 < \pi < 1$, possibly different from $\frac{1}{2}$, as proposed by van Dijk et al. (2007). The π-life consists of the number of periods required by some indicator to enter the band $[-\pi s, \pi s]$ definitively, where s is the band scale. For instance, in order to define the time before 90% of a long-run value is realized, one can set $\pi = 0.1$ in the long-run π-life. In case of non-monotone impulse responses, this definition records the last time the indicator enters the band. The definition of π-life contains half-lives as a special case; different choices of π can be envisaged in different contexts. For the derivations of the present paper, π can be taken equal to $\frac{1}{2}$ in what follows without loss of generality.

The present notion allows one to answer the question 'how fast does variable y adjust to perturbations of variable x?'. We translate the original question into two related ones, namely (a) 'what is the long-run effect of perturbations of variable x on variable y?' and (b) 'how fast is this long-run effect accomplished?'. In case of zero long-run effect in (a), the answer to the original question is that there is no adjustment, and it may be of little interest, questionable, and difficult to measure how fast this zero long-run effect is accomplished.

When the long-run effect is not equal to zero, we address question (b) with the concept of long-run π-life, which measures speed using the long-run effect as the scale s of the band. The combination of the answers to questions (a) and (b) provides a richer answer to the original question than a single measure of speed, as we illustrate in Section 5 below.

By comparison and for ease of exposition, we indicate the π-life based on IR as 'short-run π-life'. The short-run π-life and a long-run π-life are related but different times; in general they give different results, except for the univariate AR(1) process.

The long-run π-lives have direct connections to long-run Granger noncausality as defined in Dufour and Renault (1998) and Dufour et al. (2006). We show that long-run Granger noncausality implies a zero total multiplier, but not vice versa. This suggests to control if the total multiplier is significantly different from zero before estimation of long-run π-lives; this is the approach we take in the illustrative empirical example below.

The present definition of long-run π-life increments the available tools to evaluate the adjustment mechanism in cointegrated systems. For these systems the coefficients of the disequilibrium errors in the Vector Equilibrium Correction (VEC) are usually interpreted as indicators of speed of adjustment to the long-run. If these coefficients are significant, this is interpreted as implying presence of adjustment; their size (absolute value) is then associated with the strength of the speed of adjustment, large coefficients implying faster adjustment. This interpretation is motivated by 1-step ahead predictions for VAR processes of order 1, but it is simple to see that it is not applicable in complete generality. On the contrary, the present definition of long-run π-lives can be used to measure speed of adjustment to disequilibrium errors in general, and hence it fills a gap in this part of the literature.

We present likelihood-based inference on the long-run π-lives, which places the present analysis within the framework described in Johansen (1996). We derive point estimators and asymptotic confidence sets for the long-run π-lives. We exploit the fact that long-run π-lives are (discrete) functions of the autoregressive coefficients to derive (discrete) confidence sets for them from the confidence set of the autoregressive coefficients. This construction allows one to consider confidence sets for several π-lives simultaneously, without affecting the asymptotic joint global coverage probability. This approach does not require resampling, it is general and it can be applied to other half-lives, such as short-run ones.

The concept of long-run π-life can also be applied to nonlinear multivariate systems, such as the ones considered in Koop et al. (1996), Potter (2000), and by van Dijk et al. (2007). In the nonlinear case, however, the impulse responses are not time-invariant, and hence inference is more complicated than in the linear case covered by the present paper. Within the VAR case, other choices of indicators also lead to different functions of the parameters, and hence require different calculations with respect to the ones reported in this paper. However, the methods presented here can be directly extended to those situations with minor additional work.

The rest of the paper is organized as follows. Section 2 presents the definition of π-lives and discusses several special cases, including long-run and short-run π-lives. Section 3 discusses long-run π-lives for cointegrated systems of order 1. Section 4 presents likelihood-based inference on long-run π-lives. Section 5 reports a brief illustration on PPP and Section 6 concludes. Proofs are reported in the Appendix.

2. Definition and properties of π-lives

In this section we introduce notation and the general definition of π-life, following van Dijk et al. (2007). We assume that the

[2] This contrasts with the rather fast time to reach 50% of its steady state value, which took about 1 year, implying a "first quick — then very slow speed", see page 6 in Giannini (1983).