NBER WORKING PAPER SERIES

STOCK PRICES, NEWS AND ECONOMIC FLUCTUATIONS

Paul Beaudry
Franck Portier

Working Paper 10548
http://www.nber.org/papers/w10548

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
June 2004

The authors thank Susanto Basu , Larry Christiano, Roger Farmer, Robert Hall, Richard Rogerson, Julio Rotemberg and participants at seminars at CEPR ESSIM 2002, SED Paris 2003, Bank of Canada, Bank of England, the Federal Reserve of Philadelphia, the National Bureau of Economic Research, University of Berlin, Université du Québec à Montréal, Université de Toulouse and CREST for helpful comments.
The views expressed herein are those of the author(s) and not necessarily those of the National Bureau of Economic Research.

Stock Prices, News and Economic Fluctuations
Paul Beaudry and Franck Portier
NBER Working Paper No. 10547
June 2004
JEL No. E3

ABSTRACT

In this paper we show that the joint behavior of stock prices and TFP favors a view of business cycles driven largely by a shock that does not affect productivity in the short run -- and therefore does not look like a standard technology shock -- but affects productivity with substantial delay -- and therefore does not look like a monetary shock. One structural interpretation we suggest for this shock is that it represents news about future technological opportunities which is first captured in stock prices. We show that this shock causes a boom in consumption, investment and hours worked that precede productivity growth by a few years. Moreover, we show that this shock explains about 50\% of business cycle fluctuations.

Paul Beaudry
Department of Economics
University of British Columbia
997-1873 East Mall
Vancouver, BC
Canada V6T 1Z1
and NBER
paulbe@interchange.ubc.ca

Franck Portier
Manufacture des Tabacs
Universite des Sciences Sociales de Toulouse
21 allee de Brienne
31000 Toulouse, France
fportier@cict.fr

1 Introduction

There is a huge literature suggesting that stock prices movements reflect the market's expectation of future developments in the economy. As a test of standard valuation models, Fama [1990] shows that monthly, quarterly and annual stock returns are highly correlated with future production growth rates for 1953-1987. This result is confirmed on a extended sample (1889-1988) by Schwert [1990]. Both authors argue that the relation between current stock returns and future production growth reflects information about future cash flows that is impounded in stock prices. On the other hand, not all stock prices movements are informative, as Shiller [1981] noted that stock prices move too much to be justified by subsequent changes in dividends, such an evidence being confirmed by Flavin [1983] and Mankiw, Romer, and Shapiro [1985]. There is also a huge literature, and a long tradition in macroeconomics (from Pigou [1926] and Keynes [1936] to the survey of Benhabib and Farmer [1999]), suggesting that changes in expectation may be an important element driving economic fluctuations.

Given this, it is surprising that the empirical macro literature – especially the VAR based literature – rarely exploits stock prices movements to expand our understanding of the role of expectations in business cycle fluctuations. In this paper, we take a step in this direction by showing how stock prices movements, in conjunction with movements in total factor productivity (TFP), can be fruitfully used to help shed new light on the forces driving business cycle fluctuation.

The empirical strategy we adopt in this paper is to perform two different orthogonalization schemes as a means of identifying properties of the data that can then be used to evaluate theories of business cycles. Let us be clear that our empirical strategy is a purely descriptive device which becomes of interest only when its implications are compared with those of structural models. The two orthogonalization schemes we use are based on imposing sequentially, not simultaneously, either impact or long run restrictions on the orthogonalized moving average representation of the data. The primary system of variables that interests us is one composed of an index of stock market value (SP) and measured total factor productivity (TFP). Our interest in focusing on stock market

information is motivated by the view that stock prices are likely a good variable for capturing any changes in agents expectations about future economic conditions.

The two disturbances we isolate with our procedure are first, a disturbance which represents innovations in stock prices which are orthogonal to innovations in TFP and second, a disturbance that drives long run movements in TFP. The main intriguing observation we uncover is that these two disturbances– when isolate separately without imposing orthogonality – are found to be almost perfectly co-linear and to induce the same dynamics. We also show that these co-linear shock series causes standard business cycle co-movements (i.e., induces positive co-movement between consumption and investment) and explains a large fraction of business cycle fluctuations. Moreover, when we use measures of TFP which control for variable rates of factor utilization, as for example when we use the series constructed by Basu, Fernald, and Kimball [2002], we find that our shock series anticipate TFP growth by several years.

In order to interpret the result from our empirical exercise, we begin by presenting a simple model where fluctuations are driven by surprise changes in productivity as well as a temporary disturbance– which in our example is a monetary shock. This example allows us to clarify the extent to which the data on TFP and stock prices have properties that run counter to those implied by models where surprise changes in productive capacity are a central part of fluctuations. We also present a model where technological innovations only affect productive capacity with delay, and show how such a model can explain quite easily the patterns observed in the data. In particular, our evidence suggests that business cycles may be driven to a large extent by TFP growth that is heavily anticipated by economic agents; thereby leading to what might be called expectation driven booms. In effect, the original burst in economic activity associated with the shock we identify, using either the impact or the long run restriction, looks like a business cycle fluctuations which preempts future growth in productivity. Hence, our empirical results suggests that an important faction of business cycles fluctuations may be driven by changes in expectations – as is often suggested in the macro literature – but where these changes in expectations may well be based on fundamentals since they anticipate future changes in productivity.

The remaining sections of the paper are structured as follows. In Section 2, we present our empirical strategy and show how it can be used to shed light of the sources of economic fluctuation. In Section 3, we present the data and in Section 4 we implement our strategy using post-war US data. We present our empirical results in steps from a smaller dimensional system – composed only of TFP and stock prices – to a larger system that includes alternatively or jointly consumption, investment and hours. We begin by considering the bi-variate system for TFP and stock prices since it offers the most straightforward way of highlighting an intriguing property of the data. In a second stage, we consider a tri-variate system composed of TFP, stock prices and consumption. The advantage of the tri-variate system is that it allows us to easily embed a standard view about the sources of fluctuations. We also report results based on a set a four-variable systems in order to further document the robustness of our results. In Section 5, we discuss the strength and weaknesses of different models in explaining he observations presented in Section 4. Finally, Section 6 offers some concluding comments.

2 Using Impact and Long-Run Restrictions Sequentially to Learn About Macroeconomic Fluctuations

The object of this section is to present a new means of using orthogonalization techniques –i.e. impact and long run restrictions – to learn about the nature of business cycle fluctuations. Our idea is not to use these techniques simultaneously (as is now common in the literature), but is instead to use them sequentially. In particular, we will want to apply this sequencing to describe the joint behavior of stock prices (SP) and measured total factor productivity (TFP_t) in a manner that can be easily mapped into structural models. The main characteristic of stock prices that we want to exploit is that it be an unhindered jump variable, that is, a variable that can immediately react to changes in information without lag.

2.1 Two Orthogonalization Schemes

Let us begin our discussion from a situation where we already have an estimate of the reduced form moving average (Wold) representation for the bivariate system $\{TFP_t, SP_t\}$, as given below (for ease of presentation we neglect any drift terms).

$$\begin{pmatrix} \Delta TFP_t \\ \Delta SP_t \end{pmatrix} = C(L) \begin{pmatrix} \mu_{1,t} \\ \mu_{2,t} \end{pmatrix}$$

where L is the lag operator, $C(L) = I + \sum_{i=1}^{\infty} C_i L^i$, and where the variance co-variance matrix of μ is given by Ω. Furthermore, we will assume that the system has at least one stochastic trend and therefore $C(1)$ is not equal to zero. In effect, most of our analysis will be based on a moving average representation derived from estimation a vector error correction model (VECM) for TFP and stock prices.

Now consider deriving from this Wold representation alternative representations with orthogonalized errors. As is well know, there are many ways of deriving such representations. We want to consider two of these possibilities, one that imposes an impact restriction on the representation and one that imposes a long run restriction. In order to see this most clearly, let us denote these two alternative representations by:

$$\begin{pmatrix} \Delta TFP_t \\ \Delta SP_t \end{pmatrix} = \Gamma(L) \begin{pmatrix} \epsilon_{1,t} \\ \epsilon_{2,t} \end{pmatrix}, \qquad (1)$$

$$\begin{pmatrix} \Delta TFP_t \\ \Delta SP_t \end{pmatrix} = \widetilde{\Gamma}(L) \begin{pmatrix} \widetilde{\epsilon}_{1,t} \\ \widetilde{\epsilon}_{2,t} \end{pmatrix}, \qquad (2)$$

where $\Gamma(L) = \sum_{i=0}^{\infty} \Gamma_i L^i$, $\widetilde{\Gamma}(L) = \sum_{i=o}^{\infty} \widetilde{\Gamma}_i L^i$ and the variance covariance matrices of ϵ and $\widetilde{\epsilon}$ are identity matrices. In order to get such a representation, say in the case of (1), we need to find the Γ matrices that solve the following system of equations:

$$\begin{cases} \Gamma_0 \Gamma_0' &= \Omega \\ \Gamma_i &= C_i \Gamma_0 \quad \text{for} \quad i > 0 \end{cases}$$

However, since the above system has one more variable than equations, it is necessary to add a restriction to pin down a particular solution. In case (1), we will pin down a solution by imposing that the $1,2$ element of Γ_0 be equal to zero, that is, we choose an orthogonalization where the second disturbance ϵ_2 has no contemporaneous impact on TFP. In case (2), we impose that the $1,2$ element of the long run matrix $\widetilde{\Gamma}(1) = \sum_{i=0}^{\infty} \widetilde{\Gamma}_i$ equals zero, that is, we choose an orthogonalization where the disturbance $\widetilde{\epsilon}_2$ has no long run impact on TFP (the use of this type of orthogonalization was first proposed by Blanchard and Quah [1989]). Our idea now is to use these two different ways of organizing the data to help evaluate different classes of economic models and indicate directions for model reformulation. For example, a particular theory may imply that the correlation between the shocks ϵ_2 and $\widetilde{\epsilon}_1$ be close to zero and that their associated impulses responses be different. Therefore, we can evaluate the relevance of such a theory by examining the validity of its implications along such a dimension.

In order to clarify the potential usefulness of such a procedure, we will begin by presenting a simple canonical model of fluctuations driven by money shocks and surprise technology shocks. The model we chose for this illustration is in the New-Keynesian tradition in order to easily incorporate real effects of money. However, the point we want to make does not depend on the presence of nominal rigidities, as will become clear. In effect, our goal with this example is to highlight how the co-variance properties of the derived shocks ϵ_2 and $\widetilde{\epsilon}_1$ can be used to evaluate a theoretical model.

2.2 Two Simple Models

Here we illustrate the implications of sequentially using impact and long-run restrictions in a canonical New-Keynesian model driven by monetary shocks and surprise changes in technology. Later, we will present an example where technological improvements only diffuse slowly across the economy but where agents recognize the potential impact of an innovation well before it has improved productivity. We will show that these two models deliver different predictions with respect to the correlation between ϵ and $\widetilde{\epsilon}$. As we want to derive simple and explicit results, the models we present here do not aim at realism as many assumptions are made in order to allow analytical

solutions.

A New-Keynesian type model : Let us consider an economy with monetary shocks, pre-set wages and technological disturbances. Money is introduced through a cash-in-advance constraint and preferences of the representative household j are given by

$$U = E_0 \sum_{t=0}^{\infty} \beta^t \left[\log C_t^j - \Lambda \frac{(L_t^j)^\sigma}{\sigma} \right] \tag{1}$$

There is no capital in the model and only one final good y. The final good is produced by a continuum of intermediate goods z_i, and each intermediate good is produced by a composite of labor from different households as follows

$$y = \left(\int_o^1 z_i^{\rho_1} di \right)^{\frac{1}{\rho_1}}, \qquad 0 < \rho_1 < 1 \tag{2}$$

$$z_i = \theta_t \left(\int_o^1 l_j^{\rho_2} dj \right)^{\frac{1}{\rho_2}}, \qquad 0 < \rho_2 < 1 \tag{3}$$

The technology parameter θ_t is assumed to follow a random walk (in logs) with innovations $\eta_{1,t}$. Both the labor market and the intermediate goods market are assumed to be monopolistically competitive. In the labor market, households set their wages ahead of the realizations of money and technology disturbances. The log of money supply (m_t) follows a random walk with innovation $\eta_{2,t}$, with $\eta_{2,t}$ being uncorrelated with $\eta_{1,t}$. The intermediate goods market is also monopolistically competitive, but prices are set after the realization of $\eta_{1,t}$ and $\eta_{2,t}$. Hence, this is a model with flexible prices and pre-set wages. The profits of the intermediate good firms are returned to households, all of which hold the market portfolio. The value of firms (the stock market value) is the discounted sum of profits, where the discount rate is given by the intertemporal marginal rate of substitution between consumption in different periods.

The representative household decides each period how much to consume and how much save in terms of money balances. It also decides on the nominal wage at which it will supply labor

next period. At the beginning of period t, a household's money holdings carried from the previous period are multiplied by the monetary shock.

In this model, as shown in the appendix, prices will be a markup on marginal cost ($\frac{w_t}{\theta_t}$), and nominal wages will be directly proportional to the expected supply of money. In equilibrium, output and firm profits will be affect by unexpected money and the level of technology. Hence this model delivers the following simple structural moving average representation for $TFP = \log(\theta_t)$ and log stock market value (SP_t) (where we have again omitted constants)

$$\begin{pmatrix} \Delta TFP_t \\ \Delta SP_t \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ 1 & (1-L) \end{pmatrix} \begin{pmatrix} \eta_{1,t} \\ \eta_{2,t} \end{pmatrix} \tag{4}$$

Since the structural moving average representation of this system satisfies our short-run and long-run orthogonalization restrictions, we can immediate see that this model implies:

$$\epsilon_1 = \eta_1 \quad , \quad \epsilon_2 = \eta_2 \quad , \quad \widetilde{\epsilon}_1 = \eta_1 \quad , \quad \widetilde{\epsilon}_2 = \eta_2 \tag{5}$$

In particular, this type of model implies that $\epsilon_2 \perp \widetilde{\epsilon}_1$ since $\eta_2 \perp \eta_1$.

It is straightforward to understand that in this economy, the shock that has permanent effect on TFP, $\widetilde{\epsilon}_1$, is also the one that affects TFP in the short run, while the money shock does not affect TFP in the short run nor in the long run. Therefore, if such a model is the data generating process, the shock ϵ_2 recuperated using our impact restriction should be found to be orthogonal to the shock $\widetilde{\epsilon}_1$ recuperated using our long run restriction. [1]

A model with delayed response of innovation on productivity Let us now consider an alternative setting where stock prices continue to be a discounted sum of future profits but where technological innovations no longer immediately increases productivity but instead only increase productive capacity over time. The objective of this example is to emphasize what such an environment predicts regarding the correlation between ϵ_2 and $\tilde{\epsilon}_1$ derived using sequentially impact and long run restrictions. To this end, let us assume that measured TFP, denoted θ, is

[1] A similar orthogonality result can be derived for an RBC type model with temporary preference changes and permanent but unexpected changes in technological opportunities.

composed of two components: a non-stationary component D_t and a stationary component ν_t. The component ν_t can be thought of as either a measurement error or as a temporary technology shock. For the discussion, we will treat ν_t as a temporary shock to θ, although the measurement error interpretation has the same implications. In contrast, the component D_t is the permanent component of technology, and is assumed to follow the process given below:

$$\begin{cases} \theta_t &= D_t + \nu_t \\ D_t &= \sum_{i=0}^{\infty} d_i \eta_{1,t-i} \\ d_i &= 1 - \delta^i, \quad 0 \geq \delta < 1 \\ \nu_t &= \rho \nu_{t-1} + \eta_{2,t}, \quad 0 \leq \rho < 1 \end{cases} \tag{6}$$

We will call the process for D_t a diffusion process since an innovation η_1 is restricted to have no immediate impact on productive capacity ($d_0 = 0$), the effect of the technological innovation on productivity is assumed to grow over time ($d_i \leq d_{i+1}$) and the long run effect is normalized to 1. In contrast to the common random walk assumption for the permanent component of TFP, such a process allows for an S-shaped response of TFP to a technological innovation; which is consistent with many micro-based studies of the effects of technological innovation on productivity (Pakes [1985] mentions the "long and erratic lag structure between invention and the current benefits derived from it") We now want to derive the implied structural moving average for ΔTFP and ΔSP. To that end, consider a simple Lucas' tree type of model, where the ownership of the unique tree of the economy is tradable and where it pays dividend θ_t.

Households consume and trade firms shares. Preferences are represented by:

$$U = E_0 \sum_{t=0}^{\infty} \beta^t \frac{C_t^{1-\sigma}}{1-\sigma} \tag{7}$$

with $\sigma \geq 0$. The household can buy or sell shares S_t at unit price P_t. As there is a unique tree in the economy, the stock market value is $SP_t = P_t$. The household budget constraint is given by

$$C_t + P_t S_{t+1} \leq (P_t + \theta_t) S_t \tag{8}$$

Optimal behavior of the household is given by budget constraint (8), Euler equation (9)

$$P_t = \beta E_t \left[\left(\frac{C_t}{C_{t+1}} \right)^{\sigma} (P_{t+1} + \theta_{t+1}) \right] \tag{9}$$

and the transversality condition $\lim_{j\to\infty} E_t \beta^j P_{t+j} S_{t+j+1} = 0$.

At the competitive equilibrium, $S_t = 1$ and $C_t = \theta_t$ for all t. The stock market value SP_t is then given by

$$SP_t = \beta E_t \left[\left(\frac{\theta_t}{\theta_{t+1}} \right)^{\sigma} (SP_{t+1} + \theta_{t+1}) \right] \tag{10}$$

In order to obtain simple analytical results, we make the further assumption that households are risk neutral, so that $\sigma = 0$. In that case, equation (10) collapses to

$$SP_t = \beta E_t \left[(SP_{t+1} + (1-\alpha)\theta_{t+1}) \right] \tag{11}$$

Solving forward and using the transversality condition $\lim_{j\to\infty} \beta^j E_t[SP_{t+j}] = 0$, we obtain

$$SP_t = \beta E_t \left[\sum_{j=0}^{\infty} \beta^j \theta_{t+j+1} \right] \tag{12}$$

Using the process of θ_t given in (6), one can obtain the following structural moving average representation for TFP and stock prices first differences:

$$\begin{pmatrix} \Delta TFP_t \\ \Delta SP_t \end{pmatrix} = \begin{pmatrix} (1-\delta) \sum_{i=1}^{\infty} \delta^{i-1} L^i & \frac{(1-L)}{(1-\rho L)} \\ \frac{\beta(1-\delta)}{1-\beta\delta} \sum_{i=0}^{\infty} \delta^i L^i & \frac{\beta\rho}{1-\beta\rho} \quad \frac{1-L}{1-\rho L} \end{pmatrix} \begin{pmatrix} \eta_{1,t} \\ \eta_{2,t} \end{pmatrix} \tag{13}$$

From the above representation, we see that the impact matrix on levels of TFP and SP is of the form:

$$\begin{pmatrix} 0 & 1 \\ \frac{\beta(1-\delta)}{1-\beta\delta} & \frac{\beta\rho}{1-\beta\rho} \end{pmatrix} \tag{14}$$

And the long run matrix for the levels of TFP and SP is of the form

$$\begin{pmatrix} 1 & 0 \\ \frac{\beta}{1-\beta\delta} & 0 \end{pmatrix} \tag{15}$$

Hence, performing our short-run and long-run identification on this system, the relationship between the identified errors ϵ_t, $\tilde{\epsilon}_t$ and the structural errors η_t are:

$$\epsilon_1 = \eta_2 \quad , \quad \epsilon_2 = \eta_1 \quad , \quad \widetilde{\epsilon}_1 = \eta_1 \quad , \quad \widetilde{\epsilon}_2 = \eta_2 \tag{16}$$

In particular, we have that ϵ_2 is co-linear to $\widetilde{\epsilon}_1$ in this case.

Discussion : The important aspect of the first model is that the derived ϵ_2 shock, which under this theory should correspond to the money shock, is predicted to be orthogonal to $\widetilde{\epsilon}_1$, which should be the surprise increase in productivity. Moreover, in such a case, the implied impulse responses associated with these shocks would be expected to be different. Therefore, looking whether this type of pattern is found in the data provides a means of evaluating the relevance of such a class of models, that is, models where surprise technological disturbances are a potentially important source of fluctuations.

The diffusion model is different. It is an example of a model where, even before technological opportunities have actually expanded an economy's production possibility set, forward looking variables – such as stock prices– are incorporating this possibility. [2] If this class of models is relevant, the long run restriction used to derive the orthogonal moving average representation given by $\widetilde{\Gamma}_i$ and $\widetilde{\epsilon}$ still implies that $\widetilde{\epsilon}_1$ can be interpreted as a technological shock, but now it implies that this shock have zero effect on productivity on impact, that is, if productivity changes are anticipated then by definition of an anticipated shock, the actual shock has zero effect on impact on TFP. Hence, under this type of models ϵ_2 and $\widetilde{\epsilon}_1$ are predicted to be co-linear (as opposed to being orthogonal in the previous model) as they both should capture the effect of anticipated changes in technological opportunities. Moreover, the impulse responses associated with ϵ_2 and $\widetilde{\epsilon}_1$ should be identical

[2] There are at least two types of models that would exhibit this property. First, there are models where agents are informed about future technological opportunities before such opportunities fully implementable, as the one we have presented. Second, there are implementation cycle models (see Shleifer [1986]) where a sunspot coordinates the economy decision to start implementing a new technology.

In summary, in this section we suggested a way of presenting information on stock prices and TFP which maps easily to structural models The approach requires deriving two orthogonal moving average representations of the data and comparing the resulting error series and their associated impulse responses. In particular, the approach suggests examining the correlation between ϵ_2 and $\widetilde{\epsilon}_1$ as a means of evaluating whether models which emphasize surprise changes in technology as an explanation of the data, or if instead the data favors a characterization where technological innovation are first reflected in forward looking variables like stock prices and only later reflected in changes in productivity. In the remaining sections of the paper, we will often refer to this later type of shock as a news shock since it is a shock that brings news about future growth in productivity.

Before proceeding to our empirical exploration, let us immediately emphasize that the prediction of a high correlation between ϵ_2 and $\widetilde{\epsilon}_1$ can be derived from other setting than one which assumes an exogenous and slow diffusion process for productivity. For example, in the case where technological change in embodied in capital and where such embodiment is mis-measured in the capital stock, then innovations in stock prices should precede growth in TFP; therefore predicting a high correlation between ϵ_2 and $\widetilde{\epsilon}_1$. The reader may want to keep this alternative data generating process in mind while we present the empirical results, as we will discuss the strength and weakness of this interpretation, as well as others, in Section 5.

3 Data and Specification Issues

Our empirical investigation will use US data over the period 1948Q1 to 2000Q4 (the data was collected in August 2002). The two series that interest us for our bi-variate analysis are an index of stock market value (SP) and a measure of total factor productivity. Later, we will consider larger systems that also include consumption and investment and therefore we also present the source of these data here.

The stock market index we use is the quarterly Standards & Poors 500 Composite Stock Prices Index, deflated by the seasonally adjusted implicit prices deflator of GDP in the non farm private

business sector [3] and transformed in per-capita terms by dividing it by the population aged 15 to 64. As the population series is annual, it has been interpolated assuming constant growth within the quarters of the same year. We denote the log of this index by SP

The construction of our baseline TFP series is relatively standard. We restrict our attention to the non farm private business sector. From the U.S. Bureau of Labor Statistics, we retrieved two annual series: labor share (s_h) and capital services (KS) which measures the services derived from the stock of physical assets and software. The average value of the labor share is $\overline{s}_h = 67.66\%$. The capital services series has been interpolated to obtain a quarterly series, assuming constant growth within the quarters of the same year. Output (Y) and hours (H) are quarterly seasonally adjusted non farm business measures, from 1947Q1 to 2000Q4 (also from U.S. Bureau of Labor Statistics). We then construct a measure of (log) TFP as

$$TFP_t = \log\left(\frac{Y_t}{H_t^{\overline{s}_h} KS_t^{1-\overline{s}_h}}\right)$$

The consumption measure (C) we use is the per capita value of real personal consumption of non durable goods and services (obtained from Bureau of Economic Analysis), while investment (I) is the per capita value of the sum of real personal consumption of durable goods and real fixed private domestic investment (also obtained from the Bureau of Economic Analysis).

The resulting four series for SP, TFP, C, I are plotted of figure 1.

Specification: From our data on TFP and SP, we first want to recover the Wold moving average representation for ΔTFP and ΔSP. Since from unit root tests (not reported here) and cointegration tests, we found that SP and TFP are likely cointegrated I(1) processes, a natural means of recovering the Wold representation is by inverting a VECM. However, in a VECM framework, one must be careful in properly identifying the matrix of co-integration relationship in order to avoid mispecification. In effect, as emphasized in Hamilton [1994], if one is worried of potential mispecifi-

[3] Our results are robust to using the producer price index for finished goods as a deflator instead of the GDP deflator

cation it may be best to estimate the VECM allowing for the matrix of co-integrating relationship to be of full rank– which corresponds to estimating the system in level. Then one can estimate the VECM with a matrix of cointegration relationships which is of reduced rank and examine whether the resulting Wold representation is similar to that found by estimating the system in levels. In the following we will adhere to this principal by reporting [4] results based both on a Wold representation achieved by inverting a VECM, as well as results based on estimating the system in level. Since we want to avoid mispecification bias due to an omitted co-integration relationship, our approach to testing for co-integrating relationship is conservative in the sense of testing from more (H0) co-integrating relationship to less (H1). To this end, we used the test proposed by Nyblom and Harvey [2000] to test for cointegration. This procedure indicates that co-integration between SP and TFP could not be rejected at the 5% level and therefore we adopted the VECM specification as our benchmark specification. [5]

The second specification choice is related with the number of lags to include in the VECM. Again, our strategy is not to impose much to the data. According to likelihood ratio test two or five lags appear preferable – when testing in a descendant way for the optimal number of lags from two years up to one quarter. When tested one against the other, five is preferred to two. We

[4] Our approach finds support in chapter 20, section 4 of Hamilton [1994] entitled "Overview of Unit Roots – To Difference or Not to Difference". According to Hamilton, estimating in levels has the advantage that the parameters that describe the system's dynamics are estimated consistently. Differentiating should improve the small-sample performance of all of the estimates if the true process is in difference. But the drawback is that the true process may not be a VAR in difference. Working with a VECM specification, that imposes some cointegration relations between the variables, might make one falls in the same trap than with the VAR in difference. The restrictions imposed may be invalid, and alternative tests for unit root and cointegration can produce conflicting results. From this informal discussion, Hamilton concludes that the "eclectic" strategy would begin by estimating the VAR in levels without restrictions. The next step is then to make an assessment as to which series are likely nonstationary, so that a VECM of a VAR in difference could then be estimated. According to Hamilton, *"If the VAR for the data in levels form yields similar inferences to those for the VAR in stationary form, the researcher might be satisfied about the assumptions made about unit roots"*, which is essentially the approach we take.

[5] In order to perform the Nyblom and Harvey test, it is necessary to make a choice of lags length to correct for serial correlation. Our preference was for a lag length of 12 quarters and at this lag length co-integration was not rejected at the 5% level. However, the result of non-rejection of co-integration is robust to varying this lag length anywhere from 6 to 18 quarters. Let us note that if we proceeded in the inverse fashion of adopting non-cointegration as our null, we could not reject it either. Our preference for the VECM representation instead of the VAR representation in difference is that the former mimics the unconstrained level representation very closely, while the later does not which suggests mispecification.

therefore choose to work with five lags since this seemed to us large enough not to put to much restrictions on the data. We will nevertheless show the robustness of our results with a two lags specification.

4 Results in Bi-Variate System

4.1 Preliminary Results

We began by estimating a VECM for (TFP, SP) with one cointegrating relation and recover two orthogonalized shock series corresponding to the ϵ and $\widetilde{\epsilon}$ discussed in Section 2, that is, ϵ was recovered by imposing an impact restriction (a restriction on Γ_0) and $\widetilde{\epsilon}$ was recovered by imposing a long run restriction. The level impulse responses on (TFP, SP) associated with the ϵ_2 shock and the $\widetilde{\epsilon}_1$ shock are displayed on Figure 2. A first striking observations is that those responses appear very similar when comparing one orthogonalization to another. More specifically, the dynamics associated with the ϵ_1 shock–which by construction is an innovation in stock prices which contemporaneously orthogonal to TFP– seems to permanently affect TFP, while the dynamics associated with the $\widetilde{\epsilon}_1$ shock –which by construction has a permanent effect on TFP– has essentially no impact effect on TFP (the point estimate indicates a slight negative effect) but has a substantial effect on SP. On the one hand, these results suggest that ϵ_2 contains information about future TFP growth which is instantaneously and positively reflected in stock prices[6]. While on the other hand, they suggest that permanent changes in TFP are first reflected in stock prices before they actually increase productive capacity.

The similarity between the effects of these two shocks is further confirmed by the inspection of the forecast error variance decomposition plot (Figure 3). Observe that the $\widetilde{\epsilon}_1$ shock explains virtually none of the short run movements of TFP, but does explain the variability of stock prices. On the other hand, the ϵ_2 shock also explains most of the long variance of TFP. This result derives

[6]The observation in Figure 2, whereby TFP increases following an innovation in SP, indicates that stock prices Granger causes TFP. In effect, we also directly performed the test of whether SP Granger causes TFP in this system, and we found that such causality could not be rejected at the 1%.

from the quasi-identity between the ϵ_2 shock and the $\widetilde{\epsilon}_1$ shock, as shown in Figure 4 which simply plots $\epsilon_{2,t}$ against $\widetilde{\epsilon}_{1,t}$. In effect, the correlation coefficient between these two series is 0.97 (with a standard deviation of .006), that is, these two orthogonalization techniques recover essentially the same shock series. The interesting question then becomes, what kind of structural macroeconomic model is consistent with these two orthogonalization techniques generating the same shock series? As we have highlighted in Section 2.2, this observation runs counter to simple models where technological improvements are modelled as surprises since these models generally imply that ϵ_2 and $\widetilde{\epsilon}_1$ should be orthogonal. In contrast, this pattern appears consistent with the view –which we call the news view– that improvements in productivity are generally anticipated by market participants due to a lag between the recognition of a technological innovation and its eventual impact on productivity. However, before exploring this news interpretation further, we want to illustrate the robustness of the observation that ϵ_2 and $\widetilde{\epsilon}_1$ are strongly correlated and induce similar dynamics.

In Figures 5, 6 and 7, we report analogues to Figures 2 and 4 for the case where we obtain the Wold moving average representation (1) by estimating our system in levels instead of in a VECM form, and (2) by estimating our VECM with only 2 lags instead of 5. Figure 5 superimposes the impulse responses associated with ϵ_2 and $\widetilde{\epsilon}_1$ for the case where the Wold representation was obtained by estimating the system (TFP, SP) in levels versus in VECM form. In the case of ϵ_2 (top panels), the resulting impulse responses are very similar, except that in levels there is slightly more mean reversion in the index of stock prices. In the case of $\widetilde{\epsilon}_1$, the instantaneous response of TFP when derived from a level estimation is now positive (lower left panel of 5), but it is still very small and not significantly different from zero. [7] In the left panel of Figure 7, we plot ϵ_2 against $\widetilde{\epsilon}_1$ for the case where they are obtained from estimating the system in levels. As can be seen from this Figure, these two shocks series are again very highly correlated (.998 with s.d. .005) which indicates the robustness of this observation to estimating the system in VECM form or in level form. In Figure 6 we superimpose the impulse responses associated with ϵ_2 and $\widetilde{\epsilon}_1$ for the case where the Wold representation was obtained by estimating the system (TFP, SP) in VECM form

[7] The confidence bands of the level estimation responses are not reported on the figure for clarity reasons.

with only two lags versus with five lags. Furthermore, in the right side panel of Figure 7 we plot ϵ_2 against $\widetilde{\epsilon}_1$ for the case where they are obtained from estimating the system in VECM form with 2 lags. As is clear from these figures, the number of lags in the VECM form does not greatly affect the patterns we are highlighting in the data, as the correlation is again close to one (.9997 with s.d. .0017).[8] In particular, the high correlation between ϵ_2 and $\widetilde{\epsilon}_1$ observed in Figure 7 appears robust and therefore remains counter to the predictions of simple models with surprise changes in productive capacity. In contrast, this pattern suggest that agents may be able to predict future TFP growth.

4.2 Exploring the News Interpretation Further

The observation that our estimates of ϵ_2 and $\widetilde{\epsilon}_1$ are highly correlated and induce similar impulse responses suggests that news about future technological developments may be a relevant driving force behind business cycle fluctuations. In light of this possibility, we now want to go a step further and ask: How does the economy respond to such a shock, that is, do the responses to ϵ_2 (or $\widetilde{\epsilon}_1$) look like standard business cycle fluctuations in the sense of generating positive co-movements in consumption and investment? To answer this questions, we will begin by exploiting the co-movements between different variables (i.e. consumption and investment) and the ϵ shock series derived from our baseline specification for (TFP, SP). Later we will look directly at larger systems which incorporate these other variables explicitly. We can focus here on the effect of only the ϵ shocks since, as we have shown, they are essentially mirror images of the $\widetilde{\epsilon}$ shocks. Our first approach to this issue will therefore be to estimate the following truncated moving average representation for different variables Z_t:

$$\Delta Z_t = \sum_{j=0}^{J} \phi_j^1 \epsilon_{1,t-j} + \sum_{j=0}^{J} \phi_j^2 \epsilon_{2,t-j} + \mu_t \tag{17}$$

where Z will either be consumption (C) or investment (I) and where μ a variable-specific disturbance

[8] Note that the variance decompositions are also very robust to choice of lag length or to estimating system in levels or in VECM form.

that is orthogonal to ϵ_2 and to ϵ_1. The resulting sequence given by $\sum_{j=0}^{n} \phi_j^2$ provides an estimate of the impulse response function of X to a ϵ_2 shock, that is, the response to what we claim may be a news shocks. The truncation is done for $J = 40$.

Figure 8 displays the responses of consumption and investment to ϵ_2, that is, the responses to what we suggest may reflect news of a technological innovation which only diffuses slowly into productive system. As can be seen in the Figure, a positive ϵ_2 has an expansionary impact: investment and consumption increase on impact, and seem to reach a permanently higher level after 10 to 12 quarters. These results suggest that a ϵ_2 (1) creates business cycle like fluctuations, (2) does not affect TFP contemporaneously and (3) affect TFP in the long run. This pattern is consistent with the interpretation of ϵ_2 as being primarily a news shock. Such a structural interpretation is supported by the fact that the same responses for the economy are obtained from a short run identification in which we identify a news shock as ϵ_2 in our (TFP, SP) system as the innovation to stock prices that is orthogonal to current TFP, or if we examine the effects of $\widetilde{\epsilon}_1$ which by definition affects long run TFP.

Let us emphasize that, if we interpret the current results as reflecting a diffusion process from innovation to productivity, it suggest that diffusion is rather fast. In effect, in Figure 2 we observed that measured TFP starts growing quickly after the initial increase in stock prices; with the peak obtained after approximately 4 quarters. However, one potential problem with this observation is that our measure of TFP may be an improper measure of technological opportunities since it does not take account of potential changes in rates of factor utilization. Therefore, it may be the case that in response to a technological innovation, properly measured TFP does not increase for a substantial period of time, but that mis-measured TFP responds rapidly due to changes in factor utilization. Hence, in the next subsection, we explore the robustness of our observations with respect to alternative measures of TFP.

4.3 Controlling for Variable Rates of Factor Utilization

There is a vast literature regarding how best to calculate TFP in order to obtain a good reflection of changes in production opportunities. In particular, the literature on this issue emphasizes several potential problems with the type of measure of TFP we used in the previous section. For example, our previous measure may be inappropriate due to our lack of correction for variable rates of capital utilization, labor hoarding or composition bias. One attempt to control for most of these biases can be found it the TFP series produced by Basu, Fernald, and Kimball [2002] (hereafter BFK). This series has the advantage of being constructed from dis-aggregated data which controls for variable rates of factor utilization. For this reason it appears as a good alternative series to examine the robustness of our previous results. However, it also has some drawbacks. First, it is an annual series instead of a quarterly series. Second, it covers only the period 1948-89. Notwithstanding these drawbacks, we will begin this section by exploiting this series to see whether it changes any of our previous results. To this end, we estimated an annual bi-variate VECM representation for stock prices and the Basu, Fernald and Kimbal measure of TFP using 3 lags of data. The stock prices used are end-of-period prices. The results from sequential imposing our impact and long run restrictions to obtain orthogonal representations is given in Figures 9 and 10.

In Figure 9, we present the cross-plot of ϵ_2 and $\widetilde{\epsilon}_1$ recuperated from the bi-variate representation of TFP and SP using the BFK data. As can be seen, the two innovations are very highly correlated (.989 with s.d. .025), suggesting that both identification schemes isolate essentially the same shock. In Figure 10, we present the impulse responses for TFP and SP associated with the innovations ϵ_2 and $\widetilde{\epsilon}_1$. Although the response to both these shocks are once again very similar, the response of TFP is quite different from our previous observations. In effect, we now see that following an increase in stock prices, TFP does not increase for several years. The point estimates actually suggest that TFP only starts growing four years after the initial rise in the stock market. This long lag between stock prices increase and the increase in TFP is potentially consistent with a delayed impact of technological innovation on productivity, where the diffusion now appears quite

slow while it appeared to be rather quick with a less sophisticated measure of TFP.

As we previously indicated, there are two potential drawbacks with the BFK measure of TFP, that is, it is annual and its covers a limited period. As an alternative to the BFK measure, we therefore constructed an adjusted TFP measure, which we will denote by TFP^A, using the BLS's measure of capacity utilization (CU_t) to adjust our measure of capital services. This adjusted TFP measure is calculated as follows:

$$TFP_t^A = \log\left(\frac{Y_t}{H_t^{\bar{s}_h}(CU_t K S_t)^{1-\bar{s}_h}}\right)$$

Since the BLS measure of capital utilization is based mainly on manufacturing data, this correction is not above criticism. Nevertheless, it is an alternative worth exploiting to see how results based on this data compare to those based on either the BFK data or on our unadjusted TFP data. In order to make these comparisons, we first performed our orthogonalizations on annual bi-variate VAR over the period 1948-2000 using either the pair $\{TFP_t, SP_t\}$ or $\{TFP_t^A, SP_t\}$, where TFP refers to our original unadjusted TFP series while TFP^A refers to our series adjusted for variable rates of factor utilization. In Figure 11, we superimpose the responses of TFP and stock prices to the orthogonalized shocks ϵ_2 and $\widetilde{\epsilon}_1$ estimate for each system. In the case where we use the annualized unadjusted TFP data, we see that measured TFP increases quickly after the innovation in stock prices, reaching a peak after two years, decreasing slightly afterwards, and then resuming growth after about 4 years. This is quite similar to that observed when the quarterly version of this data was used. In contrast, the results based on the TFP data adjusted for variations in the rate of capacity utilization (TFP^A) are quite different than those based on unadjusted data, while interestingly they resemble the results obtained using the BFK data. In effect, we see that following the initial rise in stock prices, TFP^A does not overtake its initial level before approximately 3 or 4 years, and this whether we are examining the response to ϵ_2 and $\widetilde{\epsilon}_1$. In effect, we once again observe that the responses of the different variables to a ϵ_2 shock or to a $\widetilde{\epsilon}_1$ shock are very similar, that is, the impact and long run restrictions once again isolate essentially the same shock. This in

confirmed in Figure 12 where we provide a cross plot of ϵ_2 against $\widetilde{\epsilon}_1$ for both the cases where the system is estimated annually using either the unadjusted TFP measure (correlation .98 with s.d. .025) or the TFP measure adjusted for variables rates of capacity utilization (correlation .81 with s.d. .083). In order to further confirm the similarities and differences associated with adjusting TFP using the BLS measure of capacity utilization, Figure 13 reports results based on quarterly data. In particular, in Figure 13 we report the responses of SP and TFP to an ϵ_2 shock for both the case where TFP is un-adjusted and for the case where it is adjusted. As can be seen, the response of stock prices are almost unaffected by whether or not TFP is adjusted for variable utilization. In contrast, the short run response of TFP depends once again on whether our measure of TFP is adjusted for variable utilization. In the case where TFP is adjusted for variable utilization, the growth response is substantially delayed relative to the case where TFP is un-adjusted. The cross plot of ϵ_2 against $\widetilde{\epsilon}_1$ derived from the adjusted quarterly data is reported in Figure 12, and again exhibits a high correlation.

The results from using different measures of TFP suggest that our initial observation regarding the high correlation between ϵ_2 and $\widetilde{\epsilon}_1$ is very robust. In other words, it seems to be a robust feature of the data that the innovation in stock prices that is contemporaneously orthogonal to TFP is highly correlated with the innovations that cause permanent changes in TFP. In contrast, the timing of the response of TFP to such a shock depends heavily on whether or not TFP is adjusted for varying rates of capital utilization. In particular, when TFP is not adjusted for such a possibility, productivity appears to react quickly to the initial innovation in stock prices, which favors a quick diffusion interpretation. In contrast, when TFP is calculated, either according to the dis-aggregated method of BFK or simply adjusted using the BLS measure of capacity utilization, the response of TFP is substantial delayed which the first signs of improvement not arising before 3 years. In our opinion, the substantially delayed responses associated with the adjusted measures of productivity constitute the more believable response to the actual changes in technology. We now turn to examine whether this general pattern appears in higher dimensional systems.

5 Higher Dimension Systems

In this section, we study larger dimension systems in which –in addition to TFP and SP– consumption, investment and hours are alternatively or jointly introduced. For each system, we will show results that echo the results found in the bivariate case. Namely, the innovation in stock prices that is contemporaneously orthogonal to TFP, ϵ_2, and the shock that causes permanent changes in TFP, $\widetilde{\epsilon}_1$ are extremely highly correlated. Second, impulse responses associated with this shock show an aggregate expansion both in the short run and in the long run. Thirdly, the impulse response of TFP to this shock depends on whether the measure of TFP is adjusted or not for varying rates of capacity utilization. Fourthly, we will show that this shock explains a large fraction of macroeconomic movements at business cycle frequencies. Again, we find that these observations are robust to the model specification (lags, number of cointegration relations), but for concision the robustness exercises are not reported here. All the results we report in this section will be based, as in Section 4.1 on quarterly data over the period 1949-2000. Results based on yearly data give similar results.

5.1 A (TFP, SP, C) System

Our approach here parallels that presented in Section 2. Our objective is to sequentially impose orthogonalized restrictions on the moving average representation of (TFP, SP, C) as to derive, in one case, a shock that is contemporaneously orthogonal to TFP, while, in the other case, derive a shock that drives the long run movements in TFP. Then, given these two shock series, we can examine whether or not they are highly correlated and whether they induce similar dynamics. The result of this exercise can then be use to evaluate the plausibility of different view about the nature of fluctuations. In particular, if these shocks are found to be highly correlated, we claim that it provides evidence against models where surprise increases in productivity drive fluctuations. In contrast, we argue that a high correlation provides evidence in favor of the view whereby expectations, or news, about future improvements in the production frontier may be a relevant driving forces behind

macroeconomic fluctuations. The VECM for the system (TFP, SP, C) used in this section (i.e., the VECM used to derive the Wold representation) allows for two co-integrating relationships[9] and 5 lags.

Within this three variable system, it is easy to derive the shock series that drives the long run movements in TFP. This simply requires: (1) imposing the restriction that the 1,2 and 1,3 element of the long run matrix $(\sum_{i=0}^{\infty} \widetilde{\Gamma}_i(1))$ are equal to zero and (2) recuperating the shock $\widetilde{\epsilon}_1$.[10] In the case of recuperating the shock that is orthogonal to TFP, one must impose more structure. As in the bi-variate case, we impose the impact restriction that the 1,2 element of the impact matrix be equal to zero, and recuperate the associate shock ϵ_2. However, this is not sufficient to uniquely define ϵ_2. Having in mind that we would like our idea of a diffusion process to be embedded in an environment which allows for both a surprise technology shock and a temporary disturbance, we impose no restrictions related to the shock ϵ_1 as to let it potentially represent an unanticipated technology shock. As for the shock ϵ_3, we impose that it have no long run effect on either TFP or consumption, and therefore can capture a temporary shock. In effect, to understand this identification scheme, it is helpful to consider the following model of TFP.

$$\begin{aligned} TFP_t &= R_t + D_t + \nu_t \\ R_t &= R_{t-1} + \eta_{1,t} \\ D_t &= \sum_{i=0}^{\infty} d_i \eta_{2,t-i}, \quad d_0 = 0, \quad d_i \leq d_{i+1}, \quad \lim_{i \to \infty} d_i = 1 \\ \nu_t &= \rho \nu_{t-1} + \eta_{3,t}, \qquad 0 \leq \rho < 1 \end{aligned}$$

In the above case, TFP is driven by three component: the first component being a random walk, the second component being a diffusion process (as we modelled previously), and the third

[9] Using again the Nyblom and Harvery test, we found that these data do not reject 2 versus 1 co-integrating relationship at the 1%, but do reject it at the 5% level. Since we want to be cautious with respect to possible mispecification bias, we choose to allow for two co-integrating relationships instead of 1.

[10] In order to get a complete orthogonalized representation, it is necessary to impose one more restriction. We choose to impose that the 2,3 element of the long run matrix was also zero. However, as is well known, this additional restriction is only needed to separate the shocks $\widetilde{\epsilon}_2$, $\widetilde{\epsilon}_3$ and does not influence $\widetilde{\epsilon}_1$.

component being a temporary disturbance (possibly a measurement error). If this is the data generating process for TFP and these are the main shocks in the environment, then the structural impact matrix for a system composed of TFP, SP and consumption will have a zero for its 1,2 element (regardless of the precise theory for stock prices and consumption). Moreover, as long as the environment satisfied balance growth and that stock prices continue to follow a martingale, then the third structural shock will have a zero long run effect on both TFP and consumption. Hence, if the data generating process satisfies these conditions, then the recuperated ϵ_2 shock should correspond to the innovation to the diffusion process (news). Note however, that if this is the data generating process, then the long run restriction should isolate a combination of the surprise technology shock and the diffusion shock since both can have a long run effect on TFP. Hence, in the three variable case, if we find a very high correlation between ϵ_2 and $\widetilde{\epsilon}_1$ it suggests that the surprise component of productivity is of minor importance in fluctuations.

The impulse responses associated with the shocks ϵ_2 and $\widetilde{\epsilon}_1$ are presented in Figure 14 and the plots of ϵ_2 versus $\widetilde{\epsilon}_1$ are presented in Figure 15. In both figures, we report results associated with estimating the system using either our baseline TFP measure or using our measure adjusted for variable rates of capacity utilization.[11] As is clear from Figure 15, ϵ_2 and $\widetilde{\epsilon}_1$ are again highly correlated, regardless of which measure of TFP used: the correlation is .999 with s.d. .002 with non adjusted TFP and .92 with s.d. .03 when we adjust for variable rates of capacity utilization. Moreover, Figure 14 indicates that these shocks induce similar dynamics and that the responses of consumption and stock prices to these shocks are very little affected by the measure of TFP used. However, once again we can notice that the timing of the response of TFP to both ϵ_2 and $\widetilde{\epsilon}_1$ depends heavily on the measure of TFP used. When we use the the un-adjusted measure, TFP starts increasing after one quarter. In contrast, with the adjusted TFP series, the short run response is actually negative, and growth beyond its initial level take somewhere between 12 and 16 quarters, which is consistent with what we observed using the annual BFK data. [12]

[11] Since we choose to focus on results based on quarterly data, we cannot use the the BFK data.

[12] Note that there are at least two simple mis-measurement interpretations to the initial negative response to

For completeness, in Figure 16, we present the response of the economy to the shocks ϵ_1 and ϵ_3, where ϵ_1 is an unrestricted shock and ϵ_3 is a shock that is restricted to have no long run effects on TFP and consumption. As can be seen, the responses associated with ϵ_3 look like a the responses predicted by a temporary stock to TFP in many standard models. In contrast, the responses of variables to the ϵ_1 shock is harder to interpret, but since this shock does not account for much of the short run variance of any of the three variables, it may not be worth interpreting.

5.2 Four-Variables Systems

We now extend our analysis to a four-variables system where we begin by adding hours worked (in levels) to our system composed of TFP, stock prices and consumption. Our objective is again to recuperate from one representation a shock (denoted ϵ_2) that is an innovation in stock prices which is orthogonal to TFP, and to recuperate from another representation a shock (denoted $\widetilde{\epsilon}_1$) that is associated with permanent movements in TFP. The $\widetilde{\epsilon}_1$ shock can be isolated by imposing that the long run matrix $\widetilde{\Gamma}(1)$ be lower triangular. In order to isolate the shock ϵ_2, we do the following: (1) we impose no restriction related to the shock ϵ_1 as to allow it to potential capture a traditional surprise productivity shock, (2) we impose that the 1,2 element of the impact matrix Γ_0 be zero as to assure that ϵ_2 is not contemporaneously correlated with TFP, (3) we impose as before that the first and third element of the third column of the long run matrix be zero, as to potentially allow ϵ_3 to a temporary shock to technology and (4) we impose that ϵ_4 is an hours specific shock, i.e, that there are zeros in the first three element of the last column of the impact matrix (this last shock can be interpreted as a measurement error in hours worked).

Figure 17 displays the response of the four variables to the shocks ϵ_2 and $\widetilde{\epsilon}_1$. As in the case of the three variable system, we once again report results based on using our unadjusted TFP

adjusted TFP to either the ϵ_2 or $\widetilde{\epsilon}_1$ shock. The first is that our correction for varying capital utilization maybe excessive since it is based on high cyclical manufacturing data. Hence, the adjusted TFP series may inherit a counter-cyclical bias. The second is that some investments, in learning for example, may not be properly measured leading to counter-cyclical bias if such investment are pro-cyclical. In any case, given that all the results (adjusted or not) show that TFP is still approximately equal to it initial level 12 to 16 quarters after the innovation in stock prices, the analysis strongly suggests that the real growth in TFP does not start until a few years after the initial innovation in stock prices.

measure as well as our adjusted measure. Figure 18 displays the cross-plot of ϵ_2 against $\widetilde{\epsilon}_1$, which look similar to the previous plots as we observe a very high correlation (.993 (s.d. .008) with no adjustment of TFP, .990 (s.d. .01) with adjustment).

There are three aspects worth noticing in Figure 17. First, the response of consumption, hours and stock prices are very similar regardless of the measure of TFP used. Second, there is a substantial hump shaped response of hours to either the shock ϵ_2 or $\widetilde{\epsilon}_1$. In particular, this hump response last about 10 to 12 quarters, with the hump being echoed mildly in consumption.[13] Finally, as before, the timing of the response of TFP depends heavily on the measure of TFP used. When we use our adjusted measure of TFP (TFP^A), growth in TFP above it initial level arises only 12 to 15 quarters after the initial jump in stock prices. In contrast, in the case where we use our unadjusted measure of TFP, measured productivity appears to go through temporary boom, which is precisely what is expected if there are important cyclical variation in the rate of capital utilization. It is interesting to also note that the permanent growth in TFP arrives after the period of temporary boom in consumption and hours. In this sense, this way of looking at the data isolates a burst in economic activity that pre-dates the pickup in TFP growth. In effect, what is noticeable about the impulse responses in Figure 17 is the rich dynamics over the first 2 to 3 years. During this period, the economy appears to go through an important temporary boom, a slight recession, followed by a period of substantial TFP growth. Given a technological diffusion interpretation of this shock, this temporary boom period may results from a period of time where agents in the economy try to best position themselves to take advantage of future technological change.

In order to evaluate the importance of this phenomena in business cycles, Figure 19 reports the variance decompositions for consumption (C), investment (I), output ($C + I$) and hours worked (H) for the ϵ_2 and $\widetilde{\epsilon}_1$ shock retrieved from the system based on either the adjusted or unadjusted measure of TFP. In order to calculate the variance decomposition for output and investment, we

[13] The observed positive response of hours worked to a shock that permanently changes productivity presented in Figure 17 runs counter to the results presented in Gali [1999], but is consistent with the results presented in Christiano, Eichenbaum, and Vigfusson [2003].

replaced hours worked in the 4-variable VAR by investment or output. The impulse responses associated with these two latter exercises are reported in Figure 20 and Figure 21.[14] As can be seen, the shape of the responses are very similar regardless of whether the fourth variable in the system is hours worked, investment or output.

The variance decompositions in Figure 19 indicates that ϵ_2, and similarly for $\widetilde{\epsilon}_1$, explain a substantial fraction of fluctuations at business cycle frequencies. In effect, given the interpretation of this shock as reflecting news about technological innovations, the variance decomposition results suggest that news shocks may a major source of business cycle fluctuations even if surprise changes in productivity may not be. Let us note that the second part of this observation is consistent with the findings of Basu, Fernald and Kimbal (2002), and others, who have recently questioned the relevance of surprise changes in productivity as a driving force behind business cycles. In effect, Basu, Fernald and Kimbal arrived at their conclusion by noticing that TFP growth does not have a strong business cycle component. This observation is consistent with ours since our evidence suggests that the main growth in TFP (once corrected for utilization) arises well after the the boom period associated with the initial innovation in stock prices. However, our evidence goes a step further by suggesting that, even if surprise changes in productivity may not be an important source of fluctuations, technological change may nevertheless be at the heart of business cycle phenomena through its potential to generate expectational driven booms.[15]

[14] These results are based on estimating and inverting VECM representations where we allow for three cointegrating relationships. Note that the Nyblom and Harvey test was found not to reject three cointegrating relationships in favor of two at the 1% level. Since we prefer to err on the side of allowing for two many co-integrating relationships instead of too few, we opted for a specification with three cointegrating relationships.

[15] Our results indicate that innovations in stock prices predict future growth in productivity and that growth in productivity is proceeded by increases in stock prices. When observing such a result, one may be tempted to infer that the stock market is therefore informational efficient. However, such an inference from our results would be exaggerated for at least two reasons. First, our results do not tell us whether stock prices changes in anticipation of future productivity growth are of the "right" magnitude. In effect, our results are consistent with stock prices over or under reacting to information. Second, our results only indicate what happens on average, thereby they are consistent with the stock market making occasion large predictions errors. Hence, our result should not be seen as ruling out occasional speculative bubbles. In terms of the information content of stock prices, our results are best interpreted as suggesting that stock prices contain substantial information about future productivity, but they do not rule out that stock prices may also contain substantial noise.

6 Discussing Alternative Structural Interpretation of the Results

The results presented in Sections 4 and 5 suggest that stock prices anticipate future growth in productivity by several years, and that during the stage prior to the increase in productivity, the economy experience a boom in both consumption, investment and hours worked. In this section we want to discuss the extent to which standard dynamic general equilibrium models can explain such observations. In particular we want to emphasize the strengths and weakness of three different structural interpretations of these observations. Our objective here is mainly to indicate certain theoretical challenges that arise as one tries to explain these observation, with the hope tha these challenges may be addressed by future research.

Let us begin by considering interpretations in the spirit of the model present in Section 2.1, that is, interpretations where agent receive information –or news– about technological developments which lead to increases in productivity only in the future. Such increase in productivity could be delayed for several reasons. For example, it may be the case that the initial technological information is of a rather fundamental nature, and that actual commercial applications across a wide range of sectors only emerges with substantial delay. For such a case, we can ask what type of structural models of the economy, when perturbed by a news of future productivity growth, can generate the type of response observed in the data. In particular, we want to begin by asking how a simple Walrasian macro models, such as the standard one sector model, responds to a technological innovation that only has a delayed effect on productivity.

As is rather well known in the literature, the initial response of a standard one sector neo-classical macro model to an anticipated change in productivity is generally a fall in output. This arises due to the induced income effect on labor supply which favors an initial period with reduced work. Therefore, such a model is inconsistent with our observations. Obviously, one can modify a standard one sector model, for example by introducing adjustment costs to investment or time-to-build, in order to generate an immediate positive response of investment to a news of future productivity growth and thereby create an period of output growth that precedes the growth in

productivity. However, such modifications do not generally allow a full reconciliation with our observations. In effect, as discussed in Beaudry and Portier [2003], almost all standard neoclassical macro (that is, one sector models, two sectors model, models with or without adjustments costs to capital or variable rates of capacity utilization) have the prediction that an anticipated increase in productivity leads to an initial negative co-movement between investment and consumption, which is not what we observe in the data. The main exceptions to this rule is the case of multi-sector models with either sectorial adjustment costs or with cost complementarities across sectors (economies of scope). Although these later types of models may offer an explanation to our observations, they have yet to be extensively explored.[16] Hence, we can say with substantial confidence that most commonly used neo-classical macro models, if extended to allow for news about future productivity growth, do not offer an satisfactory explanation to our observations since they predict that a news shock should induce an initial period where consumption and investment move in opposite directions.

The second class of model we want to consider are sticky price models. In effect, we again want to ask whether news of future productivity growth in such models can reproduce our observations. The answer here is a qualified yes. The reason for the qualifier is two fold. First, in order reproduce the observations – in particular to reproduce a substantial boom in both consumption and investment that last for few years before the actual growth in productivity, requires a substantial degree of price stickiness. Second, the capacity of such a model to reproduce the observations depends heavily on the nature of monetary policy. For example, if monetary policy is conducted optimally, then the model should behave as if there were no price stickiness. However, if such is the case and the real structure of the model is akin to a one sector neo-classical model, then the predictions of a negative correlation between consumption and investment following a news shock reemerges. Hence, it is only if monetary policy is conducted sub-optimally that such a class of models can provide an explanation to our observations. Although less than optimal monetary policy is a plausible scenario, an explanation that relies entirely on sub-optimal monetary policy appears somewhat

[16] One exploration along these lines can be found in Beaudry and Portier [2000].

fragile and not entirely satisfactory. Therefore, while a sticky price model perturbed by news of future productivity growth offers a rather simple explanation to our observations, it has certain weaknesses.

As we mention at the end of Section 2, delayed impact of innovation on productivity is not the only candidate explanation to the observed high correlation we documented between ϵ_2 and $\widetilde{\epsilon}_1$. For example, if technological improvements are embodied in capital, and if measured capital does not fully account for changes in quality, then productivity improvements will only be reflected in measured TFP with a lag, even if the technology is available at the same time as we observe its reflection in the stock market. Consequently, a model where technological change is embodied in physical capital (as for example in the model of Greenwood, Herkowitz, and Krusell [1997]) may offer an alternative explanation to our observations. However, we view the capacity of such a model to explain the data as questionable for three reasons. First, in such a model, an improvement in technology implies that old capital becomes less valuable than new capital, which should imply a fall in stock prices instead of an increase.[17] Obviously, such a prediction could be overturned by the addition of adjustment cost to capital or the like. However, as pointed out by Christiano and Fisher [2003] , such a reversal of the stock prices implication is not easy to obtain without creating other counterfactuals implications. Second, in order to explain a substantial delay in TFP growth following a change in the stock price, it is necessary to introduce a very long lag (a few years!) between the time the investment is made and when it becomes productive, that is, one needs to introduce a very long time-to-build lag. Finally, if a long time-to build lag is introduced into embodied technological change model, then it is once again difficult to obtain a sustained increase in both consumption and investment prior to the improvement in measured TFP. The reason being that in the phase prior to arrival on line of the more productive capital, the model tends to behave like a one sector neo-classical model subjected to a change in expectations and consequently, it tends to produce a negative co-movement between consumption and investment.

[17] Greenwood and Jovanovic [1999] present an model of expected technological change which predicts a fall in stock prices in response to the news.

Hence, we believe that standard models of embodied technological change do not offer a simple explanation of observations, although a more elaborate version may be able to do so. For example, a model that explicitly models the process of technological implementation and of human capital adjustment may provide a more complete explanation.[18]

In summary, amongst standard macro model, the model which we believe most easily explains our observations is a sticky price model subjected to technological innovations (news) which effects productivity with a substantial delay. However, as we have pointed out, such an explanation is not above criticism since it requires sub-optimal monetary policy and very slow price adjustment. Moreover, it should be noted that we have only discuss to date models with a reversed causality between stock prices changes and productivity growth. However, it is possible that our observations could alternatively reflect a direct causal link between increases in stock prices and productivity improvement, as explored by Morck, Shleifer, and Vishny [1990]. For example, it may be that increases in stock prices reduces certain financial constraints and thereby favor risky investment in new technologies or favor entry of new more productive smaller, as in Jermann and Quadrini [2003]. Hence, in our opinion, the precise mechanism that best explains the data remains unclear and therefore suggests the need for future research along these lines.[19]

7 Conclusion

In this paper we have presented properties of the joint behavior of total factor productivity and stock prices which highlight new challenges for business cycle theory. In particular, we presented two orthogonalized moving average representation for these variables: one based on an impact restriction and one based on a long run restriction. We then examined the correlation between the innovations that drive the long run movements in TFP and the innovation which is contemporaneously orthogonal to TFP. We found this correlation to be positive and almost equal to 1,

[18]The model of Greenwood and Yorukoglu [1997] is an example of a model to incorporates both embodied technological change and implementation costs.

[19] The model of speculative growth by Caballero and Hammour [2002] is an example of an alternative approach which may help explain our observations.

indicating that permanent changes in productivity growth are proceeded by stock market booms. We showed why this observed positive correlation runs counter to that predicted by simple models where surprise changes in productivity drive fluctuations. We also discussed how the pattern could arise if agents have advanced information about future technological opportunities, or if productivity growth emerges as a delayed byproduct of a period high investment activity. In either case, the results suggests that changes in technological opportunities may be central to business cycle fluctuations even if surprise changes in productivity are not. Hence, these observations highlight the potential fruitfulness of reexamining the manner in which productivity growth is modelled in business cycle analysis. In particular, the type of model that is needed to explain the observations is one where agents recognize changes in technological opportunities well in advance of their effect on productivity, and where the recognition itself leads to a boom in both consumption and investment which precedes the growth in productivity.

References

Basu, S., J. Fernald, and M. Kimball (2002): "Are Technology Improvements Contractionary?," Discussion Paper 1986, Harvard Institute of Economic Research.

Beaudry, P., and F. Portier (2000): "An Exploration into Pigou's Theory of Cycles," mimeo, University of Bristish Columbia and Université de Toulouse, CEPR Discussion Paper 2996, October 2001, forthcoming *Journal of Monetary Economics*.

——— (2003): "Expextationally driven booms and busts in equilibrium cycle models," mimeo, University of Bristish Columbia and Université de Toulouse.

Benhabib, J., and R. Farmer (1999): "Indeterminacy and Sunspots in Macroeconomics," in *Handbook of Macroeconomics*, ed. by J. Talor, and M. Woodford, vol. 1A, pp. 387–448. North-Holland, New York.

Blanchard, O., and D. Quah (1989): "The Dynamic Effects of Aggregate Demand and Supply Disturbances," *The American Economic Review*, 79(4), 655–673.

Caballero, R., and M. Hammour (2002): "Speculative Growth," Working Paper 9381, National Bureau of Economic Research.

Christiano, L., M. Eichenbaum, and R. Vigfusson (2003): "What Happens After A Technology Shock?," International Finance Discussion Papers 768, Federal Reserve Board.

Christiano, L., and J. Fisher (2003): "Stock Market and Investment Good Prices: Implications for Macroeconomics," Working Paper 10031, National Bureau of Economic Research.

Doan, T. (1992): *Rats Manual*. Estima, Evanston, IL.

Fama, E. (1990): "Stock returns, expected returns and real activity," *The Journal of Finance*, 45(4), 1089–1108.

Flavin, M. (1983): "Excess volatility in the financial markets: A reassesment of the empirical evidence," *The Journal of Political Economy*, 91(6), 929–956.

Gali, J. (1999): "Technology, Employment and the Business Cycle: Do Technology Shocks Explain Aggregate Fluctuations?," *The American Economic Review*, 89(1), 249–271.

Greenwood, J., S. Herkowitz, and P. Krusell (1997): "Long-Run Implications of Investment-Specific Technological Change," *American Economic Review*, 87(3), 342–62.

Greenwood, J., and B. Jovanovic (1999): "The Information-Technology Revolution and the Stock Market," *The American Economic Review*, 89(2), 116–122.

Greenwood, J., and M. Yorukoglu (1997): "1974," *Carnegie-Rochester Conference Series on Public Policy*, 46(1), 49–95.

Hamilton, J. (1994): *Time Series Analysis*. Princeton University Press, Princeton, N.J.

Jermann, U., and V. Quadrini (2003): "Stock market boom and the productivity gains of the 1990s," mimeo, University of Pennsylvania and New York University.

Keynes, J. (1936): *The General Theory of Employment, Interest and Money*. London: Macmillan.

Mankiw, G., D. Romer, and M. Shapiro (1985): "An unbiased reexamination of stock market volatility," *The Journal of Finance*, 40(3), 677–687.

Morck, R., A. Shleifer, and R. Vishny (1990): "The stock market and investment: Is the market a sideshow?," *Brookings Papers on Economic Activity*, 1990(2), 157–215.

Nyblom, J., and A. Harvey (2000): "Tests of Common Stochastic Trends," *Econometric Theory*, 16, 176–199.

Pakes, A. (1985): "On Patents, R & D, and the Stock Market Rate of Return," *The Journal of Political Economy*, 93(2), 390–409.

Pigou, A. (1926): *Industrial Fluctuations*. MacMillan, London.

Schwert, G. (1990): "Stock returns and real activity: A century of evidence," *The Journal of Finance*, 45(4), 1237–1257.

Shiller, R. (1981): "Do stock prices move too much to be justified by subsequent changes in dividends?," *The American Economic Review*, 71(3), 421–436.

Shleifer, A. (1986): "Implementation Cycles," *Journal of Political Economy*, 92, 1163–1190.

Appendix

A Steps in Deriving the Structural Moving Average Representation of the New-Keynesian Model of Section 2.2

In this model, at time $t-1$, households set period t wages as to maximizes expected utility subject to a perceived demand for labor given by $(\frac{W_{j,t}}{W_t})^{\frac{1}{\rho_2}-1}L_t$, where W_t is the aggregate wage and L_t is the aggregate level of employment. In a symmetric equilibrium, this will give rise to a nominal wage given by

$$W_t = (\frac{\Lambda\rho_1^{\sigma-1}}{\rho_2})^{\frac{1}{\sigma}}(E_{t-1}M_t^{\sigma})^{\frac{1}{\sigma}}$$

where E_{t-1} is the expectation operator based on $t-1$ information. Assuming the cash-in-advance constraint is binding, at time t, consumption will be given by:

$$C_t = \frac{M_t}{P_t}$$

Given the perceived demand facing intermediate good firms, prices will be set as a markup over marginal cost as follows:

$$P_t = \frac{1}{\rho_1}\frac{W_t}{\theta_t}$$

The profits of the intermediate goods firms will be equal to $(\frac{1}{\rho_1}-1)C_t$, and the discounted sum of profits, which is the stock market value will be equal to $(\frac{1}{\rho_1}-1)\frac{C_t}{1-\beta}$. The log of the stock market value, denoted SP_t, will therefore be given by

$$SP_t = m_t - E_{t-1}m_t + \log(\theta_t)$$

Taking first differences, we have that

$$\Delta SP_t = \eta_{1,t} + \eta_{2,t} - \eta_{2,t-1}$$

while the first difference of TFP is simply equal to $\eta_{1,t}$, hence the structural moving average representation is:

$$\begin{pmatrix} \Delta TFP_t \\ \Delta SP_t \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ 1 & (1-L) \end{pmatrix} \begin{pmatrix} \eta_{1,t} \\ \eta_{2,t} \end{pmatrix} \tag{18}$$

B Main Text Figures

B.1 Figures related to section 3

Figure 1: Data

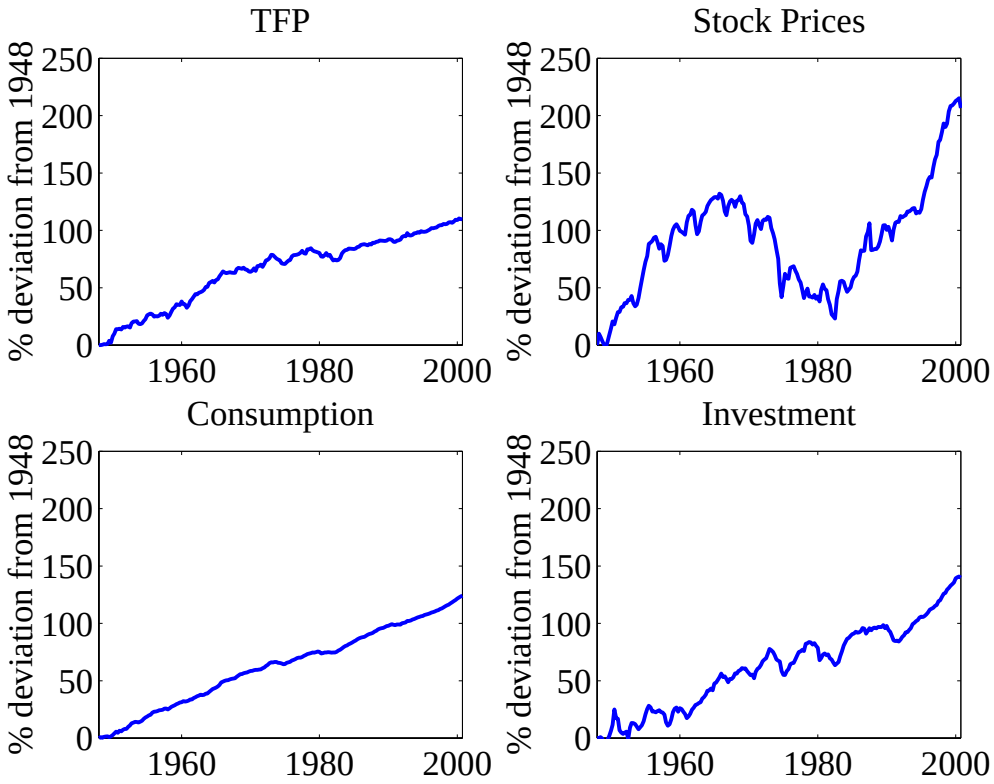

Those series are percentage deviations from 1948:Q1 level. All series have been previously divided by the 15 to 64 years old U.S. population. TFP is Total Factor Productivity in the non-farm business sector, as computed by the authors, Stock Prices is the Standard & Poors 500 index divided by the GDP deflator. Consumption is real personal consumption of non durable and services, while investment is real personal consumption of durable goods plus real fixed private domestic investment. See main text for more details.

B.2 Figures related to section 4

Figure 2: Impulse Responses to Shocks ϵ_2 and $\widetilde{\epsilon}_1$ in the (TFP, SP) VAR

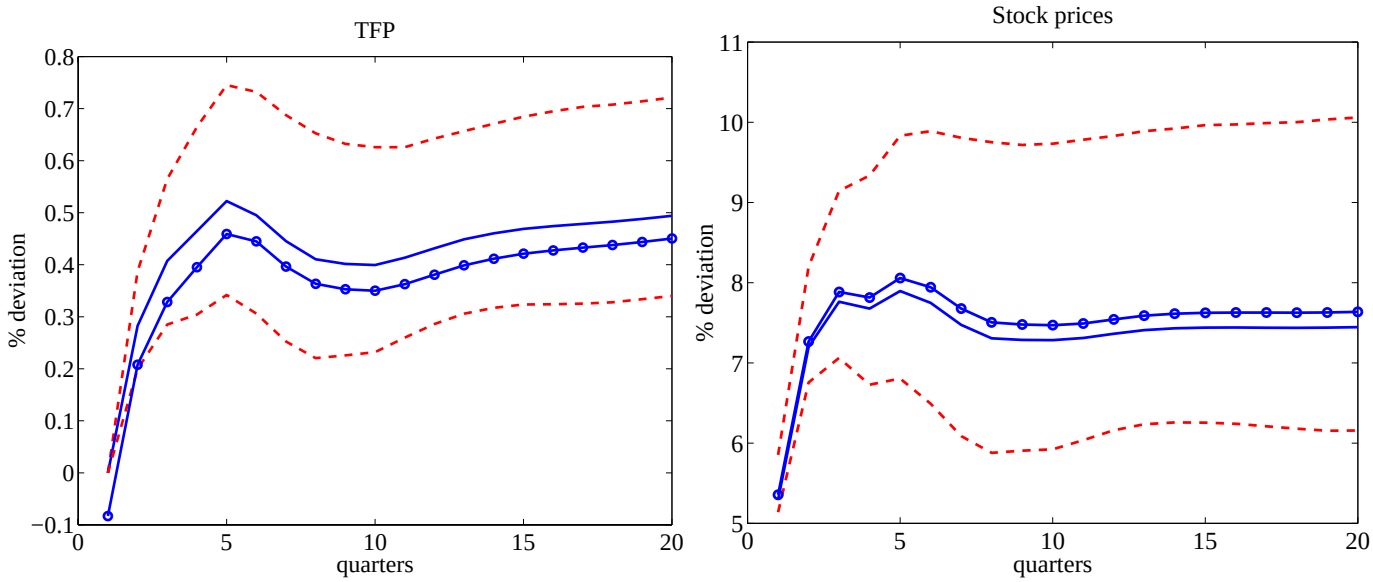

On each panel of this figure, the bold line represents the point estimate of the responses to a unit ϵ_2 shock (the shock that does not have instantaneous impact of TFP in the short run identification); the line with circles represents the point estimate of the responses to a unit $\widetilde{\epsilon}_1$ shock (the shock that has a permanent impact on TFP in the long run identification). Both identifications are done in the baseline bivariate specification. The unit of the vertical axis is percentage deviation from the situation without shock. Dotted lines represent the 10% and 90% quantiles of the distribution of the IRF in the case of the short run identification, this distribution being the bayesian simulated distribution obtained by Monte-Carlo integration with 2500 replications, using the approach for just-identified systems discussed in Doan [1992].

Figure 3: Share of the Forecast Error Variance Attributed to the ϵ_2 (left panel) or $\widetilde{\epsilon}_1$ (right panel) Shock in the (TFP, SP) VAR

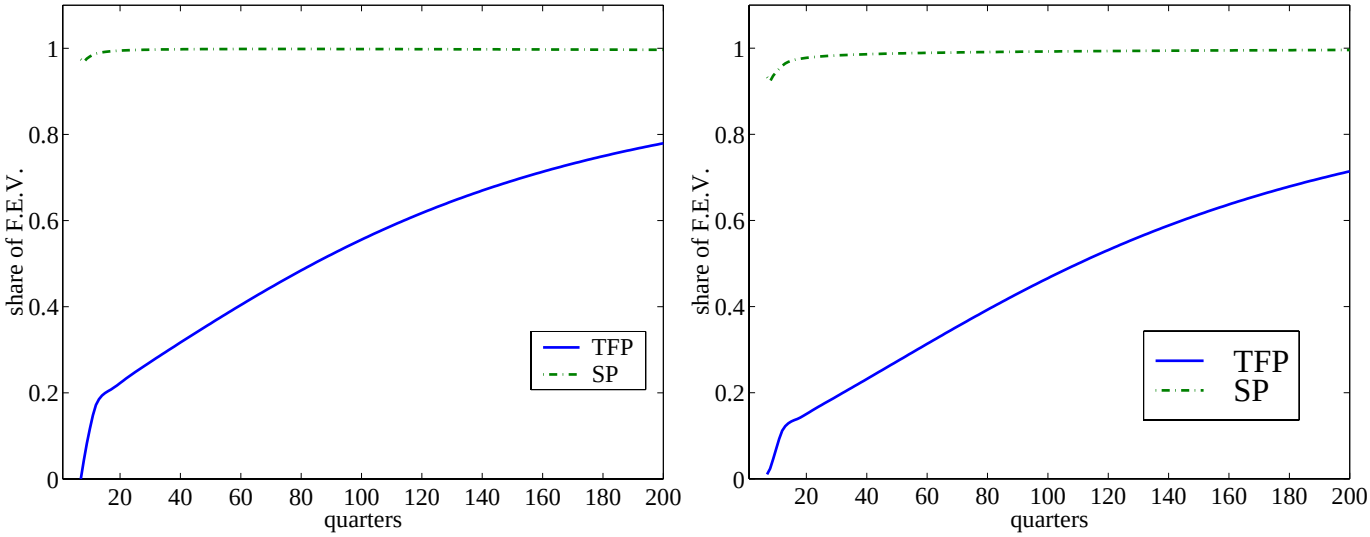

This figure displays the share of TFP and SP forecast error variance attributed to ϵ_2 (the shock that does not have instantaneous impact of TFP in the short run identification) (left panel) or to $\widetilde{\epsilon}_1$ (the shock that has a permanent impact on TFP in the long run identification)(right panel), both in the baseline bivariate specification.

Figure 4: ϵ_2 Against $\widetilde{\epsilon}_1$ in the (TFP, SP) VAR

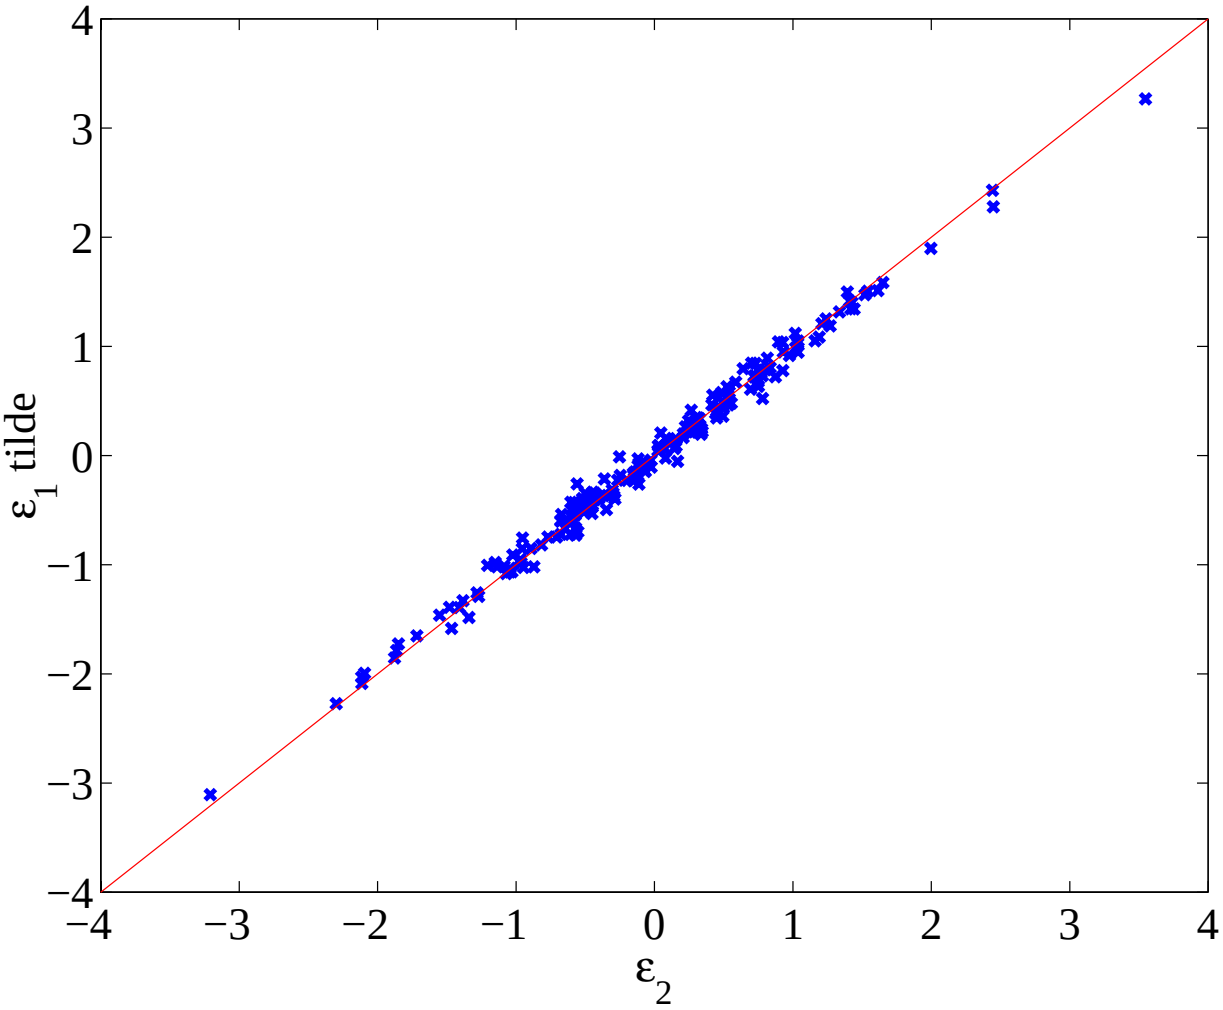

This figure plots ϵ_2 against $\widetilde{\epsilon}_1$. Both shocks are obtained from the baseline (TFP, SP) VAR, with 5 lags and one cointegrating relation. The straight line is the 45° line.

Figure 5: Robustness to Cointegration: Impulse Responses to ϵ_2 (upper panels) and $\widetilde{\epsilon}_1$ (lower panels) in the (TFP, SP) VAR

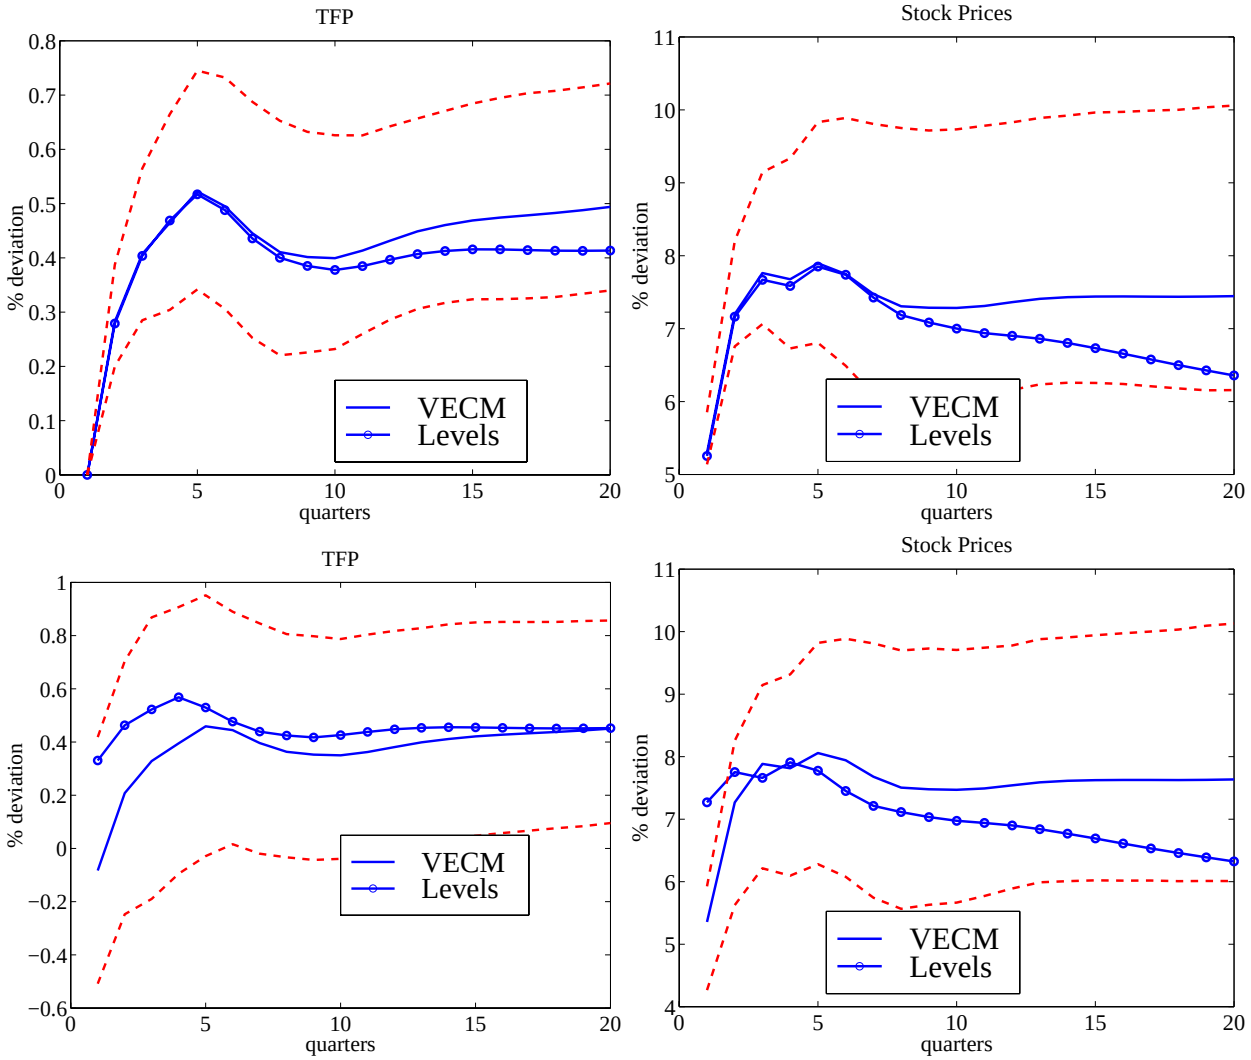

This figure displays the responses of TFP (upper left panel) and stock prices (upper right panel) to a unit ϵ_2 shock (the shock that does not have instantaneous impact on TFP in the short run identification), and the responses of TFP (lower right panel) and stock prices (lower left panel) to a unit $\widetilde{\epsilon}_1$ shock (the shock that has a permanent impact on TFP in the long run identification). The unit of the vertical axis is percentage deviation from the situation without shock. Each panel compares the responses on TFP and SP in the (TFP, SP) VAR estimated with one cointegrating relation or estimated in levels. the 10% and 90% confidence bands are computed using the VECM specification.

Figure 6: Robustness to the Lag Structure: Impulse Responses to ϵ_2 (upper panels) and $\widetilde{\epsilon}_1$ (lower panels) in the (TFP, SP) VAR

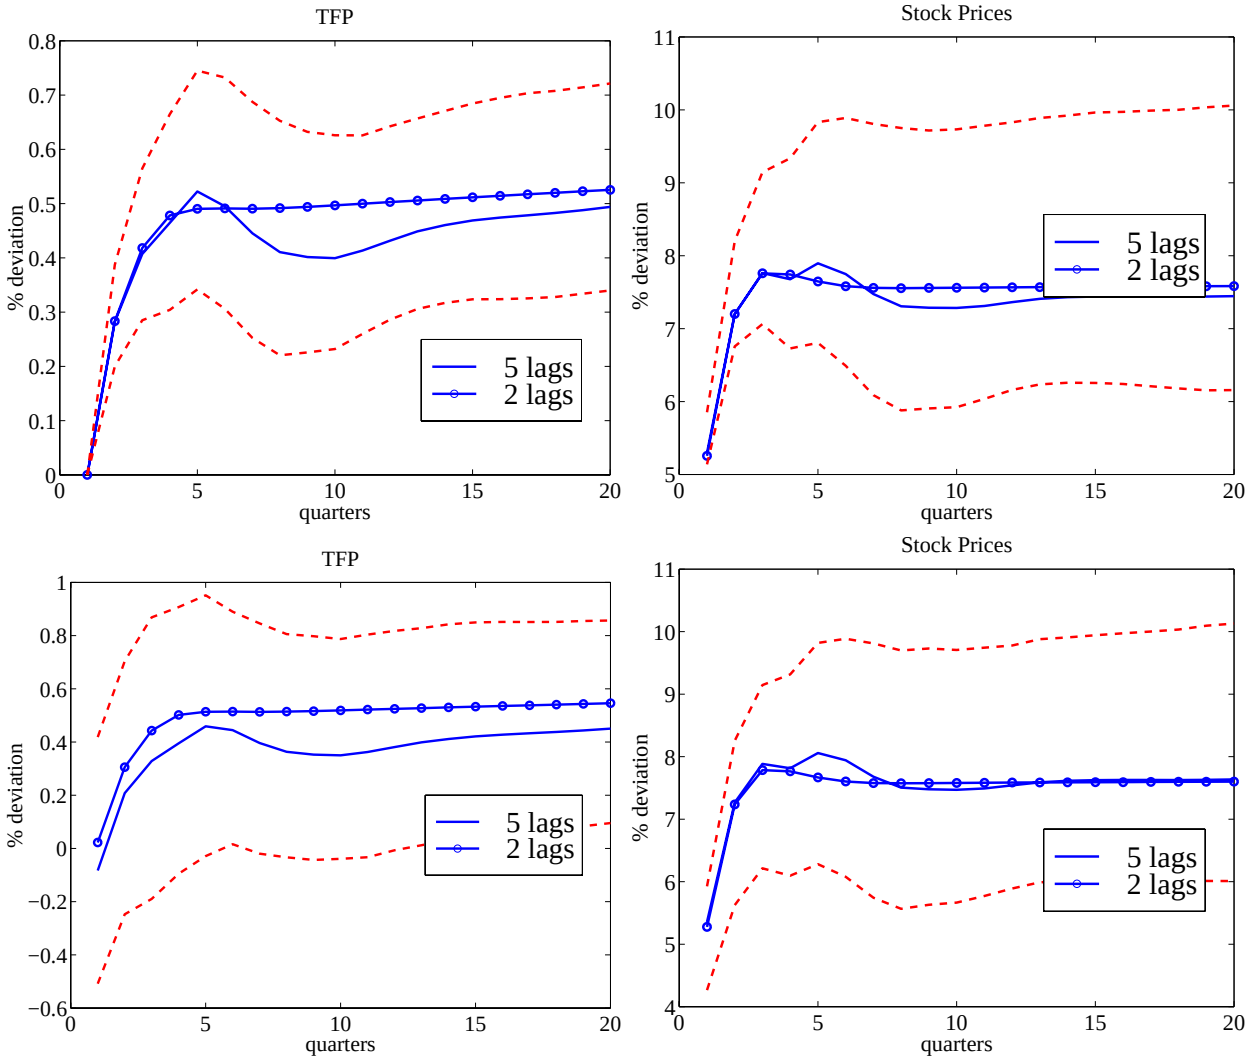

This figure displays the responses of TFP (upper left panel) and stock prices (upper right panel) to a unit ϵ_2 shock (the shock that does not have instantaneous impact on TFP in the short run identification), and the responses of TFP (lower right panel) and stock prices (lower left panel) to a unit $\widetilde{\epsilon}_1$ shock (the shock that has a permanent impact on TFP in the long run identification). The unit of the vertical axis is percentage deviation from the situation without shock. Each panel compares the responses on TFP and SP in the (TFP, SP) VAR estimated with one cointegrating relation and 5 or 2 lags . the 10% and 90% confidence bands are computed using the 5 lags VECM specification.

Figure 7: Robustness to Cointegration or Lag Structure: ϵ_2 Against $\widetilde{\epsilon}_1$ in the (TFP, SP) VAR in Levels (left panel) and with Two Lags (right panel)

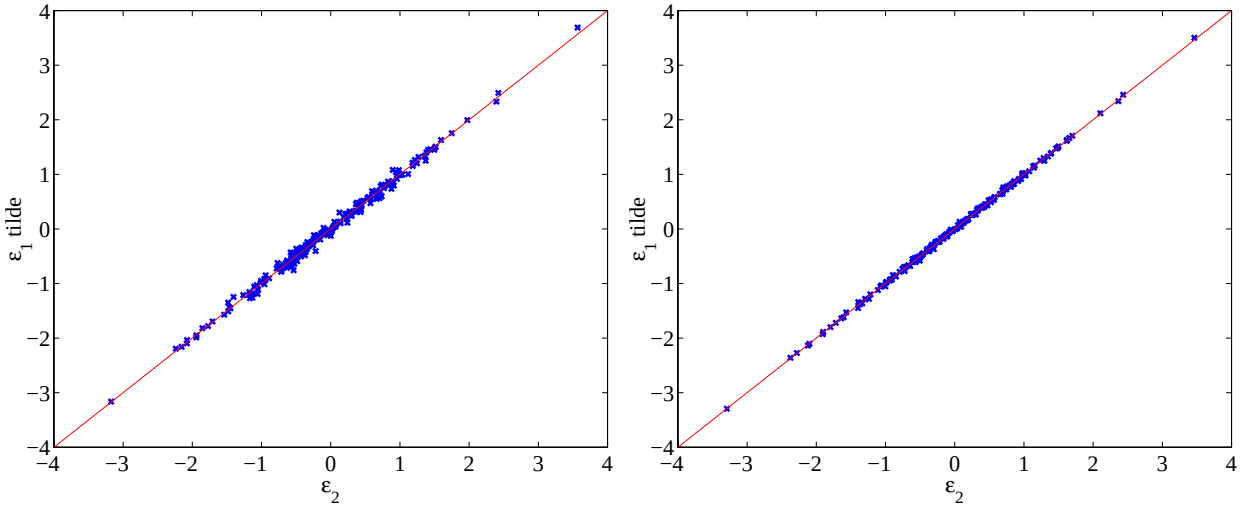

This figure plots ϵ_2 against $\widetilde{\epsilon}_1$. In the left panel, both shocks are obtained from the (TFP, SP) VAR estimated in levels, with 6 lags. In the right panel, both shocks are obtained from the (TFP, SP) VAR estimated in difference, with 2 lags and one cointegrating relation. In both panels, the straight line is the 45° line.

Figure 8: Impulse Responses to ϵ_2 in the (TFP, SP) VAR

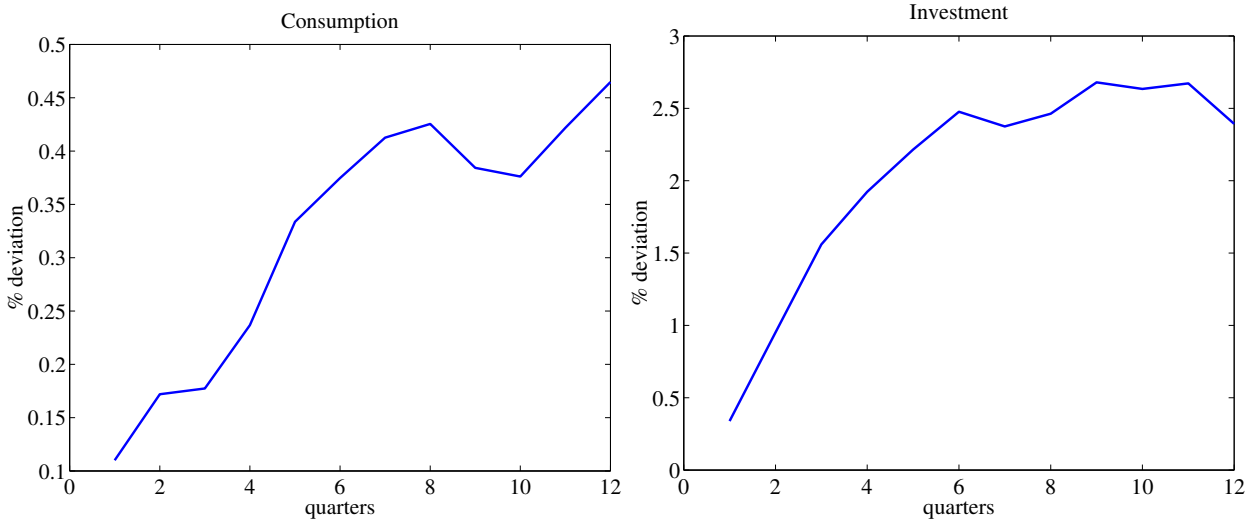

This figure displays the response of consumption and investment to a unit ϵ_2 shock (the shock that does not have instantaneous impact on TFP in the short run identification). The unit of the vertical axis is percentage deviation from the situation without shock (See the main text for more details).

Figure 9: ϵ_2 Against $\widetilde{\epsilon}_1$ in the (TFP, SP) VAR, Using Basu, Fernald, and Kimball [2002] Measure of TFP (annual, 1949-1989)

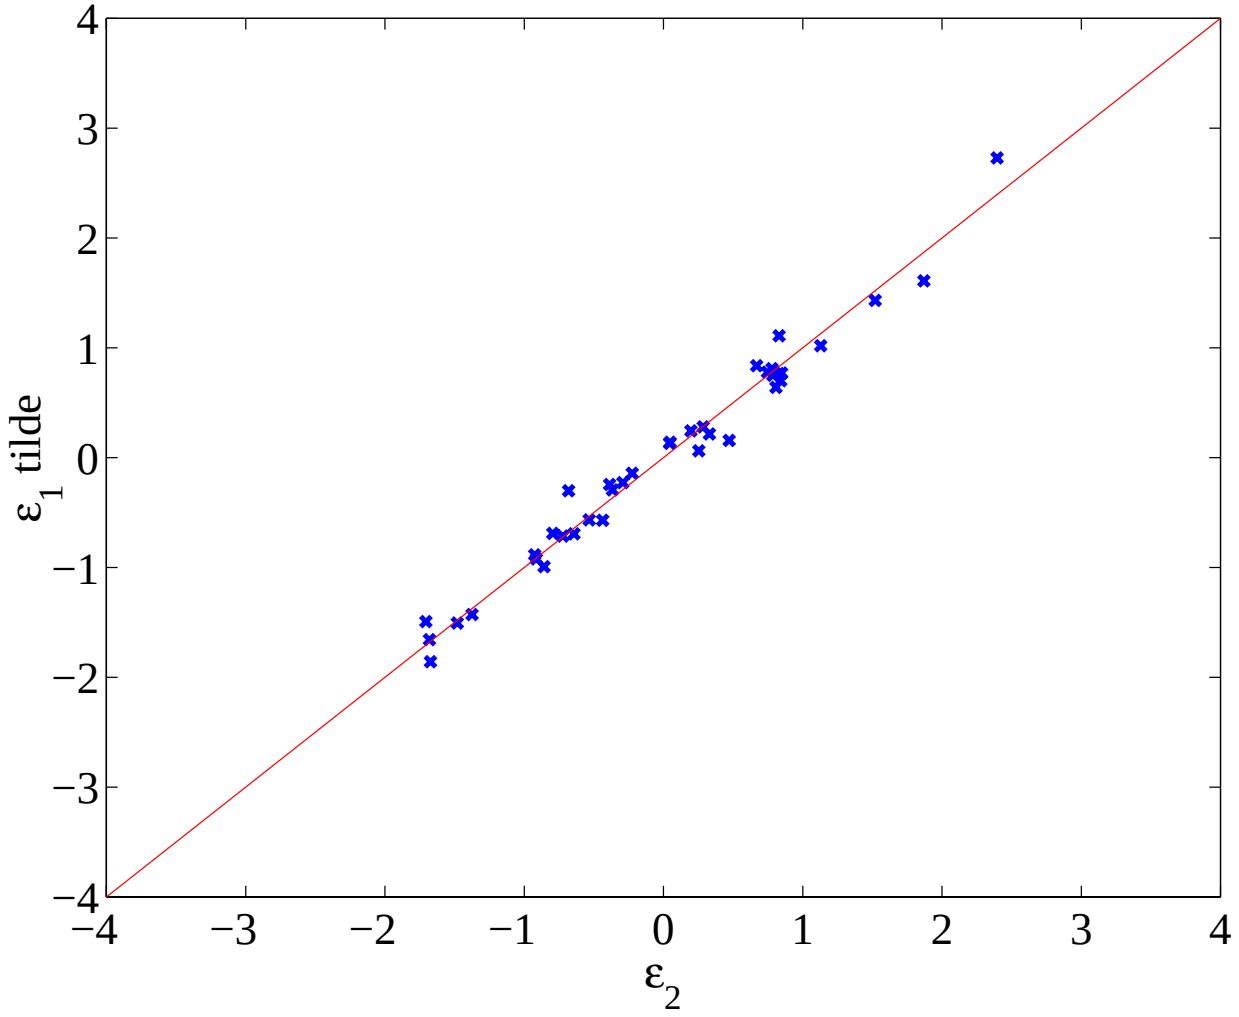

This figure plots ϵ_2 against $\widetilde{\epsilon}_1$. Both shocks are obtained from the baseline (TFP, SP) VAR, with 2 lags and one cointegrating relation. The straight line is the 45° line.

Figure 10: Impulse Responses to Shocks ϵ_2 and $\widetilde{\epsilon}_1$ in the (TFP, SP) VAR, Using citeasnounBASU/FERN/KIMB/02 Measure of TFP (annual, 1949-1989)

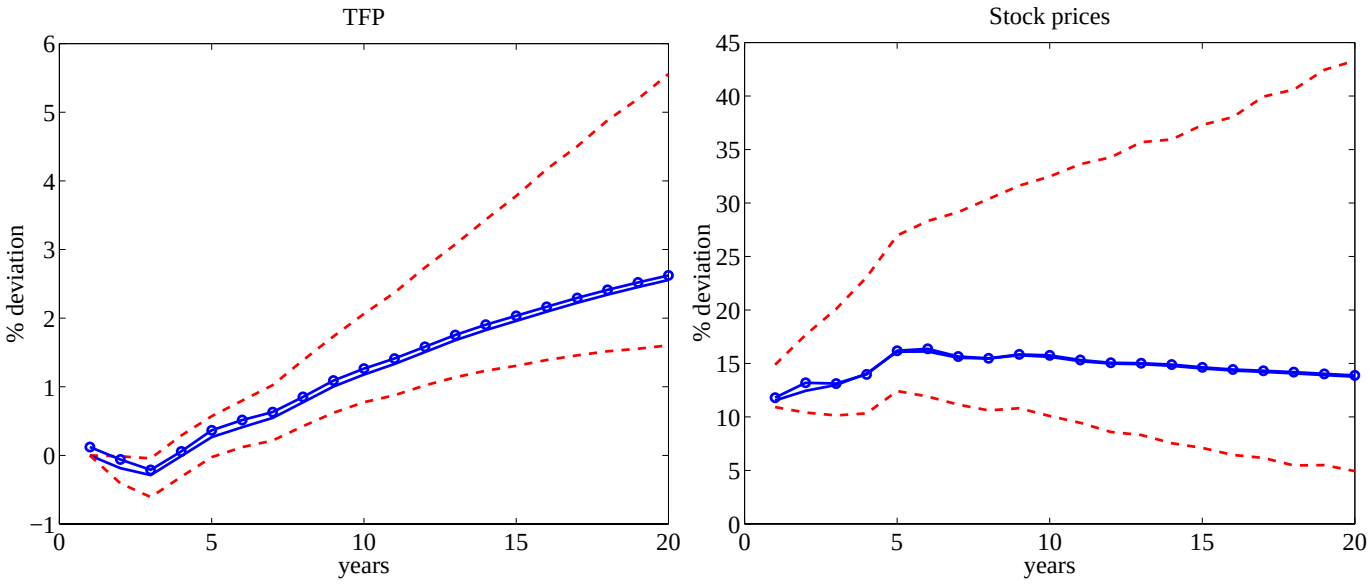

On each panel of this figure, the bold line represents the point estimate of the responses to a unit ϵ_2 shock (the shock that does not have instantaneous impact on TFP in the short run identification); the line with circles represents the point estimate of the responses to a unit $\widetilde{\epsilon}_1$ shock (the shock that has a permanent impact on TFP in the long run identification). Both identifications are done in the baseline bivariate specification. The unit of the vertical axis is percentage deviation from the situation without shock. Dotted lines represent the 10% and 90% quantiles of the distribution of the IRF in the case of the short run identification, this distribution being the bayesian simulated distribution obtained by Monte-Carlo integration with 2500 replications, using the approach for just-identified systems discussed in Doan [1992].

Figure 11: Impulse Responses to Shocks ϵ_2 and $\widetilde{\epsilon}_1$ in the (TFP, SP) VAR, Using Annual Observations (1948-2000), without Adjusting TFP for Capacity Utilization (top panels) or with TFP Adjustment (bottom panels)

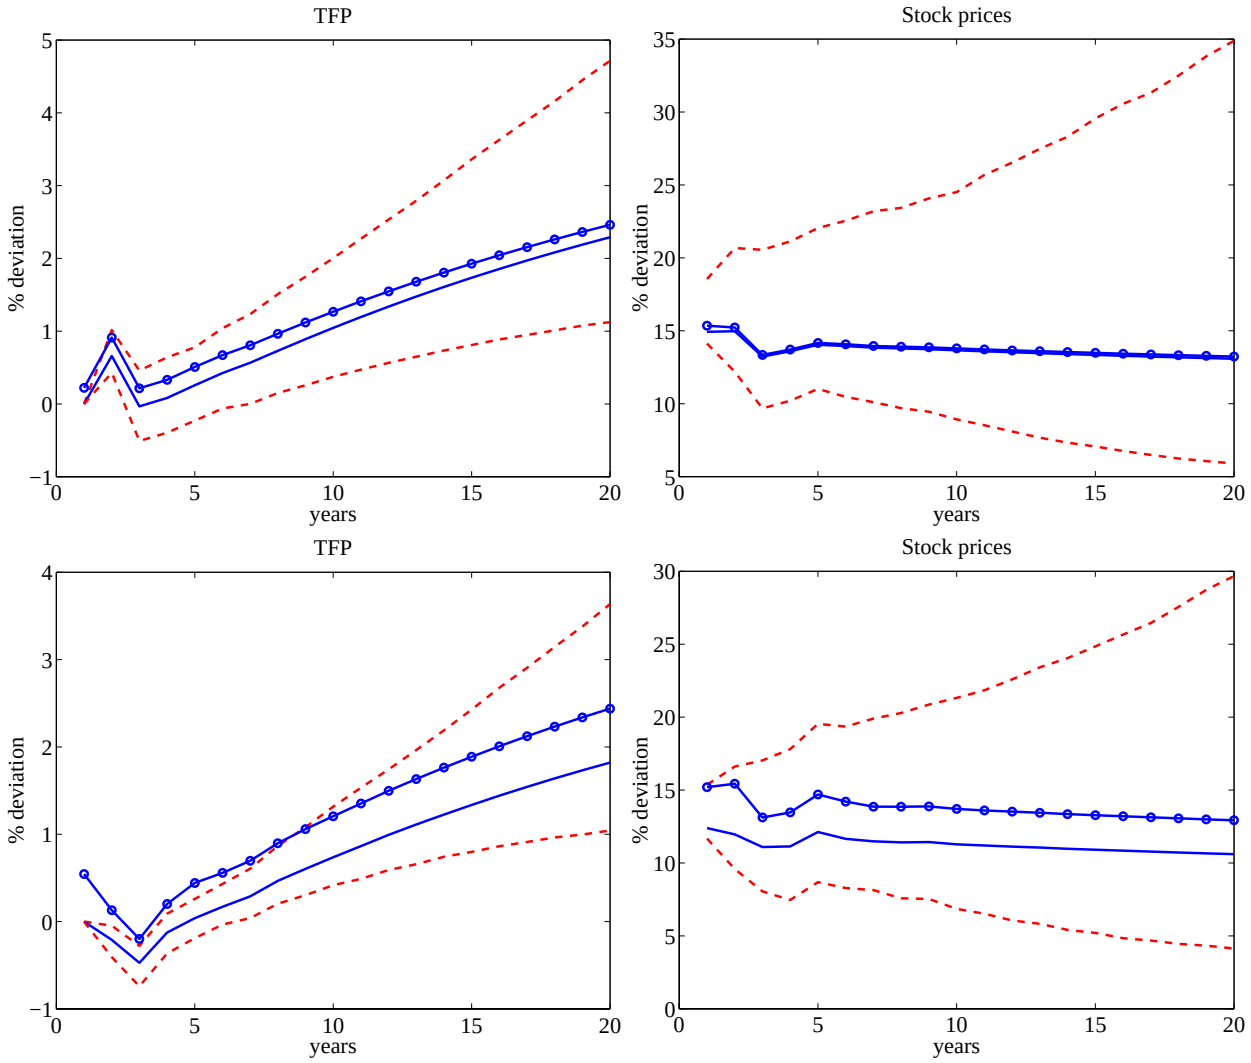

*On each panel of this figure, the bold line represents the point estimate of the responses to a unit ϵ_2 shock (the shock that does not have instantaneous impact on TFP in the short run identification); the line with circles represents the point estimate of the responses to a unit $\widetilde{\epsilon}_1$ shock (the shock that has a permanent impact on TFP in the long run identification). Both identifications are done in the baseline bivariate specification. The unit of the vertical axis is percentage deviation from the situation without shock. Dotted lines represent the 10% and 90% quantiles of the distribution of the IRF in the case of the short run identification, this distribution being the bayesian simulated distribution obtained by Monte-Carlo integration with 2500 replications, using the approach for just-identified systems discussed in Doan (1992)*49

Figure 12: ϵ_2 Against $\widetilde{\epsilon}_1$ in the (TFP, SP) VAR, Using Annual Observations (1948-2000), without Adjusting TFP for Capacity Utilization (left panel) or with TFP Adjustment

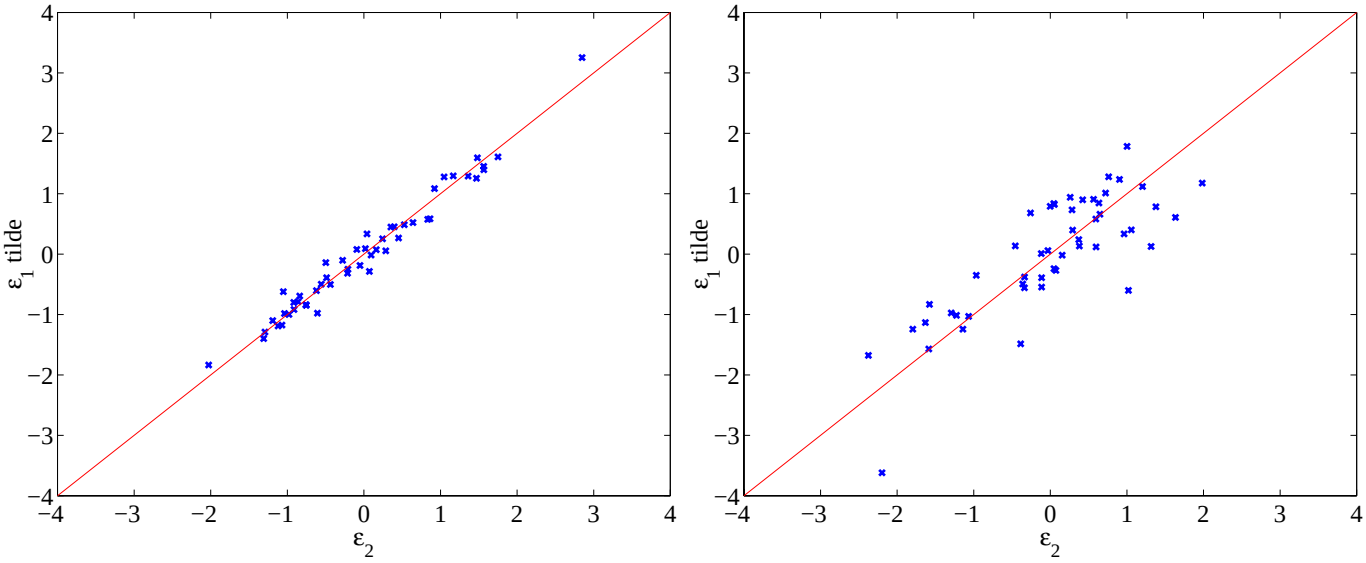

Each panel of this figure plots ϵ_2 against $\widetilde{\epsilon}_1$. Both shocks are obtained from the baseline (TFP, SP) or (TFP^A, SP) VAR, with 2 lags and one cointegrating relation. The straight line is the 45° line.

Figure 13: Impulse Responses to ϵ_2 in the (TFP, SP) VAR, Quarterly Data, with or without Correction for Variable Capacity Utilization

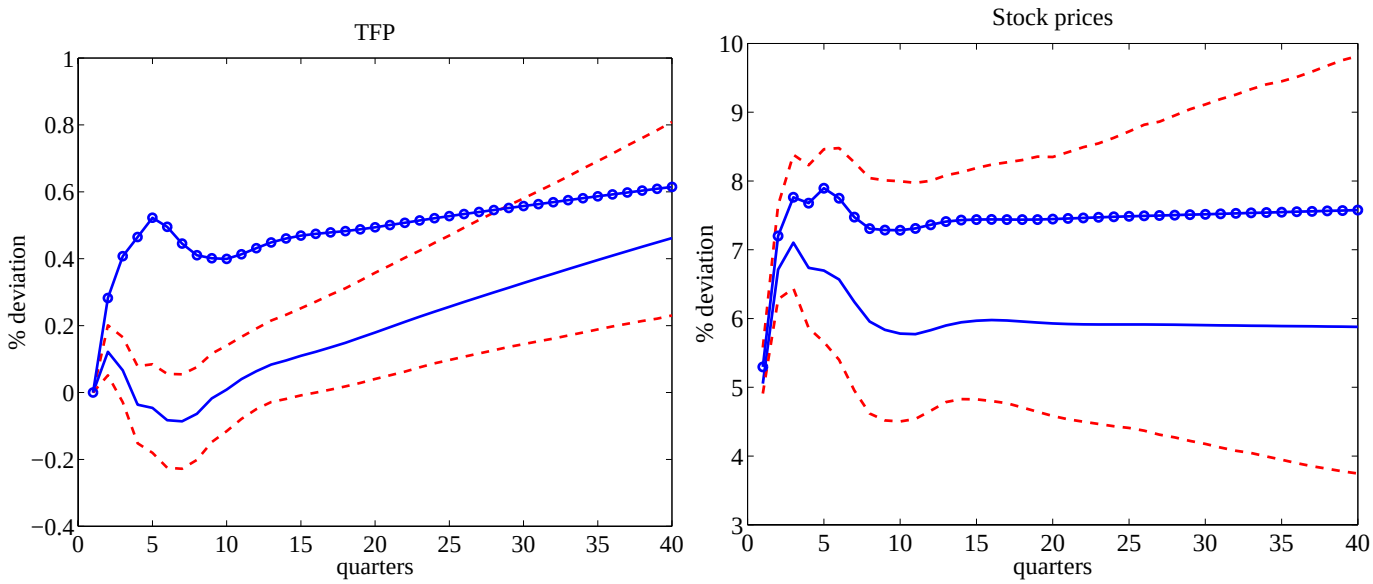

On each panel of this figure, the bold line represents the point estimate of the responses to a unit ϵ_2 shock (the shock that does not have instantaneous impact on TFP in the short run identification) in the VAR with adjusted TFP; the line with circles represents the point estimate of the responses to a unit ϵ_2 shock in the VAR with non adjusted TFP. The unit of the vertical axis is percentage deviation from the situation without shock. Dotted lines represent the 10% and 90% quantiles of the distribution of the IRF in the VAR with adjusted TYFP, this distribution being the bayesian simulated distribution obtained by Monte-Carlo integration with 2500 replications, using the approach for just-identified systems discussed in Doan [1992].

B.3 Figures related to section 5.1

Figure 14: Impulse Responses to ϵ_2 and $\widetilde{\epsilon}_1$ in the (TFP, SP, C) VAR, without Adjusting TFP for Capacity Utilization (upper panels) or with TFP Adjustment(lower panels)

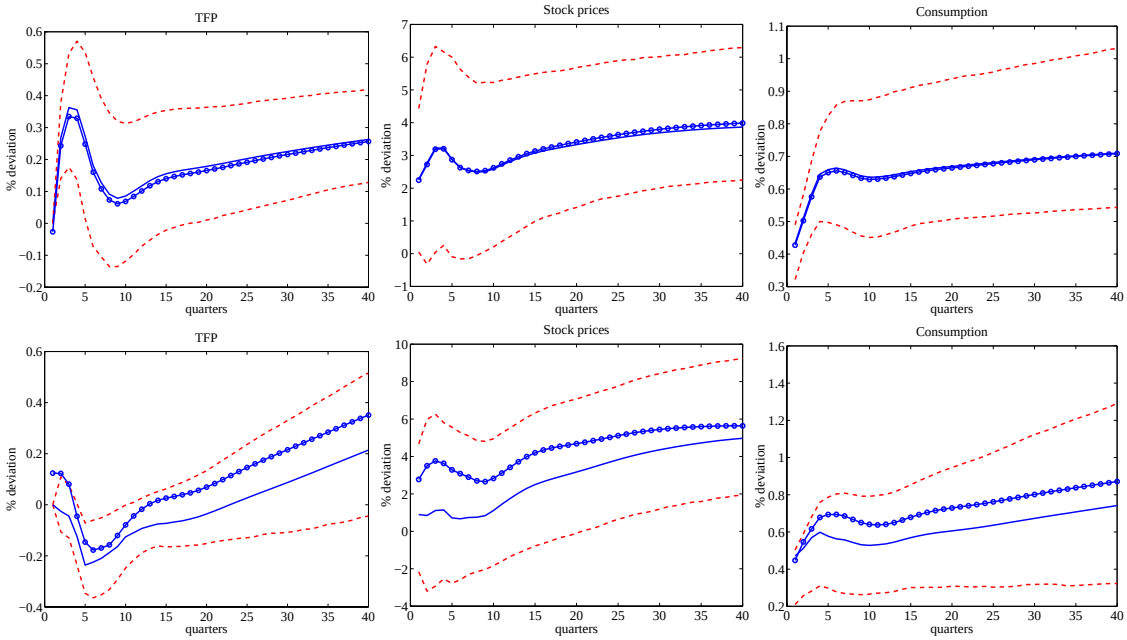

On each panel of this figure, the bold line represents the point estimate of the responses to a unit ϵ_2 shock (the shock that does not have instantaneous impact on TFP in the short run identification); the line with circles represents the point estimate of the responses to a unit $\widetilde{\epsilon}_1$ shock (the shock that has a permanent impact on TFP in the long run identification). Both identifications are done in the baseline bivariate specification. The unit of the vertical axis is percentage deviation from the situation without shock. Dotted lines represent the 10% and 90% quantiles of the distribution of the IRF in the case of the short run identification, this distribution being the bayesian simulated distribution obtained by Monte-Carlo integration with 2500 replications, using the approach for just-identified systems discussed in Doan [1992].

Figure 15: ϵ_2 Against $\widetilde{\epsilon}_1$ in the (TFP, SP, C) VAR, without (left panel) or with (right panel) Adjusting TFP for Capacity Utilization

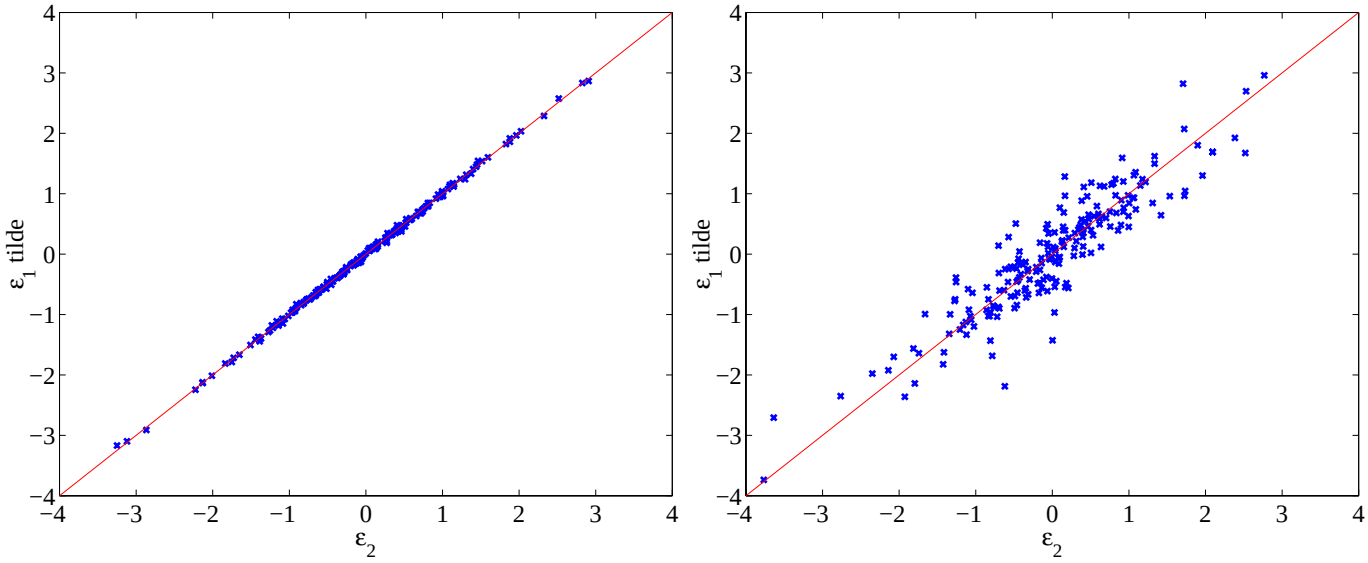

This figure plots ϵ_2 against $\widetilde{\epsilon}_1$. Both shocks are obtained from the baseline (TFP, SP, C) VAR, with 5 lags and two cointegrating relation. The straight line is the 45° line.

Figure 16: Impulse Responses to ϵ_3 (upper panels) and ϵ_1 (lower panels) in the (TFP, SP, C) VAR

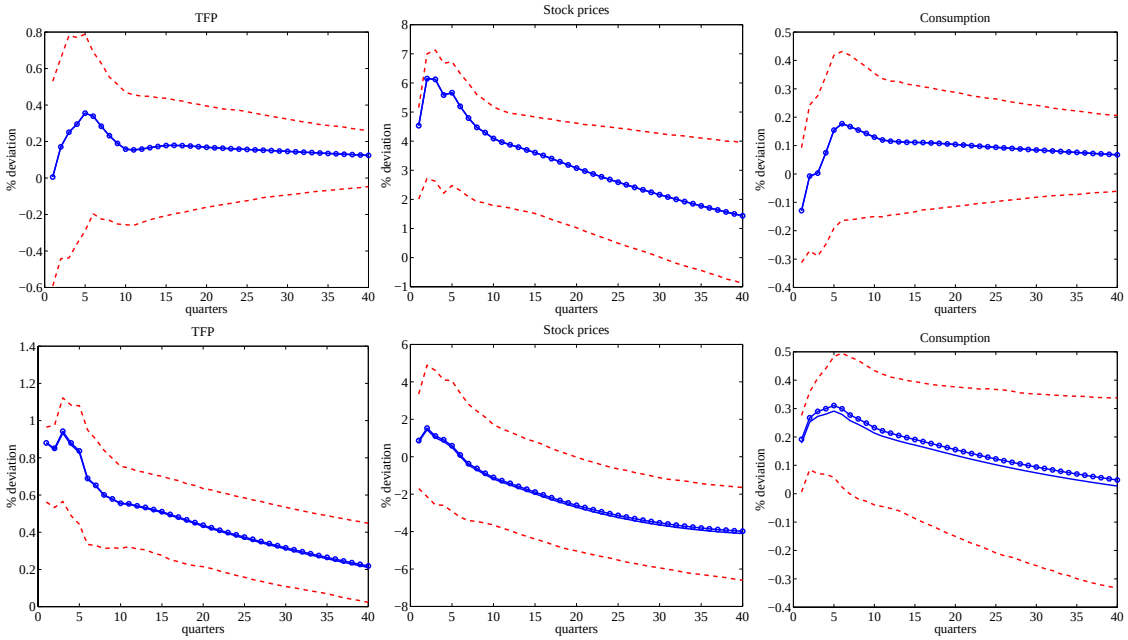

On each panel of this figure, the bold line represents the point estimate using non adjusted TFP, while the line with circles is the point estimate when adjusted TFP is used. ϵ_1 and ϵ_3 are the two "other" shocks in the short run identification. The unit of the vertical axis is percentage deviation from the situation without shock. Dotted lines represent the 10% and 90% quantiles of the distribution of the IRF when non adjusted TFP is used, this distribution being the bayesian simulated distribution obtained by Monte-Carlo integration with 2500 replications, using the approach for just-identified systems discussed in Doan [1992].

B.4 Figures related to section 5.2

Figure 17: Impulse Responses to ϵ_2 and $\widetilde{\epsilon}_1$ in the in the (TFP, SP, H) VAR, without (upper panels) or with (lower panels) Adjusting TFP for Capacity Utilization

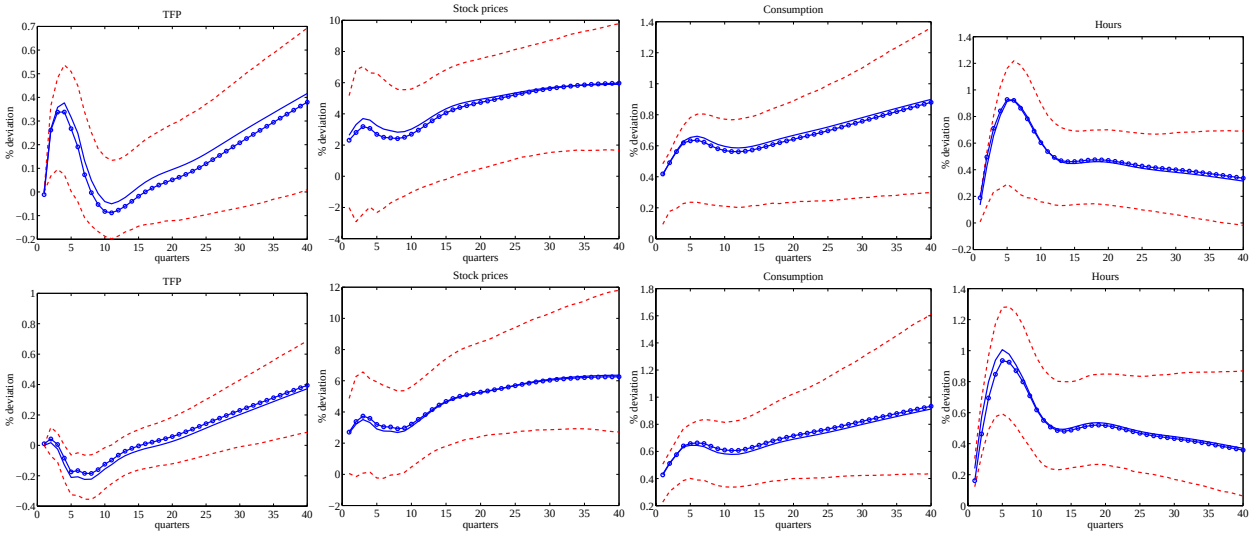

On each panel of this figure, the bold line represents the point estimate of the responses to a unit ϵ_2 shock (the shock that does not have instantaneous impact on TFP in the short run identification); the line with circles represents the point estimate of the responses to a unit $\widetilde{\epsilon}_1$ shock (the shock that has a permanent impact on TFP in the long run identification). Both identifications are done in the baseline bivariate specification. The unit of the vertical axis is percentage deviation from the situation without shock. Dotted lines represent the 10% and 90% quantiles of the distribution of the IRF in the case of the short run identification, this distribution being the bayesian simulated distribution obtained by Monte-Carlo integration with 2500 replications, using the approach for just-identified systems discussed in Doan [1992].

Figure 18: ϵ_2 Against $\widetilde{\epsilon}_1$ in the (TFP, SP, C, H) VAR, without (left panel) or with (right panel) Adjusting TFP for Capacity Utilization

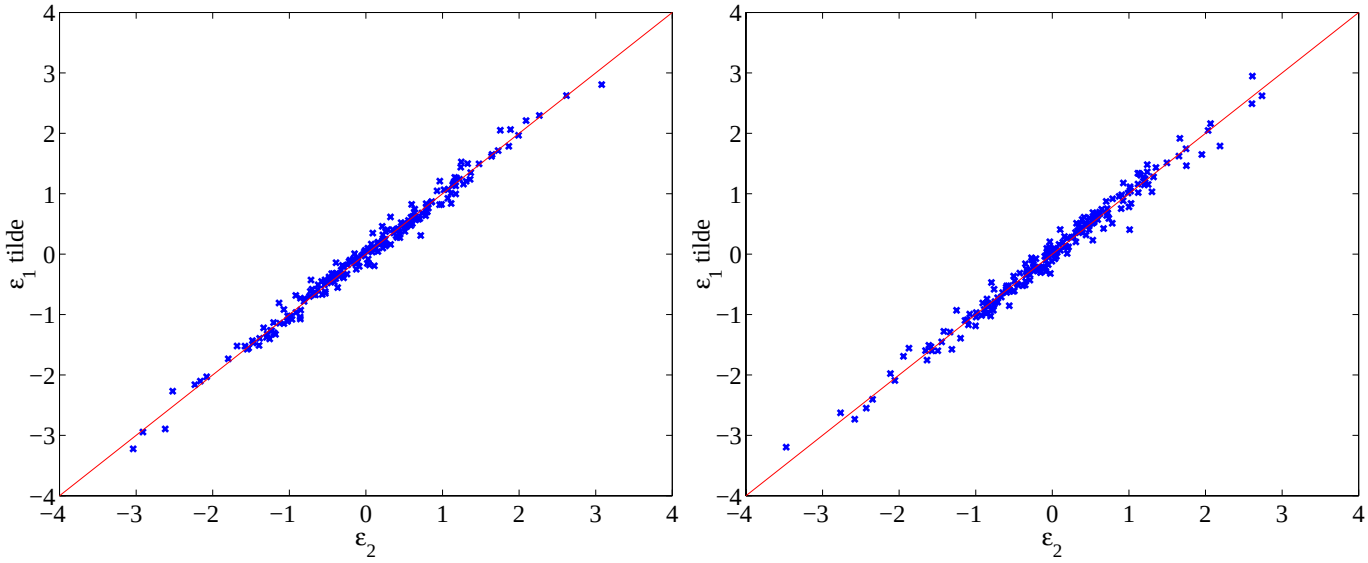

This figure plots ϵ_2 against $\widetilde{\epsilon}_1$. Both shocks are obtained from the baseline (TFP, SP, C) VAR, with 5 lags and two cointegrating relation. The straight line is the 45° line.

Figure 19: Share of the Forecast Error Variance of Consumption (C), Investment I, Output ($C+I$) and hours (H) attributable to ϵ_2 (left panel) and to $\widetilde{\epsilon_1}$ (right panel) in 4-variables VARs, with non adjusted TFP (top panels) or adjusted TFP (bottom panels)

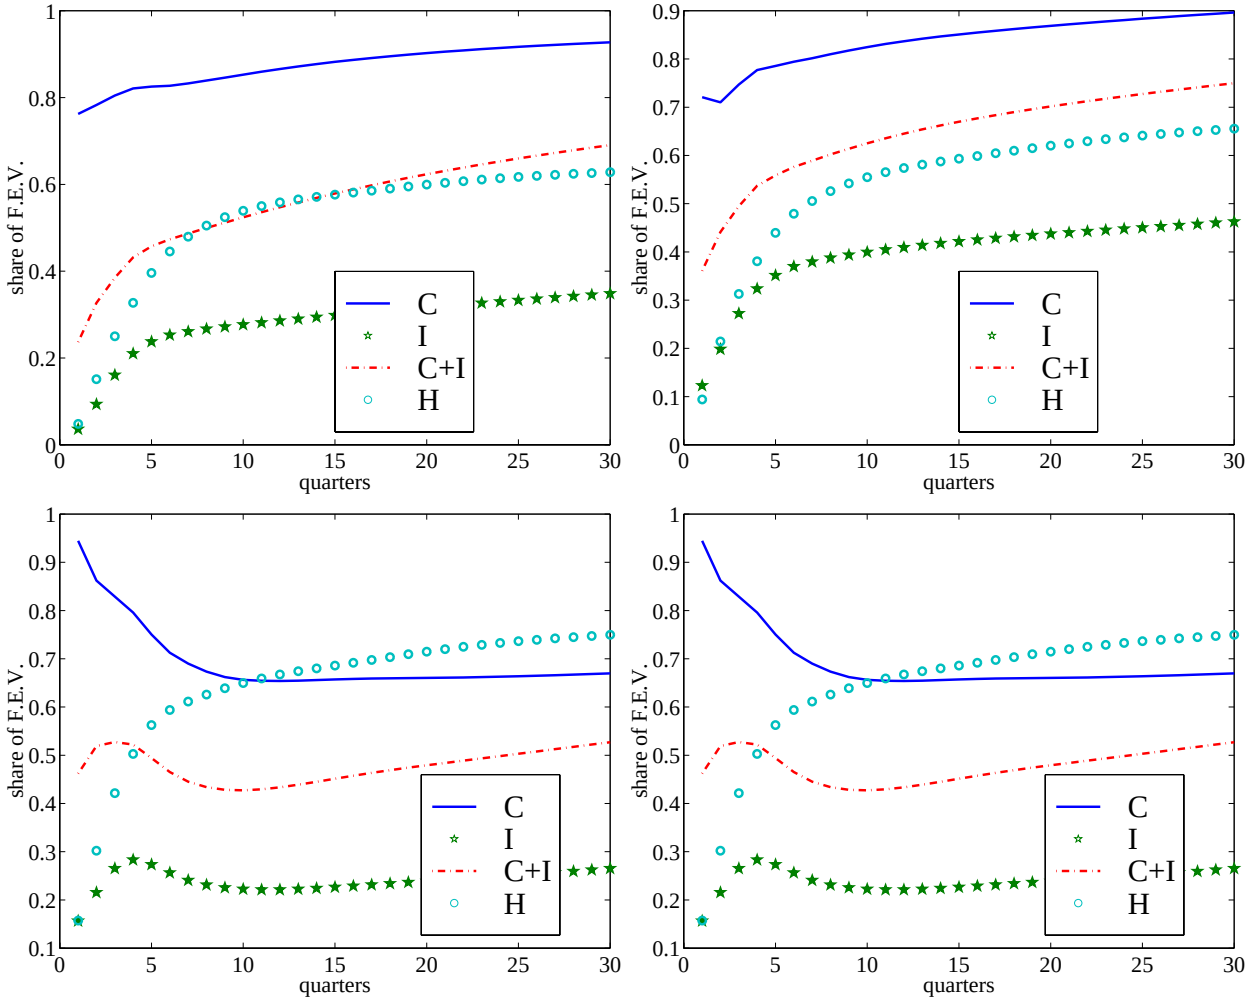

This figure has four panels. On the left panels is displayed the share of the forecast variance of consumption and investment that is attributable to ϵ_2 (short run identification) in the (TFP, SP, C, I) VAR, of output $(C+I)$ in the $(TFP, SP, C, C+I)$ VAR and of hours (H) in the (TFP, SP, C, H) VAR. The right panel presents the same information in the case of the shock $\widetilde{\epsilon}_1$ (long run identification). The top row uses a non adjusted measure of T5FP, while TFP is adjusted for variable capacity utilization in the bottom row.

Figure 20: Impulse Responses to ϵ_2 and $\widetilde{\epsilon}_1$ in the in the (TFP, SP, C, I) VAR, without (upper panels) or with (lower panels) Adjusting TFP for Capacity Utilization

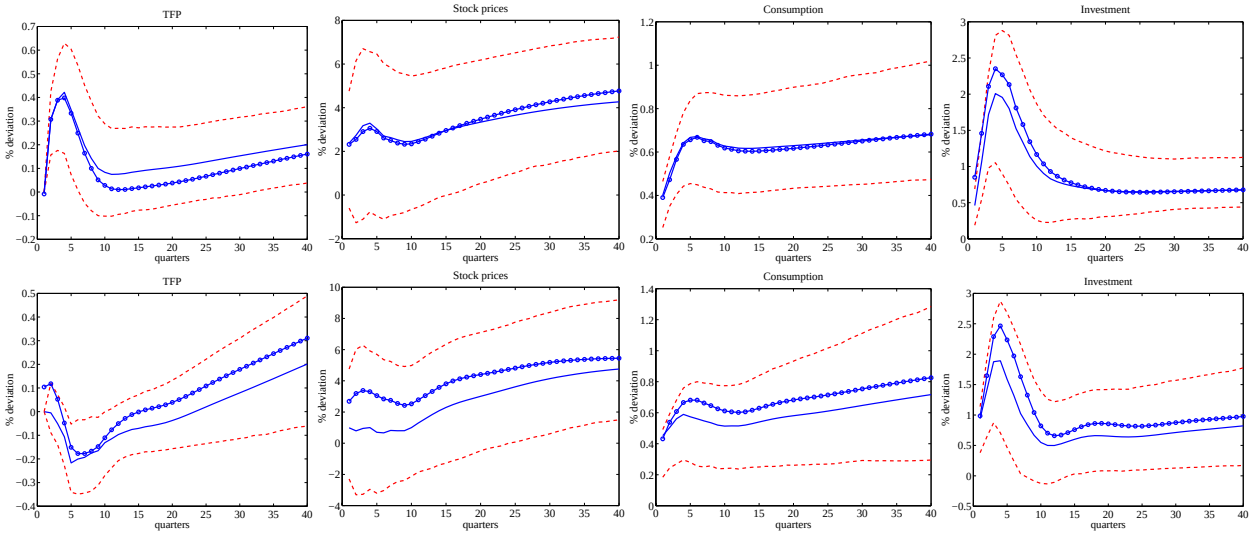

On each panel of this figure, the bold line represents the point estimate of the responses to a unit ϵ_2 shock (the shock that does not have instantaneous impact on TFP in the short run identification); the line with circles represents the point estimate of the responses to a unit $\widetilde{\epsilon}_1$ shock (the shock that has a permanent impact on TFP in the long run identification). Both identifications are done in the baseline bivariate specification. The unit of the vertical axis is percentage deviation from the situation without shock. Dotted lines represent the 10% and 90% quantiles of the distribution of the IRF in the case of the short run identification, this distribution being the bayesian simulated distribution obtained by Monte-Carlo integration with 2500 replications, using the approach for just-identified systems discussed in Doan [1992].

Figure 21: Impulse Responses to ϵ_2 and $\widetilde{\epsilon}_1$ in the in the $(TFP, SP, C+I)$ VAR, without (upper panels) or with (lower panels) Adjusting TFP for Capacity Utilization

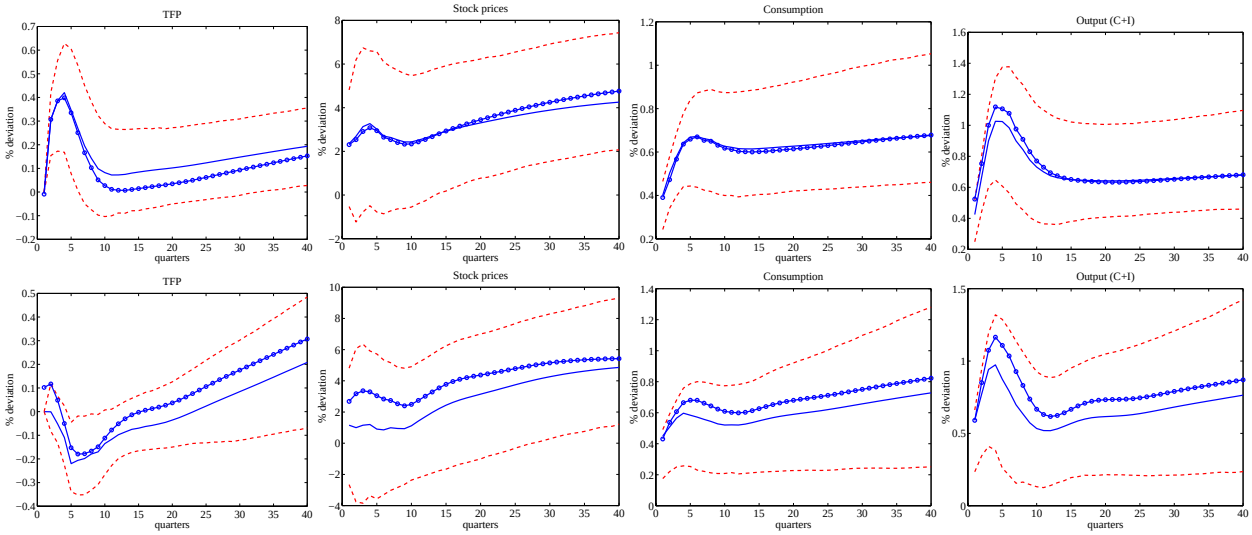

On each panel of this figure, the bold line represents the point estimate of the responses to a unit ϵ_2 shock (the shock that does not have instantaneous impact on TFP in the short run identification); the line with circles represents the point estimate of the responses to a unit $\widetilde{\epsilon}_1$ shock (the shock that has a permanent impact on TFP in the long run identification). Both identifications are done in the baseline bivariate specification. The unit of the vertical axis is percentage deviation from the situation without shock. Dotted lines represent the 10% and 90% quantiles of the distribution of the IRF in the case of the short run identification, this distribution being the bayesian simulated distribution obtained by Monte-Carlo integration with 2500 replications, using the approach for just-identified systems discussed in Doan [1992].