September 14, 2006

Inference for Impulse Responses*

Abstract

This paper derives a closed-form, analytic, asymptotic covariance matrix for structural impulse responses identified with either short-run or long-run restrictions and estimated semi-parametrically with local projections (Jordà, 2005). Based on this covariance matrix, the paper introduces time-profile bands, which are upper and lower bounds with 95% probability coverage on the set of possible impulse response trajectories. These bands are narrower than traditional two standard error bands because they account for the correlation between impulse response coefficients. In addition to several joint hypothesis tests, the paper introduces a method of counterfactual simulation based on choosing alternative policy paths from the set of trajectories in the 95% probability coverage region. Monte Carlo experiments illustrate the small sample properties of the joint hypothesis test and an application to a seven variable monetary system for the U.S. and the U.K. illustrates all of these techniques in practice.

- *Keywords:* impulse response, local projection, time-profile band, counterfactual simulation.
- *JEL Codes:* C32, E47, C53.

Òscar Jordà
Department of Economics
U.C. Davis
One Shields Ave.
Davis, CA 95616
Phone: (530) 752 7021
e-mail: ojorda@ucdavis.edu

*The hospitality of the Federal Reserve Bank of San Francisco during the preparation of this manuscript is gratefully acknowledged.

1 Introduction

Impulse responses remain the primary empirical tool for macroeconomic analysis and constitute a fundamental yardstick by which theoretical models are measured, and intuition about macroeconomic forces is gained. It is therefore surprising what little effort practitioners make in reporting statistical measures of uncertainty along with impulse response estimates. At least two explanations for this observation come to mind: the perception that impulse responses are imprecisely estimated; and the difficulty in deriving a closed-form analytic expression for their covariance matrix.

This paper derives a closed-form, analytic, asymptotic covariance matrix for structural impulse responses identified with either short-run or long-run restrictions and estimated semi-parametrically by local projections (Jordà, 2005). These asymptotic results are the cornerstone of a standard protocol for reporting statistical measures of uncertainty of impulse response estimates and a protocol for counterfactual simulation designed to address, to some extent, the Lucas critique (Lucas, 1976).

Impulse response coefficients are imprecisely estimated but the response trajectories themselves are not. The coefficients of an impulse response are highly correlated and have a natural temporal ordering: the future path of the impulse response today depends on the past through the path that brought us to our current state, but clearly the future cannot change the path already traveled.

This observation suggests that traditional, two standard-error bands vastly overstate the uncertainty about the time-profile of the impulse response. For this reason, I propose a different set of visual displays that capture the uncertainty on the trajectories more precisely and which I call time-profile bands. Time-profile bands are calculated by translating the original coordinate system of the impulse response into an orthogonal coordinate system that preserves the temporal ordering and allows simple construction of 95% confidence upper and lower bounds.

I then argue that a plot of the impulse response should display the estimates, two-standard

error bands and upper and lower time-profile bands along with p-values of the hypothesis test that the impulse response coefficients are jointly zero and the p-value of the test that their cumulative effect is zero. An empirical example with a monetary system involving the U.S. and the U.K. with seven variables and hence 49 different impulse response plots illustrates that these calculations can be easily done even for rather large systems.

Impulse responses estimated by local projections provide a natural environment for counterfactual simulation. I propose counterfactual experiments that consist of alternative policy trajectories. Agents form expectations about the policy responses and insofar as their uncertainty about the true values is well reflected by the statistical uncertainty we can measure, a counterfactual chosen from within the 95% coverage of possible time-profiles would not cause agents to revise their expectations. Conditional on such counterfactual, it is then straight-forward to derive the conditional distribution of the remaining impulse responses. An example of such a counterfactual suggests that when a central bank responds more aggressively to an inflation shock, the unemployment rate is lower than it would otherwise be, even though unemployment responds positively to interest rate shocks.

The paper provides Monte Carlo evidence on the small sample properties of the joint hypothesis tests discussed here and an empirical application illustrating all the techniques introduced in the paper. I begin by deriving the foundational asymptotic results for these tests.

2 Asymptotic Distribution of Structural Impulse Responses

This section provides the asymptotic distribution of reduced-form impulse responses calculated by local projections. The derivations come from Jordà and Kozicki (2006) and provide an asymptotic covariance matrix for all impulse response coefficients, across time and across variables under flexible assumptions about the underlying data generating process. Structural impulse responses can be obtained by imposing a set of restrictions on the nature of the contemporaneous correlations

among the variables in the system. The type of identifying restrictions imposed then determines how the asymptotic covariance matrix of the structural impulse responses should be calculated. I provide formulas for two common approaches: short-run and long-run zero coefficient restrictions.

2.1 Reduced-Form Impulse Responses from Local Projections

Suppose we are interested in the covariance-stationary $r \times 1$ vector of time series $\mathbf{y}_t$, whose Wold decomposition is given by

$$\mathbf{y}_t = \sum_{j=0}^{\infty} B_j \boldsymbol{\varepsilon}_{t-j} \tag{1}$$

where for simplicity and without loss of generality, we drop the constant and any other deterministic terms. From the Wold decomposition theorem (see e.g. Anderson, 1994):

(i) $E(\boldsymbol{\varepsilon}_t) = 0$ and $\boldsymbol{\varepsilon}_t$ are i.i.d.

(ii) $E(\boldsymbol{\varepsilon}_t \boldsymbol{\varepsilon}_t') = \underset{r \times r}{\Sigma_\varepsilon}$

(iii) $\sum_{j=0}^{\infty} \|B_j\| < \infty$ where $\|B_j\|^2 = tr(B_j' B_j)$ and $B_0 = I_r$

(iv) $\det\{B(z)\} \neq 0$ for $|z| \leq 1$ where $B(z) = \sum_{j=0}^{\infty} B_j z^j$

The process in (1) can also be written as:

$$\mathbf{y}_t = \sum_{j=1}^{\infty} A_j \mathbf{y}_{t-j} + \boldsymbol{\varepsilon}_t \tag{2}$$

such that,

(v) $\sum_{j=1}^{\infty} \|A_j\| < \infty$

(vi) $A(z) = I_r - \sum_{j=1}^{\infty} A_j z^j = B(z)^{-1}$

(vii) $\det\{A(z)\} \neq 0$ for $|z| \leq 1$.

In what follows, I take expression (2) as primitive in describing the class of models whose impulse responses we are interested in characterizing.

Jordà's (2005) local projection method of estimating impulse responses is based on the expression that results from simple recursive substitution in the $VAR(\infty)$ representation of expression (2), specifically

$$\mathbf{y}_{t+h} = A_1^h \mathbf{y}_t + A_2^h \mathbf{y}_{t-1} + ... + \boldsymbol{\varepsilon}_{t+h} + B_1 \boldsymbol{\varepsilon}_{t+h-1} + ... + B_{h-1} \boldsymbol{\varepsilon}_{t+1} \tag{3}$$

where:

(i) $A_1^h = B_h$ for $h \geq 1$

(ii) $A_j^h = B_{h-1} A_j + A_{j+1}^{h-1}$ where $h \geq 1$; $A_{j+1}^0 = 0$; $B_0 = I_r$; and $j \geq 1$.

Now consider truncating the infinite lag expression (3) at lag k

$$\mathbf{y}_{t+h} = A_1^h \mathbf{y}_t + A_2^h \mathbf{y}_{t-1} + ... + A_k^h \mathbf{y}_{t-k+1} + \mathbf{v}_{k,t+h} \tag{4}$$

$$\mathbf{v}_{k,t+h} = \sum_{j=k+1}^{\infty} A_j^h \mathbf{y}_{t-j} + \boldsymbol{\varepsilon}_{t+h} + \sum_{j=1}^{h-1} B_j \boldsymbol{\varepsilon}_{t+h-j}.$$

Jordà and Kozicki (2006) show that least-squares estimates of expression (4) produce consistent and asymptotically normal estimates of $A_1^h = B_h$ for $h \geq 1$ as long as, among other technical conditions described in that paper,

1. *k is chosen as a function of the sample size T such that*

$$\frac{k^3}{T} \to 0; T, k \to \infty$$

2. *k is chosen as a function of T such that*

$$k^{1/2} \sum_{j=k+1}^{\infty} \|A_j\| \to 0 \text{ as } T, k \to \infty$$

The main results in Jordà and Kozicki (2006) can be best reproduced with matrix algebra. Hence, define $\mathbf{y}_j$ for $j = h, ..., 1, 0, –1, ..., –k$ as the $(T-k-h)\times r$ matrix of stacked observations of the $1\times r$ vector $\mathbf{y}'_{t+j}$. Additionally, define the $(T-k-h)\times r(h+1)$ matrix $Y \equiv (\mathbf{y}_0, ..., \mathbf{y}_h)$; the $(T-k-h)\times r$ matrix $X \equiv \mathbf{y}_0$; the $(T-k-h)\times r(k-1)+1$ matrix $Z \equiv \left(\mathbf{1}_{(T-k-h)\times 1}, \mathbf{y}_{-1}, ..., \mathbf{y}_{-k+1}\right)$ and the $(T-k-h)\times(T-k-h)$ matrix $M_z = I_{T-k-h} - Z\left(Z'Z\right)^{-1}Z'$. Notice that the inclusion of $\mathbf{y}_0$ in Y is a simplifying notational trick that has no other effect than to ensure that the first block of coefficients is I_r. This will be convenient when deriving the structural impulse response function. Using standard properties of least-squares, impulse responses over h horizons can be jointly estimated as

$$\widehat{B_T}(0,h) = \begin{bmatrix} I_r \\ \widehat{B}_1 \\ \vdots \\ \widehat{B}_h \end{bmatrix} = \left[Y'M_zX\right]\left[X'M_zX\right]^{-1} \tag{5}$$

Jordà and Kozicki (2006) then show that $\widehat{b}_T = vec(\widehat{B_T}(0,h))$ converges in distribution to

$$\sqrt{T}\left(\widehat{b}_T - b_0\right) \xrightarrow{d} N\left(0, \Omega_B\right) \tag{6}$$

where Ω_B can be estimated with $\widehat{\Omega}_B = \left\{\left[X'M_zX\right]^{-1} \otimes \widehat{\Sigma}_v\right\}$. Although properly speaking the equations associated with $B_0 = I_r$ have zero variance, I find it notationally more compact and mathematically equivalent to calculate the residual covariance matrix $\widehat{\Sigma}_v$ as

$$\widehat{\Sigma}_v = \widehat{\Psi}_B\left(I_{h+1} \otimes \widehat{\Sigma}_\epsilon\right)\widehat{\Psi}'_B,$$

where $\widehat{\Psi}_B$ is

$$\underset{r(h+1)\times r(h+1)}{\widehat{\Psi}_B} = \begin{bmatrix} \mathbf{0}_r & \mathbf{0}_r & \mathbf{0}_r & ... & \mathbf{0}_r \\ \mathbf{0}_r & I_r & \mathbf{0}_r & ... & \mathbf{0}_r \\ \mathbf{0}_r & \widehat{B}_1 & I_r & ... & \mathbf{0}_r \\ \vdots & \vdots & \vdots & ... & \vdots \\ \mathbf{0}_r & \widehat{B}_{h-1} & \widehat{B}_{h-2} & ... & I_r \end{bmatrix} \tag{7}$$

with, $\widehat{\Sigma}_\epsilon = \frac{\widehat{\mathbf{v}_1}'\widehat{\mathbf{v}_1}}{T-k-h}$; $\widehat{\mathbf{v}_1} = M_z\mathbf{y}_1 - M_z\mathbf{y}_0\widehat{B}_1$. Therefore, $\widehat{\Omega}_B$ is a simple estimate of the analytic asymptotic covariance matrix of impulse responses across time and across variables. The next section provides formulas to compute *structural* impulse responses and their associated standard errors.

2.2 The Distribution of Structural Impulse Responses

The Wold decomposition for $\mathbf{y}_t$ in expression (1) does not assume that the residuals $\boldsymbol{\varepsilon}_t$ are orthogonal to each other and therefore $E(\boldsymbol{\varepsilon}_t\boldsymbol{\varepsilon}_t') = \Sigma_\varepsilon$ is a symmetric, positive-definite matrix with possibly non-zero entries in the off-diagonal terms. Let the structural residuals $\mathbf{u}_t$ be the rotation of the reduced-form residuals $\boldsymbol{\varepsilon}_t$ given by $P\mathbf{u}_t = \boldsymbol{\varepsilon}_t$, where $E(\mathbf{u}_t\mathbf{u}_t') = I_r$ and hence $\Sigma_\varepsilon = PP'$. Notice that the decomposition of Σ_ε is not unique: Σ_ε contains $r(r+1)/2$ distinct terms but P contains r^2 terms and therefore $r(r-1)/2$ additional conditions are required to achieve just-identification of the terms in P. Traditional methods of estimating P consist in exogenously imposing $r(r-1)/2$, ad-hoc, constraints. Two common approaches are identification via the Cholesky decomposition of Σ_ε (which is equivalent to imposing $r(r-1)/2$ zero restrictions on P); and identification with long-run restrictions that impose $r(r-1)/2$ zero restrictions on the long-run matrix of structural responses, $\Phi_\infty = \sum_0^\infty \Phi_j$.

Consequently, given some estimate $\widehat{P}$ the structural impulse responses Φ_i can be calculated as follows:

$$\widehat{\Phi}(0,h) = \begin{bmatrix} \widehat{\Phi}_0 \\ \widehat{\Phi}_1 \\ \vdots \\ \widehat{\Phi}_h \end{bmatrix} = \widehat{B}(0,h)\,\widehat{P} = \begin{bmatrix} I_r \\ \widehat{B}_1 \\ \vdots \\ \widehat{B}_h \end{bmatrix} \widehat{P} = \begin{bmatrix} \widehat{P} \\ \widehat{B}_1\widehat{P} \\ \vdots \\ \widehat{B}_h\widehat{P} \end{bmatrix} \tag{8}$$

Let $\widehat{\phi}_T = vec\left(\widehat{\Phi}(0,h)\right)$, we want to determine the asymptotic covariance matrix Ω_ϕ since it is clear that

$$\sqrt{T}\left(\widehat{\phi}_T - \phi_0\right) \xrightarrow{d} N(0, \Omega_\phi)$$

2.2.1 Short-Run Identification

When identification is achieved by imposing short-run identification assumptions via the Cholesky decomposition, then

$$\Omega_\phi = \frac{\partial\phi}{\partial b}\Omega_B\frac{\partial\phi}{\partial b'} + \frac{\partial\phi}{\partial vec(P)}\frac{\partial vec(P)}{\partial vech(\Sigma_\varepsilon)}\Omega_\Sigma\frac{\partial vec(P)}{\partial vech(\Sigma_\varepsilon)'}\frac{\partial\phi}{\partial vec(P)'} \tag{9}$$

with $\Omega_\Sigma \equiv E\left[vech(\Sigma_\varepsilon)\,vech(\Sigma_\varepsilon)'\right]$ and $E\left[b, vech(\Sigma_\varepsilon)\right] = 0$ since $E\left[X'M_z\mathbf{v}_1/(T-k-h)\right] \xrightarrow{p} 0$. Since $\Phi(0,h) = B(0,h)\,P$ then it is easy to see that

$$\begin{aligned} \frac{\partial\phi}{\partial b} &= (P' \otimes I_{h+1}) \\ \frac{\partial\phi}{\partial vec(P)} &= (I_r \otimes B(0,h)) \end{aligned}$$

Lütkepohl (2005), chapter 3 provides the additional results

$$\begin{aligned} \frac{\partial vec(P)}{\partial vech(\Sigma_\varepsilon)} &= L_r'\left\{L_r\left(I_{r^2} + K_{rr}\right)(P \otimes I_r)\,L_r'\right\}^{-1} \\ &\quad \sqrt{T}\left(vech\left(\widehat{\Sigma}_\varepsilon\right) - vech(\Sigma_\varepsilon)\right) \xrightarrow{d} N(0, \Omega_\Sigma) \\ \Omega_\Sigma &= 2D_r^+\left(\Sigma_\varepsilon \otimes \Sigma_\varepsilon\right)D_r^{+'} \end{aligned} \tag{10}$$

where L_r is the elimination matrix such that for any square $r \times r$, matrix A, $vech(A) = L_r vec(A)$, K_{rr} is the commutation matrix such that $vec(A') = K_{rr} vec(A)$, and $D_r^+ = (D_r' D_r)^{-1} D_r$, where D_r is the duplication matrix such that $vec(A) = D_r vech(A)$ and hence $D_r^+ vec(A) = vech(A)$. Notice that $D_r^+ = L_r$ only when A is symmetric, but does not hold for the more general case in which A is just a square (but not necessarily symmetric) matrix.

Putting together all of these results, we have,

$$\begin{aligned} & \sqrt{T}\left(\widehat{\phi}_T - \phi_0\right) \xrightarrow{d} N\left(0, \Omega_\phi\right) \\ \Omega_\phi &= \left(P' \otimes I_{h+1}\right) \Omega_B \left(P \otimes I_{h+1}\right) + \\ & 2\left(I_r \otimes B\left(0,h\right)\right) C D_r^+ \left(\Sigma_\varepsilon \otimes \Sigma_\varepsilon\right) D_r^{+'} C' \left(I_r \otimes B\left(0,h\right)'\right) \\ C &= L_r' \left\{L_r \left(I_{r^2} + K_{rr}\right)\left(P \otimes I_r\right) L_r'\right\}^{-1} \end{aligned}$$

where in practice, $\widehat{\Omega}_\phi$ can be calculated by plugging the sample estimates $\widehat{B}\left(0,h\right)$; $\widehat{\Omega}_B$; $\widehat{P}$; and $\widehat{\Sigma}_\varepsilon$ into the previous expression.

2.2.2 Long-Run Identification

If instead structural identification is based on long-run identification assumptions, the infinite order process in expression (2) can be rewritten, without loss of generality, as

$$\mathbf{y}_t = \sum_{j=1}^{\infty} \Psi_j \Delta \mathbf{y}_{t-j} + \Pi \mathbf{y}_{t-1} + \boldsymbol{\varepsilon}_t \tag{11}$$

with $\Psi_i = -\sum_{j=i}^{\infty} A_j$ and $\Pi = \sum_{j=1}^{\infty} A_j$. Under condition (v), then $(I - \Pi)^{-1}$ is the reduced-form, long-run impact matrix. For P the structural rotation matrix such that $P\mathbf{u}_t = \boldsymbol{\varepsilon}_t$, then the structural long-run impact matrix is

$$\Phi_\infty = \left(P^{-1} - P^{-1}\Pi\right)^{-1} = \left(I - \Pi\right)^{-1} P$$

Lütkepohl (2005) then shows that long-run identification assumptions can be easily imposed by applying the Cholesky decomposition to

$$\Phi_{\infty}\Phi_{\infty}' = (I-\Pi)^{-1} PP' (I-\Pi')^{-1} = (I-\Pi)^{-1} \Sigma_{\varepsilon} (I-\Pi')^{-1} = QQ' \tag{12}$$

and hence $P = (I-\Pi) Q$.

An estimate of Π can be easily obtained by redefining M_z in expression (5) as $\widetilde{Z} \equiv \left(\mathbf{1}_{(T-k)\times 1},\ \Delta\mathbf{y}_{t-1},\ ...,\ \Delta\mathbf{y}_{t-k+1}\right)$ with $\widetilde{M}_z = I_{T-k} - \widetilde{Z}\left(\widetilde{Z}'\widetilde{Z}\right)^{-1}\widetilde{Z}$ so that

$$\widehat{\Pi} = \left(\mathbf{y}_0'\widetilde{M}_z\mathbf{y}_{-1}\right)\left(\mathbf{y}_{-1}\widetilde{M}_z\mathbf{y}_{-1}\right)^{-1}$$

with

$$\sqrt{T}\left(\widehat{\pi} - \pi_0\right) \xrightarrow{d} (0, \Omega_{\pi})$$

where $\pi = vec(\Pi)$ and

$$\Omega_{\pi} = \left(\mathbf{y}_{-1}\widetilde{M}_z\mathbf{y}_{-1}\right)^{-1} \otimes \Sigma_{\varepsilon} \tag{13}$$

Σ_{ε} can then be estimated as $\widehat{\Sigma}_{\varepsilon} = \frac{\widetilde{\mathbf{v}}_1\widetilde{\mathbf{v}}_1'}{T-k}$; $\widetilde{\mathbf{v}}_1 = \widetilde{M}_z\mathbf{y}_0 - \widetilde{M}_z\mathbf{y}_{-1}\widehat{\Pi}$.

Given the estimates $\widehat{\Pi}, \widehat{\Sigma}_{\varepsilon}$, then an estimate of Q can be obtained from the Cholesky decomposition described in expression (12) and the structural impulse responses can be constructed as

$$\widehat{\Phi}(0,h) = \widehat{B}(0,h)\left(I - \widehat{\Pi}\right)\widehat{Q} \tag{14}$$

where the asymptotic normality of each element ensures that

$$\sqrt{T}\left(\widehat{\phi}_T - \phi_0\right) \xrightarrow{d} N(0.\Omega_{\phi})$$

but where now

$$\Omega_{\phi} = \frac{\partial \widehat{\phi}_T}{\partial \widehat{b}_T}\Omega_B\frac{\partial \widehat{\phi}_T}{\partial \widehat{b}_T'} + \frac{\partial \widehat{\phi}_T}{\partial \widehat{\pi}_T}\Omega_{\pi}\frac{\partial \widehat{\phi}_T}{\partial \widehat{\pi}_T'} + \frac{\partial \widehat{\phi}_T}{\partial \widehat{q}_T}\left[\frac{\partial \widehat{q}_T}{\partial \widehat{\pi}_T}\Omega_{\pi}\frac{\partial \widehat{q}_T}{\partial \widehat{\pi}_T'} + \frac{\partial \widehat{q}_T}{\partial vech\left(\widehat{\Sigma}_{\varepsilon}\right)}\Omega_{\Sigma}\frac{\partial \widehat{q}_T}{\partial vech\left(\widehat{\Sigma}_{\varepsilon}\right)'}\right]\frac{\partial \widehat{\phi}_T}{\partial \widehat{q}_T'} \tag{15}$$

with $\widehat{q}_T = vech\left(\widehat{Q}_T\right)$; Ω_{Σ} is the covariance matrix of $vech\left(\widehat{\Sigma}_{\varepsilon}\right)$ and we make use of the fact that $\widehat{q}_T$ and $vech\left(\widehat{\Sigma}_{\varepsilon}\right)$ are uncorrelated since $E\left[\mathbf{y}'_{-1}\widetilde{M}_z\widetilde{\mathbf{v}}_1/(T-k)\right] \xrightarrow{p} 0$.

The appendix explains how the following results are derived:

- $\frac{\partial \widehat{\phi}_T}{\partial \widehat{b}_T} = \widehat{Q}'\left(I - \widehat{\Pi}'\right) \otimes I$
- $\frac{\partial \widehat{\phi}_T}{\partial \widehat{\pi}_T} = -\left(\widehat{Q} \otimes \widehat{B}(0,h)\right)$
- $\frac{\partial \widehat{\phi}_T}{\partial \widehat{q}_T} = I \otimes \widehat{B}(0,h)\left(I - \widehat{\Pi}\right)L'$
- $\frac{\partial \widehat{q}_T}{\partial \widehat{\pi}_T} = \left\{\left(\widehat{Q} \otimes I\right)L_r'\right\}^{-1}\left\{\left(I - \widehat{\Pi}\right)^{-1}\widehat{\Sigma}_{\varepsilon} \otimes I\right\}\left\{\left(I - \widehat{\Pi}'\right)^{-1} \otimes \left(I - \widehat{\Pi}\right)^{-1}\right\}$
- $\frac{\partial \widehat{q}_T}{\partial vech\left(\widehat{\Sigma}_{\varepsilon}\right)} = \left\{L\left[\left(I - \widehat{\Pi}\right) \otimes \left(I - \widehat{\Pi}\right)\right](I_{r^2} + K_{rr})\left(\widehat{Q} \otimes I\right)L'\right\}^{-1}$

and L_r is the elimination matrix introduced in expression (10); Ω_B is given by expression (6); Ω_{Σ} is given by expression (10); and Ω_{π} is given by expression (13).

When condition (v) is violated, it is instructive to rewrite expression (11) as

$$\Delta\mathbf{y}_t = \sum_{j=1}^{\infty}\Psi_j\Delta\mathbf{y}_{t-j} + \Gamma\mathbf{y}_{t-1} + \boldsymbol{\varepsilon}_t$$

where $\Gamma = -(I - \Pi)$. Violation of condition (v) occurs when $rank(\Gamma) < r$, in which case Γ has a non-standard asymptotic distribution and Γ is superconsistent, i.e., convergence in distribution occurs at rate T instead of the conventional rate $\sqrt{T}$. Such rank conditions can be tested with a Johansen cointegration test (see Hamilton, 1994, chapters 19 and 20). If $rank(\Gamma) = 0$, then the system has exactly r unit roots and clearly the long-run impact matrix is simply I. The superconsistency of Γ simplifies the derivation of (15): since $\widehat{\Gamma}$ (and hence $\widehat{\Pi}$) converges at rate T,

then the distribution of $\widehat{\phi}_T$ is dominated by the terms converging at rate $\sqrt{T}$ and hence expression (15) simplifies, considerably, to

$$\Omega_\phi = \frac{\partial \widehat{\phi}_T}{\partial \widehat{b}_T} \Omega_B \frac{\partial \widehat{\phi}_T}{\partial \widehat{b}_T'} + \frac{\partial \widehat{\phi}_T}{\partial \widehat{q}_T} \left[\frac{\partial \widehat{q}_T}{\partial vech\left(\widehat{\Sigma}_\varepsilon\right)} \Omega_\Sigma \frac{\partial \widehat{q}_T}{\partial vech\left(\widehat{\Sigma}_\varepsilon\right)'} \right] \frac{\partial \widehat{\phi}_T}{\partial \widehat{q}_T'} \tag{16}$$

where the formulas for each of the terms in the previous expression are the same as those already derived above. Expression (16) is therefore parallel to expression (9) and serves to highlight that any identification scheme based on constraints that do not depend on parameter estimates (irrespective of whether these are zero coefficient restrictions or some other form of linear restriction) will generate a structural covariance matrix that can be calculated on the basis of expression (9). Even when the restrictions imposed depend on coefficient estimates (such as long-run identification restrictions but not limited to these), expression (9) is still valid as long as these coefficients are superconsistent and have no effect on the distribution of terms converging at rate $\sqrt{T}$.

Knowledge of the asymptotic distribution of the structural impulse responses and the associated covariance matrix Ω_ϕ allows one to construct formal statistical tests on a variety of hypothesis of interest that I discuss in the following sections. I begin by introducing an alternative visual representation to the traditional two standard-error bands that I call "time-profile bands." These bands provide a 95% coverage ratio about the uncertainty of the time profile itself rather than on individual coefficients. Next, I discuss a collection of joint hypothesis tests and a method for counterfactual experimentation designed to broadly satisfy the Lucas critique. Monte Carlo experimentation and an empirical application illustrate the properties of all the new techniques introduced.

3 Time-Profile Bands

Error bands for impulse responses should provide visual cues about the uncertainty of the possible time profiles that the responses can follow. Error bands based on the standard errors of individual

coefficients ignore two elementary facts: (1) that impulse response coefficients have a natural temporal ordering – the value of the impulse response today determines the possible trajectories of the response in future periods but future periods can not change the current value of the response; and (2) that impulse response coefficients are usually correlated. It should be clear that, absent any correlation between impulse response coefficients, the simple practice of reporting $\widehat{\phi}_T$ along with $\widehat{\phi}_T \pm 1.96 \times \text{diag}\left(\widehat{\Omega}_\phi\right)^{1/2}$ would indeed provide an approximate probability coverage of 95% confidence under the result that $\widehat{\phi}_T$ is asymptotically normal where $\text{diag}\left(\widehat{\Omega}_\phi\right)$ is the $r^2(h+1) \times 1$ vector that contains the diagonal elements of $\widehat{\Omega}_\phi$, that is, the individual variances of each of the elements in $\widehat{\phi}_T$.

However and as an example, consider the response of the U.S. unemployment rate to a shock in the U.S. federal funds rate from the empirical application in section 7 and displayed in figure 5. The correlation of the response coefficient in the first period with periods 2-12, for example, is 0.69, 0.61, 0.55, 0.51, 0.45, 0.41, 0.36, 0.32, 0.28, 0.23, and 0.23 respectively. That is, even for an impulse response that we would traditionally consider as being essentially zero, the response coefficients are very highly correlated even 12 periods after impact.

Therefore, it is natural to orthogonalize the coordinate system of the impulse response coefficients to ensure uncorrelatedness and preservation of the original temporal ordering. Specifically, suppose we are interested in displaying the impulse response of variable i given a shock to variable k and thus associated with the coefficients $\widehat{\phi}_{i,k}^1, ..., \widehat{\phi}_{i,k}^h$ which can be collected compactly into the vector $\widehat{\boldsymbol{\phi}}_{i,k}$. Notice that we have excluded the coefficient $\widehat{\phi}_{i,k}^0$ because its standard error is the result of the contemporaneous identification assumptions and thus uncorrelated by construction with the remaining coefficients of the impulse response profile. Let $\Omega_\phi(i,k)$ denote the rows and columns of the matrix Ω_ϕ associated with $\widehat{\boldsymbol{\phi}}_{i,k}$, effectively, the variance-covariance matrix of $\boldsymbol{\phi}_{i,k}$, given all the other impulse response coefficients. $\Omega_\phi(i,k)$ is a positive-definite symmetric matrix

and therefore admits a unique Cholesky decomposition such that

$$\Omega_{\phi}(i,k) = A_{i,k} D_{i,k} A'_{i,k}$$

where $A_{i,k}$ is a lower triangular matrix with ones along the main diagonal and $D_{i,k}$ is a diagonal matrix with positive entries. The matrix $A_{i,k}$ has an interesting interpretation: it is the matrix that collects the projection of the response at time 1 onto itself; the projection at time 2 onto the responses at time 1, and 2; and so on. In other words, we project the original responses to obtain orthogonality and to preserve the natural temporal ordering of the response trajectory. These projected coordinates can be collected in the vector $\widehat{\boldsymbol{\psi}}_{i,k} = A_{i,k}^{-1} \widehat{\boldsymbol{\phi}}_{i,k}$, where $\widehat{\boldsymbol{\psi}}_{i,k}$ has a diagonal covariance matrix given by $D_{i,k}$ and each $\widehat{\psi}_{i,k}^{s}$ is the result of a projection on to $\phi_{i,k}^{1}, ..., \phi_{i,k}^{s}$; for $s = 1, ..., h$.

Consequently, the uncorrelatedness of the elements in $\widehat{\boldsymbol{\psi}}_{i,k}$ ensures that the interval $\widehat{\boldsymbol{\psi}}_{i,k} \pm 1.96 \times \text{diag}(D_{i,k})^{1/2}$ contains approximately 95% probability coverage and provides upper and lower limits for the admissible values of $\widehat{\boldsymbol{\psi}}_{i,k}$. Therefore, upper and lower limits for the time profiles of the original impulse response coefficients $\widehat{\boldsymbol{\phi}}_{i,k}$ with 95% confidence can be calculated by translating the orthogonal coordinate system back to the original coordinate system as follows:

$$\begin{aligned} \widehat{\boldsymbol{\phi}}_{i,k}^{+} &= \widehat{\boldsymbol{\phi}}_{i,k} + A_{i,k} \left[1.96 \times \text{diag}(D_{i,k}) \right] \\ \widehat{\boldsymbol{\phi}}_{i,k}^{-} &= \widehat{\boldsymbol{\phi}}_{i,k} - A_{i,k} \left[1.96 \times \text{diag}(D_{i,k}) \right] \end{aligned}$$

Figure 5 displays these 95% probability coverage profile bands for the empirical application of section 7. Overall, the time profile bands are narrower than traditional two standard-error bands. As an example, notice that the profile bands for the response of UK unemployment to the lending rate (the panel in row four, column six of figure 5) suggest that indeed the unemployment rate responds significantly to a shock in short-term interest rates, seven months after impact. This observation finds formal confirmation in a joint hypothesis test (to be described in detail shortly) of the null that all coefficients are zero, which has a p-value of 0.075.

Moreover, a cursory look at the remainder of the panels in figure 5 serves to dispel a common misperception: that impulse responses are estimated too imprecisely to be useful. It is true that individual response coefficients are estimated imprecisely but the time profiles of the impulse responses are clearly not.

The next section introduces a number of joint hypothesis tests that complement the inference represented by the time-profile bands.

4 Joint Tests of Impulse Responses

Commonly reported, two standard-error bands are a measure of unconditional uncertainty for individual impulse response coefficients. They are not a measure of the uncertainty of the time-profile of the impulse response function because impulse response coefficients are usually correlated. Just as we do not use individual t-tests to determine the joint significance of a set of regressors, we should not use two standard-error bands to determine whether, say, an impulse response has a statistically significant profile. This is the main thrust of the argument in the previous section.

The availability of the joint distribution of the structural impulse response coefficients, $\widehat{\phi}_T$, suggests that better measures of uncertainty can be reported by defining more precisely the underlying hypothesis tests of interest, and then constructing the appropriate test statistics. Thus, we can think of traditional two standard-error bands as a graphic representation of the sequence of t-tests associated with the null hypotheses, $H_0 : \phi_i = 0 \quad i = 0, 1, ..., r^2 (h + 1)$. Below I discuss two alternative families of joint hypothesis tests: significance tests and equality tests.

4.1 Joint Significance Tests

A natural and more informative question that frequently arises consists in determining the statistical significance of the impulse response profile of a variable that responds to stimuli from another. Such a question can be cast in terms of the joint null hypothesis:

$$H_0 : \phi^0_{i,k} = \phi^1_{i,k} = ... = \phi^h_{i,k} = 0 \tag{17}$$

where $\phi^j_{i,k}$ denotes the impulse response of variable i to a shock in variable k, at horizon j. In the context of the local projection estimator I use here to derive the asymptotic distribution of the impulse response coefficients, such a null effectively translates into a test of Granger-causality along the lines of the test proposed in Sims (1972) and the more recent semi-parametric test of Angrist and Kuersteiner (2004). Testing the null hypothesis in (17) as well as other general hypotheses involving linear restrictions, can be easily computed by casting the null in terms of the auxiliary matrix Q and the auxiliary vector q, such that

$$H_0 : \underset{J\times r^2(h+1)}{Q} \phi = \underset{J\times 1}{q} .$$

Since we have shown that

$$\sqrt{T}\left(\widehat{\phi}_T - \phi_0\right) \xrightarrow{d} N\left(0, \Omega_\phi\right)$$

then a Wald test can be easily constructed. It will be useful to tailor the notation for the Wald test to see how tests of the null in expression (17) can be constructed in practice. The additional notation will come in handy when discussing some of the other tests presented below.

Let $\underset{1\times r}{R_i} = \left(0, ..0, \underset{i}{1}, 0, .., 0\right)$, that is, R is a zero row vector of dimension $1 \times r$ that contains a one in the i^{th} entry; let $\underset{1\times r}{C_k} = \left(0, ..., 0, \underset{k}{1}, 0, ..., 0\right)$, that is, C is a zero row vector of dimension r that contains a one in the k^{th} entry. Therefore, if we are interested in the response of the i^{th} variable to a shock in the k^{th} variable, all that is required is to place a one in the correct location in the R_i and C_k vectors. Next, construct the selector matrix

$$S_{i,k} \equiv \left[C_k \otimes \left(I_{h+1} \otimes R_i\right)\right]$$

This selector matrix is such that

$$S_{i,k}\widehat{\phi}_T = \widehat{\boldsymbol{\phi}}_{i,k} = \left(\widehat{\phi}^0_{i,k}, ..., \widehat{\phi}^h_{i,k}\right)'$$

so that the covariance matrix associated with these coefficients can be easily constructed from $\widehat{\Omega}_\phi$ as

$$\widehat{\Omega}_\phi(i,k) = S_{i,k}\widehat{\Omega}_\phi S'_{i,k}.$$

Now consider constructing a Wald test for the joint null in expression (17). This null can be cast as

$$\begin{aligned} H_0 \quad &: \quad Q_{i,k}\boldsymbol{\phi}_{i,k} = q_{i,k} \\ Q_{i,k} \quad &= \quad I_r; q = \mathbf{0}_{r,1} \end{aligned}$$

where $\mathbf{0}_{r,1}$ is a vector of zeroes of dimension $r \times 1$. The asymptotic joint normality of $\widehat{\phi}_T$ allows one to construct the Wald statistic associated with the null as

$$\left(Q_{i,k}S_{i,k}\phi - q_{i,k}\right)' \left(Q_{i,k}S_{i,k}\Omega_\phi S'_{i,k}Q'_{i,k}\right)^{-1} \left(Q_{i,k}S_{i,k}\phi - q_{i,k}\right) \xrightarrow{d} \chi^2_{h+1} \tag{18}$$

or more specifically for the null in (17)

$$\phi' S'_{i,k} \left(S_{i,k}\Omega_\phi S'_{i,k}\right)^{-1} S_{i,k}\phi \xrightarrow{d} \chi^2_{h+1}$$

or, in terms of the F-statistic

$$\frac{\phi' S'_{i,k} \left(S_{i,k}\Omega_\phi S'_{i,k}\right)^{-1} S_{i,k}\phi}{h+1} \xrightarrow{d} F_{h+1,T-k-h} \tag{19}$$

where Ω_ϕ can be substituted by its sample counterpart in practice. Thus, we have a method for computing a joint significance test of the null hypothesis that all the coefficients of the response of variable i to variable k are zero, for any pair $i, k \in \{1, ..., r\}$.

This framework can be used directly on other hypotheses of interest. For example, sometimes, the question of interest is whether the cumulative effect of the impulse response is statistically

significant. The null that captures this type of information can be cast by choosing

$$\begin{aligned} Q_{i,k} &= (1, ..., 1) \\ q_{i,k} &= 0 \end{aligned} \tag{20}$$

so that the Wald test in expression (18) has degrees of freedom one or $(1,\ T-k-h)$ degrees of freedom for the F-statistic in expression (19).

4.2 Joint Equality Tests

Sometimes we may be interested in evaluating whether two variables respond in a similar fashion to a given stimulus. We may want to assess whether the impulse response profiles are the same or whether the cumulative effect is the same even when the profiles may differ. Each of these nulls is, respectively:

$$H_0 : \phi^0_{i,k} = \phi^0_{j,l}; ...; \phi^h_{i,k} = \phi^h_{j,l}$$

and

$$H_0 : \phi^0_{i,k} + ... + \phi^h_{i,k} = \phi^0_{j,l} + ... + \phi^h_{j,l}$$

Using the notation introduced above and noticing that

$$\begin{aligned} S_{i,k}\phi &= \boldsymbol{\phi}_{i,k} \\ S_{j,l}\phi &= \boldsymbol{\phi}_{j,l} \end{aligned}$$

define

$$
\begin{aligned}
S &= (S_{i,k} \quad S_{j,k})' \\
Q^J &= (I_{h+1} \quad -I_{h+1}) \, ; q^J = \mathbf{0}_{h+1} \\
Q^C &= (1, ..., 1, -1, ..., -1) \, ; q^C = 0
\end{aligned}
$$

Then, the Wald test of the null that the responses of variables i, and j to a shock in variable k are the same is

$$(Q^m S\phi)' \left(Q^m S \Omega_\phi S' Q'^m\right)^{-1} (Q^m S\phi) \xrightarrow{d} \chi^2_p$$

for $m = J, C$ and $p = h+1, 1$ respectively.

Section 7 provides applications of all the tests discussed in this section whereas section 6 reports Monte Carlo experiments that speak to their small sample properties.

5 Counterfactual Experimentation

The Lucas Critique (Lucas, 1976) warns of the dangers of counterfactual experimentation with empirical models. In real economies, economic agents immediately adapt to the new environment generated by the counterfactual in ways the empirical model cannot anticipate. In essence, the parameters of the empirical model are not constant to the counterfactual so that predictions based on keeping the parameters constant will be "economically" biased. Statistically speaking, of course, there is no impediment to carrying out the calculation – the Lucas Critique simply suggests that it may not be very useful in practice.

This section provides the mechanical details of how to carry a counterfactual simulation in an impulse response analysis. There are at least two justifications why this exercise may be of interest despite the Lucas critique. One is that counterfactuals chosen from within the 95% confidence interval of possible responses would not necessarily generate a revision of the agents' expectations insofar as agents do not know with certainty what the economy's parameters are.

Such a counterfactual would not be at odds (statistically speaking) with the data.

A second justification goes directly to the heart of the assumption that agents are fully rational in practice. For example, Bernanke, Gertler and Watson (1997) have argued in favor of short-run departures of full rationality, Hoover and Jordà (2001) produce evidence that the economy can be better thought of as having a continuum of agents with varying degrees of rational and adaptive behavior, and there is now a rather extensive literature on learning and adaptive behavior (see Evans and Honkapohja, 2001).

These disclaimers notwithstanding, think of the counterfactual experiment as setting the response of the j^{th} variable to a shock in the k^{th} variable to a particular vector of values chosen by the researcher, for example, $\boldsymbol{\phi}_{j,k} = \boldsymbol{\phi}_{j,k}^{c}$. Next, we want to calculate the response of variable i to a shock in variable k conditional on the counterfactual. In other words, we want to calculate the distribution of $\boldsymbol{\phi}_{i,k}$ given $\boldsymbol{\phi}_{j,k} = \boldsymbol{\phi}_{j,k}^{c}$. Since $\sqrt{T}\left(\widehat{\phi}_T - \phi_0\right) \xrightarrow{d} N\left(0, \Omega_\phi\right)$, what we need is to take advantage of the properties of the multivariate normal distribution.

In general, we know that if $\mathbf{y}_1$ and $\mathbf{y}_2$ are two random vectors of generic dimensions with joint normal distribution

$$\begin{bmatrix} \mathbf{y}_1 \\ \mathbf{y}_2 \end{bmatrix} \sim N\left(\begin{bmatrix} \mu_1 \\ \mu_2 \end{bmatrix}; \begin{bmatrix} \Sigma_{11} & \Sigma_{12} \\ \Sigma_{21} & \Sigma_{22} \end{bmatrix}\right)$$

then the conditional distribution of $\mathbf{y}_1$ given $\mathbf{y}_2 = \mathbf{y}_2^c$ is

$$\mathbf{y}_1 | \mathbf{y}_2 = \mathbf{y}_2^c \sim N\left(\mu_{1|2}; \Sigma_{11|2}\right)$$

with

$$\begin{aligned} \mu_{1|2} &= \mu_1 + \Sigma_{12}\Sigma_{22}^{-1}\left(\mathbf{y}_2^c - \mu_2\right) \\ \Sigma_{11|2} &= \Sigma_{11} - \Sigma_{12}\Sigma_{22}^{-1}\Sigma_{21} \end{aligned}$$

Accordingly and using the same definitions of R_i, C_k and thus $S_{i,k}$ introduced previously, note that

$$\left.\begin{array}{l} S_{i,k}\widehat{\phi}_T = \boldsymbol{\phi}_{i,k} \\ S_{j,k}\widehat{\phi}_T = \boldsymbol{\phi}_{j,k} \end{array}\right\} \text{ and } \left[\begin{array}{c} S_{i,k} \\ S_{j,k} \end{array}\right] \Omega_\phi \left[\begin{array}{cc} S'_{i,k} & S'_{j,k} \end{array}\right] = \left[\begin{array}{cc} \Sigma_{ii,k} & \Sigma_{ij,k} \\ \Sigma_{ji,k} & \Sigma_{jj,k} \end{array}\right]$$

where, for example, $\Sigma_{ij,k}$ is the covariance between $\widehat{\boldsymbol{\phi}}_{i,k}$ and $\widehat{\boldsymbol{\phi}}_{j,k}$. The conditional distribution of $\boldsymbol{\phi}_{i,k}$ given $\boldsymbol{\phi}_{j,k} = \boldsymbol{\phi}^c_{j,k}$ is therefore

$$\widehat{\boldsymbol{\phi}}_{i,k} | \left(\boldsymbol{\phi}_{j,k} = \boldsymbol{\phi}^c_{j,k}\right) =$$
$$\widehat{\boldsymbol{\phi}}_{i,k} + S_{i,k}\Omega_\phi S'_{j,k} \left(S_{j,k}\Omega_\phi S'_{j,k}\right)^{-1} \left[\boldsymbol{\phi}^c_{j,k} - \widehat{\boldsymbol{\phi}}_{j,k}\right]$$

with

$$\Sigma_{i,k|j,k} = S_{i,k}\Omega_\phi S'_{i,k} - S_{i,k}\Omega_\phi S'_{j,k} \left(S_{j,k}\Omega_\phi S'_{j,k}\right)^{-1} S_{j,k}\Omega_\phi S'_{i,k}$$

Several remarks deserve comment. First, notice that the second term in the previous expression is a positive definite matrix and hence, the conditional variance given the counterfactual, $\Sigma_{i,k|j,k}$, is smaller than its unconditional counterpart, $\Sigma_{i,k}$. This is a consequence of the observation that, effectively, the counterfactual is replacing the unknown path of the estimated response $\widehat{\boldsymbol{\phi}}_{j,k}$ with the "certain" counterfactual path. Second, the counterfactual is not limited to responses originating from a shock in the same variable – the conditioning arguments do not impose any restriction in this respect. Third, when the correlation between the counterfactual response $\boldsymbol{\phi}_{j,k}$ and the response whose conditional distribution we are interested in calculating, $\boldsymbol{\phi}_{i,k}$, is zero, then $\widehat{\boldsymbol{\phi}}_{i,k} | \left(\boldsymbol{\phi}_{j,k} = \boldsymbol{\phi}^c_{j,k}\right) = \widehat{\boldsymbol{\phi}}_{i,k}$. In some situations, this may, in an of itself, constitute a hypothesis of interest that can be evaluated formally.

The next section investigates the small sample properties of the tests discussed in previous sections and the following section provides an empirical application that illustrates all of the techniques introduced previously.

6 Monte Carlo Experiments

This section investigates the small sample properties of the joint hypothesis tests introduced in section 4. The derivation of the formulas for structural identification are not, strictly speaking, an asymptotic approximation in that the small sample properties of the structural impulse responses are mostly determined by the properties of the reduced-form impulse responses estimated with local projections given the variance of whatever identification method is used. Similarly, the small sample properties of the counterfactual experiments in section 5 are the result of applying standard conditioning arguments to estimates whose small sample properties depend entirely on the statistical properties of the reduced-form impulse responses.

For these reasons, the Monte Carlo experiments in this section are based on the following model

$$\begin{aligned} \mathbf{y}_t &= A\mathbf{y}_{t-1} + \mathbf{u}_t \qquad \mathbf{u}_t \sim N(\mathbf{0}, I_3); \qquad \mathbf{y}_t = (y_{1t}, y_{2t}, y_{3t})' \\ A &= \begin{bmatrix} 0.25 & 0.375 & 0 \\ 0.375 & 0.5 & 0 \\ 0.3 & 0 & 0.75 \end{bmatrix} \end{aligned} \tag{21}$$

for sample sizes of 100, 250, and 500 observations after disregarding another 500 observations to avoid initialization problems. The impulse responses that result from simulating data from (21) are displayed in figure 1. The figure displays the true impulse responses along with error bands and joint hypothesis tests based on the Monte Carlo covariance matrix given least-squares estimates of the $VAR(1)$ model in expression (21) repeated 1,000 times with a sample size of 250 observations. These measures of uncertainty are meant to provide the reader an illustration of the type of small sample variation induced in the experiments.

Figure 1 explains the specific choice of parameter values in (21). The responses of y_{1t} and y_{2t} to a shock in y_{3t} are exactly zero and are meant to provide information about the size of the joint

hypothesis tests. The responses of $\mathbf{y}_t$ to its own shocks and displayed along the main diagonal of figure 1, are a scaled version of one another with impulse response coefficients in period 1 equal to 0.25, 0.5, and 0.75 respectively. I choose these values to obtain a power curve of the statistics of interest as a function of the distance to the null. The response of variable y_{1t} to a shock in y_{2t} is designed to be the same as the response of variable y_{2t} to a shock in y_{1t} so as to be able to evaluate the joint equality tests. Because these responses are themselves scaled versions of the responses in the main diagonal, I will be able to obtain a power curve for this test as well.

In particular, figures 2, 3, and 4 each display four graphs each, corresponding to Monte Carlo experiments with sample sizes of 100, 250, and 500 observations respectively. The top left panel in each figure displays the power curve of the joint-hypothesis test of the null that all the impulse response coefficients are jointly zero for horizons that vary between 1 to 8 periods. When the horizon is set to one period, the joint hypothesis test is equivalent to a t-test on the first coefficient. The more periods with a zero response included in the test, the lower the power of the test will be, just as in a classical regression context. The numbers on the horizontal axis are the values of the first coefficient of the impulse response to which they refer. For example, 0.5 corresponds to the response of y_{2t} to its own shock and so on. The top right panel instead examines the power curve when the null is that the joint cumulative effect is zero. The curves are derived for a nominal size of 5% for a conventional 95% confidence level test.

The bottom panels of figures 2, 3, and 4 refer to the joint equality tests. The left panel tests the null that all the coefficients between two impulse responses are equal to each other while the right panel tests only that their cumulative effect is the same. The null is based on comparing the responses of y_{1t} to a shock in y_{2t} and vice versa. Thereafter, I compare the nulls based on the response of y_{2t} to its own shock with the response of y_{2t} to a shock in y_{1t} (labeled 0.5 vs. 0.375) and the zero response of y_{1t} to a shock in y_{3t} relative to the the responses of y_{1t}, y_{2t}, and y_{3t} to their own shocks (labeled 0 vs. 0.25; 0 vs. 0.50; and 0 vs. 0.75, respectively). As the horizon grows from 1 to 8 periods, notice that the impulse responses go back to their long-run value of

zero. This has two consequences: (1) it makes rejecting the joint null that the responses are zero more difficult as coefficients in later periods are more imprecisely estimated and they are closer to zero; and (2) it makes rejecting the joint null of equality more difficult for much the same reasons.

Several results deserve comment. First, the hypothesis test of the null that all the coefficients are zero and displayed in the top left panel of figures 2, 3, and 4 is the best behaved. There are no significant size distortions (the actual size is very close to the nominal 5% at horizon 1 and increases to about 10% when eight periods are jointly considered) and as the sample increases, the power of the test increases very quickly toward 1. Second, the test on the null that the joint cumulative effect is zero and displayed in the top right panel of figures 2, 3, and 4, has good power properties as well but tends to be somewhat oversized in small samples the more horizons in the impulse response are included. The actual size is close to the nominal 5% for one period (4.4%) but increases to about 30% for eight periods. Third, the test for joint equality displayed in the bottom left panel of figures 2, 3, and 4 has weaker power properties than the two tests discussed previously and tends to be undersized – actual size is consistently below 1% for a nominal size of 5%. The cumulative equality test displayed in the bottom right panel suffers from the same problems with size but has better power properties, perhaps not surprisingly as the null is less demanding – the requirement is that the cumulative sum of coefficients be equal across impulse responses but not that each coefficient be equal to each other.

These Monte Carlo experiments suggest the joint hypothesis tests presented here have good power properties that rapidly increase with the sample size. However, the experiments also suggest that the size of the tests can be distorted very quickly when one considers impulse response coefficients at long horizons that are imprecisely estimated and/or are close to zero. While this distortion is not surprising (even in the context of more traditional joint hypothesis test settings, the result is well known), it raises a cautionary note to empirical practitioners: parsimony in the impulse response horizon considered is desirable.

Admittedly, a more exhaustive Monte Carlo investigation of the properties of the statistics

discussed in the paper would be beneficial but space considerations suggest this is better left for a different paper. For example, there is now a fair amount of work (see Pesavento and Rossi, forthcoming *a* and *b*) that investigates appropriate adjustments to the distribution of impulse response coefficients in systems that have near-unit roots. Presumably, the same distortions induced by high persistence are likely to appear in local projection estimates and it would be interesting to explore this and other features in a different paper that goes beyond the scope of this one.

7 Policy Trade-offs in the U.S. and in the U.K.

The philosophy of this section is entirely empirical: the essential goal is to provide a summary of the basic features of the data considered here by means of impulse response functions estimated semi-parametrically with local projections. The application is to a monetary system that includes two countries and their bilateral exchange rate. The exercise is not meant to be the definitive answer to the enigmas of monetary policy but rather to provide intuition during the presentation that follows.

7.1 Overview

The data is for a system of seven variables and includes the unemployment rate (in percent), consumer price inflation (in percent) and the federal funds rate (in percent) for the U.S.; the unemployment rate (in percent), retail price inflation (in percent), and the Bank of England's lending rate (in percent) for the U.K.; and the U.S. Dollar/British Pound exchange rate (in logs). The data is available monthly, beginning January 1971 and ending December 2005. Hence, for each country, we can think of the three variable sub-system that includes unemployment, inflation and the policy short-term rate as a variant of the New-Keynesian framework often used in the literature that investigates optimal policy rules (see, e.g. Walsh 2003, chapter 11). Exploring a

system of seven variables is meant to ensure the end-user that the techniques introduced in this paper can be applied to relatively large systems.

This system results in 49 impulse responses that I estimate jointly by local projections with equation (5) over a horizon of twelve periods (one year). The lag length for the projections is selected automatically by AIC_c[1] to be four. In order to obtain structural impulse responses, I apply short-run zero coefficient constraints implied by the Wold-causal ordering of the variables in the same order in which they are described above. A more thorough investigation of the appropriate contemporaneous identification assumptions seems warranted but for the purposes of keeping the discussion within the parameters of the techniques described in the paper, it is left for future research.

Figure 5 displays these 49 impulse response functions. Each panel displays the following information: (1) traditional, two standard-error bands; (2) 95% confidence level, time-profile bands as described in section 4; (3) the p-value of the joint significance test of the responses labeled "Joint" and; (4) the p-value of the cumulative joint significance test labeled "Cum," both described in section 4.

There are a number of natural questions of economic interest, for example: (1) what is the sensitivity of the policy interest rate to shocks in the unemployment, inflation, and exchange rates; (2) what is the response of unemployment, inflation and exchange rates to monetary shocks; (3) what is the sensitivity of domestic policy rates to shocks in the foreign policy rate; (4) what is the response of exchange rates to shocks in inflation; and (5) what is the response of the unemployment rate when the central bank responds more aggressively to inflation shocks.

Answers to question (1) can be used to compare the relative emphasis that each central bank places on growth and price stabilization. Question (2) establishes the effectiveness of monetary policy; question (3) determines whether policy changes in one central bank influence policy changes

[1] AIC_c refers to the correction to AIC introduced in Hurvich and Tsai (1989), which is specifically designed for autoregressive models. There were no significant differences when using SIC or the traditional AIC.

in the other; question (4) speaks loosely about the relative merits of the purchasing power parity condition; and question (5) measures how different is the response of the unemployment rate when the central bank's response to shocks in inflation is more aggressive. Each of these five questions requires different types of inference based on joint hypothesis tests and counterfactual simulations of the type introduced in previous sections.

7.2 Results

The profile bands introduced in section 3, the joint significance test in expression (17) and the joint cumulative test in expression (20) all provide relevant information that should be reported in any impulse response graph. To explain the value of each element more specifically, consider the response of the U.S. unemployment rate to a shock in the federal funds rate displayed in the first row, third column of the panel of impulse responses in figure 5. The joint significance test has a p-value of 0.326, and the joint cumulative test a p-value of 0.450. Both tests suggests that the federal funds rate has no effect on the unemployment rate at a year horizon. However, the profile bands suggest that the effect is statistically significant for the last three months. Indeed, the response of U.S. unemployment is very similar to the response of UK unemployment to a shock in the Bank of England lending rate (which plays a similar role to the federal funds rate in the U.S.) and displayed in row four, column six of figure 5. The joint significance test has a p-value of 0.075 and the joint cumulative test a p-value of 0.143 but notice that, like the U.S., there is basically no response of the unemployment rate for the first six to seven periods after impact, after which the unemployment rate steadily rises. It should be clear then that the results of the joint significance tests are dominated by the behavior of the unemployment rate response over the first six to seven periods. The profile bands articulate this case more clearly by showing what happens in the remaining six months: the time profile steers away from zero significantly.

This simple example argues in favor of using all three forms of uncertainty measures (profile bands, joint significance and cumulative tests) in a complementary manner as they convey infor-

mation on different features of the impulse response. As another illustration, consider the response of the U.S. Dollar/British Pound exchange rate to a shock in U.S. inflation and displayed in row seven, column two of figure 5. This is an example of a response where the joint significance test has a p-value of 0.344 but the joint cumulative test has a p-value of 0.055, a result that confirms the pattern displayed by the profile bands, which suggest that the U.S. Dollar significantly depreciates between the third and the eleventh/twelfth month after impact.

Returning now to the questions posed in the previous subsection, recall that in question (1) we are trying to assess the relative sensitivity of the central bank to shocks in unemployment, inflation and exchange rates. This information is summarized in row three, columns one, two and seven of figure 5 for the U.S. and row six for the U.K. In both countries, interest rates drop significantly by about 50 basis points in response to a shock in the unemployment rate, with both significance tests and profile bands indicating a statistically significant response (the p-values are essentially zero). However, tests of the joint and cumulative equality of the responses are rejected pretty decisively (with p-values of 0.000 and 0.000 respectively). Perhaps this last result is not surprising: the U.S. drops interest rates more quickly and keeps them low for a longer period than the U.K. does. In contrast, both countries do not respond to a shock in inflation: both significance tests have p-values well above the normative 0.05 value and the profile bands are generally not significant at any horizon. Finally, while joint significant tests do not indicate that interest rates respond to fluctuations in exchange rates, the U.S. displays a significant cumulative effect of interest rates (p-value = 0.015) in response to a depreciation of the U.S. Dollar, a result that is corroborated by the profile bands.

Question (2) examines the relative effectiveness of monetary policy. The relevant panels are rows one, two and seven of column three for the U.S. and rows four, five and seven in column six for the U.K. The response of unemployment to a shock in interest rates is very similar in both countries although only the joint significance test for the U.K. has a p-value bordering on significance at 0.075. Unemployment remains essentially flat six to seven months after impact

and then steadily climbs (in both countries, the profile bands suggest this climb is significant). More formally, joint and cumulative equality tests have p-values of 0.52 and 0.83 which confirm the similarity in the responses.

Inflation in both countries tends to climb in response to an interest rate shock but only significantly in the U.K. (the joint cumulative test has a p-value of 0.004). The joint equality test cannot reject the null, attaining a p-value of 0.43 although the joint cumulative equality test does, with a p-value of 0.055. Both results seem at odds with what economic theory would predict although this price puzzle has been detected many times before in the U.S. (see e.g. Sims, 1992).

Before discounting these results, it is important to examine the response of the exchange rate. The U.S. Dollar tends to appreciate somewhat throughout the year after impact when interest rates increase (the joint cumulative test has a p-value of 0.090 and the profile bands border on the zero line). The British pound significantly depreciates on impact but then appreciates for the remainder of the year, although not in a statistically significant way. Hence, although the initial responses of the exchange rate are consistent, the uncovered interest rate parity condition seems to hold only somewhat for the U.K.

In terms of the effect that each central bank has on the other (row three, column six for the U.S. and row six, column three for the U.K.), neither country exhibits a significant response although these tend to remain on the positive side, suggesting that both countries tend to move interest rates in the same direction.

The final question has to do with the response of the exchange rate to a shock in inflation and is displayed by the panels in row seven, column two for the U.S. and column five for the U.K. The U.S. Dollar tends to depreciate in response to a positive shock in inflation, as purchasing power parity would predict: the joint cumulative test has a p-value of 0.055 and the profile bands are significant for about ten out of the twelve months displayed. In contrast, the British Pound does not display any appreciable response at any horizon, an observation confirmed by the joint significance and cumulative tests with p-values of 0.78 and 0.64 respectively.

Finally, consider the counterfactual simulation that investigates the effects of a more aggressive response of monetary policy to inflation shocks in the U.S. and in the U.K. The panel in the third row, second column of figure 5 displays the response of the U.S. federal funds rate to a shock in the U.S. inflation rate. This response is not statistically significant: the joint significance and cumulative tests have p-values of 0.950 and 0.546 respectively (0.396 and 0.688 for the U.K.). Hence, I choose a counterfactual experiment that constrains the response of the fed funds rate/lending rate to an inflation shock to be that corresponding to the upper, two standard-error band. A more conservative choice (insofar as we want to avoid the Lucas critique) would have been to choose the upper boundary of the time-profile band instead. However, the former experiment provides a better illustration of the counterfactual framework introduced in section 5 although we also compute the latter counterfactual.

Before reporting on the counterfactual itself, notice that the response of the U.S. unemployment rate to a positive inflation shock results in a statistically significant response of unemployment as evinced in row one, column 2 of figure 5 by the time-profile bands and joint and cumulative tests, with p-values 0.050 and 0.048 respectively. In contrast, the UK counterpart displayed in row four, column 5 of figure 5, is virtually zero with corresponding p-values 0.552 and 0.867.

The counterfactual experiment is displayed in figure 6 and shows that a more aggressive response of the central bank to an inflation shock has similar consequences in both countries: surprisingly the unemployment rate is lower than it would otherwise be. Figure 6 displays the original and the counterfactual response along with the profile bands under the counterfactual. While the response of the unemployment rate in the UK changes little and remains statistically insignificant, the response in the US is economically and statistically meaningful: the cumulative increase of the unemployment rate in response to a 0.30% increase in inflation is 0.46% historically but only 0.12% in the counterfactual. If instead we had chosen the upper time-profile band for the counterfactual, the difference would not have been statistically significant and the cumulative effect would have declined instead to 0.36%. However a 0.10% decline with respect to 0.46% is an economically

sizeable decline. These results are surprising considering that an increase in the federal funds rate generates higher unemployment according to the panel in row one, column three of figure 5.

8 Conclusion

Impulse responses are a summary statistic of the dynamic properties of a vector time series. Associated to these statistics, there is an asymptotic distribution from which formal hypothesis tests can be constructed. This paper derives a standard set of inferential tools for empirical analysis with impulse responses. An important message of the paper is that impulse response time-profiles are estimated rather accurately; that hypothesis tests on these time-profiles have good statistical properties; that meaningful economic hypothesis can be easily tested based on these response estimates; and that economically useful counterfactual simulations can be formally conducted.

The foundation for these results is the local projection semi-parametric estimator introduced by Jordà (2005). This general estimator is not only more robust to model misspecification (in fact, the starting assumption is that the data is generated by an infinite order process), it provides a simple method to derive a closed-form, analytic expression for the covariance matrix of structural impulse responses identified through either short-run or long-run restrictions, the two most common approaches.

This structural covariance matrix is instrumental to derive joint inference based on the Wald principle for a broad set of hypothesis of interest, some of which are discussed in the paper and their small sample properties investigated with Monte Carlo experiments. The covariance matrix also allows one to construct a new graphical representation of the uncertainty surrounding the impulse response time-profiles (rather than the individual coefficients) and whose results complement the information provided by joint hypothesis tests.

The asymptotic distribution of impulse response coefficients also provides an intuitive method

for counterfactual simulation based on standard conditioning methods. The set-up consists in choosing a dynamic path for a policy variable (or variables) from the family of trajectories within the 95% probability coverage region. Such a choice is less likely to violate the precepts of the Lucas critique and is simpler to articulate: one does not require knowledge of the form of the policy equation or its underlying structural parameters.

9 Appendix

I provide here the detailed derivations required to derive the covariance of the structural impulse responses derived by imposing long-run identification assumptions. First notice that there are three different ways of expressing the vec version of expression (14), specifically

$$\widehat{\phi}_T = vec\left(\widehat{\Phi}\left(0,h\right)\right) = \begin{array}{c} \left(\widehat{Q}'\left(I-\widehat{\Pi}\right)\otimes I\right)\widehat{b}_T \\ \left(\widehat{Q}\otimes\widehat{B}\left(0,h\right)\right)vec\left(I-\widehat{\Pi}\right) \\ \left(I\otimes\widehat{B}\left(0,h\right)\left(I-\widehat{\Pi}\right)\right)vec\left(\widehat{Q}\right) \end{array}$$

from where we obtain $\frac{\partial\widehat{\phi}_T}{\partial\widehat{\pi}_T}$, $\frac{\partial\widehat{\phi}_T}{\partial\widehat{\pi}_T}$, and $\frac{\partial\widehat{\phi}_T}{\partial\widehat{q}_T}$ by realizing that $dvec\left(I-\widehat{\Pi}\right) = -dvec\left(\widehat{\Pi}\right)$ and since Q is lower triangular, $vec\left(Q\right) = L'vech(Q)$.

Next, I derive the expressions for $\frac{\partial\widehat{q}_T}{\partial\widehat{\pi}_T}$ and $\frac{\partial\widehat{q}_T}{\partial vech\left(\widehat{\Sigma}_\varepsilon\right)}$ by first noticing that

$$\left(I-\Pi\right)^{-1}\Sigma_\varepsilon\left(I-\Pi'\right)^{-1} = QQ'$$

so that

$$d\left(I-\Pi\right)^{-1}\Sigma_\varepsilon\left(I-\Pi'\right)^{-1} + \left(I-\Pi\right)^{-1}\Sigma_\varepsilon d\left(I-\Pi'\right)^{-1} + \tag{22}$$

$$\left(I-\Pi\right)^{-1}d\Sigma_\varepsilon\left(I-\Pi'\right)^{-1} = dQQ + QdQ'$$

Begin by setting $d\left(I-\Pi\right)^{-1} = 0$ then

$$(I-\Pi)^{-1} d\Sigma_{\varepsilon} (I-\Pi')^{-1} = dQQ + QdQ'$$

Taking the vec operator on both sides of this expression

$$\begin{aligned}\left[(I-\Pi)^{-1} \otimes (I-\Pi)^{-1}\right] dvec(\Sigma_{\varepsilon}) &= (Q \otimes I)\, dvec(Q) + (I \otimes Q)\, K_{rr} dvec(Q) \\ dvec(\Sigma_{\varepsilon}) &= [(I-\Pi) \otimes (I-\Pi)]\, (I_{r^2} + K_{rr})\, (Q \otimes I)\, dvec(Q)\end{aligned}$$

using the same rule for the right hand side as in the derivation of the short-run identification case. Finally, using the elimination matrix L_r introduced in expression (10) and noticing that since Q is lower triangular then $L_r' vech(Q) = vec(Q)$, we arrive at the desired result

$$\frac{\partial q}{\partial vech(\Sigma_{\varepsilon})} = \{L[(I-\Pi) \otimes (I-\Pi)]\, (I_{r^2} + K_{rr})\, (Q \otimes I)\, L'\}^{-1}$$

To derive $\frac{\partial \widehat{q}_T}{\partial \widehat{\pi}_T}$, return to expression (22) and instead set $d\Sigma_{\varepsilon} = 0$ so that

$$d(I-\Pi)^{-1} \Sigma_{\varepsilon} (I-\Pi')^{-1} + (I-\Pi)^{-1} \Sigma_{\varepsilon} d(I-\Pi')^{-1} = dQQ + QdQ'$$

Taking the vec operator on both sides of the expression and using similar manipulations as in the previous derivation, it is easy to see that we arrive at

$$\left[(I-\Pi)^{-1} \Sigma_{\varepsilon} \otimes I\right] dvec\left\{(I-\Pi)^{-1}\right\} = (Q \otimes I)\, dvec(Q)$$

where the term $(I_{r^2} + K_{rr})$ cancels on both sides of the previous expression. It is straight forward to see then that

$$dvec\left\{(I-\Pi)^{-1}\right\} = \left[(I-\Pi)^{-1} \otimes (I-\Pi')^{-1}\right] dvec(\Pi)$$

and since $L_r' vech(Q) = vec(Q)$, then we arrive at the desired result,

$$\frac{\partial q}{\partial \pi} = \{(Q \otimes I)\, L_r'\}^{-1} \left\{(I-\Pi)^{-1} \Sigma_{\varepsilon} \otimes I\right\} \left\{(I-\Pi')^{-1} \otimes (I-\Pi)^{-1}\right\}$$

References

Anderson, Theodore W. (1994) **The Statistical Analysis of Time Series Data**. New York, New York; Wiley-Interscience.

Angrist, Joshua D. and Guido Kuersteiner (2004) "Semiparametric Causality Tests Using the Policy Propensity Score," National Bureau of Economic Research, working paper 10975.

Bernanke, Ben S., Mark Gertler and Mark W. Watson (1997) "Systematic Monetary Policy and the Effects of Oil Price Shocks," *Brookings Papers on Economic Activity*, 1, 91-157.

Evans, George W. and Seppo Honkapohja (2001) **Learning and Expectations in Macroeconomics**. Princeton, New Jersey: Princeton University Press.

Hoover, Kevin D. and Òscar Jordà (2001) "Measuring Systematic Monetary Policy," *Review*, Federal Reserve Bank of st. Louis, 83(4): 113-138.

Hurvich, Clifford M. and Chih-Ling Tsai (1989) "Regression and Time Series Model Selection in Small Samples," *Biometrika*, 76(2): 297-307.

Jordà, Òscar (2005) "Estimation and Inference of Impulse Responses by Local Projections," *American Economic Review*, 95(1): 161-182.

Jordà, Òscar and Sharon Kozicki (2006) "Projection Minimum Distance: An Estimator for Dynamic Macroeconomic Models," U.C. Davis, working paper 6-23.

Lucas, Robert E., Jr. (1976) "Econometric Policy Evaluation: A Critique," *Journal of Monetary Economics*, Supplementary Series 1976, 1(7): 19-46.

Lütkepohl, Helmut (2005) **New Introduction to Multiple Time Series**. Berlin, Germany: Springer-Verlag.

Pesavento, Elena and Barbara Rossi (forthcoming, a) "Small Sample Confidence Bands for Multivariate Impulse Response Functions," *Journal of Applied Econometrics.*

Pesavento, Elena and Barbara Rossi (forthcoming, b) "Impulse Response Confidence Intervals for Persistent Data: What Have We Learned?" *Journal of Economic Dynamics and Control.*

Sims, Christopher A. (1972) "Money, Income and Causality," *American Economic Review,* 540-562.

Sims, Christopher A. (1992) "Interpreting the Macroeconomic Time Series Facts: The Effects of Monetary Policy," *European Economic Review,* 36(10): 975-1000.

Walsh, Carl E. (2003) **Monetary Theory and Policy, second edition.** Cambridge, MA: MIT Press.

Figure 1 – True Impulse Responses and Monte Carlo Standard Errors

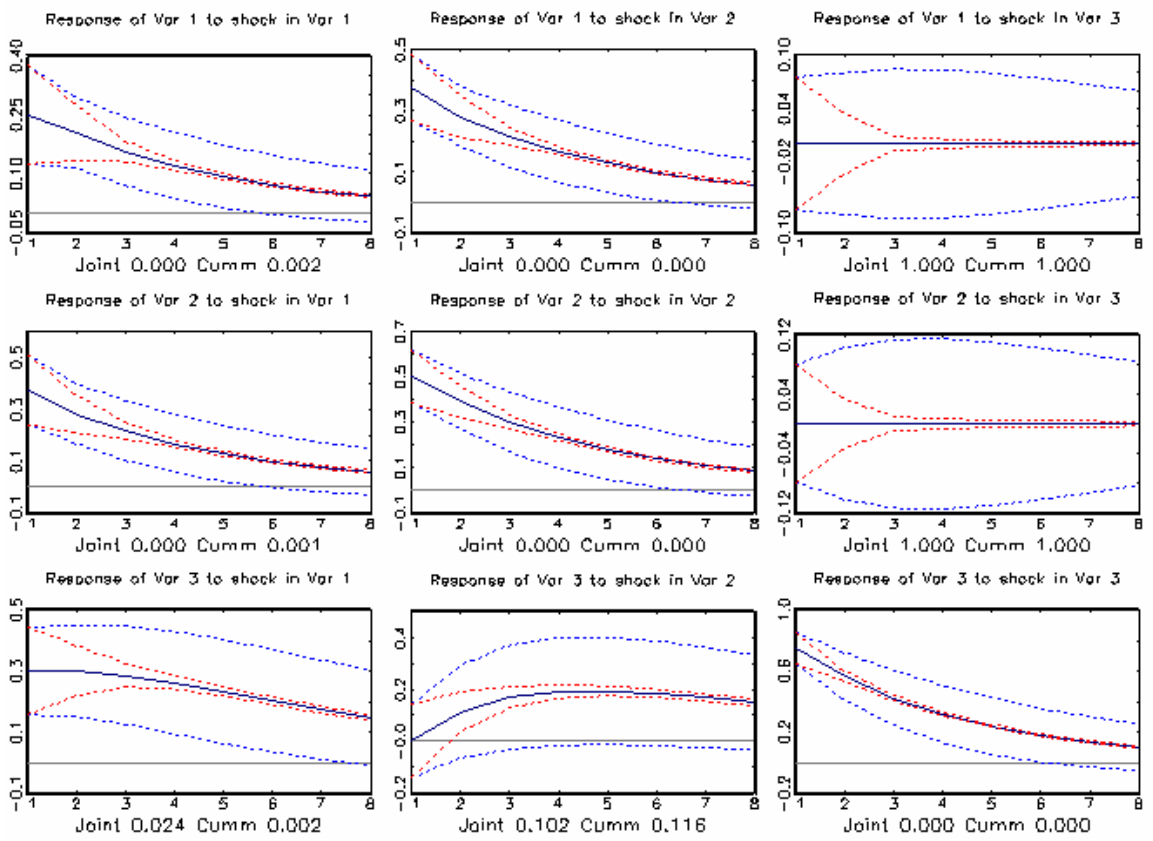

Notes: True impulse responses based on the VAR(1) model in expression (21) and used in the Monte Carlo experiments of section 6. The figure displays conventional two standard-error bands (the dashed, wide bands) and time-profile bands (the dashed but narrow bands) based on estimating the VAR(1) on a simulated sample with 250 observations and replicated for 1,000 times. "Joint" refers to the p-value of the null that all coefficients of the impulse response displayed are zero and "Cumm" refers to the p-value of the null that the cumulative effect of the impulse response displayed is zero.

Figure 2 – Power Curves of the Joint Hypothesis Test for a Sample Size = 100

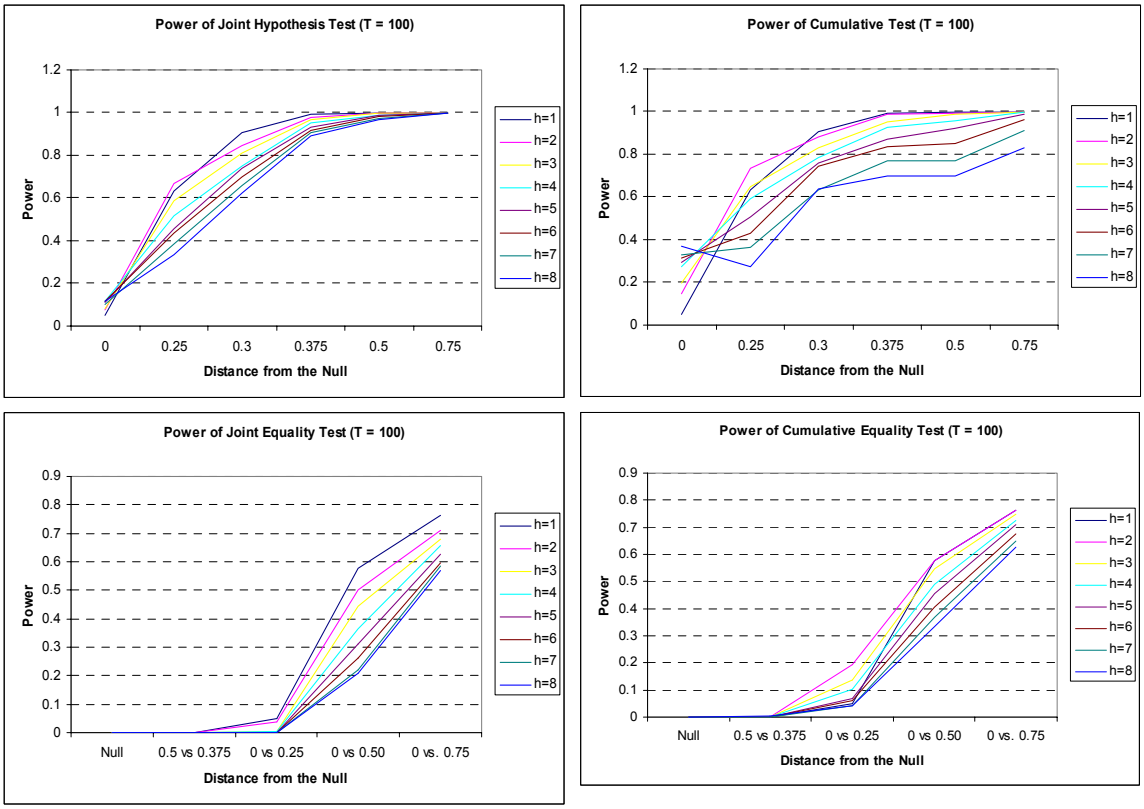

Notes: Calculations are based on drawing samples of 100 observations (after disregarding an initial 500 observations), drawn 1,000 times. The "Joint Hypothesis Tests" refers to the null that all the coefficients are zero; the "Cumulative Test" refers to the null that the sum of the coefficients of the impulse response is zero; the "Joint Equality Test" refers to the null that two impulse responses have equal coefficients; and the "Cumulative Equality Test" refers to the null that the cumulative effect of these two impulse responses is the same.

Figure 3 – Power Curves of the Joint Hypothesis Test for a Sample Size = 250

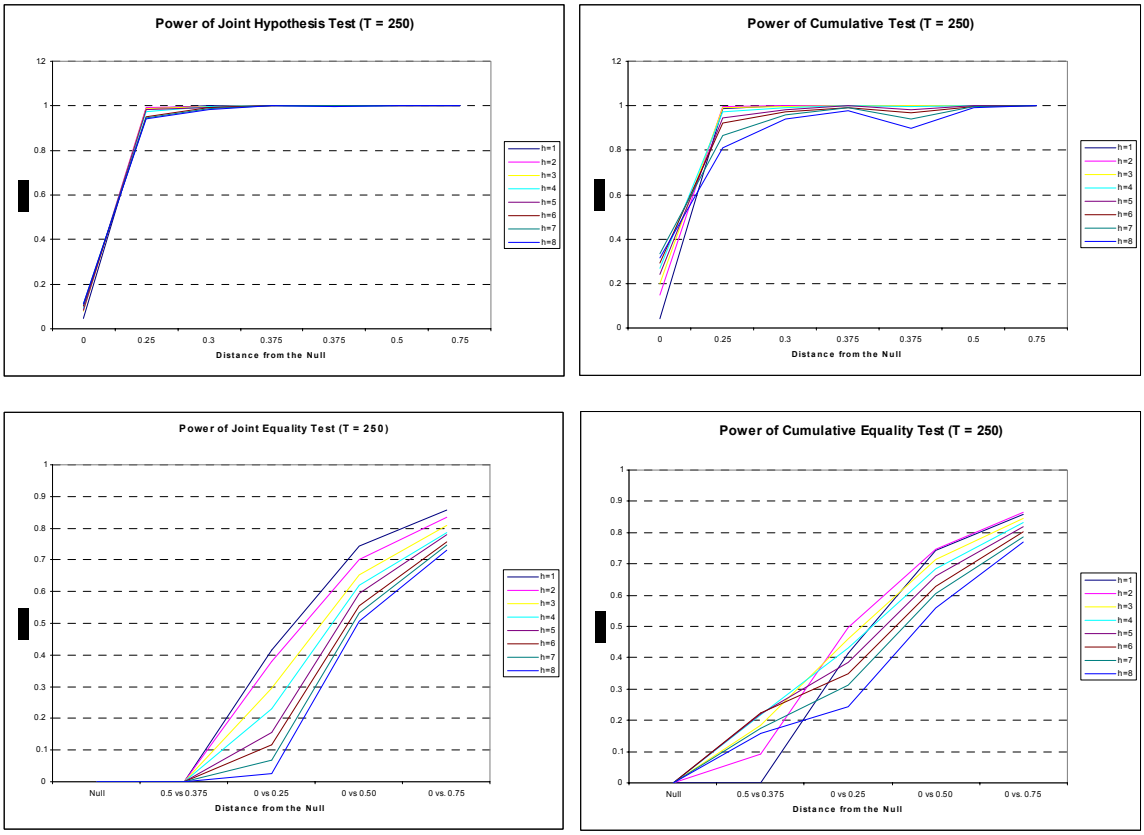

Notes: Calculations are based on drawing samples of 100 observations (after disregarding an initial 500 observations), drawn 1,000 times. The "Joint Hypothesis Tests" refers to the null that all the coefficients are zero; the "Cumulative Test" refers to the null that the sum of the coefficients of the impulse response is zero; the "Joint Equality Test" refers to the null that two impulse responses have equal coefficients; and the "Cumulative Equality Test" refers to the null that the cumulative effect of these two impulse responses is the same.

Figure 4 – Power Curves of the Joint Hypothesis Test for a Sample Size = 500

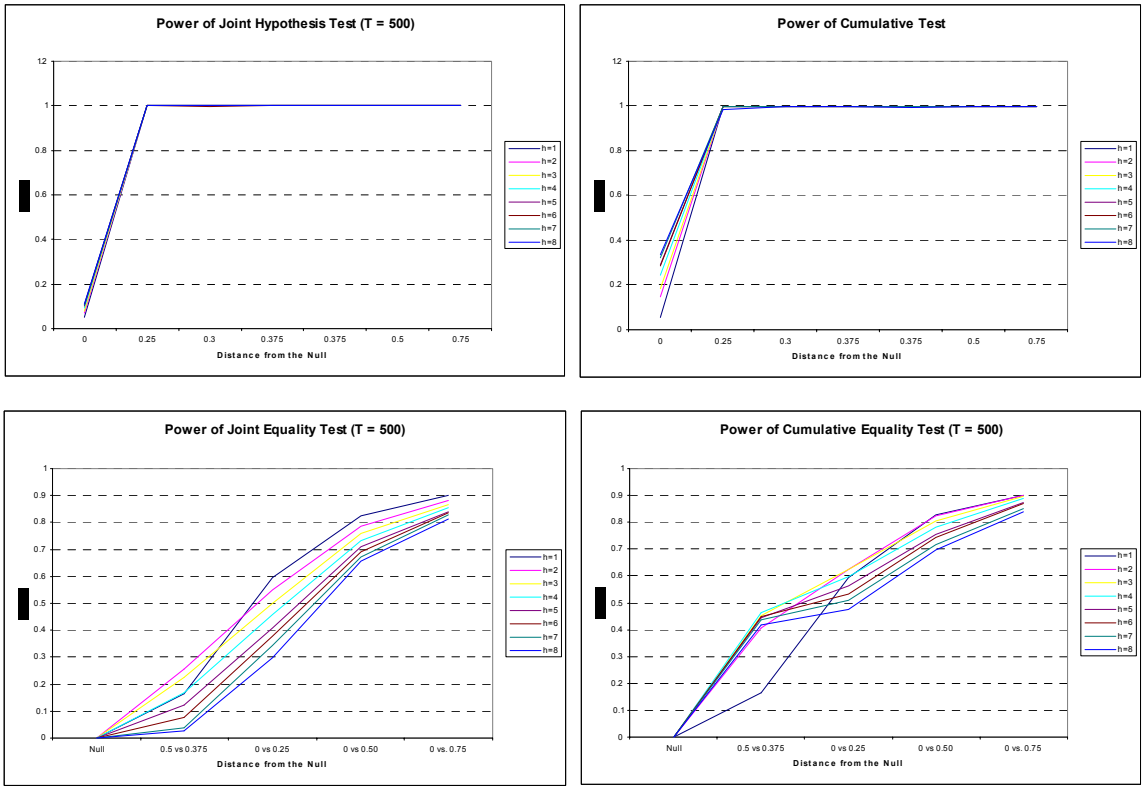

Notes: Calculations are based on drawing samples of 100 observations (after disregarding an initial 500 observations), drawn 1,000 times. The "Joint Hypothesis Tests" refers to the null that all the coefficients are zero; the "Cumulative Test" refers to the null that the sum of the coefficients of the impulse response is zero; the "Joint Equality Test" refers to the null that two impulse responses have equal coefficients; and the "Cumulative Equality Test" refers to the null that the cumulative effect of these two impulse responses is the same.

Figure 5 – Seven Variable Monetary System for the U.S. and the U.K. Impulse Responses

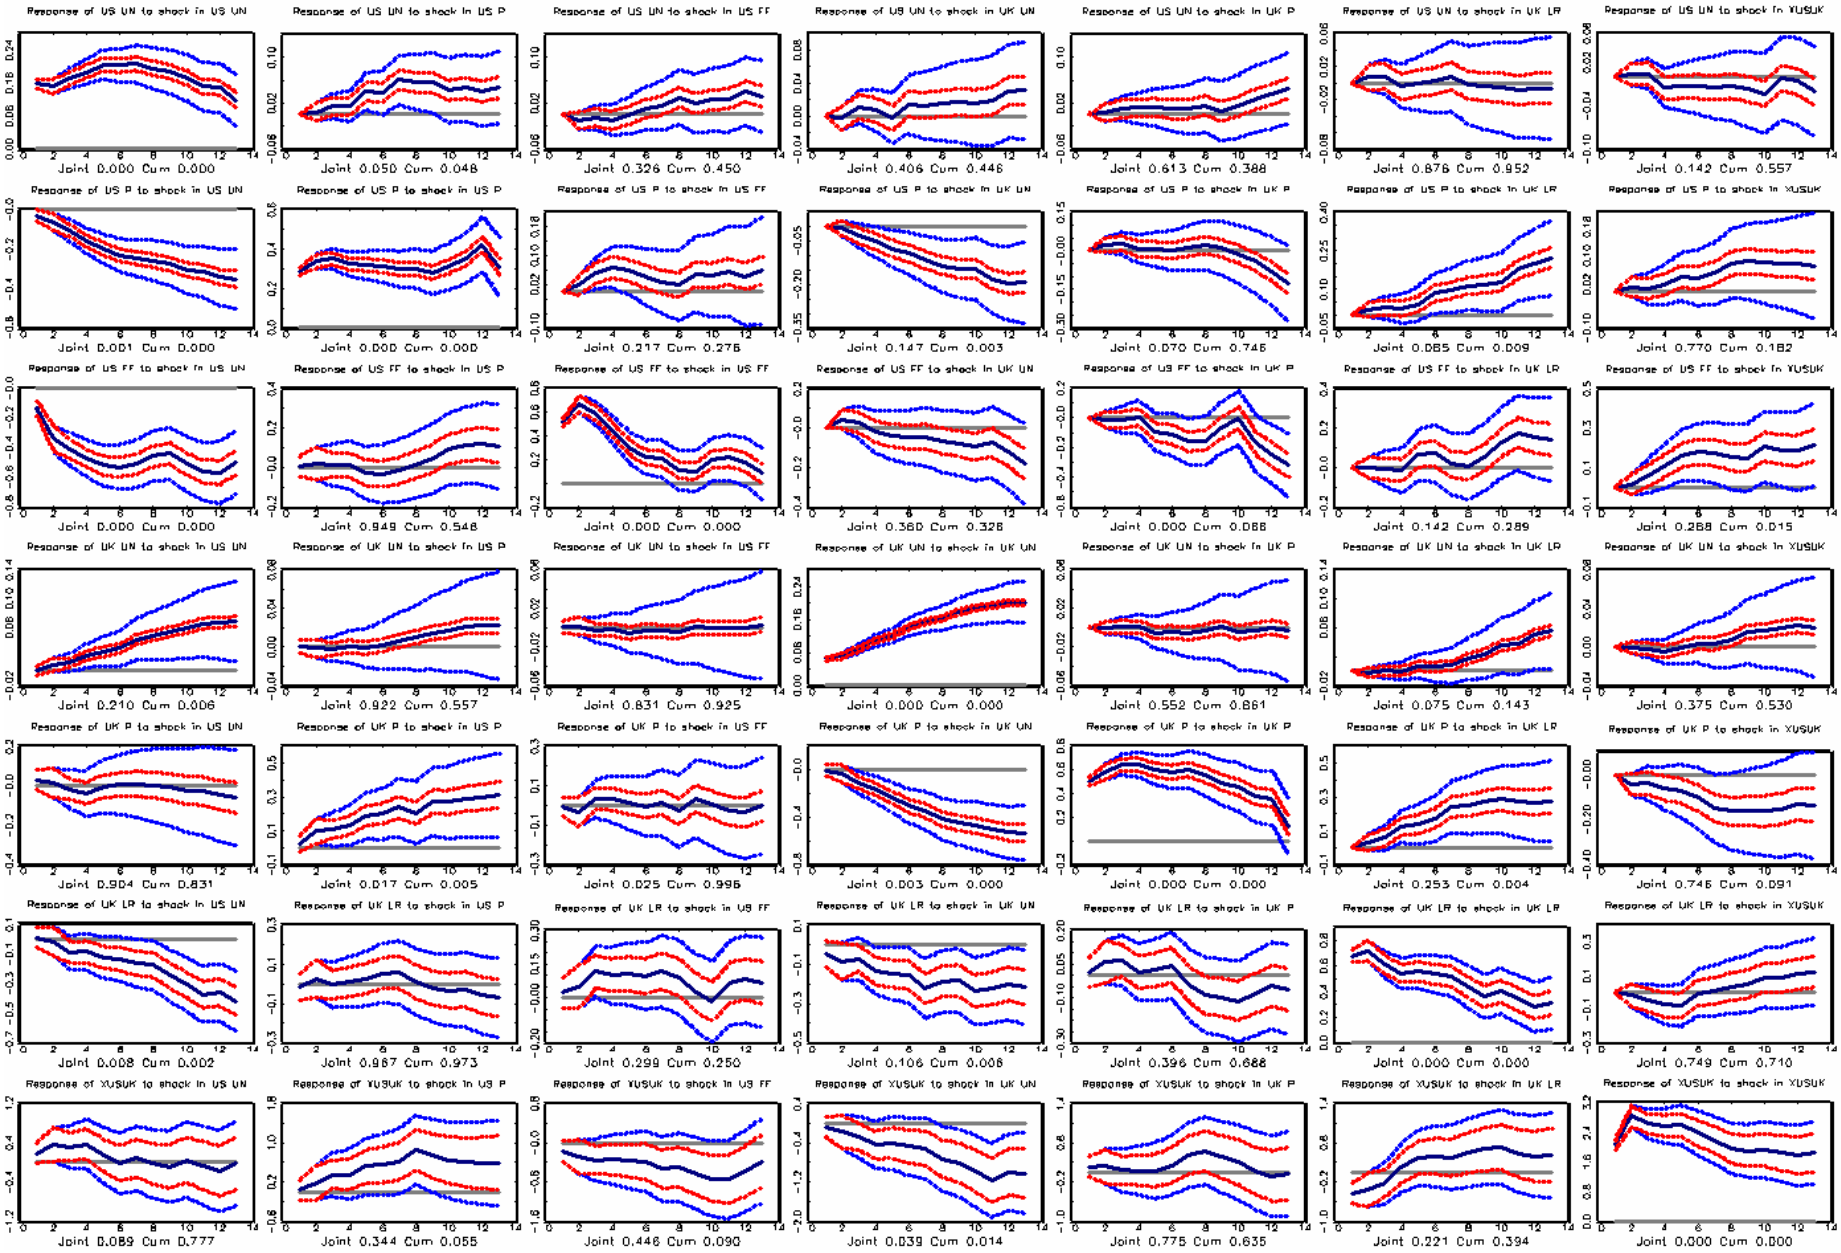

Figure 6 – Counterfactual Simulation

U.S. Counterfactual Path for the Response of Unemployment to a Price Shock

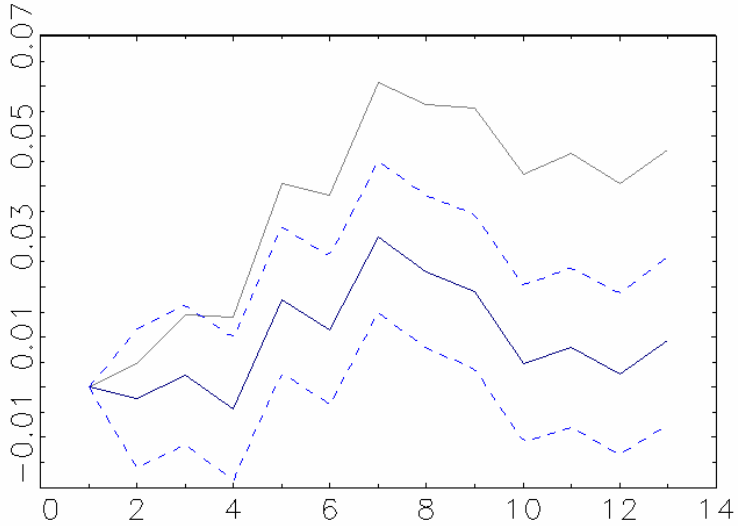

U.K. Counterfactual Path for the Response of Unemployment to a Price Shock

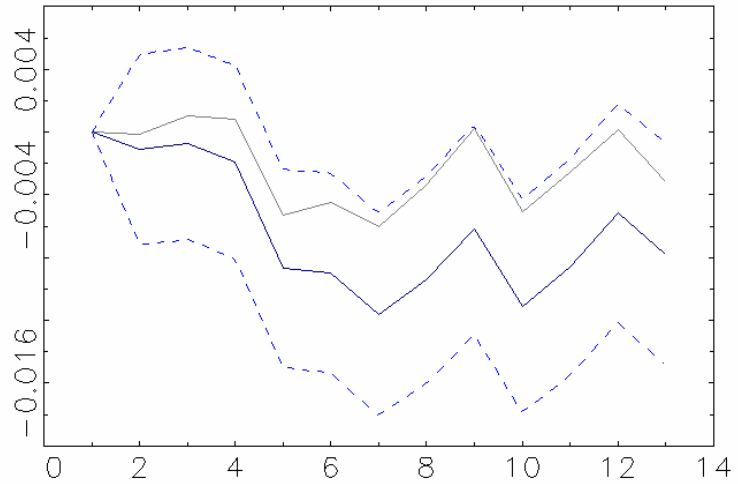

Notes: for each country, the counterfactual consists in setting the response of the policy interest rate to an inflation shock to the upper boundary of the two standard error band. Each pane; displays the original response of the unemployment rate to an inflation shock and then the counterfactual path surrounded by the time-profile bands given the counterfactual.