

FEDERAL RESERVE BANK OF SAN FRANCISCO

WORKING PAPER SERIES

**The International Dimension of Productivity and
Demand Shocks in the US Economy**

Giancarlo Corsetti

European University Institute, University of Rome III and CEPR

Luca Dedola

European Central Bank and CEPR

Sylvain Leduc

Federal Reserve Bank of San Francisco

May 2009

Working Paper 2009-09

<http://www.frbsf.org/publications/economics/papers/2009/wp09-09bk.pdf>

The views in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Federal Reserve Bank of San Francisco or the Board of Governors of the Federal Reserve System.

The International Dimension of Productivity and Demand Shocks in the US Economy¹

Giancarlo Corsetti

European University Institute, University of Rome III and CEPR

Luca Dedola

European Central Bank and CEPR

Sylvain Leduc

Federal Reserve Bank of San Francisco

Abstract

Identifying productivity and real demand shocks in the US with sign restrictions based on standard theory, we provide evidence on real and financial channels of their international propagation. Productivity gains in US manufacturing have substantial macroeconomic effects, raising US consumption, investment and the terms of trade, relative to the rest of the world, while lowering US net exports. Significant international financial adjustment occurs via a rise in the global value of the US stock market, portfolio shifts in US foreign assets and liabilities, and especially real dollar appreciation. Positive demand shocks to US manufacturing also lead to real appreciation and raise investment, but have otherwise limited effects on trade flows. This evidence suggests a fundamental role of cross-country endogenous demand and wealth movements in shaping international macroeconomic interdependence.

JEL classification: F32, F41, F42

Keywords: International transmission mechanism, structural VARs, sign restrictions, consumption risk sharing, US dollar real exchange rate.

¹A previous version of this paper was circulated under the title "Productivity and the dollar". We thank our dicussants Emmanuel Farhi and Silvana Teneyro, as well as Fabio Canova, Larry Christiano, Martin Eichenbaum, Fabio Ghironi, Wouter den Haan, Andrea Raffo, Sergio Rebelo, Martin Uribe and seminar participants at the Bank of Canada, Cornell University, Duke University, ECB, ESSIM 2007, the Federal Reserve Bank of Australia, the Federal Reserve Bank of San Francisco Pacific Basin conference, the Federal Reserve Bank of New York, the Federal Reserve Board, Georgetown University, MIT, the NBER Summer Institute 2007, Ohio State University, the University of Aix-Marseille II, University of Amsterdam, U.C. Berkeley, University of Bern, University of Lausanne, and University of Zurich for useful comments. We are grateful to Flavia Corneli and Lenno Uusküla for outstanding research assistance. Corsetti's work on this paper is part of the Pierre Werner Chair Programme on Monetary Union, at the Robert Schuman Centre of the European University Institute. The views expressed here are those of the authors and do not necessarily reflect the positions of the ECB, the Board of Governors of the Federal Reserve System, or any other institutions with which the authors are affiliated.

1 Introduction

What are the consequences of US business cycle shocks for the US trade balance and capital flows, the international value of the dollar and the dynamics of US foreign assets and liabilities? In this paper we carry out a study of the international dimension of US business cycles impulses by identifying structural shocks to productivity and demand in US manufacturing, and tracing their effects on a broad range of variables, as to encompass both real and financial channels of the international transmission.

Understanding international macroeconomic interdependence requires structural analysis, in order to provide guidance on the empirical content of different theories. While most structural investigations of the international dimension of US business cycle have been devoted to the analysis of the effects of policy shocks,² we study demand and supply shocks, providing comprehensive conditional evidence on their open-economy effects. Not only our evidence directly relates to the predictions of a broad range of models in the literature (obviously conditional on specific shocks); it also enlarges the set of empirical results relevant to further modeling, particularly as regards the response of domestic and international relative prices, as well as the financial channels of international transmission.³

We identify productivity and demand shocks in structural VARs via the sign-restriction methodology proposed by Canova and De Nicoló [2002] and Uhlig [2005].⁴ Namely, we impose intuitive, theory-based restrictions, consistent with a large class of models, on the sign of the impulse responses of a small subset of variables, mainly relative prices and quantities. Our focus is deliberately on shocks to US tradable goods production relative to the rest of the industrial world — we identify tradables with the manufacturing sector. Since manufactures account for a large share of international trade, they arguably play a key role in shaping macroeconomic interdependence across countries. Most importantly, as discussed in Section 2, focusing on manufacturing facilitates identification, and avoid potentially difficult issues in the interpretation of the results.⁵

²See e.g. Eichenbaum and Evans [1995]) for monetary policy shocks.

³In notable early work, Clarida and Galí [1994] used long-run restrictions to identify aggregate demand and supply shocks and their effects on the real exchange rate, and cross-country GDP and inflation differentials for the US vis-à-vis the other G7 countries.

⁴The method of identifying monetary policy shocks using sign restrictions on impulse responses has been introduced by Faust [1998] and Uhlig [2005]; Faust and Rogers [2003] and Scholl and Uhlig [2007] provide open-economy applications. Canova and De Nicoló [2002] identify monetary policy shocks using sign restrictions on impulse response correlations. Mountford and Uhlig [2005] extend Uhlig's method to identify fiscal policy shocks orthogonal to the business cycle and monetary policy shocks. Canova and Pappa [2006] also identify fiscal shocks using sign restrictions on impulse response correlations. Dedola and Neri [2007] have used sign restrictions to identify technology shocks and Mountford [2005], and Enders et al. [2008], among others, have addressed the identification of multiple shocks using sign restrictions — the latter in an open-economy context. Finally, Canova and Paustian [2008] provide a thorough assessment of the properties of sign restrictions using data simulated from DSGE models.

⁵In previous work (Corsetti, Dedola and Leduc [2006]), we used long-run restrictions as in Galí [1999]

In our estimates, productivity shocks to US manufacturing have significant macroeconomic effects. Consistent with the predictions of the intertemporal-trade approach to the current account, we find that positive shocks persistently increase US aggregate consumption and investment relative to the rest of the world, raising imports and worsening the US trade balance. However, we also find that over the business cycle these shocks lead to real appreciation, rather than depreciation, in all our measures of US international relative prices, despite the persistent increase in US manufacturing output relative to the rest of the world. As we consider a CPI-based, a PPI-based and an export-deflator-based real exchange rate (the latter proxying for bilateral terms of trade), the real dollar appreciation triggered by a positive productivity shock to US manufacturing is not exclusively due to a rise in the relative price of US nontradables relative to US tradables— according to the classical Harrod-Balassa-Samuelson effect — but also to an improvement in the US terms of trade. In line with Engel [1999], these findings suggest that movements in traded goods prices are major drivers of the dollar real exchange rate.

The repercussions of productivity shocks through the international financial adjustment mechanism are substantial. In addition to causing real appreciation, these shocks sharply raise the value of the US stock market relative to an aggregate index of foreign markets; they open a positive nominal interest differentials in favor of the US over time. Using the series of valuation-adjusted US Foreign assets and liabilities computed by Gourinchas and Rey [2007], we find that in addition to deteriorating net trade, asymmetric productivity gains worsen the US net foreign asset position relative to GDP. Notably, the fall in US net foreign wealth corresponds to an increase in both gross assets and gross liabilities.

Concerning demand shocks to US manufacturing, we find that they typically have the expected qualitative effects, but, quantitatively, tend to be inconsequential for several aggregate variables. In our identification, a positive shock to demand for US manufacturing is assumed to lead to an increase in both the relative price and output of US manufacturing (the latter vis-à-vis US GDP as well as manufacturing output abroad). In response to such shock, US absorption rises, reflecting an increase in aggregate investment, while international relative prices strengthen significantly. However, the response of consumption, trade flows, gross and net external asset positions, stock prices and interest differentials are all rather subdued. Despite the rise in investment and the real appreciation, we find limited crowding out of net exports.

to identify technology shocks in the manufacturing sector in a sample of five G7 countries, and found their effects on the seven US variables considered, to be similar to those of productivity shocks in this paper. However, here we first propose a novel identification strategy addressing many of the shortcomings arising from the adoption of long-run restrictions with sectoral data. Second, we study in much greater detail the US economy, identifying the effects of two shocks (productivity and demand), on as many as 17 variables, systematically documenting real and financial channels of transmission.

Overall, these findings provide conditional support for key building blocks of modern macroeconomic theory, such as intertemporal consumption smoothing. However, they also question specific aspects of the international transmission as envisioned by the vast majority of open-economy models.

First, contrary to the conventional wisdom rooted in the Mundell-Fleming tradition, real demand shocks for US manufacturing which raise domestic absorption and appreciate the dollar in real terms, do not produce appreciable crowding-out effects on net trade. Second, our results suggest that domestic productivity gains in the US do not necessarily lead to an efficient allocation at global level, as the consumption (and thus wealth) expansion in the US relative to the rest of the world is accompanied by appreciation of US international relative prices — preventing other countries from benefitting from the US-based technological improvement in terms of cheaper imports, and thus from sharing production risk. This is so in spite of the marked relative increase in US equity prices following these shocks, suggesting that widely traded assets provide potential opportunities for risk diversification. Such findings are at odds with predictions by standard business cycle models, assuming that productivity risk is efficiently shared across border, either by virtue of international price movements or through risk-hedging in financial markets.

In this respect, our results emphasize the need to deepen our understanding of asymmetric wealth and demand effects of business cycle impulses, as a key dimension of the international transmission mechanism. The traditional Mundell-Fleming-Dornbusch (MFD) framework emphasizes asymmetric real demand disturbances treating them as exogenous. As amply discussed in the literature (see e.g. Dornbusch [1980]), however, many sharp predictions in the MFD tradition obtain only because, while postulating demand shocks, this framework does not explicitly account for any endogenous changes in wealth, which may work as to largely offset the effects of the exogenous disturbances in many instances.⁶ Different from the traditional MFD analysis, modern open-economy theory makes it clear that different structural impulses, including supply shocks, can generate endogenous movements in wealth and thus demand across countries. Yet these movements can matter only in economies where national agents have limited ability to hedge country-specific risk.

Indeed, endogenous wealth divergences are ruled out by assumption in the open-economy literature positing complete markets or a high degree of cross-border risk sharing.⁷ Consistent with our findings, instead, recent theoretical contributions such as Ghironi and Melitz

⁶This issue is at the heart of the debate on the Harberger-Laursen-Metzler effect, revolving around the relationship between aggregate spending, the current account and international relative prices — see e.g. Obstfeld [1982] and Svensson and Razin [1983].

⁷This is true for many seminal contributions to the modern international real business cycle literature — e.g. see Backus, Kehoe and Kydland [1994] and Stockman and Tesar [1995]) — as well as to the modern sticky price literature — e.g., see Obstfeld and Rogoff [1995] and Chari, Kehoe and McGrattan [2002].

[2005] and Corsetti, Dedola, and Leduc [2008a] show that, under incomplete markets, supply disturbances lead to wealth and demand responses disproportionately concentrated in the economy experiencing a shock. In the former contribution, asymmetric wealth effects from a positive technology shock are associated with an appreciation of the terms of trade and inefficient adjustment at the extensive margin.⁸ In the latter contribution, persistent productivity and output booms in a large, relatively closed country like the United States can induce substantial movements in cross-country wealth and demand that cause both the trade balance to deteriorate and the international price of domestic goods to appreciate.

The paper is organized as follows. Section 2 briefly reviews the international transmission mechanism in standard theoretical and quantitative models, identifying alternative views and empirical predictions on which we base our sign restrictions and define the key questions our empirical analysis will address. Section 3 describes the data and the empirical methodology. Section 4 reports and analyzes in detail our main findings, while Section 5 presents some sensitivity analysis. Section 6 concludes.

2 Using theory insights to identify productivity and demand shocks

In this section we motivate and lay out the theoretical underpinnings of our approach to identification. Our identification strategy via sign restrictions consists of focusing on a minimal set of theoretical predictions on the direction of key variables' responses to a given shock, which are consistent with a large class of analytical frameworks. In other words, our identification scheme hinges only on those restrictions to variables whose behavior in response to shocks is both unambiguous according to theory and fairly uncontroversial. Conversely, the response of variables for which theory presents a fairly wide range of predictions is left unconstrained, so as to reconsider theoretical controversies in light of the results from our analysis.

2.1 Identification strategy

Our strategy builds on a straightforward idea: supply-side shifts should move relative prices and quantities in opposite directions, while demand shifts should move them in the same direction. Since this idea is especially powerful when applied to sector- or industry-specific

⁸Ghironi and Melitz [2005] refers to the 'terms of labor', that is, relative cost of effective units of labor across countries, which is proportional to the terms of trade. In a related setup, Corsetti, Martin and Pesenti [2007] shows that no terms of trade appreciation occurs under complete markets.

shocks,⁹ a natural focus for studies of the international transmission mechanism is on the tradable sector, which we identify with manufacturing. In most open economy models both productivity and demand shocks in this sector increase tradable output relative to GDP; however, productivity gains unambiguously decrease the price of domestic tradables in terms of domestic nontradables — a well-understood prediction of the standard Harrod-Balassa-Samuelson theory — while demand shocks unambiguously increase it. Thus, we can use the domestic relative price of tradables in terms of nontradables as the key price variable in our identification scheme. As explained in Section 2.2, this choice is superior to the alternative of using an international relative price, like the terms of trade. This is because theory’s predictions on the response of the international prices to shocks are conditional on the amount of cross-country insurance, the functioning of international financial markets, and more generally the international transmission mechanism, whose investigation is instead the main objective of our study.

Focusing on manufacturing not only facilitates identification but also avoids potentially difficult issues in the interpretation of the results, relative to the alternative of analyzing economy-wide shocks with unspecified industry origin. In this alternative, the response of many macro (internal and external) variables, including relative prices, would indeed be sensitive to the distribution of shocks across the tradable and the nontradable sector, among other things. In other words, interpreting empirical results would require additional and possibly controversial assumptions about the relative importance of aggregate disturbances in each sector.¹⁰

By the same token, consistent with our interest in tracing the external effects of shocks to isolate the channels of international transmission, our analysis naturally focuses on those shocks hitting the US asymmetrically vis-à-vis the rest of the world. Failure to focus on asymmetric disturbances would raise issues in interpretation of shocks of unspecified geographic origin, similar to the one discussed above.¹¹ Our identification scheme is detailed below.

⁹It is more problematic to identify demand and supply shifts by looking at movements in the *level* of prices and aggregate quantities, especially because this requires conditioning on a given monetary policy response to shocks.

¹⁰For instance, suppose that we find a positive association between the level of the US economy-wide labor productivity and a terms-of-trade deterioration. It would be quite difficult to infer that the depreciation is evidence in favor of a particular transmission mechanism without knowing whether the productivity increase is concentrated in tradables or nontradables, and the extent to which the distribution of shocks is stable over time.

¹¹The interpretation of the international repercussions of global shocks would again require auxiliary assumptions about their distribution and consequences across countries. For instance, suppose that we find a positive association between the level of US labor productivity and the US trade deficit. Could we infer that this is evidence in support of the intertemporal approach to the current account? Unfortunately, the answer to this question is ‘No’. Without controlling for movements in foreign productivity we could not reach this conclusion, as e.g. forcefully argued by Glick and Rogoff [1995].

Productivity in US manufacturing The first shock under consideration consists of supply shocks increasing labor productivity in the US manufacturing sector, relative to that in the rest of the world. To identify these shocks we postulate a set of four restrictions. Positive supply shocks should: (1) raise (the log of) manufacturing output relative to aggregate output in the US; (2) lower the domestic price of US manufacturing relative to US nontradables in accordance with the Harrod-Balassa-Samuelson (HBS) effect; (3) raise US manufacturing output relative to Foreign manufacturing output; and (4) raise US labor productivity relative to Foreign labor productivity in manufacturing.

As explained earlier, the theoretical underpinning of restrictions (1) and (2) is that supply shifts move price and quantity in opposite directions: according to standard HBS theory, productivity gains in tradables should increase the domestic price of nontradables in terms of tradables.¹² Restriction (3) isolates shocks with US-specific effects. The last restriction ensures that our identified supply shocks are associated with an increase in relative labor productivity, as would be the case with standard technology shocks analyzed by the international real business cycle (IRBC) literature. All other variables included in our analysis, namely, aggregate consumption and investment, trade variables, capital flows, international relative prices and asset prices, are left unconstrained, so that the external consequences of productivity shocks can be traced in the data.

Demand for US manufacturing The second shock consists of demand shifts in favor of US manufactured goods, relative to all other goods and services produced in the US. The identification of these shocks is based on a set of three restrictions. Positive demand shocks specific to US manufacturing goods should: (1) raise manufacturing output relative to aggregate output in the US; (2) raise the relative price of manufacturing in terms of other goods in the economy; and (3) raise US manufacturing output relative to Foreign manufacturing output.

Restrictions (1) and (2) now capture the theoretical prior that demand shifts move price and quantity in the same direction (rather than in opposite directions, as is the case for productivity shocks); restriction (3) instead isolates shocks that are specific to US tradable production, relative to the other countries in our sample. These restrictions are consistent

¹²The Harrod-Balassa-Samuelson effect is sometimes related to the ‘Penn Tables effect’, i.e., the observation that consumer price levels in wealthier countries are systematically higher than in poorer ones, when expressed in the same currency. The textbook rendition of the theory typically relies on small open-economy models with country-specific nontradables, but one homogeneous tradable good (so that the terms of trade are constant), abstracting from general equilibrium effects stemming from asset trading. In this case, an increase in the domestic relative price of nontradables unambiguously translates into real appreciation. As discussed in Section 2.2 below, however, when national tradables are differentiated, and the role of cross-country risk insurance is taken into account, the theoretical prediction that productivity gains in tradables lowers their price in terms of domestic nontradables (the effect at the core of HBS theory) does not necessarily translate into real exchange rate appreciation (the Penn Tables effect).

with the kind of good-specific taste shocks in Stockman and Tesar [1995], as well as with sector-specific real demand shock in the Mundell-Fleming-Dornbusch tradition. All other variables included in our analysis are, again, left unconstrained, so as to trace the external consequences of demand shocks in the data.

Caveats and qualifications The two sets of restrictions defined above can be derived from a vast majority of models in the literature. One model which can be solved in closed form is provided by Obstfeld and Rogoff [2000]; others are provided by the IRBC literature, e.g. Stockman and Tesar [1995]. However, we are well aware of the fact that no identification scheme is ironclad, and it is possible to concoct examples in which industry shocks violate our restrictions.

For our purposes, a specific concern is that different shocks with similar effects on restricted variables may elicit opposite responses of unconstrained variables. For instance, our productivity shocks clearly capture exogenous shifts in the production function of manufacturing due to technology improvements. Yet, to the extent that measured labor productivity is endogenous, e.g. because of labor hoarding, these shocks could also correspond to other supply disturbances that do not immediately shift the production function, like embodied capital shocks or labor supply shocks. Similarly, our identified demand shock could have domestic as well as foreign origin.¹³

In these cases, it could be possible that our identified shocks end up having no significant estimated effects on unconstrained external variables. Lack of detectable effects would clearly raise serious interpretation issues. However, to the extent that our analysis does recover some responses with a high degree of probability, our results provide model builders with useful evidence. For a candidate shock to be the main driver of our results, it should be able to account for both the estimated responses of constrained variables, e.g., an increase in labor productivity, as well as the response of any other unconstrained variables we find in the data.

2.2 International prices and risk sharing in the international transmission mechanism

In this subsection, we briefly review a few theoretical results supporting the notion that the domestic relative price of manufacturing goods — our measure of tradables — is a natural candidate for identification via sign restrictions, while both the real exchange rate and the terms of trade should be left unconstrained in the analysis, as key indicators to discriminate across international transmission mechanisms.

¹³Our restrictions for the demand shock, however, rule out that an increase in manufacturing TFP in the rest of the world would qualify as a demand shock for US manufacturing, as this shock should imply a violation of restriction (3).

Highly influential contributions to the modern literature placing emphasis on supply disturbances — e.g. Backus, Kehoe and Kydland [1994] and Stockman and Tesar [1995] for the IRBC literature, and Obstfeld and Rogoff [1995] and Chari, Kehoe and McGrattan [2002] for the sticky-price literature — have a specific prediction concerning the behavior of international relative prices: in a boom, the exchange rate should depreciate in real terms, and the relative price of domestic tradables should worsen in terms of foreign tradables. A key reason for this prediction is that these models have been developed either under the assumption of complete markets, or using specifications implying a high degree of international risk sharing.

The link between risk sharing and the response of international relative prices to shocks is best illustrated by using a well-known result in general-equilibrium theory (see, e.g., Gravelle and Rees [1992], Chapter 20): in an efficient allocation, across any two households the marginal utility of consumption should be lower for the one whose consumption is cheaper. For two national representative households residing in different countries, this means that, in a decentralized equilibrium with full consumption risk sharing, the ratio of their marginal utilities should be proportional to their bilateral CPI-based real exchange rate (RER) — see e.g., Backus and Smith [1993] and Obstfeld and Rogoff [2001]:

$$\frac{U_C(\cdot)}{U_{C^*}(\cdot)} = \kappa \frac{P}{P^*} = \kappa \cdot RER \quad (1)$$

Here, C denotes domestic consumption, $U_C(\cdot)$ marginal utility of consumption, and P and P^* the domestic and foreign price levels, expressed in terms of the same numéraire. Under standard assumptions about preferences, the above condition implies that domestic consumption can rise relative to foreign consumption only if its relative price is low, i.e. the real exchange rate simultaneously depreciates.¹⁴

In an environment with tradable goods only, the above stark prediction for the real exchange rate has immediate consequences for the relative price of tradables across countries $\left(\frac{PT}{PT^*}\right)$, as the former is a direct function of the latter. Therefore, real exchange rate depreciation would be reflected in a deterioration of the relative price of domestic tradables in terms of foreign tradables.¹⁵ In the more general case with both nontradables and tradables, we can rewrite the above condition distinguishing between $\left(\frac{PT}{PT^*}\right)$ and the ratio of the

¹⁴With non-separable preferences between consumption and leisure, it is not a priori obvious that under complete risk-sharing, a domestic relative consumption boom should be associated with a real depreciation. However, the quantitative literature has amply shown that this obtains under plausible parameterizations (see e.g. Backus et al. [1994]).

¹⁵With national specialization in the production of tradables, this would be the case under home bias in consumption.

within-country relative price of tradables in terms of nontradables $\left(\frac{PN/PT}{PN^*/PT^*}\right)$:¹⁶

$$\frac{U_C(\cdot)}{U_{C^*}(\cdot)} = \kappa \cdot RER \left(\frac{PT}{PT^*}, \frac{PN/PT}{PN^*/PT^*} \right). \quad (2)$$

Consider productivity gains in the tradable sector of the domestic economy. According to the within-country HBS effect, these shocks increase the domestic price of nontradables in terms of tradables, PN/PT . Holding constant the other relative prices, PT/PT^* and PN^*/PT^* , the higher price of nontradables due to the HBS effect would make domestic consumption more expensive relative to foreign consumption, thus appreciating the domestic real exchange rate. But in this case, *under perfect risk sharing*, domestic consumption would need to drop relative to the foreign one: by condition (2), relative domestic consumption can rise with a domestic productivity shock to tradables only if the real exchange rate simultaneously depreciates.¹⁷

Perfect risk sharing therefore has the following important implication: for relative consumption to increase in the country benefitting from technology gains, the international price of tradables produced in that country must necessarily worsen, consistent with efficiency of the global allocation. Hence, PT/PT^* must fall enough as to outweigh changes in the domestic relative price of nontradables PN/PT (net of changes in PN^*/PT^*), up to causing the required overall depreciation of the real exchange rate.

Absent deviations from the law of one price in tradables, a fall in their cross-country relative price PT/PT^* would directly translate into a worsening of the terms of trade of the country — defined as the the price of the country’s exports in terms of its imports. So, the above result would apply to both PT/PT^* and the terms of trade. In the presence of deviations from the law of one price, however, there is no simple link between PT/PT^* and the terms of trade: a different response of these prices to (productivity) shocks is possible, raising interesting empirical issues.¹⁸ This observation motivates the inclusion of both prices as unconstrained variables in our empirical analysis.

The above analysis goes through in models explicitly positing complete markets, as well as

¹⁶In Engel [1999], for instance:

$$\log RER \approx \log \frac{P_T}{P_T^*} + (1 - \gamma^*) \log \left(\frac{P_T^*}{P_N^*} \right) + (1 - \gamma) \log \left(\frac{P_N}{P_T} \right),$$

where γ is the share of traded goods in consumption.

¹⁷Observe that this must be so independent of the shock persistence, trade elasticity, as well as trade frictions.

¹⁸With deviations from the law of one price due to pricing to market, e.g. due to price stickyness in the importer’s currency, the tight relation among the real exchange rate, PT/PT^* and the terms of trade may be weakened and even reversed in the short run (see e.g. Corsetti et al. [2008b]).

in models assuming incomplete markets, yet envisioning allocations that are close to perfect risk sharing, in the sense that they predict a positive and high correlation between relative consumption and the real exchange rate — see e.g. Cole and Obstfeld [1991] and Obstfeld and Rogoff [2000], and the discussion in Corsetti, Dedola, and Leduc [2008a].

When markets are incomplete, however, condition (2) no longer determines the equilibrium comovements of relative consumption and the real exchange rate. Once the assumptions generating a high degree of risk sharing are relaxed, without transfers of income via insurance contracts persistent productivity shocks in the tradable sector can create sizeable wedges in relative wealth and demand. A large response of domestic absorption to productivity gains can strengthen the terms of trade, moving the international price of domestic tradables in the same direction as the domestic price of nontradables and reinforce, rather than offset, the HBS effect, thus resulting in a real exchange rate appreciation. Such an outcome would clearly be inefficient from a global perspective — reflecting limited insurance of productivity risk when markets are incomplete.

Risk sharing is equally relevant for understanding relative price movements in response to shocks to the demand for domestic tradables. Relative to the case of productivity disturbances, a demand shift in favor of these goods should increase their price relative to the price of domestic nontradables: other things equal, this would translate into a depreciation of the real exchange rate. The latter must be the outcome if markets are complete (as long as the shock raises domestic relative to foreign consumption, without otherwise shifting marginal utilities): the domestic relative price movement cannot be outweighed by movements in the terms of trade or PT/PT^* . If markets are incomplete, it is possible that positive demand shocks increase the price of domestic tradables relative to the price of foreign tradables to such an extent to result in a real exchange rate appreciation.

Taking stock of the arguments in this subsection, first, it will be important to include in the study measures of *all* the relevant international relative prices, distinguishing between the real exchange rate, the relative price of tradables and the terms of trade (a different behavior between the relative price of tradables and the terms of trade possibly reflecting deviations from the law of one price). Second, while restricting the within-country relative price PN/PT , all international prices should be left unconstrained, so as to characterize in detail the relevant channels of international propagation in the data.

3 The empirical framework

In our analysis, we adopt the VAR methods using sign restrictions similarly to Faust [1998], Uhlig [2005], Canova and De Nicoló [2002] for monetary policy shocks, Dedola and Neri [2007] for technology shocks, and Mountford and Uhlig [2005] for fiscal policy shocks. We

go over the main elements below, starting with the data sample and the VAR specification, and going into the details of the identification.

Our study analyzes the US vis-à-vis an aggregate of the other G7 countries (Japan, Germany, the UK, Italy, Canada, and France) and three other OECD countries (Australia, Sweden, and Ireland) for which we were able to build quarterly data on hourly labor productivity in manufacturing.¹⁹ As in Section 2, we refer to this aggregate as the rest of the world (ROW). All ROW's variables are built as an aggregate of the above mentioned countries (excluding the US), weighted according to their respective (time-varying) GDP shares at PPP values.²⁰

In order to identify shocks which have asymmetric effects across countries, we follow a common approach in the empirical literature: in the tradition of empirical open-economy macroeconomics, we measure all variables except trade and domestic relative prices in terms of cross-country differentials (see Clarida and Galí [1994] and Glick and Rogoff [1995]). As is well understood, the alternative of expanding the empirical system to include both US and ROW variables has the clear disadvantage of running quickly against the constraint imposed by data availability, exhausting any degree of freedom in the empirical analysis. A potential issue in working with cross-country differentials is raised by the (implicit) assumption of symmetry across economic areas — an assumption that is clearly unappealing in studies focused on small open economies. In our case, however, a symmetry assumption is not obviously consequential, as we compare a large country such as the US, with a large aggregate of OECD countries — we will return to this issue in Section 5.

So, we examine the effects of productivity and demand shocks to the US manufacturing sector, on *relative* consumption and investment, trade and capital flows, the real exchange rate, the terms of trade, and asset prices. The sample period is 1973 - 2004, covering the developments in the international monetary system after the collapse of Bretton Woods (and the longest period for which we have data). A detailed description of the data source is in the appendix.

We estimate several specifications of the following reduced-form VAR model (omitting the constant):

$$Y_t = B(L) Y_{t-1} + U_t, \quad (3)$$

where the vector Y includes the n variables of interest in levels and $B(L)$ is a lag polynomial of order p .²¹ The covariance matrix of the vector of reduced-form residuals U_t is denoted

¹⁹These 10 countries add up to roughly half of world GDP at purchasing power parity (PPP) values, so they represent a substantial sample of the global economy. Moreover, trade flows among them also amount to over a half of their respective total trade, on average. For instance, the US trade share with the other nine countries in our sample is around 60 percent of US total trade.

²⁰The GDP shares we use as trade weights were not available for all countries going back to 1973.

²¹See Hamilton [1994], chapter 20.4, for an argument in favor of this approach to model the dynamics of

by Σ . In our specifications (unless stated otherwise), the vector Y_t is 6x1. Following a common practice in open-economy VAR studies, we deal with the curse of dimensionality (due to including too many variables with relatively short samples), by keeping the first five variables in Y_t fixed, while changing the sixth and last variable across specifications.

The first five variables in Y_t are as follows: (i) (the log of) quarterly labor productivity in US manufacturing, in deviation from quarterly labor productivity in manufacturing in the ROW; (ii) the US index of manufacturing production and (iii) aggregate private consumption, both in deviation from the same variable for the ROW; (iv) (the log of) the relative US domestic producer price index over the services consumer price index; and (v) (the log of) real US manufacturing output over US real GDP.

The sixth and last variable in Y_t is, in turn, real private investment in the US relative to the ROW; the ratio of US nominal net export over US nominal GDP and US real imports and exports of goods; the ratio of US external assets and liabilities and their difference over nominal GDP; and three measures of international relative prices:

$$RER(i)_t = \frac{P_t(i)}{P_t^*(i)} \quad i = CPI, PPI \text{ and } Export \text{ Deflator}.$$

The price indexes $P_t(i)$ and $P_t^*(i)$ are alternatively (the log) of the CPI, PPI, and export-deflator in dollars. Note that $P_t^*(i)$ is built as a PPP, GDP-weighted aggregate of prices for the countries included in the ROW. Finally, we also look at the responses of relative equity prices in common currency and short-term interest rate differentials, where the ROW aggregates are computed as above.

Our empirical implementation closely follows Uhlig [2005]. As is well-known, the reduced form (3) can be estimated consistently using ordinary least squares (OLS), which, conditional on Gaussian innovations U_t and initial conditions, amount to maximum-likelihood (ML) estimation.

In the structural VAR literature, identification amounts to providing enough restrictions as to solve uniquely for the following decomposition of the $n \times n$ estimated covariance matrix of the reduced-form VAR residuals Σ (up to an orthonormal transformation Q such that $QQ' = I$):

$$\Sigma = \mathcal{A}_0 \mathcal{A}'_0.$$

This matrix equation defines a one-to-one mapping from the vector of orthogonal *structural* shocks V to the reduced form residuals U , $U = \mathcal{A}_0 V$. Because of the orthogonality assumption, and the symmetry of Σ , at least $\frac{n(n-1)}{2}$ restrictions on \mathcal{A}_0 need to be imposed.

The j -th column of the identification matrix \mathcal{A}_0 , $\mathcal{A}_{0,j}$, is called an *impulse vector* in \mathcal{R}^n ,

a vector of variables some of which may be nonstationary.

as it maps the innovation to the j -th *structural* shock v_j into the contemporaneous impact responses of all the n variables, $\Psi_{0,j}$. With the structural impulse vector $\mathcal{A}_{0,j}$ in hand, the set of all structural impulse responses of the n variables up to the horizon k , $\Psi_{1,j}, \dots, \Psi_{k,j}$ can then be computed using the estimated coefficient matrix $B(L)$ of the reduced-form VAR, B_1, B_2, \dots, B_p :

$$\begin{aligned}\Psi_{s,j} &= \sum_{h=0}^s B_{s-h} \Psi_{h,j}, & s \geq 1, B_{s-h} = 0, s-h \geq p; \\ \Psi_{0,j} &= \mathcal{A}_{0,j}.\end{aligned}$$

Proposition 1 in Uhlig [2005] shows that any structural impulse vector $\mathcal{A}_{0,j}$ arising from a given identifying matrix \mathcal{A}_0 can be represented as Pq , for an appropriate vector q belonging to the hypersphere of unitary radius $\mathcal{S}^n \subset \mathcal{R}^n$, and an arbitrary matrix P such that $PP' = \Sigma$. For instance, natural candidates for the orthogonal decomposition P are either the eigenvalue-eigenvector or the Cholesky decomposition of Σ .

Our procedure to obtain estimates of impulse responses consistent with a given set of assumed sign restrictions can be described as follows. Economic theory can be brought to bear, as in Uhlig [2005] or Dedola and Neri [2007], to attribute all the probability mass to the event that the responses of $m \leq n$ variables (e.g., relative labor productivity, relative output, and so on) to the specific structural shock of interest have a given (positive or negative) sign for $s \leq k$ quarters. For instance, Uhlig [2005] appeals to standard monetary theory and assumes that a contractionary monetary policy shock in the US uniquely brings about a hike in the federal funds rate, a drop in the price level and a contraction in money demand (nonborrowed reserves).

Thus, on the basis of theoretical insights, a-priori a non-zero probability is attributed only to structural impulse vectors $\mathcal{A}_{0,j}$ which, for a given reduced-form estimate of the VAR, yield impulse responses whose signs are consistent with the assumed restrictions. It is important to stress that, while allowing to implement in the identification procedure natural theoretical assumptions (e.g., supply-side shocks should move relative quantities and prices in the opposite direction), sign restrictions lead to a plurality of candidate structural impulse responses. Rather than as a shortcoming, this is a potentially important advantage of this approach, as it allows us to complement our results obtained with exact restrictions, such as long-run restrictions, that could be sensitive to small perturbations to model specification and parameterization.²²

²²For instance, to adapt long-run restrictions to identify manufacturing shocks in an international context, Corsetti et al [2006] had to assume that productivity shocks are exactly nonstationary and the only source of a stochastic trend in both the level of labor productivity, and its differential across countries. This amounts to assume permanent differences in international productivity levels. Most importantly, in contrast with the

As argued by Uhlig [2005], the fact that the Bayesian approach views the VAR parameters as random variables makes it particularly suited to interpreting and implementing sign restrictions. From a Bayesian point of view, on the one hand, the approach amounts to attributing zero probability to reduced-form parameter realizations for which impulse responses contravene the assumed set of sign restrictions. On the other hand, all the impulse responses from the same reduced-form realization that satisfy those restrictions are attributed the same probability a priori. Thus, we can use standard Bayesian methods for estimation and inference, obtaining measures of the uncertainty about estimated impulse responses.

Formally, (i) under a standard diffuse prior on the VAR reduced-form parameters $B(L)$ and Σ , (ii) under a uniform prior on rotation vectors q in S^n yielding impulse responses in accord with restrictions, and (iii) assuming a Gaussian likelihood for the data sample at hand, the posterior density of the reduced-form VAR parameters with the type of restrictions we implement will be proportional to a standard Normal-Wishart — whose parameters are known functions of the OLS-MLE estimates of the VAR reduced form (with the proportionality factor being an indicator function equal to one when at least one rotation q exists yielding impulse responses consistent with the restrictions). This implies that all impulse responses satisfying the restrictions obtained from a given reduced form draw are attributed the same posterior probability, according to the standard Normal-Wishart distribution.

Therefore, it is possible to simulate the posterior distribution of impulse responses consistent with our sign restrictions by jointly drawing from the Normal-Wishart posterior for Σ , $B(L)$ and the uniform for q over S^n , discarding the impulse responses that violate the restrictions. This could be operationalized by using the following algorithm suggested by Uhlig [2005]: for a given estimate of the VAR reduced-form matrices Σ and $B(L)$, and the associated decomposition P , we draw (a large number of) candidate q vectors from a uniform distribution over S^n , and compute the associated impulse vector $A_{0,j}$ and impulse response matrix Ψ , discarding those that do not satisfy the assumed sign restrictions. In practice, the q vectors are drawn from a multivariate standard normal and normalized with their Euclidean norm to make sure they have unitary length.

It should be kept in mind that, as stressed by Uhlig [2005], the sign restriction approach amounts to estimating simultaneously the coefficients of the reduced-form VAR and the impulse vector. Draws of the VAR parameters from their unrestricted posterior which do not admit any impulse vector satisfying the imposed sign restrictions are discarded as they have zero prior weight. As our sign-restriction prior is “informative” only when it attributes a zero probability to a draw, a way to check whether the prior is drastically affecting the

kind of short-run sign-restrictions we use in this paper, long-run restrictions require heavy reliance on unit root tests to buttress key specification assumptions in all the variables included in the VARs.

posterior relative to the standard unrestricted Normal-Wishart is to keep track of these rejections, as we do below.

4 The international dimension of productivity and demand shocks to US tradables

In this section, we present and discuss our empirical findings for productivity and demand shocks identified using the sign restrictions discussed in Section 2, and conveniently summarized by Table 1, for our estimates with all variables in levels. Results from extensive sensitivity analysis are discussed in the next section.

As shown in Table 1, operationally we require our sign restrictions to be in place for 20 quarters — in the case of the relative price of nontradables (proxied by the PPI over the services CPI) the restriction is in place only from the 5th quarter, as to allow for possible short-run effects of nominal rigidities. The choice of a horizon of five years over which the restrictions are imposed reflects the prior that these shocks be mildly persistent — e.g., in the case of an AR(1) with autoregressive coefficient of 0.75, a 1% shock will have all but died out after 20 quarters — but is somewhat arbitrary. Therefore, in our robustness exercises, we have also experimented with restrictions over horizons of 28 and 12 quarters. Only the variables listed in the table are restricted: the responses of all the other variables included in our analysis — US net exports over GDP, relative consumption and investment, and all international relative prices — are left unrestricted.

Table 1 Sign restrictions on VAR variables

<i>Productivity shock</i>		<i>Demand shock</i>	
Variable	Horizon in quarters	Variable	Horizon in quarters
$LP_k - LP_k^* > 0$	$k = 1, \dots, 20$	—	—
$YT_k - YT_k^* > 0$	$k = 1, \dots, 20$	$YT_k - YT_k^* > 0$	$k = 1, \dots, 20$
$YT_k - Y_k > 0$	$k = 1, \dots, 20$	$YT_k - Y_k > 0$	$k = 1, \dots, 20$
$PT_k/PN_k < 0$	$k = 5, \dots, 20$	$PT_k/PN_k > 0$	$k = 5, \dots, 20$

In our experiments, we typically consider 1000 draws from the posterior, and 5000 rotations each. It is worth stressing that the high percentage of accepted reduced form draws for which we find at least one vector q satisfying our restrictions. As shown in Table 2, in our experiments this percentage is in most cases well above 95 percent, thus confirming that our posterior is not radically different from the standard Normal-Wishart.

Table 2

Acceptance rates of reduced forms draws: Individual shocks

Sample is 1974:1-2004:4

	Productivity	Demand
Net exports over GDP	98.90%	98.40%
Real exports	79.90%	95.90%
Real imports	96.50%	95.10%
Relative investment	99.70%	98.70%
CPI-based RER	99.30%	99.10%
PPI-based RER	99.20%	99.20%
Terms of trade	99.10%	98.50%
Exp. Deflator/PPI	90.70%	95.40%
Share prices	91.70%	96.20%
Interest rate differential	98.80%	95.40%
Foreign assets over GDP	96.80%	98.20%
Foreign liabilities over GDP	94.80%	96.80%
Net foreign assets over GDP	99%	97.70%

4.1 Productivity shocks

The impulse response functions to a positive productivity shock for the specification in Table 1 are displayed in Figures 1, 2, 3, and 5. By way of example, Figure 1 displays the response of US relative productivity, manufacturing output ($YT-YT^*$), and aggregate consumption ($C-C^*$), all in log differential with ROW, along with (the log of) manufacturing output over real GDP ($YT-Y$), the (log of the) PPI relative to the services CPI, and nominal net trade over GDP (NX/Y). Each figure reports Bayesian credible intervals, by showing the 16th and 84th percentiles (the dashed lines) together with the median (the solid line) of the posterior distribution of the responses satisfying our restrictions in Table 1 for a productivity shock. Charts depicting restricted variables include vertical lines, marking the horizon over which restrictions are imposed. We discuss our main findings, analyzing each figure in turn.

Consider first the graphs corresponding to the four restricted variables in Figure 1. The median effect of the productivity shock on relative manufacturing output and labor productivity is of the order of 1% and 0.5%, respectively. The effect is quite persistent: notably, the 16th percentile of the productivity responses is above zero well beyond the 20 quarters over which the sign restriction is imposed. The increase of manufacturing output over real GDP is slightly smaller and less persistent. Interestingly, however, the response of (relative) manufacturing output peaks after the first year following the shock. Finally, the productivity shock leads to a prolonged fall in the relative price of domestic tradables already from the second quarter after the shock. This fall corresponds to the HBS effect, that we impose as restriction from the fifth to the 20th quarter, reflecting the conventional wisdom on the relative price implications of productivity gains in manufacturing. It turns out, however, that it is immaterial for our results whether this restriction is imposed from the fifth quarter or from the impact response.²³

Focusing now on the two unrestricted variables shown in Figure 1, the responses of relative consumption and net trade are also very substantial and persistent: the 16th percentile of the responses of these variables remains positive for the entire period (10 years) displayed in the figure. Relative consumption rises on impact and peaks after two years at 0.5%. The fall in net exports is more gradual: the 85th percentile of its responses falls below zero after two quarters; the median deficit gradually reaches -0.1% of GDP after four years.²⁴

²³The immediate drop in this relative prices makes us confident that our estimates are not contaminated by ‘news’ shocks of the kind analyzed by Beaudy and Portier [2005] and Jaimovich and Rebelo [2008], since there is no presumption that anticipated supply-side shifts should move current relative prices and quantities in opposite direction. Interestingly, however, as discussed below, in our identification we find strong wealth and demand effects.

²⁴A potential concern is that our identification scheme only picks measurement error in manufacturing labor productivity, output, and the PPI. If this were the case, however, our overall results would then require this measurement error to be positively correlated with a very persistent increase in relative aggregate consumption (and investment) and a deterioration of net exports — a quite far-fetched set of conditions.

The prolonged fall in net exports shown in Figure 1 may be surprising in light of some applied and policy literature postulating that a productivity increase in tradables should bring about an improvement in net trade. To investigate the source of the trade deterioration, Figure 2 reproduces the responses of relative consumption ($C-C^*$) and net trade, along with relative private investment ($I-I^*$), the CPI-based real exchange rate (RER), and the response of real imports and real exports. Note that all the variables in this figure are unrestricted.

Two findings are worth stressing. First, the deterioration of the trade balance is essentially driven by a rise in real imports: the response of real exports is inconclusive. Therefore, the deterioration in net trade is seemingly consistent with an increase in US absorption driven by a productivity shock, as predicted by the standard intertemporal-trade approach to the current account. Indeed, relative investment rises persistently, mimicking the response of consumption, though it reacts more strongly. Relative investment peaks at around 2 percent, and reaches back its previous baseline level after six years. Second, the CPI-based RER persistently appreciates (an increase is an appreciation) in the aftermath of the shock, then reverts to its baseline, showing some signs of long-run depreciation.²⁵

Because of the positive response of relative consumption, the appreciation of the CPI-based RER is at odds with standard conditions for perfect consumption risk sharing — but it is consistent with the unconditional evidence in Backus and Smith [1993]. This result is of crucial importance in light of the observation, often made in the literature, that a positive correlation between relative consumption and real appreciation in the data could be reconciled with the risk-sharing condition (2) to the extent that taste shocks weaken the link between relative marginal utility and consumption. Our results document that the risk-sharing condition (2) still fails to hold when measured conditional on productivity shocks only.²⁶

The determinants of the response of the real exchange rate are further investigated in Figure 3. This figure shows the response of three alternative measures of international relative prices, based on the CPI, the PPI, and the export deflator, respectively — the latter denoted as terms of trade (TOT) — together with the response of the export deflator for goods relative to the domestic manufacturing PPI ($EXPDEF/PPI$). Remarkably, our three measures of international relative prices display the same appreciation pattern. As two of our measures are built using PPIs (i.e. price indexes including a larger share of tradable goods than the CPI) and export deflators (including only the price of traded goods), our results suggest that the CPI-based RER appreciation reflects more than the classical HBS

²⁵The sign-restriction methodology is also adopted by Enders et al. [2008] for the study of the US real exchange rate dynamics, leading to results for this variable which — accounting for differences in identification — are broadly in line with ours.

²⁶Here we greatly expand on Corsetti et al. [2006], whereas they only looked at the effects technology shocks on relative consumption, net trade and relative prices.

effect (i.e., a rise in the price of nontradables), in line with the unconditional evidence in Engel [1999]. As shown in the figure, real appreciation also reflects substantial fluctuations in the relative price of US tradables relative to ROW tradables. In terms of our discussion in Section 2.2, both PT/PT^* and the terms of trade substantially contribute to the observed real appreciation, in stark departure from allocation efficiency. Over time, however, US international prices tend to depreciate.

Another intriguing result is that the price of exports in domestic currency appears to fall relative to domestic tradables: the 85th percentile of the responses of the US export deflator relative to the PPI is negative for almost 30 quarters. This is consistent with the vast body of evidence on deviations from the law of one price and destination-specific markup adjustment (see, e.g., the survey by Goldberg and Knetter [1997]). In fact, vis-à-vis the real appreciation of the currency, a fall in the export deflator implies that the US export price in foreign currency adjusts by less than one-to-one with the exchange rate (see e.g. Atkeson and Burstein [2008] for recent analyses).²⁷

To complete our analysis of the response of international prices, Figure 4 shows the probability of RER appreciation and of a TOT improvement, calculated as the frequency at which the impulse responses record a strengthening of these relative prices. We find that the probabilities that the real exchange rate appreciates and that the terms of trade improve are 80 percent or higher for roughly five years, before gradually declining.²⁸

Overall, the evidence displayed in Figures 1 through 4 is consistent with international business cycle models with incomplete markets, featuring suboptimal cross-country distribution of wealth and demand in response to technology shocks. An instance of model specifications with these characteristics can be found in Corsetti, Dedola, and Leduc [2008a]: when productivity shocks are persistent, and the long-run price elasticity of exports is high enough, the demand for tradables rises above supply in the short run and appreciates the price of domestic tradable goods relative to foreign ones. As investment raises the capital stock, output rises over time, reversing the movements in relative prices.²⁹ Another instance

²⁷A different (but possibly complementary) interpretation draws on recent literature in international trade and open economy macroeconomics stressing heterogeneity in productivity among firms (see, e.g., Melitz [2003], Ghironi and Melitz [2005], and Eaton and Kortum [2002]). A fall in the export deflator relative to the PPI would obtain only if the most productive firms export and/or productivity gains are stronger among exporters. According to this interpretation, the result in the figure would be driven by changes in marginal costs, rather than markup adjustment.

²⁸Likewise, we computed the joint probability of a RER appreciation and an increase in $C-C^*$, quarter by quarter, finding the probability of such an event to be equal to that of a RER appreciation.

²⁹This paper highlights the crucial role of a high trade elasticity as a precondition for this dynamic response. A high elasticity contains the adverse movements in the price of Home tradables when their quantity rises. This means that, other things equal, the present discounted value of future output is higher, as is the increase in wealth. Note, however, that according to our findings, short-run TOT and RER volatility appear more consistent with relatively low short-run elasticities, raising interesting issues in possible differences in trade elasticities over different horizons.

is provided by Ghironi and Melitz [2005], emphasizing adjustment at the extensive margin of trade. In this model, in the absence of cross-country insurance, positive technology shocks necessarily raise the terms of trade (or ‘terms of labor’) of the country experiencing it (due to endogenous wealth effects limiting the response of domestic labor) to re-equilibrate the entry of new domestic firm. Higher domestic consumption is therefore associated with stronger terms of trade and higher average prices.

An interpretation of our evidence stressing the lack of consumption risk sharing vis-à-vis productivity shocks begs the question of whether this is so because financial markets fail to provide US and foreign residents with good opportunities to hedge against the macroeconomic risks implied by such a shock. This point is addressed by Figure 5. Notably, this Figure shows that the US stock market increases relative to an aggregate index of foreign markets: the relative dollar performance of the US stock market — with a differential between 2 and 7 percent — is well beyond the estimated rate of real dollar appreciation — which is below 2 percent. Potentially, international diversification in equities could provide hedge against the macroeconomic risks implied by our identified productivity shocks.

Most interestingly, US international portfolio positions do react to the shock. Figure 5 shows that asymmetric positive productivity disturbances worsen the net foreign asset position of the US — the results in the figure are based on the times series of valuation-adjusted external assets and liabilities calculated by Gourinchas and Rey [2007]. So, the widening of the trade deficit discussed earlier is matched by an overall deterioration of the current account, defined as the difference in net foreign assets between two points in time.

However, the fall in the US net foreign wealth corresponds to an increase in the stock of both US gross liabilities and US gross assets (a result also shown, following a different methodology, by Corsetti and Konstantinou [2008]). The rise in the stock of (gross) foreign assets is somewhat surprising. Since the bulk of US foreign assets are denominated in foreign currency, other things equal, the dollar appreciation documented above tends to reduce the value of their outstanding stock. The observed increase in the value of foreign assets must then be driven by some combination of capital gains in foreign currency, and positive purchases by US residents. This evidence on the conditional comovements of foreign assets and liabilities provides an intriguing empirical benchmark for the recent literature encompassing portfolio diversification in general equilibrium dynamic models after Devereux and Sutherland [2007].

Also shown in Figure 5 is the response of the short-term interest rate differential, initially zero but then positive (in favor of the US) over time, peaking 10 to 15 quarters after the shock.³⁰ The positive interest rate differential emerges as the relative domestic demand boom

³⁰Comparing this response with that of the real exchange rate, a widening interest differential coexists with a stably appreciated exchange rate, suggesting deviations from uncovered interest parity.

in response to persistent productivity shocks drives the trade into deficit and appreciates the currency in real terms. The estimated response of interest rates is consistent with an endogenously countercyclical stance of monetary policy, leaning against the wind of the perceived demand expansion.³¹

4.2 Demand shocks

Figures 6 through 8 display the impulse response functions to a positive demand shock for our benchmark specification, with the same format as Figures 1 to 3 above — each figure reports the 16th and 84th percentiles (the dashed lines) together with the median (the solid line) of the posterior distribution of the responses satisfying our restrictions in Table 1 for a demand shock. As before, vertical lines mark the horizon of restrictions.

Consider Figure 6, presenting the same variables as in Figure 1 — US relative productivity, manufacturing output ($YT-YT^*$), and aggregate consumption ($C-C^*$), all in log differential with ROW, along with (the log of) manufacturing output over real GDP ($YT-Y$), the (log of the) PPI relative to the services CPI, and nominal net trade over GDP (NX/Y). Comparing Figures 1 and 6 makes it clear that, while the estimated effects of demand shocks on the domestic relative price of US manufacturing are similar to those of productivity shocks, the responses of $YT-YT^*$ and $YT-Y$ are smaller and less persistent; relative labor productivity in manufacturing (now an unrestricted variable) only rises, slightly, on impact — a natural interpretation points to short-run variations in capacity utilization corresponding to a demand-driven rise in production. The credible intervals for consumption and net exports always include zero.

Aggregate effects of the shock can instead be detected for variables such as investment and international relative prices. In Figure 7, the 16th percentile of investment is positive between the 15th and the 25th quarter. Yet, the movement in investment is much smaller than its counterpart after a productivity shock. The CPI-based RER mirrors the pattern of investment. It does not respond in the first few quarters, but clearly appreciates after three years, around the peak of the investment response.

In light of the positive response to investment and the appreciation of the real exchange rate, the lack of response of the trade balance is surprising, especially from the perspective of the MFD framework. However, while real imports do not move appreciably, exports in real terms do tend to fall moderately in the aftermath of the exchange rate appreciation:

³¹As shown by Chari, Kehoe and McGrattan [2002], even allowing for a high degree of price stickiness, monetary policy cannot by itself generate violations of the risk sharing condition in economies in which the natural rate allocation has consumption growth systematically associated with real depreciation. In other words, no currency appreciation is possible — with or without nominal rigidities — in the absence of strong wealth effects mirroring fundamental lack of risk sharing. These effects simultaneously drive both the private sector demand boom, and the policy response to it.

the 85th percentile of the responses of real exports becomes negative after quarter 20. This result squares well with the presumption that the origin of the shock is mostly domestic.

Figure 8 confirms our previous results, that the response of international relative prices has common determinants going beyond and even countervailing the pressure towards real depreciation due to the persistent increase in the price of manufacturing output relative to nontradables. All measures of international prices show the same pattern, appreciating between three and four years after the shock. The response of the price of export goods relative to overall manufacturing is basically zero. Actually, the slight fall in the median of this relative price is again consistent with imperfect exchange rate pass-through vis-à-vis the currency appreciation (a point already discussed at length in relation to productivity shocks).

Figure 9 shows that demand shocks appear to have little impact on relative US stock prices and interest rates. While net foreign assets — once again using valuation-corrected data — remain utterly unaffected, we detect a small increase in the stocks of both gross assets and liabilities. For the former, the 16th percentile of the responses rises above zero between the 5th and the 15th quarters.

5 Sensitivity analysis

In this section, we report on the sensitivity of our results along a number of dimensions. First, we vary the horizon over which sign restrictions are imposed, adding or subtracting eight quarters relative to the baseline specification in Table 1. Second, we study both shocks simultaneously, imposing that they are orthogonal to each others. Third, we address the symmetry assumptions implicit in defining variable as differentials between the US and the rest of the sample. Fourth, we test variants of our sign restrictions and alternative specifications of the model. As a by-product of the analysis above, we experiment with a large set of alternative specifications of the model. For instance, we include different measures of international relative price — using a PPI- instead of a CPI-based real exchange rate. None of these alternative specifications has any significant impact on our benchmark results reported above.

Varying the restriction horizon All our results are broadly unchanged when we re-estimate the VAR model imposing our set of restrictions over a horizon that is shorter than in our baseline specification. In the case of productivity shocks, for instance, the main detectable effects of shortening the (upper) restriction horizon from 20 to 12 quarters (while exactly replicating the analysis in all other dimensions) is that the response of all variables to the productivity shock becomes less persistent.

By the same token, imposing that our restrictions bind for up to 28 quarters, i.e., eight quarters more than our baseline case, increases the persistence of the estimated effects of productivity and demand shocks, as expected. However, lengthening the restriction horizon makes it more difficult to find productivity shocks in the data — we end up rejecting a somewhat larger fraction of draws from the reduced-form posterior of our VAR. Similar considerations apply to the case of demand shocks. To save space, we do not report figures for these exercises.

Orthogonal productivity and demand shocks In our analysis, we require the two shocks to have opposite effects on the U.S. relative price of tradables vis-à-vis CPI services, while always increasing U.S. manufacturing output relative to real GDP. In this respect, our identification strategy limits the risk that our results confound their effects. Yet, as we identify productivity and demand shocks individually, without requiring that both be present in the data and be orthogonal to each other, our estimated effects of either shock could potentially be biased — for instance, if the two shocks happen to be negatively correlated with each other. Specifically, in the case of the demand shock, it may happen that some of the estimated responses, while picking up the assumed *positive* demand shock, could be contaminated by a (weaker) *negative* productivity shock. For instance, this could have the effect of strengthening the positive response of the relative price of tradables, while attenuating that of other variables, like relative output or consumption.

The change in the procedure described in Section 3 needed to produce two sets of candidate impulse responses which are orthogonal on impact is straightforward. We now need to find two vectors, q_1 and q_2 , both belonging to the hypersphere of unitary radius $\mathcal{S}^n \subset \mathcal{R}^n$, which also satisfy the additional orthogonality condition $q_1'q_2 = 0$; then, we can compute the two impulse vectors Pq_1 and Pq_2 and the related impulse responses, verifying that they satisfy the sign restrictions for productivity and demand shocks. In practice, the vectors are again drawn from a multivariate standard normal, then orthogonalized and normalized with their Euclidean norm to make sure they have unitary length.

Orthogonality may have a further important consequence for our results. Since the number of restrictions imposed simultaneously is larger, it may be more difficult to find the two shocks in the data, leading us to reject a much higher number of reduced forms in our estimation procedure. This could affect our findings, that each shock individually is very likely to be present in the data. Yet, when we estimate orthogonal productivity and demand shocks imposing the restrictions in the two columns of Table 1 simultaneously, we find virtually no difference from the results for each shock in isolation. Once again, to save space, we do not report figures for these exercises.

Cross-country symmetry An important robustness check consists in testing whether the symmetry assumption implicit in estimating VARs in cross-country differentials be a source of bias in our results. To address this specification issue we first investigated potential sources of bias by regressing the residuals from our baseline six-variable specification of the VAR model on the levels of all our variables, for both the U.S. and the ROW. Second, we re-estimate our model using a larger VAR, in which we include the variable(s) that turn out to be significant in the auxiliary regressions estimated in the first step.

Interestingly, the only such variable in the auxiliary regressions is productivity in manufacturing in the U.S. — suggesting that its specification in the benchmark VAR in differential with the ROW could be a potential source of bias. Thus we re-estimate a seven-variable VAR, including this variable, left unconstrained, in order to keep our identification unaltered. The results from this exercise are broadly similar to those from our benchmark. This is clearly shown by Figure 10, which displays a selection of our variables, using the same format as before — each chart displays the 16th and 84th percentiles (the dashed lines) together with the median (the solid line) of the posterior distribution of the responses satisfying the restrictions. Notably, the level of U.S. labor productivity in manufacturing, which is left unconstrained, rises persistently with a high probability, closely mirroring the response of the (constrained) productivity differential; all the other variables, both restricted (YT-YT*, YT-Y and PPI/CPI) and unrestricted, behave in a similar way as before under the symmetric specification. Overall, this is evidence that the symmetry assumption does not appear to drive our findings in a significant way.

Changing sign restrictions As seen above, the response of many macro variables to demand shocks identified according to our benchmark scheme tends to be subdued. As a robustness exercise, we check the implications of imposing the additional restriction that relative consumption increases in response to demand disturbances for eight quarters. The goal of this experiment is to verify whether adding a restriction, aiming at isolating domestic demand shocks through their effects on a large component of absorption, would provide more conclusive results for a larger set of variables.

We find that, when demand shocks are associated with an increase in relative consumption, they tend to have stronger effects on trade variables, as the response of gross and net exports becomes clearly negative. However, these shocks have weaker effects on investment and international relative prices, as shown in Figures 11 and 12, reproducing the responses of the same variables as in Figures 6 and 7, respectively.

Specifically, the estimated effects of demand shocks on relative consumption (restricted in Figure 11 to increase for the first two years) and net exports are much more clear cut, with the latter turning negative for some periods. Figure 12 shows that the trade deficit

mainly reflects a persistent drop in real exports, while the response of real imports appears inconclusive. Relative to Figure 7, however, the response of investment and the CPI-based RER is much smaller and also inconclusive.

This robustness experiment shows that the demand shock identified with further restrictions, rather than leading to more conclusive results across the board, has quite different effects relative to our benchmark. This supports the notion that the two sets of restrictions are isolating two distinct shocks to U.S. manufacturing, whose aggregate effects materialize mainly in international relative prices and investment in the case of our benchmark model — in consumption and trade in the other case.

6 Conclusions

In this paper, we provide empirical evidence on the international dimension of productivity and demand shocks in the U.S., vis-à-vis an aggregate of other large industrial countries. First, we find that sectoral real demand shocks in US manufacturing, while having limited aggregate effects on absorption and net trade, still affect specific demand components, especially investment, and appreciate the real value of the dollar. Second, sectoral productivity shocks in US tradables have a non-trivial impact on aggregate demand across countries, suggesting that an important dimension of their transmission operates via general equilibrium effects on cross-country wealth.

Our findings point to endogenous demand and wealth effects as a fundamental aspect of international interdependence in the data. They provide evidence at odds with a key transmission channel typically postulated by standard textbook open-economy models — that a higher supply of tradables must be matched by a fall in a country’s terms of trade —, and have important implications for the efficiency of global allocations. As pointed out by Cole and Obstfeld [1991], a fall in international relative prices of domestic tradables partly offsetting productivity and output differentials would provide consumption risk insurance by containing suboptimal differences in national demand and wealth: international price movements would substitute for asset income payments from internationally diversified portfolios. In contrast, our estimates suggest that, at least over the business cycle, the movements in the relative prices of US traded goods are far from providing risk insurance opportunities against US productivity fluctuations, while movements in relative prices of assets such as stocks apparently do.

Specifically, our results complement the findings in Backus and Smith [1993] that domestic consumption, on average, rises when the currency appreciates, by showing that this feature of the data also emerges conditional on identified productivity shocks only. As stressed by Obstfeld and Rogoff [2001], “consumption correlations puzzles” are key to as-

sessing alternative general equilibrium models. Our evidence not only clearly points to lack of consumption risk sharing vis-à-vis productivity disturbances; together with the rest of our analysis, our evidence also emphasizes imperfect risk sharing as an essential element in the characterization of the transmission mechanism, and suggests that asymmetric wealth effects play an important role in generating aggregate demand fluctuations across countries.

In this sense, our findings provide an empirical benchmark for the development of open-economy models and a base for further theoretical and empirical investigations. Open issues, for instance, include the extent to which our conditional evidence could be reconciled with recent dynamic models featuring endogenous portfolio diversification, as well as with quantitative models allowing for goods quality and/or variety adjustment over the business cycle.

References

- [1] Atkeson, Andrew, and Ariel Burstein, [2008]. “Pricing-to-Market in a Ricardian Model of International Trade,” *American Economic Review*, forthcoming.
- [2] Backus, David K., Patrick J. Kehoe, and Finn E. Kydland [1994]. “Dynamics of the Trade Balance and the Terms of Trade: The J Curve?” *American Economic Review* 82, pp. 864-88.
- [3] Backus, David K., and Gregor W. Smith, [1993]. “Consumption and Real Exchange Rates in Dynamic Economies with Non-traded Goods,” *Journal of International Economics* 35, pp. 297-316.
- [4] Beaudry, Paul and Franck Portier, [2006]. “Stock Prices, News, and Economic Fluctuations,” *American Economic Review*, 96, pp. 1293-1307.
- [5] Canova, Fabio, and Gianni De Nicoló [2002]. “Monetary Disturbances Matter for Business Fluctuations in the G-7,” *Journal of Monetary Economics* 49, pp. 1131-59.
- [6] Canova, Fabio and Evi Pappa, [2006]. “The Elusive Costs and the Immaterial Gains of Fiscal Constraints,” *Journal of Public Economics* 90, pp. 1391-1414.
- [7] Canova, Fabio, and Matthias Paustian [2008]. “Measurement with some theory: Using sign restrictions to evaluate business cycle models,” *mimeo*, Universitat Pompeu Fabra.
- [8] Chari, V.V., Patrick J. Kehoe, and Ellen McGrattan [2002]. “Can Sticky Prices Generate Volatile and Persistent Real Exchange Rates?,” *Review of Economic Studies* 69, pp. 633-63.

- [9] Clarida, Richard, and Jordi Galí [1994]. “Sources of Real Exchange Rate Fluctuations: How Important Are Nominal Shocks?” *Carnegie-Rochester Series in Public Policy* 41, pp. 1-56
- [10] Cole, Harold L., and Maurice Obstfeld [1991]. “Commodity Trade and International Risk Sharing: How Much Do Financial Markets Matter?” *Journal of Monetary Economics* 28, pp. 3-24.
- [11] Corsetti, Giancarlo, Philippe Martin, and Paolo Pesenti [2007]. “Productivity, Terms of Trade and the ‘Home Market Effect’,” *Journal of International Economics* 73, pp. 99-127.
- [12] Corsetti Giancarlo, and Panagiotis Konstantinou [2008]. “The Dynamics of US Foreign Borrowing: Evidence on External Adjustment to Transitory and Permanent Shocks,” *mimeo*, European University Institute.
- [13] Corsetti Giancarlo, Luca Dedola and Sylvain Leduc [2006]. “Productivity, External Balance and Exchange Rates: Evidence on the Transmission Mechanism Among G7 Countries,” *NBER International Studies on Macroeconomics*.
- [14] Corsetti Giancarlo, Luca Dedola and Sylvain Leduc [2008a]. “International Risk-Sharing and the Transmission of Productivity Shocks,” *Review of Economic Studies* 75, pp. 443-473.
- [15] Corsetti Giancarlo, Luca Dedola and Sylvain Leduc [2008b]. “High exchange rate volatility and low pass-through,” *Journal of Monetary Economics* 55 , pp. 1113–1128.
- [16] Dedola, Luca and Stefano Neri [2007]. "What does a technology shock do: A VAR analysis with model-based sign restrictions," *Journal of Monetary Economics* 54, pp. 512-549.
- [17] Devereux, Michael B. and Alan Sutherland [2007]. “Country Portfolios in Open Economy Macro Models,” manuscript.
- [18] Dornbusch Rudiger [1980], “Exchange Rate Economics: Where Do We Stand”, *Brookings Papers on Economic Activity* 1, pp. 143-185.
- [19] Eaton, Jonathan and Samuel Kortum [2002]. “Technology, Geography and Trade,” *Econometrica* 70, pp. 1741-79.
- [20] Eichenbaum, Martin and Charles L. Evans [1995]. “Some Empirical Evidence on the Effects of Shocks to Monetary Policy on Exchange Rates,” *Quarterly Journal of Economics* 110, pp. 975-1009.

- [21] Enders Zeno, Gernot J. Müller and Almut Scholl [2008], “How do Fiscal and Technology Shocks affect Real Exchange Rates? New Evidence for the United States”, Center for Financial Studies, Working Paper 2008/22.
- [22] Engel, Charles [1999]. “Accounting for Real Exchange Rate Changes,” *Journal of Political Economy* 107, pp. 507-38.
- [23] Faust, Jon [1998]. “The Robustness of Identified VAR Conclusions about Money,” *Carnegie-Rochester Conference Series on Public Policy* 49, pp. 207-44.
- [24] Faust, Jon and John H. Rogers [2003]. “Monetary policy’s role in exchange rate behavior,” *Journal of Monetary Economics* 50 (7), pp. 1403–1622.
- [25] Gali, Jordi [1999]. “Technology, Employment, and the Business Cycle: Do Technology Shocks Explain Aggregate Fluctuations?” *American Economic Review* 89, pp. 249-71.
- [26] Ghironi, Fabio and Marc Melitz [2005]. “International Trade and Macroeconomics Dynamics with Heterogenous Firms,” *Quarterly Journal of Economics*, pp. 865-915.
- [27] Glick, Reuven, and Kenneth Rogoff [1995]. “Global versus Country-Specific Productivity Shocks and the Current Account,” *Journal of Monetary Economics* 35, pp. 159-92.
- [28] Goldberg, Pinelopi and Michael M. Knetter, [1997]. “Goods Prices and Exchange Rates: What Have We learned?,” *Journal of Economic Literature* 35, pp. 1243-1272.
- [29] Gourinchas, Pierre Olivier, and Helene Rey [2007]. “International Financial Adjustment,” *Journal of Political Economy* 115, pp. 665-703.
- [30] Gravelle, Hugh, and Ray Rees [1992]. *Microeconomics*, London: Longman.
- [31] Hamilton, James D. [1994]. *Time series analysis*, Princeton: Princeton University Press.
- [32] Jaimovich J. and Sergio Rebelo, “Can News about the Future Drive the Business Cycle?,” *American Economic Review*, forthcoming.
- [33] Melitz, Marc [2003]. “The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity,” *Econometrica* 71, pp. 1695-1725.
- [34] Mountford, Andrew [2005]. “Leaning into the Wind: A Structural VAR Investigation of UK Monetary Policy,” *Oxford Bulletin of Economics and Statistics* 67, pp. 597-621.
- [35] Mountford , Andrew, and Harald Uhlig [2005], “What are the effects of fiscal policy shocks? Results from an agnostic identification procedure,” CEPR Discussion Papers 3338.

- [36] Obstfeld, Maurice [1982]. “Aggregate Spending and the Terms of Trade: Is There a Laursen-Metzler Effect,” *Quarterly Journal of Economics* 97, pp. 251-270.
- [37] Obstfeld, Maurice, and Kenneth Rogoff [1995]. “Exchange Rate Dynamics Redux,” *Journal of Political Economy* 103, pp. 624-655.
- [38] Obstfeld, Maurice, and Kenneth Rogoff [2000]. “New Directions for Stochastic Open Economy Models.” *Journal of International Economics* 50, February, pp. 117-153.
- [39] Obstfeld, Maurice, and Kenneth Rogoff [2001]. “The six major puzzles of international macroeconomics: Is there a common cause?” *NBER Macroeconomics Annual*.
- [40] Scholl, Almuth and Harald, Uhlig, [2007]. “New Evidence on the Puzzles: Results from Agnostic Identification on Monetary Policy and Exchange Rates,” *Journal of International Economics* 76, pp. 1-13.
- [41] Stockman, Alan C., and Linda Tesar [1995]. “Tastes and Technology in a Two-Country Model of the Business Cycle: Explaining International Comovements,” *American Economic Review* 83, pp. 473-486.
- [42] Svensson Lars and Assaf Razin [1983], “The Terms of Trade and the Current Account: the Harberger-Laursen-Metzler Effect”, *Journal of Political Economy* 91, pp. 97-125.
- [43] Uhlig, Harald [2005], “What are the effects of monetary policy on output? Results from an agnostic identification procedure,” *Journal of Monetary Economics* 52, pp. 381-419.

Appendix A Data description and sources

United States

Labor productivity: Index of output per hour of all persons in manufacturing sector, seasonally adjusted, 1992 = 100 (Bank of International Settlements and Dept. of Labor).

Manufacturing output: Index of industrial production in manufacturing, seasonally adjusted, 2000 = 100 (Federal Reserve Board)

Consumption: Private final consumption expenditure, volume in national currency, seasonally adjusted (OECD, Economic Outlook Database).

Nominal GDP: Gross domestic product, value, market prices in national currency, seasonally adjusted (OECD, Economic Outlook Database)

Net exports: Nominal net exports of goods & services, value in national currency, seasonally adjusted (OECD, Economic Outlook Database)

Real imports and exports: Real imports and exports of goods, national currency, seasonally adjusted, 2000 = 100 (NIPA, Table 4.2.3)

PPI index: Producer price index of manufactured products, seasonally adjusted, 2000 = 100 (OECD, Main Economic Indicators Database)

CPI total: Consumer price index all items, seasonally adjusted, 2000 = 100 (OECD, Main Economic Indicators Database)

CPI services: Consumer price index for services less energy services, seasonally adjusted; 1982-84 = 100, monthly converted to quarterly averages (BLS)

Export deflator: Exports of goods, deflator, seasonally adjusted, national accounts basis; 2000 = 100 (OECD, Economic Outlook Database)

Short-term rate: Federal Fund Rate, quarterly (IMF, International Financial Statistics)

Stock prices: Share Prices, quarterly (IMF, International Financial Statistics)

US Foreign assets and liabilities: Valuation adjusted US net foreign assets, gross foreign assets and gross foreign liabilities (Gourinchas and Rey [2007], Appendix B)

CPI-based real exchange rate: Index of ratio of US CPI (total) to aggregate CPI (total) of 9 OECD countries, all in current US dollars, weighted with GDP shares at annual PPP values, 1970q1 = 100 (authors calculations based on OECD, Economic Outlook Database)

PPI-based real exchange rate: Index of ratio of US PPI (manufacturing) to aggregate PPI (manufacturing) of 9 OECD countries, all in current US dollars, weighted with GDP shares at annual PPP values, 1971q1 = 100 (authors calculations based on OECD, Economic Outlook Database)

Terms of trade: Index of ratio of US export deflator (goods and services) to aggregate export

deflator (goods and services) of 9 OECD countries, all in current US dollars, weighted with GDP shares at annual PPP values, 1970q1 = 100 (authors calculations based on OECD, Economic Outlook Database)

Relative stock prices: Index of ratio of US stock prices to aggregate stock prices of 9 OECD countries, all in current US dollars, weighted with GDP shares at annual PPP values, 1970q1 = 100 (authors calculations based on IMF, International Financial Statistics)

Rest of the world

The rest of the world comprises Japan, Germany, UK, Italy, France, Canada, Australia, Sweden and Ireland. This choice was dictated by data availability regarding hourly productivity in manufacturing.

Individual country's variables were aggregated by first taking quarterly growth rates to remove national basis effects; then cross-country average growth rates were computed with weights based on each country's GDP share in the 9-country aggregate calculated at annual purchasing power parity (PPP) values. Average growth rates were then cumulated starting from the initial base year to obtain levels.

Annual PPP based GDP shares are from the IMF's World Economic Outlook Database from 1980; before 1980 they were computed directly on the basis of annual GDP at PPP values from OECD's Economic Outlook Database.

Labor productivity: Aggregate of country-specific indexes of output per hour of all persons in manufacturing sector, seasonally adjusted, 1970q1 = 100 (authors calculations based on national statistical sources, BIS and IMF)

Manufacturing output: Aggregate of country-specific indexes of industrial production, manufacturing, seasonally adjusted, 1970q1 = 100 (authors calculations based on national statistical sources, BIS and IMF)

Consumption and investment: Aggregate of country-specific private final consumption expenditure, volumes in national currency, seasonally adjusted, 1970q1 = 100 (authors calculations based on OECD, Economic Outlook Database).

Stock prices: Aggregate of country-specific share prices, 1970q1 = 100 (authors calculations based on IMF, International Financial Statistics and BIS (for UK only)).

Short-term rates: Weighted average with PPP GDP weights of country-specific short-term interest rates (authors calculations based on IMF, International Financial Statistics).

Figure 1

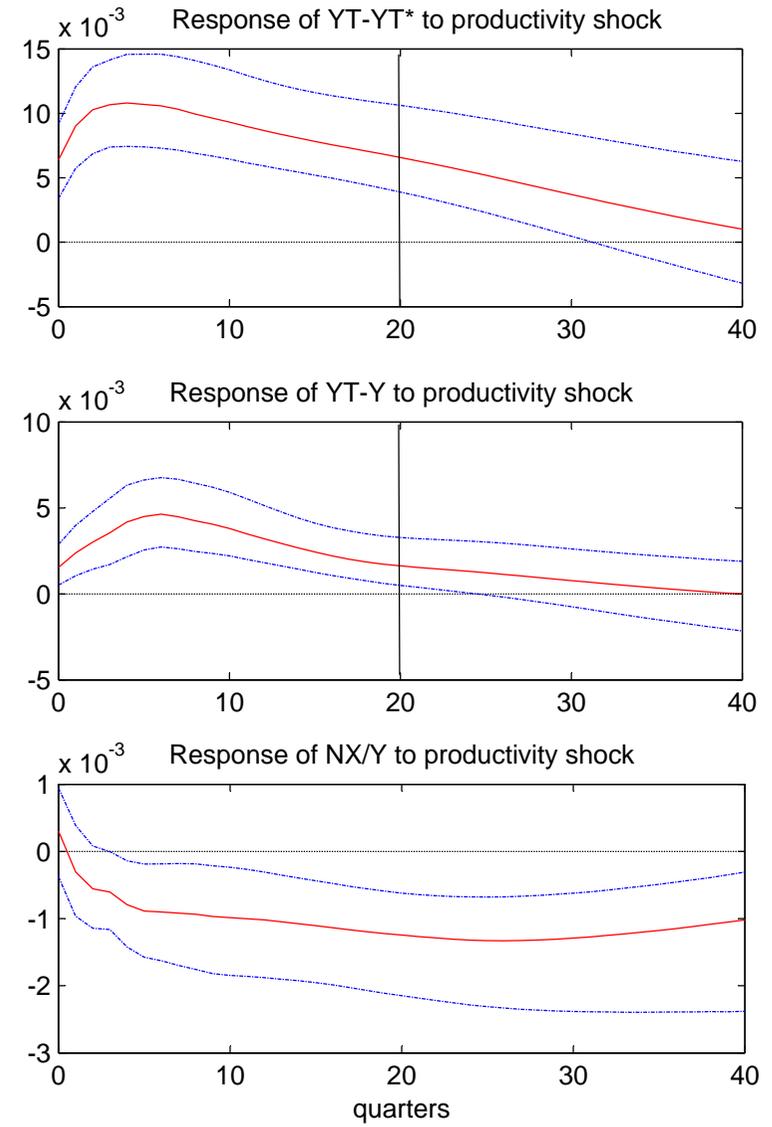
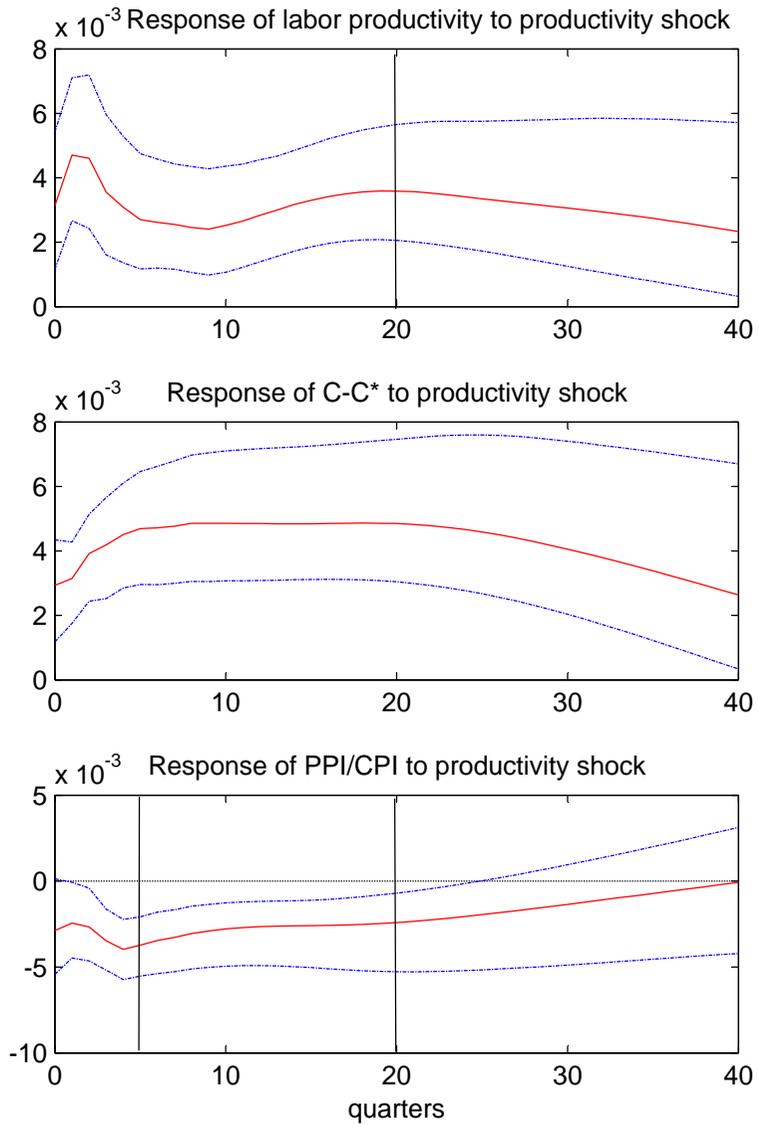


Figure 2

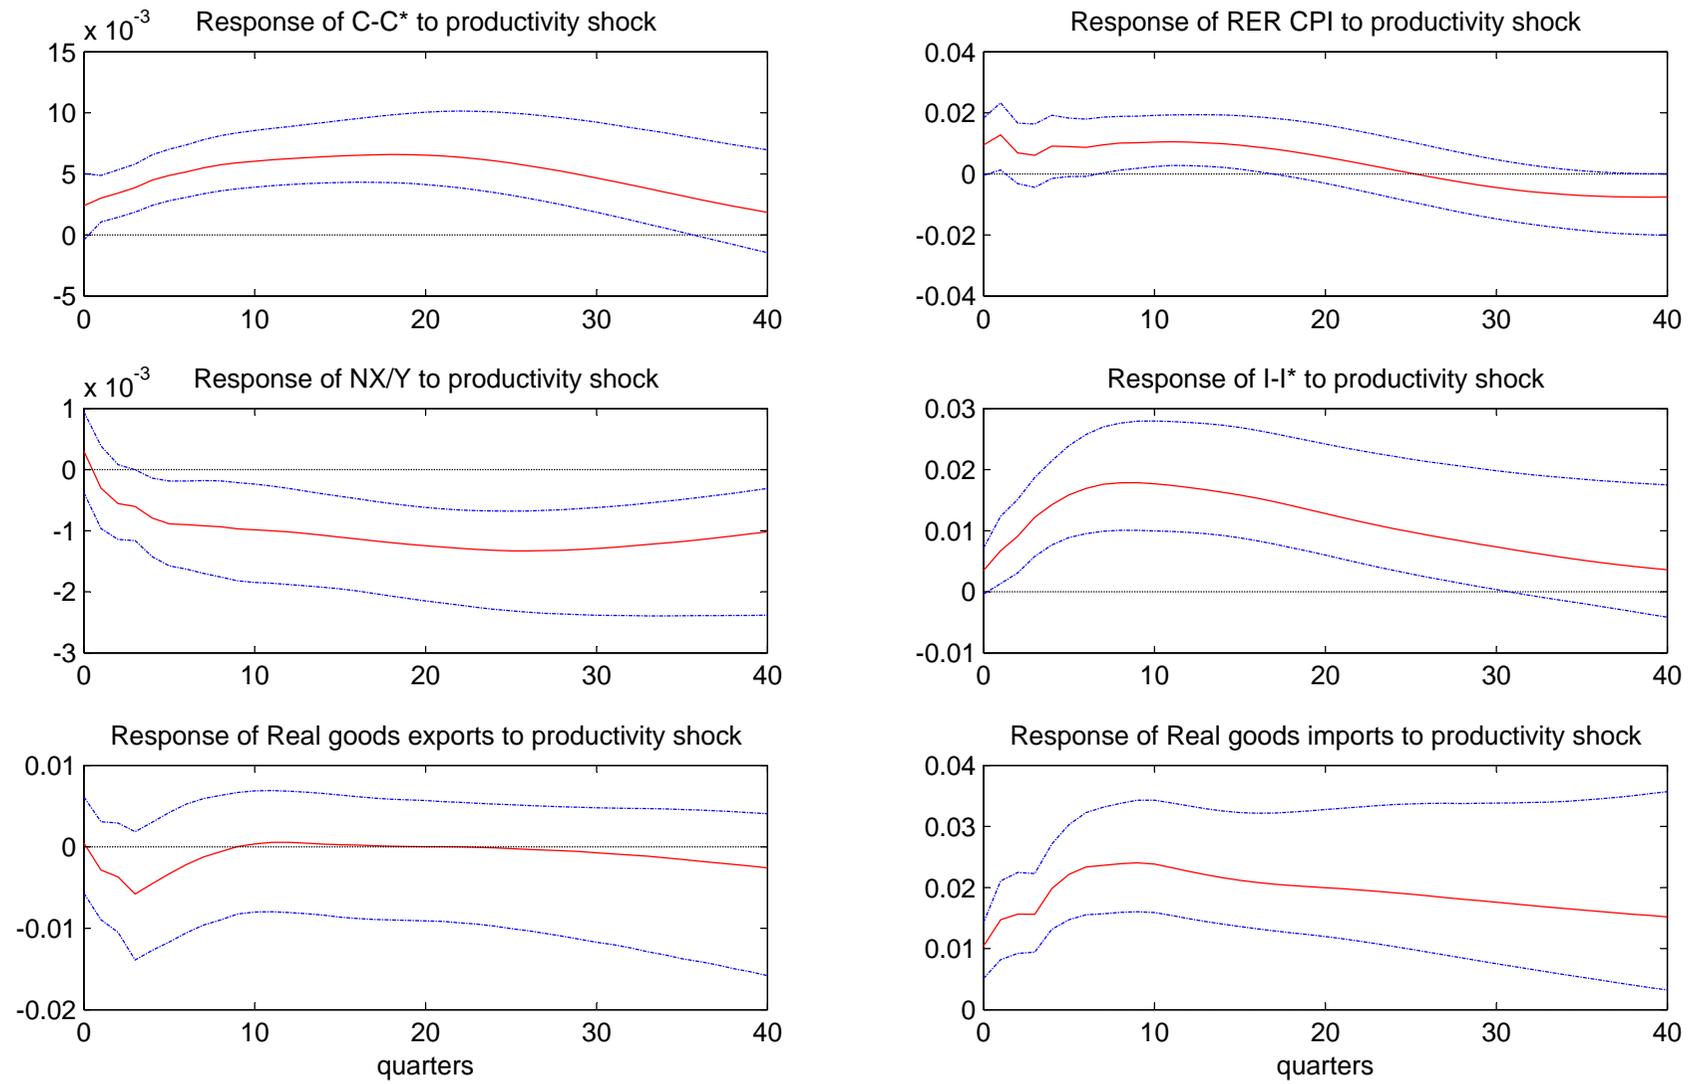


Figure 3

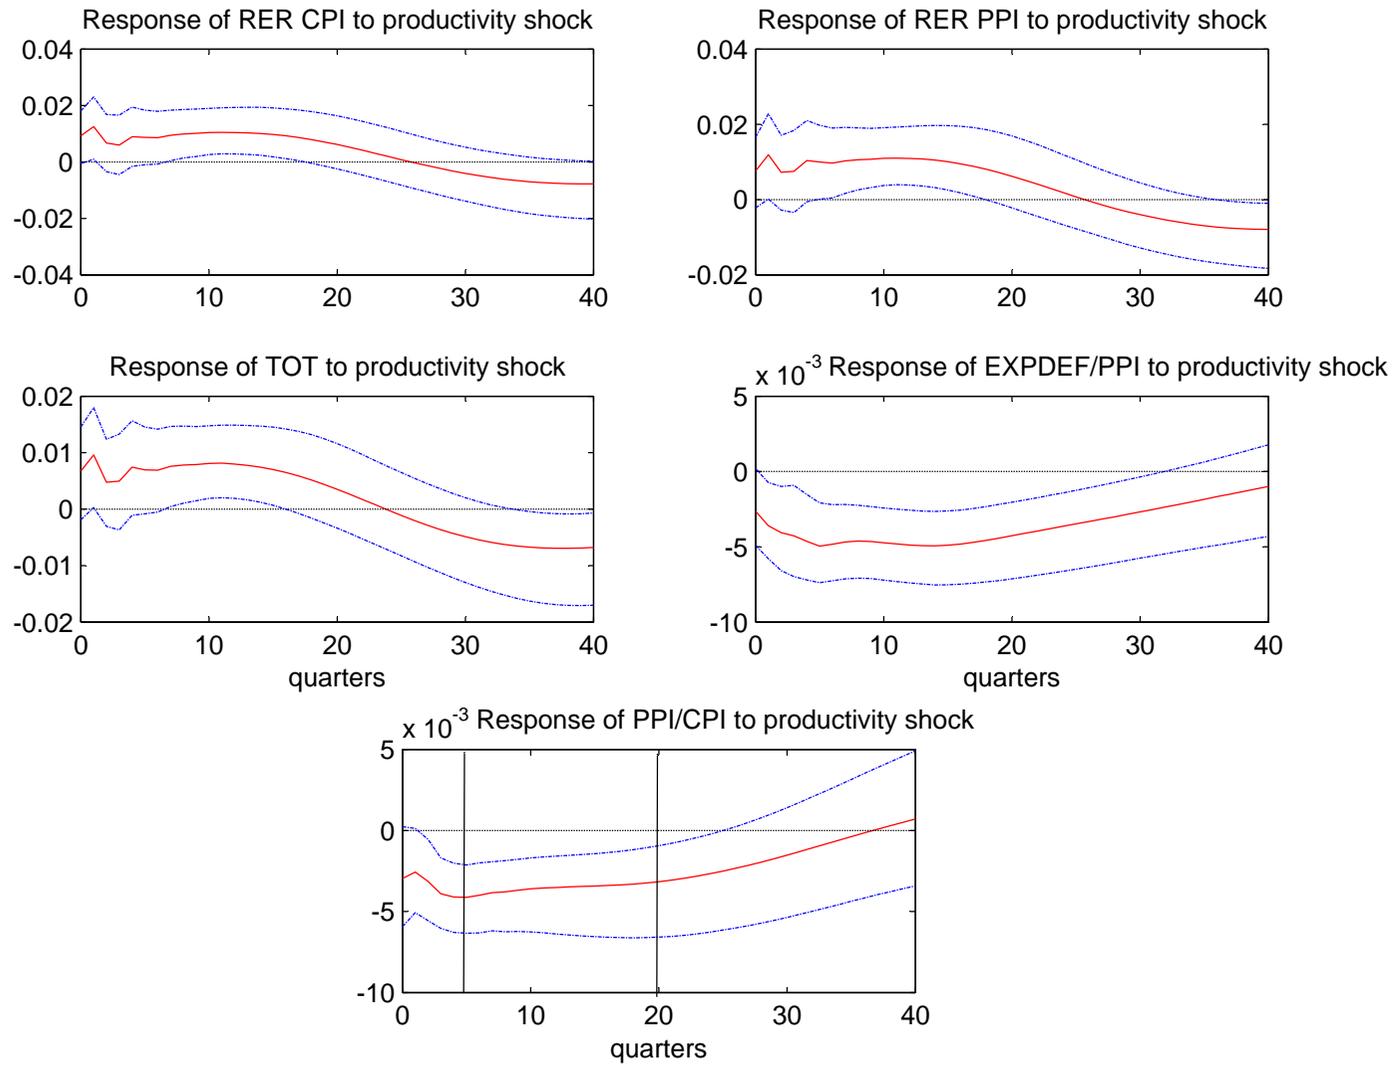


Figure 4

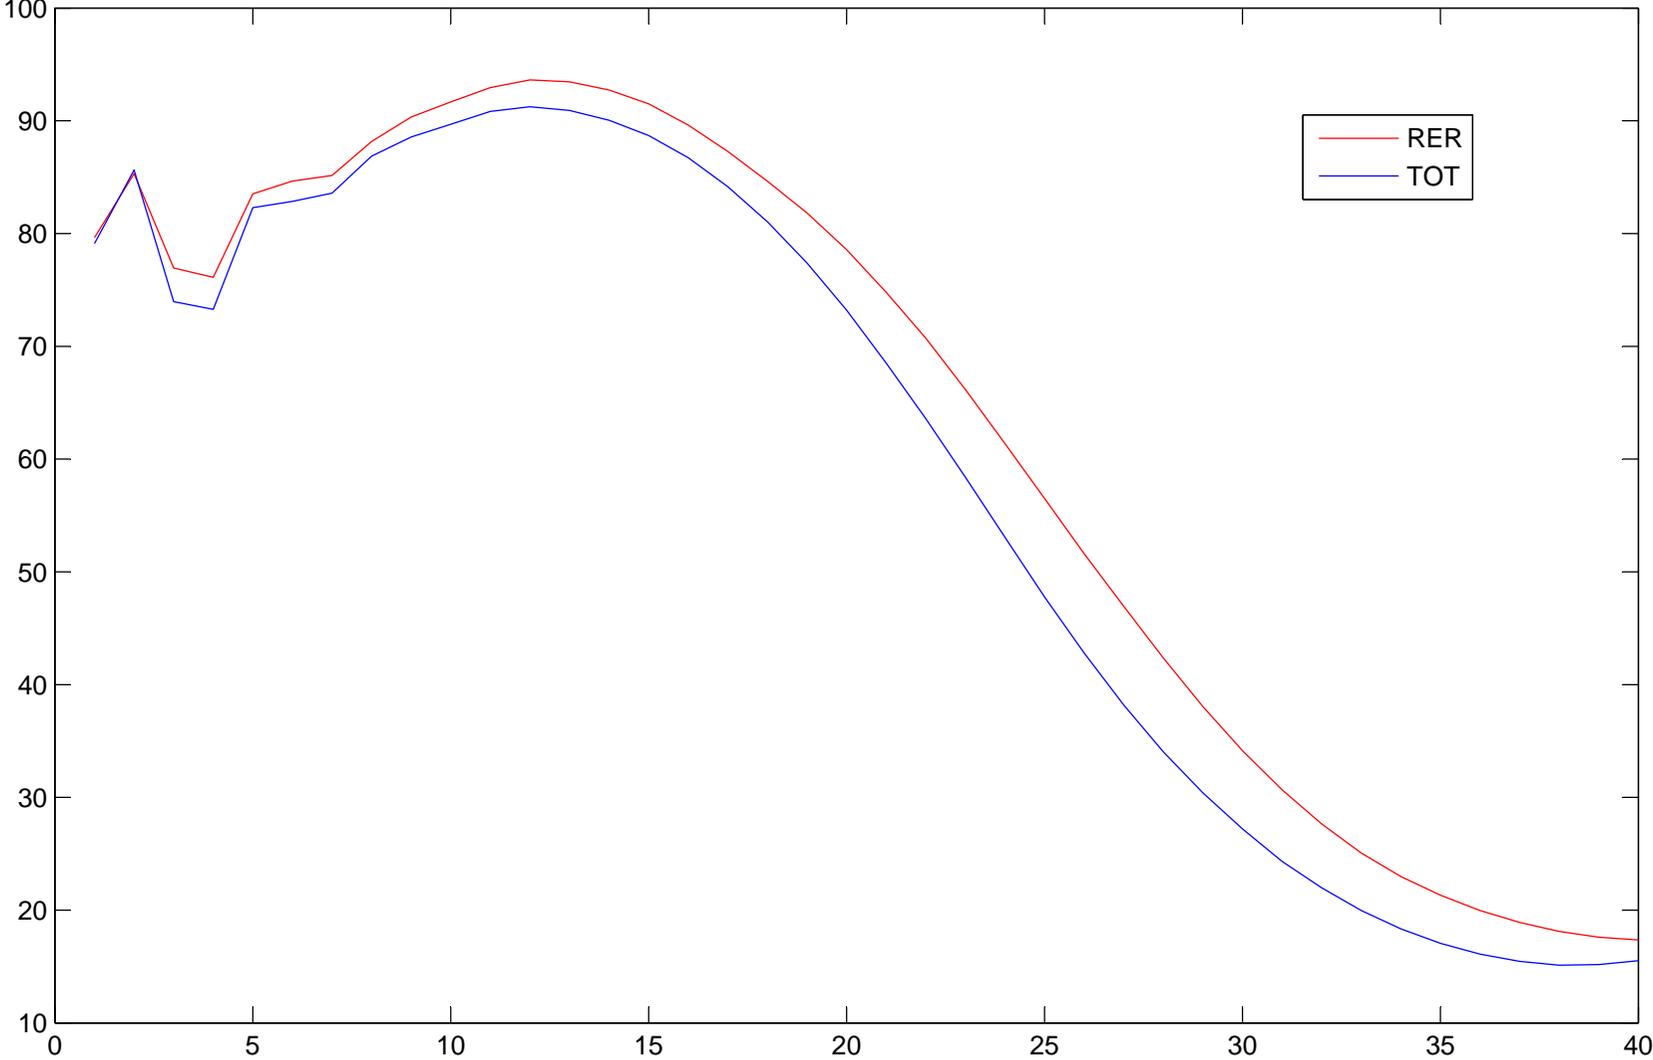


Figure 5

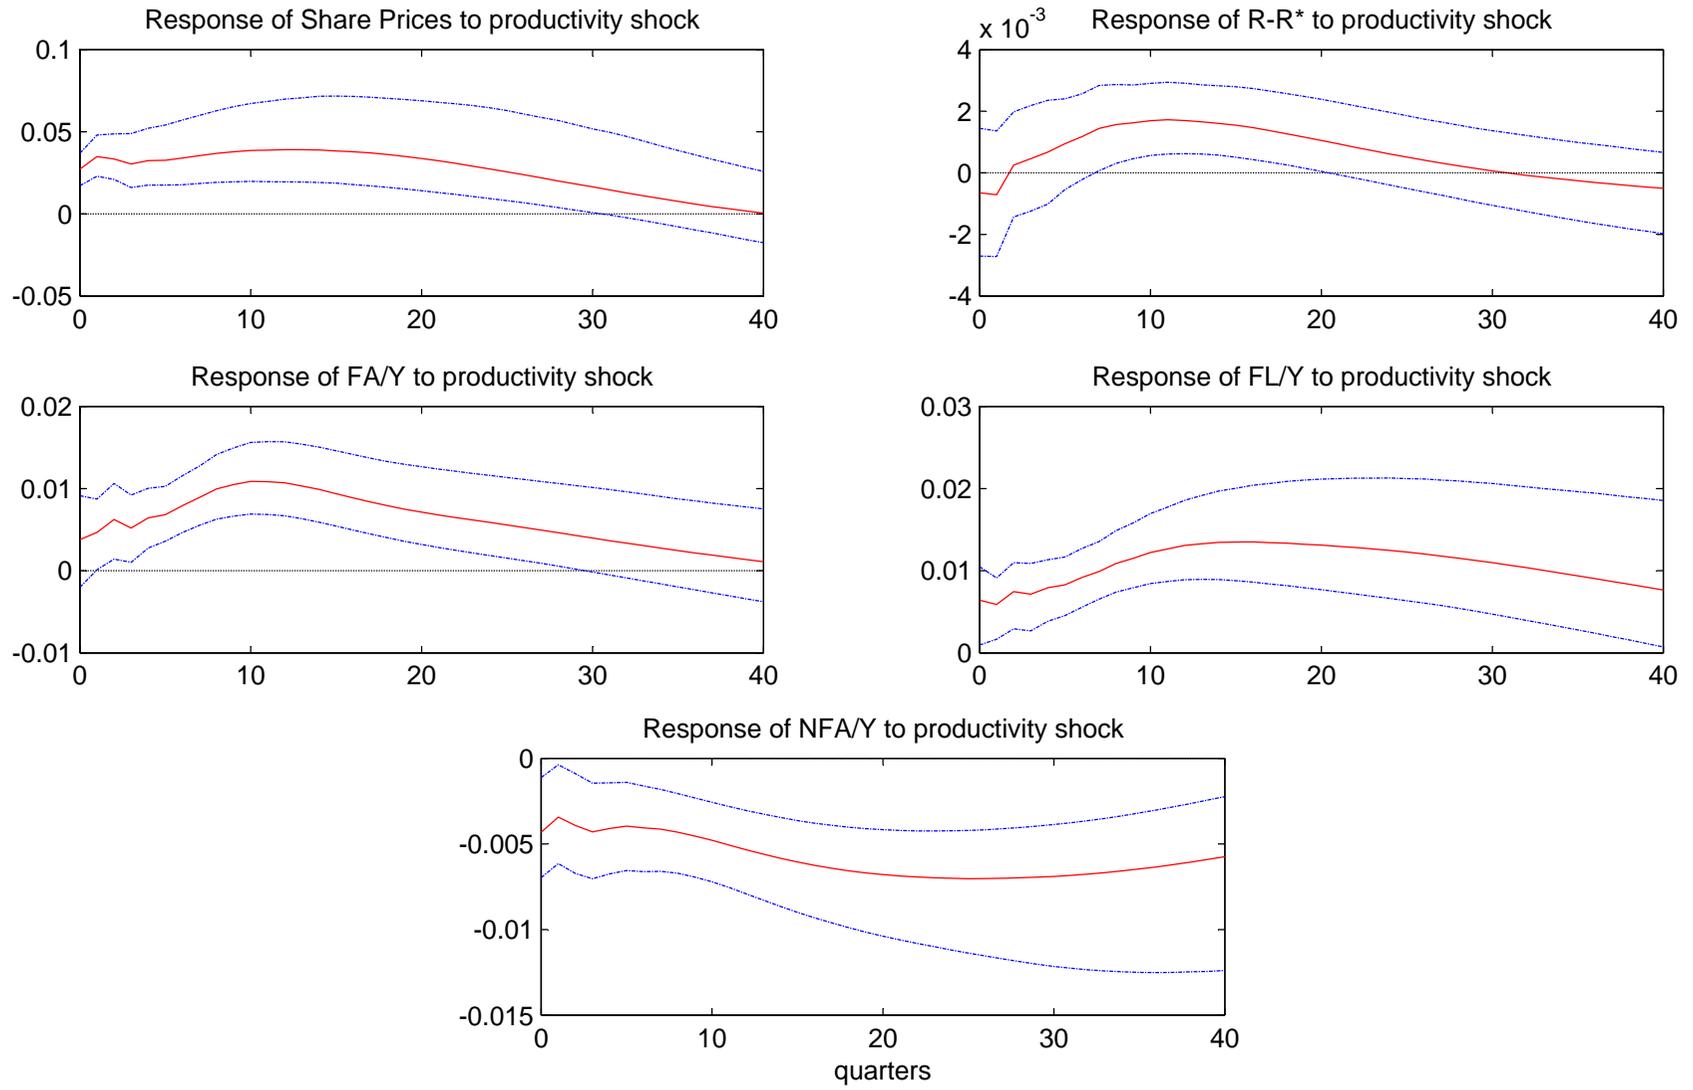


Figure 6

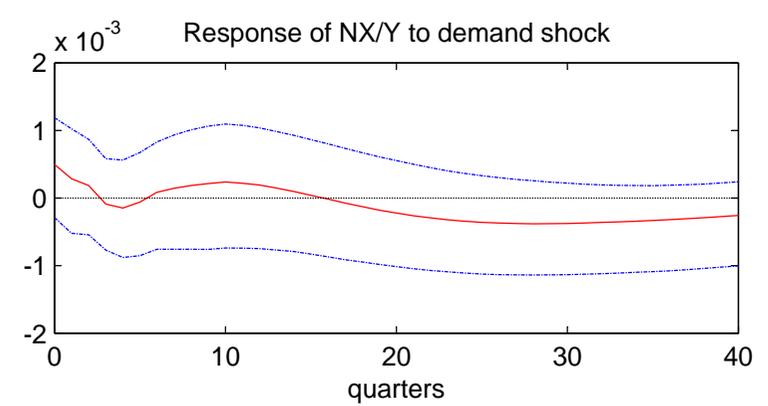
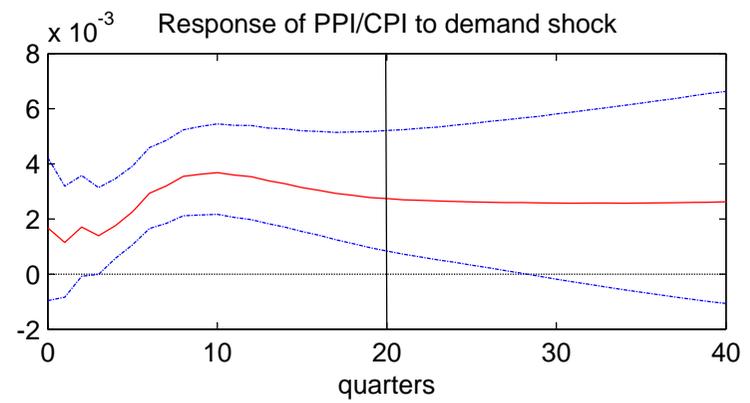
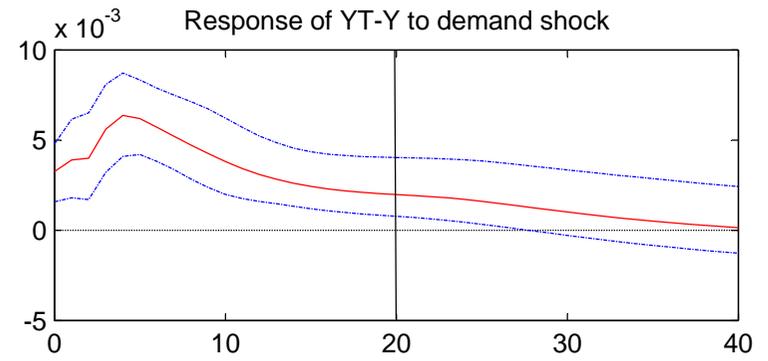
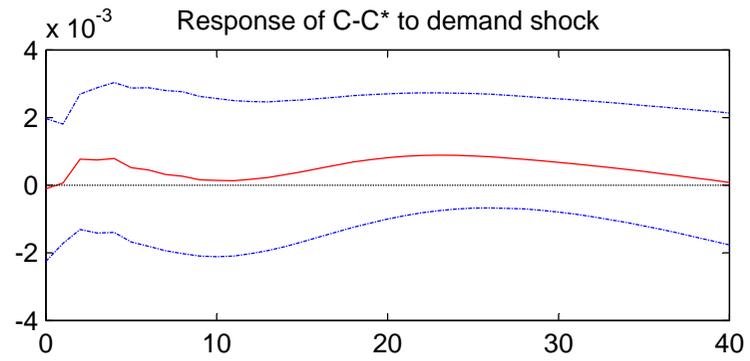
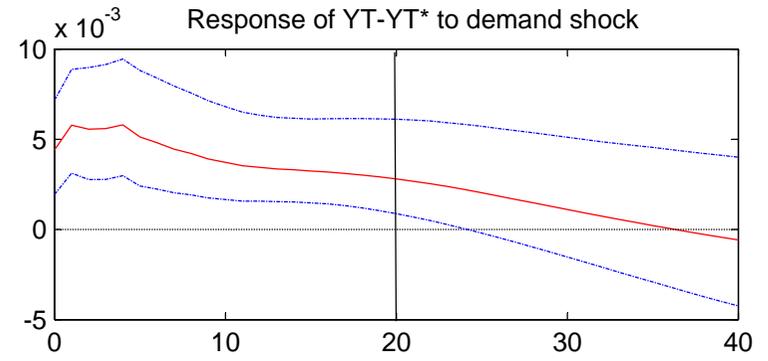
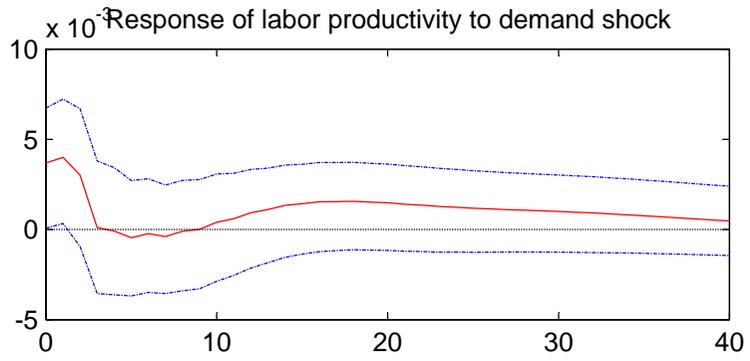


Figure 7

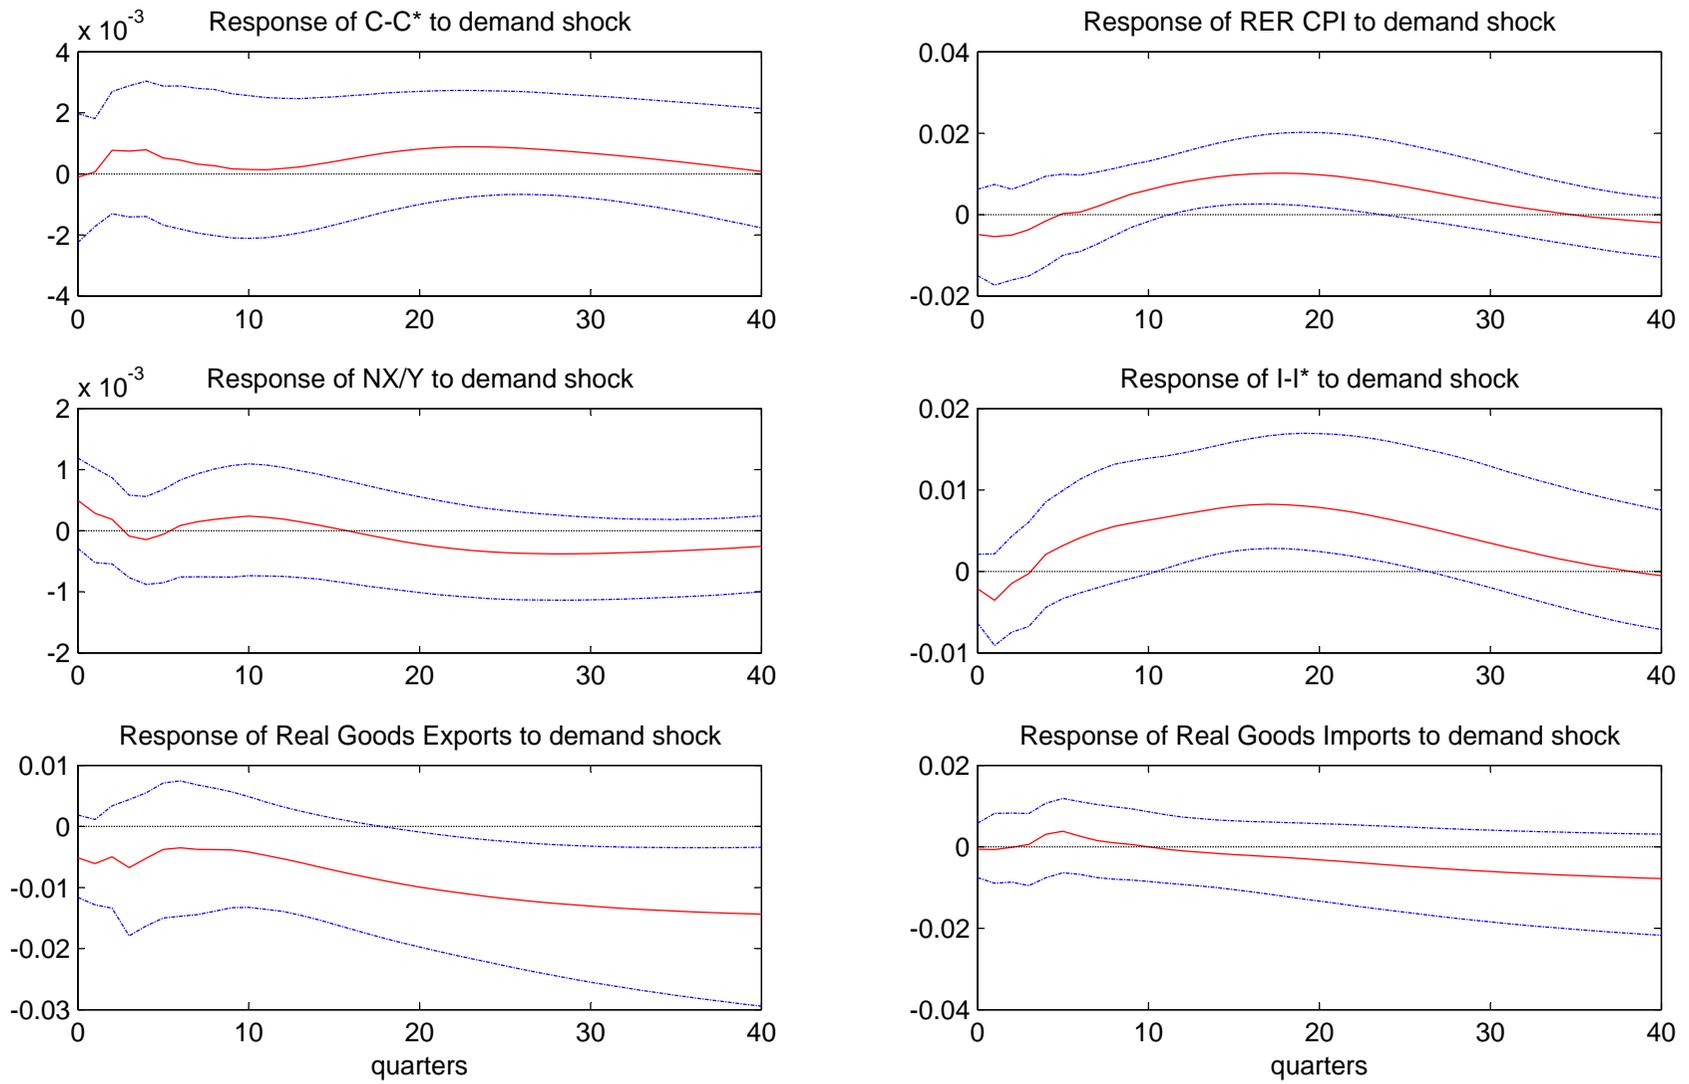


Figure 8

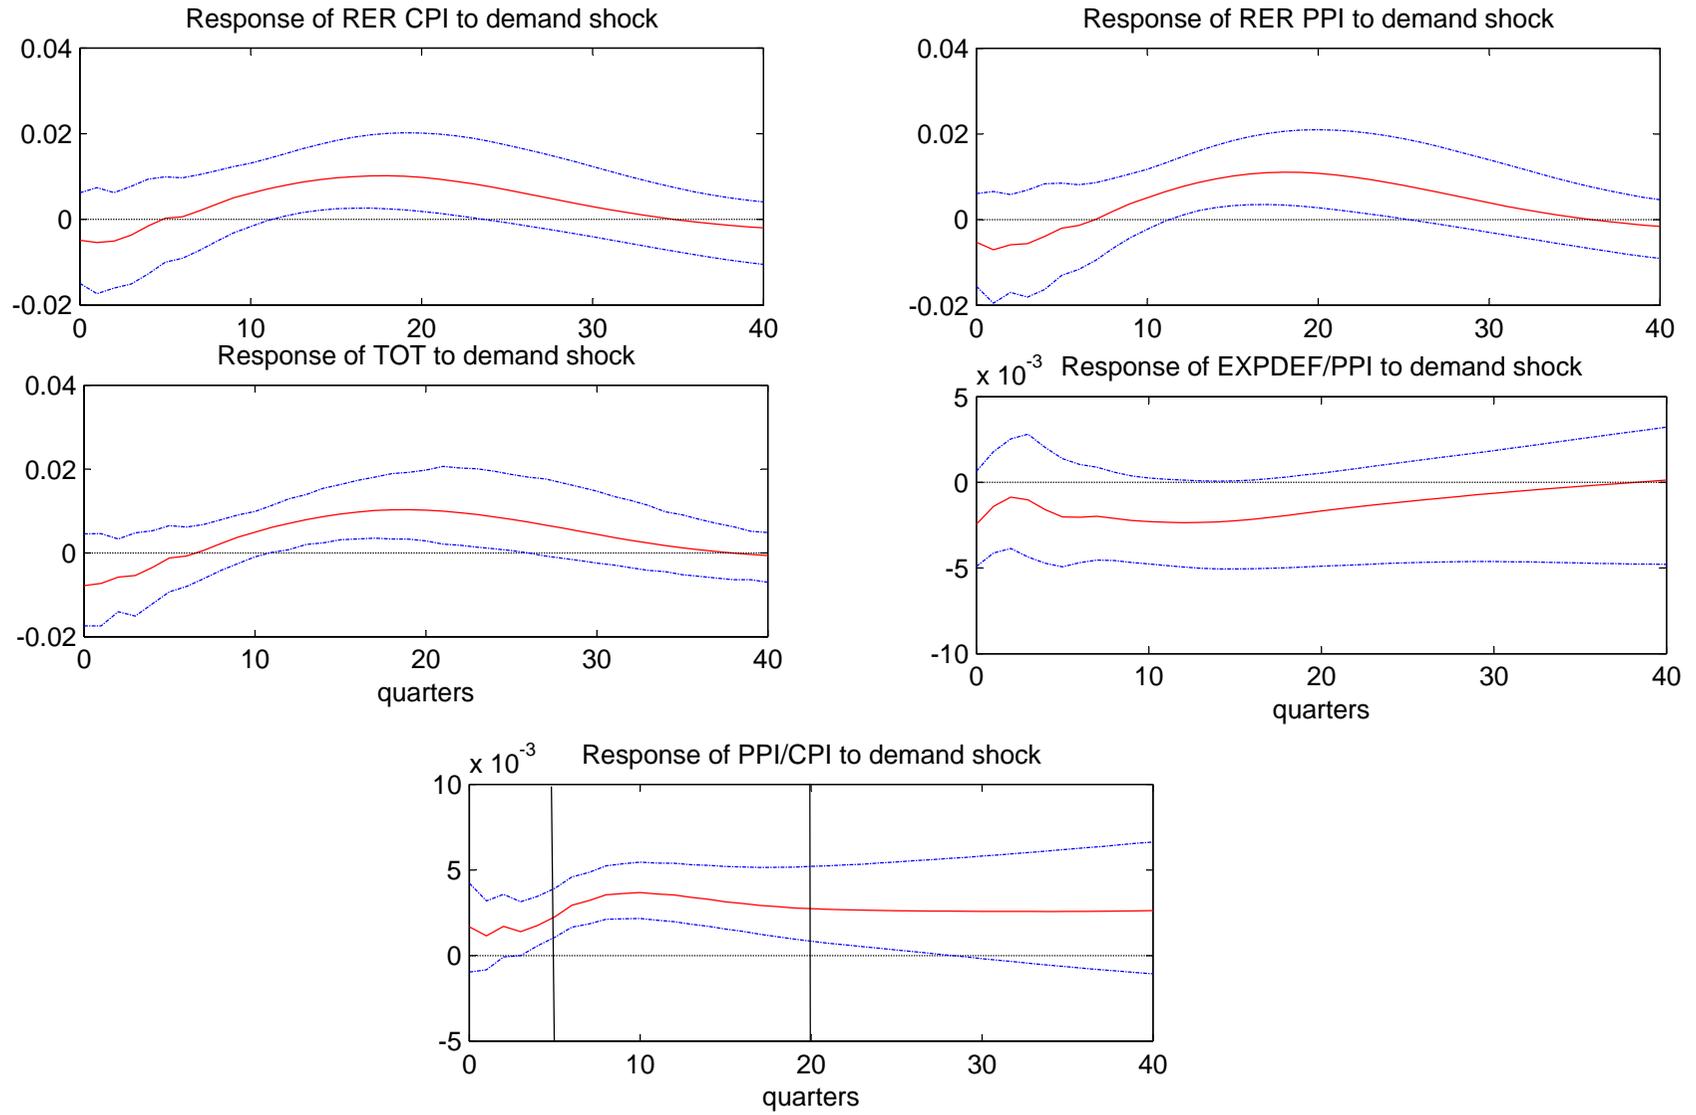


Figure 9

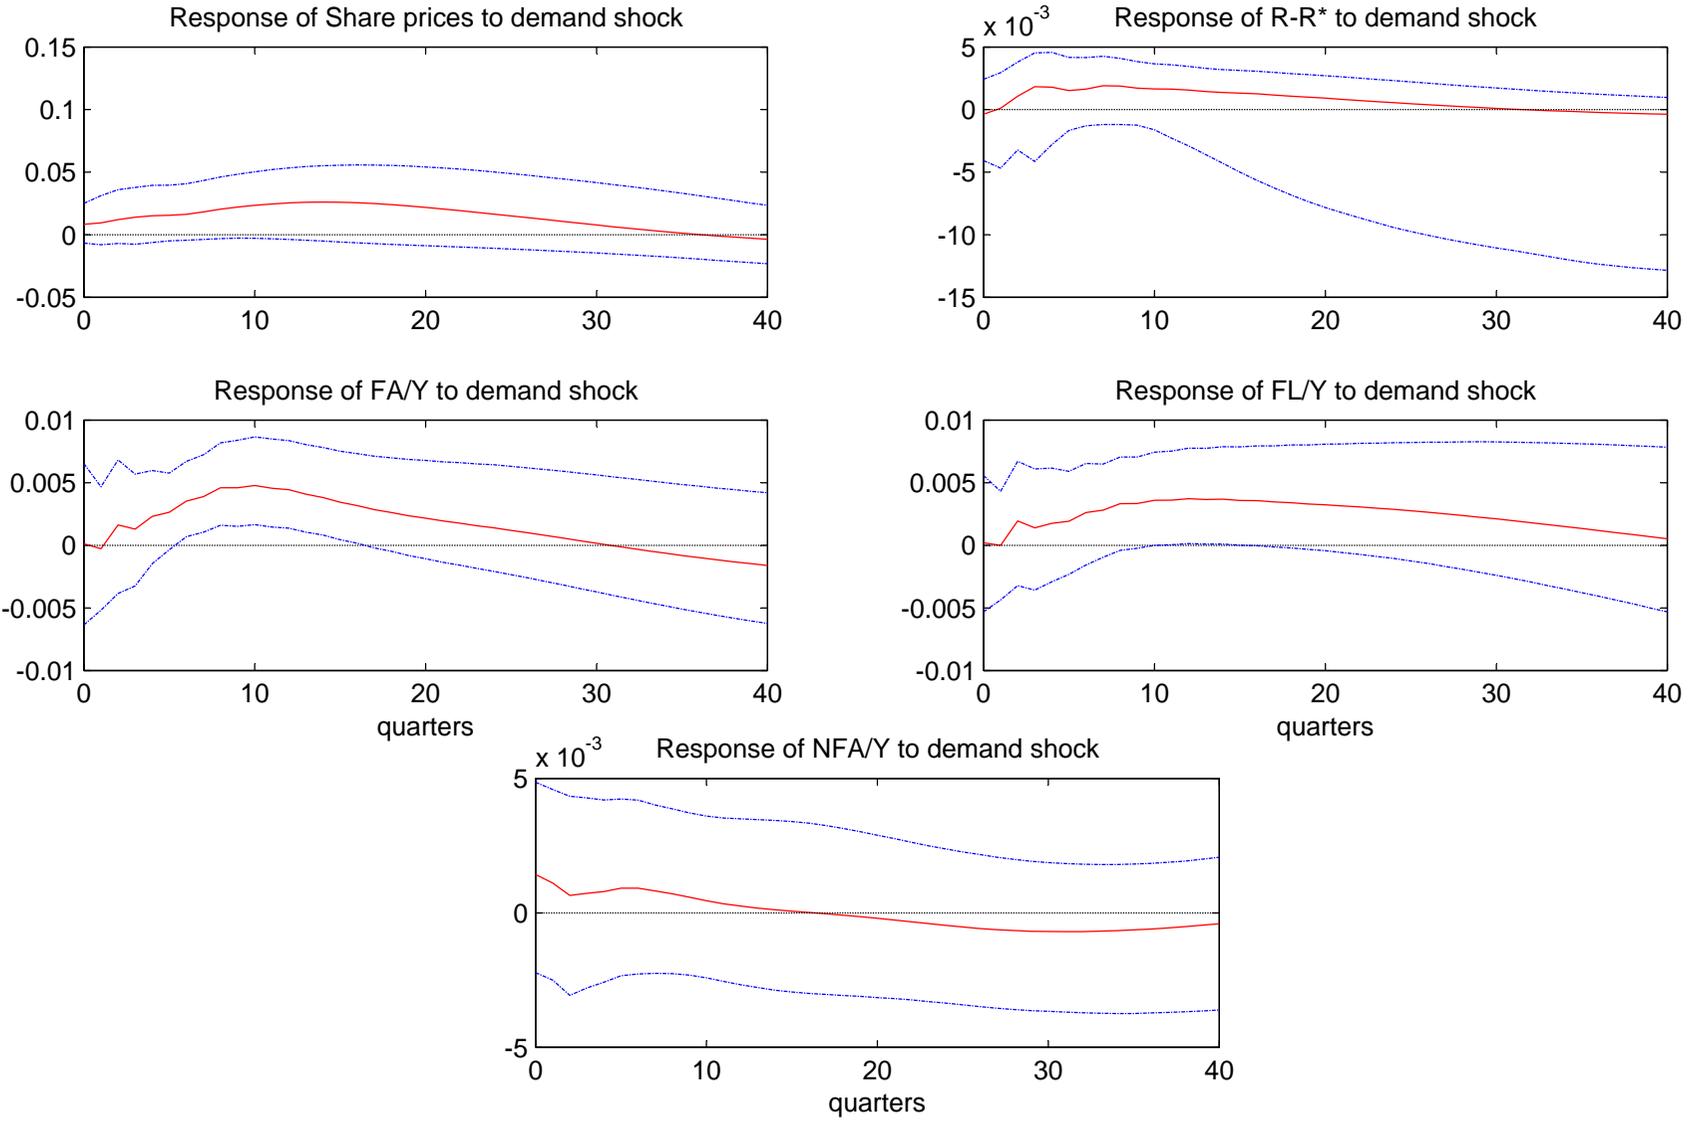


Figure 10

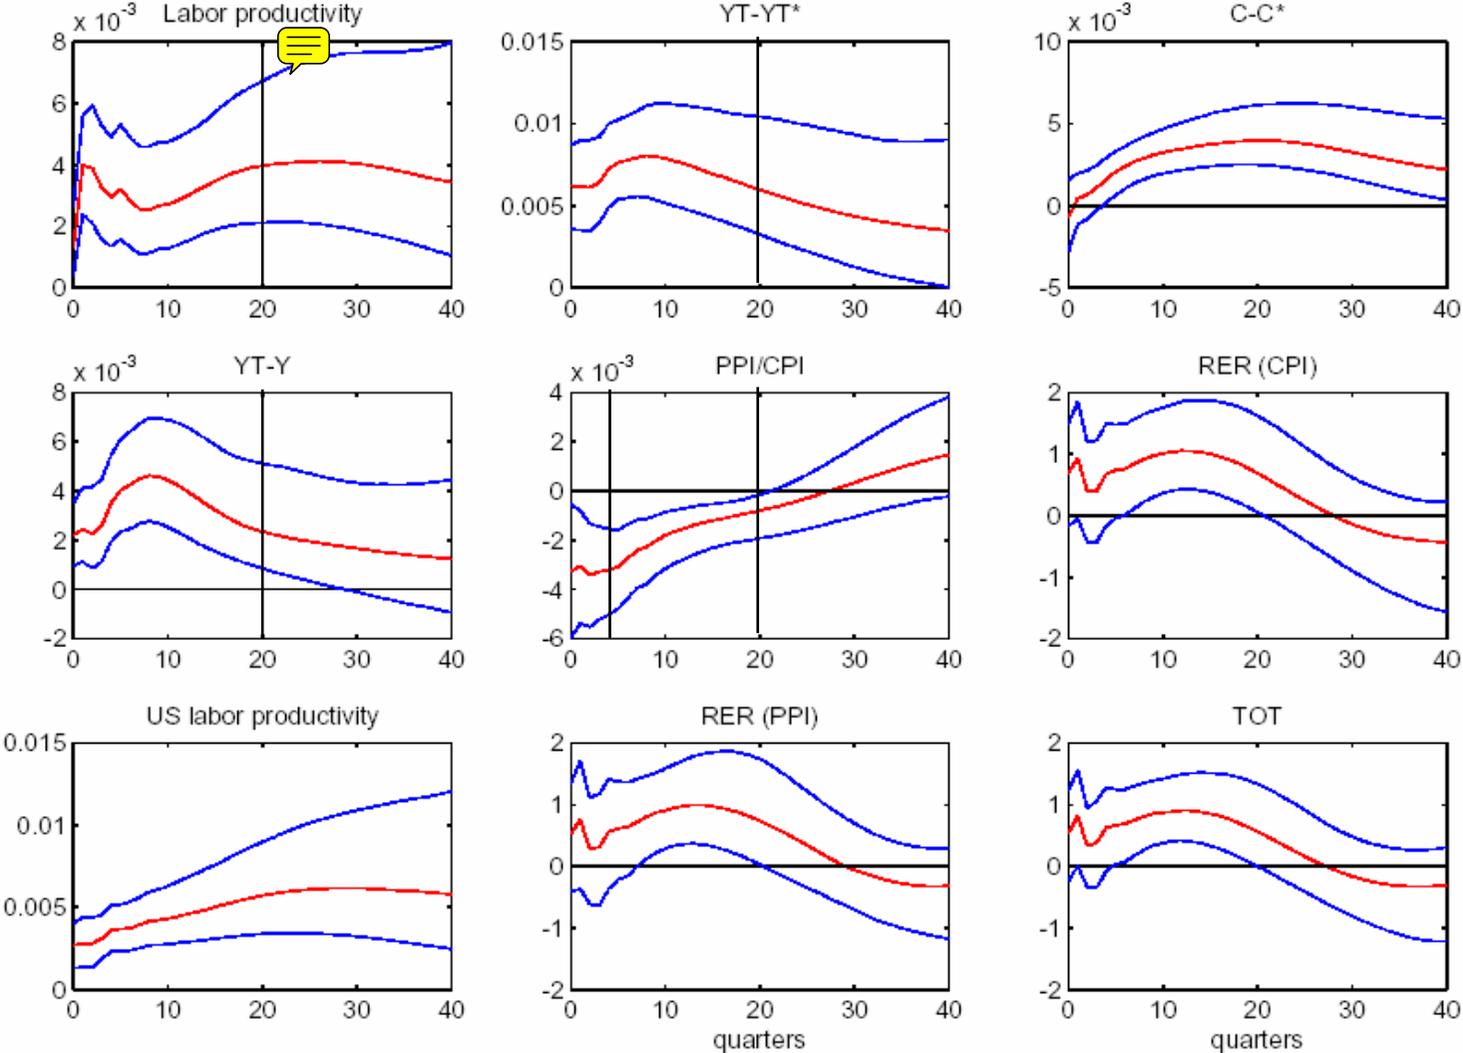


Figure 11

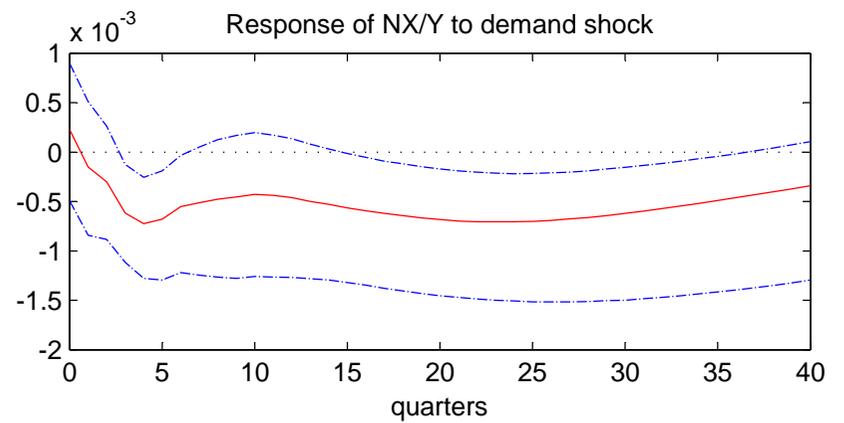
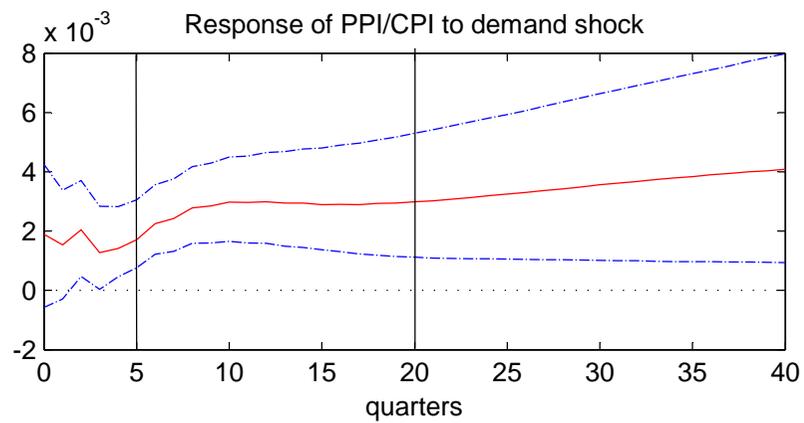
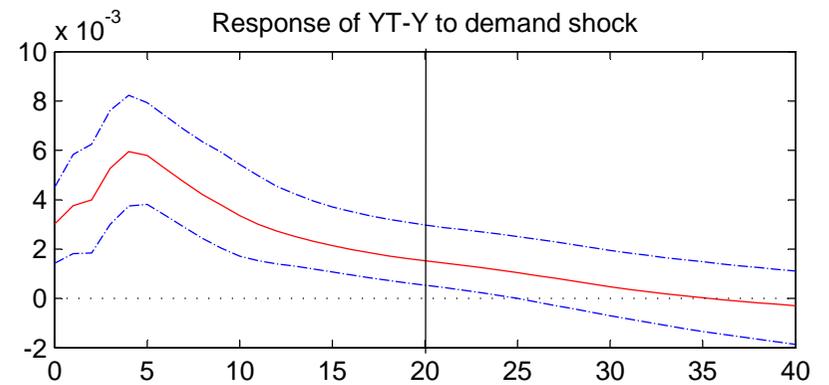
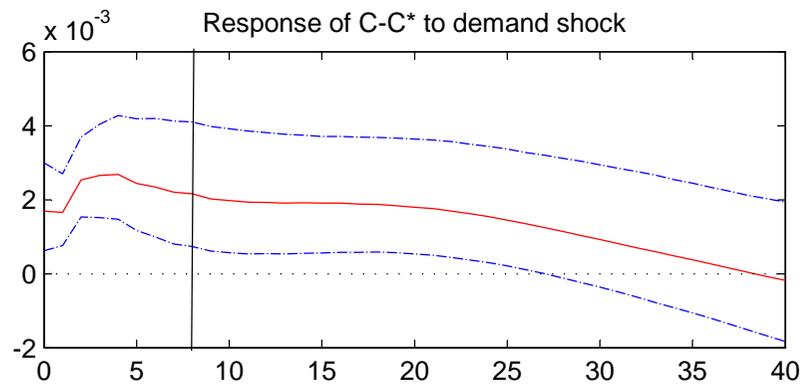
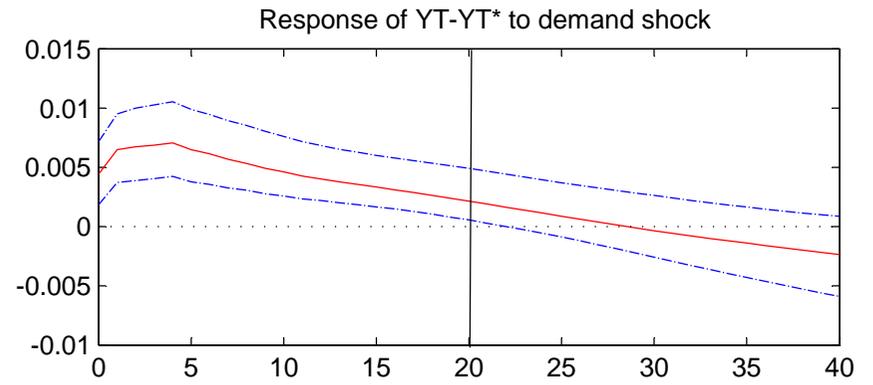
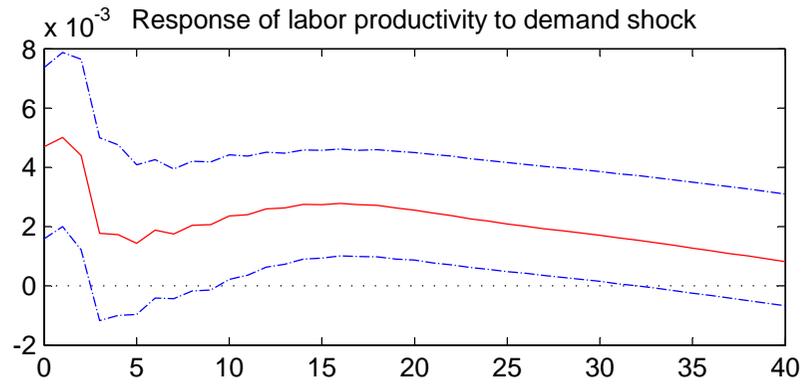


Figure 12

