

Indirect Inference: A Local Projection Approach ^{*}

Juan Castellanos [†]

Russell W. Cooper [‡]

This draft: May 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 ask: *how should one pick between VARs and LPs when choosing an auxiliary model for indirect inference?* To answer this question, we make use of 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, for which we borrow from the literature the empirically estimated responses of key macro variables to technology, fiscal and monetary shocks.

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

JEL classification: C13, C15, E00

^{*}We thank Jesus Bueren, Andrea Gazzani, Lukas Nord, Sebastian Rast, and Barbara Rossi for useful comments and suggestions. Any views expressed herein are those of the authors and do not necessarily coincide with those of the ECB.

[†]European University Institute and European Central Bank, 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 auxiliary 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) estimated values that Plagborg-Møller and Wolf (2021) use for their graphical illustration (section 2.4. of their paper).¹ 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 preliminary results show that there are some discrepancies, e.g. we obtain a lower intertemporal elasticity of substitution and higher frequencies of wage and price adjustment when we match the responses to technology and 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. During an expansion, the model is able to match these responses at our and their parameter estimates. Overall, our parameter estimates and those estimated by Smets and Wouters (2007) are very close, which in turn result in very similar theoretical responses to the various shocks.

¹ These values correspond to those obtained using Johannes Pfeifer Dynare replication package: https://github.com/JohannesPfeifer/DSGE_mod/tree/master/Smets_Wouters_2007. Notice that these differ from the values reported in Smets and Wouters (2007) paper because a different sample is used.

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. It was estimated using Bayesian methods. 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

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

for $t = 1, 2, \dots, T$. 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:

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

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 $h = 0, 1, 2, \dots, H$. 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:

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

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*:

$$A(L)w_t = c + B\eta_t \quad (5)$$

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:

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

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 Òscar Jordà in the Monte-Carlo study of his seminal paper “Estimation and Inference of Impulse Responses by Local Projections” (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). Moreover, we will also show that estimated impulse responses have particularly large variability in small samples.³ Nevertheless, we haven’t yet explored how this choice influences our parameter estimates, but we conjecture that the LP approach to indirect inference may also be very sensible to changes in T .

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
1.2679	0.8056	2.5201	6.3144	0.7668	0.5304	0.5345	0.1779

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 his paper “Estimating Nonlinear Time-Series Models Using Simulated Vector Autoregressions”, in which 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.

³ Many applied macroeconomists are aware of the erratic behavior of the impulse responses estimated by LPs, which has even lead to Barnichon and Brownlees (2019) to propose a new estimation method that smooth out the response. We’ll show that in our set-up, in which we have access to an infinite amount of data, the variability disappears as one increases T .

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.


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, see for example [Christiano et al. \(2005\)](#). The subtle, but important difference, between IRF matching and our approach to indirect inference is that we are estimating the IRF using the model's simulated data either via LPs or via a VAR regression followed by a Cholesky decomposition; while in an IRF matching exercise the responses are directly recovered from the model. In other words, 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 true one.⁵

We construct 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. The counterparts of [Figure 1](#) for each of these two other shocks can be seen in [Appendix C.1](#).

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.

⁵ In our specific application, model misspecification will depend upon the choices of T^s and p as shown in [Appendix B](#).

Figure 1: Technology Shock



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.⁶ 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 moments (see figures in Appendix C.2) and try to match the response of investment to the three aggregate shocks as well as the response of consumption to a technology shock.

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 than the “observed” one as in Smith (1993). To minimize the effects of initial conditions, we simulate $0.1 \times T^s$ periods which are then discarded.⁷ 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 8 parameters. As we have an over-identified model we use the optimal weighting matrix, that is the inverse

⁶ 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.

⁷ This is also done for the “observed” sample size.

of the variance covariance matrix of the moments, to solve the minimization problem stated in equation (6).

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:

$$\text{Bias}_i \equiv \mathbb{E} [\hat{\Theta}_i] - \Theta_i^* \quad (7)$$


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

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


where expectations are taken over the S Monte Carlo draws. Tables 6 and 7 in Appendix D 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. Notice that the RMSE 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 most parameters are better estimated through the responses to a monetary shock. That is the ratio of RMSEs is greater than 1 in most cases. There are though a few exceptions: the elasticity of labor supply (σ_l) and the Calvo price adjustment probability ($1 - \xi_p$) have a ratio of RMSEs below 1. It seems surprising that the Calvo price adjustment probability is better picked up by a non-monetary 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.

Figure 2: Relative performance in terms of RMSE



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



(b) LP - IRFs vs. SVAR - IRFs

Panel (b) focuses on the comparison across moment generating functions. The message is a lot more clear in this case. The SVAR - IRF approach to indirect inference seems to be superior for most parameter values. This result 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.⁸

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: Overall performance: Estimated Impulse Responses

	Local Projections			Vector Autoregression		
	Avg. J	Max. J	Time (in min.)	Avg. J	Max. J	Time (in min.)
<i>Technology Shock</i>	86.93	117.61	17.38	83.71	245.40	69.90
<i>Fiscal Policy</i>	88.02	127.94	21.63	83.62	250.59	68.03
<i>Monetary Policy</i>	89.68	123.26	20.31	82.69	210.79	68.42
<i>Selected Responses</i>	86.28	123.03	20.11	81.97	237.03	62.49

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.

⁸ It is still needs to be studied how much of this result is explained by the bias and how much by the variance since RMSE gives them equal weights.

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 is consistent with the discussion about 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.⁹ Thus, these last two observations favor our proposed approach of using the LP - IRFs as moments.

Table 3: Overall performance: Model Impulse Responses

	Local Projections		Vector Autoregression	
	Avg. J^*	Max. J^*	Avg. J^*	Max. J^*
<i>Technology Shock</i>	2.90	7.84	226.54	1014.14
<i>Fiscal Policy</i>	3.36	14.94	289.26	1893.02
<i>Monetary Policy</i>	3.14	12.63	821.66	3380.27
<i>Selected Responses</i>	9.37	75.05	686.72	9307.71

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/theoretical 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 shows the average and the maximum value of this metric across all the $\hat{\Theta}^s$. Results differ substantially. According to this metric, the LP - IRFs approach is able to match the theoretical IRFs much better than its SVAR counterpart, which 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 it underperformance in terms of the J-statistic and the RMSE for most parameters.

⁹ Notice that the lower computing time need 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.

5. An Empirical Application

Finally, we illustrate our local projection approach to indirect inference by estimating a subset of the parameters of the [Smets and Wouters \(2007\)](#) model using empirically estimated responses 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. \(2014\)](#).¹⁰ [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:

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

where $\tilde{z}_{t+h}$ is the variable of interest, shock_t is the innovation to the growth rate of TFP and 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.¹¹ 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

¹⁰ They identify the shock through medium-run restrictions.

¹¹ [Stock and Watson \(2018\)](#) explain very neatly all these issues in empirical applications and how to tackle them by means of external instruments.

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 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, however, the data moments are now fixed and the distribution comes only from the different draws of the shocks.

Parameter estimates are reported in the first subtable of Table 4, along with confidence intervals. The estimates are relatively close to the estimates in [Smets and Wouters \(2007\)](#). In fact, their point estimates fall within our error bands. We estimate a lower curvature in utility, with a median estimate of $\hat{\sigma}_c = 0.85$; and higher Calvo adjustment probabilities, with median estimates of $1 - \hat{\xi}_w = 0.66$ and $1 - \hat{\xi}_p = 0.41$.


Table 4: SMM estimates using the LP - IRFs approach

	$\hat{\sigma}_c$	$\hat{h}$	$\hat{\sigma}_l$	$\hat{\varphi}$	$\hat{\xi}_w$	$\hat{\xi}_p$	$\hat{\iota}_w$	$\hat{\iota}_p$
<i>S&W 07</i>	1.26	0.80	2.52	6.31	0.70	0.66	0.58	0.24
Technology Shocks								
<i>Median</i>	0.85	0.69	3.28	8.20	0.44	0.59	0.47	0.14
<i>10th pctl.</i>	0.76	0.48	1.51	3.79	0.42	0.40	0.35	0.14
<i>90th pctl.</i>	1.36	0.89	3.28	8.20	0.84	0.86	0.75	0.31
Fiscal Policy								
	$\hat{\sigma}_c$	$\hat{h}$	$\hat{\sigma}_l$	$\hat{\varphi}$	$\hat{\xi}_w$	$\hat{\xi}_p$	$\hat{\iota}_w$	$\hat{\iota}_p$
<i>Median</i>	1.23	0.85	1.51	5.14	0.52	0.59	0.41	0.15
<i>10th pctl.</i>	0.90	0.55	1.51	3.79	0.42	0.40	0.35	0.14
<i>90th pctl.</i>	1.57	0.97	2.18	7.89	0.85	0.82	0.72	0.30

Here : *S&W 07* denote the mean estimates from [Smets and Wouters \(2007\)](#), see their Table 1A.

We also compute the estimated responses on simulated data at our median estimates. In particular, we compute the simulated moments over a sample of 300 periods where we use our mean estimates to generate such sample. Then, to clear up simulation error we repeat this procedure 1000 times and then take the mean of the simulated moments over all those draws. The resulting impulse responses are depicted in red in Figure 3.

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



Here : The blue dash line corresponds with the estimates in [Ramey \(2016\)](#) with confidence bands shaded in light blue. The red dash line corresponds to the simulated moments at the median parameter estimates. Finally, the black dotted line corresponds to the simulated moments at the [Smets and Wouters \(2007\)](#) mean parameter estimates. When computing the simulated moments we average out across 100 draws of the shocks.

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.

Finally, we also report the estimated impulse responses at the [Smets and Wouters \(2007\)](#) mean estimates, the black dotted line in Figure 3. It is important to note that these parameters

were obtained using Bayesian estimation techniques on a different sample. Nevertheless, we think the comparison is still interesting. Notice for example how the responses of output, investment and hours worked are very similar. At these parameter estimates the model is also not able to produce the delayed peak in consumption and in addition generates a smaller response in the short run paired with smoother hump.

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 + 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 [Blanchard and Perotti \(2002\)](#) strategy. 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 for 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 second subtable in [Table 4](#). In [Figure 4](#), we also plot (in red) the estimated impulse responses at the median parameter estimate. 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.

Finally, we compare our results to those obtained using Bayesian techniques by [Smets and Wouters \(2007\)](#). The black dotted line in [Figure 4](#) corresponds to the estimated responses

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



Here : The blue dash line corresponds with the estimates in [Ramey \(2016\)](#) with confidence bands shaded in light blue. The red dash line corresponds to the simulated moments at the median parameter estimates. Finally, the black dotted line corresponds to the simulated moments at the [Smets and Wouters \(2007\)](#) mean parameter estimates. When computing the simulated moments we average out across 1000 draws of the shocks.

at their mean estimated parameters. The responses of output and hours worked are almost identical, and those of consumption and investment are also closely related. Therefore, at [Smets and Wouters \(2007\)](#) parameters the model also 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.

5.3. Monetary Policy Shocks

The estimation of parameters based upon monetary policy shocks rests upon the local projection estimates in [Tenreiro 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.¹² Their local projection estimation is based upon:

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

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 and investment combined with a slow return. From the third row, the effects of a monetary contraction during a recession are quite different.¹³


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


¹²They derives economic expansions and contractions in terms of output growth, not levels.

¹³One 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.


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



(a) Linear Model



(b) Non-Linear Model: Expansion



(c) Non-Linear Model: Recession

Here : The blue dash line corresponds with the mean estimate in [Tenreiro and Thwaites \(2016\)](#) with confidence bands shaded in light blue. The red dash line corresponds to the median estimate at our estimated parameters with confidence intervals shaded in light red. Finally, the black dotted line corresponds to the model IRFs at the [Smets and Wouters \(2007\)](#) mean parameter estimates.

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.

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 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

Table 5: SMM estimates using the LP - IRFs approach

	$\hat{\sigma}_c$	$\hat{h}$	$\hat{\sigma}_l$	$\hat{\varphi}$	$\hat{\xi}_w$	$\hat{\xi}_p$	$\hat{l}_w$	$\hat{l}_p$
<i>SEW 2007</i>	1.26	0.80	2.52	6.31	0.70	0.66	0.58	0.24
Linear Model								
<i>Median</i>	1.26	0.91	3.15	7.89	0.46	0.32	0.32	0.10
<i>10th pctl.</i>	0.85	0.80	1.51	5.94	0.46	0.32	0.32	0.10
<i>90th pctl.</i>	1.57	0.98	3.15	7.89	0.76	0.66	0.66	0.21
Non-Linear Model: Expansion								
<i>Median</i>	1.57	0.76	1.51	4.06	0.72	0.66	0.32	0.10
<i>10th pctl.</i>	0.76	0.64	1.51	3.79	0.46	0.32	0.32	0.10
<i>90th pctl.</i>	1.57	0.94	3.15	7.89	0.90	0.66	0.66	0.21
Non-Linear Model: Recession								
<i>Median</i>	1.57	0.91	3.15	7.89	0.46	0.32	0.66	0.21
<i>10th pctl.</i>	0.90	0.79	1.51	4.83	0.46	0.32	0.32	0.10
<i>90th pctl.</i>	1.57	0.98	3.15	7.89	0.77	0.56	0.66	0.21

Here : Our set of data moments, that is the LP - IRFs, come from the different specifications of [Tenreyro and Thwaites \(2016\)](#).

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.

6. Conclusion

This paper studies the use of an indirect inference approach to structural estimation using the coefficients of a LP as moments. The LP method of estimating impulse responses is now commonly used. The contribution here is to use study how informative the moments are for estimation of underlying parameters of preferences, technology and the flexibility of wages and prices.

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 estimated model parameters using LP moments generated from other studies. Here we looked individually at technology, fiscal policy and monetary policy induced fluctuations, comparing our parameter estimates to those reported in [Smets and Wouters \(2007\)](#).

Further work will estimate the parameters from the LP jointly. One interesting dimension of the analysis is that though, for example, wage and price stickiness might matter the most for the analysis of monetary policy shocks, these parameters might be best identified through the effects of a technology shock.

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

- Barnichon, R. and Brownlees, C. (2019). Impulse response estimation by smooth local projections. *Review of Economics and Statistics*, 101(3):522–530.
- 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.
- 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 Techniques):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:

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

- The consumption Euler equation:

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

- The investment Euler equation:

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

- The arbitrage equation for the value of capital:

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

- The aggregate production function:

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

- Capital services:

$$\hat{k}_t^s = \hat{k}_{t-1} + \hat{z}_t \quad (\text{A.6})$$

- Capital utilization:

$$\hat{z}_t = \frac{1-\psi}{\psi} \hat{r}_t^k \quad (\text{A.7})$$

- The accumulation of installed capital:

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

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

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

- New Keynesian Phillips curve:

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

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

$$\hat{r}_t^k = \hat{l}_t + \hat{w}_t - \hat{k}_t^s \quad (\text{A.11})$$

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

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

- Wage adjustment:

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

- Monetary policy reaction function:

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

B. The Role of the Sample Size

We have already stressed the importance of the sample size. In this section, however, we want to be more specific about what we actually mean. Start from looking at panel (a) of Figure B.0.1, where we plot the estimated LP - IRFs (red line) and SVAR - IRFs (blue line) on a sample with $T = 300$ observations, as well as the true impulse response function coming from the model (black dash line).


First, notice that LP - IRFs and SVAR - IRFs approximately agree until horizon $p = 4$, and after that, their behavior diverges. This is consistent with the results in [Plagborg-Møller and Wolf \(2021\)](#) and [Li et al. \(2021\)](#). The result is explained by the different extrapolation at horizons $h > p$: a VAR with p lags extrapolates its long-run responses from the first p sample autocovariances, and therefore it yields lower variance but a high bias. On the other hand, LPs uses the sample covariances more flexibly by directly projecting at future horizons h on current covariates (see (3) for more clarity). The similar behavior of the two estimators up to horizon p , and their disagreement beyond it, is very nicely seen in the response of output. Notice how the LP - IRF follows more closely the shape of the true response (low bias), but with a lot more ups and downs around it (high variance). On the other hand, the SVAR - IRF follows the true response only up to horizon p and then misses consistently the dynamics (high bias, low variance).

Second, as we move towards panels (b) and (c), which depict the same estimated impulse responses but on databases with a much longer time dimension ($T = 3,000$ and $T = 30,000$, respectively), we see how the LP - IRF matches more and more closely the true IRF and the wiggles at further horizons disappear. On the other hand, sample size doesn't seem to matter for the SVAR - IRFs. The response is still similar to that obtained through LPs until horizon p but then it diverges. To get closer to the true IRF with the SVAR approach one most likely needs to increase p . Recall that we picked p in a heuristic way, i.e. we haven't rely on the AIC or any other statistical criteria to make that choice.

What are then the possible implications of these two results for estimation? First, one may think that using LP - IRFs as moments may result in estimates of the economic parameters that are closer to the truth but with larger confidence intervals than if one uses SVAR - IRFs as moments. And second, if one simulates its model long enough, the estimates obtained through

LP - IRFs will coincide with those obtained by a maximum likelihood or IRFs matching approach since the auxiliary model is well-specified (estimated IRF = true IRF); while for the SVAR - IRFs approach to indirect inference it will require the researcher to fine-tune the optimal lag length and/or the inclusion of more variables in the system to avoid model misspecification.


Figure B.0.1: The Role of T for the Estimated Impulse Responses to a Technology Shock



C. Additional Figures

C.1. Fan Charts: Fiscal and Monetary Policy Shocks


Figure C.1.1: Fiscal Policy Shock



(a) Local Projections

(b) Vector Autoregression

Figure C.1.2: Monetary Policy Shock



(a) Local Projections

(b) Vector Autoregression

C.2. Sensitivity of the Moments to Changes in the Parameters

Figure C.2.1: LP-IRFs elasticities





Figure C.2.2: SVAR-IRFs elasticities



C.3. The Role of Sample Size for Fiscal and Monetary Policy Shocks

Figure C.3.1: The Role of T for the Estimated Impulse Responses to Fiscal Policy





Figure C.3.2: The Role of T for the Estimated Impulse Responses to Monetary Policy



D. The Monte Carlo Study: Results in Detail

D.1. The Local Projection Approach to Indirect Inference

Table 6: SMM estimates using LP - IRFs

	$\hat{\sigma}_c$	$\hat{h}$	$\hat{\sigma}_l$	$\hat{\varphi}$	$\hat{\xi}_w$	$\hat{\xi}_p$	$\hat{l}_w$	$\hat{l}_p$
Technology shock, ε_t^a								
<i>Mean</i>	1.35	0.79	2.66	5.73	0.67	0.59	0.48	0.16
<i>Bias</i>	0.08	-0.02	0.14	-0.58	-0.10	0.06	-0.05	-0.02
<i>Std dev.</i>	0.28	0.06	0.70	1.79	0.17	0.08	0.16	0.05
<i>RMSE</i>	0.29	0.06	0.71	1.89	0.19	0.10	0.17	0.06
Fiscal Policy, ε_t^g								
<i>Mean</i>	1.40	0.80	2.74	6.41	0.63	0.43	0.50	0.15
<i>Bias</i>	0.13	-0.01	0.22	0.10	-0.14	-0.10	-0.03	-0.03
<i>Std dev.</i>	0.14	0.13	0.68	1.84	0.12	0.12	0.17	0.05
<i>RMSE</i>	0.19	0.13	0.71	1.84	0.18	0.16	0.18	0.06
Monetary Policy, ε_t^m								
<i>Mean</i>	1.44	0.79	2.57	5.68	0.70	0.43	0.48	0.18
<i>Bias</i>	0.17	-0.02	0.05	-0.63	-0.06	-0.10	-0.06	-0.00
<i>Std dev.</i>	0.20	0.05	0.77	1.54	0.10	0.14	0.17	0.05
<i>RMSE</i>	0.26	0.05	0.77	1.66	0.12	0.17	0.18	0.05
Selected Responses to All Shocks								
<i>Mean</i>	1.36	0.80	2.11	5.45	0.60	0.51	0.49	0.15
<i>Bias</i>	0.10	-0.01	-0.41	-0.86	-0.17	-0.02	-0.04	-0.02
<i>Std dev.</i>	0.26	0.06	0.73	1.38	0.17	0.12	0.16	0.05
<i>RMSE</i>	0.28	0.06	0.84	1.63	0.24	0.12	0.17	0.06

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

Table 7: SMM estimates using SVAR - IRFs

	$\hat{\sigma}_c$	$\hat{h}$	$\hat{\sigma}_l$	$\hat{\varphi}$	$\hat{\xi}_w$	$\hat{\xi}_p$	$\hat{l}_w$	$\hat{l}_p$
Technology shock, ε_t^a								
<i>Mean</i>	1.24	0.81	2.59	6.78	0.67	0.52	0.45	0.15
<i>Bias</i>	-0.03	0.01	0.07	0.47	-0.10	-0.01	-0.09	-0.03
<i>Std dev.</i>	0.20	0.04	0.72	1.50	0.12	0.08	0.16	0.05
<i>RMSE</i>	0.21	0.04	0.73	1.57	0.16	0.08	0.18	0.06
Fiscal Policy, ε_t^g								
<i>Mean</i>	1.35	0.78	2.50	5.90	0.67	0.48	0.56	0.17
<i>Bias</i>	0.08	-0.03	-0.02	-0.41	-0.09	-0.05	0.02	-0.01
<i>Std dev.</i>	0.16	0.15	0.73	1.92	0.12	0.14	0.15	0.05
<i>RMSE</i>	0.18	0.16	0.73	1.96	0.15	0.15	0.16	0.05
Monetary Policy, ε_t^m								
<i>Mean</i>	1.29	0.80	2.57	5.70	0.69	0.37	0.53	0.17
<i>Bias</i>	0.02	-0.01	0.05	-0.62	-0.08	-0.16	-0.00	-0.00
<i>Std dev.</i>	0.25	0.05	0.72	1.32	0.09	0.08	0.16	0.05
<i>RMSE</i>	0.26	0.05	0.73	1.46	0.12	0.18	0.16	0.05
Selected Responses to All Shocks								
<i>Mean</i>	1.25	0.81	2.14	5.59	0.60	0.51	0.49	0.16
<i>Bias</i>	-0.02	0.00	-0.38	-0.72	-0.17	-0.02	-0.04	-0.02
<i>Std dev.</i>	0.23	0.06	0.75	1.38	0.14	0.10	0.16	0.05
<i>RMSE</i>	0.23	0.06	0.84	1.56	0.22	0.10	0.17	0.06