Stock Prices, News and Economic Fluctuations

Paul Beaudry* and Franck Portier‡

January 2003

Abstract

A common view in macroeconomics is that business cycles can be meaningfully decomposed into fluctuations driven by demand shocks – which are shocks that have no short or long run effects on productivity – and fluctuations driven by unexpected changes in technology. In this paper we propose a means of evaluating this view and we show that it is strongly at odds with the data. In contrast, we show that the data favors a view of business cycles driven primarily by a shock that does not affect productivity in the short run —therefore it looks like a demand shock – but affects productivity in the long run. The structural interpretation we suggest for this shock is that it represents news about future technological opportunities. We show that this shock explains about 50% of business cycle fluctuations and therefore deserves to be acknowledged and further understood by macroeconomists.

Key Words : Business Cycle – News – Productivity Shocks

JEL Classification : E3

*CRC University of British Columbia and NBER.

†Université de Toulouse (GREMAQ, IDEI, LEERNA, Institut Universitaire de France and CEPR)

‡The authors thank Susanto Basu , Larry Christiano, Roger Farmer and participants at CEPR ESSIM 2002, NBER Summer Institute 2002, University of Berlin, Université du Québec à Montréal, Université de Toulouse for helpful comments.

Stock Prices, News and Economic Fluctuations

January 2003

Abstract

A common view in macroeconomics is that business cycles can be meaningfully decomposed into fluctuations driven by demand shocks – which are shocks that have no short or long run effects on productivity – and fluctuations driven by unexpected changes in technology. In this paper we propose a means of evaluating this view and we show that it is strongly at odds with the data. In contrast, we show that the data favors a view of business cycles driven primarily by a shock that does not affect productivity in the short run —therefore it looks like a demand shock – but affects productivity in the long run. The structural interpretation we suggest for this shock is that it represents news about future technological opportunities. We show that this shock explains about 50% of business cycle fluctuations and therefore deserves to be acknowledged and further understood by macroeconomists.

Key Words : Business Cycle – News – Productivity Shocks

JEL Classification : E3

1 Introduction

What drives business cycle fluctuations? Much of the modern debate on this issues revolves around whether business cycle fluctuations are driven primarily by demand shocks (such as monetary shocks, which have no or little long run effects) or by technology shocks (i.e, unexpected change in technological opportunities). In contrast, several empirical papers, such as Cochrane [1994], question the relevance of either of these views by suggesting that business cycle fluctuations may be driven by shocks that look like neither of the two traditionally studied shocks. In this paper, we further pursue this line of reasoning by presenting new evidence suggesting that business cycle fluctuations may be primarily (or at least largely) driven by a shock which is neither a traditional demand or technology shock, but is instead a type of hybrid which admits a simple structural interpretation as a news shock. 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 different 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 measured total factor productivity (TFP) and an index of stock market value (SP). Our interest in focusing on stock market information in our analysis is motivated by the view that stock prices are likely to be a good variable for capturing any changes in agents expectations about future economic growth. An important empirical result of the paper is to show that the innovation in stock prices which is contemporaneously orthogonal to TFP is actually extremely correlated with the shock that explains long run movements in TFP.

In order to interpret the result from our empirical exercise, we use simple models in the New-Keynesian and RBC tradition to highlight the predictions of these models along a particular dimension. We also present a simple model with news shocks as an alternative way of interpreting the data. The main claim of the paper is that the data on TFP and stock market value have properties that run counter to the demand and supply type dichotomy inherent to most New-Keynesian and RBC models. In contrast, we argue that the observed pattern is easily understood as the result of news shocks, that is, innovations in agents expectations of future technological opportunities that arise before these opportunities are actually productive in the market.

The two reduced form disturbances we isolate are first, one which represents innovations in stock prices which are orthogonal to innovations in TFP and second, one that drives long run movements in TFP. The intriguing observation is that these two disturbances– when isolate separately without imposing orthogonality (as is implicit in many other approaches)– are found to be almost perfectly co-linear and induce the same dynamics. Moreover, we will show that this particular 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. Since this shock is initially captured in forward looking variables like stock prices and not TFP, it may be reasonably interpreted as a change in expectations. And, since this shock explains long run TFP, we think of this shock as likely reflecting innovations to agents' information set (news) regarding future technological opportunities.

However, alternative structural interpretations of this shock are possible. For example, this shock may alternatively reflect a sunspot shock in a model with an implementation coordination problem. Regardless of the precise interpretation of this shock, the main claim of this paper is that the data on stock prices and TFP suggest that macroeconomic research may need to orient effort away from models which emphasize monetary shocks and unexpected changes in technological opportunities as primitive sources of fluctuations, and instead direct effort towards models that explain fluctuations in the short run as the result of expectational changes that in the long run are supported by improvements in productivity.

Our goal in presentation is to introduce a extremely simple empirical strategy that can be used to evaluate the relevance of many commonly used macroeconomic models, as well as help identify alternative classes of models which may better match the properties of the data. In addition, we will argue that the shock series we identify through this methodology is not just an empirical curiosity, but instead is a shock with easy economic interpretation as either a news shock or as a some type of coordination shock. To this end, we present our 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, output and hours worked. We begin by considering the bi-variate system for TFP and stock prices since it offers the most straightforward way of evaluating the relevance of many standard macroeconomic models. 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 explicitly allows for the possibility of two standard shocks which can have the interpretation as a standard demand and standard technology shock. Finally, we also report results based on a set a four-variable system in order to further document the robustness of our results.

2 Using impact and long-run restriction 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 these techniques sequentially as a means of evaluating different classes of economic models. In particular, we will want to use this technique to (1) evaluate the appropriateness of models which divide business cycle fluctuations into orthogonal demand driven versus supply driven components, and (2) to evaluate the relevance of an alternative view whereby fluctuations are driven by expected changes in future technological opportunities. To this end, let us begin in the simplest case of a bi-variate system. The bivariate system we want to consider is one composed of measured total factor productivity, denoted TFP_t, and a forward looking economic decision variable X_t. The only characteristic of X_t that is important for our argument is that it be an unhindered jump variable, that is, a variable that can immediately react to changes in information without lag. One macroeconomic variable that nicely fits this requirement is a stock price index. For this reason, in our empirical analysis, we will use a stock price index as our preferred measure of X, although for now we do not need to be so precise about the identify of X.

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

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

where L is the lag operator, $C(L) = I + \sum_{i=1}^{N} C_i L^i$, and 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.

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 alternative representation and one that imposes a long run restriction. In order to see this most clearly, let us denote these two alternative representations by:

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

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

where $\Gamma(L) = \sum_{i=0}^{\infty} \Gamma_i L^i$, $\tilde{\Gamma}(L) = \sum_{i=o}^{\infty} \tilde{\Gamma}_i L^i$ and the variance covariance matrices of ϵ and $\tilde{\epsilon}$ 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:

$$\Gamma_0 \Gamma_0' = \Omega$$

and for $i > 0$

$$\Gamma_i = C_i \Gamma_0$$

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 $\tilde{\Gamma}(1) = \sum_{i=0}^{\infty} \tilde{\Gamma}_i$ equals zero, that is, we choose an orthogonalization where the disturbance $\tilde{\epsilon}_2$ 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. For example, a particular theory may imply that the correlation between the resulting errors ϵ_2 and $\tilde{\epsilon}_1$ 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.

We will now present a simple New-Keynesian model and a simple RBC model to illustrate the predictions of such models with regards to the theoretical co-variance properties of ϵ_2 and $\tilde{\epsilon}_1$. In both cases, the supply disturbances in the model will be represented by exogenous and permanent changes in technological opportunities. In the New-Keynesian model, the second source of shock will be a monetary disturbance. In the RBC model the second source of fluctuations will be a temporary shock to preferences. The main implication of these two types of models we will highlight is that they imply that business cycle fluctuations can be decompose into structurally meaningful (orthogonal) supply driven and demand driven components. In this sense, their bi-variate structural moving average representation is very similar in that, in either case, the non technological disturbance – which we will refer to as the demand disturbance – should be contemporaneously orthogonal to innovations in TFP and should not cause long run movements in TFP. Hence, these two classes of models suggest that the two orthogonalized representations of the data discussed above should be quasi-identical. In particular, it suggests that ϵ_2, which under these theories can be referred to as a demand disturbance, should be orthogonal to $\tilde{\epsilon}_1$, which can be refereed to as a supply shock. Therefore, looking whether this type of pattern is found in the data provides a means of evaluating the relevance of this class of models, that is, models that dichotomize fluctuations between orthogonal supply versus demand disturbances.

2.1 Two simple and canonical models

Here we illustrate the implications of using impact and long-run restrictions in two canonical New-Keynesian and RBC models. Later, we will also present an example of a model where agents receive advanced news about future technological opportunities. We will show that the two canonical models deliver similar predictions about ϵ and $\widetilde{\epsilon}$, while the news delivers different results. As we want to derive simple and explicit results, the models do not aim at realism as many assumptions are made in order to allow explicit 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) (omitting the 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 $\epsilon_2 \perp \widetilde{\epsilon}_1$.

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 "demand" shock does not affect TFP in the short run nor in the long run. Therefore, ϵ_2 is orthogonal to $\widetilde{\epsilon}_1$.

A simple RBC model with technology and preference shocks: Let us now consider an economy in which preferences of the representative household are given by

$$U = E_0 \sum_{t=0)}^{\infty} \beta^t \left[\log C_t - \Lambda_t \frac{L_t^\sigma}{\sigma} \right] \tag{6}$$

where C is consumption, L labor and Λ a stationary preference shock.

$$\Lambda_t = e^{\eta_{2,t}} \tag{7}$$

This preference shock acts here as a “demand” shock. A government spending shock would be a more natural candidate for a demand shock, but the present formulation has the advantage of analytical tractability, and for our purpose, is equivalent to a government spending shock. The household accumulates capital, and we assume full depreciation, so that

$$K_{t+1} = I_t \tag{8}$$

where K is capital and I investment. The budget constraint of the household, that rents capital and labor services to the representative firm, is given by

$$C_t + I_t = w_t L_t + \kappa_t I_{t-1} \tag{9}$$

where κ is the rental rate of capital services and w the wage rate.

The representative firm in this economy produces according to the CRS technology

$$Y_t = \theta_t K_t^{\gamma} L_t^{1-\gamma} \tag{10}$$

where θ is again a random walk technology shock.

$$\theta_t = \theta_{t-1}\ e^{\eta_{1,t}} \tag{11}$$

$\eta_{1,t}$ and $\eta_{2,t}$ are assumed to be *iid* processes with identity covariance matrix and zero mean.

We assume that agents behave competitively, maximize utility or profit at given prices and that markets clear. In such an economy, as shown in the appendix, the solution is log-linear. With this solution, one can perform the short-run and long-run orthogonalizations we presented above, and recover the shocks ϵ and $\widetilde{\epsilon}$ as functions of the structural shocks $\eta_{1,t}$ and $\eta_{2,t}$. Since firms make zero profits every period, the stock market value of firms is uninteresting in this model, but there are still asset price fluctuations in the bond market. Hence, here we will focus on the joint behavior of TFP and the bond price as the system of interest, that is, the bond price will play the role of the variable X_t introduced in the preceding section.

In this model, the equilibrium joint behavior of TFP and the log bond price (denoted p^b) has a structural moving average given by:

$$\begin{pmatrix} \Delta TFP_t \\ \Delta p_t^b \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ \frac{(1-\gamma)}{1-\gamma L} - 1 & -\frac{(1-L)(1-\gamma)^2}{\sigma(1-\gamma L)} \end{pmatrix} \begin{pmatrix} \eta_{1,t} \\ \eta_{2,t} \end{pmatrix} \tag{12}$$

Performing short-run and long-run identification on this system, we obtain

$$\epsilon_1 = \eta_1 \quad , \quad \epsilon_2 = \eta_2 \quad , \quad \widetilde{\epsilon}_1 = \eta_1 \quad , \quad \widetilde{\epsilon}_2 = \eta_2 \tag{13}$$

In particular, we have $\epsilon_2 \perp \widetilde{\epsilon}_1$.

As in the case with the New-Keynesian model, we observed that the simple RBC model also predict the shock that is orthogonal to current TFP should also be orthogonal to the disturbance which drives long run movements in TFP. We now present an alternative model, which has a very different prediction in terms of the correlation between ϵ_2 and $\tilde{\epsilon}_1$.

2.2 A model with news shocks

Let us now consider a small deviation from the RBC model we just presented. In particular assume that θ_t has both a permanent component – $\bar{\theta}_t$ – and a temporary component – ν_t, and we disregard preference shock. The important additional assumption is that permanent innovation to technology are known to agents 1 period before they actually impact TFP. The process for TFP can therefore be expressed as follows.

$$\begin{aligned} TFP_t &= \bar{\theta}_t + \nu_t \\ \bar{\theta}_{t+1} &= \bar{\theta}_t + \eta_{1,t} \\ \nu_t &= \rho \nu_{t-1} + \eta_{2,t}, \qquad 0 < \rho < 1 \end{aligned}$$

In this model, the structural moving average for TFP and bond prices is given by:

$$\begin{pmatrix} \Delta TFP_t \\ \Delta p_t^b \end{pmatrix} = \begin{pmatrix} L & \frac{\gamma(1-L)}{(1-\rho L)} \\ \frac{(1-\gamma)L}{1-\gamma L} - 1 & \frac{(1-L)(1-\rho(1-\gamma L))}{(1-\rho L)(1-\gamma L)} \end{pmatrix} \begin{pmatrix} \eta_{1,t} \\ \eta_{2,t} \end{pmatrix} \tag{14}$$

Performing short-run and long-run identification on this system, we obtain

$$\epsilon_1 = \eta_2 \quad , \quad \epsilon_2 = \eta_1 \quad , \quad \widetilde{\epsilon}_1 = \eta_1 \quad , \quad \widetilde{\epsilon}_2 = \eta_2 \tag{15}$$

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

Observe that the model dynamics are now different, and that it gives a very contrasted view of fluctuations. The shock that moves the measure of technology in the long-run does not in the short-run, while the shock that moves technology in the short run is transitory.

Discussion : The important aspect of the two first models is that they imply that business cycle fluctuations that can be decomposed into structurally meaningful supply driven and demand driven components. In this sense, their bi-variate structural moving average representation is very similar in that, in either case, the non technological disturbance – which we will refer to as the demand disturbance – should be contemporaneously orthogonal to innovations in TFP and should not cause long run movements in TFP. Hence, these two classes of models suggest that the two orthogonalized representations of the data discussed above should be quasi-identical. As we have shown, ϵ_2, which under these theories can be referred to as a demand disturbance, is orthogonal to $\tilde{\epsilon}_1$, which can be referred to as a supply shock. Moreover, the implied impulses of these shocks are different. Therefore, looking whether this type of pattern is found in the data provides a means of evaluating the relevance of this class of models, that is, models that dichotomize fluctuations between supply versus demand disturbances.

The news 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 and consumption– already incorporate this possibility. [1] If this class of models is relevant, the long run restriction used to derive the orthogonal moving average representation given by $\tilde{\Gamma}_i$ and $\tilde{\epsilon}$ still implies that $\tilde{\epsilon}_1$ 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 result from anticipated TFP shocks then by definition the anticipated shock has zero effect on impact on TFP.

[1] 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.

Hence, under this class of models, our two orthogonal representation discussed above are no longer identical but instead are mirror images. In the news model, ϵ_2 and $\tilde{\epsilon}_1$ are now co-linear (as opposed to being orthogonal in the previous class of models) as they both should capture the effect of anticipated changes in technological opportunities. Moreover, the impulse responses associated with ϵ_2 and $\tilde{\epsilon}_1$ are identical. The main feature of this class of models that we want to emphasize is that it does not validate a decomposition in terms of demand versus supply effect. In particular, in the short run an anticipated technological improvement (a news shock) looks like a demand effect, while in the long run it looks like a supply effect.

In summary, in this section we suggest a simple but potentially powerful way evaluating different classes of 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 $\tilde{\epsilon}_1$ as a means of evaluating whether models which presuppose a dichotomy between supply and demand disturbances fits the data, or if instead the data rejects this characterization in favor of an alternative characterization which emphasizes the importance of news shocks (that is, shocks that are first reflected in forward looking variables like stock prices and only later reflected in changes in technological opportunities).

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 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 (SP), deflated by the seasonally adjusted implicit prices deflator of GDP 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.

The construction of our 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 mispecification 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. Alternatively, one can estimate the VECM with a matrix of cointegration relationships which is of reduced rank and then examine whether the resulting Wold representation is similar to that found by estimating the system in levels. We choose to follow this route. [2] Therefore our main results will be based on a Wold representation achieved by inverting a VECM, but we will also show that the results are robust to 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 Andrew [2000] to test for cointegration. This procedure indicated that co-integration between SP and TFP could not be rejected at the 1% level and therefore we adopted the VECM specification. [3]

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 (not reported) 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 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 Main Results

We estimated a VECM for (TFP, SP) with one cointegrating relation and recovered two orthogonalized shock series corresponding to the ϵ and $\tilde{\epsilon}$ discussed in Section 2, that is, ϵ was recover by imposing

[2]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.

[3] 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.

an impact restriction (a restriction on Γ_0) and $\tilde{\epsilon}$ was recovered by imposing a long run restriction. The level impulse responses on (TFP, SP) associated with the ϵ_2 shock and the $\tilde{\epsilon}_1$ 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 $\tilde{\epsilon}_1$ shock –which by construction has a permanent effect on TFP– has statistically no impact effect on TFP (if anything, the impact effect is negative) but has a significant effect on SP. On the one hand, this suggests that ϵ_2 contains information about future TFP growth which is instantaneously and positively reflected in stock prices. While on the other hand, it suggests that all 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 $\tilde{\epsilon}_1$ 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 from the quasi-identity between the ϵ_2 shock and the $\tilde{\epsilon}_1$ shock, as shown in figure 4 which simply plots $\epsilon_{2,t}$ against $\tilde{\epsilon}_{1,t}$. In effect, the correlation coefficient between these two series is 0.9717, 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 2.1, this observation runs counter to simple RBC and New-Keynesian models since these models imply a dichotomy between supply and demand shocks which should result in ϵ_2 being orthogonal to $\tilde{\epsilon}_1$. In contrast, this pattern appears consistent with the view –which we call the news view– that technological change is generally anticipated by market participants as these agents may receive information about future technological opportunities before they are actually available in production. However, before explore this news interpretation further, we want to illustrate the robustness of the observation that ϵ_2 and $\tilde{\epsilon}_1$ 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 $\tilde{\epsilon}_1$ 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 $\tilde{\epsilon}_1$, 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. [4] In the left panel of Figure 7, we plot ϵ_2 against $\tilde{\epsilon}_1$ 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 which indicates the robustness of this observation to estimating the system in VECM form or in level form. In Figure 6 we now superimpose the impulse responses associated with ϵ_2 and $\tilde{\epsilon}_1$ for the case where the Wold representation was obtained by estimating the system (TFP, SP) in VECM form with only two lags versus with five lags. Furthermore, in the right side panel of Figure 7 we plot ϵ_2 against $\tilde{\epsilon}_1$ 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.[5] In particular, the high correlation between ϵ_2 and $\tilde{\epsilon}_1$ observed in Figure 7 appears robust and therefore remains counter to the predictions of most RBC and New-Keynesian models, while this pattern may be easily understood as the result of agents having news about future technological opportunities .

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

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

4.2 Exploring the News Interpretation Further

The observation that our estimates of ϵ_2 and $\tilde{\epsilon}_1$ 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: (1) How does the economy respond to such a shock, that is, do the responses to ϵ_2 (or $\tilde{\epsilon}_1$) look like standard business cycle fluctuations in the sense of generating positive co-movements in consumption, investment and hours worked? (2) Could this type of disturbance be a major source of fluctuations? To answer these questions, we will begin by exploiting the co-movements between different variables (i.e. consumption, investment and hours worked) and the ϵ shock series derived from our baseline specification for (TFP, SP). Later we look directly at larger systems which incorporate these other variables explicitly. We only need to focus here on the ϵ shocks since, as we have shown, they are essentially mirror images of the $\tilde{\epsilon}$ shocks. Our first approach to this issue will therefore be to estimate the following truncated moving average representation for different variables Z_t:

$$\Delta Z_t = \sum_{j=0}^{J} \phi_j^1 \epsilon_{1,t-j} + \sum_{j=0}^{J} \phi_j^2 \epsilon_{2,t-j} + \nu_t \tag{16}$$

where Z will either be consumption (C), investment (I), output (defined as $C + I$) or worked hours (H), and where ν a variable-specific disturbance that is orthogonal to ϵ_2 (the news or the permanent shock) and to ϵ_1 (the other shock). 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 $J = 40$. We also compute the contribution of ϵ_1 and ϵ_2 to the forecast error variance of X. For example, we compute the fraction of variance of X at horizon h due to ϵ_2 as

$$\frac{\sum_{n=0}^{h} (\sum_{j=0}^{n} \phi_j^2)^2}{\sum_{n=0}^{h} (\sum_{j=0}^{n} \phi_j^2)^2 + \sum_{n=0}^{h} (\sum_{j=0}^{n} \phi_j^2)^2}$$

Note that such a share represents the share of the variance attributed to ϵ_2 relative to the sum of ϵ_1 and ϵ_2. Given that some of the variance of X is left unexplained by these two structural (since the residual ν is not identically null), it is useful to note that in the regressions (16), we obtain a R-square of .68 for consumption, .70 for investment, .72 for output (consumption plus investment) and .67 for hours.

Figure 8 displays the responses of consumption, investment, output (measured as consumption plus investment) and hours to ϵ_2, that is, the responses to what we are interpreting as a news shock. 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. Worked hours respond relatively little on impact, but also increase substantially increase after on quarter. Hence, it appears that this news shock induces business cycle type phenomena.

Figure 9 plots the variance decomposition of output and hours associated with ϵ_1 and ϵ_2. As approximated by this simple projection methodology, we can see that up to 45% of output variance due to the epsilons at a horizon of 8 to 12 quarters is attributable to ϵ_2 (the news shock) and not less than 50% of hours. These results suggest that a ϵ_2 (1) creates business cycle like fluctuations, (2) may account for a non negligible share of aggregate fluctuations, (3) does not affect TFP contemporaneously and (4) 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 $\tilde{\epsilon}_1$ which by definition affects long run TFP.

4.3 Relationship with previous work

In a highly influential paper, Gali [1999] examined the response associated with permanent changes in labor productivity ($\frac{Y}{H}$) and interpreted the responses as resulting from unexpected changes in technology. [6] Since we have been working with TFP instead of labor productivity, it is of potential interest to examine how our results change if we consider the system $(\frac{Y}{H}, SP)$ instead of (TFP, SP). In

[6] See also Francis and Ramey [2002]

particular, our results to date suggest that the shock isolated by Gali may not be a typical technological shock emphasized in the RBC literature, since we have shown that permanent shocks to TFP are highly correlated with shocks that are contemporaneously orthogonal to TFP. In Figures 10, 11 and 12, we report the analogue of Figure 2, 3 and 4 for the case where we focus on orthogonalizing the Wold representation for $(\frac{Y}{H}, SP)$ using both our impact restriction and our long run restriction. The only big difference between Figures 9, 10, and 11 relative to Figures 2,3, and 4 is in the bottom left panel of Figure 9. In this case, we see that a permanent change $\frac{Y}{H}$ is associated with an increase in $\frac{Y}{H}$ on impact. This is not surprising since we previously saw that a permanent change in TFP is associated with an initial period where C and I increase but where H and TFP remains essentially unchanged. Hence, it must be the case that some unmeasured aspect of capacity utilization is increasing on impact, which explains the observed increase in $\frac{Y}{H}$. Once again, in Figure 11, we can see that shocks that are contemporaneously orthogonal to $\frac{Y}{H}$ are highly correlated with shocks that affect long run movements in $\frac{Y}{H}$ (moreover, the correlation between the two innovations $\tilde{\epsilon}_1$, one estimated with TFP and one with Y/H, is .95). Adopting a short run identification shows that the ϵ_2 shock again captures almost all the long run dynamics of Y/H, which furthermore provide credence to the view that this shock may reflect news about future productivity growth, rather than being a typical technology shock. This suggests that much of the debate around the effects noted by Gali may be misguided.

4.4 Summary

We modeled the joint behavior of TFP and Stock Prices, and identify the permanent innovation to TFP. Surprisingly, this shock does not move TFP on impact, while other macroeconomic aggregates (consumption, investment, hours) respond positively to it on impact. This shock is observed to be collinear to a shock obtained by assuming zero impact on TFP. An interpretation for this shock is that it reflects news about permanent level of TFP, before this new level is effectively reached. This view that a third shock (which is neither "demand" nor "supply") may be an important source of short run fluctuations needs to be explored further. To do so, we study now larger dimension systems, in which we explicitly allow for the presence of traditional demand and supply shocks.

5 Higher Dimension Systems

In this section, we study larger dimension systems in which –in addition to TFP and SP– consumption, investment, output and hours are alternatively or jointly introduced. For each system, we will show results that echo the ones of the bivariate case. Namely, the shock that is orthogonal to TFP, ϵ_2, and the shock that causes permanent changes in TFP, $\tilde{\epsilon}_1$ are extremely highly correlated if not identical. Second, impulse responses associated with this shock show an aggregate expansion both in the short run and in the long run. Thirdly, 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 those robustness exercises are not reported here.

5.1 A (TFP, SP, C) system

Our approach here is identical to 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 macroeconomic models. In particular, if these shocks are highly correlated, we claim that it provides evidence against many of the popular macroeconomic models which embed a clear dichotomy between demand driven and supply driven fluctuations. In contrast, we argue that a high correlation provides evidence in favor of the view that news about future changes in technological opportunities 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[7] 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 ($\sum_{i=0}^{\infty} \tilde{\Gamma}_i(1)$) are equal to zero and (2) recuperating the shock $\tilde{\epsilon}_1$.[8] 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 the other two shocks to potentially represent a standard demand shock and a standard supply shock, we impose no restrictions related to the shock $\tilde{\epsilon}_1$ as to let it potentially represent an unanticipated technology shock. As for the shock $\tilde{\epsilon}_3$, we impose that it have no long run effect on either TFP or Consumption, and therefore it could potentially capture traditional demand effects. One of the advantages of the three variable system over the two variable system is that we can now explicitly consider the possibility of two fundamental shocks, like a standard demand and supply shock, and a measurement error in TFP that is potentially correlated with these shocks (due for example to unmeasured variable rates of utilization). The important aspect to note about this case is that, if the data is generated by the type of New-Keynesian models or RBC models we illustrated in the previous section, we still would not expect ϵ_2 and $\tilde{\epsilon}_1$ to be highly correlated since ϵ_1 should capture the technology shock and both ϵ_2 and ϵ_3 should represent temporary disturbances.

The impulse responses associated with the shocks ϵ_2 and $\tilde{\epsilon}_1$ are presented in Figure 13, the associated variance decompositions are presented in Figure 14, and the plot of ϵ_2 versus $\tilde{\epsilon}_1$ is presented in Figure 15. As is clear from Figure 15, ϵ_2 and $\tilde{\epsilon}_1$ are highly correlated (correlation of almost 1).

[7] 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.

[8] 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 know, this additional restriction is only needed to separate the shocks $\square_3$, $\square_3$ and does not influence $\square_1$.

Moreover, Figure 14 indicates that these shocks explain most of the short run variance of consumption and most of the long run variance of TFP. Figure 14 shows that in the long run, the ϵ_1 explains 60% of TFP variance and 50% of the variance of stock prices. Those proportions are very similar for $\tilde{\epsilon}_1$ (in the very long run, it is 100% by construction). Note that the ϵ_2 and $\tilde{\epsilon}_1$ explain more than 80% of the short run variability of consumption.

Finally, in Figure 13 we see that the impulse response associated with these shocks are again very similar. Although these impulse responses are imprecisely estimated, it is worth noting that the point estimate of the impact effect of $\tilde{\epsilon}_1$ on TFP (bottom left panel) is almost exactly equal to zero even though this shock was recuperated without imposing any impact restrictions. The one new observation in Figure 13 relative to Figure 2 is the fact that ϵ_2 and $\tilde{\epsilon}_1$ appear to cause hump shaped responses in TFP, consumption and stock prices in the first 2 to 3 years. The fact that TFP displays this hump shape behavior, suggests that it is not a perfectly clean measure of technology. In effect, these results suggest that a permanent change in TFP may be associated with an up to 10 quarters long period where there may be no actual change in technological opportunities (since the response returns to zero at around 10 quarters), but where there is a measured change in TFP with reflects endogenous utilization associated with a temporary boom. Viewed in this manner, the results in Figures 13 to 15 appear to provide even stronger evidence against the class of models driven by a combination of unanticipated technological change and a demand shock. In contrast, the patterns displayed in these figures still appear consistent with the view that agents may receive news of future technological opportunities well in advance of these improvement actually materializing, and that this type of news may cause a period where expectation cause a demand driven boom.

Using the same technique than for the bivariate VAR, we also compute output and hours response and variance decomposition associated with ϵ_2. Figures 16 and 17 show that ϵ_2 (results are similar for $\tilde{\epsilon}_1$) may be a key impulse in the business cycle. In effect, over 50% of the short run variance of output (C+I) and hours that is related to the epsilons are explained by this shock, and those two variables respond positively and persistently to ϵ_2.

For completeness, in Figure 18, we present the response of the economy to the shocks ϵ_1 and ϵ_3 which we had constructed such that they could potentially represent respectively an unanticipated technology shock and a standard demand shock. It turns out that ϵ_1 looks like a temporary TFP shock: TFP, stock prices and consumption increase on impact, but they all go back to zero quite rapidly. According to this result, it seems that permanent shock to TFP are all orthogonal to current TFP, as they do not show up in this shock. Let us stress the fact that the shock ϵ_1 is not restrained to have no long run impact on TFP. The effects of the potential "demand shock" ϵ_3 (lower panels of figure 18) displays a familiar pattern: consumption and TFP increase in the short run, but not in the long run. We take this evidence as suggesting that the data is not very supportive of a role for unanticipated permanent increase in TFP as important driving forces in macroeconomic fluctuations, but it does support the view that news shocks may be important.

5.2 Four variables system

We now extend our analysis to a four-variables system which includes– in addition to TFP, stock prices and consumption– investment, output (C+I) or hours worked. Our objective is again to recuperate from one representation a shock (denoted ϵ_2) that is orthogonal to TFP and to recuperate from another representation a shock (denoted $\tilde{\epsilon}_1$) that is associated with permanent movements in TFP. The $\tilde{\epsilon}_1$ shock can be isolated by imposing that the long run matrix $\tilde{\Gamma}(1)$ 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 technology 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 that the first, second and third element of the third column of the long run matrix be zero, as to potentially allow ϵ_3 to capture tradition demand shocks and (4) we impose that ϵ_4 is the fourth variable specific shock, i.e, that there are zeros in the first three element of the last column of the impact matrix. The results in this section are based on estimating and inverting VECM representations where we allow for three cointegrating relationships[9]. The only exception is the case where hours worked are included in the system in which case the system is estimated in levels.

In Figure 19, we plot the correlation between ϵ_2 and $\tilde{\epsilon}_1$ for three cases corresponding to different fourth variable in the system. The first panel corresponds to the case where investment is the fourth variable. The second and third panel correspond to the cases where output (C+I) and hours worked are respectively the fourth variable. As before, the correlation between ϵ_2 and $\tilde{\epsilon}_1$ is extremely high, suggesting that whatever drives permanent changes in TFP affect behavior before it affect TFP. Figure 20 reports the variance decompositions for $C + I$ and hours for all four shock in ϵ. The striking aspect of the figure is the extent to which the shock ϵ_2 explains most of the variance. Given our preferred interpretation of this shock as a news shock, this figures clearly indicates that news shocks may be the most important source of business cycle fluctuations.

Finally, in Figure 21 we report the impulse responses associated with ϵ_2 and $\tilde{\epsilon}_1$ for the case where hours worked is the fourth variable included in the system. We only report this case since the impulse responses for the other cases are very similar. What is noticeable about these impulse responses is once again the rich dynamics over the first 2 to 3 years. During this period, the economy appears to go through an important temporary boom, followed by a period of substantial TFP growth. Given a news interpretation of this shock, this temporary boom period can be interpreted as a period of time where agents in the economy try to best position themselves to take advantage of future technological change. Obviously, for this interpretation to make sense, we must accept the possibility that our measure of TFP be a measure of technological opportunities that is contaminated by in the short run due to incomplete account of variable rates of utilization. Also, as we mentioned in the introduction, this type of shock cannot be easily classified as either a demand or a supply shock, since in the short run it may reflect only the effects of anticipations and therefore be classified as a demand shock, while in the long run it clearly looks like a supply shock.

[9]The Nyblom and Harvey test does not reject three cointegrating relationships in favor of two at the 1% level. Since we once again 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.

From this analysis, we conclude that although permanent shocks to productivity (measure as TFP) are likely an important source of short run fluctuations, they do not contribute to current movements of productivity. A possible interpretation is that this shock is indeed a news one, that brings some information about the long run level of productivity, but that is not yet implemented in the production process. [10]

6 Conclusion

Our approach in this paper has been to highlight certain properties of the joint behavior of total factor productivity and stock prices (as well as some other variables) as a means of evaluating different theories of macroeconomic fluctuations. 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, and found this correlation to be very positive and almost equal to 1. We argued that this observed positive correlation runs counter to that predicted by many of the currently popular macroeconomic models — whether they be in the Real Business Cycle tradition or the New-Keynesian tradition – and therefore suggests a need to rethink the appropriateness of these theories. In contrast, we have argued that this type of pattern appear consistent with a view that emphasizes the role of expectations about future technological change (news) as a main driving force behind macroeconomic fluctuations. Accordingly, we believe that this later view deserves more attention.

[10] The existence and prevalence of such a shock in the business cycle has some important theoretical implications. As we show in Beaudry and Portier [2002] , many of the models we use in applied macroeconomic cannot display a aggregate boom following a shock to expectations, unless non convexities or rigid prices are assumed. In that paper, we derive a set of necessary and sufficient conditions for expectationally driven fluctuations to exhibit positive co-movements of consumption, investment and hours. The simple interpretation of these conditions is that some complementarity between input factors in multi-sectoral models (keeping convexity) is needed.

References

Beaudry, P., and F. Portier (2002): "Expextationally driven booms and busts in equilibrium cycle models," mimeo, University of Bristish Columbia and Université de Toulouse.

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

Cochrane, J. (1994): "Shocks," *Carnegie-Rochester Conference Series on Public Policy*, 41, 295–364.

Francis, N., and V. Ramey (2002): "Is the Technology-Driven Real Business Cycle Hypothesis Dead? Shocks and Aggregate Fluctuations Revisited," working paper 8726, NBER.

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

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

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

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

Appendix

A Steps in deriving the structural moving average representations of section 2.1

Throughout this section, unless otherwise indicated, lower case letters will denote the log of variables while the upper case letters will denote levels.

The New-Keynesian Model: 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:

$$\begin{pmatrix} \Delta TFP_t \\ \Delta SP_t \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ 1 & (1-L) \end{pmatrix} \begin{pmatrix} \eta_{1,t} \\ \eta_{2,t} \end{pmatrix} \tag{17}$$

The RBC model: In order to maximize its profit, the firm equalizes marginal products of inputs with their marginal costs

$$(1-\gamma)\frac{Y_t}{L_t} = w_t \tag{18}$$
$$\gamma\frac{Y_t}{K_t} = \kappa_t \tag{19}$$

Denoting χ_t the lagrange multiplier of the households budget constraint, maximization of intertemporal utility leads to the following fist order conditions:

$$\frac{1}{C_t} = \chi_t \tag{20}$$
$$\Lambda_t L_t^{\sigma-1} = \chi_t w_t \tag{21}$$
$$\chi_t = \beta E_t \chi_{t+1}\kappa_{t+1} \tag{22}$$

Combining those equation, we obtain the following recursion:

$$\frac{I_t}{C_t} = \alpha\beta + \alpha\beta E_t \frac{I_{t+1}}{C_{t+1}} \tag{23}$$

Solving forward and imposing the usual transversality condition leads to

$$C_t = (1-\beta\gamma)Y_t \tag{24}$$
$$I_t = \beta\gamma Y_t \tag{25}$$
$$\tag{26}$$

On the other hand, equilibrium labor is obtained from labor demand and intratemporal first order conditions of the household program:

$$L_t = \left(\frac{1-\gamma}{1-\beta\gamma}\right)^{1/\sigma} \Lambda_t^{-1/\sigma} \tag{27}$$

The equilibrium law of motion of consumption can be easily computed and is given by (omitting constant terms):

$$c_t = log(\theta_t) + \gamma c_{t-1} - \frac{(1-\gamma)}{\sigma}\eta_{2,t} \tag{28}$$

Since the price of bonds P^B must satisfy the equation $P_t^B E_t \frac{C_{t+1}}{\beta C_t} = 1$, the structural moving for ΔTFP_t and Δp_t^b is given by the following, as in the text.

$$\begin{pmatrix} \Delta TFP_t \\ \Delta p_t^b \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ \frac{(1-\gamma)}{1-\gamma L} - 1 & -\frac{(1-L)(1-\gamma)^2}{\sigma(1-\gamma L)} \end{pmatrix} \begin{pmatrix} \eta_{1,t} \\ \eta_{2,t} \end{pmatrix} \tag{29}$$

The News model : In the model with news, first order conditions are identical to the ones of the RBC model, the only change being the dating of the technological innovation. As in the RBC model, consumption is given by

$$C_t = \Omega Z_t \Lambda^{-(1-\gamma)/\sigma} C_{t-1}^{\gamma} \tag{30}$$

where Ω is a constant term. We then have in logs, omitting constant terms:

$$c_t = \gamma c_{t-1} + \log(\theta_{t-1}) + \eta_{1,t-1} - \frac{1-\gamma}{\sigma}\eta_{2,t} \tag{31}$$

Since the price of bonds P^B must satisfy the equation $P_t^B E_t \frac{C_{t+1}}{\beta C_t} = 1$, the structural moving for ΔTFP_t and Δp_t^b is approximately given by the following, as in the text.

$$\begin{pmatrix} \Delta TFP_t \\ \Delta p_t^b \end{pmatrix} = \begin{pmatrix} L & \frac{\gamma(1-L)}{(1-\rho L)} \\ \frac{(1-\gamma)L}{1-\gamma L} - 1 & \frac{(1-L)(1-\rho(1-\gamma L))}{(1-\rho L)(1-\gamma L)} \end{pmatrix} \begin{pmatrix} \eta_{1,t} \\ \eta_{2,t} \end{pmatrix} \tag{32}$$

B Main Text Figures

B.1 Figures related to section 3

Figure 1: Data

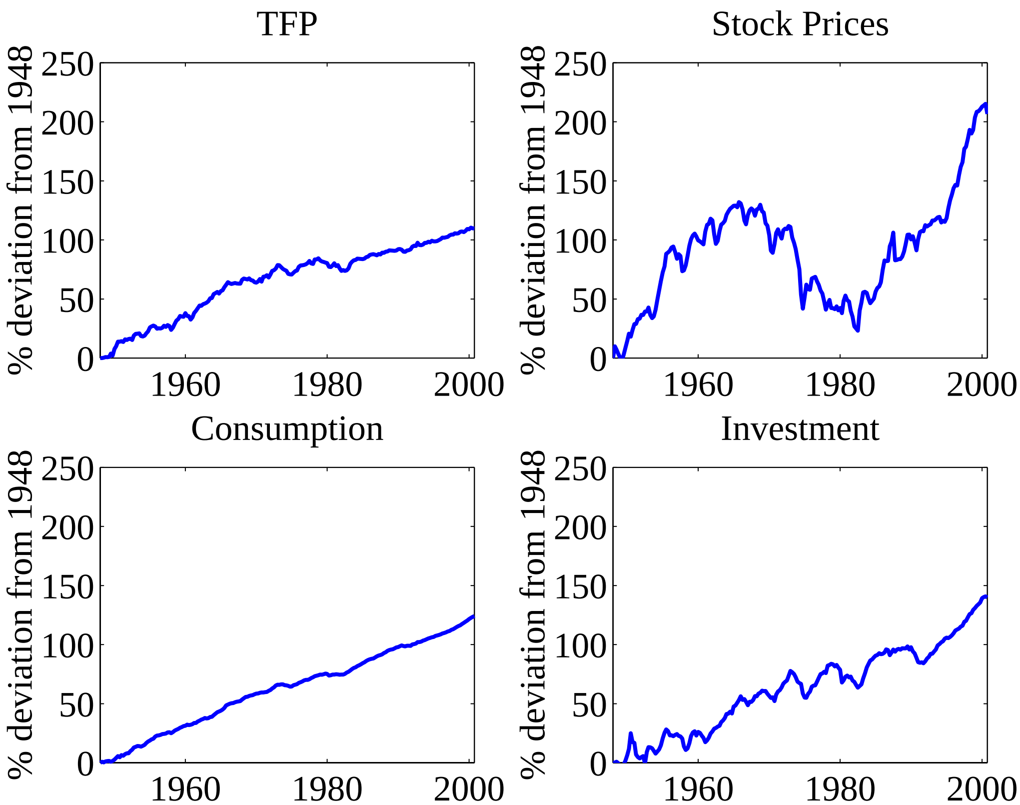

Those series are percentage deviations from 1948:Q1 level. All series have been previously divided by the 15 to 64 years old US. population. TFP is Total Factor Productivity in the non-farm business sector, as computed by the authors, Stock Prices is the Standard & Poors 500 index divided by the GDP 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 ϵ_2 (upper panels) and $\widetilde{\epsilon}_1$ (lower panels) in the (TFP, SP) VAR

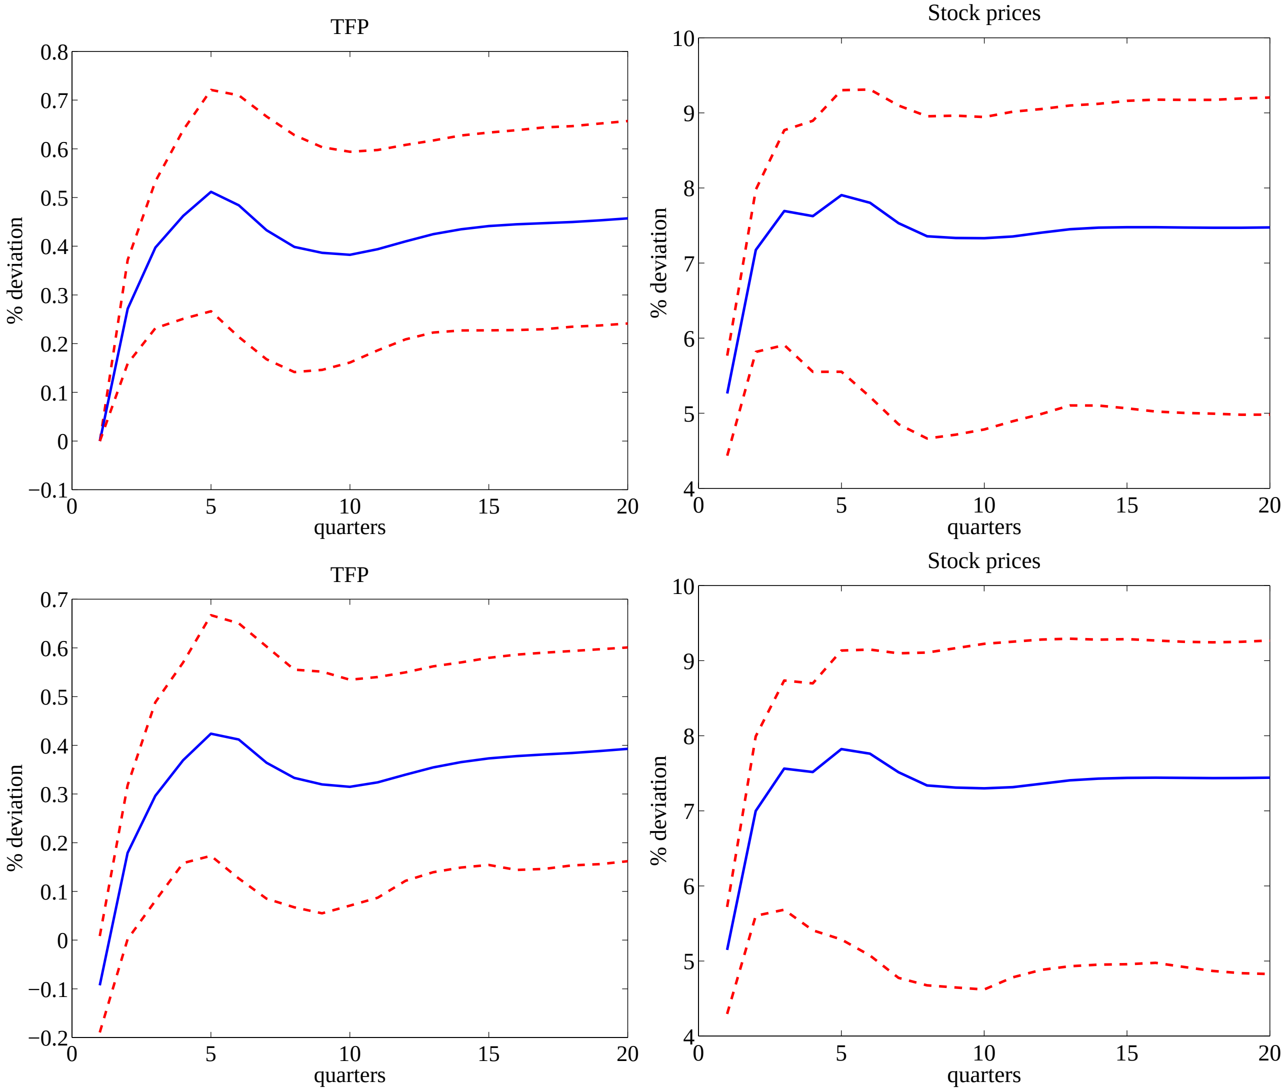

This figure displays the responses of TFP (upper left panel) and stock prices (upper right panel) to a unit ϵ_2 *shock (the shock that does not have instantaneous impact of* TFP *in the short run identification), and the responses of TFP (lower right panel) and stock prices (lower left panel) to a unit* $\widetilde{\epsilon}_1$ *shock (the shock that has a permanent impact on TFP in the long run identification). 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 5% and 95% quantiles of the distribution of the IRF, this distribution being simulated by bootstrapping 1000 times the residuals of the VAR.*

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

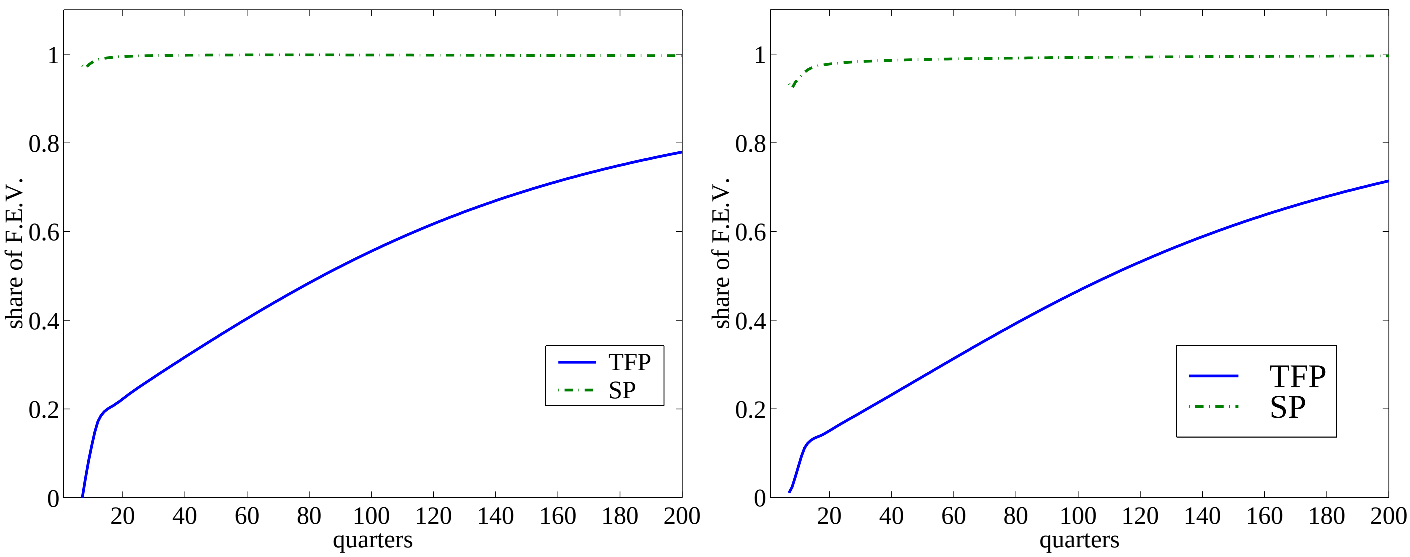

This figure displays the share of TFP and SP forecast error variance attributed to ϵ_2 *(the shock that does not have instantaneous impact of* TFP *in the short run identification) (left panel) or to* $\widetilde{\epsilon}_1$ *(the shock that has a permanent impact on TFP in the long run identification)(right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

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

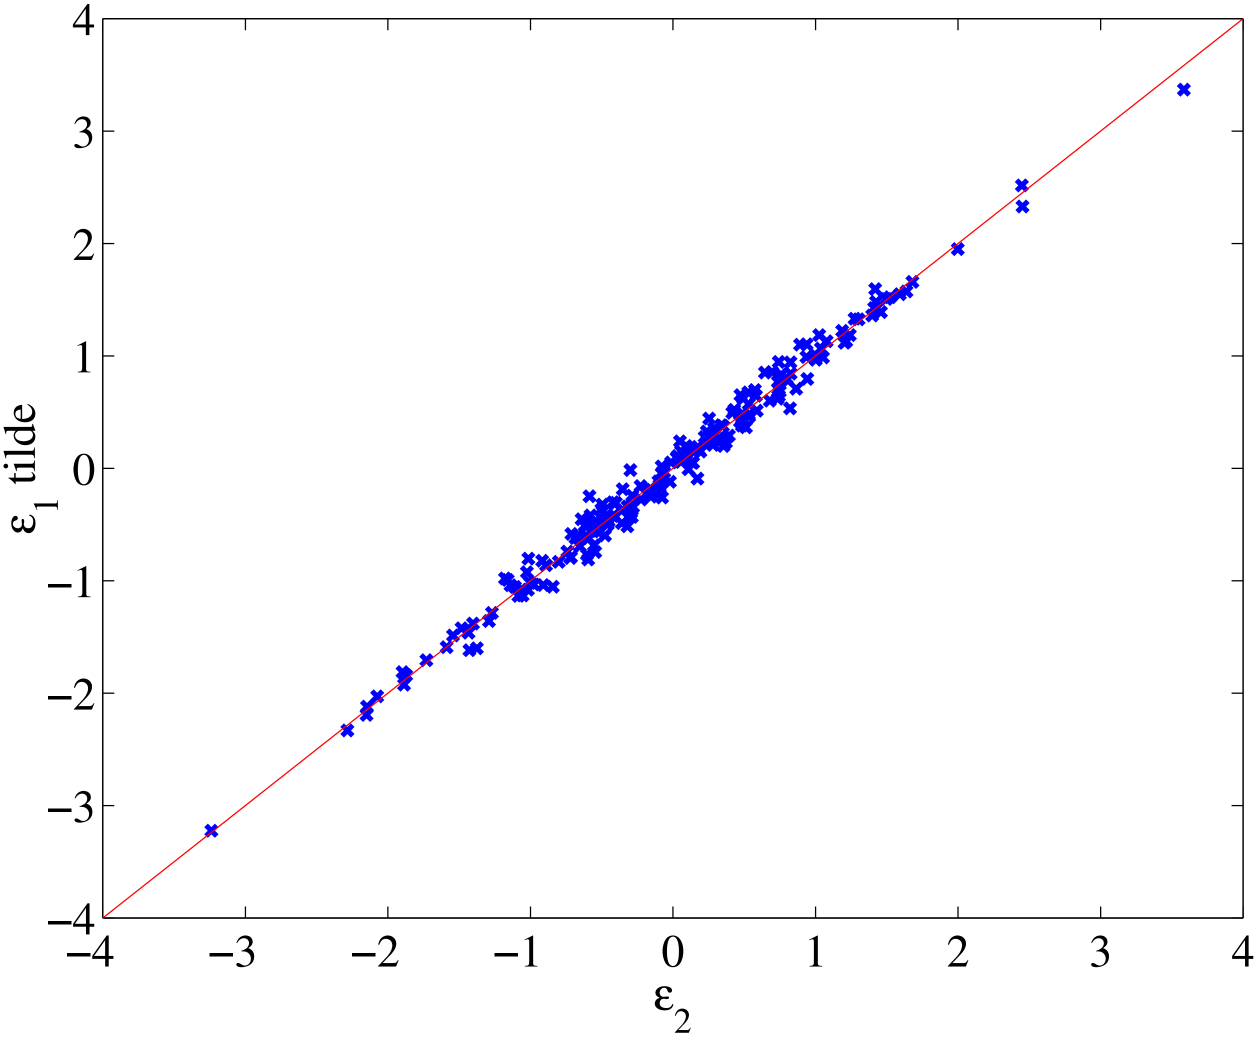

This figure plots ϵ_2 against $\widetilde{\epsilon}_1$. Both shocks are obtained from the baseline (TFP, SP) VAR, with 5 lags and one cointegrating relation. The straight line is the 45° line.

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

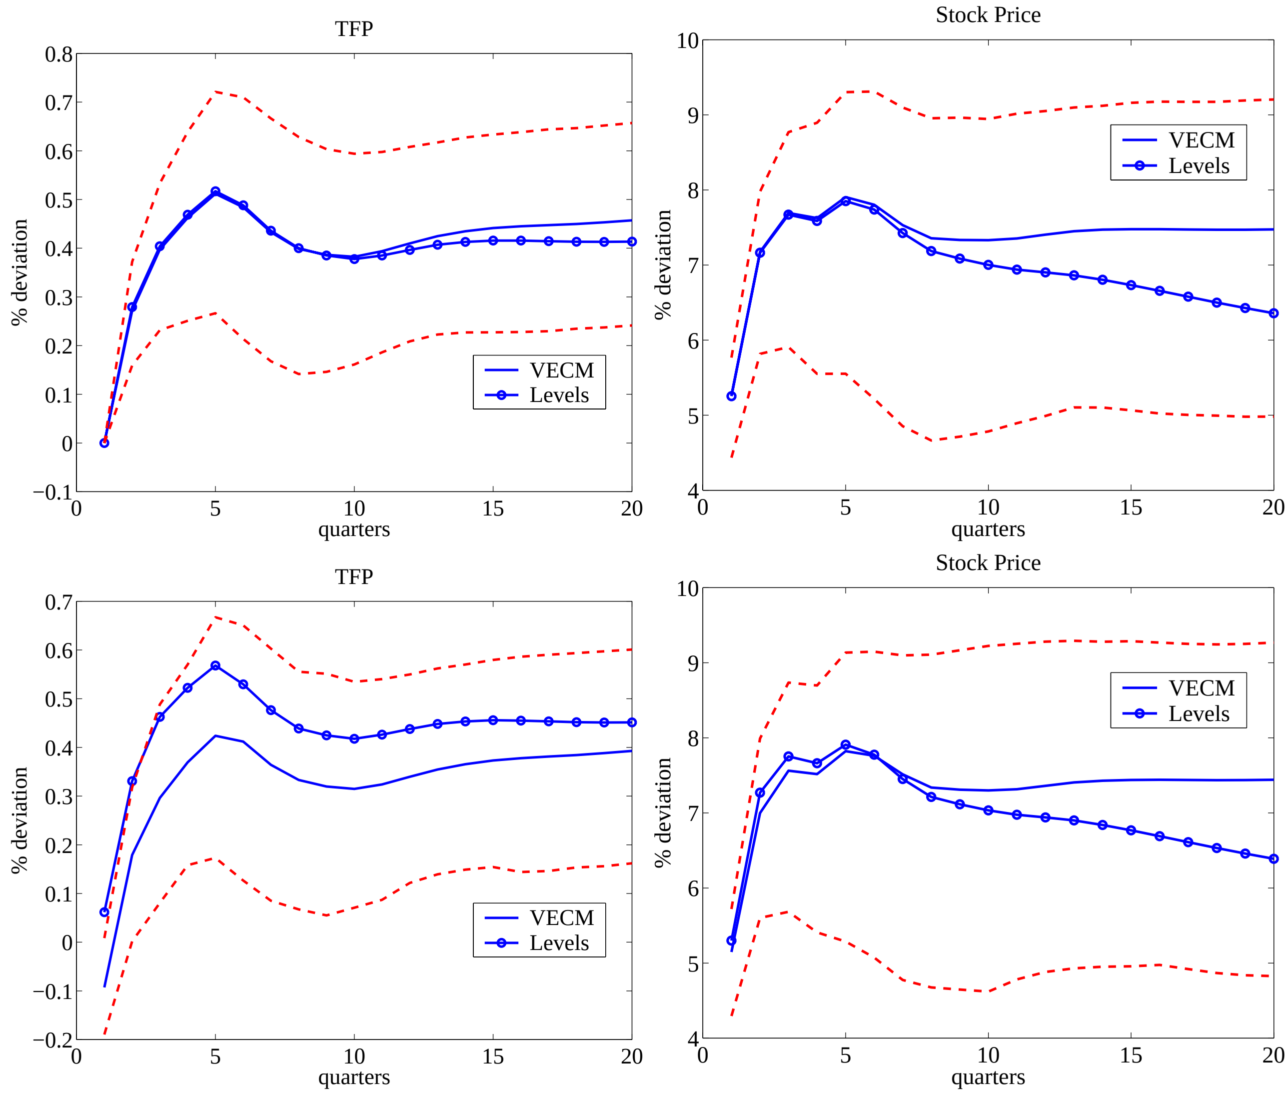

This figure displays the responses of TFP (upper left panel) and stock prices (upper right panel) to a unit ϵ_2 shock (the shock that does not have instantaneous impact of TFP in the short run identification), and the responses of TFP (lower right panel) and stock prices (lower left panel) to a unit $\widetilde{\epsilon}_1$ 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 of TFP and SP in the (TFP, SP) VAR estimated with one cointegrating relation or estimated in levels. The 5% and 95% confidence bands are computed using the short run or long run VECM specification respectively.

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

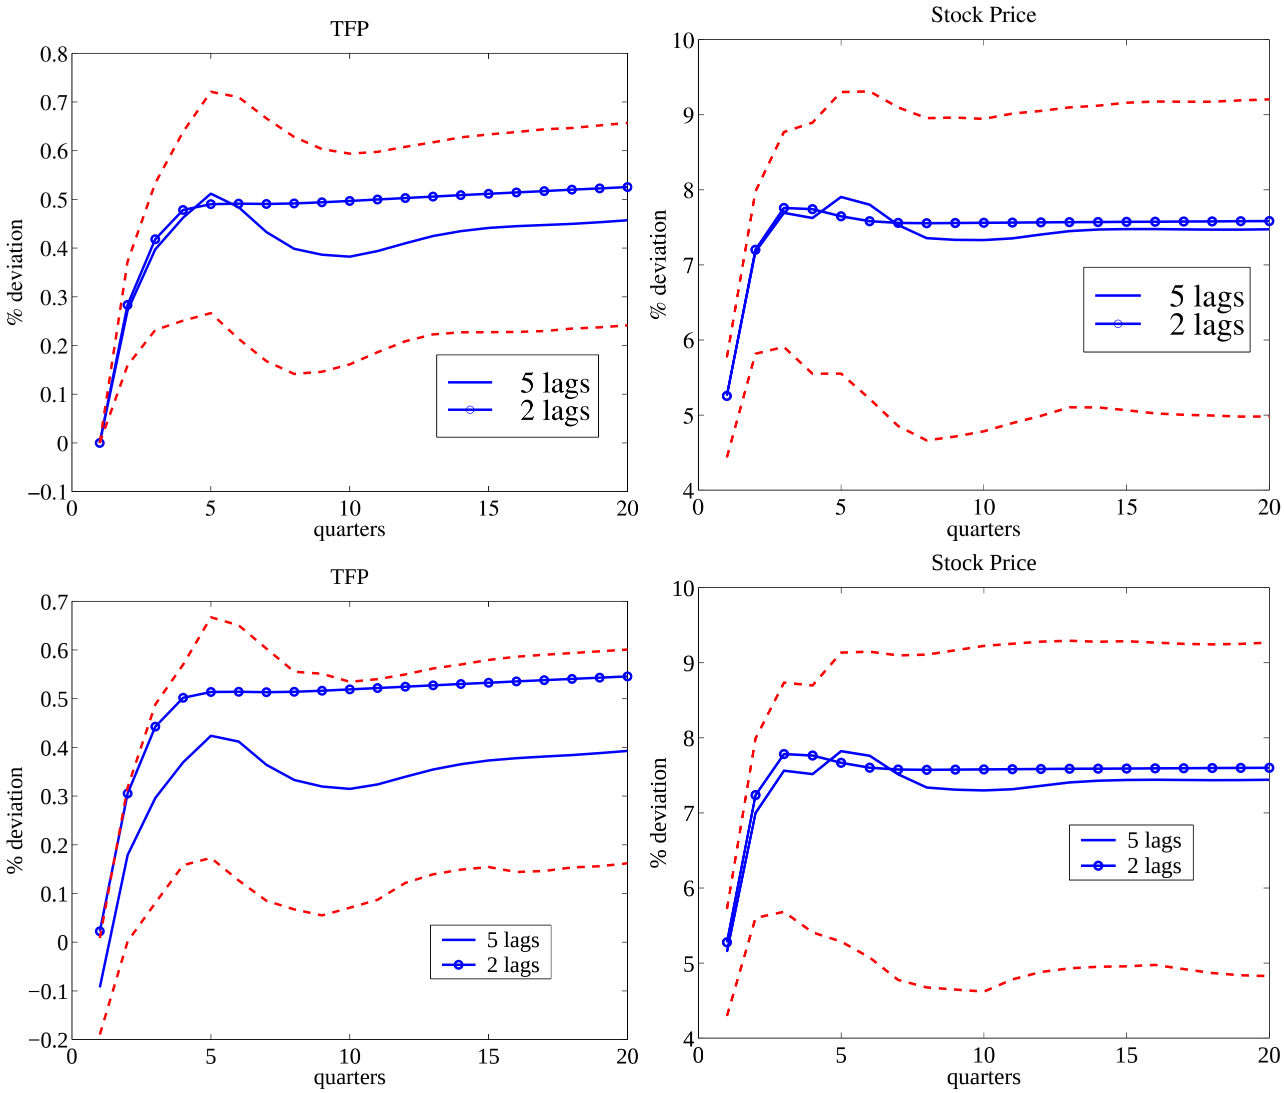

This figure displays the responses of TFP (upper left panel) and stock prices (upper right panel) to a unit ϵ_2 shock (the shock that does not have instantaneous impact of TFP in the short run identification), and the responses of TFP (lower right panel) and stock prices (lower left panel) to a unit $\widetilde{\epsilon}_1$ 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 of TFP and SP in the (TFP, SP) VAR estimated with one cointegrating relation and 5 or 2 lags . The 5% and 95% confidence bands are computed using the short run or long run 5 lags VECM specification respectively.

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)

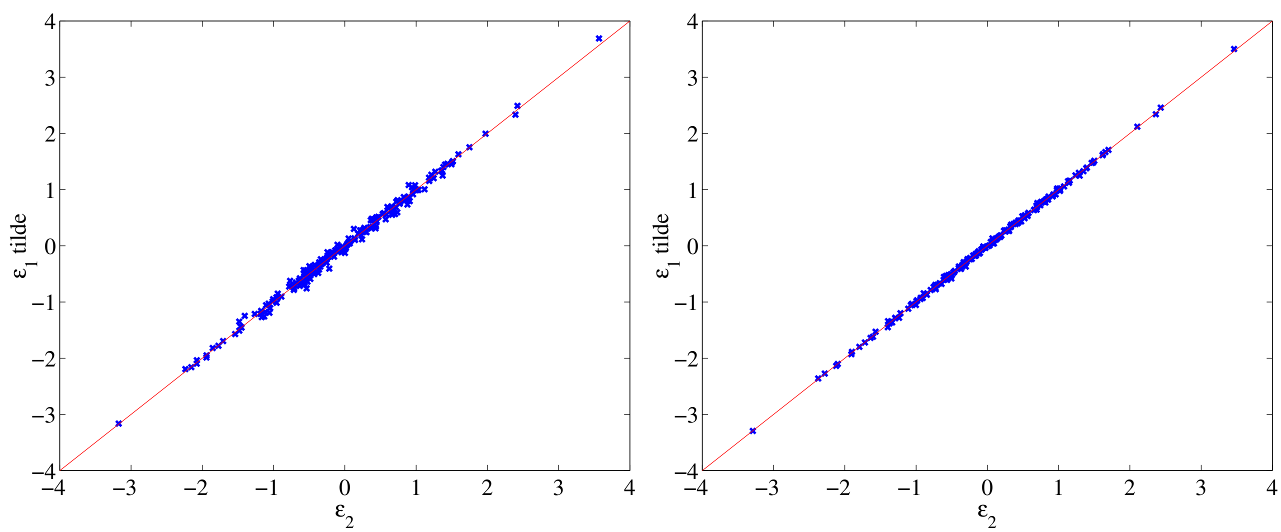

This figure plots ϵ_2 against $\widetilde{\epsilon}_1$. In the left panel, both shocks are obtained from the (TFP, SP) VAR estimated in levels, with 6 lags. In the right panel, both shocks are obtained from the (TFP, SP) VAR estimated in difference, with 2 lags and one cointegrating relation. In both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

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

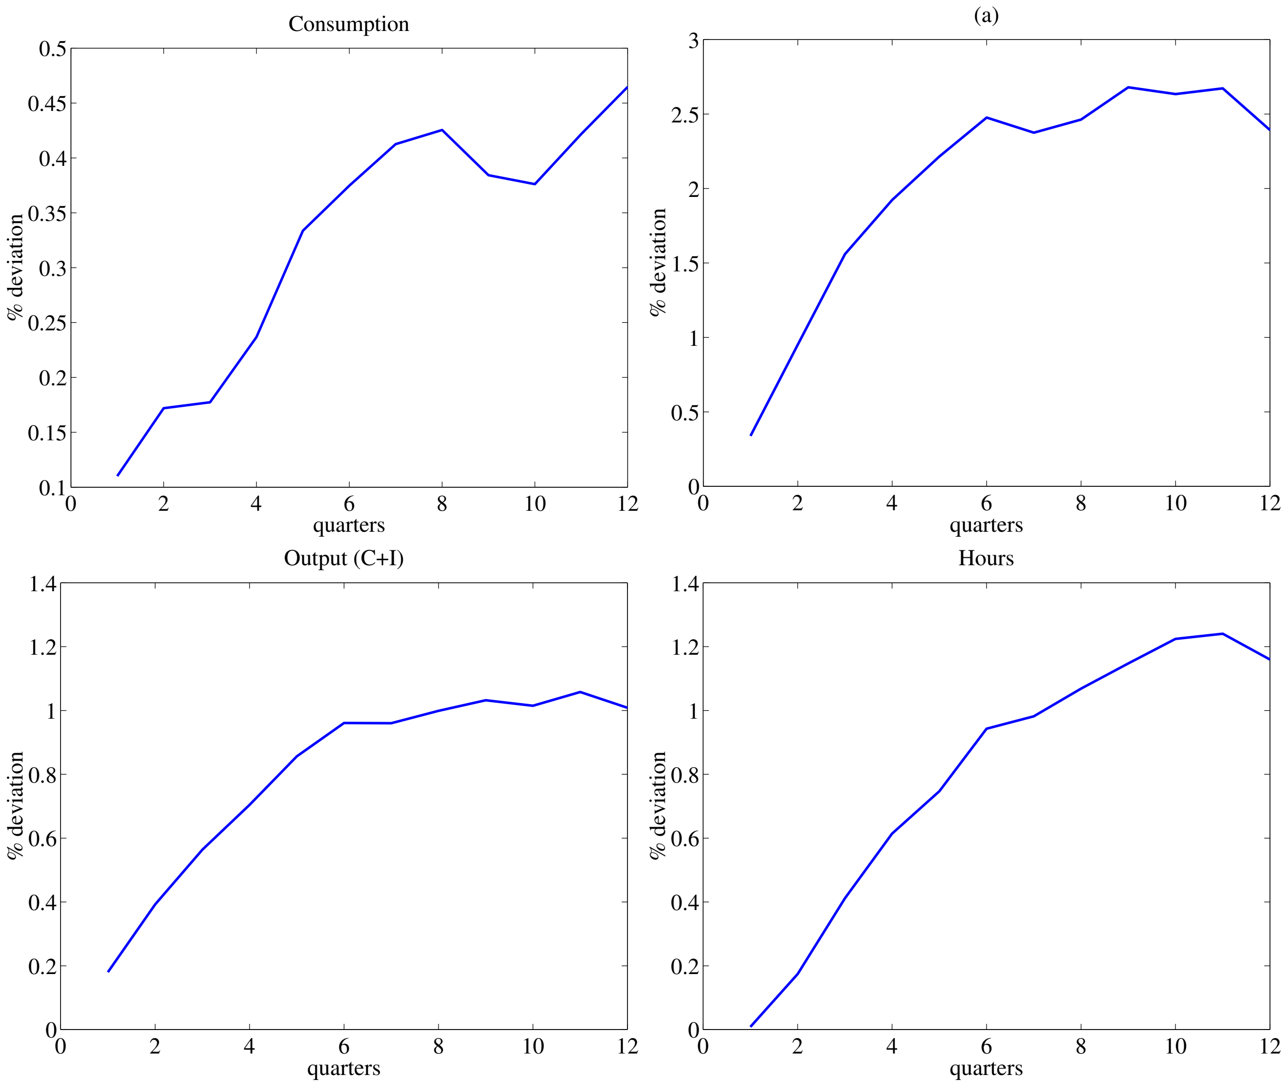

This figure displays the response of consumption, investment, output (measured as $C + I$) and hours to a unit ϵ_2 shock (the shock that does not have instantaneous impact on TFP in the short run identification). The unit of the vertical axis is percentage deviation from the situation without shock (See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

Figure 9: Variance Decomposition in the Baseline (TFP, SP) VAR

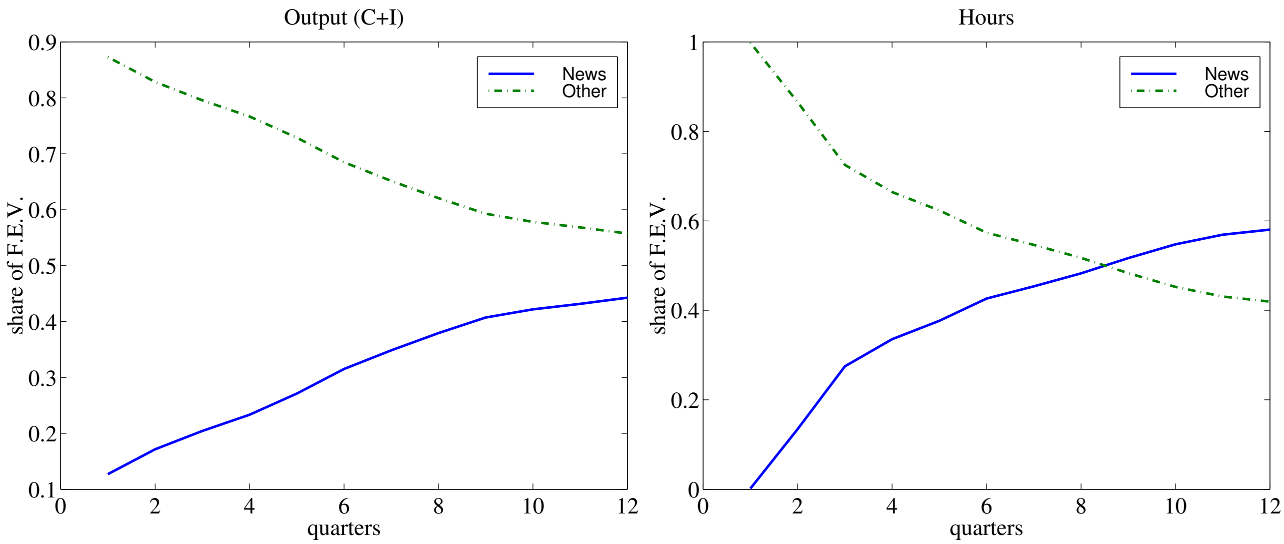

This figure can be roughly interpreted as displaying the share of output (consumption plus investment) and hours forecast error variance attributed to the shock ϵ_2 ("news") and to the other shock (See the main text for more details).

Figure 10: Impulse Responses to ϵ_2 (upper panels) and $\widetilde{\epsilon}_1$ (lower panels) in the in the $(Y/H, SP)$ VAR

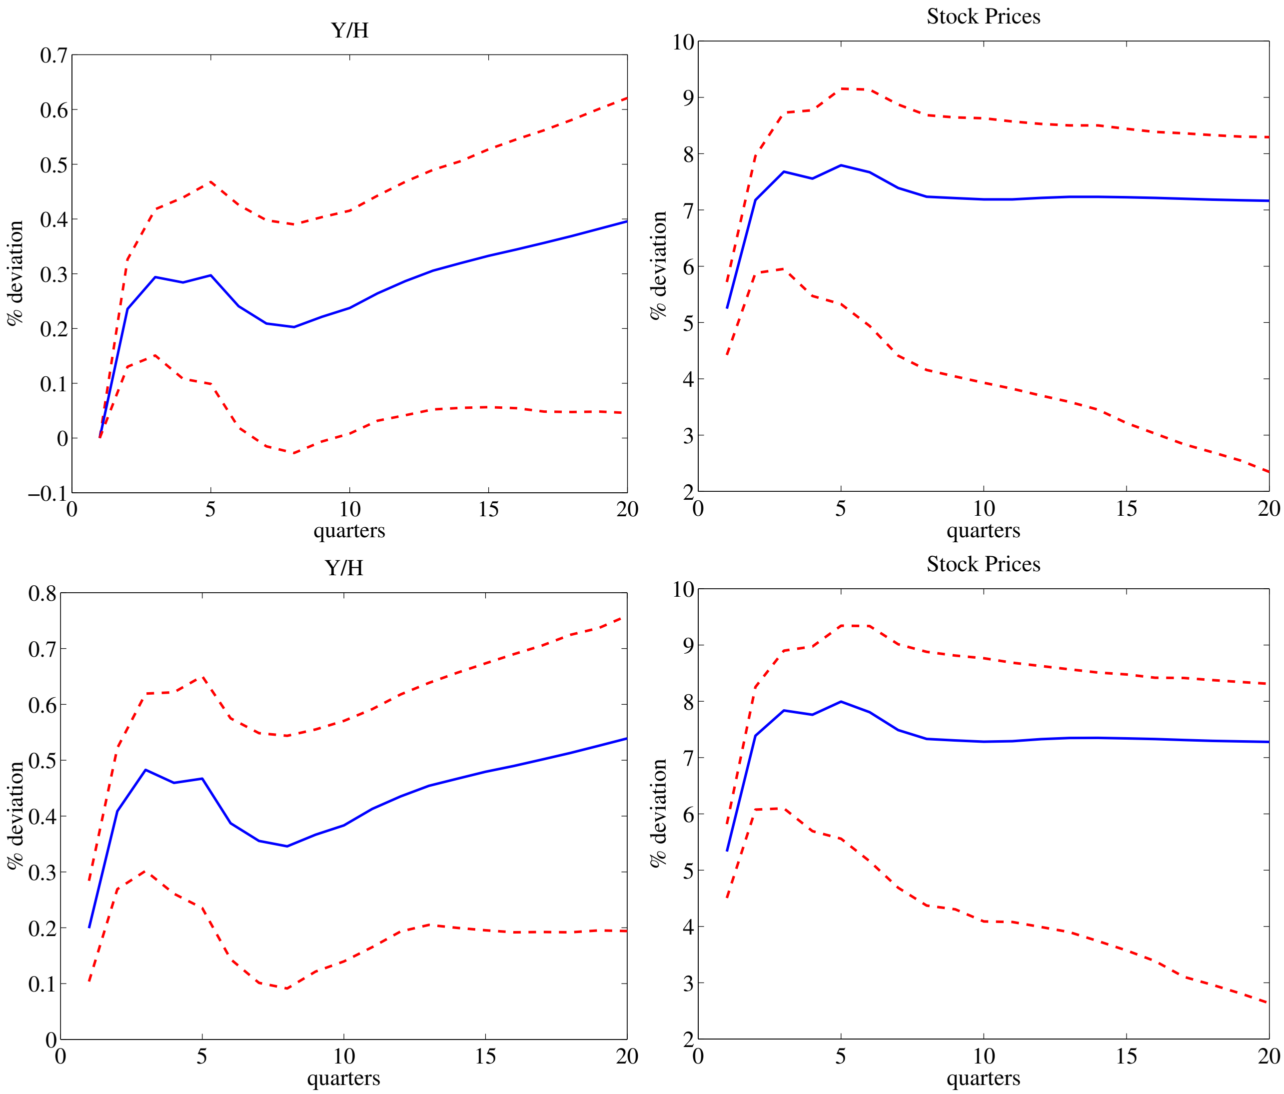

This figure displays the responses of average productivity of hours (Y/H) (upper left panel) and stock prices (upper right panel) to a unit ϵ_2 shock (the shock that does not have instantaneous impact on Y/H in the short run identification), and the responses of Y/H (lower right panel) and stock prices (lower left panel) to a unit $\widetilde{\epsilon}_1$ shock (the shock that has a permanent impact on Y/H 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 5% and 95% quantiles of the distribution of the IRF, this distribution being simulated by bootstrapping 1000 times the residuals of the VAR.

Figure 11: Share of the Forecast Error Variance Attributed to the ϵ_2 (left panel) or $\widetilde{\epsilon}_1$ (right panel) Shock in the baseline $(Y/H, SP)$ VAR

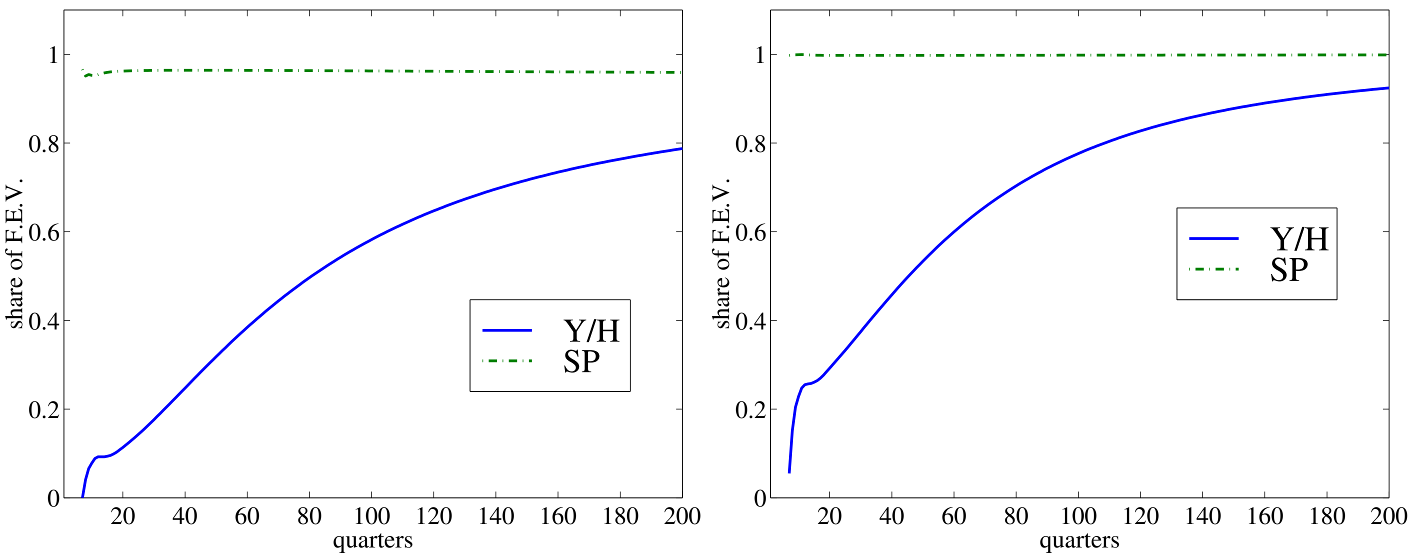

This figure displays the share of Y/H and SP forecast error variance attributed to ϵ_2 (the shock that does not have instantaneous impact of Y/H in the short run identification) (left panel) or to $\widetilde{\epsilon}_1$ (the shock that has a permanent impact on TFP in the long run identification)(right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

Figure 12: ϵ_2 Against $\widetilde{\epsilon}_1$ in the $(Y/H, SP)$ VAR, baseline specification

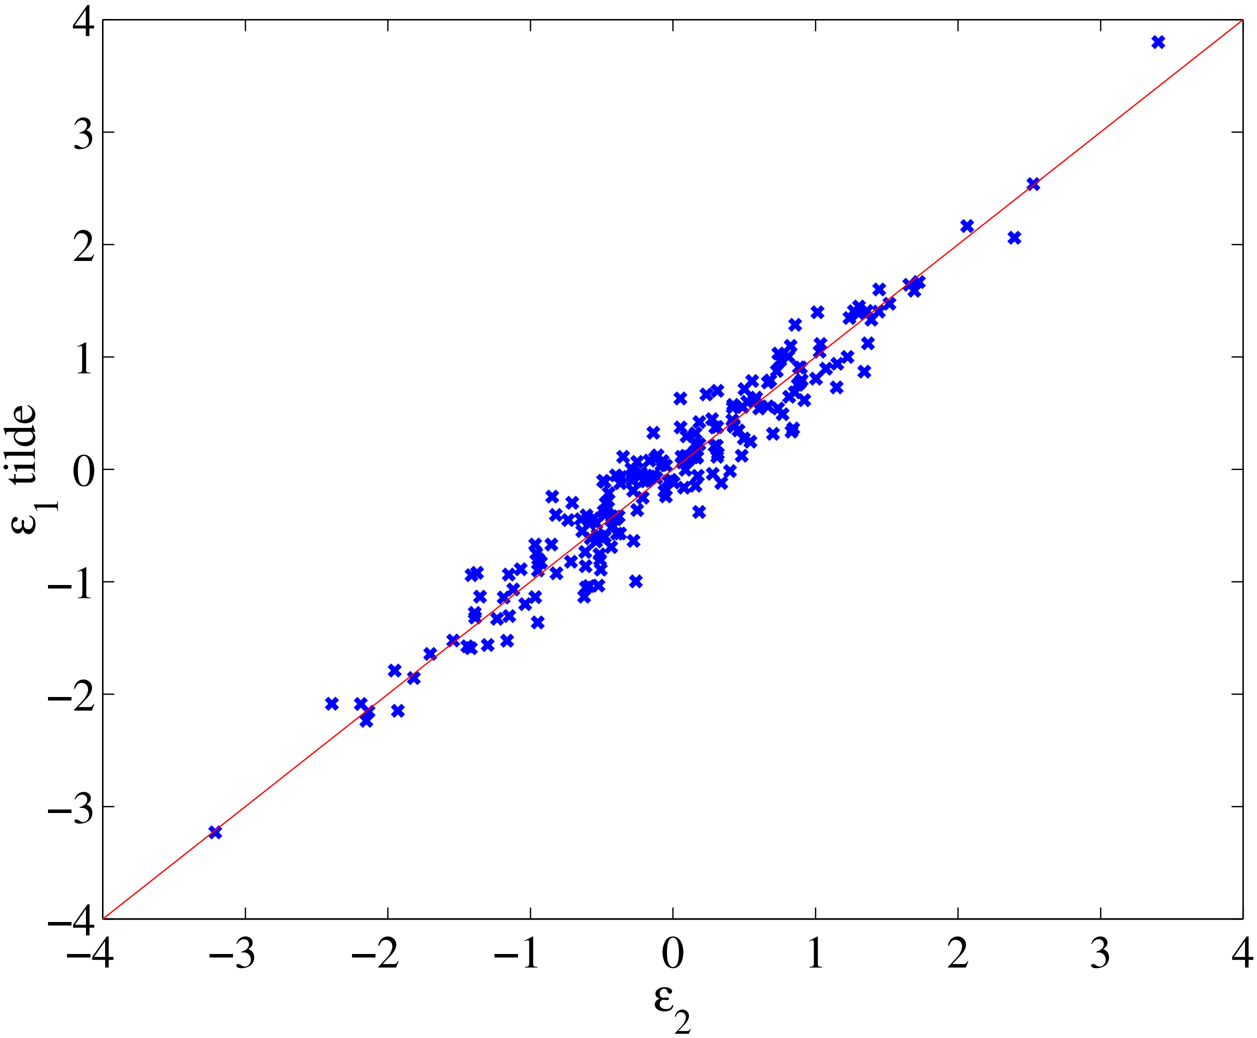

This figure plots ϵ_2 against $\widetilde{\epsilon}_1$. Both shocks are obtained from the baseline $(Y/H, SP)$ VAR, with 5 lags and one cointegrating relation. The straight line is the 45° line.

B.3 Figures related to section 5.1

Figure 13: Impulse Responses to ϵ_2 (upper panels) and $\widetilde{\epsilon}_1$ (lower panels) in the in the (TFP, SP, C) VAR, baseline specification

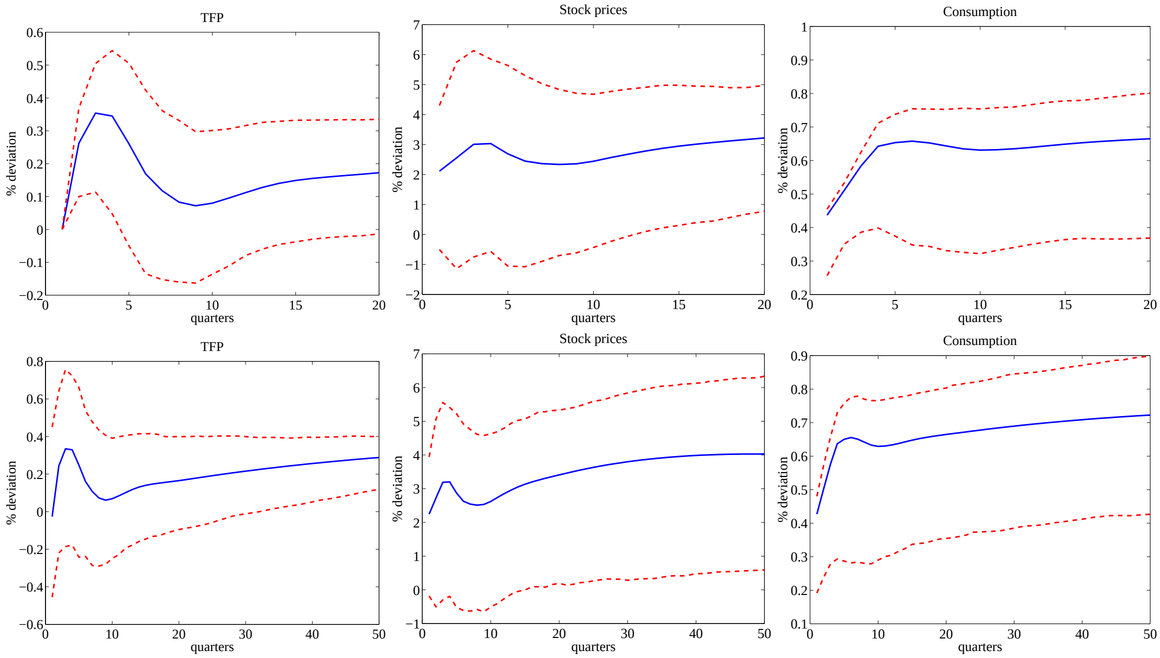

This figure displays the responses of TFP (upper left panel), stock prices (upper center panel) and consumption (upper right panel) to a unit ϵ_2 shock (the shock that does not have instantaneous impact on TFP in the short run identification), and the responses of TFP (lower right panel)), stock prices (lower center panel) and consumption (lower right panel)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 trivariate specification. The unit of the vertical axis is percentage deviation from the situation without shock. Dotted lines represent the 5% and 95% quantiles of the distribution of the IRF, this distribution being simulated by bootstrapping 1000 times the residuals of the VAR.

Figure 14: Share of the Forecast Error Variance Attributed to the ϵ_2 (left panel) or $\widetilde{\epsilon}_1$ (right panel) Shock in the baseline (TFP, SP, C) VAR

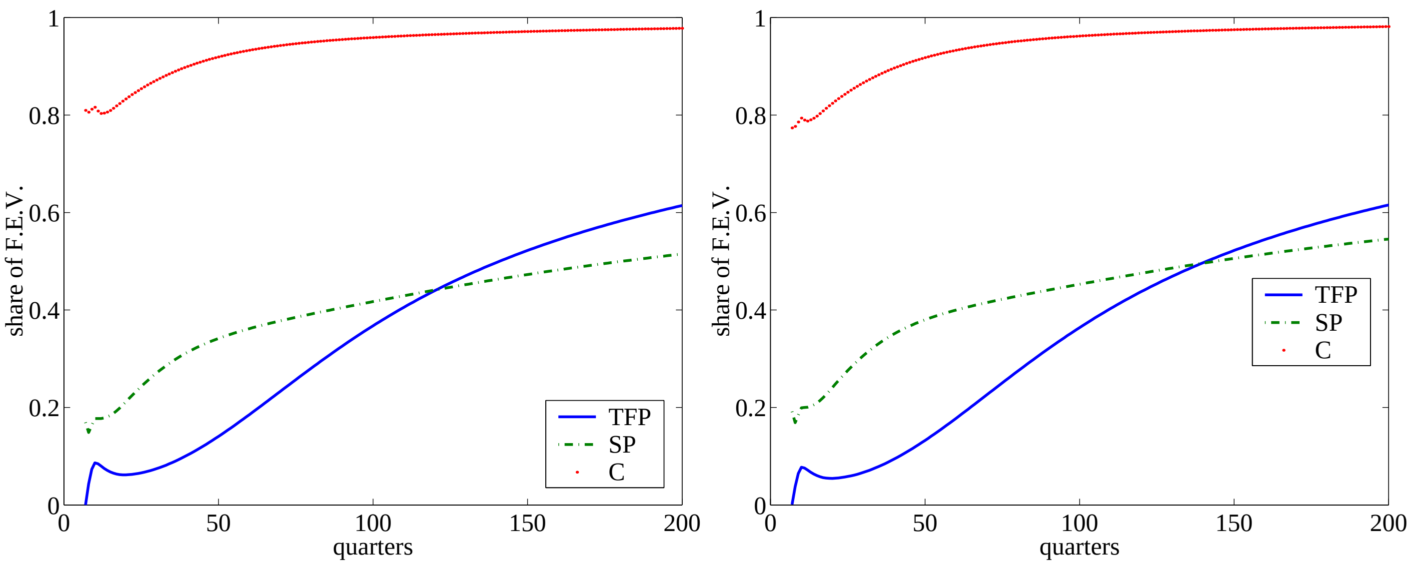

This figure displays the share of TFP, SP and C forecast error variance attributed to ϵ_2 (the shock that does not have instantaneous impact of Y/H in the short run identification) (left panel) or to $\widetilde{\epsilon}_1$ (the shock that has a permanent impact on TFP in the long run identification)(right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

Figure 15: ϵ_2 Against $\widetilde{\epsilon}_1$ in the (TFP, SP, C) VAR, baseline specification

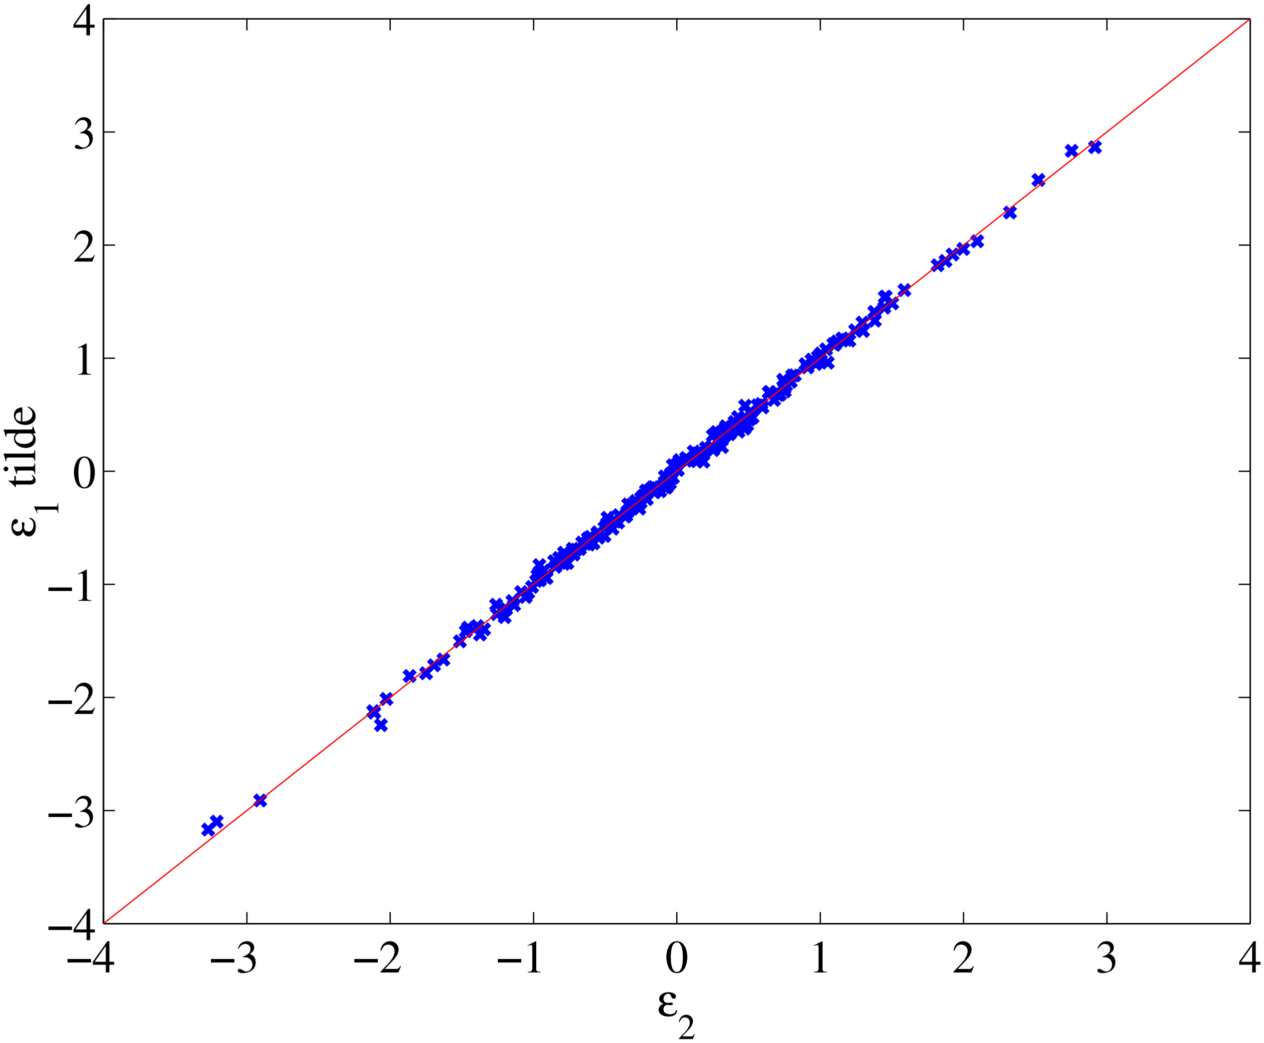

This figure plots ϵ_2 against $\widetilde{\epsilon}_1$. Both shocks are obtained from the baseline (TFP, SP, C) VAR, with 5 lags and two cointegrating relation. The straight line is the 45° line.

Figure 16: Output (defined as $C + I$) Response (a) and Variance Decomposition (b) in the Baseline (TFP, SP, C) VAR

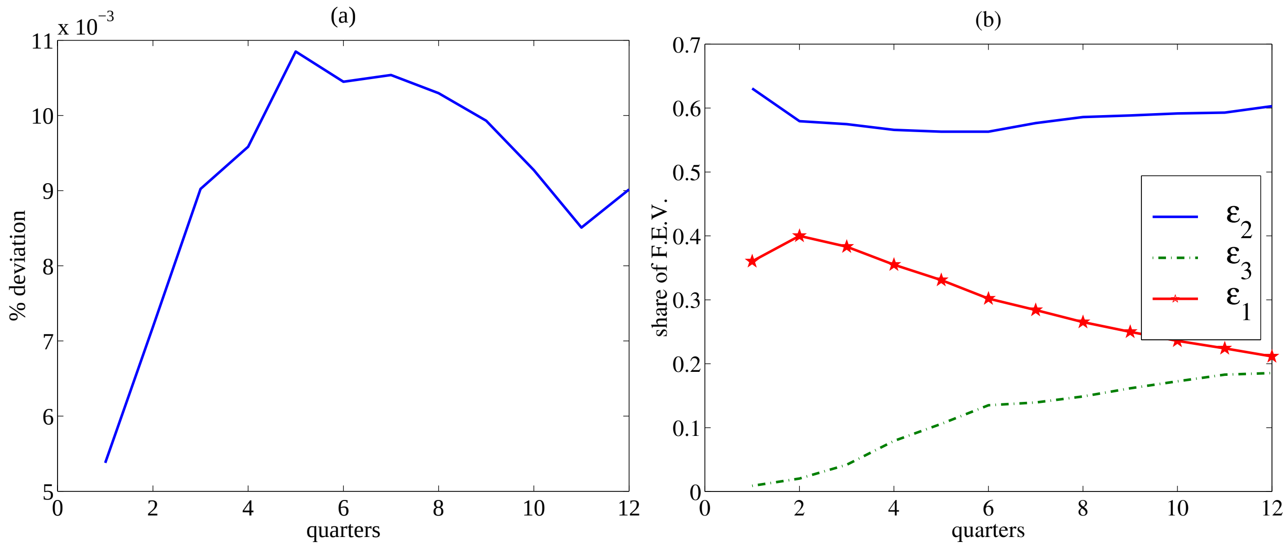

Panel (a) of this figure displays the response of output (defined as $C + I$) to a unit ϵ_2 shock (the shock that does not have instantaneous impact on TFP in the short run identification). The unit of the vertical axis is percentage deviation from the situation without shock. Panel (b) can be roughly interpreted as displaying the share of output (defined as $C + I$) forecast error variance attributed to each of the three ϵ shocks in the baseline specification (See the main text for more details).

Figure 17: Hours Response (a) and Variance Decomposition (b) in the Baseline (TFP, SP, C) VAR

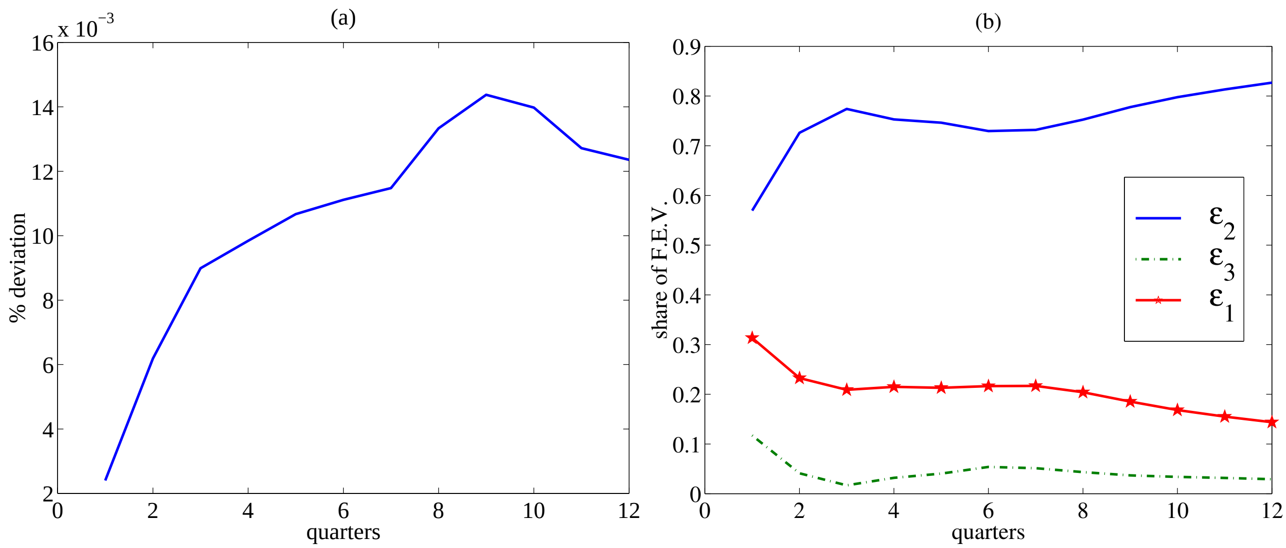

Panel (a) of this figure displays the response of worked hours to a unit ϵ_2 shock (the shock that does not have instantaneous impact on TFP in the short run identification). The unit of the vertical axis is percentage deviation from the situation without shock. Panel (b) can be roughly interpreted as displaying the share of hours forecast error variance attributed to each of the three the three ϵ shocks in the baseline specification (See the main text for more details).

Figure 18: Impulse Responses to ϵ_3 (upper panels) and ϵ_1 (lower panels) in the (TFP, SP, C) VAR, baseline specification

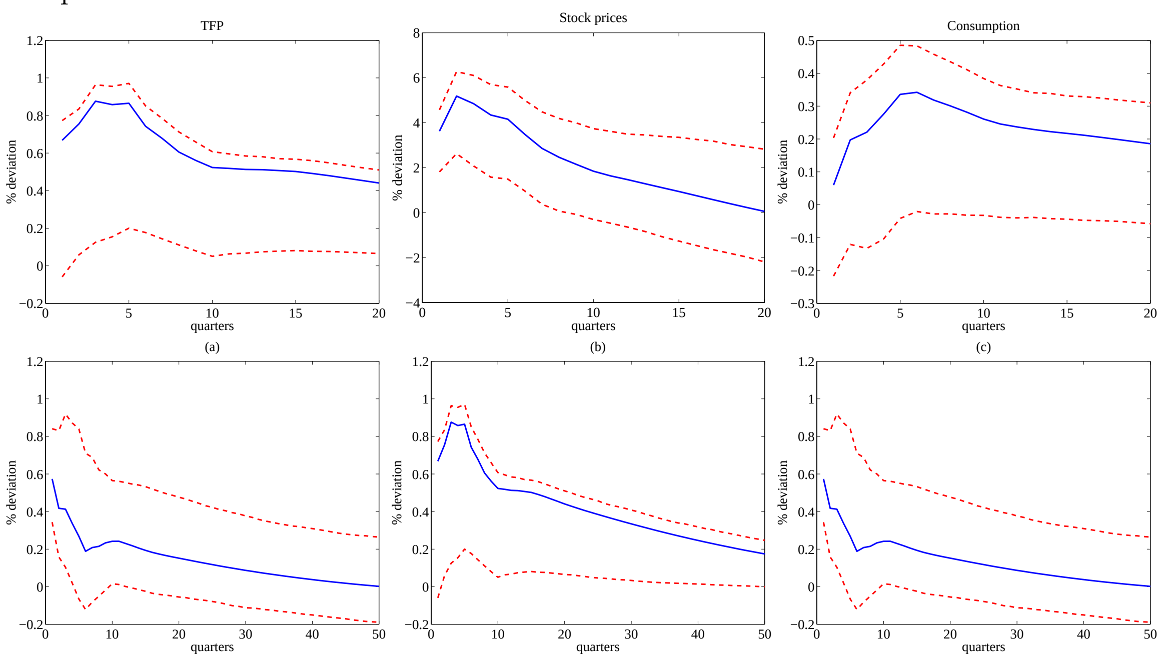

This figure displays the responses of TFP (upper left panel), stock prices (upper center panel) and consumption (upper right panel) to a unit ϵ_3 in the short run identification, and the responses of TFP (lower right panel)), stock prices (lower center panel) and consumption (lower right panel) to a unit ϵ_1 in the short run identification. The unit of the vertical axis is percentage deviation from the situation without shock. Dotted lines represent the 5% and 95% quantiles of the distribution of the IRF, this distribution being simulated by bootstrapping 1000 times the residuals of the VAR.

B.4 Figures related to section 5.2

Figure 19: ϵ_2 Against $\widetilde{\epsilon}_1$ in the (TFP, SP, C, I) VAR (left panel), in the $(TFP, SP, C, C+I)$ VAR (center panel) and in the (TFP, SP, C, H) VAR (right panel), baseline specification

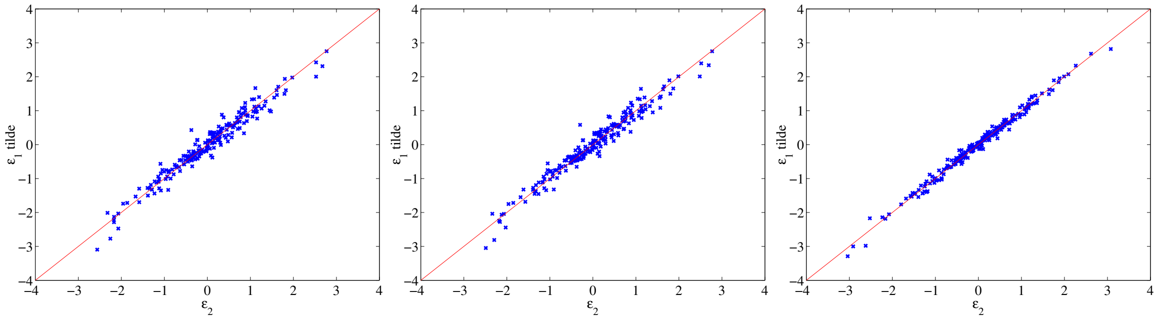

This figure plots ϵ_2 against $\widetilde{\epsilon}_1$. Both shocks are obtained from the baseline (TFP, SP, C, I) VAR, with 5 lags and three cointegrating relation (left panel), from the baseline $(TFP, SP, C, C+I)$ VAR, with 5 lags and three cointegrating relation (center panel) and from the baseline (TFP, SP, C, H) VAR, estimated in levels (right panel). In each panel, the straight line is the 45° line.

Figure 20: Share of the Forecast Error Variance of 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

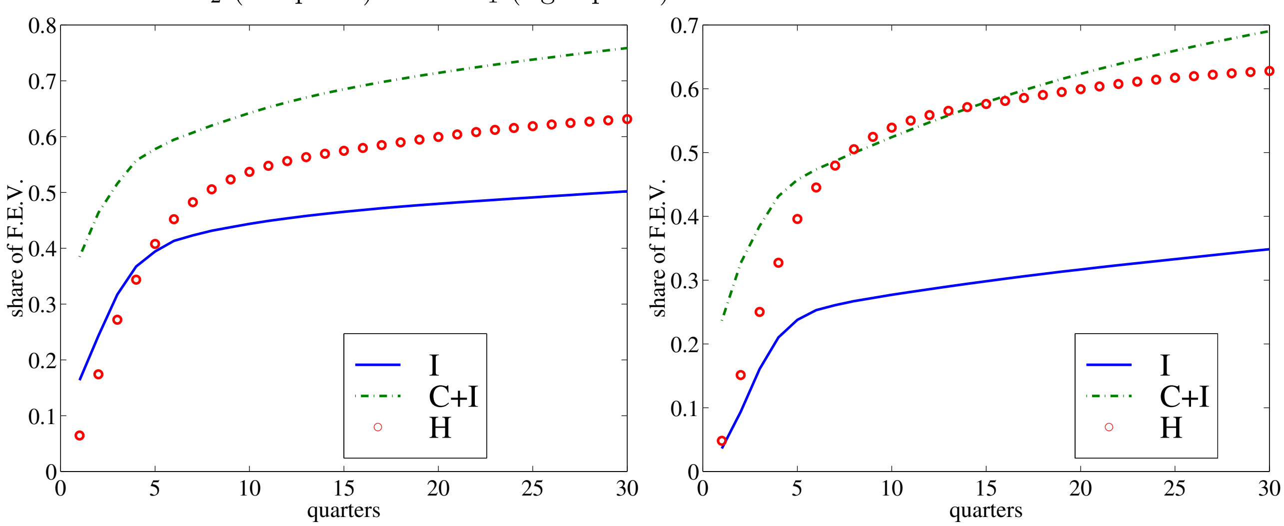

This figure has two panels. The first one displays the share of the forecast variance of investment that is $(TFP, SP, C, C+I)$ VAR and attributable to $\widehat{\epsilon}_1$ (short run identification) in the (TFP, SP, C, I) VAR, of output $(C+I)$ in the $(TFP, SP, C, C+I)$ VAR and of hours in the (TFP, SP, C, H) VAR. The right panel presents the same information in the case of the shock $\widetilde{\epsilon}_2$ (long run identification).

Figure 21: Responses to ϵ_2 (upper panels) and $\widetilde{\epsilon_1}$ (lower panels) in the Baseline (TFP, SP, C, H) VAR

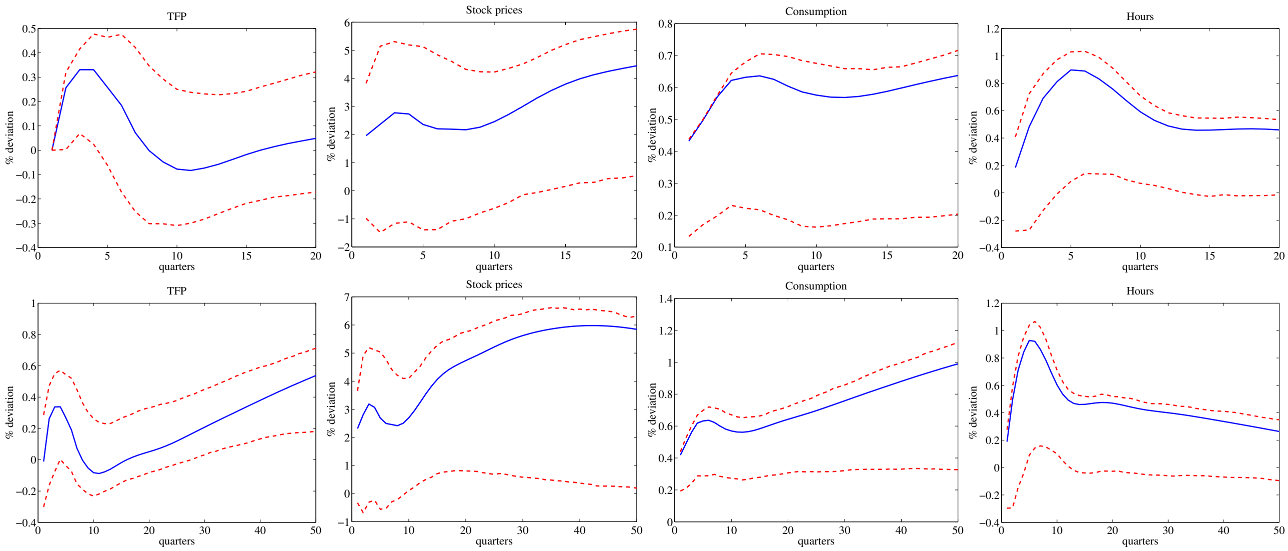

This figure displays the response of TFP (upper-left panel), stock prices (upper-center-right panel), consumption (upper-center-left panel) and hours (upper-right panel) to a unit ϵ_1 in the short run identification, together with response of TFP (lower-left panel), stock prices (lower-center-right panel), consumption (lower-center-left panel) and hours (lower-right panel) to a unit $\widetilde{\epsilon_1}$ in the long run identification. The unit of the vertical axis is percentage deviation from the situation without shock. Dotted lines represent the 5% and 95% quantiles of the distribution of the IRF, this distribution being simulated by bootstrapping 1000 times the residuals of the (TFP, SP, C, H) VAR.