Indirect Inference: A Local Projection Approach *

Juan Castellanos † Russell W. Cooper ‡

This draft: December 2022

First draft: April 2022

Abstract

This paper develops a new method for estimating the structural parameters of any DSGE model. The method uses Local Projection (LP) coefficients in an indirect inference exercise. Monte Carlo analysis is employed to examine its small sample performance and to compare it to the traditional approach in macroeconomics that relies on the use of VARs. Our results reveal that our approach produces consistent and computationally efficient estimates and that it outperforms the VAR approach in capturing the shape of true IRFs. This methodology is then used to re-estimate the Smets and Wouters (2007) model in an attempt to reconcile it with the new evidence on how the economy responds to aggregate shocks. We show that: i) our parameter estimates are similar to those obtained under full information, ii) the small differences in parameterizations are not enough to capture some of the effects of fiscal and monetary policy shocks, and iii) the model does a good job on capturing the responses to technology shocks even if we do not target them.

Keywords: *Indirect Inference, DSGE, Local Projections, VAR, Monte Carlo, Estimation*

JEL classification: C13, C15, E00

---

*We thank Wouter Den Haan, Lukas Nord, Giorgio Primiceri, Sebastian Rast, Barbara Rossi, Valerio Scalone and several participants at the EUI Macro Working Group and the ECB RCC5 Brownbag Seminar for useful comments and suggestions.

†European University Institute, juan.castellanos@eui.eu.

‡European University Institute, russellcoop@gmail.com.

1. Introduction

The Local Projection (LP) approach to understanding the dynamic effects of exogenous shocks, originating in Jordà (2005), has become a common tool for economic analyses. Though originally illustrated as an empirical model to study the impact of monetary shocks, this approach has become very widespread and used, as illustrated in Ramey (2016), to study the effects of fiscal and technology shocks as well. This paper, instead, uses the LP approach to estimate the structural parameters of a large scale macroeconomic model.

The motivation starts from the Smets and Wouters (2007) model, a leading representation of the aggregate economy. Given the prominence of this model, it is relevant to evaluate how well it matches the new evidence on how the economy responds to aggregate shocks. To that end we compare the impulse response functions (IRFs) from the Smets and Wouters (2007) at their mean parameter estimates to their empirical counterparts. In particular, we focus on the LP responses to technology and fiscal policy shocks as estimated in Ramey (2016) and to monetary policy shocks as estimated in Tenreyro and Thwaites (2016). Our first finding is that the Smets and Wouters (2007) IRFs do not match those from the data very well.

There are three possible explanations to this finding: i) the new empirical evidence points towards a different set of structural parameters, ii) the model misses certain key dimensions that makes impossible reconciling it with the data, and iii) the empirical IRFs are not well identified. This paper thoroughly explores the first one.

We re-estimate the parameters of the Smets and Wouters (2007) through an indirect inference exercise. In doing so, we make a methodological contribution by illustrating how to estimate the structural parameters of the model through the use of LP coefficients. This indirect inference approach is similar to that in Smith (1993) since he uses VAR coefficients as moments for the estimation.

In light of the theoretical result in Plagborg-Møller and Wolf (2021), i.e. VARs and LPs estimating the same impulse responses in population, one may think that using either VAR or LP should not matter for an indirect inference exercise. However, the finite sample properties of these two estimators differ. In particular, when p lags of the data are included in the VAR and as controls in the LP, IRFs approximately agree out to horizon p, but at longer horizons $h > p$ there is a bias-variance trade-off (Li et al., 2022).

This paper is complementary to Plagborg-Møller and Wolf (2021) and Li et al. (2022) in that our interest is as well on the performance of the LP approach. However, our focus is on the estimation of the structural parameters governing, for example, tastes and technology, rather than on the impulse response functions *per se*. Consequently, the paper is organized around evaluating the use of LP as a basis for indirect inference.

We first present a brief overview of the Smets and Wouters (2007) model in Section 2. This is followed by a comparison of the LP and VAR approaches to generating IRFs in Section 3. To evaluate the properties of these two econometric models as auxiliary ones in an indirect inference approach, in Section 4 we present Monte Carlo evidence.

For these experiments, the data generating process is a parameterized version of the Smets and Wouters (2007) model. In particular, the model parameters are fixed at the Smets and Wouters (2007) mean estimated values reported in Tables 1A and 1B in their paper. The structural parameters we focus on characterize household preferences, capital adjustment costs and the determination of wage and price rigidities. The model is then simulated to create time series. Using these time series, we estimate bivariate VARs which are used as one auxiliary model, as in Smith (1993). We also use the simulated data to estimate the LP from a variety of shocks: (i) monetary (ii) technology and (iii) fiscal. In both approaches, we use the observed values of the innovation to these shocks from the simulation of the Smets and Wouters (2007) model to identify the impulse responses.

Monte Carlo results show that there is a tradeoff between these two approaches. The VAR approach generally has a lower RMSE and its average fit is better. However, the theoretical responses at estimated parameters are much closer to truth under the LP approach. Overall, our findings seem to suggest that the LP approach is better at picking those parameters that have a bigger impact for the shape of the theoretical IRF at horizons h > p.

Building on this evidence, Section 5 presents our estimation results, using LP estimates from these three types of shocks from actual data, rather than from the model. The parameters can be estimated for each of these shocks independently or jointly. Compared to the estimates in Smets and Wouters (2007), our results show that there are some discrepancies, e.g. we obtain a lower intertemporal elasticity of substitution, higher frequencies of wage and price adjustment, and a lower degree of indexation to past wages and prices when we match either the responses to technology or fiscal shocks. These parameter values allow us to slightly improve the fit of the model when assessing the output, consumption, investment and hours worked responses to technology shocks, as well as the responses of output and hours worked to fiscal policy shocks. However, the model generates a negative consumption response and a large crowding out effect on investment when the economy is hit with a fiscal policy shock. These two features are at odds with the data. The poor performance of the model along these dimensions is not related to its parameterization, but rather to its structure. In fact, Galí et al. (2007) shows that one needs to introduce hand to mouth households into this type of models to be able to generate a rise in consumption in response to a government shock. Nevertheless, we cannot rule out that the disagreement between the model and the data is driven by the identifying assumptions used to empirically estimate the responses to a government spending shock.[1]

In regards to the monetary shocks, we find that the Smets and Wouters (2007) model is unable to match Tenreyro and Thwaites (2016) estimated responses of consumption, output and investment when we consider their linear LP model. Nevertheless, the disagreement between data and model doesn't seem to be related to the parameterization in general or our local projection approach in particular, but rather to the structure of the model and the empirical strategy, i.e. the linear LP model. As shown in Tenreyro and Thwaites (2016) paper, these responses are poor representations of reality since monetary policy is state dependent. It goes without saying that the log-linearized Smets and Wouters (2007) model is hopeless in generating such state-dependent responses to monetary shocks. In any case, we find interesting that the Smets and Wouters (2007) model is able to match the responses of consumption, output and investment to a contractionary shock during an expansion at both their estimates and ours; while it fails to do so during a recession. To the best of our knowledge, there is no economic model to date that is able to match the evidence in Tenreyro and Thwaites (2016). Based on our parameter estimates, which differ depending on which set of responses we are matching (boom vs. recession), we believe that an interesting avenue to reconcile the model with the data is through the firm's pricing decision. We left such extension of the model for future work.

[1] As shown in Ramey (2016), if one relies on government spending being pre-determined within a quarter to identify the shock, as in Blanchard and Perotti (2002), then government spending rises consumption; while when one uses a narrative news approach, such as the one developed by Ramey (2011), the effect is often reversed and government spending lowers consumption.

We also estimate parameters from the joint response of the model to all three types of shocks, rather than individually. In this case, we focus on the response of the most informative variables to estimate the structural parameters.[2] Estimation results are slightly different to those obtained by matching responses individually, nevertheless, the re-estimated Smets and Wouters (2007) model still misses the same dimensions of the data.

Finally, Section 6 concludes and summarizes our insights of what we believe are two promising lines of research. First, the development of a economic model that is able to capture the state dependent responses of monetary policy. And second, the application of our local projection approach to estimating such state-dependent models.

2. Economic Model

The analysis builds on the model formulated and estimated in Smets and Wouters (2007).[3] While other models may serve the same purpose, this structure captures many of the central channels of monetary and fiscal policy. For the first part of our analysis, we treat their estimated parameters as truth and see how close we come to them through the indirect inference approach. For the second part of the analysis, our estimated parameters are compared to those reported by Smets and Wouters (2007).

The Smets and Wouters (2007) model has become one, if not, the workhorse model in the DSGE literature. The model is based on Christiano et al. (2005) who added various frictions to a basic New Keynesian DSGE in order to capture the dynamic response to a monetary policy shock as measured by a structural vector autoregression (SVAR). In fact, price and wage stickiness paired with adjustment costs for investment, capacity utilization costs, habit formation in consumption, partial indexation of prices and wages as well as autocorrelated disturbance terms are able to generate a rich autocorrelation structure, which is key for capturing the joint dynamics of output, consumption, investment, hours worked, wages, inflation and the interest rate. These features of the model are crucial for our study since we are interested in matching the dynamic response of key macro aggregates to various shocks, as described in Section 3.

[2] We find that the investment response to shocks is very informative about parameters and use that feature to structure our estimation.

[3] Its parameters were estimated using Bayesian techniques.

In what follows, we summarize the model components with an emphasis on key parameters. The summary builds upon Smets and Wouters (2003).

2.1. Households

Households are infinitely lived, working and consuming in each period of life. Their lifetime utility is given by:

$$\sum_{t=0}^{\infty} \beta^t u(c_t, \bar{c}_{t-1}, n_t) \tag{1}$$

where c_t is the consumption of the representative agent, $\bar{c}_{t-1}$ represents an external habit and n_t is the amount of labor supplied. There are shocks associated with household impatience and the marginal rate of substitution between consumption and work. Household income comes from working, renting capital and dividends from the firms. Households save by holding bonds and also have access to state contingent securities which allow the household to smooth over taste shocks. Households also own the capital stock. This is rented to firms. Households incur an adjustment cost for changes in the capital stock. There is also a shock to this adjustment cost.

In the model, there are 4 parameters associated with the household problem: (i) the intertemporal elasticity of substitution σ_c, (ii) the strength of the habit h, (iii) the elasticity of labor supply σ_l, and (iv) the capital adjustment cost φ.

Further, each household acts as a wage setter, with a differentiated source of labor supply. There are associated parameters governing the likelihood of wage adjustment, $1 - \xi_w$ and the indexation of non-adjustment wages to past inflation, ι_w. These parameters are estimated as well.

2.2. Firms

The firm side of the economy entails a perfectly competitive final goods sector. Output is produced using differentiated intermediate goods. In the intermediate goods markets sellers have market power modelled as a monopolistic competition. The production of the intermediate goods requires capital and labor.

There are sticky prices in the market for intermediate goods. Price setting is not state dependent but rather firms are randomly granted an opportunity to adjust their prices, as in

Calvo (1983). For those firms not adjusting their prices, there is partial indexation to past inflation.

There is an aggregate productivity shock associated with intermediate goods production. The TFP shock is assumed to follow an AR(1) process.

From this specification there are a couple of key parameters. One is the probability that a firm can adjust its price, denoted $1 - \xi_p$ and the other one is the partial indexation parameter, denoted ι_p. These parameters are surely important for the effects of monetary policy. But, as we shall see, they also matter for the effects of other shocks and can be identified from impulse responses associated with non-monetary innovations.

2.3. Government: Fiscal and Monetary Authorities

Government spending is another source of stochastic variation in the model. Spending is specified as an AR(1) process with an iid normal error term and is also affected by the productivity shock.[4]

Monetary policy is modeled through a generalized Taylor rule that gradually adjusts the policy interest rate in response to changes in inflation and output gap. Innovations to this rule play a key role in the analysis.

3. Indirect Inference: LP vs. VAR impulse responses

This section fixes some basic ideas in order to clarify our approach and results. This includes the LP and VAR frameworks as well as the indirect inference approach.

Start by consider any economic model, M. Assume that the endogenous variables of this model, y_t, depend on its own lags y_{t-1} (endogenous states), some exogenous variables z_t (exogenous states) and some random errors u_t (shocks). Further assume that the model is parameterized by an ex-ante unknown vector of parameters Θ. That is, let

$$y_t = M(y_{t-1}, z_t, u_t; \Theta) \tag{2}$$

for $t = 1, 2, \ldots, T$. Given an initial value for the endogenous variable y_{-1} and a sequence for

[4] There is no apparent mention of taxation in either Smets and Wouters (2003) or Smets and Wouters (2007). Taxes appear to be lump sum. Distortionary taxes can change the impact of fiscal policy shocks.

the shocks $\{u_t\}_{t=1}^T$, it is possible to generate infinite data sequences $\{y_t\}_{t=1}^T$. This is a generic way to represent the Smets and Wouters (2007) model that provides the "sandbox" for our experiments.

3.1. LP

To understand the **LP** approach, consider the following regression:

$$\tilde{y}_{t+h} = \mu_h + \beta_h \tilde{x}_t + \sum_{\ell=1}^{p} \delta'_{h,\ell} w_{t-\ell} + \xi_{h,t}. \tag{3}$$

where $\tilde{y}_t$ is one of the variables of interest, $\tilde{x}_t$ denotes an innovation associated with a particular form of an aggregate shock. Finally, there are p lags of a vector of controls $w_t = \{\tilde{x}_t, \tilde{y}_t\}$.

The parameters in (3) are estimated at each horizon $h = 0, 1, 2, \ldots, H$. This is simply an OLS regression of leads of $\tilde{y}_t$ on past innovations. For each horizon, $(\mu_h, \beta_h, \{\delta'_{h,\ell}\}_{\ell=1}^p)$ are the projection coefficients.

Definition 1. *The LP - IRFs of $\tilde{y}_t$ with respect to $\tilde{x}_t$ are given by $\{\beta_h\}_{h\geq 0}$ in (3). Note that there are $H+1$ coefficients generated for each of the variables of interest, $\tilde{y}_t$ for each type of innovation, $\tilde{x}_t$.*

In our study, we focus on $\tilde{y}_t \in \{y_t, c_t, i_t, n_t\}$, being output, consumption, investment and hours worked respectively. Further $\tilde{x}_t \in \{\varepsilon_t^a, \varepsilon_t^g, \varepsilon_t^m\}$, so that we consider shocks to technology, government spending and monetary policy.

3.2. VAR

The starting point for the multivariate linear **VAR(p) projection** is:

$$w_t = c + \sum_{\ell=1}^{p} A_\ell w_{t-\ell} + u_t \tag{4}$$

where u_t is the projection residual and $(c, \{A_\ell\}_{\ell=1}^p)$ are the projection coefficients. Here p indicates the longest lag, matching the lag in the LP controls. Notice that given the definition of w_t, we are considering bivariate VAR(p) projections with the innovation ordered first (Plagborg-Møller and Wolf, 2021).

Let $\Sigma_u \equiv \mathbb{E}[u_t u_t']$ and define a *Cholesky decomposition* $\Sigma_u = BB'$ where B is lower triangular with positive diagonal entries. With this, consider the corresponding recursive *SVAR representation*:

$$A(L)w_t = c + B\eta_t \tag{5}$$

where $A(L) \equiv I - \sum_{\ell=1}^{p} A_\ell L^\ell$ and $\eta_t \equiv B^{-1}u_t$. Define the lag polynomial $\sum_{\ell=0}^{p} C_\ell L^\ell = C(L) \equiv A(L)^{-1}$.

Definition 2. *The SVAR - IRFs of $\tilde{y}_t$ with respect to an innovation in $\tilde{x}_t$ is given by $\{\theta_h\}_{h\geq 0}$ with $\theta_h \equiv C_{2,\bullet,h}B_{\bullet,1}$ where $\{C_\ell\}$ and B are defined in (5).*

3.3. Indirect Inference

Smith (1993) refers to the indirect inference approach as an extended method of simulated moments (EMSM). In fact, indirect inference is very similar to the simulated method of moments (SMM) approach since it also constructs an estimate of the parameters by minimizing the distance between data and simulated moments. The subtle difference between the two is that SMM uses unconditional moments, while in an indirect inference exercise these come from an auxiliary econometric model. Thus, the indirect inference estimator of a $q \times 1$ vector of structural parameters Θ, $\hat{\Theta}$ solves

$$\min_{\Theta} \left(\beta - \beta(\Theta)\right)' W \left(\beta - \beta(\Theta)\right) \tag{6}$$

where β is a $m \times 1$ vector containing the estimates of the econometric model from the actual data, and $\beta(\Theta)$ is its synthetic counterpart from the artificial data generated by the economic model. In this quadratic form, W is a weighting matrix.

In our application, Θ are the parameters characterizing household preferences, wage setting and price setting of firms in the Smets and Wouters (2007) model; while β are the LP coefficients associated with the impulse responses, i.e. β_h from Definition 1. In our Monte Carlo study, we also consider the case in which β are the coefficients from a VAR (or to be more precise their associated impulse responses, denoted θ_h in Definition 2).[5]

[5] We only use the coefficients associated to the impulse responses to guarantee the same number of moments across both econometric models.

4. A Monte Carlo Study

This section studies the small sample properties of the VAR and LP approaches to indirect inference under the hypothesis that the DGP and the estimated model are the same. Unlike other studies, we do not study such properties under the alternative hypothesis that the model is misspecified because moment-based methods tend to be robust to misspecification, as shown in Ruge-Murcia (2007).

4.1. The Data Generating Process

We solve the model in its log-linearized version, and then simulate it to generate an artificial database consisting of time series paths of four key macro aggregates: output, consumption, investment and hours worked $\{y_t, c_t, i_t, n_t\}$, as well as time paths for the innovations to technology, fiscal and monetary policy shocks $\{\varepsilon_t^a, \varepsilon_t^g, \varepsilon_t^g\}$.[6]

The Monte Carlo experiments are based on 100 replications using a sample size of 300 observations. Since LPs can be biased in smaller samples, as shown by Herbst and Johannsen (2021), we also repeat these experiments for a smaller sample of 100 observations. These results are shown in Appendix B.3.

We focus on the 8 structural parameters discussed above. The "true" values of these structural parameters are listed in Table 1, while the remaining ones are set and fixed at the estimated values from Smets and Wouters (2007). This reduces the computational burden in the Monte Carlo.

Table 1: True values of structural parameters

σ_c	h	σ_l	φ	ξ_w	ξ_p	ι_w	ι_p
1.26	0.80	2.52	6.31	0.70	0.66	0.58	0.24

4.2. The Moment Generating Functions

We consider two auxiliary econometric models to summarize the key features of the data: the LP and the VAR. However, as anticipated in the previous section, we only use those

[6] The log-linearized equilibrium conditions of the Smets and Wouters (2007) model are reproduced in Appendix A.

coefficients that identify the IRFs. The motivation relies in the excessive parameterization of the LP, which although useful, results in an absurd number of moments. Moreover, matching impulse responses is more intuitive as they clearly summarize the dynamics of the model and the data.

Canova and Sala (2009) argue that identification problems may arise because impulse responses at long horizons are noisy and contain little information about the parameters. Their results are based on an estimation strategy that uses the theoretical, rather than the estimated, impulse responses. As we shall see, it is precisely the different behavior of LP- and SVAR-IRFs at long horizons that is going to drive the results in our Monte Carlo study. In any case, we set the horizon to 20 quarters and verify that local identification is achievable at the true vector of parameters. That is, we verify that the Jacobian matrix of the moments is full rank, as suggested in Ruge-Murcia (2012).

Figure 1 depicts the distribution of the LP-IRFs (left panel) and SVAR-IRFs (right panel) to a technology shock over the 100 draws of the DGP. The black solid line corresponds to the median response of output, consumption, investment and hours to the shock; while the black dashed line depicts the true/model generated IRFs.

As demonstrated in Plagborg-Møller and Wolf (2021), the LP and VAR responses approximately agree up to horizon $h = 4$, which is the number of lags used in the VAR and the LP regression, $p = 4$. It is also clear from the figure that the variability of the estimated LP-IRFs is much larger than that of the SVAR-IRFs, specially at horizons $h > p$. Yet, from these figures it is also evident that the bias is much smaller for the LP-IRFs compared to the SVAR-IRFs. That is, the LP median response to the various shocks is closer to the true/model generated IRFs compared to the SVAR median response. In short, there is a bias-variance trade-off at horizons $h > p$ in the IRFs estimated form the Smets and Wouters (2007) model.[7] These three results are also present if one looks at the responses to fiscal and monetary shocks.[8]

One may expect that the increased variability in the LP-IRFs is also going to lead to more variation in the estimated economic parameters. But, at the same time, it can be useful if the response coefficients also change by more when one of the parameters changes, which in turn will help identifying the true parameter vector.

[7] This bias-variance trade off has been documented in Li et al. (2022) for other, more simple, DGP.

[8] The counterpart of Figure 1 for each of the two other shocks can be seen in Appendix B.5.1.

Figure 1: Technology Shock

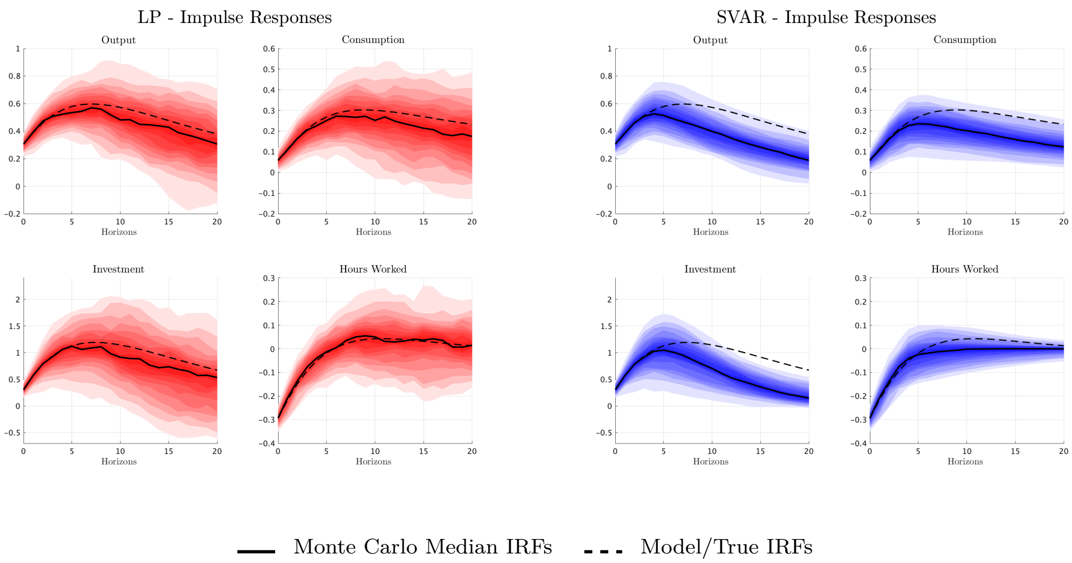

4.3. Other Aspects of the Design

We consider eight scenarios which only differ in the set of coefficients we try to match. Six of these eight scenarios correspond to optimization routines that target either the LP or SVAR responses of output, consumption, investment and hours worked to technology, fiscal or monetary shocks. In the other two scenarios, we exploit the differences in responsiveness of these coefficients to changes in the structural parameters. Thus, based on the sensitivity analysis reported in Appendix B.5.2, we target the LP or SVAR responses of investment to the three aggregate shocks as well as that of consumption to a technology shock. In sum, in each scenario we target a total of 84 ($= 21 \times 4$) moments to estimate the 8 parameters in Table 1.

The importance of each of these coefficients is weighted according to the inverse of the variance-covariance matrix of the moments, as suggested in Smith (1993). We compute it by estimating the IRF coefficients on simulated data at the true parameter vector and for 250 different draws of the innovation to the shocks.

Finally, we inflate the simulated sample size by a factor of 10 when computing $\beta(\Theta)$. In theory, we know that the asymptotic distribution of the estimates depends on this choice as simulation uncertainty decreases as the length of the simulated series to the sample size increases. However, in practice, having very long simulated series increases the computational cost and is not necessary to obtain accurate estimates. Ruge-Murcia (2012) shows how this choice affects the parameter estimates in the context of DSGE models estimated by SMM.

4.4. Results

This section reports the results of our Monte Carlo experiments. Tables 7 and 8 in Appendixes B.1 and B.2 report bias, standard deviation and root mean squared error for each of the 8 parameters. These are computed as follows:

$$\text{Bias}_i \equiv \mathbb{E}\left[\hat{\Theta}_i\right] - \Theta_i^* \tag{7}$$

$$\text{Std dev}_i \equiv \sqrt{\text{Var}(\hat{\Theta}_i)} \tag{8}$$

$$RMSE_i \equiv \sqrt{\text{Bias}_i^2 + \text{Var}(\hat{\Theta}_i)} \tag{9}$$

where expectations are taken over the 100 Monte Carlo draws and i indicates a specific parameter in the vector of estimated parameters.

We try to summarize all these information through Figure 2. Panel (a) analyses how the two moment generating functions, LPs and SVARs, pick up the different parameters when matching the same responses to the various shocks. In each of the four subplots a value greater than one indicates that a specific parameter is better identified using the SVAR as auxiliary model since its $RMSE$ is smaller. This is typically true for the Calvo adjustment and indexation parameters; while for the curvature, habit, elasticity of labor supply and capital adjustment parameters the relationship is reversed and identification is better via the LP approach, i.e. $RMSE$ ratio smaller than one. In any case, since this pattern is not general enough we also investigate whether some of these parameters may be better identified via different shocks.

Panel (b) compares the RMSE for each estimated parameter depending on the type of shock used to estimate the IRF coefficients. Here we restrict our attention to the LP approach. A few interesting patterns arise. On the one hand, there are some parameters that are better pinned down by a single source of variation. For example, the habit formation parameter h or the Calvo wage adjustment probability $(1 - \xi_w)$ are much better identified by monetary shocks alone (see subplot in the second column, third row). On the other hand, other parameters, such as the curvature in utility σ_c or the capital adjustment parameter φ are better pinned

Figure 2: Relative performance in terms of RMSE

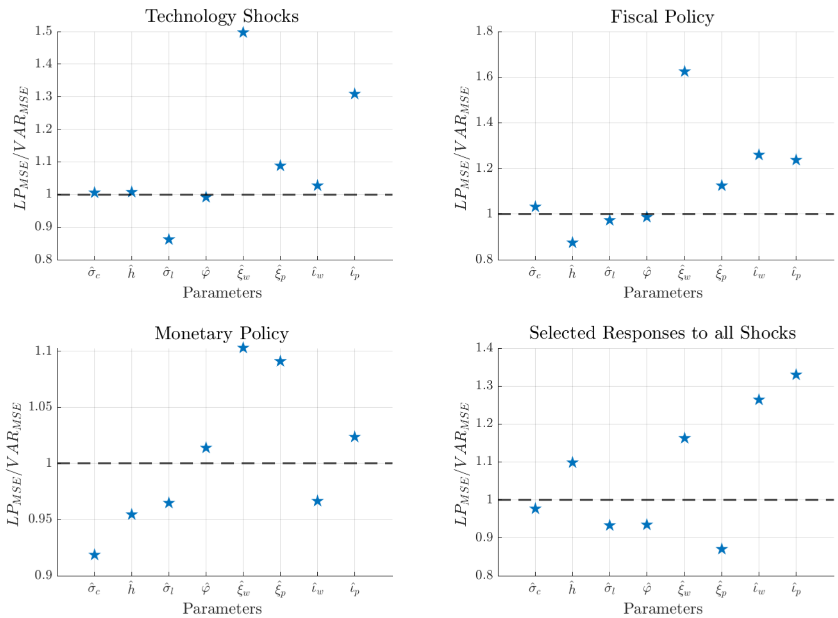

(a) LP - IRFs vs. SVAR - IRFs

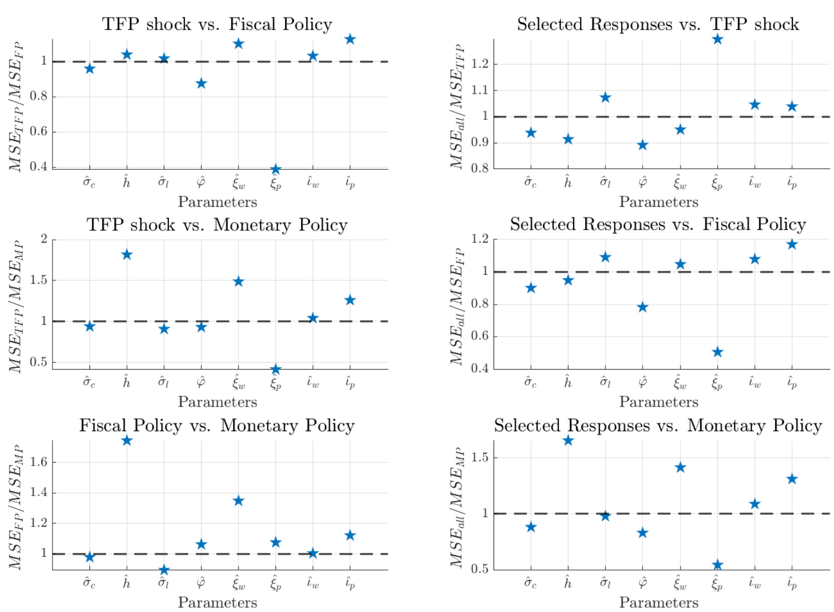

(b) Across Different Shocks (LP - IRFs approach) down when targeting the responses to all shocks. Finally, we also find surprising that the Calvo adjustment probability of prices $(1 - \xi_p)$ is better identified by a technology shock given its importance for inflation dynamics and therefore monetary policy. We do not interpret this result as problem but instead as an opportunity. Why? Because one can exploit the variability coming from a non-monetary shock to estimate some of the key parameters for monetary policy, such as $1 - \xi_p$, and then evaluate the validity of the model "out of sample" with respect to an untargeted response to a monetary shock.

These parameter-by-parameter comparisons are a good way to assess the relative performance of these two moment generating functions for each of the estimated parameters in a case-by-case basis. However, one would also like to know how these two distinct approaches to indirect inference perform when all the parameters involved in the estimation are considered as a whole. Thus, we evaluate the value of the criterion function, equation (6), at the optimal/estimated parameter values. This metric, usually refer to as the J-statistic, gives an overall measure of how well we match our targeted moments. Table 2 reports the average and the maximum value of this statistic across the 100 Monte Carlo draws as well as the time needed to complete the estimation for each of the eight different scenarios.

When comparing across rows, it easy to notice that the tailored scenario, in which we picked the most responsive IRF coefficients, results in the lowest average J statistic. A result that is probably expected. If one instead compares across moment generating functions, one finds that the LP approach does worse on average, but for very bad draws it does much better, i.e. lower maximum J. Moreover, the computational time needed to reach a minimum is significantly lower.[9]

Table 2: Overall performance: Estimated Impulse Responses

	Local Projections			Vector Autoregression		
	Avg. J	Max. J	Time (in min.)	Avg. J	Max. J	Time (in min.)
Technology Shock	87.23	117.00	28.62	82.80	247.94	67.98
Fiscal Policy	87.72	129.96	24.68	86.28	251.90	52.49
Monetary Policy	88.58	121.48	23.38	82.77	221.87	53.04
Selected Responses	86.56	128.65	20.03	82.63	240.07	72.42

[9] Notice that the lower computing comes from the lower number of iterations needed to reach the minimum since per iteration computing the LP-IRFs takes longer as one has to estimate more coefficients.

Finally, we consider a third measure to evaluate the performance of the LP and SVAR - IRFs approaches to indirect inference since neither the RMSE or the J-statistic inform us about how close we are from the true/model-implied impulse response functions, the ultimate object of interest. Therefore, we look at the weighted distance between the theoretical IRFs coming from the model at the estimated parameter values $\hat{\Theta}$ and at the true values Θ^*. Table 3 summarizes our findings.

Table 3: Overall performance: Model Impulse Responses

	Local Projections		Vector Autoregression	
	Avg. J^*	Max. J^*	Avg. J^*	Max. J^*
Technology Shock	2.57	9.43	34.67	228.41
Fiscal Policy	3.05	13.88	58.12	692.14
Monetary Policy	2.71	16.89	178.17	853.72
Selected Responses	8.37	44.69	230.46	1130.58

Results are striking. According to this metric, the LP - IRFs approach is able to match the theoretical IRFs much better than its SVAR counterpart. Why? The LP approach does a significantly better job in picking those parameters that are relevant for capturing the shape of the true impulse responses. In light of this evidence, we argue that using the LP - IRFs approach to indirect inference is the better alternative despite the mixed findings regarding the J-statistic and the RMSE.

5. An Empirical Application: Re-Estimating the Model

This last section re-estimates the Smets and Wouters (2007) model using our local projection approach to indirect inference. We target the empirical LP impulse responses to technology and fiscal shocks from Ramey (2016), and to monetary policy shocks from Tenreyro and Thwaites (2016). We estimate the same 8 parameters from the Monte Carlo analysis by exploiting the variation from each shock individually since our Monte Carlo results revealed that some of these parameters were better identified this way. Nevertheless, it wasn't true for all them. Therefore, we also re-estimate the model using the response coefficients from all three shocks jointly. For these estimation exercises, we weight the importance of each coefficient via the inverse of a diagonal matrix comprised of the standard errors of the LP coefficients. Finally, we compute the standard errors of the estimated parameters through bootstrapping since Ruge-Murcia (2012) showed that there are discrepancies between the asymptotic and finite sample distributions of the estimates obtained via SMM.

5.1. Technology Shocks

Technology shocks are the most important type of non-policy shocks. In fact, there is a vast literature on identification of these shocks on time series models. From all these, we use Francis et al. (2014) approach and identify an unanticipated TFP shock through medium-run restrictions.[10]

The identified shock is then used as dependent variable in a local projection regression to estimate the impulse response of the variable of interest. Following Ramey (2016), we estimate the following series of regressions:

$$\tilde{z}_{t+h} = \alpha_h + \beta_h \cdot \text{shock}_t + \varphi_h(L) \cdot \text{control}_{t-1} + \text{ quadratic trend } + \varepsilon_{t+h} \tag{10}$$

where $\tilde{z}_{t+h}$ are the variables of interest (real GDP, consumption, non-residential investment, and hours worked), $shock_t$ is the innovation to the growth rate of TFP, $control_{t-1}$ includes two lags each of the shock, real GDP, stock prices, labor productivity and the dependent variable, and ε_{t+h} is the error term. As in equation (3), β_h gives the response of $\tilde{z}$ at time $t + h$ to a shock at time t.

However, the regressions in (3), which we use in the simulated data, differ from the regressions in (10) in that i) we do no control for the trend in the regression on model simulated data since variables are already in deviations from steady state, and ii) we only control for lags of the dependent variable and the shock. Since we use the actual innovation of the shock in the simulated data, we think that these differences won't affect our results.[11]

Thus, we target the β_h coefficients associated with the responses of real GDP per capita, consumption, non-residential investment and hours in our inderect inference estimation. Results

[10] A review of the literature on TFP shock identification can be found in Section 5 of Ramey (2016) handbook chapter.

[11] The identified shock form the data may still suffer from measurement error and be correlated with the error term. Thus, researchers typically use additional controls, which sometimes are not observable in the model. On the other hand, the innovation of shock in the model is purely exogenous, thus there is no need for additional controls.

are reported in the second subtable of Table 4. Our median estimates are reasonably close to the mean estimates in Smets and Wouters (2007), which are reproduced again at the top of the table for convenience. In fact, their point estimates fall within our confidence intervals. Nevertheless, we estimate a lower intertemporal elasticity of substitution, habit parameter, and degrees of indexation to past wage and price inflation. Moreover, we obtain higher labor supply elasticity, adjustment costs and Calvo adjustment probabilities.

Table 4: Estimates using LP coefficients in Ramey (2016)

	$\hat{\sigma}_c$	$\hat{h}$	$\hat{\sigma}_l$	$\hat{\varphi}$	$\hat{\xi}_w$	$\hat{\xi}_p$	$\hat{\iota}_w$	$\hat{\iota}_p$
S&W 07	1.26	0.80	2.52	6.31	0.70	0.66	0.58	0.24
	Technology Shocks							
Median	0.85	0.69	3.28	8.20	0.44	0.59	0.47	0.14
10th pctl.	0.76	0.48	1.51	3.79	0.42	0.40	0.35	0.14
90th pctl.	1.36	0.89	3.28	8.20	0.84	0.86	0.75	0.31
	Fiscal Policy							
Median	1.01	0.85	1.51	3.90	0.42	0.40	0.36	0.14
10th pctl.	0.81	0.48	1.51	3.79	0.42	0.40	0.35	0.14
90th pctl.	1.57	0.96	3.00	7.89	0.81	0.82	0.72	0.30

To assess the performance of our estimation, we compute the estimated responses on simulated data at our median estimates. To reduce simulation error we report the mean β_h coefficients across 1,000 draws. These responses (red dashed line) are compared to the targeted ones (blue dashed line) in Figure 3.

Overall, the Smets and Wouters (2007) model at our parameter estimates does a good job in capturing these responses. The match is particularly close for output, investment and hours. However, the model is not able to produce the large effect on consumption upon impact nor its delayed peak. These discrepancies with the data are partially solved through the parameterization of the model. In fact, the same impulse response at the Smets and Wouters (2007) parameters, the black dotted line, generates a smaller effect on consumption upon impact and a much smoother hump.

Figure 3: TFP shocks – Empirical vs. Model Estimated IRFs

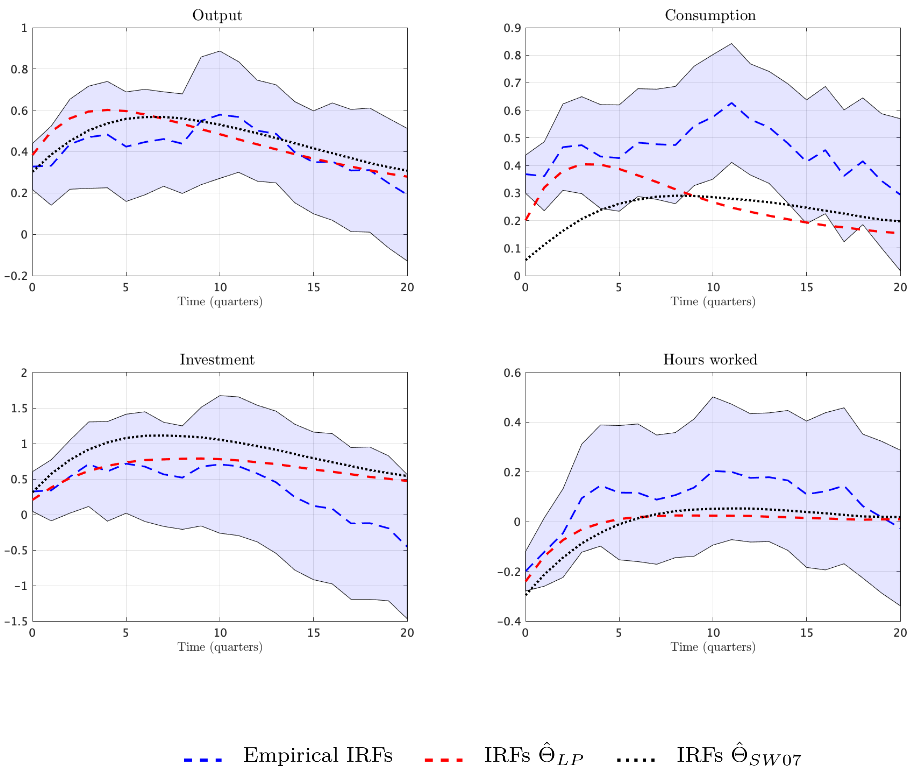

5.2. Government Spending Shocks

Now we turn to fiscal policy shocks. As for the technology shock there are many identification strategies.[12] We identify this shock as in Blanchard and Perotti (2002), that is we order government spending first in a Cholesky decomposition. This is our preferred strategy because it does not rely on the news of a fiscal intervention, an aspect that is not present in the Smets and Wouters (2007) model.[13]

We target the responses of GDP, non-durables and services consumption, non-residential investment and hours worked to the Blanchard and Perotti (2002) shock. These LP coefficients are estimated by means of regressions (10) but now with the shock being the identified government spending shock, and the controls being two lags each of the shock, real GDP, real

[12] Ramey (2016) summarizes some of the most prominent identification strategies for government spending shocks in Section 4 of her handbook chapter.

[13] Ramey (2011) or Ben Zeev and Pappa (2017) are leading examples that rely on the news to identify government spending shocks. In the Appendix we also estimate the model by matching the responses associated to Ramey (2011) shock since there the response of consumption to a government spending shock is negative as predicted by the model.

Figure 4: Fiscal Policy – Empirical vs. Model Estimated IRFs

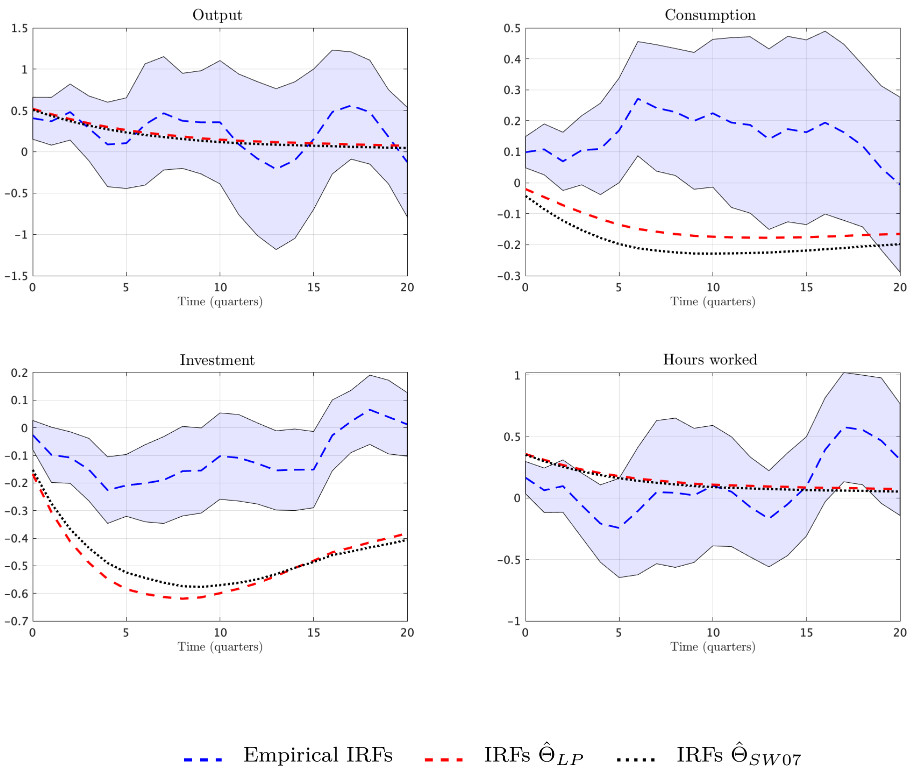

government purchases, and the tax rate. These responses are depicted in blue in Figure 4, where we also plot the estimated impulse responses at our median parameter estimates (red dashed line) as well as at Smets and Wouters (2007) mean estimates (black dotted line). The parameters used to generate these responses are reported in the third subtable of Table 4.

As with the technology shock, we also obtain lower curvature in utility, higher wage and price adjustment probabilities and lower degrees of indexation to past inflation than Smets and Wouters (2007). Moreover, trying to match only the response to government spending leads to higher consumption habit and lower elasticity of labor supply and adjustment costs.

These new estimated parameters only marginally change the impulse responses. Thus, the model at either our or their estimated parameters is only able to match the response of output. As for hours worked, the model has a bigger response than in the data in the first five quarters, it predicts a much larger crowding out effect on investment, and predicts a drop in consumption. Thus, the model is unable to match many dimensions of the dynamic response of the economy to a government spending shock.

The most worrisome is the response of consumption given its importance for the calculation of fiscal multipliers. It is known that in these class of models households anticipate future increases in (lump-sum) taxes which makes them increase their labor supply to compensate for the negative wealth effect which in turn brings consumption down. Galí et al. (2007) demonstrate that the inclusion of hand to mouth consumers breaks the Ricardian equivalence. Thus, introducing household heterogeneity into the model helps in generating a rise of aggregate consumption in response to an unexpected increase in government spending.[14]

Nevertheless, there is no consensus in the empirical literature on the response of consumption. As illustrated in Ramey (2016), the identification strategies that rely on the assumption of government spending being predetermined within a quarter, as done for example in Blanchard and Perotti (2002), find that government purchases rise consumption; while those that rely on the narrative news approach estimate the opposite effect, i.e. a reduction in consumption as a result of an increasing in government spending.

5.3. Monetary Policy Shocks

The estimation of parameters based upon monetary policy shocks rests upon the local projection estimates in Tenreyro and Thwaites (2016). They identify the monetary policy shock using a non-linear Romer and Romer (2004) regression on 40 years of quarterly data. Their specification is of particular interest as it allows a state dependent response to monetary innovations. In particular, they find that a monetary contraction during a boom creates responses in key macroeconomic variables, such as output, (nondurable) consumption and investment, that are quite different from the responses to a monetary contraction during a recession.[15] Their local projection estimation is based upon:

$$y_{t+h} = \tau t + F\left(z_t\right)\left(\alpha_h^b + \beta_h^b \varepsilon_t + \gamma^{\mathbf{b}'}\mathbf{x_t}\right) + \left(1 - F\left(z_t\right)\right)\left(\alpha_h^r + \beta_h^r \varepsilon_t + \gamma^{\mathbf{r}'}\mathbf{x_t}\right) + u_t \tag{11}$$

where τ is a time trend, α_h^j is a constant and $\mathbf{x_t}$ are controls. $F(z_t)$ is a smooth increasing function of an indicator of the state of the economy. This is the way in which the state dependence of monetary policy is captured.

[14] Another important limitation of the Smets and Wouters (2007) model for the study of fiscal policies and the propagation of shocks is the absence of distortionary taxation.

[15] They derive economic expansions and contractions in terms of output growth, not levels.

Figure 5: Monetary Policy – Empirical vs. Model Estimated IRFs

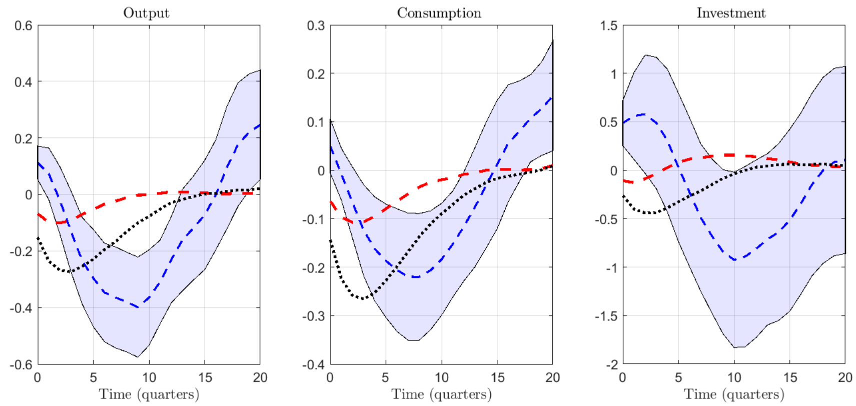

(a) Linear Model

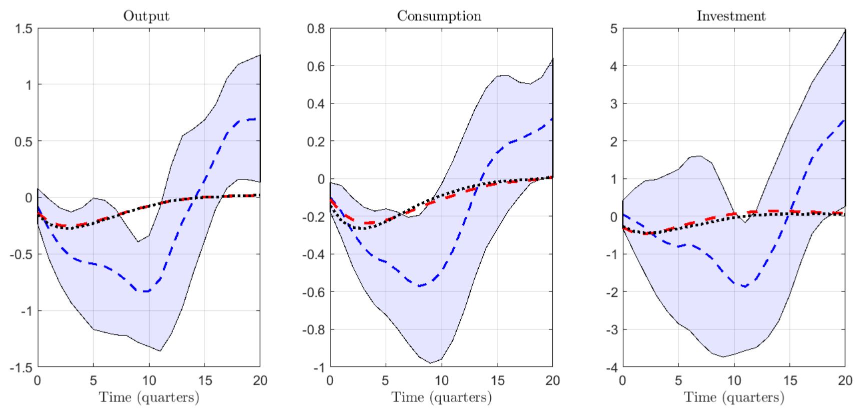

(b) Non-Linear Model: Expansion

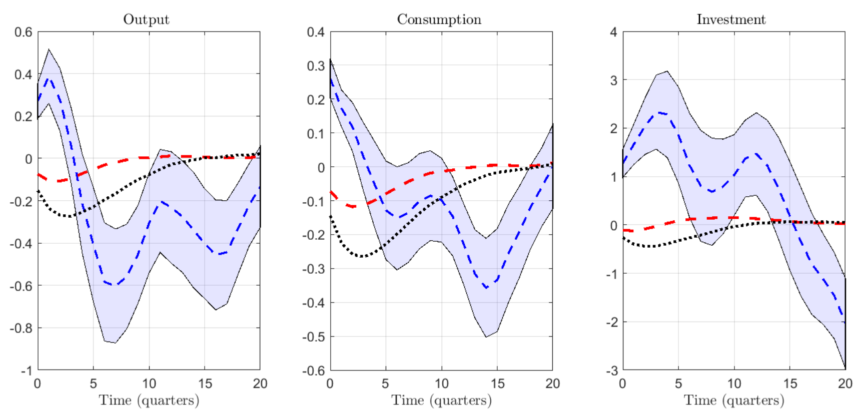

(c) Non-Linear Model: Recession

Empirical IRFs IRFs $\hat{\Theta}_{LP}$ IRFs $\hat{\Theta}_{SW07}$

Figure 5 illustrates the state dependent estimated effects of monetary policy shocks on output, consumption and investment. For the linear model (thus imposing no state dependence), output, investment and consumption show a slight increase on impact of the contractionary shock. By the second quarter, output and consumption fall, while investment does not fall until 5 quarters. The response to the shock is prolonged, including an overshooting after approximately 15 quarters.

These impulse responses combine the effects of monetary contractions in booms and recessions. Focusing on boom times, the second row of the figure shows more "conventional" effects of a contraction, including an immediate and prolonged fall in output, consumption and investment combined with a slow return. From the third row, the effects of a monetary contraction during a recession are quite different. One natural interpretation is that the contractionary policy during a recession is occurring during a period of stagflation, with monetary policy focusing on combating inflation.

The Smets and Wouters (2007) model clearly does not contain the necessary non-linearities that might give rise to the state dependent effects of monetary policy. Nonetheless it is instructive to see how well we can match the moments from the state dependent estimation as well as the linear model. Note, however, that when matching the moments form the state dependent estimation we still use a linear local projection on the model simulated data. That is, we allow the parameters to vary with the treatment to provide a sense of how the state dependent responses translate into different parameter estimates. Clearly if we had a model with nonlinearities, we would fix the parameters and try to match the state dependent responses. Table 5 reports our parameter estimates.

Looking first at the linear model, our median parameter estimates are generally close to the mean estimates in Smets and Wouters (2007) and the reported confidence intervals contain their parameter estimates. Our estimated elasticity of substitution is about 25% larger than theirs. The habit is stronger in our model and labor is more elastic. Our point estimates of the price adjustment parameter, $\hat{\xi}_p$, is much lower than theirs, indicating that in our estimated model prices adjust more frequently.

Focusing on the effects of monetary policy contractions during booms, our estimate of price flexibility match theirs, though the wage flexibility and the elasticity of labor supply is a bit lower than theirs. In contrast, the estimated habit is stronger as is the curvature in utility.

Table 5: Estimates using LP coefficients in Tenreyro and Thwaites (2016)

	$\hat{\sigma}_c$	$\hat{h}$	$\hat{\sigma}_l$	$\hat{\varphi}$	$\hat{\xi}_w$	$\hat{\xi}_p$	$\hat{\iota}_w$	$\hat{\iota}_p$
S&W 2007	1.26	0.80	2.52	6.31	0.70	0.66	0.58	0.24
	Linear Model							
Median	1.57	0.88	3.15	7.89	0.46	0.32	0.63	0.11
10th pctl.	0.92	0.71	1.51	3.79	0.46	0.32	0.32	0.10
90th pctl.	1.57	0.95	3.15	7.89	0.80	0.66	0.66	0.21
	Non-Linear Model: Expansion							
Median	1.46	0.83	2.20	4.98	0.64	0.66	0.33	0.21
10th pctl.	0.81	0.51	1.51	3.79	0.46	0.32	0.32	0.10
90th pctl.	1.57	0.97	3.15	7.89	0.84	0.66	0.66	0.21
	Non-Linear Model: Recession							
Median	1.57	0.86	3.15	7.89	0.46	0.32	0.66	0.21
10th pctl.	0.87	0.60	1.51	3.79	0.46	0.32	0.32	0.10
90th pctl.	1.57	0.98	3.15	7.89	0.77	0.66	0.66	0.21

The impulse response function at our median estimates in each of these three cases are shown (in red) in Figure 5. The estimated model can match the basic pattern of the responses to monetary contractions during an expansion. Output, consumption and investment fall on impact and recover slowly, with a slight hump shape. The impulse responses from Tenreyro and Thwaites (2016) show a more pronounced hump-shape: it is both larger and more delayed. We also do not capture their overshooting.

The estimated model is unable to capture the effects of monetary contractions during recessions. The underlying Smets and Wouters (2007) model is evidently unable to reproduce the initial positive output response to a contractionary shock. This same point applies to the linear model, thought to a lesser degree.

The impulse responses at Smets and Wouters (2007) mean parameters are similar in the linear model except that at their parameter estimates the model predicts a more pronounced drop in output, consumption and investment, which is more similar to the data. However, the hump is more delayed in the data, around Q10, than in the model, around Q3. The model at the two set of parameter estimates are quite close when studying the effects of contractionary policy during an expansion. Neither of the models can match the effects of monetary contractions during recessions.

5.4. All Shocks: The Response of Investment

As a final exercise, we consider all three sources of variation simultaneously, rather than individually. In doing so, we are uncovering the structural parameters that best match the responses jointly. For the previous exercises, this restriction to a single vector of parameters across estimation exercises was not imposed. So, for example, the median estimate of σ_l was more that 2 times larger when matching the response to technology shocks compared to fiscal policy. The current exercise restricts these parameter to be the same across sources of variation.

As noted earlier, the response of investment is very sensitive to a change in the parameters. Accordingly, this estimation uses the LP investment response to the three types of shocks as targeted moments.[16] The results are shown in Table 6. The first block reports the estimates of Smets and Wouters (2007) which allowed multiple shocks, the middle block shows the parameter estimates from the joint shock case, and the last block recalls the estimates by type of shock.

Comparing the last two blocks, for the most part, the estimates for the cases of the individual shocks lie within the intervals created by the joint estimation. The estimated curvature of 0.91 lies between the other estimates. The habit parameter, $\hat{h}$ is smaller than any of the other point estimates. The frequency of wage and price adjustment is similar to the estimates based upon the response to fiscal shocks. It is in this sense the joint estimation exercise combines the responses to the individual shocks.

Compared to the estimates of Smets and Wouters (2007), the median estimate of the utility function is lower and the frequency of wage and price adjustment is considerably higher. That said, the point estimates of Smets and Wouters (2007) do lie in the interval between the 10th and 90th percentile.

It is important to note, however, that all these parameters have been estimated using information from different samples. For example, the responses to technology and fiscal shocks are recovered from data spanning from the late 1940s to the early 2010s, while the monetary shocks are only available from 1969 to 2007. Therefore, it is possible that some of the discrepancies we found, specially those related to the Calvo parameters, arise due to sample

[16] An alternative to pursue would be to include other variables but only at short/medium term horizons such that the moments are even more responsive to parameters. Note that the IRFs tend to die out at long horizons.

Table 6: SMM estimates combining information

	$\hat{\sigma}_c$	$\hat{h}$	$\hat{\sigma}_l$	$\hat{\varphi}$	$\hat{\xi}_w$	$\hat{\xi}_p$	$\hat{\iota}_w$	$\hat{\iota}_p$
S&W 2007	1.26	0.80	2.52	6.31	0.70	0.66	0.58	0.24
	Jointly							
Median	0.91	0.59	2.95	5.65	0.42	0.40	0.36	0.14
10th pctl.	0.76	0.48	1.51	3.79	0.42	0.40	0.35	0.14
90th pctl.	1.57	0.90	3.15	7.89	0.82	0.82	0.72	0.30
	Independently							
Technology	0.85	0.69	3.28	8.20	0.44	0.59	0.47	0.14
Fiscal Policy	1.01	0.85	1.51	3.90	0.42	0.40	0.36	0.14
Monetary Policy	1.26	0.91	3.15	7.89	0.46	0.32	0.32	0.10

selection. As noted by Fernández-Villaverde et al. (2007), there is some evidence that certain DSGE parameters, such as those characterizing the pricing behavior of firms and households, change depending on the sample used for estimation. In fact, Smets and Wouters (2007) find a higher degree of price and wage stickiness when their model is estimated only using data from the "Great Moderation" period (1984Q1 - 2004Q4).

Do the different sources of variation help in reconciling the model with the data? Figure 6 shows the targeted response of investment to the three types of shocks while Figure 7 shows the untargeted output, consumption, and hours response.

Figure 6: Targeted Investment Response to All Three Shocks

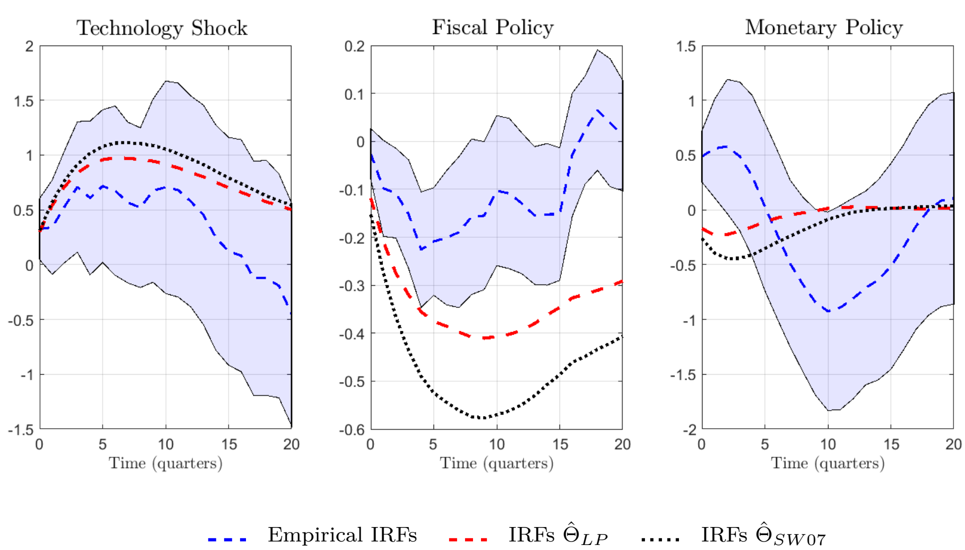

Figure 7: Untargeted Response to All Three Shocks

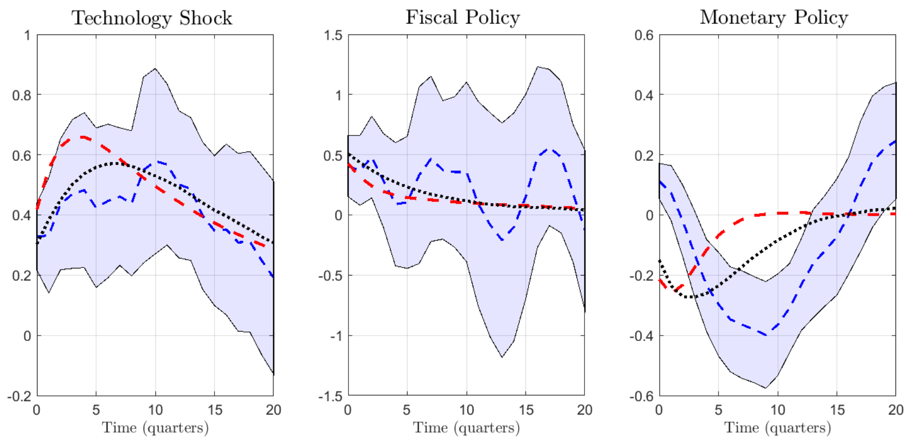

(a) Output

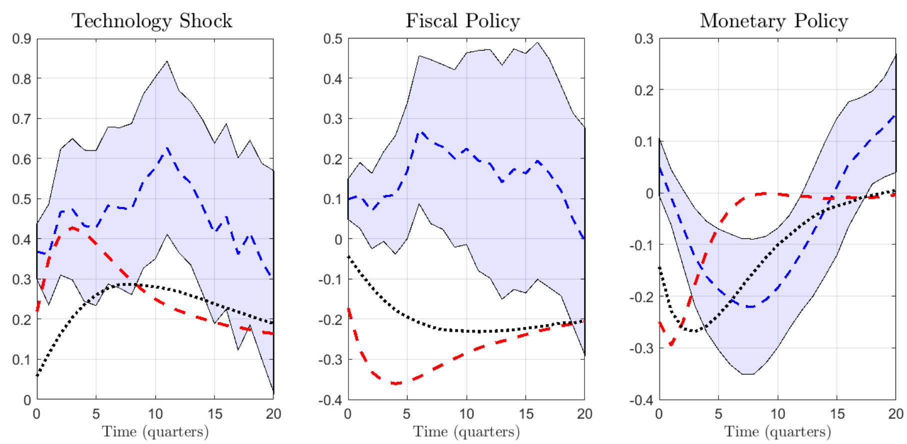

(b) Consumption

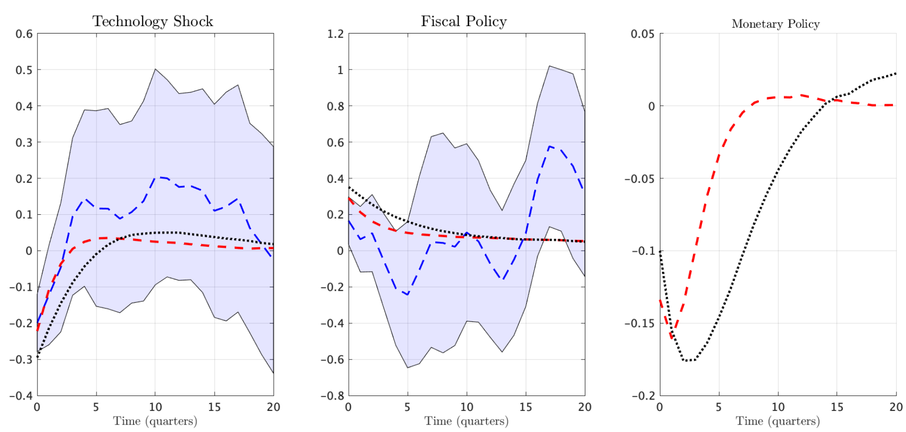

(c) Hours Worked

Empirical IRFs IRFs $\hat{\Theta}_{LP}$ IRFs $\hat{\Theta}_{SW07}$

As noted earlier, the estimated model does well in matching the investment response to technology shocks but not so well to fiscal and monetary innovations. These findings remain in the case of the joint estimation. However, through exploiting the variation of the three shocks, the estimated model is able to improve the match of the investment response to a fiscal policy shock, although it is not sufficient to reconcile it completely with the data. Moreover, the investment does not longer increase in the medium run as a result of a monetary contraction.

Looking at the untargeted output response, the Smets and Wouters (2007) parameterization again produces a deeper and longer reduction in output in response to a monetary contraction compared to these estimates based upon the shocks together. A key difference is in the frequency of wage and price adjustment. The responses to technology and fiscal policy shocks are similar to those obtained in the shock-by-shock estimations.

Regarding consumption, the model at our parameter estimates is still able to capture its response to a technology shock better than at Smets and Wouters (2007) parameters, even when we do not target it directly. However, its response to the government spending shock is better captured when estimated individually. This is probably reflected in the lower intertemporal elasticity of substitution obtained at the joint estimation.

Finally, the untargeted response of hours worked is similar to those obtained when targeting its response to either the technology or the fiscal shock. For the monetary policy shock, we do not have an empirical counterpart. In any case, it is reassuring that the response is similar to that obtained at Smets and Wouters (2007) parameters. That is, hours drop temporarily and then converge towards zero. Nevertheless, the initial drop is less prolonged and the rate at which the effect dies out is bigger at our parameters than at Smets and Wouters (2007).

5.4.1. What do the parameter differences imply for the responses?

To dig a bit deeper into differences in estimates compared to Smets and Wouters (2007), Figure 8 shows the impulse responses to the three shocks (columns) for alternative parametrizations (rows). We construct each of these four alternatives by changing one at a time the curvature of utility, the habit parameter and the Calvo's probabilities to be equal to the Smets and Wouters (2007) estimates instead of ours. So, for example, in the first row the IRFs depicted by the blue dash line are labeled $\hat{\Theta}_{-\sigma_c}$ indicating that these are the median estimated parameters from the joint estimation except for the treatment of $\hat{\sigma}_c$, which corresponds to the elasticity estimated by Smets and Wouters (2007) instead of ours. The model is simulated (not re-estimated) and the resulting IRFs are plotted along with the baseline IRFs, i.e. those generated at our median estimates from the joint estimation (solid black lines).

Looking across these combinations, we see some cases where differences in the estimated parameters matter for the economic responses. Looking first at the estimated wage and price flexibility, the output and investment responses to monetary policy are different in the two parameterizations because of the increased frequency of adjustment in our estimated model. The investment response is about half that of Smets and Wouters (2007) in our model. Both wage and price flexibility impact the magnitude of the response. Interestingly, the magnitude of the effects of technology shocks on hours is not very sensitive, at least for the first 5 quarters, to differences in price flexibility in the two models.

Figure 8: Parameter Decomposition: y_t, c_t, i_t, n_t IRFs

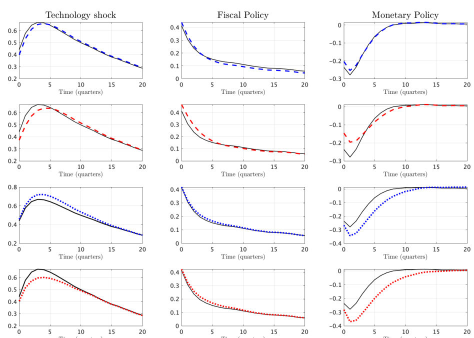

(a) Output

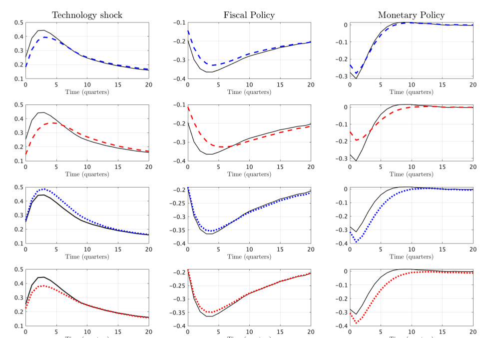

(b) Consumption

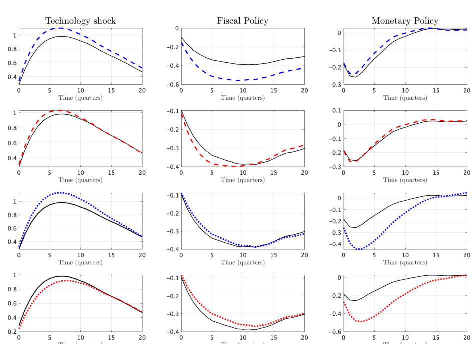

(c) Investment

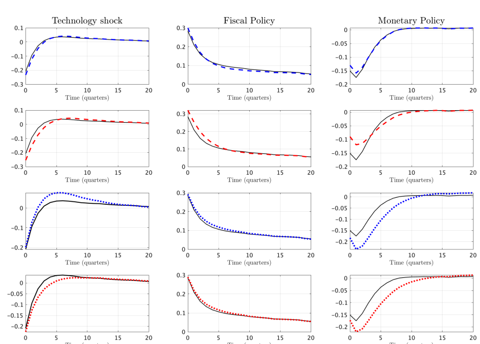

(d) Hours Worked

IRFs $\hat{\Theta}$ IRFs $\hat{\Theta}_{-\sigma_c}$ IRFs $\hat{\Theta}_{-h}$ IRFs $\hat{\Theta}_{-\xi_w}$ IRFs $\hat{\Theta}_{-\xi_p}$

The differences in preference parameters, both the elasticity of substitution and the habit, matter most for the response of consumption to all three shocks as well as the response of investment to a government spending shock. Recall that our estimates are much closer to log utility so that consumption smoothing is less important to the household. Accordingly, the consumption responses produced by our model are generally larger except for the response of consumption to a monetary innovation, which seems insensitive to this elasticity. Perhaps this reflects that channels of monetary policy other than intertemporal substitution are operative. Both models produce hump-shaped responses. Our predicted response has a peak that is earlier and more pronounced mostly due to differences in the estimate of h.

6. Conclusion

This paper studies the use of LP coefficients in an indirect inference approach to structural estimation. Monte Carlo analysis shows that the theoretical responses at the estimated parameters are much closer to the truth than if one relies in the more traditional approach that uses VAR coefficients. Moreover, the time spent in the estimation is significantly reduced, an important consideration for large scale DSGE models.

The application of this approach to the estimation of the Smets and Wouters (2007) model, despite successful, has revealed some shortcomings of the model. For example, the model is not able to replicate the large initial response and the more delayed hump in consumption in response to a technology shock, even if we target it. The responses of investment and consumption to a government spending shock are also at odds with the data whenever the empirical estimated shock is identified recursively. Nonetheless, its most relevant and obvious flaw is its inability to capture the state dependent effects of monetary innovations, specially during a recession. Our parameter estimates suggest that Calvo (1983) pricing may be behind it since we obtain a much lower degree of wage and price stickiness in that case.

Based on these findings, we believe that state-dependent rather than time-dependent pricing will help in reconciling the model with the data. In fact, if the model is truly state-dependent one should also be able to infer its parameters by jointly targeting the state-dependent LP coefficients, for example those coming from the state-dependent responses to monetary policy documented in Tenreyro and Thwaites (2016) or Ascari and Haber (2022).

References

Ascari, G. and Haber, T. (2022). Non-Linearities, State-Dependent Prices and the Transmission Mechanism of Monetary Policy. *The Economic Journal*, 132(641):37–57.

Ben Zeev, N. and Pappa, E. (2017). Chronicle of A War Foretold: The Macroeconomic Effects of Anticipated Defence Spending Shocks. *The Economic Journal*, 127(603):1568–1597.

Blanchard, O. and Perotti, R. (2002). An Empirical Characterization of the Dynamic Effects of Changes in Government Spending and Taxes on Output. *The Quarterly Journal of Economics*, 117(4):1329–1368. Publisher: Oxford University Press.

Calvo, G. A. (1983). Staggered prices in a utility-maximizing framework. *Journal of Monetary Economics*, 12(3):383–398.

Canova, F. and Sala, L. (2009). Back to square one: Identification issues in DSGE models. *Journal of Monetary Economics*, 56(4):431–449.

Christiano, L., Eichenbaum, M., and Evans, C. (2005). Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy. *Journal of Political Economy*, 113(1):1–45.

Fernández-Villaverde, J., Rubio-Ramírez, J. F., Cogley, T., and Schorfheide, F. (2007). How Structural Are Structural Parameters? [with Comments and Discussion]. *NBER Macroeconomics Annual*, 22:83–167. Publisher: The University of Chicago Press.

Francis, N., Owyang, M. T., Roush, J. E., and DiCecio, R. (2014). A flexible finite-horizon alternative to long-run restrictions with an application to technology shocks. *The Review of Economics and Statistics*, 96(4):638–647. Publisher: The MIT Press.

Galí, J., López-Salido, J. D., and Vallés, J. (2007). Understanding the Effects of Government Spending on Consumption. *Journal of the European Economic Association*, 5(1):227–270.

Herbst, E. P. and Johannsen, B. K. (2021). Bias in Local Projections. *Finance and Economics Discussion Series*, 2020(010r1):1–62.

Jordà, O. (2005). Estimation and Inference of Impulse Responses by Local Projections. *The American Economic Review*, 95(1):161–182. Publisher: American Economic Association.

Li, D., Plagborg-Møller, M., and Wolf, C. (2022). Local Projections vs. VARs: Lessons From Thousands of DGPs. page w30207.

Plagborg-Møller, M. and Wolf, C. K. (2021). Local Projections and VARs Estimate the Same Impulse Responses. *Econometrica*, 89(2):955–980. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.3982/ECTA17813.

Ramey, V. A. (2011). Identifying Government Spending Shocks: It's all in the timing. *The Quarterly Journal of Economics*, 126(1):1–50. Publisher: Oxford University Press.

Ramey, V. A. (2016). Chapter 2 - Macroeconomic Shocks and Their Propagation. In Taylor, J. B. and Uhlig, H., editors, *Handbook of Macroeconomics*, volume 2, pages 71–162. Elsevier.

Romer, C. D. and Romer, D. H. (2004). A New Measure of Monetary Shocks: Derivation and Implications. *The American Economic Review*, 94(4):1055–1084. Publisher: American Economic Association.

Ruge-Murcia, F. (2012). Estimating nonlinear DSGE models by the simulated method of moments: With an application to business cycles. *Journal of Economic Dynamics and Control*, 36(6):914–938.

Ruge-Murcia, F. J. (2007). Methods to estimate dynamic stochastic general equilibrium models. *Journal of Economic Dynamics and Control*, 31(8):2599–2636.

Smets, F. and Wouters, R. (2003). An Estimated Dynamic Stochastic General Equilibrium Model of the Euro Area. *Journal of the European Economic Association*, 1(5):1123–1175. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1162/154247603770383415.

Smets, F. and Wouters, R. (2007). Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach. *American Economic Review*, 97(3):586–606.

Smith, A. A. (1993). Estimating Nonlinear Time-Series Models Using Simulated Vector Autoregressions. *Journal of Applied Econometrics*, 8:S63–S84. Publisher: Wiley.

Tenreyro, S. and Thwaites, G. (2016). Pushing on a String: US Monetary Policy Is Less Powerful in Recessions. *American Economic Journal: Macroeconomics*, 8(4):43–74. Publisher: American Economic Association.

A. Log-Linearized Equilibrium Conditions

- The aggregate resource constraint:

$$\hat{y}_t = c_y \hat{c}_t + i_y \hat{i}_t + z_y \hat{z}_t + \varepsilon_t^g \tag{A.1}$$

- The consumption Euler equation:

$$\begin{aligned} \hat{c}_t = \frac{h/\gamma}{1+h/\gamma}\hat{c}_{t-1} + \frac{1}{1+h/\gamma}\mathbb{E}_t\hat{c}_{t+1} + \frac{wl_c\,(\sigma_c-1)}{\sigma_c(1+h/\gamma)}\left(\hat{l}_t - \mathbb{E}_t\hat{l}_{t+1}\right) + \\ - \frac{1-h/\gamma}{(1+h/\gamma)\sigma_c}\left(\hat{r}_t - \mathbb{E}_t\hat{\pi}_{t+1}\right) - \frac{1-h/\gamma}{(1+h/\gamma)\sigma_c}\varepsilon_t^b \end{aligned} \tag{A.2}$$

- The investment Euler equation:

$$\hat{i}_t = \frac{1}{1+\beta\gamma^{(1-\sigma_c)}}\hat{i}_{t-1} + \frac{\beta\gamma^{(1-\sigma_c)}}{1+\beta\gamma^{(1-\sigma_c)}}\mathbb{E}_t\hat{i}_{t+1} + \frac{1}{\varphi\gamma^2\left(1+\beta\gamma^{(1-\sigma_c)}\right)}\hat{q}_t + \varepsilon_t^i \tag{A.3}$$

- The arbitrage equation for the value of capital:

$$\hat{q}_t = \beta(1-\delta)\gamma^{-\sigma_c}\mathbb{E}_t\hat{q}_{t+1} - \hat{r}_t + \mathbb{E}_t\hat{\pi}_{t+1} + \left(1-\beta(1-\delta)\gamma^{-\sigma_c}\right)\mathbb{E}_t\hat{r}_{t+1}^k - \varepsilon_t^b \tag{A.4}$$

- The aggregate production function:

$$\hat{y}_t = \Phi\left(\alpha\hat{k}_t^s + (1-\alpha)\hat{l}_t + \varepsilon_t^a\right) \tag{A.5}$$

- Capital services:

$$\hat{k}_t^s = \hat{k}_{t-1} + \hat{z}_t \tag{A.6}$$

- Capital utilization:

$$\hat{z}_t = \frac{1-\psi}{\psi}\hat{r}_t^k \tag{A.7}$$

- The accumulation of installed capital:

$$\hat{k}_t = \frac{(1-\delta)}{\gamma}\hat{k}_{t-1} + (1-(1-\delta)/\gamma)\hat{i}_t + (1-(1-\delta)/\gamma)\varphi\gamma^2\left(1+\beta\gamma^{(1-\sigma_c)}\right)\varepsilon_t^i \tag{A.8}$$

- Cost minimization by firms implies that the price mark up:

$$\hat{\mu}_t^p = \alpha\left(\hat{k}_t^s - \hat{l}_t\right) - \hat{w}_t + \varepsilon_t^a \tag{A.9}$$

- New Keynesian Phillips curve:

$$\begin{aligned}\hat{\pi}_t = & \frac{\beta\gamma^{(1-\sigma_c)}}{1+\iota_p\beta\gamma^{(1-\sigma_c)}}\mathbb{E}_t\hat{\pi}_{t+1} + \frac{\iota_p}{1+\beta\gamma^{(1-\sigma_c)}}\hat{\pi}_{t-1} + \\ & - \frac{\left(1-\beta\gamma^{(1-\sigma_c)}\xi_p\right)(1-\xi_p)}{\left(1+\iota_p\beta\gamma^{(1-\sigma_c)}\right)(1+(\Phi-1)\varepsilon_p)\,\xi_p}\hat{\mu}_t^p + \varepsilon_t^p\end{aligned} \tag{A.10}$$

- Cost minimization by firms implies that the rental rate of capital:

$$\hat{r}_t^k = \hat{l}_t + \hat{w}_t - \hat{k}_t^s \tag{A.11}$$

- In the monopolistically competitive labor market, the wage mark-up

$$\hat{\mu}_t^w = \hat{w}_t - \sigma_l\hat{l}_t - \frac{1}{1-h/\gamma}\left(\hat{c}_t - h/\gamma\hat{c}_{t-1}\right) \tag{A.12}$$

- Wage adjustment:

$$\begin{aligned}\hat{w}_t = & \frac{\beta\gamma^{(1-\sigma_c)}}{1+\beta\gamma^{(1-\sigma_c)}}\left(\mathbb{E}_t\hat{w}_{t+1} + \mathbb{E}_t\hat{\pi}_{t+1}\right) + \frac{1}{1+\beta\gamma^{(1-\sigma_c)}}\left(\hat{w}_{t-1} - \iota_w\hat{\pi}_{t-1}\right) + \\ & - \frac{1+\beta\gamma^{(1-\sigma_c)}\iota_w}{1+\beta\gamma^{(1-\sigma_c)}}\hat{\pi}_t - \frac{\left(1-\beta\gamma^{(1-\sigma_c)}\xi_w\right)(1-\xi_w)}{\left(1+\beta\gamma^{(1-\sigma_c)}\right)\left(1+(\lambda_w-1)\,\epsilon_w\right)\xi_w}\hat{\mu}_t^w + \varepsilon_t^u\end{aligned} \tag{A.13}$$

- Monetary policy reaction function:

$$\hat{r}_t = \rho\hat{r}_{t-1} + (1-\rho)\left(r_\pi\hat{\pi}_t + r_y\left(\hat{y}_t - \hat{y}_t^*\right)\right) + r_{\Delta y}\left(\left(\hat{y}_t - \hat{y}_t^*\right) - \left(\hat{y}_{t-1} - \hat{y}_{t-1}^*\right)\right) + \varepsilon_t^r \tag{A.14}$$

B. The Monte Carlo Study: Results in Detail

B.1. The Local Projection Approach to Indirect Inference

Table 7: SMM estimates using LP - IRFs

	$\hat{\sigma}_c$	$\hat{h}$	$\hat{\sigma}_l$	$\hat{\varphi}$	$\hat{\xi}_w$	$\hat{\xi}_p$	$\hat{\iota}_w$	$\hat{\iota}_p$
	Technology shock, ε_t^a							
Mean	1.23	0.82	2.81	5.91	0.57	0.62	0.48	0.15
Bias	-0.03	0.02	0.29	-0.40	-0.13	-0.04	-0.10	-0.09
Std dev.	0.26	0.10	0.61	1.74	0.15	0.07	0.17	0.05
RMSE	0.26	0.10	0.67	1.78	0.20	0.08	0.19	0.10
	Fiscal Policy, ε_t^g							
Mean	1.40	0.80	2.60	5.90	0.54	0.46	0.52	0.17
Bias	0.14	0.00	0.08	-0.41	-0.16	-0.20	-0.06	-0.07
Std dev.	0.23	0.09	0.70	1.85	0.11	0.14	0.17	0.05
RMSE	0.27	0.09	0.70	1.89	0.19	0.25	0.18	0.09
	Monetary Policy, ε_t^m							
Mean	1.38	0.79	2.36	5.52	0.62	0.53	0.47	0.16
Bias	0.12	-0.01	-0.16	-0.79	-0.08	-0.13	-0.11	-0.08
Std dev.	0.26	0.06	0.77	1.60	0.14	0.14	0.17	0.05
RMSE	0.28	0.06	0.79	1.78	0.15	0.19	0.20	0.09
	Selected Responses to All Shocks							
Mean	1.29	0.81	2.56	5.75	0.56	0.59	0.47	0.15
Bias	0.03	0.01	0.04	-0.56	-0.14	-0.07	-0.11	-0.09
Std dev.	0.25	0.09	0.74	1.40	0.14	0.10	0.17	0.05
RMSE	0.25	0.10	0.74	1.50	0.19	0.12	0.20	0.10

B.2. The Structural Vector Autoregression Approach to Indirect Inference

Table 8: SMM estimates using SVAR - IRFs

	$\hat{\sigma}_c$	$\hat{h}$	$\hat{\sigma}_l$	$\hat{\varphi}$	$\hat{\xi}_w$	$\hat{\xi}_p$	$\hat{\iota}_w$	$\hat{\iota}_p$
	Technology shock, ε_t^a							
Mean	1.28	0.82	2.36	6.81	0.64	0.62	0.43	0.14
Bias	0.02	0.02	-0.16	0.50	-0.06	-0.04	-0.15	-0.10
Std dev.	0.20	0.08	0.77	1.46	0.10	0.07	0.15	0.05
RMSE	0.20	0.08	0.78	1.54	0.12	0.08	0.21	0.11
	Fiscal Policy, ε_t^g							
Mean	1.32	0.82	2.65	5.44	0.68	0.47	0.52	0.18
Bias	0.06	0.02	0.13	-0.87	-0.02	-0.19	-0.06	-0.06
Std dev.	0.25	0.11	0.68	1.82	0.07	0.16	0.16	0.05
RMSE	0.25	0.11	0.70	2.02	0.08	0.25	0.17	0.07
	Monetary Policy, ε_t^m							
Mean	1.32	0.79	2.39	5.54	0.66	0.48	0.50	0.17
Bias	0.06	-0.01	-0.13	-0.77	-0.04	-0.18	-0.08	-0.07
Std dev.	0.27	0.06	0.80	1.63	0.09	0.14	0.17	0.05
RMSE	0.27	0.06	0.81	1.80	0.10	0.23	0.18	0.08
	Selected Responses to All Shocks							
Mean	1.21	0.85	2.57	6.45	0.58	0.58	0.50	0.15
Bias	-0.05	0.05	0.05	0.14	-0.12	-0.08	-0.08	-0.09
Std dev.	0.24	0.08	0.75	1.39	0.12	0.10	0.16	0.05
RMSE	0.24	0.09	0.75	1.40	0.17	0.13	0.18	0.10

B.3. The Role of the Sample Size

We set $T = 300$ observations because that is the sample size chosen by Jordà (2005) in the Monte Carlo study of his seminal paper. However, most empirical applications that used identified shocks as regressors within the LP framework employ fewer observations. For such sample sizes LPs suffer from small sample bias (Herbst and Johannsen, 2021). Thus, we check if our Monte Carlo results still hold in small samples.

We generate a new repeated dataset consisting of time series of length 100, which corresponds to 25 years of quarterly data. This is roughly the most common sample size used in applied macroeconometric papers.

The new simulated dataset is used to generate S vectors of the true data moments. Simulated moments are computed on a sample that is 10 times as large, i.e. $T^s = 1,000$. All remaining hyper-parameters used in the optimization stage are unchanged. Notice, however, that the optimal weighting matrix is now computed over repeated samples of length 100.

Figure B.3.1: Relative performance in terms of RMSE with T = 100

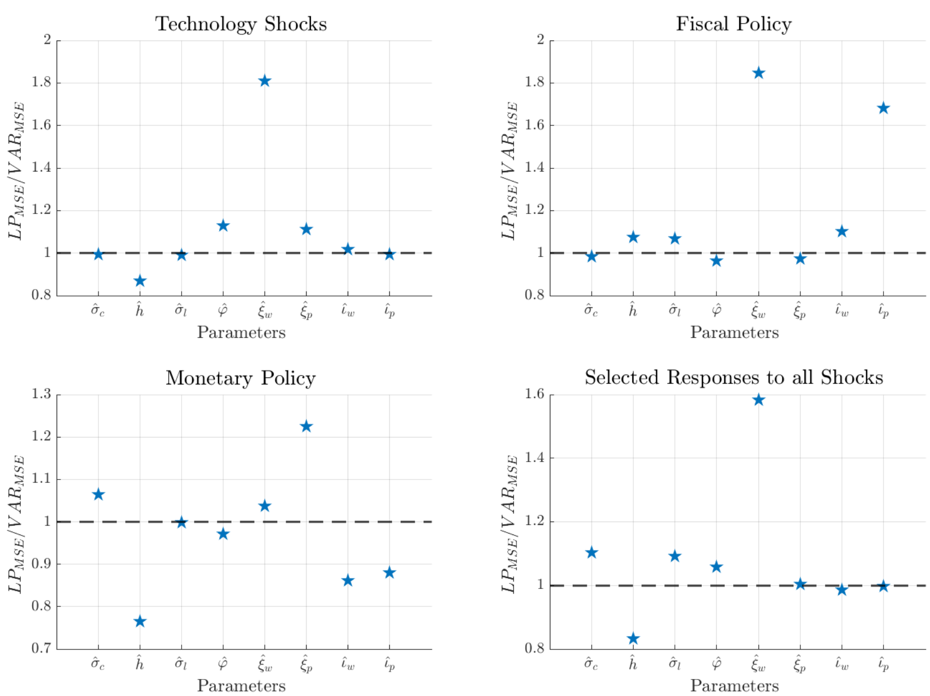

Table 9: Overall performance: Estimated Impulse Responses

	Local Projections			Vector Autoregression		
	Avg. J	Max. J	Time (in min.)	Avg. J	Max. J	Time (in min.)
	Observed Sample Size, T = 100					
Technology Shock	86.70	142.01	6.63	76.27	251.59	22.96
Fiscal Policy	93.01	134.29	7.07	88.15	258.36	24.61
Monetary Policy	90.73	134.65	7.19	92.86	257.67	25.39
Selected Responses	88.77	162.77	6.62	87.09	260.48	23.99
	Observed Sample Size, T = 300					
Technology Shock	87.23	117.00	28.62	82.80	247.94	67.98
Fiscal Policy	87.72	129.96	24.68	86.28	251.90	52.49
Monetary Policy	88.58	121.48	23.38	82.77	221.87	53.04
Selected Responses	86.56	128.65	20.03	82.63	240.07	72.42

Table 10: Overall performance: Model Impulse Responses

	Local Projections		Vector Autoregression	
	Avg. J^*	Max. J^*	Avg. J^*	Max. J^*
	Observed Sample Size, T = 100			
Technology Shock	2.33	7.60	668.75	18116.41
Fiscal Policy	1.38	5.82	193.71	3841.47
Monetary Policy	2.33	10.17	1633.33	12961.19
Selected Responses	15.70	65.83	3079.22	40779.98
	Observed Sample size, T = 300			
Technology Shock	2.57	9.43	34.67	228.41
Fiscal Policy	3.05	13.88	58.12	692.14
Monetary Policy	2.71	16.89	178.17	853.72
Selected Responses	8.37	44.69	230.46	1130.58

Results are summarized in Figure B.3.1, and tables 9 and 10. In light of this evidence, we can confirm that the small sample bias associated with LPs is not an issue for indirect inference. In fact, all our results from the Monte Carlo hold for $T = 100$. That is, RMSE is smaller for the majority of parameters and the overall fit in terms of the J statistic is better when using the SVAR approach, however, when looking at the distance with respect to the true impulse responses, the local projection approach does a much better job.

B.4. The Role of Lag Length

Another crucial choice is the number of lags used in the VAR and the number of lags used as controls in the LP regression. Following the equivalence results in Plagborg-Møller and Wolf (2021), we set them to be equal because LP-IRFs and SVAR-IRFs approximately agree until horizon p in finite samples.

In the baseline we set $p = 4$ because it is a common parctice for VARs that are estimated on quarterly data.[17] However, this is an heuristic choice. Consequently, we also repeat the Monte Carlo study for alternative lag lengths, $p \in \{2, 4, 8, 12\}$. For computational ease, and given that our results still hold in smaller samples, we set $T = 100$ for this tests.

Intuitively as one increases the lag length, the two moment generating functions will be more similar to each other. Therefore, the estimated parameters should also get closer as one increases p. In the limit, when $p \geq H$, then the two approaches should estimate the same parameters. In that case, the RMSE ratio reported in Figure 2 panel (b) should be equal to 1.

Figure B.4.1: Relative performance in terms of RMSE for different lag lengths

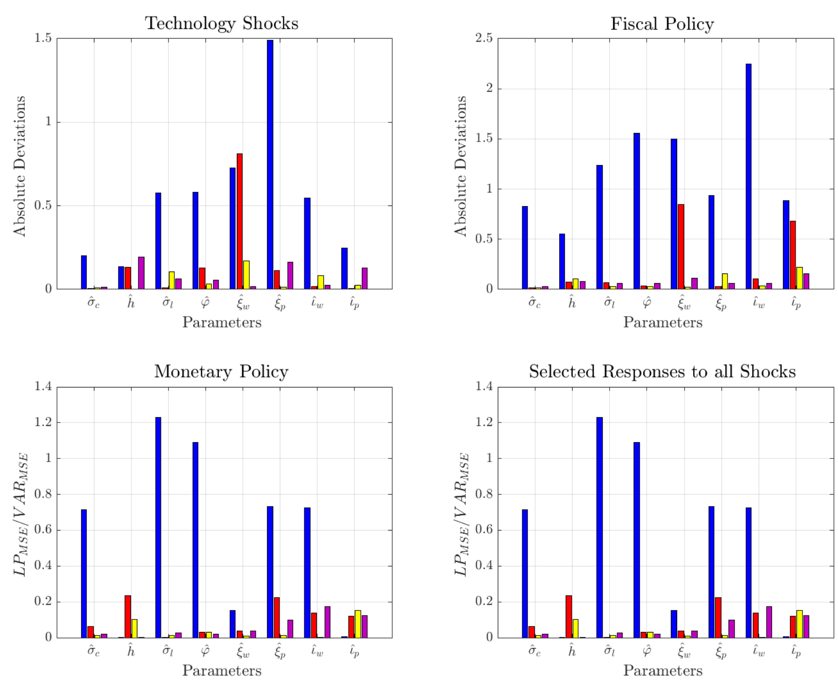

[17] Recall that in the Smets and Wouters (2007) model one period corresponds to one quarter.

Following this intuition, we compute the absolute deviation from 1 of the RMSE ratio for each of the estimated parameters and lag lengths. Results are depicted in Figure B.4.1. For $p = 8$ (yellow) and $p = 12$ (magenta), this statistic is very close to 0 for all parameters and scenarios; while for $p = 2$ (blue) and $p = 4$ (red) this statistic is far from 0, reflecting the fact that the differences across the two moment generating functions are decreasing in p.

B.5. Additional Figures

B.5.1. Fan Charts: Fiscal and Monetary Policy Shocks

Figure B.5.1: Fiscal Policy Shock

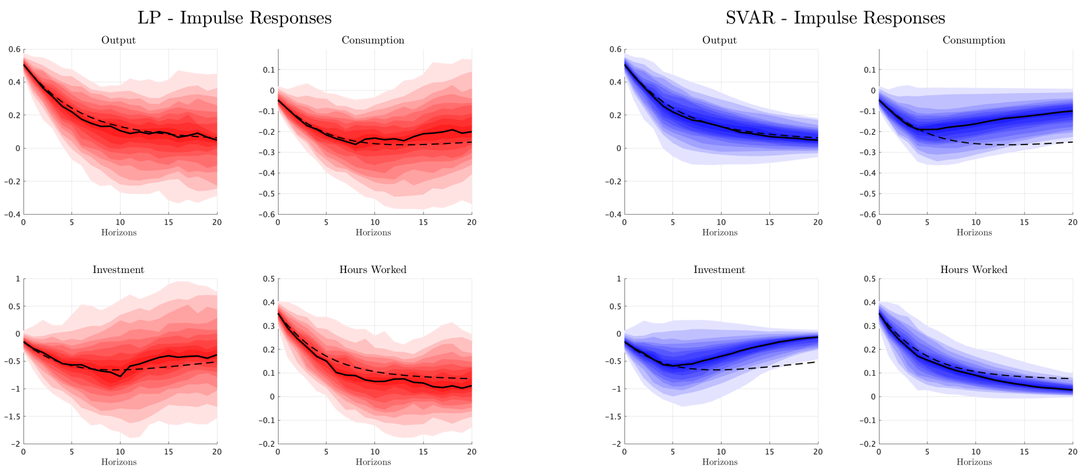

Monte Carlo Median IRFs   Model/True IRFs

Figure B.5.2: Monetary Policy Shock

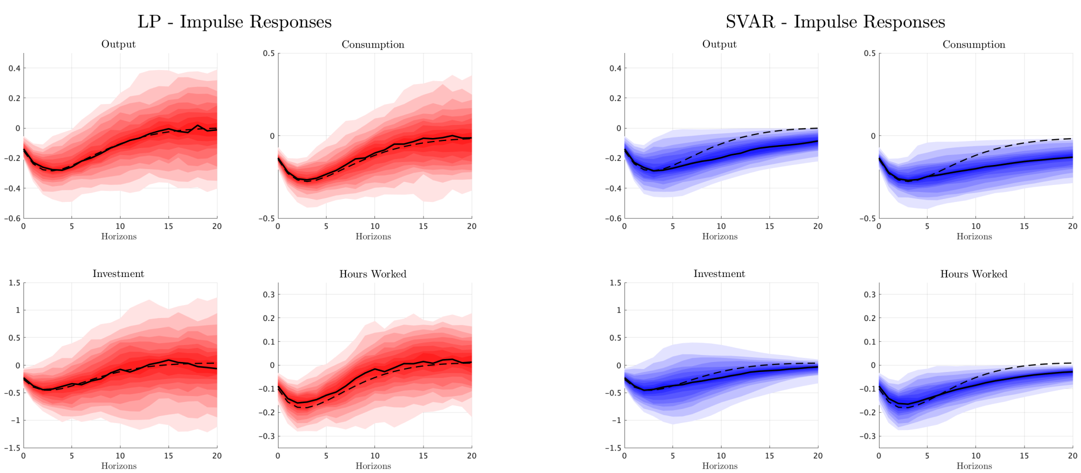

Monte Carlo Median IRFs   Model/True IRFs

B.5.2. Sensitivity of the Moments to Changes in the Parameters

$$\text{sensitivity}(\theta) = \frac{M(\theta+\Delta) - M(\theta)}{\theta+\Delta-\theta} = \frac{M(\theta+\Delta) - M(\theta)}{\Delta} \tag{B.1}$$

Figure B.5.1: LP-IRFs sensitivity

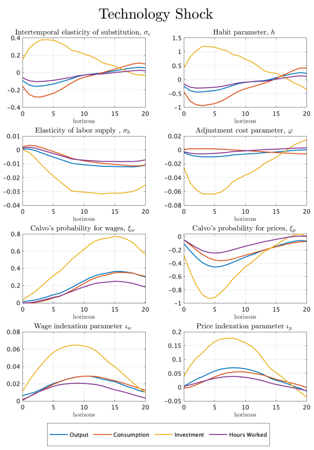

(a) Technology shock

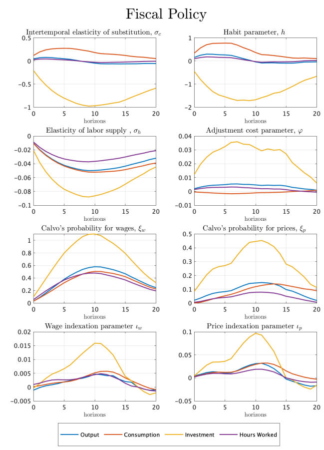

(b) Fiscal Policy

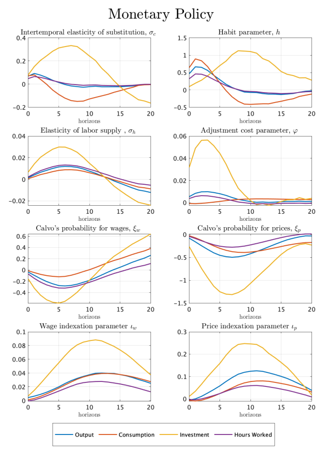

(c) Monetary Policy

Figure B.5.2: SVAR-IRFs sensitivity

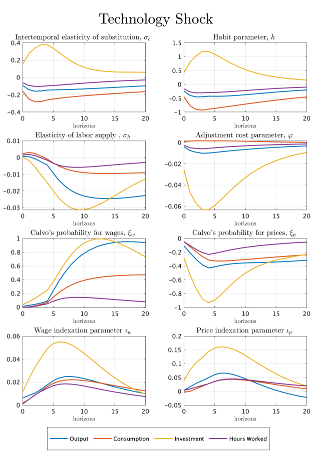

(a) Technology shock

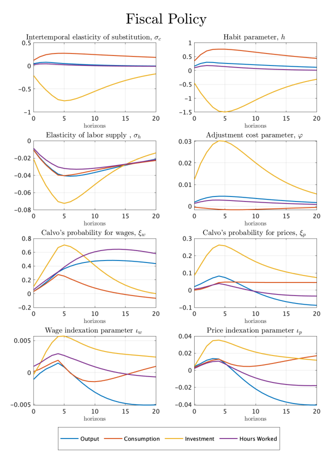

(b) Fiscal Policy

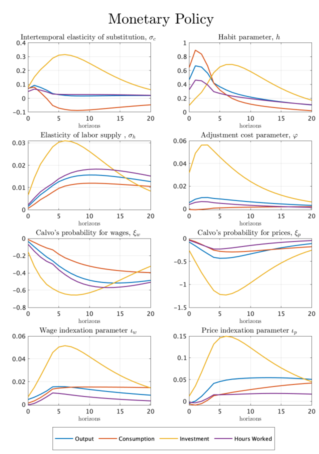

(c) Monetary Policy

B.5.3. Elasticity of the Moments to Changes in the Parameters

$$\text{elasticity}(\theta) = \frac{\frac{M(\theta+\Delta)-M(\theta)}{M(\Theta)}}{\frac{\theta+\Delta-\theta}{\theta}} = \frac{\frac{M(\theta+\Delta)-M(\theta)}{M(\Theta)}}{\Delta/\theta} \tag{B.2}$$

Figure B.5.3: LP-IRFs elasticities

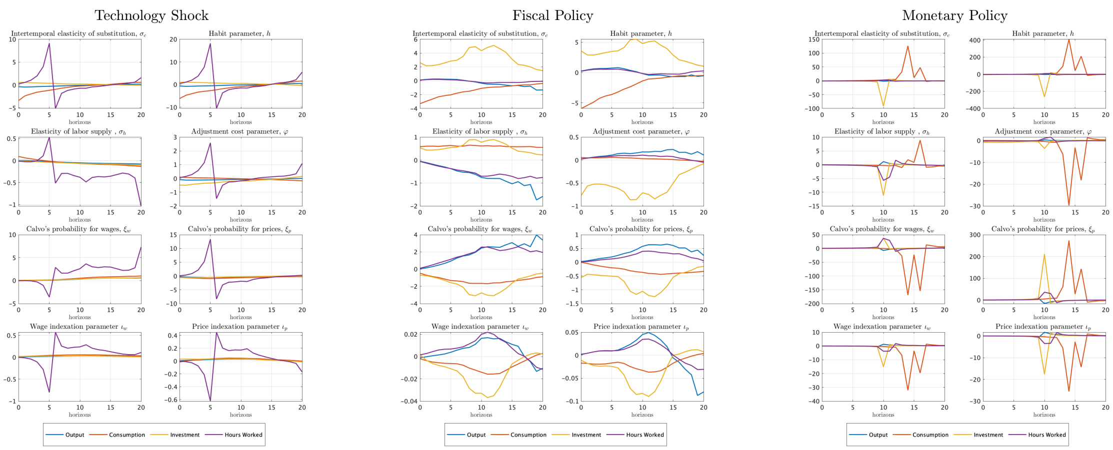

(a) Technology shock (b) Fiscal Policy (c) Monetary Policy

Figure B.5.4: SVAR-IRFs elasticities

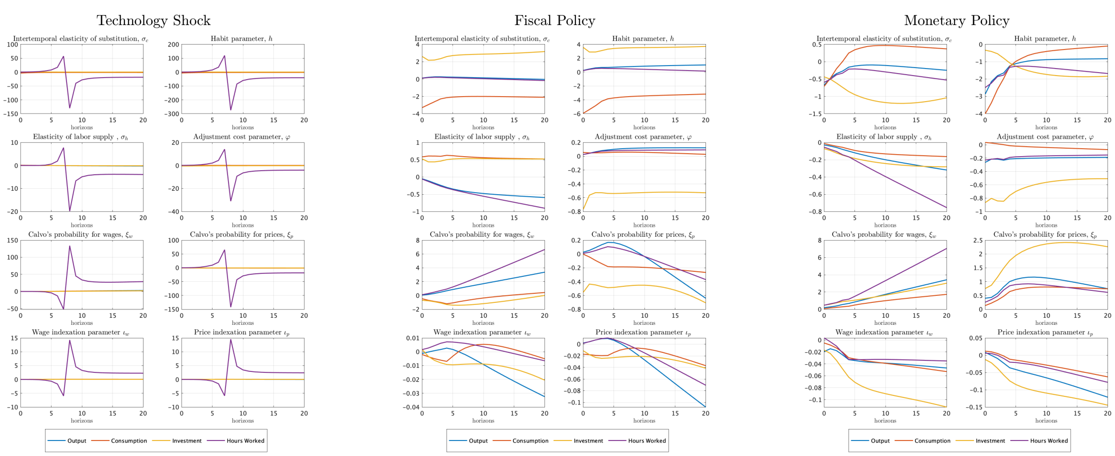

(a) Technology shock (b) Fiscal Policy (c) Monetary Policy