Macroeconomic Dynamics, **13** (2009), 580–604. Printed in the United States of America.
doi:10.1017/S1365100509080249

TECHNOLOGY SHOCKS AND HOURS WORKED: A FRACTIONAL INTEGRATION PERSPECTIVE

LUIS ALBERIKO GIL-ALANA AND ANTONIO MORENO
University of Navarra

Previous research has found that the dynamic response of hours worked to a technology shock crucially depends on whether the hours variable is assumed to be an $I(0)$ or an $I(1)$ variable ex ante. In this paper we employ a multivariate fractionally integrated model that allows us to simultaneously estimate the order of integration of hours worked and its dynamic response to a technology shock. Our evidence lends support to the hypothesis that hours fall in response to a positive technology shock.

Keywords: Technology Shocks, Impulse Response Functions, Hours Worked, Fractional Integration, Multivariate Analysis.

1. INTRODUCTION

What is the effect of a technology shock on the number of hours worked by the labor force at business cycle frequencies? This question lies at the heart of modern macroeconomics. The reason is that technological innovation is a perennial candidate as a source of dynamics for the aggregate economy and employment. At the same time, full employment remains the main goal for policy makers. The question that naturally arises in this context is then, How compatible are these two, at least in the short run? A host of macroeconomic models have tackled this issue from both theoretical and empirical perspectives, but no consensus has emerged in the literature yet. This paper uses fractional integration techniques to give an answer to this important question.

Authors disagree on the empirical implications of a technology shock on hours worked per capita. Galí (1999) ignited an empirical literature on the issue when he contradicted the tenets of the real business cycle (RBC) theory, whereby technology shocks are key for business cycle dynamics. Galí (1999) and later on Galí and Rabanal (2004; henceforth GR) and Francis and Ramey (2005a) showed not only that technology shocks were unimportant for business cycle fluctuations but that, contrary to the implications of RBC models, hours worked declined in response

Comments of the Editor and of two anonymous referees are gratefully acknowledged. We thank seminar participants in the Bank of Spain, Yonsei University, and the Bank of Sweden. Luis Alberiko Gil-Alana acknowledges financial support from ECO2008-03035 ECON Y FINANZAS, Spain. Antonio Moreno acknowledges financial support from the CICYT Grant SEJ2005-06302. Address correspondence to: Antonio Moreno, School of Economics, University of Navarra, 31080 Pamplona, Spain; e-mail: antmoreno@unav.es.

 1365-1005/09

to a technology shock. Galí's results have recently been challenged by Christiano et al. (2003) (henceforth CEV). In an empirical framework very similar to that of Galí, these authors find that hours actually increase after a technology shock. The crucial difference between the two sets of studies is that CEV treat the variable hours as stationary $I(0)$, whereas the former authors, such as GR, treat it as a nonstationary $I(1)$ variable. From this perspective, the main issue that remains to be determined is the exact order of integration of hours worked per capita: Is it one or zero? Using an asymptotic local-to-unity approach, Pesavento and Rossi (2005) also propose an agnostic method and find that a positive productivity shock has a negative impact effect on hours. In their framework, the researcher does not have to choose between levels and first differences in hours worked.

The main contribution of the present paper is to show that there is an alternative way to resolve the technology-hours issue without assuming a given order of integration for standard measures of hours ex ante. We derive a simple method in a fractional integration framework, which lets the data simultaneously determine the response of hours to a technology shock and the order of integration of hours worked. This method presents some advantages with respect to previous approaches. First, the fractional integration approach makes it possible to discern the order of integration of a given variable without restricting the econometrician to choose a priori between one and zero. The order of integration could be zero, a fraction of one, one, or even above one. Thus, our approach is agnostic with respect to the order of integration of the variables before their inclusion in a vector autoregressive (VAR) framework. As a result, pretests on the order of integration of the variables are not required. As previously mentioned, Pesavento and Rossi (2005) already addressed this problem in the context of a local-to-unity approach. Our approach is, however, essentially different from theirs in that we allow the order of integration to be any real number, whereas they remain in the $I(0)/I(1)$ framework.

Second, we show that, in our fractional setting, the implied impulse response functions are invariant to estimating the model with the variables in either levels or first differences. Because the order of integration is estimated from the data, the responses of the series in first differences are exactly the same as those implied by the variables in levels by construction. Moreover, the multivariate fractionally integrated model employed in this paper permits us to identify the structural impulse response functions in a way similar to the classic VAR systems with the additional interaction of the binomial expansions implied by the fractional polynomials involved in the model.

For our two data specifications, we find that hours worked decline on impact in response to a technology shock. This response is statistically significant using the CEV measure (total business hours worked per capita) but not under the GR measure (nonfarm business hours worked per capita). We also find that the orders of integration of hours worked identified by the more general fractionally integrated multivariate systems are uniformly lower than for their univariate counterparts. Whereas all the univariate frameworks point at orders of integration of hours close to 1 or even larger, multivariate models that allow richer and more realistic dynamics identify orders of integration lower than 1. Thus, it seems that the cross-sectional dependence allowed in the multivariate framework reduces the degree of persistence in aggregate hours worked. Finally, our multivariate model implies statistically different orders of integration for hours worked across data specifications. Whereas the variable used by GR has an order of integration of 0.67, the order of integration of the hours variable used by CEV is slightly above zero, though with a higher level of dependence with respect to the short-run structure.

Section 2 revisits the controversial issue in hand, the divergence between the responses of hours worked to a technology shock depending on the assumed order of integration of hours worked. Section 3 develops our econometric framework, intended to simultaneously identify the orders of integration of the macroeconomic variables and the impulse responses to the structural shocks. It also compares our framework with the local-to-unity approach. Section 4 performs univariate tests for the orders of integration of productivity and hours from a fractionally integrated perspective. Section 5 employs the multivariate fractionally integrated model derived in Section 3 to determine the response of hours worked to a technology shock across data specifications. Section 6 concludes.

2. THE CONTROVERSY

In this section we revisit the empirical evidence regarding the effect of a technology shock on hours worked. We first describe the data used throughout the paper. Then we report the impulse responses for both GR and CEV's specifications and comment on the differences across responses.

Both GR and CEV work with quarterly data, which is commonplace in the business cycle literature. Whereas GR use productivity and hours data from the nonfarm business sector in their bivariate VARs, CEV use data from all businesses, including farming activities. We perform our analysis throughout the paper with both data sets in order to uncover potential discrepancies across data specifications. Both the nonfarm business data and total business data were collected from the Federal Reserve Bank of St. Louis database (FRED). Nonfarm business sector productivity is measured as output per hour of all persons (OPHNFB is the ID of the series). Nonfarm business hours are computed as the ratio between the nonfarm business sector hours of all persons (HOANBS) and the civilian noninstitutional population over the age of 16 (CNP16OV). Total business productivity is measured as the output per hour of all persons (OPHPBS) and total business hours per capita are measured as the business hours of all persons (HOABS) divided by the civilian noninstitutional population over the age of 16 (CNP16OV). We apply natural logarithms to the resulting productivity and hours series. Our data set runs from the first quarter of 1948 to the fourth quarter of 2004. We compared our total business productivity and hours per capita series with those employed by CEV.[1] We compared the data, and the differences between our total business series and

https://doi.org/10.1017/S1365100509080249 Published online by Cambridge University Press

theirs were indeed minimal. Moreover, the impulse responses were essentially the same, despite the fact that their database ends on the fourth quarter of 2001.

Throughout the paper we work with bivariate VARs, because as Galí (1999) and CEV show, introducing additional variables into the vector autoregressive systems does not change the direction of the key impulse responses qualitatively. Our empirical framework is similar to that of GR and CEV. This framework is based upon the existence of an infinite moving average representation for the first differences of the productivity ($\triangle x_t$) and the hours series ($\triangle^i n_t$), where $i = 0$ corresponds to the CEV specification of hours in levels, and $i = 1$ corresponds to the GR specification with the first difference of hours. In matrix notation,

$$\begin{bmatrix} \triangle x_t \\ \triangle^i n_t \end{bmatrix} = \begin{bmatrix} C^{11}(L) & C^{12}(L) \\ C^{21}(L) & C^{22}(L) \end{bmatrix} \begin{bmatrix} \varepsilon_t^x \\ \varepsilon_t^n \end{bmatrix}, \tag{1}$$

where the $C^{ik}(L)$ $(i, k = 1, 2)$ elements are polynomials of infinite order dependent on the lag operator L. ε_t^x and ε_t^n are the technology and hours i.i.d. shocks, respectively. In order to recover the structural macroeconomic shocks, we first estimate bivariate VAR systems. The order of the VAR is chosen to minimize the Schwarz information criterion. In all cases the order chosen was 2. With the estimates of the bivariate VAR(2), we obtain the infinite joint moving average representation of the first differences in productivity and of hours worked as in (1). We then apply the Blanchard and Quah (1989) (BQ) technique to identify the structural shocks. Following both GR and BQ, the identification assumption is that a shock to the hours worked does not affect productivity in the long run, i.e. that $C^{12}(1) = 0$. This identification strategy is implemented by means of a standard Choleski decomposition.

Figure 1 displays the impulse responses of the level of hours worked to a technology shock for the GR and CEV's data specifications with their associated centered 90% asymptotic confidence intervals. We estimated the VAR systems with hours in levels and first differences. The size of the technology shock is normalized to one. Both figures confirm the results reported in the literature [see, for instance, Fernald (2007)]. When the hours variable is treated as an $I(0)$ stationary process, it increases after a technology shock and displays a persistent hump-shaped trajectory, with a slow decay to the steady state value. When hours are treated as a unit root, a different picture emerges, as hours decrease after a technology shock. Then, hours increase and reach their steady-state after several quarters. At the 10% significance level, the responses of hours in levels are statistically positive after two or three quarters, whereas the responses of hours in first differences are only significant on impact and during the following two or three quarters.

As shown in Figure 1, within this framework, the key issue that remains to be elucidated is then the order of integration of the variable hours worked per capita. Standard unit root tests, such as Dickey and Fuller (1979; ADF), Phillips and Perron (1988; PP), or Kwiatkowski et al. (1992; KPSS), are unable to reject the

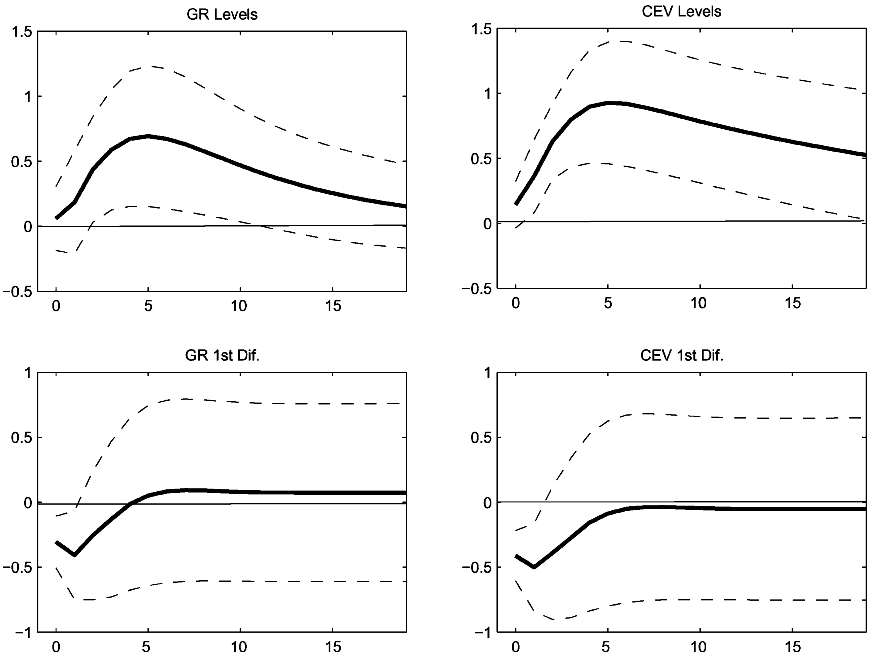

FIGURE 1. Impulse response functions of hours to a technology shock. *Note*: This figure shows the impulse response functions of hours per capita worked to a technology shock. Units are in percentages. The top two panels show the responses of the level of hours to a technology shock using the VAR specification with hours in levels. GR is the Galí and Rabanal (2004) data specification with data for nonfarm businesses and CEV is the Christiano et al. (2003) data specification with data for total businesses. The bottom panels show the analogous responses of hours in levels implied by the VAR specification with hours in first differences. The centered 90% coverage intervals in dashed lines were constructed using a Monte Carlo experiment with 500 replications.

hypothesis of a unit root for the level of the series, and cannot reject that the series is stationary in first differences (see both GR and CEV). Although this result is robust across data specifications, it is well known that the power of these tests is small under meaningful alternatives. Diebold and Rudebusch (1991), Hassler and Wolters (1994), and Lee and Schmidt (1996), among others, show that standard unit-root tests have extremely low power if the alternatives are close to the unit-root circle, but also if they are of a fractional form.

In this paper we circumvent the problem of pretesting the order of integration of the series object of study. Instead of testing for the order of integration of hours before estimating the impulse response of hours to a technology shock, we perform both tasks simultaneously. To do so, we develop a simple method for estimating empirical macroeconomic systems in a multivariate fractional integration setting.

https://doi.org/10.1017/S1365100509080249 Published online by Cambridge University Press

3. A FRACTIONAL INTEGRATION APPROACH

In this section, we first develop a simple method for obtaining the impulse response functions in a multivariate framework when the variables follow $I(d)$ processes. Then we compare the fractional integration framework with the local-to-unity approach, recently employed in multivariate analysis.

3.1. A General Method for Computing Impulse Response Functions in a Multivariate Fractional Integration Framework

In a fractional integration setting, if a variable y_t has an order of integration d ($d \in R$), it is denoted as $y_t \sim I(d)$ and can be expressed as

$$(1-L)^d y_t = \mu_t \quad t = 1, 2, \ldots, \tag{2}$$

with $y_t = 0$, $t \leq 0$. μ_t is an $I(0)$ process, defined as a covariance stationary process, with a spectral density function that is positive and finite at the zero frequency. Thus, μ_t may be a stationary ARMA process. We can express $(1-L)^d$ as the following binomial expansion:

$$\begin{aligned}(1-L)^d &= \sum_{j=0}^{\infty} \binom{d}{j} (-1)^j L^j \\ &= \left[1 - dL + \frac{d(d-1)}{2!}L^2 - \frac{d(d-1)(d-2)}{3!}L^3 \ldots\right].\end{aligned} \tag{3}$$

The representation of y_t in (2) can then be approximated for any real d as

$$\left[1 - dL + \frac{d(d-1)}{2!}L^2 - \frac{d(d-1)(d-2)}{3!}L^3 \ldots\right] y_t = \mu_t. \tag{4}$$

Whereas d captures the long-memory component of the series, μ_t describes the short-run dynamics through its ARMA structure. The literature on fractional models such as (2) has recently emerged in macroeconomics and finance. Some examples are Diebold and Rudebusch (1989), Baillie and Bollerslev (1994), and Gil-Alana and Robinson (1997).[2]

At a theoretical level, an argument employed to justify fractional integration in macroeconomic series is the aggregation of heterogeneous AR processes [see, for instance, Robinson (1978) and Granger (1980)]. Moreover, the fractional integration framework nests the two standard cases documented in the vast majority of applied work in time series. If $d = 0$, as is the case for hours worked in CEV, the series is a covariance stationary process and possesses "short memory," with the autocorrelations decaying fairly rapid. If $d = 1$, as is the case for hours worked in GR, the series is a non stationary $I(1)$ process. But in a fractional framework there are more alternatives available for the order of integration of y_t. If d belongs to the interval (0, 0.50), y_t is still covariance stationary, but both the autocorrelations and the response of a variable to a shock take much longer to disappear than in the standard ($d = 0$) stationary case.[3] If $d \in [0.50, 1)$, the series is no longer covariance stationary, but is still mean reverting, with the effect of the shocks dying away in the long run. Thus, the fractional differencing parameter d plays a crucial role in our understanding of both the structure of the economy and the macroeconomic dynamics. For instance, as d increases, a stronger policy action is required to bring a variable back to its steady state.

There exist many procedures for estimating and testing the fractional differencing parameter d in a univariate framework. They can be parametric, semiparametric, or even nonparametric and they can be specified in either the time or the frequency domain. In Section 4 we employ some of them. However, the main goal of our study is the identification of the structural macroeconomic shocks and the associated impulse response functions in a multivariate setting. We now show how the fractional integration framework captures the joint behavior of a set of macroeconomic variables. We first describe the structural multivariate model and then show how the structural shocks can be recovered from an estimable reduced-form model under standard identifying assumptions.

A set of jointly related macroeconomic variables can be described as

$$ADY_t = \nu_t, \tag{5}$$

where A is an $n \times n$ matrix, Y_t is an $n \times 1$ vector of observable macroeconomic variables, and ν_t is an $n \times 1$ vector of possibly correlated errors. D is an $n \times n$ diagonal matrix that has the following form:

$$D = \begin{bmatrix} (1-L)^{d_1} & 0 & 0 & \dots & 0 \\ 0 & (1-L)^{d_2} & 0 & \dots & 0 \\ \vdots & \vdots & \vdots & \dots & \vdots \\ 0 & 0 & 0 & \dots & (1-L)^{d_n} \end{bmatrix}, \tag{6}$$

where d_i is the order of integration of the variable y_i. We assume, without loss of generality, that the vector of errors ν_t follows a VAR(1) process,

$$\nu_t = G\nu_{t-1} + \varepsilon_t, \tag{7}$$

where G is an $n \times n$ matrix and ε_t is an $n \times 1$ vector of structural macroeconomic shocks i.i.d. with diagonal variance-covariance matrix Σ. Substituting (5) into (7), one can obtain the following infinite moving average representation for the macroeconomic system:

$$Y_t = D^{-1}[I - (A^{-1}GA)L]^{-1}A^{-1}\varepsilon_t, \tag{8}$$

where I is the identity matrix of order n. D^{-1} can be computed easily from the binomial expansion in (3), valid for any real d. We therefore need to identify $2n + 2n^2$ structural parameters: $2n$ from D and Σ and $2n^2$ from A and G. Equation (8) makes clear that the (potentially) fractional orders of integration of the

https://doi.org/10.1017/S1365100509080249 Published online by Cambridge University Press

macroeconomic variables (D) will directly affect the impulse response functions to the structural shocks. Notice that our setting generalizes the standard impulse response function framework, where the diagonal values in D are restricted to be 1 or $1 - L$, depending on the choice of integration order for a variable, $I(0)$ or $I(1)$ respectively. Moreover, we do not have to impose any a priori assumption about the order of integration of the variables because, as we show below, they are simultaneously estimated with the remaining system parameters.

Unlike the standard $I(0)/I(1)$ VAR framework displayed in Section 2, the impulse response functions in our fractional specification are not sensitive to the choice between levels and first differences for the macroeconomic variables. The reason is that the orders of integration are directly obtained from the data using a method that is valid even in nonstationary contexts. Suppose, for instance, that in a bivariate framework, d_1 and d_2 are, respectively, 1.5 and 0.6. If we take first differences before estimating the model, this could be expressed as

$$A\tilde{D}ZY_t = \nu_t, \quad (\mathbf{9})$$

$$\nu_t = G\nu_{t-1} + \varepsilon_t, \quad (\mathbf{10})$$

where Z is a diagonal matrix with $(1 - L)$ elements on the diagonal. Thus we should expect estimates around $\tilde{d}_1 = 0.5$ and $\tilde{d}_2 = -0.4$, so that the model dynamics are equivalent to those in (5) and (7). As a result, the impulse response functions are invariant to the decision on differencing the variables previous to the estimation. Indeed, the impulse response functions of the variables in first differences are exactly the same as those implied by the series in levels.

In a multivariate setting, the number of estimation procedures for fractional integration is very limited. Nielsen (2004) proposed a computationally simple maximum likelihood procedure for multivariate $I(d)$ models. Gil-Alana (2003a, 2003b) proposed an extension of the univariate tests of Robinson (1994) in the frequency domain, whereas Nielsen (2005) developed time-domain versions of Gil-Alana's (2003a, 2003b) tests. These methods make it possible to estimate a reduced-form system such as

$$DY_t = \zeta_t, \quad (\mathbf{11})$$

where ζ_t is an $n \times 1$ stationary $I(0)$ vector of errors. We can further assume that ζ_t follows a stationary VAR(1) process, such as

$$\zeta_t = F\zeta_{t-1} + \eta_t, \quad (\mathbf{12})$$

where F is an $n \times n$ matrix and η_t is a vector of reduced-form errors with variance-covariance matrix V. Substituting (12) into (11), the following infinite moving average representation can be derived:

$$Y_t = D^{-1}(I - FL)^{-1}\eta_t. \quad (\mathbf{13})$$

The relation between structural and reduced-form error terms in (8) and (13) is then given by

$$A^{-1}\varepsilon_t = \eta_t, \tag{14}$$

whereas the relation between structural and reduced-form coefficient matrices is given by

$$A^{-1}GA = F. \tag{15}$$

The reduced-form model has $n + n^2 + n(n+1)/2$ parameters: n from D, n^2 from F, and $n(n+1)/2$ from V. As a result, we need $n(n+1)/2$ restrictions in the structural system (5)–(7) so that the structural errors can be identified. It is standard in the literature to assume that the variance-covariance of the structural error vector is the identity matrix. Therefore, we need $n(n-1)/2$ additional restrictions. A standard identification strategy is based upon imposing long-run restrictions, as in BQ, GR, or CEV. Notice that we can express (8) and (13), respectively, as

$$Y_t = C_0\varepsilon_t + C_1\varepsilon_{t-1} + C_2\varepsilon_{t-2} + \ldots, \tag{16}$$

$$Y_t = \eta_t + R_1\eta_{t-1} + R_2\eta_{t-2} + \ldots. \tag{17}$$

We let $C(1) = \sum_{j=0}^{\infty} C_j$ and $R(1) = \sum_{j=0}^{\infty} R_j$, so that placing restrictions on $C(1)$ is equivalent to setting long-run restrictions. To see this, note that equalizing the variances across expressions, we have $R(1)VR(1)' = C(1)C(1)'$. Once $C(1)$ is identified, so are $C_0, C_1, C_2, \ldots$ (since $C_0 = R(1)^{-1}C(1), C_1 = R_1C_0, \ldots$), and thus the structural shocks. In our framework $n = 2$, so that we only need one long-run restriction.[4] In agreement with both GR and CEV, we assume that a structural shock to hours worked does not affect productivity in the long run, i.e., $C^{12}(1) = 0$. Notice finally that D, while capturing additional long-memory dynamics, does not alter the standard impulse response identification techniques.

3.2. Fractional Integration and Local-to-Unity Approaches

In a closely related article, Pesavento and Rossi (2005) employ an alternative local-to-unity approach to elucidate the response of hours to a technology shock. Although the fractional integration and local-to-unity approaches are essentially different, they also present some interesting relationships. The local-to-unity approach is agnostic in the sense that it does not impose a priori an $I(0)$ or an $I(1)$ process for hours worked. However, similarly to the classic methods, the model is encompassed in the classic AR framework, whereas ours uses a different fractionally integrated model. We illustrate this in the univariate case for hours. The model in Pesavento and Rossi (2005) is

$$(1-\rho L)n_t = \zeta_t, \quad \rho = 1 + \frac{1}{T}c, \tag{18}$$

whereas ours is

$$(1-L)^d n_t = \zeta_t \tag{19}$$

https://doi.org/10.1017/S1365100509080249 Published online by Cambridge University Press

and the unit-root null corresponds respectively to

$$H_0 : \rho = 1 \quad (\text{or} \quad c = 0) \tag{20}$$

in (18) and

$$H_0 : d = 1 \tag{21}$$

in (19). However, fractional and AR departures from (20) and (21) have very different long-run implications. In (19), n_t is nonstationary but nonexplosive for all $0.5 \leq d < 1$. As d increases beyond 0.5 and through 1, n_t can be viewed as becoming "more nonstationary," but it does so gradually, unlike the case of (18) around (20). The dramatic long-run change in (18) around $\rho = 1$ has the attractive implication that rejection of (20) can be interpreted as evidence of either stationarity or explosivity. However, rejection of the null does not necessarily warrant acceptance of any particular alternative. Alternatively, the $I(d)$ class comprises many stationary, nonstationary, invertible, and noninvertible processes. This is in sharp contrast to asymptotic theory for statistics directed against AR alternatives where different asymptotic theory is obtained for $|\rho| < 1$, for $\rho = 1$, and for $|\rho| > 1$.

The fact that we do not impose an $I(0)/I(1)$ specification for hours worked is consistent with what Pesavento and Rossi (2005) did. They, however, assume a VAR specification as in (1), with $\triangle^i$ replaced by $(1 - \rho L)$ and $\rho = 1 + \frac{1}{T}c$, such that the process is stationary and highly persistent (with $c < 0$). An advantage of the $I(d)$ approach is that d can be any real value, encompassing then stationary $I(0)$ and nonstationary $I(1)$ models. Nevertheless, a drawback compared to the local-to-unity approach is that the results for long-run dynamics rely on a single parameter, d. The two approaches lead to a quite different decay of the impulse responses; the local-to-unity being exponential, the fractional integration being hyperbolical.

4. UNIVARIATE ANALYSIS

This section presents empirical evidence on the fractional orders of integration of both labor productivity and hours worked per capita in a univariate setting. This evidence is relevant for two reasons. First, although researchers have applied standard unit root tests to the productivity and hours variables in order to decide whether to introduce these variables in either levels or first differences into their VARs, no study has investigated the fractional orders of integration of these variables. The fractional setting is clearly more general, as it allows a given variable to display an order of integration different from one and zero. We apply parametric and semiparametric fractional integration tests in the frequency domain. Second, we will be able to use the univariate results in a multivariate setting for two purposes: First, if the fractional integration tests manage to pin down the orders of integration of some of the variables clearly, then we can directly assume them to be the right ones in a multivariate analysis. Second, we can assess whether the estimates of the order of integration of a given variable differ in univariate and multivariate contexts.

We first use a parametric method proposed by Robinson (1994) to test for the order of integration of the productivity and hours series. This method is based on the Lagrange multiplier (LM) principle and uses the Whittle function, which is an approximation to the likelihood function. One advantage of this method is that it allows us to consider fractional orders of integration at any real value d, including thus both stationary and nonstationary processes. In fact, other parametric methods such as that of Sowell (1992) only allowed $-0.5 < d < 0.5$, i.e., in the stationary region. Another advantage of the fractional approach presented here is that it does not display an abrupt change in the limit behavior of the tests against the unit root. In fact, the limit distribution is a standard normal for any real value d. In contrast, the classic ADF, PP, and KPSS methods have a nonstandard limit distribution in the sense that the critical values must be tabulated case by case by means of a Monte Carlo simulation study. For ease of exposition, we rewrite the standard expression for a fractionally integrated process y_t,

$$(1 - L)^d y_t = \mu_t, \tag{22}$$

with $I(0)$ μ_t. Following the approach of Robinson (1994), we test

$$H_0 : d = d_0, \tag{23}$$

for any given real value d_0, in a model given by

$$x_t = \alpha + \beta t + y_t, \tag{24}$$

with t as a time trend and y_t given by (22). Note that x_t is the observable macroeconomic variable and y_t is now the regression error series, which might be fractionally integrated according to (22). We first assume that $\alpha = \beta = 0$ in (24); i.e., there are no deterministic terms, implying that $x_t = y_t$. We also consider the cases of an unknown α and $\beta = 0$ (with an intercept) and both α and β unknown (a linear time trend). The results for the four series are given in Tables 1 and 2.[5] In Table 1 we assume that the μ_t disturbances are white noise. In Table 2 we permit autocorrelation patterns in the error term. Across these tables, we report the confidence intervals of those values of d_0 where the null hypothesis cannot be rejected at the 5% level.[6] We also display in the tables the value of d_0 producing the lowest statistic (in absolute value) across d's. This value should be an approximation to the maximum likelihood estimate.

Starting with the case of white noise for μ_t, we see in Table 1 that if we do not include regressors, the unit root null hypothesis (i.e., $d_0 = 1$) cannot be rejected for any series. This hypothesis cannot be rejected for either of the two productivity series when an intercept and/or a linear trend is included in the regression model. For the number of hours, the unit root is rejected in all cases in favor of higher orders of integration. In what respects to the model with autocorrelated residuals, we first estimated autoregressive (AR) models. Modeling μ_t in terms of an AR(1)

TABLE 1. Robinson's (1994) univariate test for fractional integration: White noise disturbances

Series	No regressors			Intercept			Linear time trend		
GP	[0.91	**0.98**	1.08]	[0.91	**1.01**	1.13]	[0.94	**1.01**	1.10]
GH	[0.91	**0.98**	1.08]	[1.45	**1.60**	1.77]	[1.45	**1.60**	1.77]
CP	[0.92	**0.99**	1.09]	[0.93	**1.03**	1.13]	[0.96	**1.02**	1.10]
CH	[0.91	**0.98**	1.08]	[1.36	**1.49**	1.65]	[1.36	**1.49**	1.65]

Note: This table shows the 95% confidence intervals for the order of integration of a given time series computed through the Robinson (1994) model,

$$x_t = \alpha + \beta t + y_t,$$
$$(1 - L)^d y_t = \mu_t,$$

where x_t is the macroeconomic variable: GP is the productivity variable used by Galí and Rabanal (2004), GH is the hours variable used by Galí and Rabanal (2004), CP is the productivity variable used by Christiano et al. (2003), and CH is the hours variable used by Christiano et al. (2003). α and β are constants, d is the order of integration of each process, and μ_t is assumed to be a white noise process. The value of d corresponding to the lowest statistics of the Robinson (1994) test appears in bold in the middle of the confidence interval.

TABLE 2. Robinson's (1994) univariate test for fractional integration: Autocorrelated disturbances

Series	No regressors			Intercept			Linear time trend		
GP	[0.83	**0.97**	1.13]	[0.77	**1.04**	1.27]	[0.92	**1.01**	1.17]
GH	[0.84	**0.97**	1.13]	[0.84	**1.05**	1.36]	[0.85	**1.05**	1.36]
CP	[0.85	**0.96**	1.14]	[0.98	**1.16**	1.33]	[1.00	**1.10**	1.25]
CH	[0.85	**0.97**	1.14]	[1.07	**1.21**	1.35]	[1.05	**1.14**	1.27]

Note: This table shows the 95% confidence intervals for the order of integration of a given time series computed through the Robinson (1994) model,

$$x_t = \alpha + \beta t + y_t$$
$$(1 - L)^d y_t = \mu_t,$$

where x_t is the macroeconomic variable. GP is the productivity variable used by Galí and Rabanal (2004), GH is the hours variable used by Galí and Rabanal (2004), CP is the productivity variable used by Christiano et al. (2003), and CH is the hours variable used by Christiano et al. (2003). α and β are constants, d is the order of integration of each process, and μ_t is assumed to be an autocorrelated process. The value of d corresponding to the lowest statistics of the Robinson (1994) test appears in bold in the middle of the confidence interval. These statistics are computed assuming that μ_t follows the nonparametric autocorrelated model of Bloomfield (1973).

process produced some inconsistencies in the interpretation of the results. For instance, the null hypothesis of $d = 0$ was not rejected in some cases; it was rejected for values of d between 0 and 1 and it was again not rejected for values of d close to 1. This lack of consistency can be explained by the fact that the AR coefficients, though lower than 1 in absolute value, can be arbitrarily close to 1 and thus they might be competing with d in describing nonstationarity. Note that other standard unit root testing procedures face the same problem. We solved

https://doi.org/10.1017/S1365100509080249 Published online by Cambridge University Press

this problem by using the method of Bloomfield (1973). This method, which can be flexibly applied in the context of Robinson's (1994) tests, does not impose a given parametric model for the $I(0)$ disturbances but implies autocorrelations for μ_t which decay exponentially as in the ARMA case. Moreover, this model is stationary across the whole range of values for the parameter set, unlike the AR case. Using this model, the results in Table 2 are very similar across series and most of the nonrejection values for d oscillate around 1. However, for the number of hours, although the unit root is not rejected in the case of the GR series, d is larger than 1 using CEV's definition and the null of a unit root is rejected in two of the three cases. Two additional important features observed across the tables are worth commenting. First, if we do not include regressors, the lowest statistics occur in all series at values of d lower than 1. However, including deterministic terms, they occur at values slightly higher than 1. Second, the estimated orders of integration of hours are substantially lower if autocorrelated errors are permitted, especially in the models including intercept and a linear trend. This suggests that the addition of short-run dynamics, even in a nonparametric form, tends to lower the orders of integration. This is clearly a consequence of the inclusion of an additional alternative way to describe the dependence across observations.

To confirm the above results we also display in Figure 2 the estimates of d based on a semiparametric "local" Whittle method proposed by Robinson (1995). We use this method because of its computational simplicity, noting that it simply requires a single bandwidth number (m), and no additional user-chosen numbers are required in the estimation, as is the case with other semiparametric methods. The top panel of Figure 2 shows the results for the GR series whereas the bottom panel presents the results for the CEV's counterparts. For both series, we display the estimates of d across the whole range of values for the bandwidth number ($m = 1, 2, \ldots, T/2$), along with the 95% confidence interval corresponding to the $I(1)$ hypothesis.[7] Notice that the estimate of d is asymptotically normally distributed. Starting with the number of hours, we see that the results are quite unstable. Thus, if the bandwidth number is lower than $T/4$, most of the estimates of d are within the $I(1)$ interval; however, if it is higher than $T/4$, the values of d are significantly above 1. Alternatively, the results for the productivity series strongly support the hypothesis of a unit root in the two cases.

To sum up, our univariate fractional integration results strongly support the hypothesis of a unit root for the productivity series and lead to some ambiguous conclusions about the order of integration for hours worked. In the next section we estimate the order of integration of the hours series in a multivariate context.

5. MULTIVARIATE ANALYSIS

This section applies the multivariate model derived in Section 3 to give an answer to the technology-hours question. What is interesting about this framework is that it allows the econometrician to jointly estimate the orders of integration of

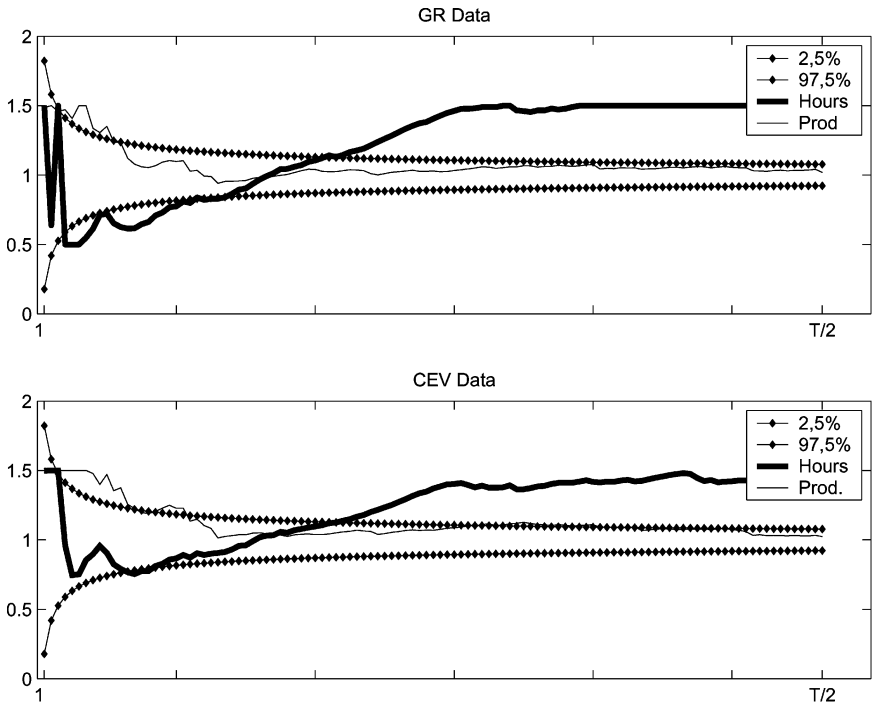

FIGURE 2. Fractional orders of integration using the local Whittle semiparametric estimates of d [Robinson (1995)]. *Note*: This figure shows the fractional orders of integration of productivity and hours worked for the data specifications in Galí and Rabanal (2004) (GR) and Christiano, Eichenbaum, and Vigfusson (2003) (CEV). The orders of integration are computed according to the model proposed by Robinson (1995). The horizontal axis identifies the amplitude of the bandwidth, which goes from 1 to $T/2$, where T is the sample size. The vertical axis identifies the order of integration (d). The 95% confidence intervals of the null for $d = 1$ appear in diamonds.

the macroeconomic variables and the impulse response functions to the structural shocks.

As noted in Section 2, we estimate bivariate systems with labor productivity and hours worked. We will assume that the order of integration of the productivity series is 1 throughout the following analysis. Our motivation for this assumption is threefold. First, the univariate tests decisively pointed at 1 as the order of integration of productivity, unlike the hours case (see, e.g., Figure 2). Second, this assumption is uncontroversial for all of the papers in the technology-hours literature. Indeed, GR, CEV, and all related papers assume that productivity is integrated of order 1. Third, by assuming that the order of integration of productivity is 1, our approach will estimate more efficiently the order of integration of hours in a multivariate model, a very important object of study in the present paper. Nevertheless, we also

https://doi.org/10.1017/S1365100509080249 Published online by Cambridge University Press

computed the procedure allowing both orders of integration to be unknown, and the value for the productivity series was very close to 1 in all cases.

The bivariate models are estimated following Gil-Alana (2003a), who derives an extension of the Robinson (1994) univariate frequency domain tests, leading to exactly the same estimates as in the maximum likelihood procedure proposed by Nielsen (2004). This method yields normally distributed values for d, which allows us to perform inference on impulse response function analysis. The multivariate approach presents two important advantages. First, we do not need to impose a priori any assumption about the orders of integration of the series since they are freely estimated from the real line. Second, with respect to the univariate case, the order of integration of hours is estimated more efficiently, because it makes use of additional information in both the parameters of the variance-covariance of the residuals and those in the productivity equation.

We estimated the fractionally integrated model

$$\begin{bmatrix} (1-L) & 0 \\ 0 & (1-L)^{d_H} \end{bmatrix} \begin{pmatrix} x_t \\ n_t \end{pmatrix} = \begin{pmatrix} \zeta_{1,t} \\ \zeta_{2,t} \end{pmatrix}, \tag{25}$$

letting the differenced series ($\zeta_t = [(1-L)x_t, \quad (1-L)^{d_H} n_t]'$) follow a VAR(1) in order to accommodate both short-run persistence and long memory, as described in equations (11) and (12).[8] Gil-Alana's (2003a) method is based on testing the null hypothesis $d_H = d_{H_0}$ for any real d_{H_0} value in (25) and the functional form of the test statistic and its limiting distribution is described in Appendix A.1. We note that this method is based on the LM principle and uses the Whittle function in the frequency domain. Under the null hypothesis, we estimate the VAR coefficients in the model given by

$$\begin{pmatrix} \zeta_{1,t} \\ \zeta_{2,t}^0 \end{pmatrix} = \begin{pmatrix} f_{11} & f_{12} \\ f_{21} & f_{22} \end{pmatrix} \begin{pmatrix} \zeta_{1,t-1} \\ \zeta_{2,t-1}^0 \end{pmatrix} + \begin{pmatrix} \eta_{1,t} \\ \eta_{2,t} \end{pmatrix}, \tag{26}$$

where $\zeta_{2,t}^0 = (1-L)^{d_{H0}} n_t$, and the parameters in F are estimated by least squares. On the other hand, using Nielsen's (2004) MLE approach, the vector of long-memory parameters d follows the distribution

$$\sqrt{T}(\hat{d} - d) \to_d N(0, \Gamma^{-1}), \tag{27}$$

with

$$\Gamma = \frac{\pi}{6} \Omega \otimes \Omega^{-1} - (\Omega \Phi' V^{-1} \Phi \Omega) \otimes \Omega^{-1}. \tag{28}$$

Ω and V are the variance-covariance matrices of the reduced-form error terms ζ_t and η_t in (12), respectively. Φ is the matrix of coefficients in the Wold representation of the ζ_t process.

Table 3 lists the orders of integration for hours worked across data specifications. The order of integration of the GR measure (0.67) is much higher than that of CEV (0.04). It also displays the associated 95% asymptotic confidence intervals, revealing that most of the integration orders' probability mass lies around the

TABLE 3. Fractional order of integration of hours worked: Multivariate model

Series			
GR	[0.62	**0.67**	0.74]
CEV	[0.00	**0.04**	0.07]

Note: This table shows the results for the order of integration of per capita hours worked obtained under the multivariate model described in Section 5. The model is expressed as

$$DY_t = \zeta_t,$$
$$\zeta_t = F\zeta_{t-1} + \eta_t,$$

where

$$D = \begin{bmatrix} (1-L) & 0 \\ 0 & (1-L)^{d_H} \end{bmatrix}$$

and Y_t is a bivariate vector including labor productivity and hours worked in levels. The table shows the 95% confidence interval along with the value of d_H producing the lowest value statistic (in bold) for the order of integration of hours across data specifications [Galí and Rabanal (2004) (GR) and Christiano et al. (2003) (CEV)].

TABLE 4. Multivariate autocorrelated model: VAR(1) matrix for structural residuals

$$G^{\mathrm{GR}} = \begin{bmatrix} 0.0087 & 0.0045 \\ 0.0213 & 0.7984 \end{bmatrix}$$
$$G^{\mathrm{CEV}} = \begin{bmatrix} -0.0209 & -0.0057 \\ 0.1498 & 0.9800 \end{bmatrix}$$

Note: This table shows the implied VAR(1) matrices for the structural residuals of the multivariate model described in Section 5. The model is expressed as

$$\mathrm{ADY}_t = \nu_t,$$
$$\nu_t = G\nu_{t-1} + \varepsilon_t.$$

G^{GR} is the G matrix obtained with the data in Galí and Rabanal (2004), whereas G^{CEV} is the G matrix obtained with the data in Christiano et al. (2003).

point estimates. Table 3 also shows that the orders of integration for hours worked across data specifications are both economically and statistically lower than in the univariate frameworks. This finding suggests that introducing additional cross-sectional and time series information reduces the estimated order of integration. Table 4 shows the implied VAR(1) matrix of coefficients for the structural error terms [matrix G in equation (7)]. Interestingly, it shows that although the autoregressive coefficient in the productivity equation is close to zero across data specifications, its counterpart in the hours equation is close to one in both cases,

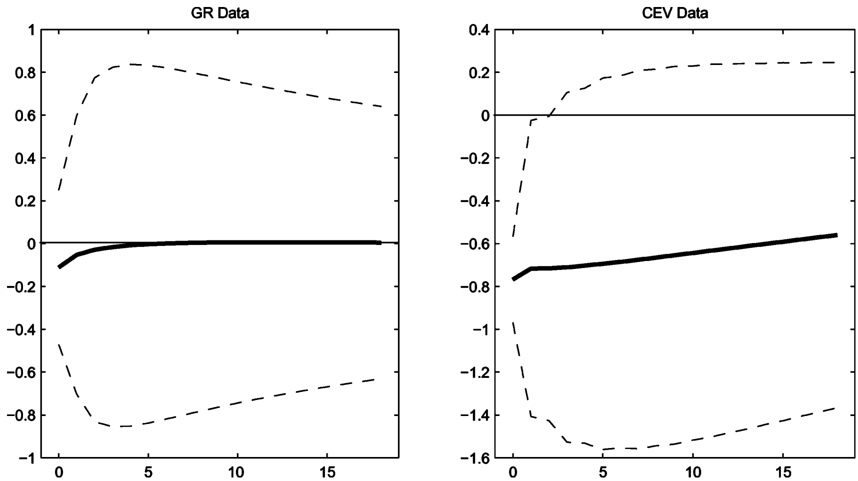

FIGURE 3. Dynamic response of hours worked to a technology shock. *Note*: This figure shows the response of per capita hours worked to a structural technology shock in the fractional integration multivariate model

$$\begin{aligned} ADY_t &= \nu_t, \\ \nu_t &= G\nu_{t-1} + \varepsilon_t, \end{aligned}$$

where Y_t is a 2×1 vector including productivity and hours, ν_t follows a VAR(1) law of motion, and ε_t follows a white noise process. Units are in percentages. The left-hand panel shows the response of hours with the Galí and Rabanal (2004) (GR) data set, whereas the right-hand panel shows the response of hours with the Christiano et al. (2003) (CEV) data set. The centered 90% asymptotic coverage intervals in dashed lines were constructed using a Monte Carlo experiment with 500 replications, described in Appendix A.2.

especially with the CEV data. In other words, most of the time dependence in the CEV hours variable is now captured by the autoregressive structure, so that the degree of integration of the series is close to zero.

The next step in our analysis is to report the associated impulse response functions of hours to a technology shock. To do so, we plot in Figure 3 the impulse response functions implied by the point estimates of the model in (25) under both data specifications. Both panels include the centered 90% asymptotic confidence intervals computed through a Monte Carlo exercise described in Appendix A.2. Notice that the confidence intervals become larger after the initial impact, which should reflect the additional uncertainty entailed by the fractional integration parameter.

The left-hand panel of Figure 3 shows that the response of the level of hours in the GR specification to a technology shock is negative on impact, although not statistically significant. Following the initial reaction, hours increase toward the steady-state level, which is reached after six quarters. Then hours slightly overshoot the

https://doi.org/10.1017/S1365100509080249 Published online by Cambridge University Press

steady-state level and after twelve quarters start to slowly revert to it. Our point-estimate response of hours to the technology shock is qualitatively similar to (and quantitatively slightly smaller than) that of GR and Francis and Ramey (2005a). With respect to Pesavento and Rossi (2005), our impulse response point-estimate is slightly less negative on impact and less positive after several quarters.

Regarding the response of hours in the CEV specification, the initial impact is negative and statistically significant at the 10% confidence level. Since then, hours remain negative, but they increase slowly through time. Notice that the initial negative response of per capita hours is larger in the case of the CEV specification. After the initial impact, hours remain negative but they increase very slowly through time. This impulse response is qualitatively similar to the one obtained by Fernald (2007) in his Figure 2, although slightly larger and more slowly mean-reverting.

It may appear surprising to obtain a strong negative response of hours worked to the technology shock in the CEV specification, when the order of integration of hours is close to zero. Moreover, the response is quantitatively larger than in the GR data specification, where the order of integration was estimated to be considerably higher. To develop some intuition as to why this may be the case, we perform the following exercise. We fractionally differentiate the CEV hours series over a grid of values on 0.01 increments, including 0 and 1. Then we proceed to estimate a VAR(1) in the resulting $I(0)$ framework, imposing the BQ long-run restriction in order to recover the response of hours to the technology shock. Notice that this is a two-step exercise, as opposed to the one-step estimation performed earlier. However, it can give us an intuition of how the response of hours may change depending on its fractional order of integration. Figure 4 shows the response of hours on impact to the technology shock for the grid of values described. Interestingly, the series is clearly nonmonotonic. As expected, when $d_H = 0$ the response is positive, whereas when $d_H = 1$ the response is negative. Nevertheless, notice that the response of hours becomes negative for very low values of d_H. Indeed, for $d_H = 0.04$, the initial response of hours is -0.12%, almost twice as negative as when $d_H = 0.67$. The series also exhibits several local maxima and a fast decline of the impact effect on hours as d_H approaches unity. Finally, it is interesting to note that the impact response of hours to the technology shock is negative for all values of d_H smaller than zero. Thus, a slight differentiation of the original series yields a decline of hours following a technology shock.

As a corollary to this exercise, we would like to emphasize that all of the model's parameter estimates, as well as the identified structural shocks, are conditional on our fractional integration structure, so that differences with respect to the $I(0)/I(1)$ cases can naturally arise. Moreover, an important difference between the two approaches is that the fractional integration framework makes a variable depend on an infinite number of lags through the binomial expansions of the long-memory coefficients.

There can also be economic explanations behind the statistical differences of the responses across data specifications. The nature of the hours and productivity

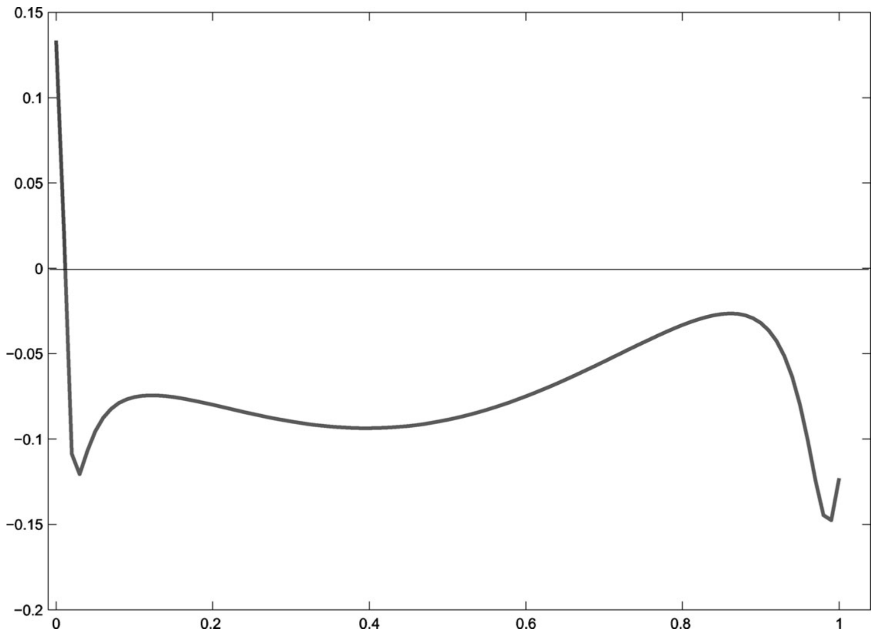

FIGURE 4. Hours response (on impact) to the technology shock: Sensitivity analysis. *Note*: This figure displays a sensitivity analysis showing the impact response of hours to the technology shock for a grid of values for the fractional differencing parameter d_H on the number of hours. This series is the result of a two-step procedure performed with the CEV data specification, where hours are first fractionally differentiated and then the resulting bivariate $I(0)$ VAR(1) system is estimated for each order of integration of hours.

series differs across specifications. The CEV measure includes hours worked in the agricultural sector, which can be more sensitive to productivity shocks. Indeed, Figure 5 plots the logs of the two hours series, showing that nonfarm hours did not converge with total hours until the mid-1970s. Before that date, total hours were clearly above their nonfarm counterpart. With respect to the productivity series, total business productivity growth is larger on average and in standard deviation than its nonfarm counterpart, especially for the first part of the sample. The combination of these two facts can make the CEV measure more sensitive to labor productivity shocks.[9]

In summary, our two data specifications detect a decline of hours worked on impact in response to a technology shock. In the CEV specification, the results are statistically significant.

6. DISCUSSION AND CONCLUSIONS

The goal of this paper was to determine the response of hours worked to a technology shock, a debated question in macroeconomics today. Our contribution is to

https://doi.org/10.1017/S1365100509080249 Published online by Cambridge University Press

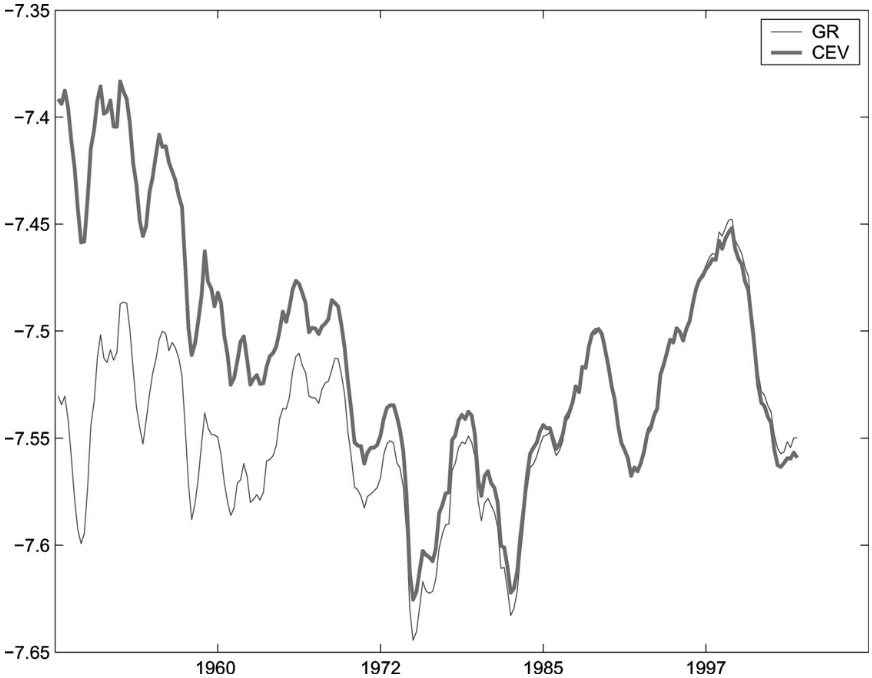

FIGURE 5. Hours across data specifications. *Note*: This figure compares the plots of the natural logarithms of per capita hours worked across data specifications. GR stands for nonfarm business hours, and CEV stands for total business hours.

derive a new unified econometric framework which simultaneously determines the order of integration of hours worked and the dynamic impulse response function of hours to a technology shock. Our results lend support to the hypothesis that hours worked fall in response to a positive technology shock.

In a recent paper, Francis and Ramey (2005b), building on the intuition of Fernald (2007), construct a measure of per capita hours worked that removes some of the low-frequency fluctuations of the standard measures of hours worked.[10] They find that removing these low-frequency dynamics renders the variable stationary and that this variable decreases on impact in response to a technology shock. As mentioned above, our implied impulse responses are consistent with their findings, despite of the clear methodological differences. We believe that the fractional integration framework presented in this paper is well suited to account for low-frequency dynamics, because it controls for the long memory of the stochastic processes implied by macroeconomic aggregates. Nevertheless, it would be interesting to perform univariate and multivariate fractional integration analysis with the new hours variable proposed by Francis and Ramey (2005b).

Recent works by Fisher (2006) and Fernald (2007) also show that the responses of hours to a neutral technology shock, such as the one studied in this paper, are sensitive to the existence of breaks in the series (both in productivity and in hours worked). We perform subsample analysis in our fractional integration framework and indeed verify that this may also be the case in our setting. We split the sample in the first quarter of 1973, when there is a decline in productivity growth, as shown by Fernald (2007). An interesting result of this subsample analysis is that although the estimated orders of integration of hours worked are very similar for the GR measure across subsamples (0.70 and 0.66), they are statistically and qualitatively different for the CEV measure (0.19 and 0.80). This implies that during the first subsample, which coincides with a greater weight of farm business activities, the hours variable used by CEV displays a lower degree of long memory compared with the post-1973 period. It is thus this subsample that seems to trigger the low integration order of hours worked in the full sample.

The present article raises a number of interesting questions for future research. A first issue is related to the difference in the order of integration of the variable hours estimated in univariate and multivariate contexts. We found that it was lower in the case of the multivariate models. This finding, in itself, suggests that conditioning on additional information may reduce the memory of a given process. Although multivariate tests are not often used to determine the level of integration of a given variable, they are most interesting for macroeconomists, because the macroeconomic literature often focuses on the dynamic properties in systems of variables. In this sense, the issue of pretesting for the order of integration of a given variable in univariate frameworks may be of secondary importance once we control for the fractional order of integration in a multivariate framework. Moreover, our setup solves the potential problem of unbalanced orders of integration in standard time series regression frameworks.[11] The study of fractionally co-integrated systems, allowing for a nondiagonal matrix D in (6), also seems a fruitful avenue for future research in this area.

Finally, exploring the relation between fractionally integrated systems and structural macroeconomic models remains another important unresolved topic. Macroeconomic models typically imply stationary systems and co-integrating relationships, but not fractional integration. However, discrete shifts in either behavioral or policy parameters can give rise to time-varying coefficients in reduced-form representations, which, in turn, can be reconciled with fractional integration [see Ding and Granger (1996)].

NOTES

1. We are very grateful to Elena Pesavento and Barbara Rossi for kindly providing the data. They, in turn, received the data directly from CEV.

2. The fractional integration literature was pioneered by Granger (1980) and Granger and Joyeaux (1980). See Baillie (1996) and Gil-Alana and Hualde (2009) for a complete review of $I(d)$ processes.

3. Note that if $d = 0$ and μ_t follows an AR process, the decay in the autocorrelations is exponentially rapid compared with the $I(d, d > 0)$ case, where the decay is hyperbolic.

4. In other contexts, an alternative strategy followed by researchers such as Christiano et al. (1999) is to place short-run restrictions by assuming that the matrix A in the structural model (5) is lower triangular.

5. The inclusion of a quadratic trend in (24) for hours worked [as suggested by GR and Fernald (2007)] does not significantly change the conclusions reported in the paper. The results in this case are available from the authors upon request.

6. These intervals were constructed as follows: First we choose a value of d from a grid, $d_0 = 0, 0.01, \ldots, 2$. Then we compute the test statistic testing the null for this value. If the null is rejected at the 5% level, we discard this value of d. Otherwise, we keep it. An interval is then obtained after consideraction of all the values of d in the grid.

7. In the case of the "local" Whittle estimator, the use of optimal bandwidth values has not been theoretically justified. Some authors, such as Lobato and Savin (1998), use an interval of values for m.

8. Multivariate versions of the Bloomfield's (1973) model have not been yet developed. Moreover, they would be of no use to compute impulse responses, given that the Bloomfield model does not display a parametric formula for the disturbances μ_t. We also estimated a VAR(2) specification for the residuals in (25), and the results did not reveal significant differences.

9. Nonfarm (industrial and services) hours may display greater persistence because of other reasons, such as mobility or substitution effects. Marelli (1994) shows that persistence in unemployment is higher in industrialized sectors than in agriculture. The same is found in the case of Spanish unemployment [see García del Barrio and Gil-Alana (2007)].

10. Fernald (2007) removes the means of identified subsamples in labor productivity growth and finds a fall of hours in response to a technology shock. He finds the same effect removing a quadratic trend from labor productivity. A differential feature of our fractional integration approach is that it lets the data stochastically determine the low frequencies present in hours.

11. See Baillie and Bollerslev (1994) for an exposition of this problem in the context of the forward premium puzzle.

REFERENCES

Baillie, Robert T. (1996) Long memory processes and fractional integration in econometrics. *Journal of Econometrics* 73, 5–59.

Baillie, Robert T. and Tim Bollerslev (1994) The long memory of the forward premium. *Journal of International Money and Finance* 15, 565–571.

Blanchard, Olivier J. and Danny Quah (1989) The dynamic effects of aggregate demand and supply disturbances. *American Economic Review* 79, 655–673.

Bloomfield, Peter (1973) An exponential model in the spectrum of a scalar time series. *Biometrika* 60, 217–226.

Christiano, Lawrence, Martin Eichenbaum, and Charles L. Evans (1999) Monetary policy shocks: What have we learned and to what end? In John B. Taylor and Michael Woodford (eds.), *Handbook of Macroeconomics*, Vol. 1A, pp. 65–148. Amsterdam: North Holland.

Christiano, Lawrence, Martin Eichenbaum, and Robert Vigfusson (2003) What Happens After a Technology Shock? NBER Working Paper 9819.

Dickey, David A. and Wayne A. Fuller (1979) Distribution of the estimator for autoregressive time series with a unit root. *Journal of the American Statistical Association* 74, 427–431.

Diebold, Francis X. and Glenn D. Rudebusch (1989) Long memory and persistence in aggregate output. *Journal of Monetary Economics* 24, 189–209.

Diebold, Francis X. and Glenn D. Rudebusch (1991) On the power of Dickey-Fuller tests against fractional alternatives. *Economic Letters* 35, 155–160.

Ding, Zhuanxin and Clive W. J. Granger (1996) Modelling volatility persistence of speculative returns. *Journal of Econometrics* 73, 185–215.

Fernald, John (2007) Trend breaks, long-run restrictions, and the contractionary effects of technology shocks. *Journal of Monetary Economics* 54, 2467–2485.

Fisher, Jonas (2006) The dynamic effects of neutral and investment-specific shocks. *Journal of Political Economy* 114, 413–451.

Francis, Neville and Valerie A. Ramey (2005a) Is the technology-driven real business cycle hypothesis dead? Shocks and aggregate fluctuations Revisited. *Journal of Monetary Economics* 52, 1379–1399.

Francis, Neville and Valerie A. Ramey (2005b) Measures of Per Capita Hours and Their Implications for the Technology-Hours Debate. NBER Working Paper 11694.

Galí, Jordi (1999) Technology, employment, and the business cycle: Do technology shocks explain aggregate fluctuations? *American Economic Review* 89, 249–271.

Galí, Jordi and Pau Rabanal (2004) Technology shocks and aggregate fluctuations: How well does the RBC model fit postwar U.S. data? *NBER Macroeconomics Annual*, pp. 225–288.

García-del-Barrio, Pedro and Luis A. Gil-Alana (2007) New revelations about unemployment persistence in Spain: Time series and panel data approaches using regional data. *Applied Economics* 38(21), 1–18.

Gil-Alana, Luis A. (2003a) A fractional multivariate long memory model for the US and the Canadian real output. *Economic Letters* 81, 355–359.

Gil-Alana, Luis A. (2003b) Multivariate tests of fractionally integrated hypotheses. *South African Statistical Journal* 37, 1–28.

Gil-Alana, Luis A. and Javier Hualde (2009) Fractional integration and cointegration: An overview and an empirical application. In Terence C. Mills and K. Patterson (eds.), *Handbook of Applied Econometrics*, Vol. 2, pp. 434–470. New York: Palgrave MacMillan, Ltd.

Gil-Alana, Luis A. and Peter M. Robinson (1997) Testing of unit roots and other nonstationary hypotheses in macroeconomic time series. *Journal of Econometrics* 80, 241–268.

Granger, Clive W. J. (1980) Long memory relationships and the aggregation of dynamic models. *Journal of Econometrics* 14, 227–238.

Granger, Clive W. J. and Roselyne Joyeaux (1980) An introduction to long memory time series and fractional differencing. *Journal of Time Series Analysis* 1, 15–29.

Hamilton, James D. (1994) *Time Series Analysis*. Princeton, NJ: Princeton University Press.

Hassler, Uwe and Jurgen Wolters (1994) On the power of unit root tests against fractional alternatives. *Economic Letters* 45, 1–5.

Kwiatkowski, Denis, Peter C. B. Phillips, Peter Schmidt, and Yongcheol Shin (1992) Testing the null hypothesis of stationarity against the alternative of a unit root. *Journal of Econometrics* 54, 159–178.

Lee, Junsoo and Peter Schmidt (1996) On the power of the KPSS test of stationarity against fractionally integrated alternatives. *Journal of Econometrics* 73, 285–302.

Lobato, Ignacio N. and Nathan E. Savin (1998) Real and spurious long-memory properties of stock market data. *Journal of Business and Economic Statistics* 16, 261–283.

Marelli, Enrico (1994) Evolution of unemployment structures and regional specialisation in the EU. *Economic Systems* 28, 35–59.

Nielsen, Morten Ørregaard (2004) Efficient inference in multivariate fractionally integrated time series models. *Econometrics Journal* 7, 63–97.

Nielsen, Morten Ørregaard (2005) Multivariate Lagrange multiplier tests for fractional integration, *Journal of Financial Econometrics* 3, 372–398.

Pesavento, Elena and Barbara Rossi (2005) Do technology shocks drive hours up or down? A little evidence from an agnostic procedure. *Macroeconomic Dynamics* 9, 478–488.

Phillips, Peter C. B. and Pierre Perron, (1988) Testing for a unit root in time series regression. *Biometrika* 75, 335–346.

Robinson, Peter M. (1978) Statistical inference for a random coefficient autoregressive model. *Scandinavian Journal of Statistics* 5, 163–168.

Robinson, Peter M. (1994) Efficient tests of nonstationary hypotheses. *Journal of the American Statistical Association* 89, 1420–1437.

Robinson, Peter M. (1995) Gaussian semi-parametric estimation of long range dependence. *Annals of Statistics* 23, 1630–1661.

Sowell, Fallaw (1992) Maximum likelihood estimation of stationary univariate fractionally integrated time series models. *Journal of Econometrics* 53, 165–188.

APPENDIX

A.1. GIL-ALANA'S MULTIVARIATE FRACTIONAL INTEGRATION TEST

A simple version of the procedure proposed in Gil-Alana (2003a) consists of testing the null hypothesis,

$$H_0: \quad d \equiv (d_1, d_2, \ldots, d_n)' = (d_{10}, d_{20}, \ldots, d_{n0})' \equiv d_0, \tag{A.1}$$

for any real vector d_0, in the model given by (11), where ζ_t is assumed to be an $I(0)$ vector process with positive definite spectral density matrix $f(\lambda)$. Thus ζ_t may be white noise, but it can also accommodate stationary VAR structures. We assume that ζ_t in (11) is generated by a parametric model of the form

$$\zeta_t = \sum_{j=0}^{\infty} A_j(\tau)\omega_{t-j} \quad t = 1, 2, \ldots, \tag{A.2}$$

where ω_t is white noise and W is the unknown variance-covariance matrix of ω_t. The spectral density matrix of ζ_t is then

$$f_\zeta(\lambda; \tau) = \frac{1}{2\pi}\theta(\lambda; \tau)W\theta(\lambda; \tau)^*, \tag{A.3}$$

where $\theta(\lambda; \tau) = \sum_{j=0}^{\infty} A_j(\tau)e^{i\lambda j}$, and θ^* is the complex-conjugate transpose of θ. A number of conditions are required on A and f_ζ to derive the test statistic. The main practical implication is that the spectral density matrix must be finite, with eigenvalues bounded away from zero. It can be shown that a Lagrange multiplier (LM) of the H_0 in (A.1) for (11) takes the form

$$\tilde{S} = T\tilde{b}^T[\tilde{C} - \tilde{D}^T\tilde{E}^{-1}\tilde{D}]^{-1}\tilde{b}, \tag{A.4}$$

where T is the sample size and

$$\tilde{b} = -\frac{1}{T}\sum_{r=1}^{T-1}\psi(\lambda_r)\,\mathrm{tr}\left[I_\zeta(\lambda_r)\tilde{f}(\lambda_r; \tilde{\tau})\right], \tag{A.5}$$

$$\tilde{C} = \frac{4}{T}\sum_{r=1}^{T-1}\psi(\lambda_r)\psi(\lambda_r)^T, \tag{A.6}$$

$$\tilde{D}^T = -\frac{1}{T}\sum_{r=1}^{T-1}\psi(\lambda_r)\left\{\mathrm{tr}\left[\tilde{f}^{-1}(\lambda_r; \tilde{\tau})\frac{\partial\tilde{f}(\lambda_r; \tilde{\tau})}{\partial\tau_1}\right]; \ldots; \mathrm{tr}\left[\tilde{f}^{-1}(\lambda_r; \tilde{\tau})\frac{\partial\tilde{f}(\lambda_r; \tilde{\tau})}{\partial\tau_q}\right]\right\}, \tag{A.7}$$

$$\tilde{E}^{uv} = \frac{1}{2T}\sum_{r=1}^{T-1}\mathrm{tr}\left[\tilde{f}^{-1}(\lambda_r; \tilde{\tau})\frac{\partial\tilde{f}(\lambda_r; \tilde{\tau})}{\partial\tau_u}\tilde{f}^{-1}(\lambda_r; \tilde{\tau})\frac{\partial\tilde{f}(\lambda_r; \tilde{\tau})}{\partial\tau_v}\right], \tag{A.8}$$

where $I_\zeta(\lambda_r)$ is a matrix with the (u, v)th element

$$I_{uv}(\lambda_r) = W_u(\lambda_r)\bar{W}_v(\lambda_r); \quad W_u(\lambda_r) = \frac{1}{\sqrt{2\pi T}} \sum_{t=1}^{T} \tilde{\zeta}_{ut} e^{i\lambda_r t}; \quad \lambda_r = \frac{2\pi r}{T}, \qquad \textbf{(A.9)}$$

where $\bar{W}$ denotes the complex conjugate and $\tilde{f}$ is the estimated spectral density matrix of $\tilde{\zeta}_t$, and where $\tilde{\zeta}_t$ are the reduced-form errors. Finally,

$$\tilde{\tau} = \arg\min_{\tau \in T^*} \left\{ \frac{T}{2} \log|\tilde{f}(\lambda_r; \tau)| + \frac{1}{2} \sum_{r=1}^{T-1} \operatorname{tr}\left[\tilde{f}^{-1}(\lambda_r; \tau) I_\zeta(\lambda_r)\right] \right\}, \qquad \textbf{(A.10)}$$

where T^* is a compact subset of the q-dimensional Euclidean space. Extending the conditions in Robinson (1994), Gil-Alana (2003a) shows that, under H_0 (A.1),

$$\tilde{S} \to_d \chi^2_n \quad \text{as} \quad T \to \infty. \qquad \textbf{(A.11)}$$

A.2. ASYMPTOTIC CONFIDENCE INTERVALS

In this section, we explain how to obtain the confidence intervals of the impulse response of hours worked to the structural technology shock implied by our fractional integration VAR model. For ease of exposition, we reproduce here the model in matrix form:

$$DY_t = \zeta_t,$$
$$\zeta_t = F\zeta_{t-1} + \eta_t,$$

where V is the covariance matrix of η_t. The Monte Carlo exercise can be summarized in three steps:

1. We first perform 500 draws from the independent multivariate normal distributions to build the above equations and obtain the matrix of fractional differencing parameters (D), the matrix of vector autocorrelation of the errors F, and the covariance matrix of η_t (V). The distribution of D, derived by Nielsen (2004) and given in (27) and (28), ensures the joint distribution of the $I(0)$ VAR parameters F and V, as specified by Proposition 11.2 in Hamilton (1994).
2. Given the 500 parameter sets, we compute the 500 impulse response functions of hours worked associated with the structural model in equations (5) and (7) by imposing the restriction that the hours shock does not affect productivity in the long-run.
3. With the standard deviations of the asymptotically distributed 500 impulse response functions, we form the centered 90% asymptotic coverage intervals.