Indirect Inference: A Local Projection Approach

Juan Castellanos
European University Institute

Russell Cooper
European University Institute

ECB RCC5 Brownbag Seminar

July 5, 2022

Motivation

- Maximum likelihood estimation, despite providing a full-characterization of the DGP when successful, can be quite challenging.

- **Indirect Inference** is an alternative to *full-information* estimation techniques, which is more transparent and often more robust.

- A key feature of indirect inference is the use of an **auxiliary (econometric) model** to form a criterion function.
  - The most popular approach is a *vector autoregression* (VAR) for the variables of interest.

- Starting with Jordà (2005), **local projections** (LP) have become an increasing widespread alternative to study the propagation of structural shocks.

- ***How should we choose between VAR and LP*** *when estimating the structural parameters of our DSGE model?*

What We Do

- We use **Monte Carlo** methods to compare the performance of the two moment generating functions in a specific application involving the estimation of the parameters of a dynamic general equilibrium model (DSGE).
  - ∗ We use the log-linearized version of *Smets and Wouters (2007) model* as our data generating process.
  - ∗ We use the Root Mean Squared Error (RMSE) as a measure to assess the performance parameter by parameter and the J-statistic for the overall performance.

- We use our Local Projection approach to Indirect Inference in an actual **empirical application**.
  - ∗ We study how well the Smets and Wouters model is able to capture the response of key macro aggregates to technology, fiscal and monetary shocks.
  - ∗ We compare how the resulting parameter estimates differ from those in their paper.

Where We Stand in the Literature

- There are two papers that stand out and that we use as our starting point:
  * **Smith (1993)** $\Longrightarrow$ VAR coefficients as moments in a SMM exercise.
  * **Plagborg-Møller and Wolf (2021)** $\Longrightarrow$ LP with p lags as controls and VAR(p) estimators approximately agree at impulse response horizons $h \leq p$.
- However, we **differ from these two papers** in a few important matters:
  * Smith (1993) $\Longrightarrow$ we are using local projections coefficients.
  * Plagborg-Møller & Wolf (2021) $\Longrightarrow$ we are interested in the structural model parameters.
- The paper is also related to the literature on:
  * *Estimation of DSGE models* $\Longrightarrow$ Christiano et al. (2005), Smets and Wouters (2007) or Scalone (2018)
  * *Simulation studies of LP & VAR methods* $\Longrightarrow$ Li, Plagborg-Møller and Wolf (2021), Herbst and Johannsen (2020)

A Monte-Carlo Study

Setting up the Monte Carlo Study

Following the footsteps of Smith (1993, J Appl. Econom.)

- We solve and simulate the log-linearized version of the Smets and Wouters (2007) model to obtain repeated samples of macroeconomic aggregates $\{x_t^s\}_{s=1}^S$ that then we use to estimate a *subset of the parameters* of the model, Θ.
- We focus on the following **8 parameters**:
  - σ_c : intertemporal elasticity of substitution
  - h : habit parameter
  - σ_l : elasticity of labor supply
  - φ : investment adjustment cost parameter
  - ξ_w, ξ_p : Calvo adjustment probabilities
  - ι_w, ι_p : Degree of indexation to past inflation
- The remaining parameters are set at the estimated values from Smets and Wouters (2007).
- Hyper-parameter choices:
  - Observed sample size, $T = 300$
  - Number of "observed" time series, S = 100
  - Simulated sample size, $T^S = 3,000 \implies \tau = T^S/T = 10$

THE MOMENT GENERATING FUNCTIONS

The Auxiliary Econometric Model

- We focus on the ***estimated* impulse responses** of four variables:
  - y_t : output
  - c_t : consumption
  - i_t : investment
  - hw_t : hours worked

  to one (or a selection) of three following **shocks**:
  - ε_t^a : total factor productivity (TFP) shock
  - ε_t^g : fiscal policy (FP) shock
  - ε_t^m : monetary policy (MP) shock
- The IRFs are estimated using the traditional ***VAR + Cholesky decomposition*** (SVAR - IRFs) or the more recent Jordà (2005) approach which relies on ***Local Projections*** (LP - IRFs).
- In either case, the econometrician still needs to decide on at least two more things:
  - The impulse response horizon, H. We set $H = 20$.
  - The number of lags, p. We set $p = 4$.
- **Weighting matrix**: we use the inverse of the variance-covariance matrix of the moments $W = \Sigma^{-1}$

LP - IRFs SVAR - IRFs

Monetary Policy: Estimated IRFs for $S = 100$, $T = 300$

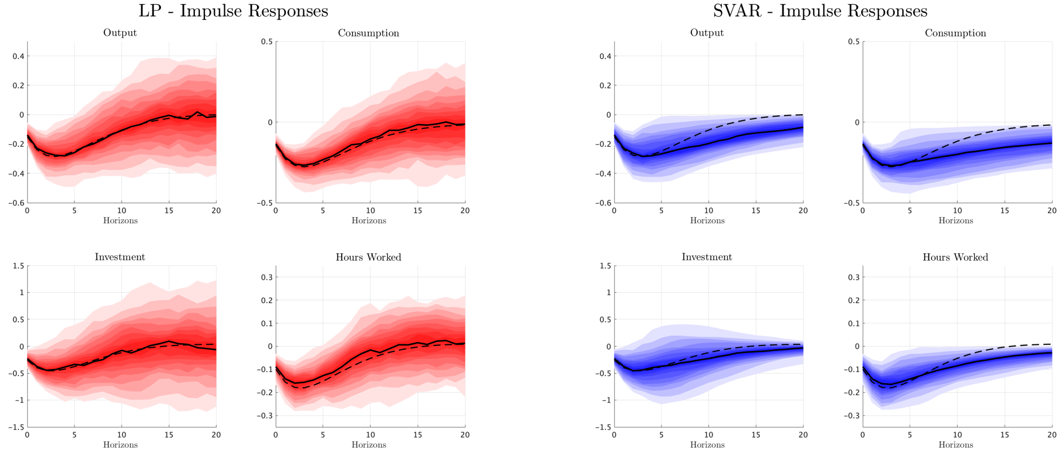

— Monte Carlo Median IRFs   - - Model/True IRFs

Technology Shock   Fiscal Policy

RESULTS

The Local Projection Approach

EUROPEAN UNIVERSITY INSTITUTE

	$\hat{\sigma}_c$	$\hat{h}$	$\hat{\sigma}_l$	$\hat{\varphi}$	$\hat{\xi}_w$	$\hat{\xi}_p$	$\hat{\iota}_w$	$\hat{\iota}_p$
	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

Definitions VAR Results Identification

Relative Performance: Local Projections vs SVAR

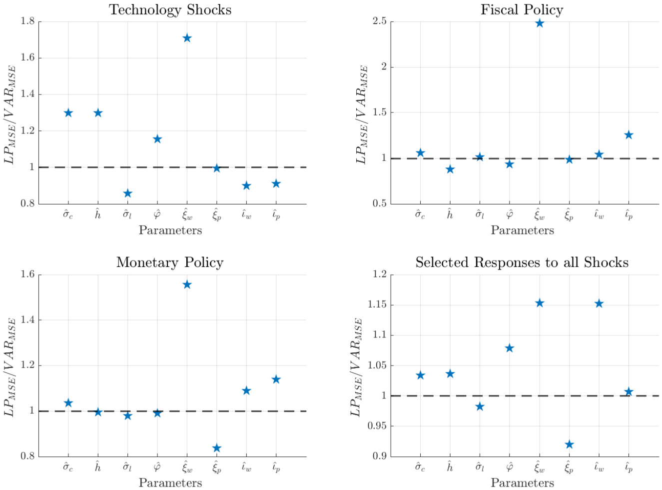

Across Different Shocks

Overall Performance: Local Projections vs SVAR

- The value of the loss function at the estimated parameters $\hat{\Theta}$, which is given by:

$$J(\hat{\Theta}) = \left(\mu^S(x_t;\hat{\Theta}) - \mu(x_t)\right)' W \left(\mu^S(x_t;\hat{\Theta}) - \mu(x_t)\right) \quad (1)$$

is a good measure to assess the overall performance of the estimation.

	Local Projections			Vector Autoregression		
	Avg. $J(\hat{\Theta})$	Max. $J(\hat{\Theta})$	Elapsed Time	Avg. $J(\hat{\Theta})$	Max. $J(\hat{\Theta})$	Elapsed Time
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

An alternative measure of overall performance

- Neither of the two previous measures, RMSE and J-statistic, can inform us about how close we are from the true/theoretical impulse responses.

- Thus, we also look at the weighted distance between theoretical IRFs coming from the model at the estimated parameter values $\hat{\Theta}$ and at the true values Θ^*.

	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

$\Longrightarrow$ **The LP-IRF approach to indirect inference does a significantly better job at picking those parameters that are relevant for the shape of the true impulse response function.**

Estimation: Empirical Applications

Applying the Local Projection Approach with Real Data

- We take from the literature some **empirically estimated IRFs** that used Jordà's (2005) **local projection** method.

- *Technology shocks* $\Longrightarrow$ Francis, Owyang, Roush, and DiCecio (2014)
    * We use Ramey's (2016) estimates of the responses of *real GDP*, *consumption*, *non-residential investment* and *hours* to an unanticipated TFP shock.

- *Fiscal policy shocks* $\Longrightarrow$ Blanchard and Perotti's (2002)
    * We also use Ramey's (2016) estimates of the responses of *GDP*, *non-durables + services consumption* and *non-residential investment* to a government spending shock.

- *Monetary policy shocks* $\Longrightarrow$ Romer and Romer (2005)
    * We use the estimates reported in Tenreyro and Thwaites (2016). They report the responses of *GDP*, *non-durable and services consumption* and *fixed business investment*.

Ramey (2016) Tenreyro and Thwaites (2016)

All shocks

The Response of Investment

Estimation Strategy & Results

1. **Strategy**
   * We focus on the response of only investment and try to match the response to all 3 shocks: i) technology, ii) fiscal policy and iii) monetary policy shocks.
   * We use the a diagonal weighting matrix whose entries are the inverse of the IRFs' standard deviation.
   * CI for the structural parameters are obtained through bootstrapping.
2. **The Distribution of the Parameter Estimates**

	$\hat{\sigma}_c$	$\hat{h}$	$\hat{\sigma}_l$	$\hat{\varphi}$	$\hat{\xi}_w$	$\hat{\xi}_p$	$\hat{\iota}_w$	$\hat{\iota}_p$
S&W 2007	1.26	0.80	2.52	6.31	0.70	0.66	0.58	0.24
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

Targeted Empirical vs. Estimated Investment IRFs

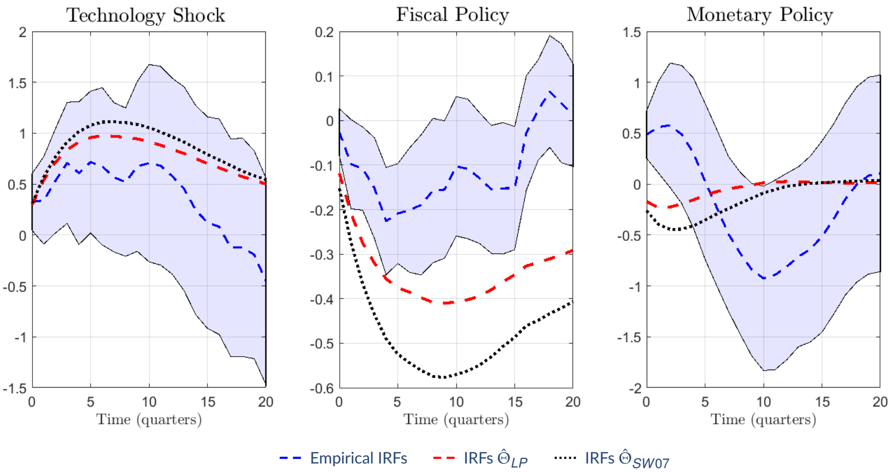

Untargeted Empirical vs. Estimated Output IRFs

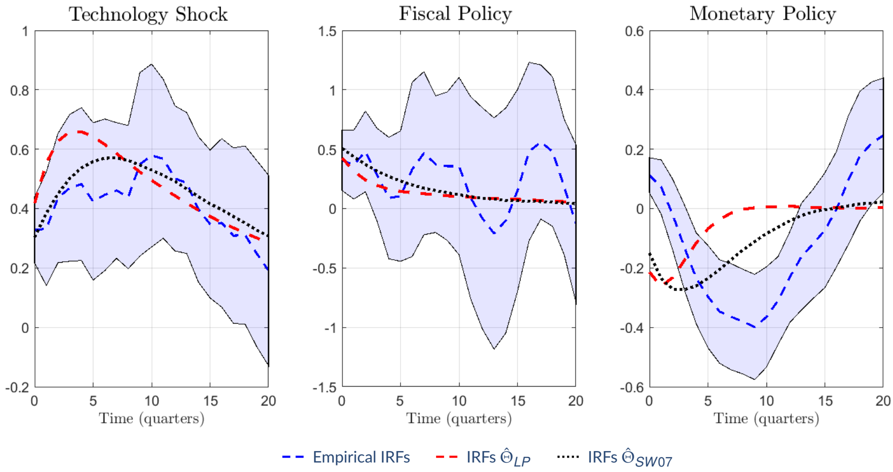

Untargeted Consumption Response

Concluding Remarks

What We Find

- Our Monte-Carlo analysis point to a trade-off between the two approaches:
  - The *VAR approach* performs better in terms of RMSE and overall fit (J-statistic)
  - However, the *LP approach* does a much better job in matching the true/theoretical impulse responses.

$\Longrightarrow$ **The *LP approach* picks up much better the relevant parameters for the IRFs.**

- Our empirical application show that
  - Either looking at shocks individually or jointly, we obtain similar parameter estimates than Smets and Wouters (2007).
  - Impulse responses to technology, fiscal and monetary shocks at these parameter estimates are similar to those obtained in the data.

$\Longrightarrow$ **The *LP approach* is a valid alternative for DSGE estimation.**

Thank You!

Local Projections (LP - IRFs)

- Some notation:
  - Let $\tilde{y}_t \in \{y_t, c_t, i_t, hw_t\}$ denote one of response variables of interest.
  - Let $\tilde{x}_t \in \{\varepsilon_t^a, \varepsilon_t^g, \varepsilon_t^m\}$ denote the innovation of one of the three aggregate shocks.
  - Define the vector of controls $w_t = \{\tilde{x}_t, \tilde{y}_t\}$.
- Then, consider for each horizon $h = 0, 1, 2, \ldots, H$ the *linear projections*:

$$\tilde{y}_{t+h} = \mu_h + \beta_h \tilde{x}_t + \sum_{\ell=1}^{p} \delta'_{h,\ell} w_{t-\ell} + \xi_{h,t} \tag{2}$$

where $\xi_{h,t}$ is the projection residual and $\mu_h, \beta_h, \{\delta'_{h,\ell}\}_{\ell=1}^{p}$ are the projection coefficients.

- **Definition**. *The LP - IRFs of $\tilde{y}_t$ with respect to $\tilde{x}_t$ is given by $\{\beta_h\}_{h\geq 0}$ in the equation above.*

Structural Vector Autoregression (SVAR - IRFs)

EUROPEAN UNIVERSITY INSTITUTE

- Consider the multivariate linear *VAR(p) projection*:

$$w_t = c + \sum_{\ell=1}^{p} A_\ell w_{t-\ell} + u_t \tag{3}$$

where u_t is the projection residual and c, $\{A_\ell\}_{\ell=1}^{p}$ are the projection coefficients.

- Let $\Sigma_u \equiv \mathbb{E}\left[u_t u_t'\right]$ and define the *Cholesky decomposition* $\Sigma_u = BB'$ where B is lower triangular with positive diagonal entries.

- Consider the corresponding recursive *SVAR representation*:

$$A(L)w_t = c + B\eta \tag{4}$$

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

- **Definition**. *The SVAR - IRFs of $\tilde{y}_t$ with respect to $\tilde{x}_t$ 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 above.*

Back

Technology Shock: Estimated IRFs for $S = 100$, $T = 300$

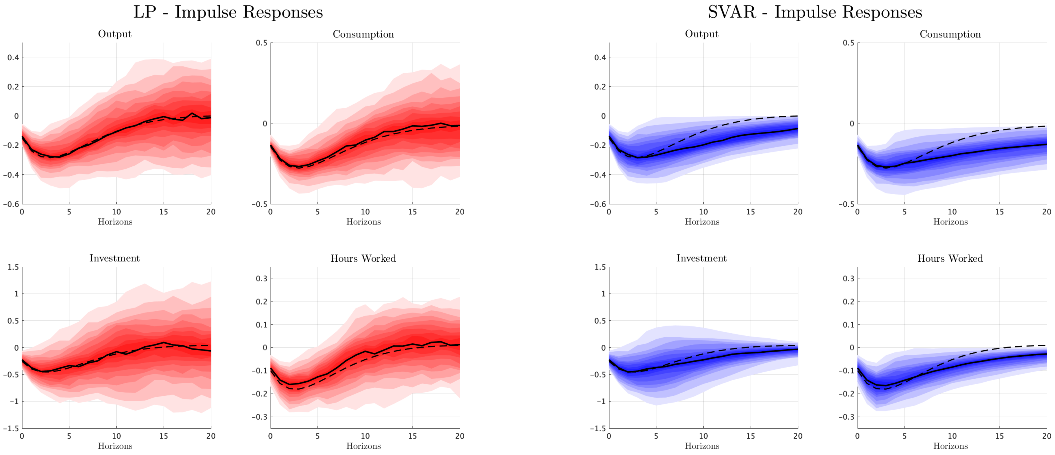

Monte Carlo Median IRFs — — Model/True IRFs

Fiscal Policy: Estimated IRFs for $S = 100$, $T = 300$

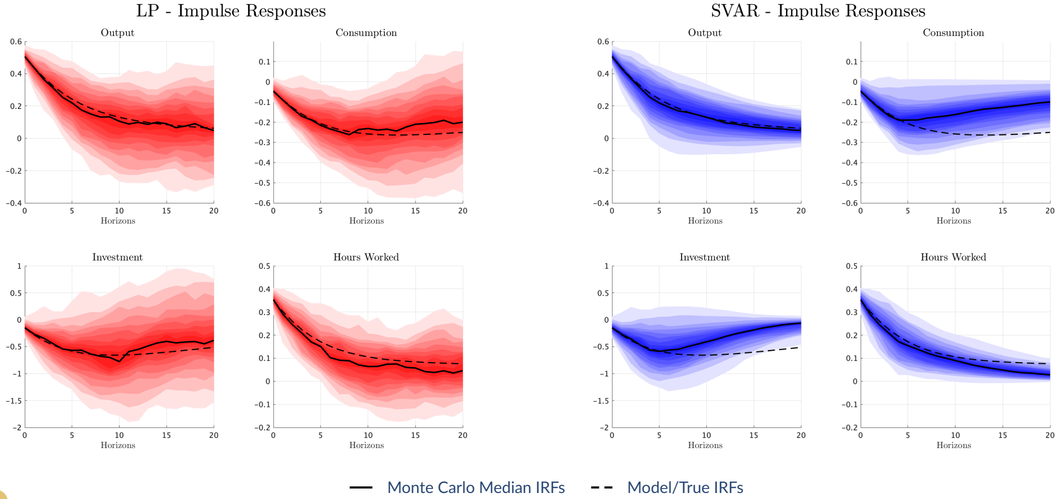

Back

Relative Performance Across Different Shocks

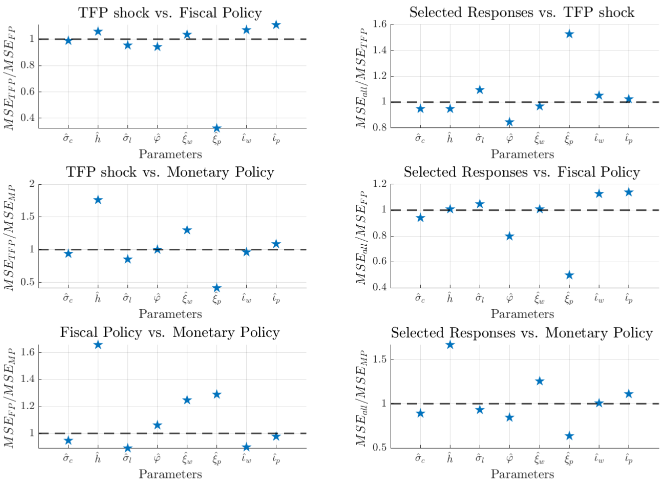

Back

Monte Carlo Results: What do we report?

EUROPEAN UNIVERSITY INSTITUTE

- Following Smith (1993), we compute for each of the estimated parameters $\hat{\theta}_h \in \hat{\Theta}$, the following *statistics*:

$$\text{Bias}_h \equiv \mathbb{E}\left[\hat{\theta}_h\right] - \theta_h$$

$$\text{Std dev}_h \equiv \sqrt{\text{Var}(\hat{\theta}_h)}$$

$$MSE_h^{1/2} \equiv \sqrt{\text{Bias}_h^2 + \text{Var}(\hat{\theta}_h)}$$

  where expectations are taken over the S Monte Carlo draws.

- In future work, it may be interesting to explore whether the **bias-variance trade off** documented in Li, Plagborg-Møller and Wolf (2021) for the parameters of our auxiliary models (VAR or LP) is also present in the estimated structural parameters. For example, by means of the following *loss function*:

$$\mathcal{L}_\omega\left(\theta_h, \hat{\theta}_h\right) = \omega \times \left(\mathbb{E}\left[\hat{\theta}_h - \theta_h\right]\right)^2 + (1-\omega) \times \text{Var}\left(\hat{\theta}_h\right)$$

Monte Carlo Results: Vector Autoregression

	$\hat{\sigma}_c$	$\hat{h}$	$\hat{\sigma}_l$	$\hat{\varphi}$	$\hat{\xi}_w$	$\hat{\xi}_p$	$\hat{\iota}_w$	$\hat{\iota}_p$
	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

Moments’ Sensitivity: Vector Autoregression IRFs

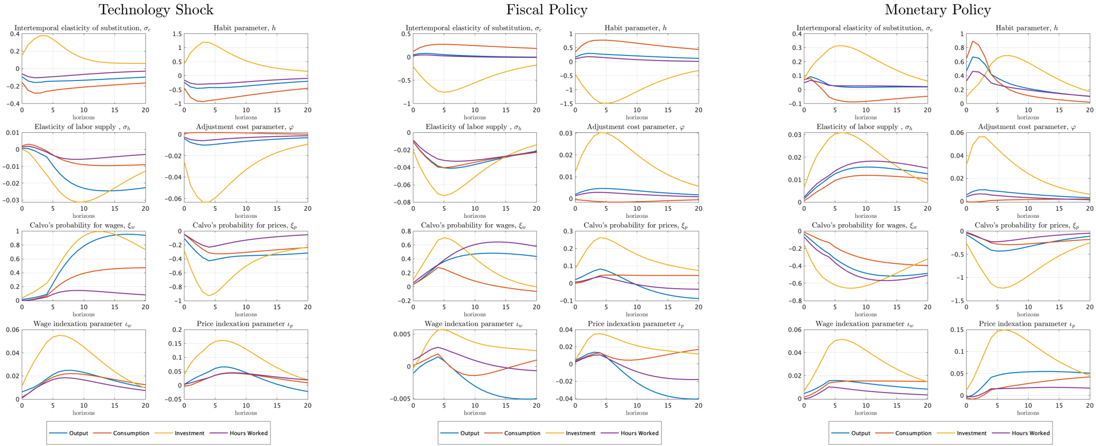

Moments' Sensitivity: Local Projection IRFs

Technology Shock

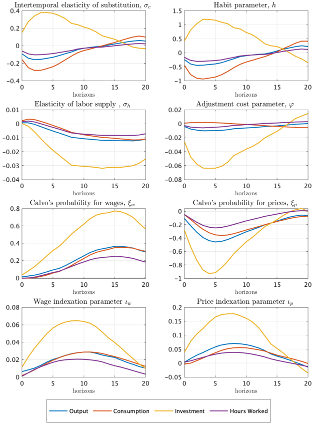

Fiscal Policy

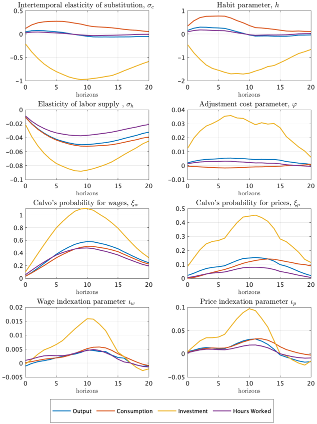

Monetary Policy

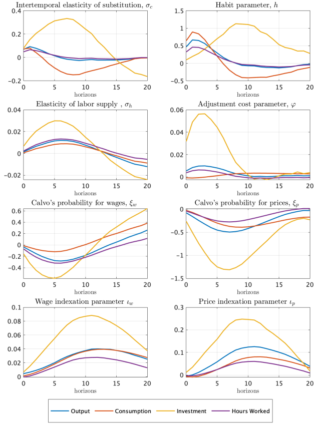

Back

Ramey (2016)

Technology & Fiscal Policy Shocks

Estimation Strategy

- **Fitted Parameters.** We try to identify the same eight parameters from the Monte Carlo study, that is we optimize over $\Theta = \{\sigma_c, h, \sigma_l, \varphi, \xi_p, \xi_w, \iota_p, \iota_w\}$.

- **Targeted Data Moments.**
  * *Technology shocks*: response of real GDP, consumption, non-residential investment and hours to an unanticipated TFP shock.
  * *Fiscal policy*: response of GDP, non-durables + services consumption, non-residential investment and hours worked to an unanticipated government spending shock.

  Local projection regression for each dependent variable:

$$z_{t+h} = \alpha_h + \theta_h \cdot \text{ shock }_t + \varphi_h(L) y_{t-1} + \text{ quadratic trend } + \varepsilon_{t+fh} \tag{5}$$

  $\implies 84 = 4 \times 21$ *moments*.

- **Auxiliary Econometric Model**. We simulate the model $T^s = 3,000$ times and use the simulated paths for output, consumption, investment and hours worked as well as the two shocks to estimate equation (1), in which we set $\tilde{y}_t \in \{y_t, c_t, i_t, hw_t\}$, $\tilde{x}_t = \{\varepsilon_t^a, \varepsilon_t^g\}$ and $p = 2$. The coefficients $\{\beta_h\}_{h=0}^{20}$ are then our simulated moments.

- **Weighting Matrix**. We use a diagonal matrix whose entries coincide with the inverse of the IRFs' standard deviation.

Estimation Results

- **Confidence intervals** for the structural parameters are obtained through **bootstrapping**. That is we repeat the estimation described above 100 times to get a *distribution* over the structural parameter estimates.

	$\hat{\sigma}_c$	$\hat{h}$	$\hat{\sigma}_l$	$\hat{\varphi}$	$\hat{\xi}_w$	$\hat{\xi}_p$	$\hat{\iota}_w$	$\hat{\iota}_p$
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

Model Validation: Empirical vs. Model Estimated IRFs

(a) Technology Shocks

(b) Fiscal Policy

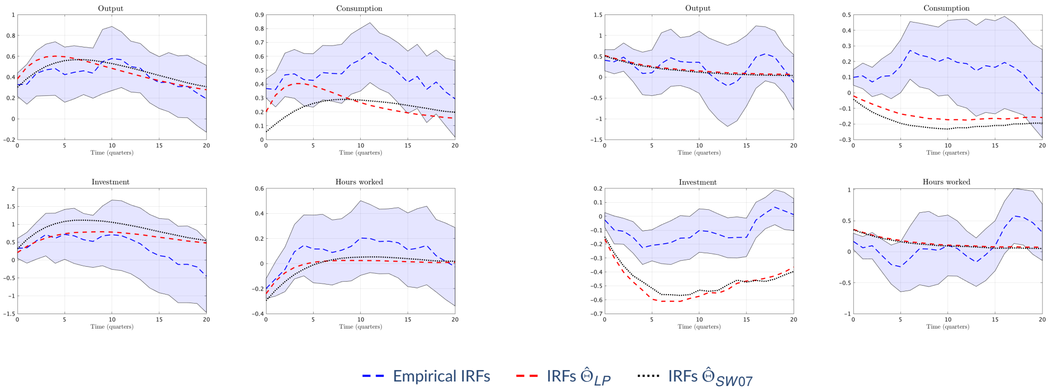

Tenreyro & Thwaites (2016)

Monetary Policy Shocks

Estimation Strategy

- The response of output, consumption and investment to monetary policy is state-dependent.
- The Smets and Wouters model does not contain the relevant non-linearities to capture these effects, so why should we care?
- We think it is still constructive to match the non-linear responses because it may inform us about which parts of the model are more reactive to this state-dependance, i.e. those in which the parameters differ the most across the two scenarios.
- Tenreyro and Thwaites (2016) empirical specification:

$$y_{t+h} = \tau t + F(z_t)\left(\alpha_h^b + \beta_h^b \varepsilon_t + \gamma^{\mathbf{b}'} \mathbf{x_t}\right) + (1 - F(z_t))\left(\alpha_h^r + \beta_h^r \varepsilon_t + \gamma^{\mathbf{r}'} \mathbf{x_t}\right) + u_t \tag{6}$$

- The estimation routine is very similar to that used for matching the responses to technology and fiscal shocks.

Estimation Results

	$\hat{\sigma}_c$	$\hat{h}$	$\hat{\sigma}_l$	$\hat{\varphi}$	$\hat{\xi}_w$	$\hat{\xi}_p$	$\hat{\iota}_w$	$\hat{\iota}_p$
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

Linear Model: Empirical vs. Estimated IRFs

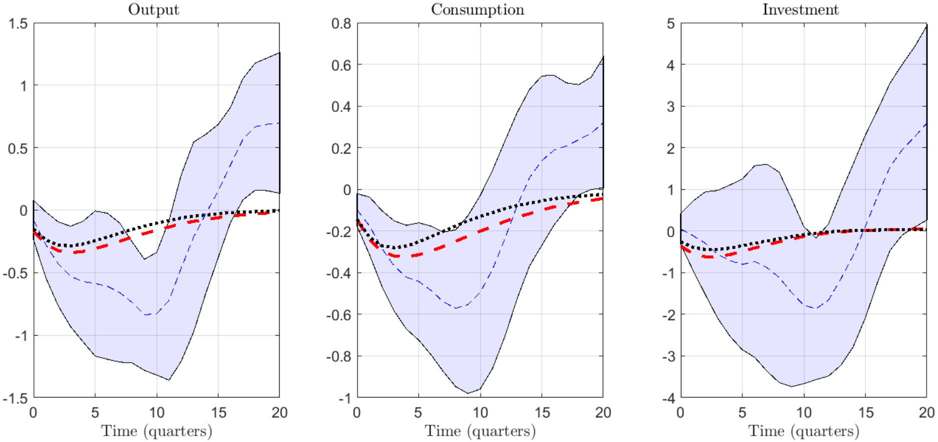

Empirical IRFs — IRFs $\hat{\Theta}_{LP}$ — IRFs $\hat{\Theta}_{SW07}$

Non-Linear Model: Empirical vs. Estimated IRFs

(a) Non-Linear: Expansion

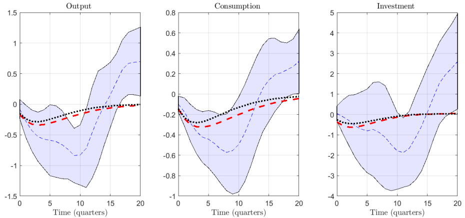

(b) Non-Linear: Recession

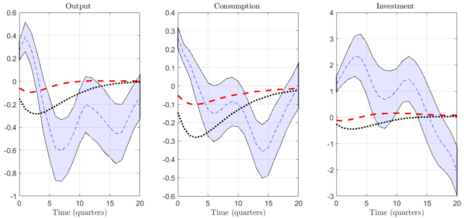

Empirical IRFs — IRFs $\hat{\Theta}_{LP}$ — IRFs $\hat{\Theta}_{SW07}$

Back

Untargeted Empirical vs. Estimated Consumption IRFs

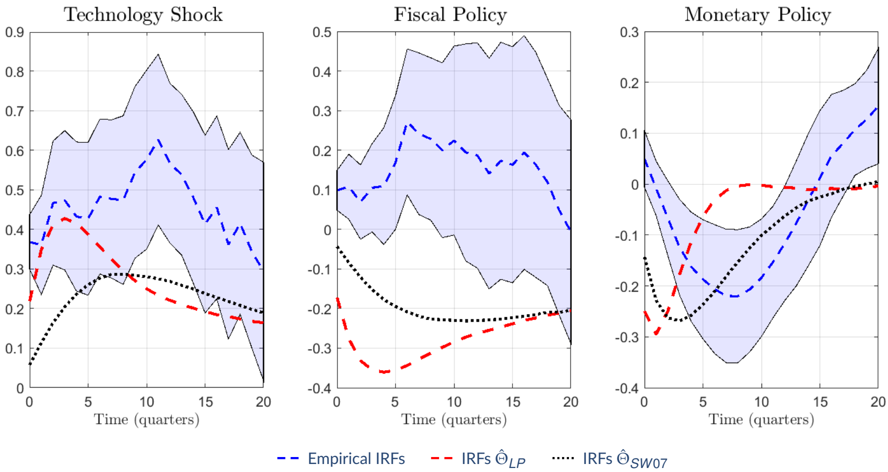

Back