

Impulse Response Functions for Self-Exciting Nonlinear Models*

Neville Francis[†] Michael T. Owyang[‡] Daniel Soques[§]

November 15, 2023

Abstract

We calculate impulse response functions from regime-switching models where the driving variable can respond to the shock. Two methods used to estimate the impulse responses in these models are generalized impulse response functions and local projections. Local projections depend on the observed switches in the data, while generalized impulse response functions rely on correctly specifying regime process. Using Monte Carlos with different misspecifications, we determine under what conditions either method is preferred. We then extend model-average impulse responses to this nonlinear environment and show that they generally perform better than either generalized impulse response functions and local projections. Finally, we apply these findings to the empirical estimation of regime-dependent fiscal multipliers and find multipliers less than one and generally small differences across different states of slack.

Keywords: generalized impulse response functions, local projections, threshold models, model averaging

JEL Codes: C22, C24, E62

*The authors benefited from helpful comments from Graham Elliott, Ana María Herrera, Oscar Jorda, Valerie Ramey, seminar participants at Northeastern University, Indiana University, Kansas University, and participants at the CFE 2021 Conference. The views expressed here are the authors' alone and do not reflect the opinions of the Federal Reserve Bank of St. Louis or the Federal Reserve System.

[†]Department of Economics, University of North Carolina, Chapel Hill and NBER

[‡]Corresponding author. Research Division, Federal Reserve Bank of St. Louis.
michael.t.owyang@stls.frb.org

[§]Department of Economics and Finance, University of North Carolina Wilmington

1 Introduction

Recent empirical analyses evaluate the differential effect of macroeconomic shocks across the business cycle. An important aspect of this is determining how these shocks affect future business cycle regimes [Owyang et al. (2013) and Jackson et al. (2018)]. This type of analysis computes impulse response functions (IRFs) from models where the regimes respond to economic conditions that, in turn, respond to the macroeconomic shocks (i.e., the so-called self-exciting regime switching models).¹

Computing IRFs for self-exciting models, however, is not trivial and is typically accomplished using one of two methods: generalized impulse responses functions [GIRFs, Koop et al. (1996)] and regime-dependent local projections [LP, Jordà (2005)]. The former constructs IRFs using Monte Carlo methods that simulate possible paths for the responding variables by iterating the model forward. If the regime is parametrically-related to these responding variables, the regime path can be determined for each response horizon. The latter constructs the IRF at each horizon as a direct multistep forecast that can account for future regime switches observed in the data.

While there are papers in the literature that use either method, LP has become increasingly popular—in part, because of its computational simplicity. Starting with Auerbach and Gorodnichenko (2013), a number of papers have argued that LP is robust to the types of nonlinearities associated with state-dependence [see also Owyang et al. (2013); Ramey and Zubairy (2018, hereafter RZ) among many others]. LPs are still linear but capture future switches by projecting onto future data where switches occur. In this sense, LP accounts for regime-switching *empirically*, rather than based on the shock. GIRFs, on the other hand, let the transition variable respond to the shock but are computed parametrically and, therefore, rely heavily on the model being correctly specified.

One might reasonably wonder under what conditions each method’s liabilities dominate.

¹Another option found in the literature computes IRFs holding the state of the economy fixed throughout the response horizon [e.g., Auerbach and Gorodnichenko (2012)]. Fixing the regimes, however, may overestimate the duration of cyclical downturns in the presence of expansionary policy shocks.

We consider the advantages of each of the methods for computing IRFs in various self-exciting model frameworks under different forms of misspecification. These misspecifications include estimating the wrong transition function, using the wrong transition variable, and choosing the wrong number of lags in the VAR. We then propose the use of model-average GIRFs (MA-GIRFs) that use weights based on the BIC.

Based on simulation exercises, we arrive at a number of conclusions. First, if one knows the true model, one should estimate GIRFs from the true model. Second, in the presence of uncertainty about the transition process, the MA-GIRF that weights all the GIRFs from possible models dominates both LP and the use of any single model. Third, LP only dominates when the lag order of the VAR is truncated in the estimated model. However, computing the MA-GIRF with VARs with a large lag order in the model choice set can mitigate this issue.

Our methodological contribution lies at the intersection of two rapidly expanding strands of literature. First strand compares IRFs from a VAR and LP.² While the majority of this literature explores the linear case, [Gonçalves et al. \(2023\)](#) use a class of nonlinear models similar to ours and consider the effect of varying the magnitude of the structural shock. They show that the conditional average IRF estimated with regime-dependent LP is asymptotically biased when the shock is large [of the order of the military news shock identified by [Ramey \(2011\)](#) and [Ramey and Zubairy \(2018\)](#)] and the state process is endogenous. The bias stems from LP's failure to capture how the shock affect the evolution of the state of the economy over the response period. The bias is likely to be larger with larger innovations.³ However, by directly modeling the transition processes, regime-dependent GIRF estimators would overcome this bias.

The second strand uses model averaging as a solution to model uncertainty problems when computing IRFs [e.g., [Ho et al. \(2023\)](#) and [Li et al. \(2021\)](#)]. [Ho et al. \(2023\)](#) averages

²[Kilian and Kim \(2011\)](#), [Plagborg-Møller and Wolf \(2021\)](#), and [Li et al. \(2021\)](#) among others consider these differences for linear models.

³[Gonçalves et al. \(2023\)](#) demonstrates that LP performs well if the economy is more likely to remain in the same state after a shock, often with infinitesimal shocks.

over a variety of structural models using prediction pool in the spirit of [Geweke and Amisano \(2011\)](#); [Li et al. \(2021\)](#) constructs averages using the Stein combination estimator proposed in [Hansen \(2016\)](#). These papers differ from ours in at least two dimensions. First, our application extends this literature to self-exciting nonlinear models. Second, our weighting function is based on in-sample fit, which both prevents us from including the LP in our model average and limits us to having constant model weight across horizons.

To illustrate the usefulness of our approach, we estimate regime-dependent fiscal multipliers from local projections, the various nonlinear VARs, and our model average approach [[Auerbach and Gorodnichenko \(2012, 2013\)](#); [Fazzari et al. \(2015, 2021\)](#); [Owyang et al. \(2013\)](#); [Ramey and Zubairy \(2018\)](#); [Laumer and Phillips \(2020\)](#), among many others]. Accounting for future state dynamics is vital to the estimation of the fiscal multiplier, as noted by [RZ](#) and [Caggiano et al. \(2015\)](#). We estimate fiscal multipliers across states of the business cycle and find relatively small multipliers (less than one) in both slack and nonslack phases across most model specifications. We find scant evidence that the fiscal spending multiplier depends upon the degree of economic slack. The only cases when the multiplier is substantially different across the business cycle is when local projections are used in tandem with the timing assumptions from [Blanchard and Perotti \(2002\)](#). Our results suggest that previous studies that find state dependent multipliers likely are not robust to model misspecification or alternative shock identification.

The balance of the paper is outlined as follows: Section 2 presents the class of models we study and the methods we will use to compute the impulse responses. Section 3 describes the Monte Carlo experimental design, the metrics used to gauge model-performance, and the simulation results. Section 4 applies the various methods to estimate state-dependent fiscal multipliers. Section 5 concludes.

2 Methodology

Computing impulse response functions (IRFs) in linear settings is straightforward and poses little challenge to researchers. However, nonlinearities can manifest in a variety of ways, each of which may present different problems for the computation of IRFs. To limit the scope of our investigation, we restrict attention to the class of models to self-exciting regime-switching VAR models that vary based on the specification of their state transition processes. A unifying characteristic of the models is that the variable driving the switching process—hereafter denoted z_t —will be an element of the VAR. To remove one source of uncertainty, we assume that the structural shock, u_t , is observed.

2.1 Self-Exciting Models

We consider three models: (i) the Threshold VAR (TVAR), (ii) the Time-Varying Transition Probability Markov-Switching VAR (MSVAR), and (iii) the Smooth Transition VAR (STVAR). These models are regime-switching, produce state-dependent impulse responses, and are common in the literature. To fix ideas, a vector of endogenous variables Y_t propagates according to:

$$Y_t = (1 - S_t) Y_{0t} + S_t Y_{1t} + \varepsilon_t, \quad (1)$$

where $S_t \in \{0, 1\}$ in the case of the TVAR and the MSVAR; $S_t = [0, 1]$ in the case of the STVAR; and $\varepsilon_t \sim N(0, \Sigma_t)$, where Σ_t is the state-dependent variance-covariance matrix

$$\Sigma_t = (1 - S_t) \Sigma_0 + S_t \Sigma_1. \quad (2)$$

In each state $k \in \{0, 1\}$, the state-dependent VAR is:

$$Y_{kt} = C_k + B_k(L)Y_{t-1} + D_k(L)u_t,$$

where C_k is a state-dependent intercept, $B_k(L)$ and $D_k(L)$ are matrix polynomials in the lag operator.

This class of models has several key features. First, the variable governing the state transition process, denoted z_t , is an element of Y_t . Thus, the structural shocks, u_t , that propagate through Y_t produce future changes in regime. Second, the economy takes on two alternative dynamics (or, in the case of the STVAR, a convex combination of two alternative dynamics), dictated by the realization of the underlying state S_t . When $S_t = 0$, the economy has $C_0, B_0(L), D_0(L)$ dynamics and when $S_t = 1$, the economy has $C_1, B_1(L), D_1(L)$ dynamics; thus, the model is linear, conditional on S_t being known.

The main difference between the three data generating processes is the specification of the state transition process. In the TVAR, the regime process is determined by the transition variable, z_{t-d} , relative to a fixed threshold, z^* , as:

$$S_t = \begin{cases} 1 & \text{if } z_{t-d} > z^* \\ 0 & \text{if } z_{t-d} \leq z^* \end{cases} .$$

The regime dynamics are discrete, meaning that either regime 0 or regime 1 prevails at any time t .

The TVTP-MSVAR [see [Diebold et al. \(1993\)](#) and [Filardo \(1994\)](#) for univariate and [Billio et al. \(2016\)](#) for VARs] has state transition probabilities that depend on a lag of z_t via a logistic transition function, where

$$\Pr [S_t = j | S_{t-1} = i] = \frac{\exp(\bar{\gamma}_{ji} + \gamma_{ji} z_{t-d})}{\sum_k \exp(\bar{\gamma}_{ki} + \gamma_{ki} z_{t-d})} \quad (3)$$

for each of the regimes with $\sum_k p_{ki}(z_{t-d}) = 1$ for all i, t .

The STVAR [see [Granger and Teräsvirta \(1993\)](#)] has a continuous state process, $S_t(z_{t-d}) \in [0, 1]$, taking the form of a first-order logistic function:

$$S_t(z_{t-d}; \gamma, c) = [1 + \exp(-\gamma(z_{t-d} - c))]^{-1}, \quad (4)$$

where $\gamma \geq 0$ is the speed of transition (As $|\gamma| \rightarrow \infty$, the transition becomes sharper and the regime switches resemble a pure threshold model: At $\gamma = 0$, the model collapses to a linear model) and c is a fixed threshold. In equation (4), the regime process is determined by the sign and magnitude of the deviation of z_{t-d} from the threshold c .

2.2 Methods for Computing Impulse Responses

IRFs can be viewed as the difference between two forecasts—one conditional on a structural shock at time t and one conditional on no shock—evaluated over a number of horizons. Leaving aside the problem of identifying the structural shock, IRFs can be based on either iterative multistep forecasts (conventional IRFs computed by propagating a VAR(1)) or direct multistep forecasts (local projections). The literature comparing methods for computing IRFs in linear models is large [Kilian and Kim (2011), Plagborg-Møller and Wolf (2021), Li et al. (2021)]; relatively fewer papers consider use nonlinear VARs [Gonçalves et al. (2023)].

2.2.1 Generalized Impulse Response Functions

One method that accounts for future regime changes is generalized impulse response functions [see Koop et al. (1996); GIRFs].⁴ GIRFs use Monte Carlo (MC) methods to draw from (i) the simulated historical paths of the data and (ii) future reduced-form shocks to obtain the difference between one path where a structural shock size, $u_t = \delta$, occurs at time t and a path where no structural shock occurs ($u_t = 0$).⁵ To account for the regime switching, we compute the state-dependent GIRF as the average over the set of histories where $S_{t-1} = i$.

⁴If the transition probabilities are constant, Krolzig (2006) suggests computing a weighted average of the regime-dependent responses, where the weights are determined by the transition probabilities. These models are also called “exogenous” switching models and are explored in detail in Gonçalves et al. (2023). Because our interest is in exploring the effect of the interaction between variables in the VAR and the state process, we forgo analyzing exogenous switching models here. One can, however, think of the exogenous models as a special case of endogenous switching models.

⁵In the linear model, δ is a scaling factor that leaves the shape of the IRF unchanged. In nonlinear models, differences in the size of the shock have different effects on the future state process. This aspect of size asymmetry is one of the issues explored in Gonçalves et al. (2023) and we defer analysis of variation in δ to that paper. We note, however, that one can easily augment the methods used here with those in Gonçalves et al. (2023).

Define $\{\mathbf{Y}_\tau^i\}$ as the set of all histories that correspond to $S_\tau = i$, with $i = 0, 1$. Let $R_1 = \sum_{t=1}^T S_t$ represent the number of occurrences of $S_t = 1$ in our dataset and let $R_0 = T - \sum_t S_t$ represent the number of occurrences of $S_t = 0$. Conditional on a history, \mathbf{Y}_{t-1}^i , we can simulate two paths for Y_t, \dots, Y_{t+h} —one based on a structural shock $u_t = \delta$ occurring at time t and one with no structural shock, $u_t = 0$. Both paths are subject to the same path of future reduced-form shocks, $\{\varepsilon_{t+l}^{[q]}\}_{l=0}^h$. We then average over Q Monte Carlo draws of future shocks to obtain the expectation of the two paths, differentiated by the presence of the period- t shock:

$$E_t [Y_{t+h} | \Theta, \mathbf{Y}_{t-1}^i, u_t] = \frac{1}{Q} \sum_q \left[Y_{t+h} | \Theta, \mathbf{Y}_{t-1}^i, u_t, \{\varepsilon_{t+l}^{[q]}\}_{l=0}^h \right],$$

where Θ represents the set of model parameters. The value of the GIRF at horizon h after a structural shock of size δ occurs at time t with previous regime $S_{t-1} = i$ is obtained by averaging across all $S_{t-1} = i$ histories:

$$\Phi_{GIRF}^i(h, \delta, \Theta) = \frac{1}{R_i} \sum_{\mathbf{Y}_{t-1}^i \in \{\mathbf{Y}_{t-1}^i\}} \{E_t [Y_{t+h} | \Theta, \mathbf{Y}_{t-1}^i, u_t = \delta] - E_t [Y_{t+h} | \Theta, \mathbf{Y}_{t-1}^i, u_t = 0]\}. \quad (5)$$

Allowing the endogenous variables to start in different states and accounting for the future paths of the driving variable captures future regime switches over the response periods.

2.2.2 Regime-Dependent Local Projections

GIRFs compute the difference between iterated multistep forecasts, while LPs compute the difference between direct multistep forecasts. Generating the forecasts that comprise the GIRFs requires assumptions about the parametric model. LP, on the other hand, may mitigate the need for such assumptions, making it more robust to the presence of nonlinearities and possible misspecification [[Jordà \(2005\)](#); [Plagborg-Møller and Wolf \(2021\)](#); [Montiel Olea](#)

and Plagborg-Møller (2021)].

For regime-switching models, RZ propose computing regime-dependent LPs by conditioning on the regime at the shock incidence. This version of LP regresses the responding variable on the value of the period- t observed structural shock and a period- t vector of controls separately for each horizon h .

$$Y_{t+h} = (1 - S_t) (\alpha_{0h} + \beta_{0h}u_t + \gamma_{0h}X_t) + S_t (\alpha_{1h} + \beta_{1h}u_t + \gamma_{1h}X_t) + v_{t+h},$$

where S_t is the prevailing regime at the time of the shock and X_t is the vector of controls.⁶ Note that the regressions are horizon specific. Over the entire response horizon, we compute the LP IRFs as:

$$\Phi_{LP}^i(h, \delta) = \beta_{ih}\delta,$$

where the horizon-specific coefficient, β_{ih} , depends on the state at the time of the shock. The idea is that the use of Y_{t+h} as the regressand empirically accounts for the possible regime changes between periods t and $t + h$.

2.2.3 Model-Average Generalized Impulse Responses

An alternative method of addressing model uncertainty is the use of model averaging to construct a weighted average of the object of interest (in this case, the estimated IRFs) of each model [e.g., Ho et al. (2023) and Li et al. (2021) for linear models; see Steel (2020) for a recent overview of model averaging].

Let $m = 1, \dots, M$ represent the different models under consideration. The MA-GIRF is:

$$\Phi_{BMA}(h, \delta) = \sum_{m=1}^M w_m \Phi_m(h, \delta)$$

⁶We consider only discrete S for the starting conditions. Extension to non-discrete S in the spirit of Ruisi (2019), Lusompa (2021), and Inoue et al. (2022) is straightforward.

where the model weights, w_m , are approximate posterior model probabilities as in [Raftery \(1995\)](#):⁷

$$w_m = \frac{\exp(-BIC_m/2)}{\sum_{j=1}^M \exp(-BIC_j/2)}.$$

Several features of the MA approach proposed here are worth highlighting. Our weights do not vary across the impulse response horizon [as they do in [Ho et al. \(2023\)](#) and [Li et al. \(2021\)](#)] and do not vary across regimes. They are constructed using an in-sample measure of fit, and, thus, preclude use of the LP IRFs in the model set. While the weights are applied to the IRFs directly and not to the coefficients, they could be applied directly to any object of interest (e.g., the multiplier in our application). Finally, our weights can also accommodate priors information either on the models or on the impulse responses themselves.

3 Simulation Study

We compare methods of generating IRFs in Monte Carlo exercises where the “true” IRFs are known. Given a data-generating model, at each Monte Carlo iteration, we simulate $\mathbf{Y}_T = \{Y_1, \dots, Y_T\}$, $\mathbf{S}_T = \{S_1, \dots, S_T\}$, and a time series of exogenous structural shocks. IRFs to the known structural shocks are then computed using the different methods discussed above.

For each draw of the data, GIRFs are obtained from the parametric models for various forms of misspecification [e.g., lag lengths, conditioning variables, and the driving variable]. At each draw of the Gibbs sampler, 1000 artificial histories are used to compute the GIRFs. For LP, the IRFs are computed using lags of Y_t and lags of the shocks as control vectors. The state of the economy S_t is explicitly modelled in the three VAR based methods, whereas nonlinear LP is estimated conditional on knowing S_t . One could estimate LP using the true states which would provide the LP with an unfair advantage. Instead, we define the state for the LP based on whether the threshold variable is above or below its historical average,

⁷[Hansen \(2007\)](#) refers to this form of model averaging as Smoothed BIC. Using the BICs to weight the IRFs can lead to better out-of-sample characteristics [[Swanson and Zeng \(2001\)](#)]

\bar{z} :

$$\hat{S}_t^{LIP} = \begin{cases} 1 & \text{if } z_{t-d} > \bar{z} \\ 0 & \text{if } z_{t-d} \leq \bar{z} \end{cases}.$$

For the MA-GIRFs, we set the model weights as above with a flat model prior.

3.1 Defining the True Impulse Responses

In a self-exciting model, solving for the true IRF analytically is intractable [see [Gonçalves et al. \(2023\)](#)]. The true state-dependent IRF are obtained using simulation methods, similar to estimating GIRFs as in [Koop et al. \(1996\)](#). The true IRF, $\Phi_{GIRF}^i(h, \delta, \tilde{\Theta})$, is obtained as a GIRF from the correct model specification and true parameters, $\tilde{\Theta}$. The simulated histories, $\{\mathbf{Y}_t^i\}$, are drawn from a longer sample $Y_{T^L} = \{Y_t\}_{t=1}^{T^L}$ to more accurately approximate the true IRF, while only the last $T < T^L$ observations are used in estimation.

The true parameters of the DGP are calibrated to a fiscal VAR that includes government spending, output, and the Ramey military news shock for the period 1949 to 2015 [see [Ramey and Zubairy \(2018\)](#)]. The results below are based on 200 simulated data paths from 25 parameterizations drawn from the posterior distributions of each model. We set $T = 250$ and $T^L = 1000$ in our simulations.

3.2 Performance Metric

We evaluate the estimated IRFs using the Continuously Ranked Probability Score (CRPS, [Matheson and Winkler \(1976\)](#)), either taken as a function of h or averaged over the response horizon, $h = 1, \dots, H$.⁸ The CRPS compares the fit of the entire posterior distribution with the point-valued true response at each horizon h , using the empirical CDF-based approxi-

⁸We also compared the estimated IRFs using the Deviation in State Asymmetry (DSA), which measures the difference in the responses across the initial regime at each horizon. Results for DSA are comparable to those with CRPS and are available in the online appendix.

mation from Krüger et al. (2016) from the MCMC output of each model:

$$CRPS(F, h) = \int_{-\infty}^{\infty} \left[F(\chi) - \mathbb{1}(\chi > \Phi_{GIRF}^i(h, \delta, \tilde{\Theta})) \right]^2 d\chi,$$

where F is the forecast CDF implied by the posterior distribution of $\Phi_{GIRF}^i(h, \delta, \Theta)$. The indicator function, $\mathbb{1}(\chi > \Phi_{GIRF}^i(h, \delta, \tilde{\Theta}))$, captures the true distribution, equaling 0 for all χ values up until the true value $\Phi_{GIRF}^i(h, \delta, \tilde{\Theta})$, then 1 for greater values. Therefore, if our posterior were precisely estimated, $F(\chi)$ would follow exactly this pattern. A lower CRPS indicates a better fit as the probabilistic forecast of the CDF is closer to the true response. In comparing one model's overall fit to another, we will consider $CRPS(F)$ which is $CRPS(F, h)$ integrated over impulse response horizons $h = 1, \dots, H$.

3.3 Simulation Results

Table 1 shows the CRPS for the IRFs obtained using various methods, where each panel depicts a different potential source of misspecification. Each number presented in the table is the average over each horizon up to and including $h = 20$. Within a panel, the rows show the true model and the columns show the method used to compute the impulse response. The first three columns provide results for GIRFs using the indicated estimated model; the final two columns show results for LP and the MA-GIRFs weighted across the three GIRFs. In the top panel, the form of the transition process is the only source of misspecification; thus, the diagonal of the top panel represents the use of the true (correct) transition process. Subsequent panels also include misspecification in (i) the lag used for the transition function, (ii) the transition variable, and (iii) the VAR lag.⁹ The bottom panel shows results when the true model is a linear VAR but the estimated impulse response are computed assuming state asymmetry.

⁹For (i), true model uses z_{t-1} as the transition variable but the estimated model uses z_{t-4} . For (ii), the transition variable in the estimated model is uncorrelated with the true z_t . For (iii), the true model is a VAR(4) but the estimated model is a VAR(1).

A broad evaluation of the results suggest a number of key takeaways. First, except for when the STVAR is the true model, the GIRF from a model estimated with the true transition function generally exhibits less bias than alternative methods, including LP. This result obtains even when misspecifying the transition function lag or variable, so long as the functional form is correct. In these cases, LP exhibits more bias, possibly because it does not model the state transition process directly. Structural shocks that affect the driving variable put the state process on a counterfactual path; regime-dependent LP uses (an average of) empirical paths that can differ from this counterfactual path. Thus, if one is relatively certain about the functional form of the state transition process, one is better off computing GIRFs from estimates of that model.

Figure 1 shows the median CRPS for each model at each h to demonstrate how the relative performance of each model varies over the 20-period horizon. For each model, the bias is smaller at short horizons relative to long horizons, with LP’s performance deteriorating faster than the GIRFs as h increases. The latter result likely obtains because LP’s approximation of the underlying state worsens as h increases. Thus, estimates of shorter-run effects will be less biased than longer-term estimates. [Gonçalves et al. \(2023\)](#) similarly find in simulations that the asymptotic bias is larger when computing four-year responses compared to two-year estimates.¹⁰

Second, LP can exhibit less bias than GIRFs when the lag order of the VAR is truncated. The result obtains even for LP relative to MA-GIRFs, suggesting that the VAR dynamics themselves—not the form of the transition process—are of first-order interest. Including more lags than the true DGP, however, mitigates the problem for the GIRFs. Thus, we conclude that MA-GIRFs dominate LP if the dynamics of the true model can be approximated from the model choice set.¹¹ Thus, prior model uncertainty about the transition process and/or the VAR lag order can be resolved by estimating the MA-GIRFs and including a

¹⁰The appendix presents CRPS estimates by horizon for the other misspecification cases.

¹¹Note that the MA does not always place the largest weight on the true transition process.

VAR with a high lag order.¹²

Third, we note that GIRFs from the estimated STVAR are typically the most biased of the three models, even when the STVAR is the true model. This result suggest that the STVAR should only be used if one has strong prior belief that it is the true model.¹³ Conversely, the TVAR—the most sparsely parameterized model of the three—generally performs well, even when it is not the true transition process. Further, the MA-GIRF can perform better than even the true model. Taken together, these results highlight that IRF computation may be thought of more as a forecasting problem rather than a model specification problem.

Finally, state-dependent models are often used to determine if an impulse response is larger in one regime compared to another. One common example is whether the fiscal multiplier depends upon the degree of slack in an economy. In results presented in the appendix, we show that a model’s ability to capture state asymmetry is in line with its overall performance in terms of bias. The key exception is that the MA-GIRF captures state asymmetry better in all misspecification cases, even when the VAR’s lag order is truncated.

4 Application: Fiscal Multiplier

The self-exciting models considered above have a number of macroeconomic applications. Previous studies have used these models to estimate the the state-dependent effects of uncertainty and monetary and fiscal policy shocks, among others. For the government spending multiplier, in particular, results have varied in important ways. Some studies find that the multiplier is larger in periods of economic slack than in nonslack periods when crowding out is more costly. Other studies find that the multiplier does not vary across states of the busi-

¹²If one suspects the true model dynamics can only be captured by a VAR with a high lag order making the problem computationally difficult, our result suggests that LP would be a suitable alternative.

¹³This result might seem puzzling at first as the STVAR nests the TVAR for $\gamma \rightarrow \infty$. It may be that estimation of γ introduces enough uncertainty in the transition process that STVAR exhibits more bias. However, if one places substantial prior weight on the true model being linear, the STVAR remains a viable option as it nests the linear model for $\gamma \rightarrow 0$.

ness cycle. Figure 2 reviews the state-dependent multiplier estimates from a small fraction of the studies on fiscal multipliers.¹⁴

Results from our simulation study suggest that misspecification in either the VAR dynamics or the transition process can create substantial bias in the IRFs in nonlinear models. We consider whether these types of biases might explain the variation in the state-dependent fiscal multipliers found in the literature. The multiplier is typically computed from the state-dependent IRFs; thus, differences in assumptions about the underlying VAR model, the transition function, or the driving variable can produce variation in the multiplier. One proposed solution has been to compute the multipliers using LP, which has been claimed to be more robust to these types of misspecification. We also estimate the MA-GIRFs with a flat model prior, which we showed in the previous section are also robust to model uncertainty.

4.1 Data

The baseline sample, based on RZ, includes real GDP and real government spending and runs from 1890Q1-2015Q4. We also consider a truncated subsample from 1969Q1 to 2015Q4. Output and government spending are scaled by potential output, which is computed as a sixth-order polynomial trend in output [Gordon and Krenn (2010)]. All data are seasonally adjusted.

We consider two commonly-used methods to identify exogenous fiscal shocks: (i) the narrative shock to military spending identified by Ramey (2011) and (ii) the timing assumption from Blanchard and Perotti (2002). We calculate the impulse response of output and government spending to a military news shock or unexpected spending shock equal to 1 percent of GDP. The multiplier is calculated as the cumulative response of output divided by the cumulative response of government spending over a five-year period.

¹⁴See Gechert and Rannenberg (2018) for a recent survey on state-dependent fiscal multipliers. See Ramey (2019) and Castelnuovo and Lim (2019) for recent surveys on the fiscal multiplier in general.

4.2 Results

The top and bottom panels of Table 2 show the cumulative five-year state-dependent fiscal multiplier estimates for the various models when using the full and short samples, respectively. The first two columns show the multipliers using the Ramey news shock series for the slack and nonslack regimes, respectively. The next two columns show the corresponding results using Blanchard-Perotti (BP) shock identification. Each GIRF is estimated with four lags of Y_t and the output gap as the transition variable. LP is estimated at each horizon, conditional on four lags of Y_t , and requires a definition of the slack and nonslack states across time. We define the states based on whether the output gap is above (nonslack) or below (slack) its historical average.¹⁵

The results in top panel of Table 2 largely match those found by Ramey and Zubairy (2018), Barnichon et al. (2021), and Alloza (2022): (i) multipliers tend to be lower when using the BP identification compared to the Ramey news shocks; (ii) point estimates for the state-dependent multipliers are both below unity; and (iii) significant differences in multipliers across states can be found using the BP identification. Statistically relevant differences across states are only found when using the BP identification with LP impulse responses.

The bottom panel of Table 2 shows the results for the shorter sample, which differs in a number of important dimensions from the full sample results. First, the error bands are substantially larger than in the longer sample. This result is consistent with the results of Ramey (2011) and Ramey and Zubairy (2018) and likely reflects both increased parameter uncertainty and weaker identification of the Ramey news shocks when excluding the Korean War. Second, in the shorter sample, the estimated multipliers occasionally have point values larger than 1. Third, when using LP, the cumulative multipliers implied by the Ramey shocks are now negative implying a government spending shock contracts the economy. This result matches the negative post-WW2 multipliers of Ramey (2011) and contrasts with the

¹⁵The results are robust to other choices for the output gap threshold, such as 0.

relatively high multipliers from [Fisher and Peters \(2010\)](#).¹⁶ Fourth, as in the full sample, using the BP identification with LP produces multipliers that are statistically higher in slack compared to nonslack states.

The bottom row of each panel shows the multipliers computed using a weighted average across models and identification schemes. These combined MA-GIRF multipliers are slightly larger in the small sample versus the full sample. However, both state-dependent multipliers are below unity with no distinguishable difference across slack and nonslack states in both samples. Similar to the MA multipliers computed for each identification scheme, the combined MA multipliers have wider error bands when using the shorter sample.

Figure 3 shows the state-dependent MA-GIRFs for output and government spending using the Ramey news shock the full sample. The top row shows the responses of government spending and output in slack states while the bottom row shows the corresponding responses in nonslack states. Despite variation across the nonlinear VAR estimates, the state-dependent MA multiplier estimates are relatively close to the LP estimates when using the Ramey news shocks. BP identification finds a similar multiplier during nonslack states when using either the MA or LP. However, the MA finds a substantially lower multiplier of 0.46 for slack states compared to the LP estimate of 0.79.

4.3 Robustness

We explore two variations of the baseline specification. First, our simulation results suggest that truncating the VAR lags can induce significant bias in the impulse responses and, subsequently, in the multiplier. Second, much of the heterogeneity in the fiscal multipliers has been attributed to the use of different transition variables. For example, [Auerbach and Gorodnichenko \(2012\)](#) use a seven-quarter moving average of the output gap; [RZ](#) use the

¹⁶One potential explanation put forth by [Ramey \(2011\)](#) is that the military news shocks are temporary, while other identifications [e.g., [Fisher and Peters \(2010\)](#) shocks to defense stocks' excess returns] are more permanent. Also, a number of papers note that the Ramey news shock series is a poor instrument for government spending shocks after the Korean War, especially compared to shocks identified using Blanchard-Perotti [see [Ramey \(2016\)](#), [Kang and Kim \(2022\)](#), and [Jørgensen and Ravn \(2022\)](#), among others].

unemployment rate, NBER recession dates, and the output gap relative to fixed thresholds; Fazzari et al. (2015) uses an adjusted measure of capacity utilization; and Fazzari et al. (2021) uses the output gap. We check the sensitivity of the multiplier estimates to each of these model specifications. Finally, we estimate a “comprehensive multiplier” based on the MA-GIRF over all of the specifications considered.

4.3.1 VAR Lags

The first six columns of Table 3 and Table 4 examine the sensitivity of the estimated multipliers when varying the number of VAR lags. Table 3 considers the full sample whereas Table 4 uses the short sample discussed above. Because of the aforementioned identification concerns with the military news variable when excluding the Korean War, we present only the BP identification. Results for the Ramey news variable are available in the Appendix.

For the long sample, there is variation in the estimated values of the multiplier but they generally have magnitudes less than one for all lag orders. These results are consistent regardless of the identification of the fiscal shock. In the short sample estimated with more controls, we find more instances of multipliers greater than one for models with VARs with higher lag order.

While there is some variation in the value of the multiplier for different VAR lag orders, we may be able to resolve this uncertainty by examining the MA-GIRFs computed across the various models. This is because the MA-GIRFs weight the GIRFs from any given model by its relative fit. The bottom rows of each table show the MA-GIRFs for the appropriate time sample. In each case, we find that the multiplier in times of slack is slightly larger than in normal times but not significantly so. Moreover, in neither sample are the MA-GIRF multipliers larger than 1.

4.3.2 Transition Variable

Our baseline results use the output gap as the transition variable to differentiate periods of slack and nonslack. In this subsection, we check for robustness using the unemployment rate as alternative transition variable and a baseline VAR(4). The last two columns of Table 3 and Table 4 contain these results for the full sample and the short sample, respectively. The most direct comparison is the third and fourth columns that use the same baseline VAR(4) but with the output gap. Here, again, we report results using the BP identification; results obtained using the Ramey news shocks are in the Appendix.

For the full sample, the estimated multipliers are less than one and are only significantly different across business cycle regimes for LP. For the shorter sample, using the unemployment rate does appear to affect the point estimates of the state-dependent multipliers. For the MSVAR, the multipliers are larger; for the STVAR, the multipliers are smaller. For LP, the point values of the multipliers are nearly identical.

How do we resolve the apparent differences across the transition variables? We compute the MA-GIRFs including all models with both transition variables. Here, we find results more consistent with our baseline results. For the full sample, the multipliers are small (on the order of 0.4) and not statistically significant across regimes. For the shorter sample, the multipliers are relatively larger than in the full sample (around 0.8 in times of slack and 0.4 in normal times) and have larger error bands.

4.3.3 Comprehensive MA Multiplier

Given the copious model and identification uncertainty in the literature, we compute a comprehensive model average multiplier derived from variations of the four model specifications discussed above.¹⁷ We estimate models varying (i) the transition process (TVAR, MSVAR,

¹⁷As we noted above, the weights in the MA are applied directly to the IRFs. Because the object of interest is the multiplier, we could apply the weights directly to the multiplier. The appendix presents the estimated multipliers when using this approach and they are quantitatively the same as when averaging over the underlying IRFs.

or STVAR), (ii) the shock identification (Ramey News or Blanchard-Perotti), (ii) the VAR lag length (one, four, or six), and (iv) the transition variable (output gap or unemployment rate).

We estimate the state-dependent multipliers averaging over combinations of the four variations above when using, alternately, the full and short samples.¹⁸ For the full sample, the multiplier in the slack and no-slack regimes are 0.64 and 0.54, respectively. In both regimes, the multipliers are statistically below unity. Consistent with the narrative espoused by some papers in the literature, the point estimates of the multiplier in the slack regime is larger but not statistically different from the non-slack regime. In the shorter sample, the multipliers are closer to zero (0.40 in slack and 0.26 in non-slack), still seemingly below unity, and, again, not importantly different from each other. The comprehensive multipliers across self-exciting VARs reinforce the findings of [Ramey and Zubairy \(2018\)](#).

In summary, the comprehensive multiplier finds no significant difference across states whereas LP finds significant differences when using BP identification. Our simulation results suggest the LP estimates will be more biased as long as a sufficiently long lag length is used in the VAR specifications. Together with the findings of [Gonçalves et al. \(2023\)](#), care should be taken when considering estimates from state-dependent LP.

5 Conclusion

We compare the accuracy of GIRFs and RD-LP in self-exciting nonlinear models in the presence of model uncertainty and other possible misspecifications. In particular, we evaluate, using MC simulation experiments, how these misspecifications affect each method’s ability to approximate the true IRF and capture state asymmetry. In addition, we introduce a MA-GIRF that weights the IRFs of based on the BIC.

We arrive at a few main conclusions based on our MC experiments. If one has a strong model prior about the true form of the underlying state process, estimating that model is

¹⁸Error bands for the comprehensive multipliers are presented in the online appendix.

typically first best. Any significant uncertainty about model specification, however, can be ameliorated by estimating an MA-GIRF over a model set that includes VARs with sufficient lag order to capture the relevant dynamics. In the case that the appropriate VAR is suspected to have an inestimable number of parameters, RD-LP is a viable alternative. Our results complement those in [Gonçalves et al. \(2023\)](#), who explore the effect of variation in the shock size in similar models. Methods used here can be augmented with those in [Gonçalves et al. \(2023\)](#), providing a comprehensive approach to modeling IRFs in the presence of nonlinearities and model uncertainty.

Based on these findings, we address the problem of state-dependency in the fiscal spending multiplier. We estimate a variety of nonlinear VARs to obtain GIRFs and the MA-GIRF to compare with RD-LP responses. We find that fiscal multipliers are consistently below unity, irrespective of the shock, the transition variable, or number of VAR lags used in estimation, when estimated over a long sample dating back to 1890. However, the state-dependent multipliers are sensitive to the use of only post-Korean war data. Therefore, studies, using postwar data, which find significant differences in fiscal multipliers across business cycle regimes will likely not be robust to model misspecification or alternative shock identifications.

References

- Alloza, Mario (2022) ‘Is fiscal policy more effective during recessions?’ *International Economic Review* 63(3), 1271–1292
- Auerbach, Alan J., and Yuriy Gorodnichenko (2012) ‘Measuring the output responses to fiscal policy.’ *American Economic Journal: Economic Policy* 4(2), 1–27
- Auerbach, Alan J., and Yuriy Gorodnichenko (2013) ‘Fiscal multipliers in recession and expansion.’ In ‘Fiscal policy after the financial crisis’ (University of Chicago Press) pp. 63–98
- Barnichon, Regis, Davide Debortoli, and Christian Matthes (2021) ‘Understanding the Size of the Government Spending Multiplier: It’s in the Sign.’ *The Review of Economic Studies* 89(1), 87–117
- Billio, Monica, Roberto Casarin, Francesco Ravazzolo, and Herman K. Van Dijk (2016) ‘Interconnections between eurozone and us booms and busts using a bayesian panel markov-switching var model.’ *Journal of Applied Econometrics* 31(7), 1352–1370
- Blanchard, Olivier, and Roberto Perotti (2002) ‘An empirical characterization of the dynamic effects of changes in government spending and taxes on output.’ *The Quarterly Journal of Economics* 117(4), 1329–1368
- Caggiano, Giovanni, Efrem Castelnuovo, Valentina Colombo, and Gabriela Nodari (2015) ‘Estimating Fiscal Multipliers: News from a Non-linear World.’ *The Economic Journal* 125(584), 746–776
- Castelnuovo, Efrem, and Guay Lim (2019) ‘What do we know about the macroeconomic effects of fiscal policy? a brief survey of the literature on fiscal multipliers.’ *Australian Economic Review* 52(1), 78–93

- Diebold, Francis, Joon-Haeng Lee, and Gretchen Weinbach (1993) ‘Regime switching with time-varying transition probabilities.’ Working Papers 93-12, Federal Reserve Bank of Philadelphia
- Fazzari, Steven M., James Morley, and Irina Panovska (2015) ‘State-dependent effects of fiscal policy.’ *Studies in Nonlinear Dynamics & Econometrics* 19(3), 285–315
- (2021) ‘When is discretionary fiscal policy effective?’ *Studies in Nonlinear Dynamics & Econometrics* 25(4), 229–254
- Filardo, Andrew J. (1994) ‘Business-cycle phases and their transitional dynamics.’ *Journal of Business & Economic Statistics* 12(3), 299–308
- Fisher, Jonas D.M., and Ryan Peters (2010) ‘Using stock returns to identify government spending shocks.’ *The Economic Journal* 120(544), 414–436
- Gechert, Sebastian, and Ansgar Rannenberg (2018) ‘Which fiscal multipliers are regime-dependent? a meta-regression analysis.’ *Journal of Economic Surveys* 32(4), 1160–1182
- Geweke, John, and Gianni Amisano (2011) ‘Optimal prediction pools.’ *Journal of Econometrics* 164(1), 130–141
- Gonçalves, Sílvia, Ana María Herrera, Lutz Kilian, and Elena Pesavento (2023) ‘State-dependent local projections.’ Working Paper 2302, Federal Reserve Bank of Dallas
- Gordon, Robert J, and Robert Krenn (2010) ‘The end of the great depression 1939-41: Policy contributions and fiscal multipliers.’ Working Paper 16380, National Bureau of Economic Research, September
- Granger, Clive, and Timo Teräsvirta (1993) *Modelling Non-Linear Economic Relationships* (Oxford University Press)
- Hansen, Bruce E. (2007) ‘Least squares model averaging.’ *Econometrica* 75(4), 1175–1189

- Hansen, Bruce E (2016) ‘Stein combination shrinkage for vector autoregressions.’ *Manuscript, University of Wisconsin-Madison*
- Ho, Paul, Thomas A. Lubik, and Christian Matthes (2023) ‘Averaging impulse responses using prediction pools.’ *Federal Reserve Bank of Richmond Working Paper 23-04*
- Inoue, Atsushi, Barbara Rossi, and Yiru Wang (2022) ‘Local projections in unstable environments: How effective is fiscal policy?’
- Jackson, Laura E., Michael T. Owyang, and Daniel Soques (2018) ‘Nonlinearities, smoothing and countercyclical monetary policy.’ *Journal of Economic Dynamics and Control* 95, 136–154
- Jordà, Òscar (2005) ‘Estimation and inference of impulse responses by local projections.’ *American economic review* 95(1), 161–182
- Jordà, Òscar (2005) ‘Estimation and inference of impulse responses by local projections.’ *American Economic Review* 95(1), 161–182
- Jørgensen, Peter L., and Søren H. Ravn (2022) ‘The inflation response to government spending shocks: A fiscal price puzzle?’ *European Economic Review* 141, 103982
- Kang, Jihye, and Soyoung Kim (2022) ‘Government spending news and surprise shocks: It’s the timing and persistence.’ *Journal of Macroeconomics* 73, 103446
- Kilian, Lutz, and Yun Jung Kim (2011) ‘How Reliable Are Local Projection Estimators of Impulse Responses?’ *The Review of Economics and Statistics* 93(4), 1460–1466
- Koop, Gary, M. Hashem Pesaran, and Simon M. Potter (1996) ‘Impulse response analysis in nonlinear multivariate models.’ *Journal of Econometrics* 74(1), 119–147
- Krolzig, Hans-Martin (2006) ‘Impulse-response analysis in markov switching vector autoregressive models.’ Mimeo, University of Kent

- Krüger, Fabian, Sebastian Lerch, Thordis L. Thorarinsdottir, and Tilmann Gneiting (2016) ‘Predictive inference based on markov chain monte carlo output’
- Laumer, Sebastian, and Collin S. Phillips (2020) ‘Does the government spending multiplier depend on the business cycle?’ *Journal of Money, Credit, and Banking*
- Li, Dake, Mikkel Plagborg-Møller, and Christian K. Wolf (2021) ‘Local projections vs. vars: Lessons from thousands of dgps’
- Lusompa, Amaze (2021) ‘Local projections, autocorrelation, and efficiency.’ *Federal Reserve Bank of Kansas City Working Paper*
- Matheson, James E., and Robert L. Winkler (1976) ‘Scoring rules for continuous probability distributions.’ *Management Science* 22(10), 1087–1096
- Montiel Olea, José Luis, and Mikkel Plagborg-Møller (2021) ‘Local projection inference is simpler and more robust than you think.’ *Econometrica* 89(4), 1789–1823
- Owyang, Michael T., Valerie A. Ramey, and Sarah Zubairy (2013) ‘Are government spending multipliers greater during periods of slack? evidence from twentieth-century historical data.’ *American Economic Review* 103(3), 129–34
- Plagborg-Møller, Mikkel, and Christian K Wolf (2021) ‘Local projections and vars estimate the same impulse responses.’ *Econometrica* 89(2), 955–980
- Raftery, Adrian E. (1995) ‘Bayesian model selection in social research.’ *Sociological Methodology* 25, 111–163
- Ramey, V.A. (2016) ‘Chapter 2 - macroeconomic shocks and their propagation.’ vol. 2 of *Handbook of Macroeconomics* (Elsevier) pp. 71–162
- Ramey, Valerie A. (2011) ‘Identifying Government Spending Shocks: It’s all in the Timing*.’ *The Quarterly Journal of Economics* 126(1), 1–50

- (2019) ‘Ten years after the financial crisis: What have we learned from the renaissance in fiscal research?’ *Journal of Economic Perspectives* 33(2), 89–114
- Ramey, Valerie A., and Sarah Zubairy (2018) ‘Government spending multipliers in good times and in bad: Evidence from us historical data.’ *Journal of Political Economy* 126(2), 850–901
- Ruisi, Germano (2019) ‘Time-varying local projections.’ Working Paper 891, London
- Steel, Mark F. J. (2020) ‘Model averaging and its use in economics.’ *Journal of Economic Literature* 58(3), 644–719
- Swanson, Norman R., and Tian Zeng (2001) ‘Choosing among competing econometric forecasts: Regression-based forecast combination using model selection.’ *Journal of Forecasting* 20(6), 425–440

Table 1: **CRPS Across Simulation Specifications**

This table shows the median Note: This table shows the mean CRPS for each estimation method for a given true data generating process and specification error. We compute the overall CRPS as the average over 20 impulse response horizons. The 68% posterior interval is shown in brackets beneath the median.

Specification Error:	True Model:	Estimated Model:				
		<i>T</i> VAR	<i>MS</i> VAR	<i>ST</i> VAR	<i>LP</i>	<i>MA</i>
<i>Transition Process Only</i>	TVAR	0.20 [0.09,0.55]	0.42 [0.17,0.75]	0.66 [0.40,1.03]	0.66 [0.33,1.15]	0.24 [0.12,0.56]
	MSVAR	0.16 [0.09,0.35]	0.12 [0.06,0.31]	0.21 [0.12,0.41]	0.20 [0.13,0.37]	0.14 [0.08,0.30]
	STVAR	0.12 [0.06,0.27]	0.13 [0.06,0.28]	0.14 [0.08,0.28]	0.25 [0.15,0.55]	0.12 [0.06,0.25]
<i>Incorrect Lag on Transition Variable</i>	TVAR	0.38 [0.20,0.80]	0.49 [0.22,1.31]	0.71 [0.39,1.45]	0.53 [0.30,1.06]	0.44 [0.22,1.00]
	MSVAR	0.15 [0.09,0.29]	0.12 [0.06,0.23]	0.16 [0.11,0.28]	0.20 [0.13,0.33]	0.12 [0.07,0.22]
	STVAR	0.13 [0.07,0.23]	0.10 [0.05,0.22]	0.13 [0.07,0.24]	0.28 [0.14,0.59]	0.11 [0.06,0.21]
<i>Incorrect Transition Variable</i>	TVAR	0.42 [0.23,0.82]	0.49 [0.23,1.28]	0.77 [0.42,1.40]	0.57 [0.33,1.09]	0.47 [0.24,1.05]
	MSVAR	0.17 [0.10,0.33]	0.14 [0.07,0.28]	0.21 [0.13,0.35]	0.20 [0.14,0.34]	0.14 [0.08,0.28]
	STVAR	0.18 [0.09,0.34]	0.11 [0.06,0.27]	0.13 [0.08,0.34]	0.31 [0.16,0.54]	0.12 [0.06,0.28]
<i>Incorrect VAR Lag (1)</i>	TVAR	1.10 [0.73,1.91]	1.25 [0.84,2.08]	1.51 [0.93,2.34]	0.60 [0.36,1.18]	1.12 [0.76,1.95]
	MSVAR	0.63 [0.51,0.75]	0.63 [0.53,0.78]	0.65 [0.55,0.81]	0.21 [0.14,0.33]	0.60 [0.49,0.71]
	STVAR	0.25 [0.08,0.49]	0.26 [0.09,0.49]	0.23 [0.09,0.51]	0.34 [0.16,0.58]	0.25 [0.08,0.49]
<i>Incorrect VAR Lag (6)</i>	TVAR	0.25 [0.10,0.69]	0.45 [0.16,1.18]	0.69 [0.38,1.41]	0.54 [0.31,1.07]	0.31 [0.13,0.72]
	MSVAR	0.28 [0.16,0.54]	0.22 [0.14,0.42]	0.42 [0.28,0.60]	1.31 [0.23,4.23]	0.26 [0.15,0.45]
	STVAR	0.16 [0.08,0.32]	0.17 [0.09,0.33]	0.14 [0.09,0.29]	0.35 [0.18,0.67]	0.15 [0.08,0.28]
<i>True Model is Linear</i>	VAR	1.67 [1.14,2.49]	1.36 [0.78,2.09]	1.17 [0.56,1.84]	1.54 [0.85,2.05]	1.49 [0.96,2.21]

Table 2: **Fiscal Spending Multipliers**

This table shows the median posterior five-year fiscal multiplier for each model under each business cycle regime. The 68% highest posterior density interval is shown below in brackets. The multiplier is computed as the cumulative response of output divided by the cumulative response of government spending. The first and second columns show the multiplier based on the Ramey news shock series under slack and nonslack regimes, respectively. The third and fourth columns show similar estimates when using Blanchard-Perotti shock identification. In each case the shock size is a one-percent of GDP increase in government spending. Four lags of Y_t are used in both the VAR and LP specifications. The last row shows the multiplier implied by the model average impulse response when averaging over all three model specifications (TVAR, MSVAR, and STVAR) as well as both shocks (Ramey News and Blanchard-Perotti).

(a) Full Sample

Sample:	1890Q1-2015Q4			
Shock ID:	Ramey News		Blanchard-Perotti	
	Slack	Nonslack	Slack	Nonslack
TVAR	0.49 [0.15,0.87]	0.59 [0.37,0.83]	0.42 [0.13,0.68]	0.34 [0.18,0.51]
MSVAR	0.93 [0.57,1.34]	0.67 [0.33,1.00]	0.59 [0.45,0.73]	0.32 [0.13,0.50]
STVAR	0.98 [0.45,1.54]	0.90 [0.35,1.41]	0.49 [0.35,0.66]	0.29 [0.20,0.40]
Model Avg	0.53 [0.17,0.92]	0.59 [0.38,0.83]	0.46 [0.23,0.65]	0.31 [0.18,0.46]
Local Proj	0.53 [0.46,0.61]	0.68 [0.55,0.81]	0.79 [0.72,0.87]	0.31 [0.24,0.37]
		Slack	Nonslack	
Model Avg (Both Shock IDs)		0.46 [0.24,0.66]	0.33 [0.20,0.50]	

(b) Short Sample

Sample:	1969Q1-2015Q4			
Shock ID:	Ramey News		Blanchard-Perotti	
	Slack	Nonslack	Slack	Nonslack
TVAR	1.49 [-3.68,7.44]	1.88 [-7.94,11.37]	0.83 [-0.05,1.71]	0.02 [-1.19,1.12]
MSVAR	0.67 [-8.33,11.00]	0.97 [-5.62,7.74]	0.95 [-0.11,2.07]	1.06 [-0.01,2.20]
STVAR	0.97 [-3.93,5.72]	2.26 [-11.25,14.68]	1.53 [0.68,2.45]	0.33 [-0.55,1.11]
Model Avg	1.27 [-4.24,7.37]	1.65 [-8.39,11.63]	0.84 [0.02,1.71]	0.30 [-0.85,1.37]
Local Proj	-0.24 [-0.95,0.48]	-2.00 [-8.15,2.68]	0.56 [0.34,0.78]	-2.59 [-4.04,-1.49]
		Slack	Nonslack	
Model Avg (Both Shock IDs)		0.83 [-0.13,1.80]	0.27 [-1.09,1.50]	

Table 3: **Fiscal Spending Multipliers – Robustness Checks for Full Sample**

This table shows the median posterior five-year fiscal multiplier for each model under each business cycle regime when using the full sample. The first six columns show the multiplier estimates across different VAR lags when using the output gap as the transition variable. The last two columns show similar estimates when using four VAR lags and the unemployment rate as the transition variable. The last two rows of each panel show the multiplier implied by the model average impulse response when averaging over all three model specifications (TVAR, MSVAR, and STVAR) as well as either the VAR lags (second to last row) or the two transition variables (last row). The 68% highest posterior density interval is shown below in brackets.

Sample:		1890Q1-2015Q4							
Shock ID:		Blanchard-Perotti							
Transition Var:	Output Gap	Output Gap		Output Gap		UR			
VAR Lags	1	4		6		4			
	Slack	Nonslack	Slack	Nonslack	Slack	Nonslack	Slack	Nonslack	
TVAR	0.47 [0.28,0.88]	0.44 [0.34,0.55]	0.42 [0.13,0.68]	0.34 [0.18,0.51]	0.58 [0.34,0.79]	0.47 [0.33,0.61]	0.65 [0.40,0.89]	0.44 [0.28,0.60]	
MSVAR	0.01 [-0.37,0.36]	0.51 [0.36,0.65]	0.59 [0.45,0.73]	0.32 [0.13,0.50]	0.54 [0.32,0.74]	0.56 [0.38,0.73]	0.27 [0.08,0.44]	-0.02 [-0.39,0.25]	
STVAR	0.40 [0.30,0.63]	0.28 [0.23,0.33]	0.49 [0.35,0.66]	0.29 [0.20,0.40]	0.72 [0.49,1.08]	0.44 [0.33,0.55]	0.67 [0.57,0.78]	0.49 [0.40,0.58]	
Model Avg	0.42 [0.28,0.75]	0.39 [0.28,0.52]	0.46 [0.23,0.65]	0.31 [0.18,0.46]	0.55 [0.37,0.74]	0.48 [0.34,0.63]	0.65 [0.40,0.89]	0.44 [0.28,0.60]	
Local Proj	0.44 [0.38,0.50]	0.69 [0.64,0.74]	0.79 [0.72,0.87]	0.31 [0.24,0.38]	0.28 [0.19,0.36]	0.52 [0.45,0.58]	0.69 [0.61,0.77]	0.25 [0.18,0.32]	
			Slack	Nonslack					
Model Avg (All VAR Lags)			0.43 [0.30,0.65]	0.37 [0.27,0.49]					
Model Avg (Both Transition Variables)			0.45 [0.24,0.64]	0.33 [0.19,0.50]					

Table 4: **Fiscal Spending Multipliers – Robustness Checks for Short Sample**

This table shows the median posterior five-year fiscal multiplier for each model under each business cycle regime when using the short sample. The first six columns show the multiplier estimates across different VAR lags when using the output gap as the transition variable. The last two columns show similar estimates when using four VAR lags and the unemployment rate as the transition variable. The last two rows of each panel show the multiplier implied by the model average impulse response when averaging over all three model specifications (TVAR, MSVAR, and STVAR) as well as either the VAR lags (second to last row) or the two transition variables (last row). The 68% highest posterior density interval is shown below in brackets.

Sample:	1969Q1-2015Q4							
Shock ID:	Blanchard-Perotti							
Transition Var:	Output Gap		Output Gap		Output Gap		UR	
VAR Lags	1		4		6		4	
	Slack	Nonslack	Slack	Nonslack	Slack	Nonslack	Slack	Nonslack
TVAR	0.52 [0.10,0.95]	0.30 [-0.12,0.71]	0.83 [-0.05,1.71]	0.02 [-1.19,1.12]	0.83 [-0.16,1.87]	0.63 [-0.47,1.92]	0.33 [-0.53,1.09]	0.44 [-0.68,1.47]
MSVAR	0.18 [-0.45,0.79]	0.28 [-0.33,0.88]	0.95 [-0.11,2.07]	1.06 [-0.01,2.20]	0.08 [-1.03,1.09]	0.10 [-1.01,1.09]	1.49 [0.77,2.18]	1.70 [0.94,2.40]
STVAR	0.70 [-0.03,1.37]	-0.29 [-0.92,0.22]	1.53 [0.68,2.45]	0.33 [-0.55,1.11]	1.41 [0.36,2.43]	0.27 [-0.83,1.37]	0.06 [-0.83,0.85]	-0.18 [-1.19,0.69]
Model Avg	0.36 [-0.13,0.79]	0.26 [-0.16,0.67]	0.84 [0.02,1.71]	0.30 [-0.85,1.37]	0.64 [-0.34,1.63]	0.49 [-0.56,1.62]	0.80 [-0.14,1.69]	1.00 [-0.20,1.99]
Local Proj	0.41 [0.22,0.62]	-1.60 [-2.39,-0.92]	0.56 [0.33,0.79]	-2.56 [-4.07,-1.49]	0.53 [0.28,0.77]	-1.60 [-2.96,-0.52]	0.58 [0.37,0.79]	-2.02 [-3.34,-1.04]
			Slack	Nonslack				
Model Avg (All VAR Lags)			0.40 [-0.09,0.83]	0.26 [-0.18,0.69]				
Model Avg (Both Transition Variables)			0.87 [0.04,1.74]	0.42 [-0.68,1.43]				

Figure 1: **CRPS By Horizon**

This figure shows the median CRPS for each impulse response horizon for each estimation method. The data generating process for the top, middle, and bottom panels are a TVAR, MSVAR, and STVAR, respectively.

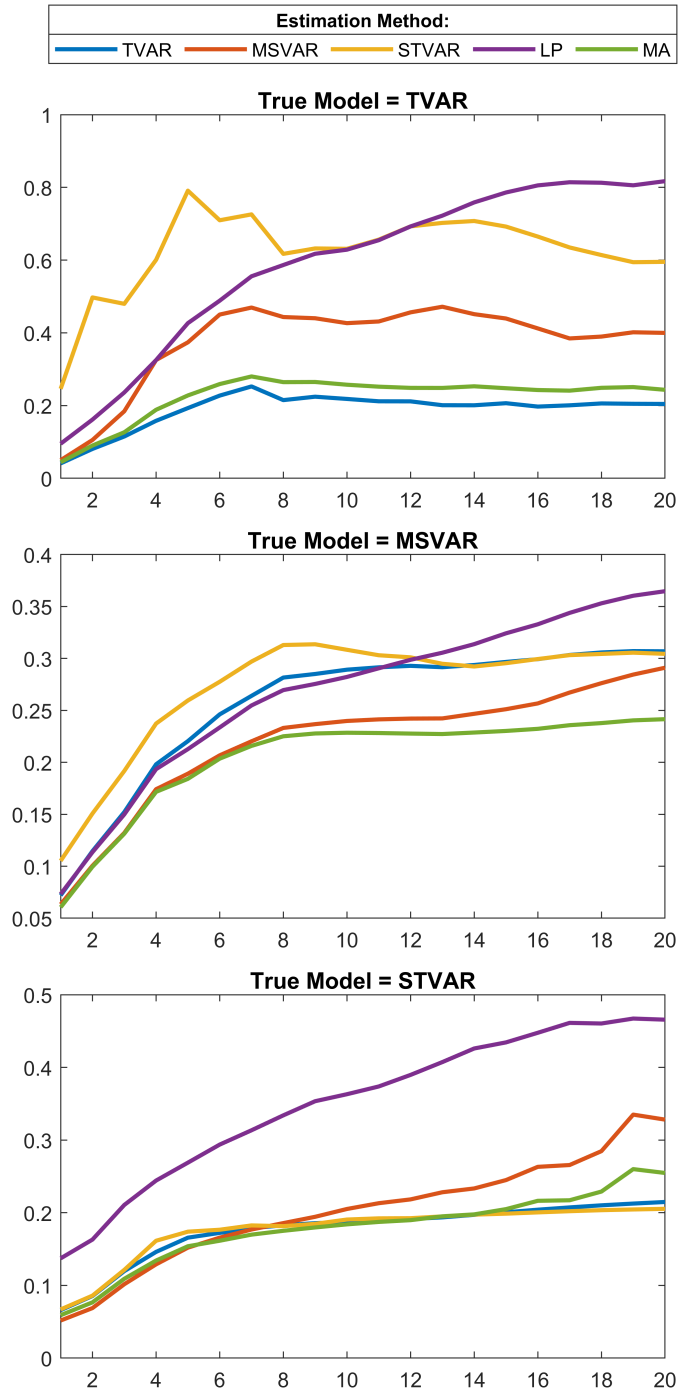


Figure 2: **Histogram of Reported State-Dependent Fiscal Multipliers.**

This histogram shows the reported regime-dependent multipliers from the meta-analysis conducted by [Gechert and Rannenberg \(2018\)](#) which covers 98 previous studies. The blue bars show the estimated fiscal multipliers in normal economic regimes and the red bars show the multipliers in slack regimes. The dashed vertical lines show the average multiplier in each regime across the previous studies.

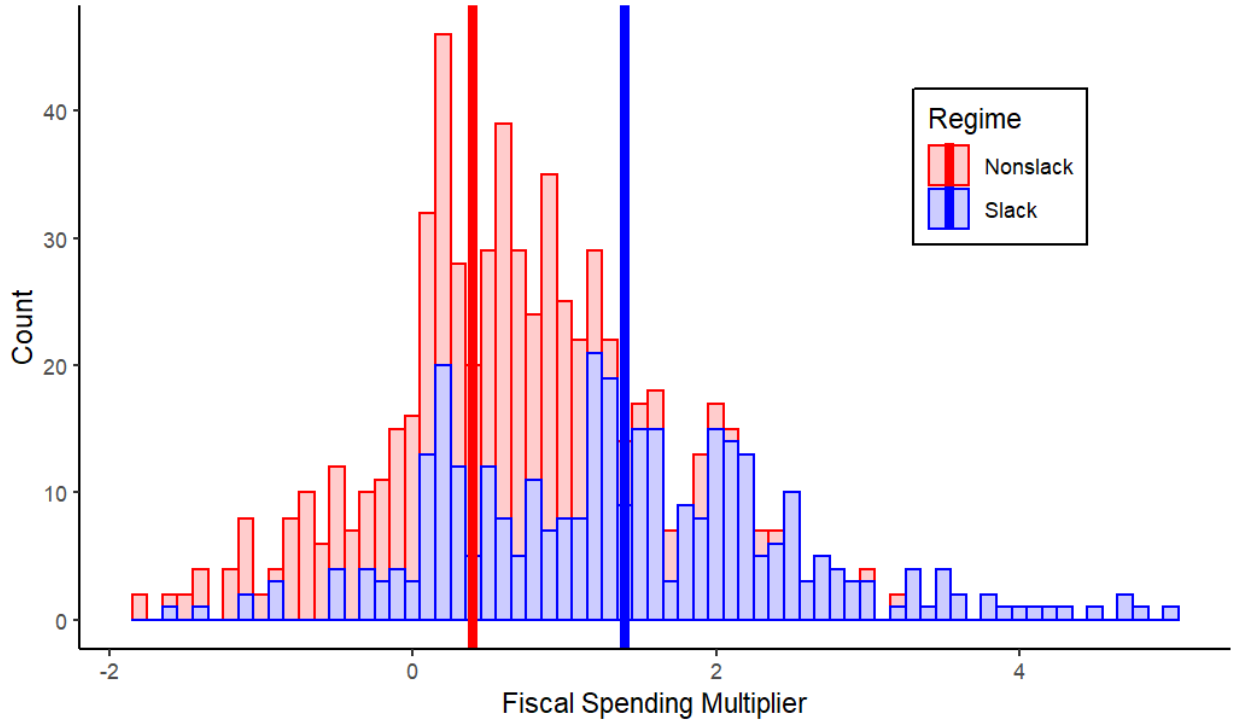


Figure 3: Model Average GIRFs

This figure shows the model average impulse responses for government spending and output when using the Ramey news shock series for the full sample (1890Q1-2015Q4). For each variable, the model average weighs the GIRF from the TVAR, MSVAR, and STVAR based on their respective BIC. Four lags of Y_t are used in each VAR specification.

