Indirect Inference: A Local Projection Approach *

Juan Castellanos † Russell W. Cooper ‡

This draft: August 2022

First draft: April 2022

Abstract

This paper proposes an alternative way of generating moments to estimate the economic parameters of a Dynamic Stochastic General Equilibrium (DSGE) model using indirect inference methods. Until now, the most common auxiliary models that have been used to form a criterion function are Vector Autoregressions (VARs). However, given the increased popularity of Local Projections (LPs) to study the propagation of structural shocks, we study the use of LPs as auxiliary models for indirect inference. We use Monte Carlo methods to compare the performance of these two moment generating functions in a specific application involving the estimation of the parameters of the Smets and Wouters (2007) model. Then, we test our LP approach to indirect inference in a real world application, using empirically estimated responses of key macro variables to technology, fiscal and monetary shocks from the literature.

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.

We propose to use LPs as a way of summarizing key features of the data when estimating the structural parameters of a model in an indirect inference exercise. One issue with this approach is its comparison to the more common structural VAR approach. Plagborg-Møller and Wolf (2021) prove that VARs and LPs estimate the same impulse responses in population. Therefore, it may seem that using either VAR or LP should not matter for indirect inference. However, the finite sample properties of these two estimators differ. In particular, they show that when p lags of the data are included in the VAR and as controls in the LP, impulse response functions (IRFs) approximately agree out to horizon p, but at longer horizons $h > p$ there is a bias-variance trade-off (Li et al., 2021). This paper is complementary to Plagborg-Møller and Wolf (2021) and Li et al. (2021) in that our interest is as well in the performance of the LP approach. But, our focus is on the estimation of structural parameters governing, for example, tastes and technology, rather than the impulse response functions *per se*. Smith (1993) illustrates the use of an autoregressive structure as an auxiliary model in a SMM estimation exercise.

Our goal is to study the comparable properties of parameter estimation using as moments the coefficients from an LP generated impulse response function. Given the findings of Plagborg-Møller and Wolf (2021) for finite samples, it is of interest to study the implications for estimation of underlying parameters using the LP rather than the VAR as the auxilary model for indirect inference.

Accordingly, we ask two questions. First, how do the properties of the parameter estimates depend on the use of a VAR as opposed to a LP as an auxiliary model? Second, how do the estimates based upon the LP auxiliary model from actual data compare to other methods?

The first question is answered through a Monte Carlo exercise, described in Section 4. For that experiment, the data generating mechanism is a parameterized version of the Smets and Wouters (2007) model. The structural parameters we focus on characterize household preferences, capital adjustment costs and the determination of wage and price rigidities. These parameters are fixed at the Smets and Wouters (2007) mean estimated values reported in Tables 1A and 1B in their paper. The model is then simulated to create time series. Using these time series, we estimate a VAR and use this 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.

From the Monte Carlo exercise, we find 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 $h > p$.

The second question is answered by using LP estimates from these same three types of shocks from actual data, rather than from the Smets and Wouters (2007) 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. 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 to a contractionary shock during a recession, neither at their estimates nor at ours. On the other hand, the model is able to match these responses during an expansion either at our or their parameter estimates.

We also estimate parameters from the joint response to 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.[1] These estimates are slightly different to those obtained by matching responses individually. Nevertheless, we are able to match both targeted and untargeted responses to the three shocks.

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

2. Economic Model

The analysis builds on the model formulated and estimated in Smets and Wouters (2007). While other models may serve the same purpose, this structure captures many of the central channels of monetary and fiscal policy. Its parameters were estimated using Bayesian techniques. For the first part of our analysis, we take those 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. Consequently, this explains why we use this model and not the traditional real business cycle as in the seminal paper by Smith (1993).

This section summarizes 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.[2]

Monetary policy is modelled 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 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 $w_t = \{\tilde{x}_t, \tilde{y}_t\}$.

[2] 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 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 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.

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

The indirect inference approach is a form of minimum distance estimation. The parameter estimates are given by:

$$\hat{\Theta} = \arg\min_{\Theta} \left(M^d - M^s(\Theta)\right)' W \left(M^d - M^s(\Theta)\right) \tag{6}$$

In (6), M^d are moments from the data and $M^s(\Theta)$ are the counterparts from simulated data, where Θ is the vector of structural parameters. In our application, these are the parameters characterizing household preferences, wage setting and price setting of firms in the Smets and Wouters (2007) model. In this quadratic form, W is a weighting matrix.

In the indirect inference approach, the moments taken from the data are regression coefficients of an auxiliary model. As in Smith (1993), the auxiliary model could be the coefficients from a VAR (or the associated impulse response, denoted θ_h in Definition 2). An alternative, explored here, is to use the LP as the source of moments, matching the β_h coefficients from Definition 1.

4. A Monte Carlo Study

This section uses the VAR and LP approaches to indirect inference in order to estimate a subset of the parameters of the Smets and Wouters (2007) model using repeated samples from the data generating process associated with their DSGE model. The goal of this Monte Carlo study is to compare the performance of these two moment generating functions in a controlled setting where we know the true parameter vector, $\Theta = \Theta^*$.

4.1. Setting up the Monte Carlo: The Data Generating Process

For the Monte Carlo study, 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 $\{\varepsilon_t^a, \varepsilon_t^g, \varepsilon_t^g\}$. The log-linearized equilibrium conditions of the Smets and Wouters (2007) model are reproduced in Appendix A.

The "observed" sample size is set at $T = 300$, which coincides with the sample size chosen by Jordà (2005). We are aware that sample size in the time dimension is typically much smaller in empirical applications and more importantly that LPs can be biased for small samples as shown by Herbst and Johannsen (2021).[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).[4]

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

Here: σ_c : intertemporal elasticity of substitution, h : habit parameter, σ_l : elasticity of labor supply, φ : investment adjustment cost parameter, ξ_w, ξ_p : Probabilities of no adjustment for wage and prices, ι_w, ι_p : Degree of wage and price indexation to past inflation.

4.2. The Moment Generating Functions

Indirect inference applications typically use an auxiliary econometric model that help researchers summarize key features of the data that they wish to match with their economic model. In most cases, and specially for representative agent models (like most DSGEs), a VAR is chosen. This approach was popularized by Smith (1993). In particular, he proposes to match all the coefficients of the VAR as well as the associated variance-covariance matrix of the error term. Alternatively, as we propose in this paper, one could also use LP coefficients.

When comparing these two approaches in our indirect inference exercise, we depart from Smith (1993) in that we simply use the estimated impulse responses instead of all the estimated coefficients. We prefer this approach because IRFs are excessively parameterized in the LP framework, resulting in an absurd number of moments. However, one should not confuse this with IRFs matching, another popular approach to estimate the parameters of DSGE models, as for example used in Christiano et al. (2005) or Christiano et al. (2015). The subtle, but important difference, between IRF matching and our approach to indirect inference is that we

[3] Our results are robust to T = 100. See Appendix B.3 for a detailed discussion.

[4] The main reason why we estimate a rather limited subset of the parameters is to reduce the computational burden. Our LP approach can be very easily adapted to estimate a large number of parameters as it is very easy to generate moments, e.g. by matching the response to a different variable, shock or over a longer horizon.

are estimating the IRF using the model's simulated data. In other words, in the canonical IRF matching approach there is a model-implied mapping from the set of parameters to the IRFs; while in the indirect inference approach this mapping also comes from the auxiliary econometric model used to estimate the IRFs.[5]

We construct the data moments as follows. For each of the 100 simulated datasets at $\Theta = \Theta^*$, which are of length $T = 300$, we estimate either the LP or the SVAR-IRFs to technology, fiscal, or monetary shocks as described in Section 3. We set $p = 4$ and $H = 20$. Figure 1 depicts the distribution of these LP-IRFs (left panel) and SVAR-IRFs (right panel) to a technology shock. 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.

Notice from Figure 1 that the variability of the estimated LP-IRFs is much larger than that of the SVAR-IRFs, specially at horizons $h > p$. As a result, one may expect that the variance of the estimated economic parameters is also going to be larger when using the LP approach to indirect inference. And indeed, this is true for the majority of the estimated parameters as we will discuss in Section 4.3. This result, i.e. the wider distribution around the LP-IRFs, also holds for the fiscal and monetary shocks.[6]

Yet, from these figures it is also clear 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.

4.3. Results

The structural model is estimated $S = 100$ times, that is once for each of the 100 vectors of data moments corresponding to one of our eight scenarios.[7] In particular, six of these eight scenarios corresponds to estimation routines that match LP or SVAR responses to only one of the three shocks. While the remaining two correspond to a selection of the LP or SVAR responses to all of the three shocks. We base this selection on the sensibility analysis of the

[5] Note that these two approaches are only equivalent when the econometric model used to estimate the IRF is not misspecified, and therefore the estimated IRF coincides with the model-implied one. In our specific application, model misspecification will depend upon our choices of T^s and p.

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

[7] To be clear, the same auxiliary model is used for each of these 100 vector of moments. The only source of variability comes from the different draws of the shocks.

Figure 1: Technology Shock

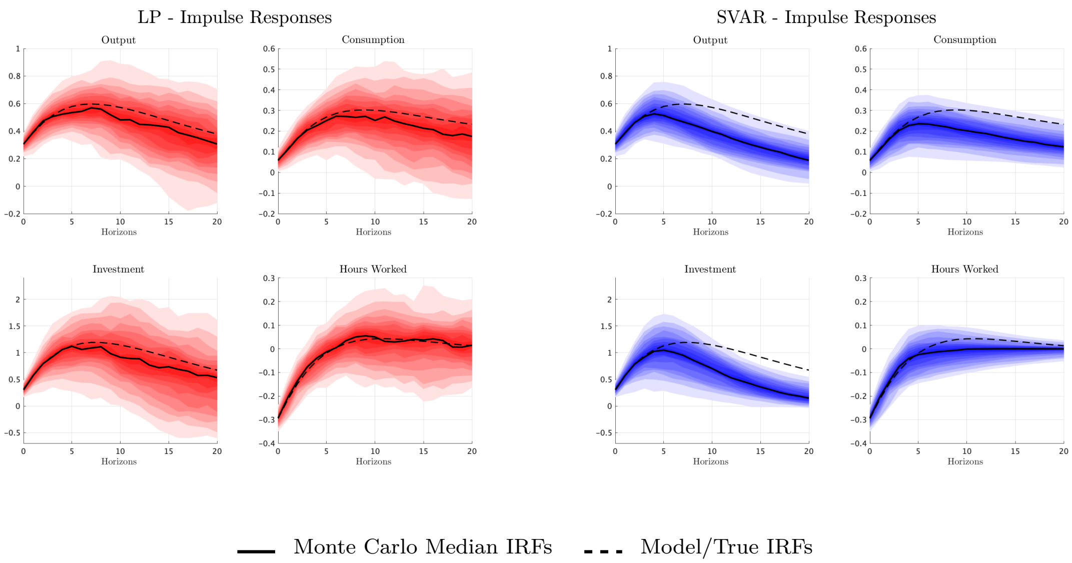

moments and try to match the response of investment to the three aggregate shocks as well as the response of consumption to a technology shock.[8]

To construct the counterpart of the data moments, we simulate the model at $\Theta = \Theta^{guess}$ for $T^s = 3,000$ periods, so that the simulated database is 10 times larger that the "observed" one (Smith, 1993). To minimize the effects of initial conditions, we simulate $0.1 \times T^s$ periods which are then discarded.[9] Then, we estimate (3) and (5) on this simulated dataset using the same auxiliary model as for the data moments.

Thus, in each scenario we target a total of 84 ($= 21 \times 4$) moments to estimate the 8 parameters in Table 1. As we have an over-identified model, we use the optimal weighting matrix, that is the inverse of the variance covariance matrix of the moments, to solve the minimization problem stated in equation (6).[10]

The solution to that problem is a vector of estimated parameters for each of the 100 vectors of data moments, i.e $\hat{\Theta} = \{\hat{\Theta}^s\}_{s=1}^{S}$. Following Smith (1993), we compute for each of the estimated parameters $\hat{\Theta}_i \in \hat{\Theta}$, the following statistics:

[8] This choice is based on the responsiveness of the moments to a one percent increase of each parameter at a time. See Appendixes B.5.2 and B.5.3 for further details.

[9] This is also done for the "observed" sample size.

[10] We compute the variance covariance matrix of the moments by computing the moments at the true parameter vector over 250 different samples. Variability comes from the different draws used for the innovation to the shocks.

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

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

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

where expectations are taken over the S Monte Carlo draws. Tables 7 and 8 in Appendixes B.1 and B.2 report these statistics for all the considered scenarios. These two tables contain a lot of information, thus, we will mainly focus on the Root Mean Squared Error (RMSE) going forward since it combines the information of both bias and variance. Therefore, it seems natural to start the discussion about the relative performance of the two moment generation functions as well as the performance across shocks by looking at this metric parameter by parameter.

Figure 2 tries to summarize this comparison. Panel (a) assesses how well responses to different shocks pick up different parameters when the LP IRFs have been used as the source of moments. For example, focus on the comparison between the estimates obtained after matching the responses to a technology shock versus those coming from targeting the responses to a monetary shock (subplot in the first column, second row). We can see that half of the parameters are better estimated through the responses to a monetary shock. That is the ratio of RMSEs is greater than 1. There is though a very interesting exception: the Calvo price adjustment probability $(1 - \xi_p)$, which has a RMSE ratio below 1. It seems surprising that this parameter is better picked up by a technology shock given its importance for inflation dynamics and therefore for monetary policy. In any case, this can be seen as an opportunity, since one can exploit the variability coming from a non-monetary shock to estimate some of the key parameters for monetary policy, e.g. $1 - \xi_p$, and then evaluate the validity of the model "out of sample" with respect to an untargeted response to a monetary shock.

Another interesting observation from panel (a) is that there are certain parameters such as the habit formation parameter h or the Calvo wage adjustment probability $(1 - \xi_w)$ that are much better pinned down by monetary shocks alone (see subplot in the second column, third row). This points out to the fact that certain parameters are better identified through various responses to the same shock rather than by different shocks. On the other hand, there are other parameters such as the intratemporal elasticity of labor supply, σ_l, that are better

Figure 2: Relative performance in terms of RMSE

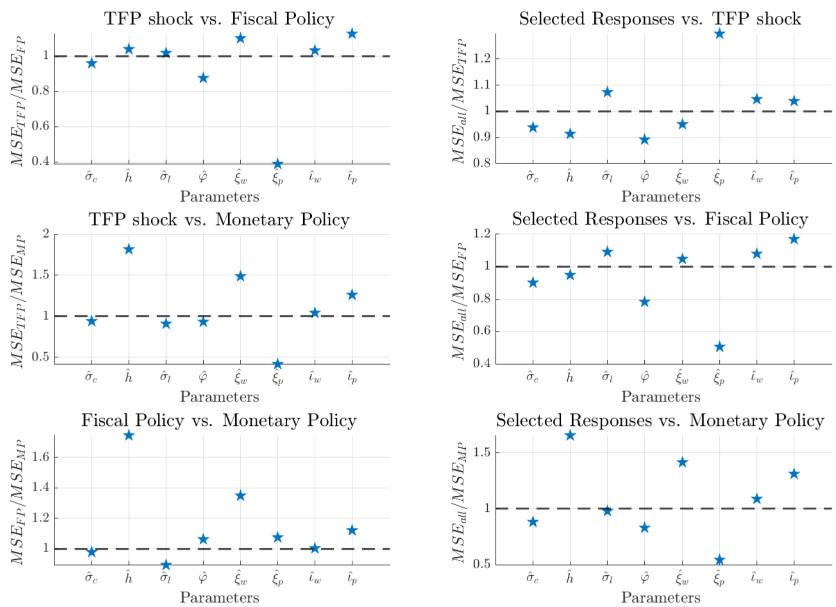

(a) Across Different Shocks (LP - IRFs approach)

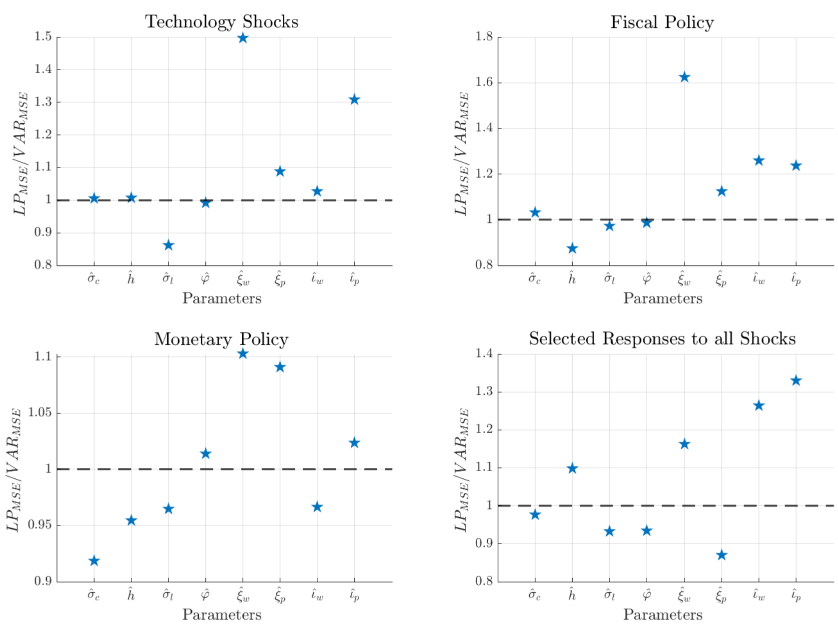

(b) LP - IRFs vs. SVAR - IRFs

identified by the variation coming from the responses to different shocks as shown in the first two subplots of the second column.

Panel (b) analyses how the two moment generating functions, LPs and SVARs, pick up the different parameters when matching the same responses to the various shocks. Here the focus is on the comparison across moment generating functions. A value greater than 1 in each of these 4 subplots indicates that a specific parameter is better identified using the SVAR as auxiliary model. Overall, this is true for approximately 80% of the cases. Thus, the SVAR - IRF approach to indirect inference seems to be superior, at least in terms of the RMSE. This result, although surprising, could be explained by the well-known fact that model misspecification is not an issue for indirect inference. Therefore, the relatively less variance of the SVAR - IRFs across different shock draws leads also to better estimates in terms of RMSE.[11]

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. To do so we evaluate the value of the criterion function, i.e. the RHS of equation (6), at the optimal/estimated parameter values $\hat{\Theta}^s$. This metric, usually refer to as the J-statistic, gives an overall measure of how well we match our targeted moments. Thus, the closer J is to zero, the better the match. However, also notice that even at Θ^* these metric won't be zero because the simulated moments are obtained in a different sample ($T \neq T^s$) with different draws for the shock.

Table 2 reports the average and the maximum value of this statistic across the S Monte Carlo draws for each of the eight different scenarios, i.e. four different sets of targeted IRFs estimated under the two alternative moment generating functions. We start by fixing the moment generating function and comparing across targeted IRFs. Notice how the tailored scenario, for which we picked the responses that were more sensitive to small changes in the parameters, gives the lowest average J. In any case, values are similar to each other and there is still some room for improvement as the J-statistic is far away from zero.

[11] How much of this result is explained by the bias and how much by the variance is still an open question since RMSE gives them equal weights.

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

If we now fix the targeted responses and compare across moment generating functions, then, we see how the LP - IRFs approach does worse on average, which points into the same direction as the evidence found for the RMSE. However, the maximum value of J is always larger, and by a significant margin, for the SVAR - IRFs scenario. Moreover, it is also important to notice the time needed to estimate the parameters over the S repeated samples is much lower for the LP - IRFs, something that it's worth taking into account for models in which the solution is time-intensive.[12] Thus, these last two observations favor our proposed approach of using the LPs as the auxiliary model in an indirect inference application.

Finally, we consider a third measure to evaluate the performance of the LP and SVAR - IRFs approaches to indirect inference since neither RMSE or J-statistics informs 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: 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

[12] Notice that the lower computing time needed to estimate the parameters comes from the lower number of iterations needed to reach the minimum. The number of iterations is significantly smaller since per iteration computing the LP-IRFs takes longer since one has to estimate more coefficients.

Table 3 shows the average and the maximum value of this metric across all the $\hat{\Theta}^s$. Results are striking. According to this metric, the LP - IRFs approach is able to match the theoretical IRFs much better than its SVAR counterpart. It may seem inconsistent with our previous findings; however, this measure is simply telling us that the LP - IRFs approach does a significantly a better job in picking those parameters that are relevant for capturing the shape of the true impulse responses. Consequently, it may still be desirable to use the LP - IRFs approach to indirect inference despite its underperformance in terms of the J-statistic and the RMSE for most parameters.

5. An Empirical Application

This last section illustrates the local projection approach to indirect inference by estimating a subset of the parameters of the Smets and Wouters (2007) model using empirically estimated esponses to technology, fiscal and monetary shocks as our data moments. These responses are estimated within the Jordà (2005) local projection framework.

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. A review of the literature on this topic can be found in Ramey (2016).

We borrow from Ramey (2016) the LP estimated responses to an unanticipated TFP shock, measured as in Francis et al. (2014).[13] Ramey (2016) reports the responses of various variables among which we select the response of real GDP, consumption, non-residential investment, and hours. Figure 3 depicts the median responses of these four variables (solid blue line) during 21 periods. Thus, we target a total of 84 (21×4) data moments during the estimation. These moments come from the following LP regression:

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

where $\tilde{z}_{t+h}$ is the variable of interest, $shock_t$ is the innovation to the growth rate of TFP and

[13] They identify the shock through medium-run restrictions.

ε_{t+h} the error term. Simulated moments are computed using equation (3), that is regressing the variable of interest on the TFP innovation over H horizons while controlling for p lags of itself and the shock. In particular, we estimate this LP regression on model simulated data of length $T^s = 300$.

Notice how equation (10) looks very similar to (3) with the important nuance that the shock may have some associated measurement error which is potentially correlated with the error term.[14] Thus, we argue that since the true model shock is *iid*, then there is no need to add extra controls to recover the correct IRFs. This simplifies things, since many times we don't have access to the model counterpart of those control variables. For example, Ramey (2016) includes the log real stock prices per capita as a control, a variable that is not available within the Smets and Wouters (2007) model. We do, however, control for the same number of lags of the shock and the dependent variable.

We estimate the same eight parameters from the Monte Carlo study. That is, we look over the following parameter space $\Theta = \{\sigma_c, h, \sigma_l, \varphi, \xi_p, \xi_w, \iota_p, \iota_w\}$ to minimize the quadratic distance between data and simulated moments. For this exercise, the weighting matrix is the inverse of a diagonal matrix comprised of the standard errors of the LP coefficients from the data analysis.

These are all the ingredients necessary to solve (6) and estimate the parameters of the model using our LP - IRFs approach. We repeat this estimation 100 times to get a distribution around our parameter estimates, i.e. confidence intervals are obtained by bootstrapping.[15]

The median parameter estimates are reported in the second subtable of Table 4, along with 10th and 90th percentiles. The estimates are reasonably close to the mean estimates in Smets and Wouters (2007), reproduced again in the first subtable of Table 4 for convinence. 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.

[14] Stock and Watson (2018) explain very neatly all these issues in empirical applications and how to tackle them by means of external instruments. For our particular application we do not rely on external instruments and keep the specification in Ramey (2016) as it is.

[15] Unlike in the case of the Monte Carlo analysis, the data moments are now fixed and the distribution comes only from the different draws of the shocks.

Table 4: SMM estimates using results 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	1.23	0.85	1.51	5.14	0.52	0.59	0.41	0.15
10th pctl.	0.90	0.55	1.51	3.79	0.42	0.40	0.35	0.14
90th pctl.	1.57	0.97	2.18	7.89	0.85	0.82	0.72	0.30

We also compute the estimated responses on simulated data at our median estimates. In particular, we solve and simulate the Smets and Wouters (2007) model using our median parameter estimates. Then we use this artificial dataset to compute the LP - IRFs. Since these are noisy due to simulation error, we repeat this procedure 1,000 times and take the mean of the simulated moments.[16] The resulting impulse responses are depicted in red in Figure 3.

Overall, the Smets and Wouters (2007) model does a good job in capturing these responses at the estimated parameters by our LP - IRFs approach. The match is particularly close for output, investment and hours. However, the model is not able to produce the delayed peak in consumption found in the data moments.

How much these results are driven by our parameter estimates and how much can be attributed to the specification of the model itself? To answer this question we report the estimated impulse responses, computed in identical fashion, at the Smets and Wouters (2007) mean estimates, the black dotted line in Figure 3.

The responses of output, investment and hours worked are very similar at both parameter estimates. Moreover, the model is not able to produce the delayed peak in consumption and in addition generates a smaller response in the short run and a smoother hump when using Smets and Wouters (2007) mean estimates. It is important to note that these estimates were obtained on a different sample.

[16] Notice that such mean estimates are equivalent to the model-implied IRFs because we include the true innovation of the shock in the regression and we have a sufficiently large sample (T = 300, S = 1,000). See results in Plagborg-Møller and Wolf (2021) for more clarity.

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

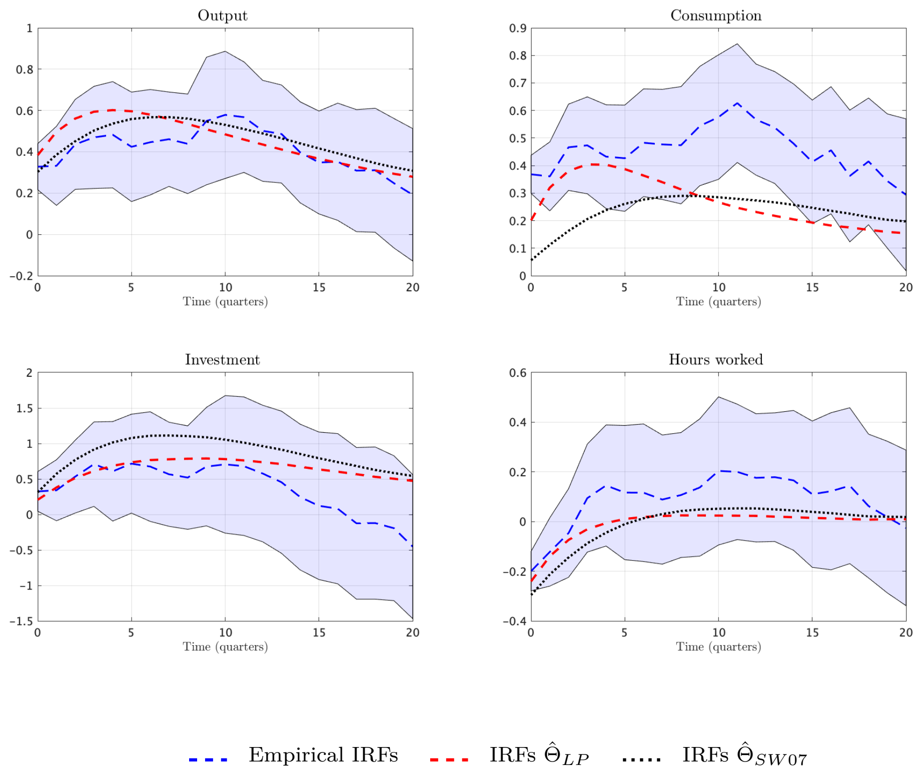

5.2. Government Spending Shocks

Now we turn to fiscal policy shocks, for which we also borrow the responses from Ramey (2016). In particular, we target the responses of GDP, non-durables and services consumption, non-residential investment and hours worked to an unanticipated government spending shock. As for the technology shock, Ramey (2016) reports these responses to shocks identified in various forms. We focus on those that were generated using the Blanchard and Perotti (2002) approach. That is, the shock is identified by assuming that government purchases were pre-determined within the quarter. In other words, the shock comes from a standard Cholesky decomposition with government spending ordered first. These responses are depicted in blue in Figure 4.

As with the technology shock, we use these moments to estimate the same eight parameters from the Monte Carlo study. That is, we look over Θ to minimize the quadratic distance between data and simulated moments. Simulated moments are computed in a similar fashion but now using the innovation of the government spending shock (instead of that of the TFP shock) in equation (3).

The distribution of the parameters estimates that we obtain using our LP - IRFs indirect inference approach are summarized in the third subtable in 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.

In Figure 4, 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 Smets and Wouters (2007) model doesn't match the empirical responses well besides that of output. As for hours worked, the model has a bigger response than in the data in the initial 10 quarters. The model also generates a much pronounced fall in investment over the entire horizon. Moreover, it also fails to generate a positive response of consumption.

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

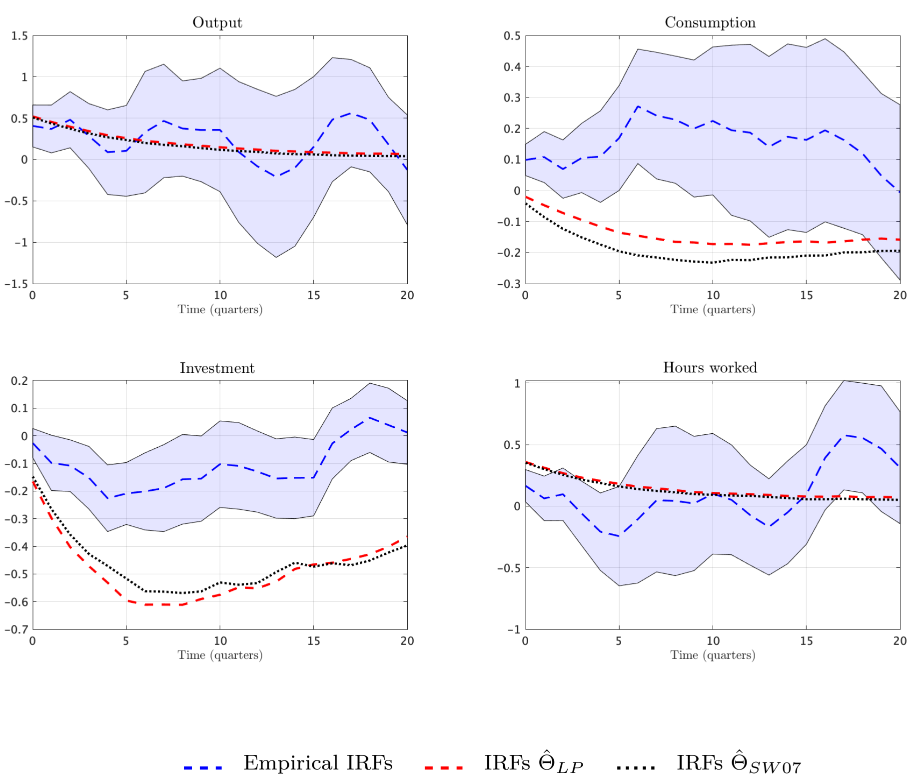

Overall, the model fails to match the responses to a government shock. We believe that this poor performance of the model could come from the lack of distortionary taxation and it is not related to the estimation strategy. At the end, the Smets and Wouters (2007) model was designed to inform policymakers about monetary policy at a time when its interactions with fiscal policy were out of the scope of Central Banks.

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). 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.[17] Their local projection estimation is based upon:

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

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.

Figure 5 illustrates the state dependent estimated effects of monetary policy shocks on output, consumption and investment for the three econometric models. The blue-dashed line in the figure are the mean estimates reported in Tenreyro and Thwaites (2016). 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 15 quarters or so.

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

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

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

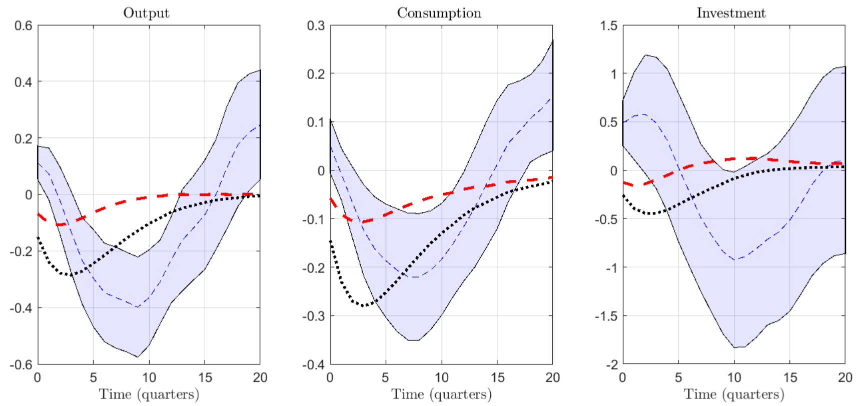

(a) Linear Model

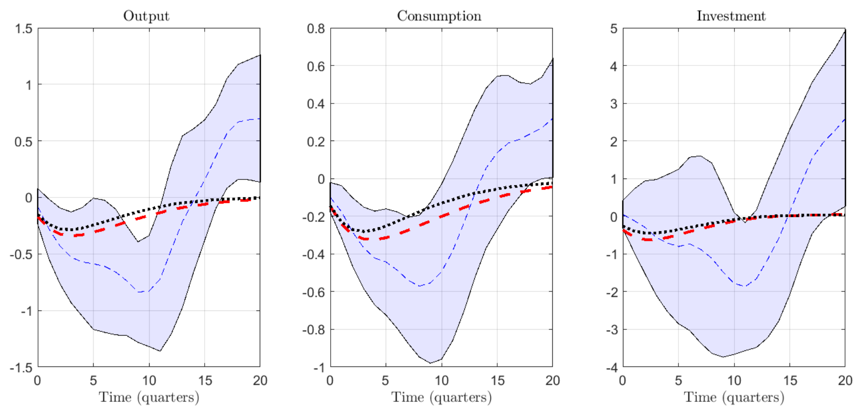

(b) Non-Linear Model: Expansion

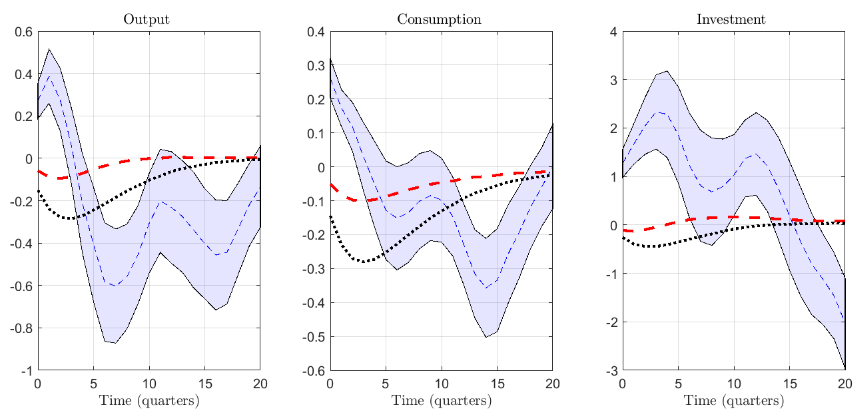

(c) Non-Linear Model: Recession

and investment combined with a slow return. From the third row, the effects of a monetary contraction during a recession are quite different.[18]

The estimated local projection parameters that identify these impulse responses are used as moments in our estimation. Of course, the underlying model of Smets and Wouters (2007) does not contain the nonlinearities that might give rise to the state dependence reported by Tenreyro and Thwaites (2016). Nonetheless it is instructive to see how well we can match the moments from the state dependent estimation as well as the linear model.

Table 5 reports our parameter estimates. The impulse response functions at our median estimates are shown (in red) in Figure 5. The estimation is undertaken for the three sets of impulse responses reported in Tenreyro and Thwaites (2016). 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: SMM estimates using results 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	1.26	0.91	3.15	7.89	0.46	0.32	0.32	0.10
10th pctl.	0.85	0.80	1.51	5.94	0.46	0.32	0.32	0.10
90th pctl.	1.57	0.98	3.15	7.89	0.76	0.66	0.66	0.21
	Non-Linear Model: Expansion							
Median	1.57	0.76	1.51	4.06	0.72	0.66	0.32	0.10
10th pctl.	0.76	0.64	1.51	3.79	0.46	0.32	0.32	0.10
90th pctl.	1.57	0.94	3.15	7.89	0.90	0.66	0.66	0.21
	Non-Linear Model: Recession							
Median	1.57	0.91	3.15	7.89	0.46	0.32	0.66	0.21
10th pctl.	0.90	0.79	1.51	4.83	0.46	0.32	0.32	0.10
90th pctl.	1.57	0.98	3.15	7.89	0.77	0.56	0.66	0.21

Looking at the impulse responses, the estimated model can match the basic pattern of the responses to monetary contractions during an expansion. Output, consumption and investment

[18] One natural interpretation is that the contractionary policy during a recession is occurring during a period of stagflation, with monetary policy focusing on the combating inflation.

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.

Finally, we compare our estimated model with that of Smets and Wouters (2007). The first row of Table 5 reports the parameter estimates from that paper. Looking first at the linear model, our median parameter estimates are generally close to theirs and the reported confidence intervals contain their parameter estimates. Our estimated elasticity of substitution matches theirs exactly. 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 expansions, our estimates of wage and price flexibility match theirs, though in this specification the intertemporal elasticity of substitution is a bit higher than their estimates. In contrast, the estimated habit is weaker as is the elasticity of labor supply.

Looking at the impulse responses in Figure 5, the ones generated at our parameter estimates and those at Smets and Wouters (2007)'s 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_c was almost 50% larger when matching the response to fiscal policy compared to technology shocks. 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.[19]

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 1.00 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 monetary shocks. It is in this sense the joint estimation exercise combines the responses to the individual shocks.

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	1.00	0.62	3.15	5.89	0.42	0.40	0.43	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	0.86	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	1.23	0.85	1.51	5.14	0.52	0.59	0.41	0.15
Monetary Policy	1.26	0.91	3.15	7.89	0.46	0.32	0.32	0.10

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.

[19] 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.

It is important to note, however, that all these parameters have been estimated using information generated 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 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).

Figures 6 shows the response of investment to the three types of shocks and Figure 7 shows the untargeted output response. As in previous figures, the black dotted line is from Smets and Wouters (2007), the red dashed line is the model response at the median estimates and the blue dashed line represents the data. As noted earlier, the estimated model does well matching the responses to the technology and fiscal shocks but not so well to the monetary innovations. That finding remains in the case of the joint estimation.

Figure 6: Targeted Investment Response to All Three Shocks

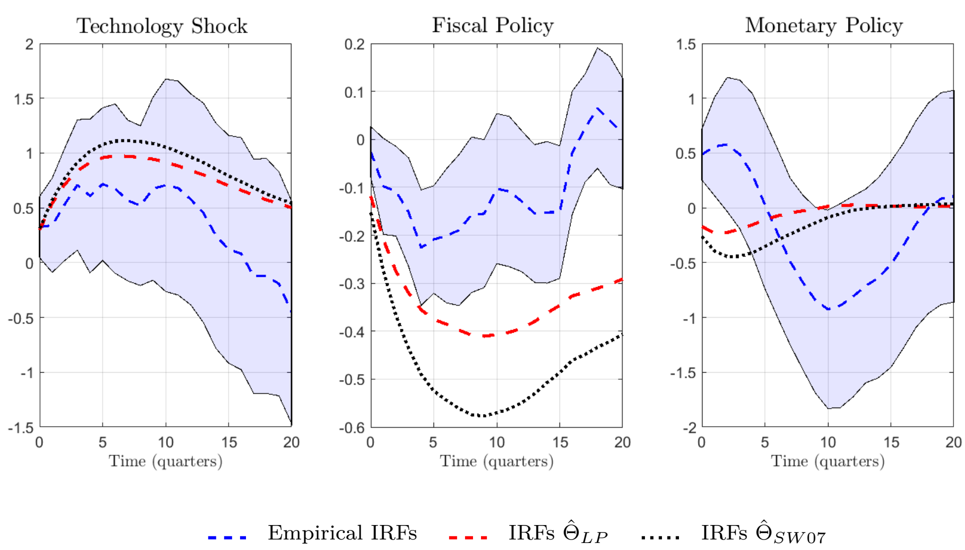

This same point holds for the output response even though we do not target it for the estimation. Looking at the output response, the Smets and Wouters (2007) parameterization again produces a deeper and longer reduction in output compared to these estimates based upon the shocks together. A key difference is in the frequency of wage and price adjustment.

Figure 7: Untargeted Output Response to All Three Shocks

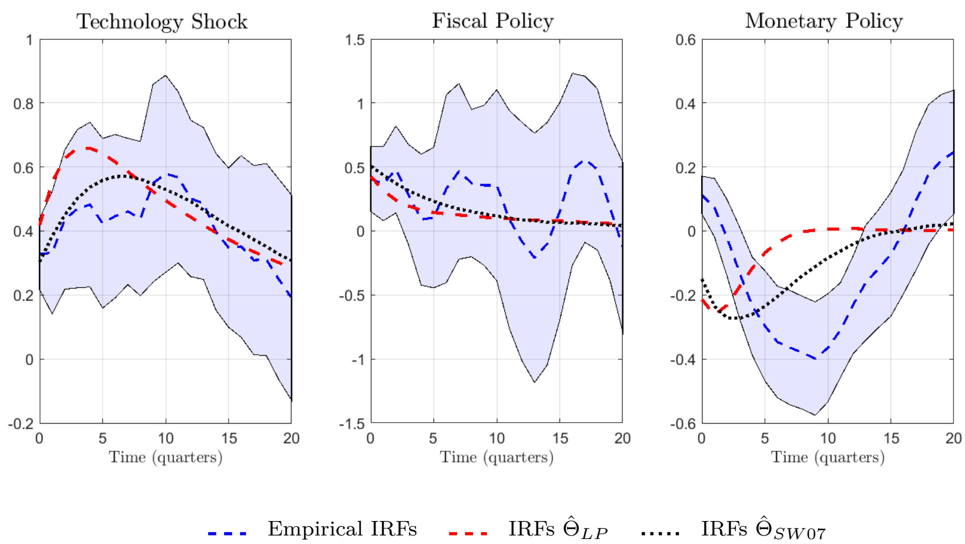

To dig a bit deeper into differences in estimates compared to Smets and Wouters (2007), Figure 8 shows the impulse responses to the three shocks (columns) for alternative parametrizations (rows). We construct each of these four alternatives by changing one at a time the curvature of utility, the habit 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

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

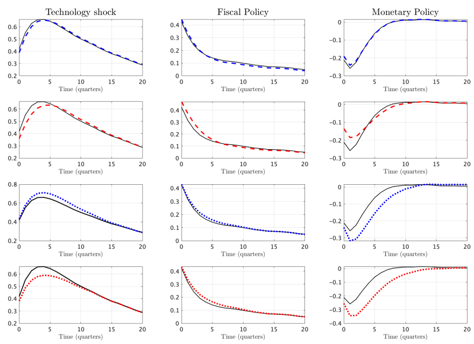

(a) Output

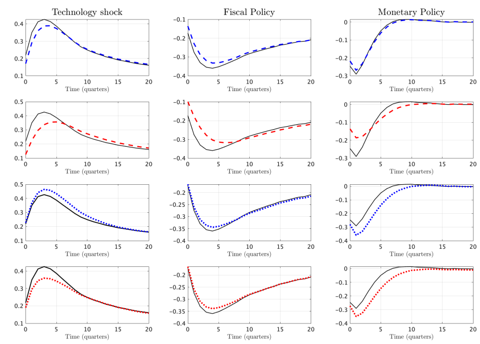

(b) Consumption

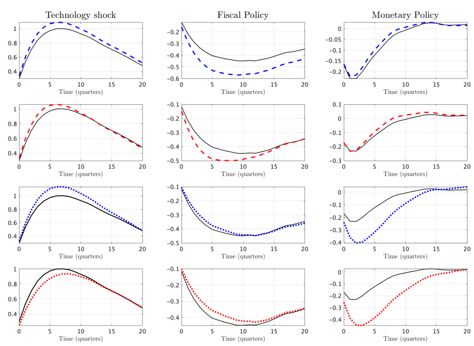

(c) Investment

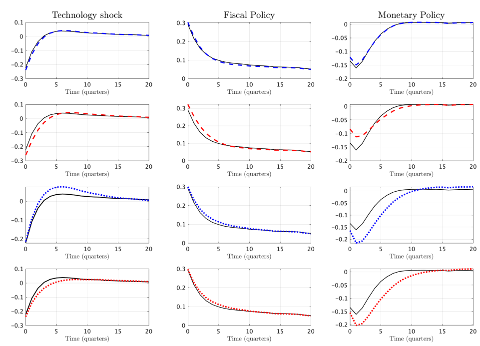

(d) Hours Worked

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

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.

The differences in preference parameters, both the elasticity of substitution and the habit, matter most for the response of consumption to all three shocks. 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 an indirect inference approach to structural estimation using the coefficients of a LP regression as moments. The LP method of estimating impulse responses is now commonly used. The contribution here is to use it to study how informative the moments are for the estimation of the underlying preferences and technology parameters as well as the degree of wage and price stickiness.

Two exercises are performed. The first is a Monte Carlo exercise, comparing the estimation using LP moments from those based upon a VAR, as in Smith (1993). Here we find that 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. This implies that the LP - IRF approach does a better job in picking up those parameters that matter the most for the shape of the IRFs, which are the ultimate object of interest.

The second estimates the model parameters using LP moments generated from other studies. Here we first looked individually at technology, fiscal policy and monetary policy induced fluctuations, comparing our parameter estimates to those reported in Smets and Wouters (2007). Then, we estimated the model parameters jointly using the LP responses to the various shocks. In particular, we focused on the investment response as it is the most sensitive to changes in parameters, but we also looked at the untargeted responses of other variables. Here we find that the model was able to match fairly well the responses to technology, fiscal shocks as well as to contractionary monetary policy shocks during expansions. However, it is not suited to match the responses of contractionary monetary shocks during recessions. In regards to point estimates, our results imply a much lower degree of wage and price stickiness than those obtained by Smets and Wouters (2007). This is true regardless of the targeted sources of variation.

The analysis uses Smets and Wouters (2007) as a baseline model. This is a natural starting point given the prominence of that model in the literature and its use by Plagborg-Møller and Wolf (2021). Nonetheless, the model is unable to produce the types of non-linear responses to monetary innovations highlighted by Tenreyro and Thwaites (2016). Clearly consideration of a non-linear model using the LP approach is a natural next step.

References

Blanchard, O. and Perotti, R. (2002). An empirical characterization of the dynamic effects of changes in government spending and taxes on output. *Quarterly Journal of Economics*, 117(4):1329–1368.

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

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

Christiano, L. J., Eichenbaum, M. S., and Trabandt, M. (2015). Understanding the Great Recession. *American Economic Journal: Macroeconomics*, 7(1):110–167.

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. *Review of Economics and Statistics*, 96(4):638–647.

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

Jordà, Ò. (2005). Estimation and inference of impulse responses by local projections. *American Economic Review*, 95(1):161–182.

Li, D., Plagborg-Møller, M., and Wolf, C. K. (2021). Local Projections vs. VARs: Lessons From Thousands of DGs. *Working Paper*.

Plagborg-Møller, M. and Wolf, C. K. (2021). Local Projections and VARs Estimate the Same Impulse Responses. *Econometrica*, 89(2):955–980.

Ramey, V. A. (2016). Macroeconomic Shocks and Their Propagation. *Handbook of Macroeconomics*, 2:71–162.

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.

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(Special Issue on Econometric Inference Using Simulation Techninques):S63–S84.

Stock, J. H. and Watson, M. W. (2018). Identification and Estimation of Dynamic Causal Effects in Macroeconomics Using External Instruments. *Economic Journal*, 128(610):917–948.

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.

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

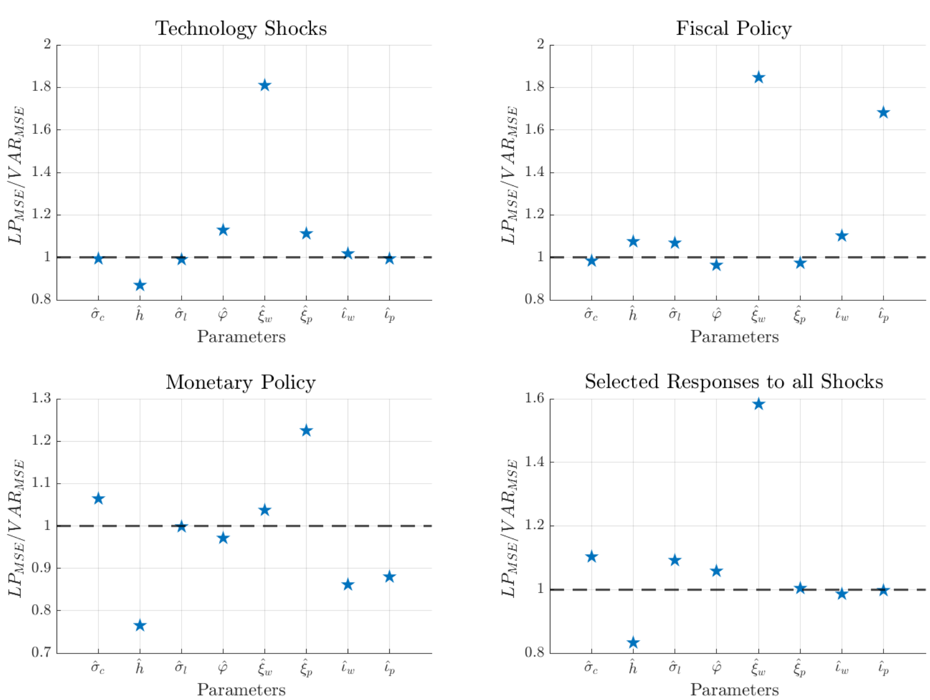

Table 9: Overall performance: Estimated Impulse Responses

	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.[20] 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

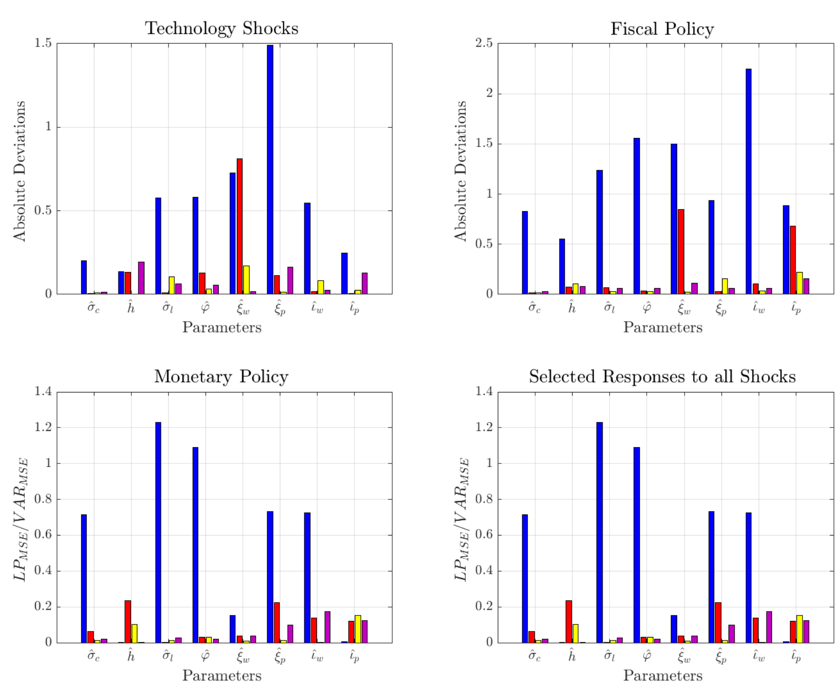

[20] Recall that in the Smets and Wouters (2007) model one period corresponds to one quarter.

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

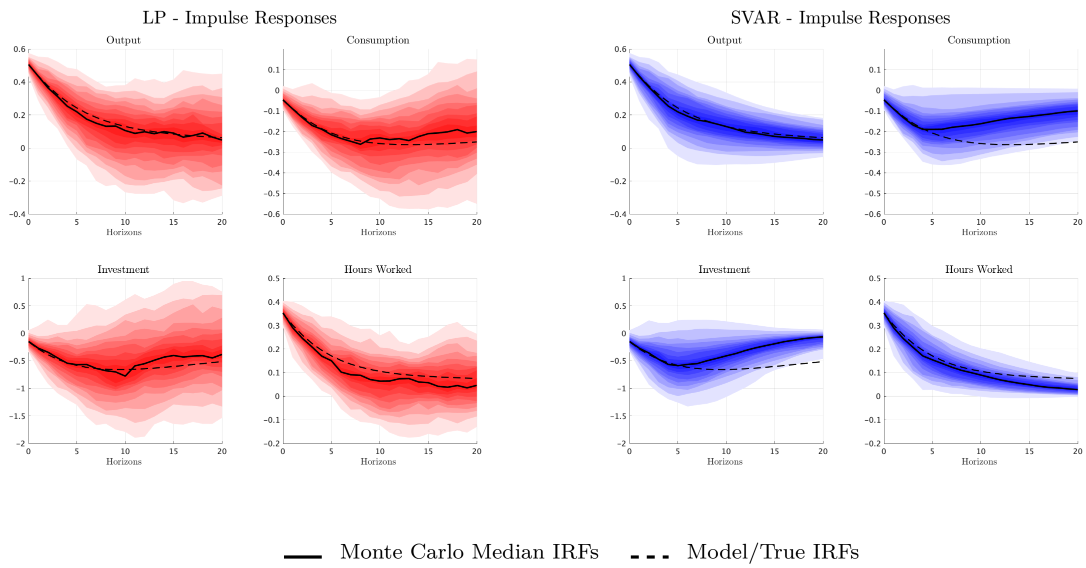

Figure B.5.2: Monetary Policy Shock

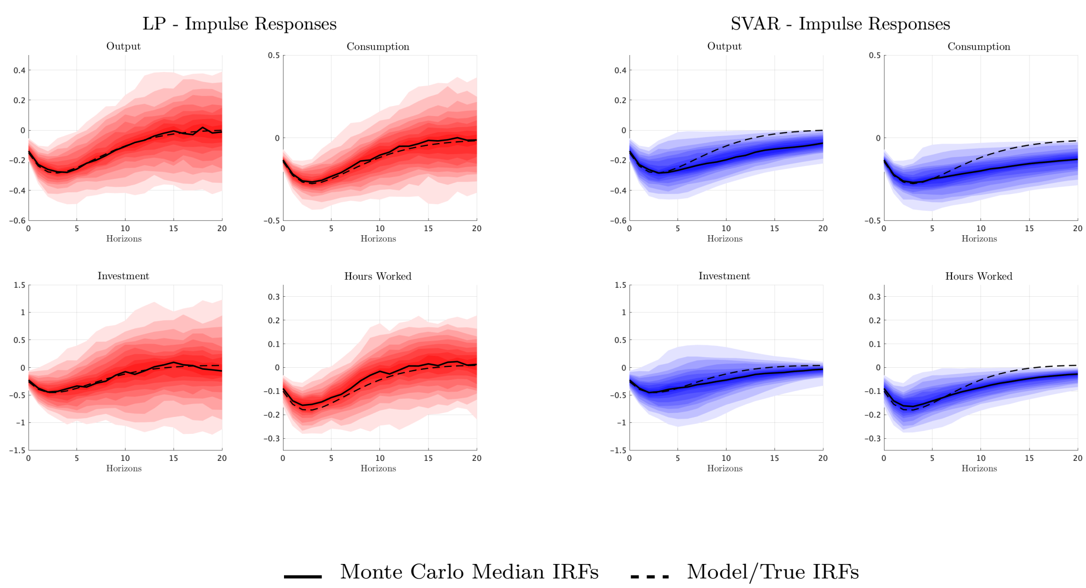

B.5.2. Sensitivity of the Moments to Changes in the Parameters

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

Figure B.5.1: LP-IRFs sensitivity

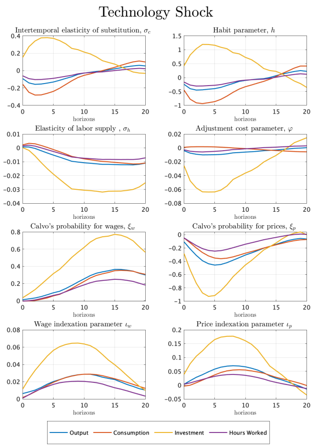

(a) Technology shock

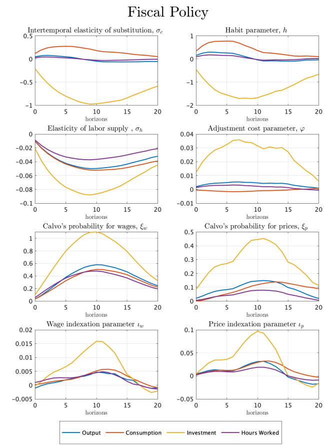

(b) Fiscal Policy

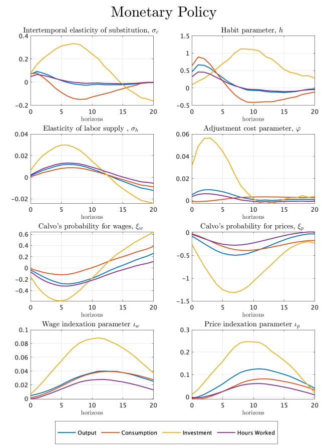

(c) Monetary Policy

Figure B.5.2: SVAR-IRFs sensitivity

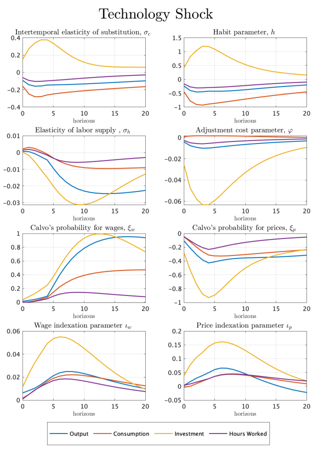

(a) Technology shock

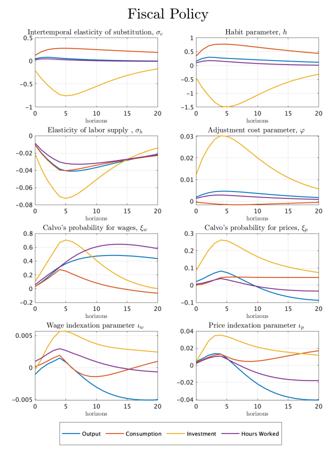

(b) Fiscal Policy

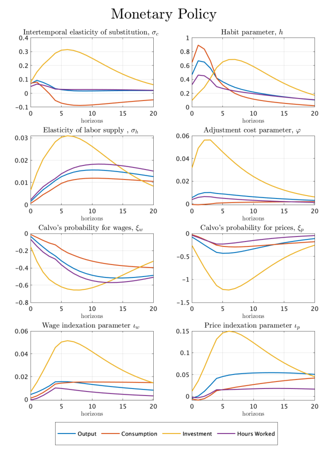

(c) Monetary Policy

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

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

Figure B.5.3: LP-IRFs elasticities

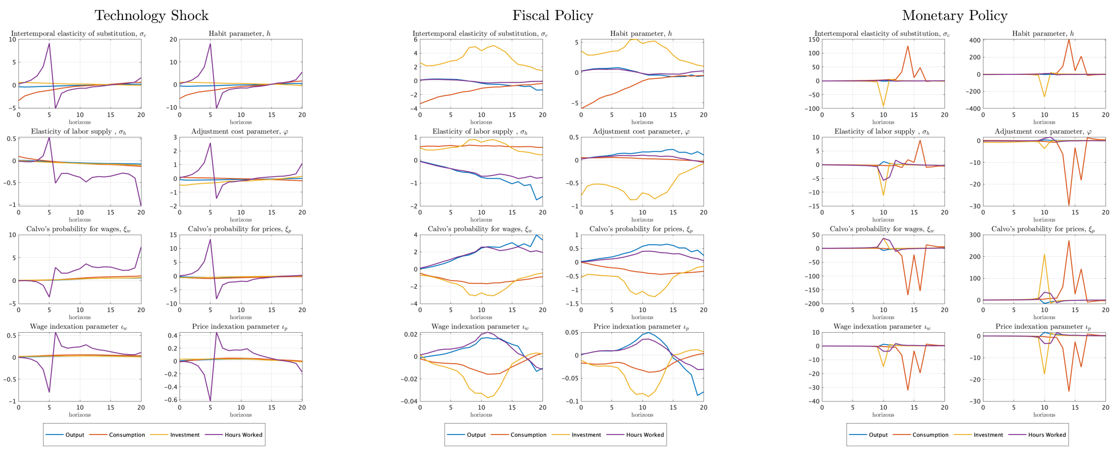

(a) Technology shock (b) Fiscal Policy (c) Monetary Policy

Figure B.5.4: SVAR-IRFs elasticities

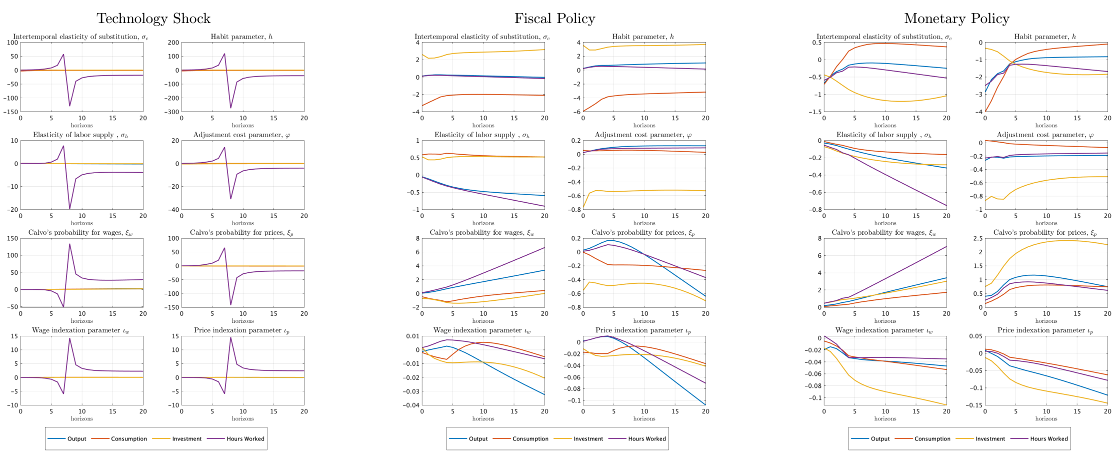

(a) Technology shock (b) Fiscal Policy (c) Monetary Policy