State-Dependent Local Projections: Understanding Impulse Response Heterogeneity*

James S. Cloyne† Òscar Jordà‡ Alan M. Taylor§

August 2023

Abstract

Much attention has been paid to estimates of how fiscal policy affects the macroeconomy, but the standard impulse response is the dynamic average effect of an intervention across horizons. We use the well-known Kitagawa-Blinder-Oaxaca decomposition to explore a response's heterogeneity over time and over states of the economy. This can be implemented with a simple extension to the usual local projection specification that nevertheless keeps the model linear in parameters. Using our new decomposition-based approach, we show how to unpack heterogeneity in the fiscal multiplier, an object that at any point in time may depend on a number of potentially correlated factors, including existing economic conditions and the monetary response. In our application, the fiscal multiplier varies considerably with monetary policy: it can be as small as zero, or as large as 2, depending on the degree of monetary offset.

JEL classification codes: C54, C99, E32, E62, H20, H5.

Keywords: fiscal multiplier, monetary offset, Kitagawa-Blinder-Oaxaca decomposition, local projections, interest rates, fiscal policy, state-dependence, balance, identification.

*We thank Helen Irvin and Chitra Marti for excellent research assistance. We are also grateful to Alan Auerbach, Olivier Blanchard, Roger Farmer, Emi Nakamura, Gernot Müller and participants at the NBER Summer Institute 2020, the Econometric Society World Congress 2020, the Society for Economic Dynamics Annual Conference 2021 and the WEAI Annual Conference 2021 for very helpful comments and suggestions. The views expressed in this paper are the sole responsibility of the authors and do not necessarily reflect the views of the Federal Reserve Bank of San Francisco or the Federal Reserve System.

†Department of Economics, University of California, Davis; NBER; and CEPR (jcloyne@ucdavis.edu).

‡Federal Reserve Bank of San Francisco; Department of Economics, University of California, Davis; and CEPR (oscar.jorda@sf.frb.org; ojorda@ucdavis.edu).

§Department of Economics and Graduate School of Management, University of California, Davis; NBER; and CEPR (amtaylor@ucdavis.edu).

1. Introduction

This paper introduces a new framework to decompose impulse responses. It allows us to explicitly unpack the heterogeneity generated by the interactions of policy treatments and state variables, which implicitly lurk behind the average effects often estimated in the literature. We see the approach having broad application to a wide range of applied macroeconomic and policy problems.

Many in the literature have previously explored interactions of policy treatments and state variables (see, e.g., Cogley and Sargent, 2005; Primiceri, 2005; Tenreyro and Thwaites, 2016; Angrist, Jordà, and Kuersteiner, 2018; Ramey and Zubairy, 2018, to cite a few of the many studies we cite below). So what is new in our paper? Our contribution is to provide a formal and highly tractable framework that encompasses previous studies as special cases. Along the way we will learn several things, the most notable takeaways being: how an agnostic specification requires interactions with all controls to avoid bias; how identification is required for state variables; and how to estimate interpretable responses that vary over time depending on observable variation in the controls.

We first show that the local projection (LP) approach in Jordà (2005) can be easily extended to carry out the well-known Kitagawa-Blinder-Oaxaca decomposition (Kitagawa, 1955; Blinder, 1973; Oaxaca, 1973). Henceforth we refer to the latter as the *KBO decomposition* for brevity. This decomposition is standard in applied microeconomics (see Fortin, Lemieux, and Firpo, 2011, for a review), but has not found equivalent acceptance in applied macroeconomics. The other main takeaway of this paper is that we argue that it should.

The KBO decomposition of an LP response allows us to evaluate three separate effects following an exogenous policy intervention. First, the *direct* effect of an intervention on outcomes: This effect embeds the typical dynamic response of all policies and state variables in the sample and is an "average effect" of the intervention, the analog to the *average treatment effect*. It is what a standard LP estimates. Second, and most important for us, the *indirect* effect: Policy interventions can themselves modify how other variables influence the outcomes. As an example, consider monetary-fiscal interactions, as we do later: the effect of fiscal treatment is less effective if there is a monetary offset (e.g., a fiscal easing triggers a monetary tightening). Third, the *composition* effect: This piece showcases the bias due to imperfect identification and the importance of appropriate control for observables.

What is gained from the KBO decomposition? We see at least two main benefits. The first benefit concerns two neglected estimation biases which are resolved by the new approach. States that are jointly determined with the intervention are clearly not exogenous to the intervention. Thus one needs sources of exogenous variation for both the intervention *and* for any jointly determined state variable of interest. However, the identification requirement goes beyond jointly determined policy variables. Even states that are pre-determined with respect to the intervention (say, the output gap in the previous period) require a source of exogenous variation since pre-determined macroeconomic variables are likely to be correlated with each other (the output gap is probably correlated with inflation, the trade balance, and so on). As Fortin, Lemieux, and Firpo (2011) note, the KBO decomposition follows a partial equilibrium approach and, without further structural identifying assumptions, it is not necessarily correct to infer how much more or less effective a policy would be if one could also manipulate a jointly determined state variable. Furthermore, because heterogeneity can be generated by states other than those considered, it is important to include all interactions with the policy intervention to avoid potential omitted variable bias.

The second benefit concerns the simple estimation of contingent responses for policy analysis. By recognizing that all state variables can generate heterogeneity (even when exogenous variation in state variables is not available and thus a causal interpretation is not possible) we can still characterize the most likely response to the policy intervention conditional on the recent history observed in the data. Thus, with a simple extension of the typical LP, we are able to estimate traditional regressions, linear in parameters, yet generate variation in the impulse response for each period in the sample. Armed with that time-varying response, the policymaker can evaluate how effective a policy intervention is likely to be, conditional on current circumstances.[1] Thus, we show that the KBO decomposition provides a natural, well-founded, and tractable unifying framework for causally analyzing state and policy dependent impulse responses, as well as computing time varying impulse responses.

Application to the fiscal multiplier What is the fiscal multiplier? In principle the definition is clear: The multiplier tells us how many extra dollars of additional economic output are gained or lost by changing government expenditure or taxation (or a mix of the two) by one dollar. Given the turbulent economic events and dramatic policy actions of the last decade or so there continues to be much interest in empirical estimates of this object.

However, there is no such thing as *the* fiscal multiplier. For example, the multiplier may vary depending on the state of the economy and other policy interventions. One obvious reason is that monetary policy may not offset the effects of fiscal policy in the same way across states of the world, countries, or time. This insight, of course, exists in many macroeconomic theories and has been noted in policy debates.[2] But there is still relatively little evidence quantifying the importance of the "monetary offset" empirically. As a result, much policy advice has been given using multiplier estimates that are likely to depend on the particular average response of monetary policy in the past.

As an application of our new empirical approach, we therefore examine this interaction of monetary and fiscal policy. As motivation, Figure 1 shows how the fiscal multiplier varies in a standard version of the well-known New Keynesian model (described in the Appendix) depending on how aggressively the monetary authority responds to inflation. One goal of this paper is to

[1]This echoes Ramey (2019) who notes "to understand whether a particular estimate of fiscal effects is suitable for use in predicting the effects of a proposed policy, one must understand how the current circumstances differ from those present in the sample used to generate that estimate."

[2]See Woodford (2011) for an analysis of this point in the standard closed economy New Keynesian model and Leeper, Traum, and Walker (2017) in the context of a larger medium-scale DSGE model. Larger multipliers are also possible at the Zero Lower Bound when monetary policy is unable or unwilling to offset the effects of a fiscal stimulus (Christiano, Eichenbaum, and Rebelo (2011) and Eggertsson (2011)).

Figure 1: *Variation in the fiscal multiplier by monetary offset*

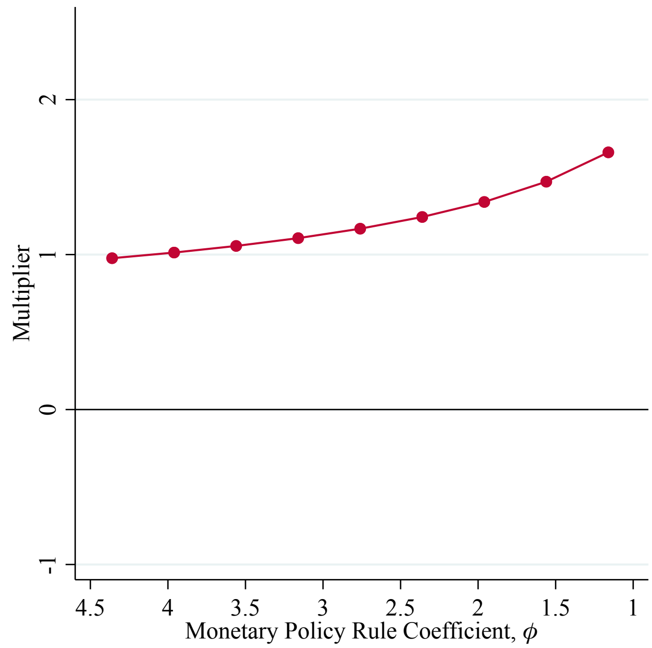

Notes: This chart shows how the cumulative fiscal multiplier (the cumulative effect on GDP scaled by the cumulative increase in government spending) at a horizon of 2 years varies with the monetary policy response in our New Keynesian model. The horizontal axis shows the inflation coefficient in model's monetary policy Taylor Rule, ϕ. This axis inverted so that moving from left to right on the horizontal axis implies a "less active" monetary policy which results in a larger multiplier.

show that we can recover a similar figure from the data using the empirical methods that we will introduce below.

For our analysis we use the dataset from an influential and established study from the IMF which identifies periods of exogenous fiscal treatment across 17 OECD countries from 1978 to 2009. This study by Guajardo, Leigh, and Pescatori (2014) employs the Romer and Romer (2010) "narrative" approach which uses historical information to isolate episodes of exogenous fiscal policy changes unrelated to current or prospective economic conditions. There are a few key reasons for using the Guajardo, Leigh, and Pescatori (2014) (GLP) dataset. First, our contribution is not a new identification of fiscal shocks. Rather, we take the existing GLP consolidation episodes off the shelf, and then show how the fiscal multiplier varies with monetary policy. Second, the cross-country coverage of the IMF study will allow us to exploit the panel nature of the dataset for identification of the monetary offset. Third, studying non-linear effects and state-dependence naturally asks more of the data and larger sample sizes are preferable, and the large GLP panel dataset is helpful in this respect.

Given that the monetary response is determined jointly with the fiscal treatment, we need to address the causality issues discussed above. In our application, we will exploit the fact that different countries may exhibit different degrees of accommodation on average.[3] This heterogeneity makes interest rates differentially sensitive to fiscal policy on average and generates *cross-sectional* variation that is useful for identification. This differential sensitivity allows us to construct a proxy for the monetary response that we can vary to undertake counterfactual policy experiments. Using this feature of the data, we show that fiscal interventions have very different effects on GDP depending on whether the intervention occurs in an environment with a more or less accommodative monetary response.[4] Exploiting the KBO decomposition, we can then quantify how the fiscal multiplier varies with the degree of monetary accommodation.

Using our new approach, we show that fiscal multipliers are around or below 1 on average, but there is considerable heterogeneity. Following a fiscal contraction—when the degree of monetary accommodation is limited—fiscal multipliers can become large. In our policy experiments, fiscal multipliers can be as low as zero, or as high as 2 depending on the monetary policy configuration. The high estimate (corresponding to a limited monetary offset) is similar to the original multiplier estimate of 2.5 posited by Keynes (1936) under the U.S. gold standard of the early 1930s.

Literature Naturally, this paper —and especially its focal application— is related to a wider literature estimating state dependent impulse responses, and to the large literature on the empirical fiscal multiplier. But translated into the language of the KBO decomposition, recent decades have seen intense research on the direct effect, but far less attention has been paid to the identification of the indirect effect.

In the fiscal policy literature early contributions, Blanchard and Perotti (2002) and Mountford and Uhlig (2009), sought to identify the average effect of fiscal interventions by imposing restrictions in a vector autoregression (VAR) framework. Numerous applications have followed. On the other hand, Romer and Romer (2010) pioneered the "narrative" approach (mentioned above) which essentially looks for historical natural experiments. A number of papers have applied or refined this method including Barro and Redlick (2011); Cloyne (2013); Mertens and Ravn (2013); Guajardo, Leigh, and Pescatori (2014); Hayo and Uhl (2014); Cloyne and Surico (2017); Gunter, Riera-Crichton, Végh, and Vuletin (2018); Nguyen, Onnis, and Rossi (2021); Hussain and Lin (2018); Cloyne, Dimsdale, and Postel-Vinay (2018). All these papers focus on identifying the *direct* effect in our KBO vernacular.

We obviously also connect to the recent literature on the state dependent effects of policy interventions. Prominent examples in the fiscal policy literature examine whether the impact of fiscal policy could vary depending on circumstances (Auerbach and Gorodnichenko, 2012; DeLong and Summers, 2012; Bachmann and Sims, 2012; Riera-Crichton, Vegh, and Vuletin, 2015; Jordà and Taylor, 2016; Ramey and Zubairy, 2018; Banerjee and Zampolli, 2019; Ferrière and Navarro, 2020;

[3] Because monetary-fiscal interactions are about the *subsequent* response of monetary policy rates, this environment is different to typical exercises studying state dependence, which tend to focus on how the effect of policy varies with some initial condition, e.g. the lagged output gap. This is why we often use the term *policy* dependence to differentiate our dimension of heterogeneity from more traditional state dependence.

[4] In exploiting the differential sensitivity of countries to shocks, our method has a connection to the approach in Nakamura and Steinsson (2014) and Guren, McKay, Nakamura, and Steinsson (2020).

Barnichon, Debortoli, and Matthes, 2020; Ghassibe and Zanetti, 2022; Jo and Zubairy, 2021). But this literature has often focused on a particular dimension of state dependence in isolation (e.g., expansions versus recessions) or dimensions one at a time.[5]

In terms of monetary-fiscal interactions, a number of papers have considered whether the fiscal multiplier is larger at the Zero Lower Bound (e.g., Ramey and Zubairy, 2018; Crafts and Mills, 2013; Kato, Miyamoto, Nguyen, and Sergeyev, 2018; Miyamoto, Nguyen, and Sergeyev, 2018). Canova and Pappa (2011) use sign-restrictions in a VAR framework and find that imposing a no-monetary response generates a larger multiplier. Wolf (2023) examines how empirical fiscal multipliers might have differed under a strict inflation targeting regime. Some papers also exploit regional variation to examine the fiscal multiplier. In these settings the aggregate effects of monetary policy are held constant (for examples, see Acconcia, Corsetti, and Simonelli, 2014; Nakamura and Steinsson, 2014; Corbi, Papaioannou, and Surico, 2019). Chodorow-Reich (2019) concludes that these "cross-sectional" multipliers are consistent with an aggregate "no-monetary-policy-response" multiplier of 1.7 or above.[6] Finally, some papers find that the exchange rate regime affects the size of the multiplier (e.g., Corsetti, Meier, Müller, and Devereux, 2012; Born, Juessen, and Müller, 2013; Ilzetzki, Mendoza, and Végh, 2013; Born, D'Ascanio, Müller, and Pfeifer, 2021), which is obviously related to whether policymakers are willing and able to use monetary tools.

These existing papers highlight the importance of the monetary offset, but often refer to a particular environment (e.g., the zero lower bound) and it is hard to know the right benchmark against which to measure the "usual" monetary response.[7] Relative to all these papers, our focus is therefore different. We aim to directly quantify the importance of this monetary-fiscal interaction on the aggregate fiscal multiplier more generally, and not just in certain episodes, and thus map out a range for how the fiscal multiplier varies with the monetary offset.

Summary This paper therefore makes three main contributions.

First, we show how to introduce decomposition methods in empirical macroeconomics more generally. Our decomposition approach turns out to be straightforward to implement and allows for a great deal of unspecified heterogeneity.

Second, using our approach, we provide new estimates of how the fiscal multiplier depends on the response of monetary policy. While fiscal multipliers are around 1 on average, there is considerable heterogeneity. In our policy experiments, fiscal multipliers can be as low as zero, or as high as 2 depending on the monetary policy response. These empirical results also have wider

[5]Moving beyond the fiscal literature, a number of papers have also considered similar questions about state dependence and time variation in the context of changes in monetary policy. See, for example, Primiceri (2005), Lo and Piger (2005), Peersman and Smets (2001), Tenreyro and Thwaites (2016), Angrist et al. (2018), Alpanda and Zubairy (2019), Jordà, Schularick, and Taylor (2020), Alpanda, Granziera, and Zubairy (2021).

[6]Guajardo, Leigh, and Pescatori (2011) also discuss how the degree of monetary accommodation might explain differences between their estimated spending and tax multipliers but do not formally attempt to estimate this interaction more generally.

[7]ZLB regimes are also potentially correlated with other factors affecting the size of the multiplier, e.g. the size of the output gap.

theoretical implications since an interaction effect is only present in models with nominal rigidities and where fiscal policy, at least partly, affects GDP through aggregate demand.

Third, our approach is naturally multi-dimensional. In this framework we revisit other popular forms of state dependence simultaneously, and show how to examine time-variation in the multiplier. We find that the change in the fiscal deficit and the size of the fiscal consolidation do not materially affect the size of the fiscal multiplier. But, like some other papers in the literature, we find that fiscal multipliers are larger in slumps (cyclically-low output states). Our approach will hopefully help researchers interested in estimating the non-linear, state-dependent, or time-varying effects of policy interventions using straightforward linear estimators.

The structure of the paper is as follows. Section 2 formally discusses the decomposition methods we use and how these can be introduced into macroeconomic analysis using local projections. Section 3 applies the new methods to study the interaction of monetary policy and the fiscal multiplier, showing that the approach recovers the theoretical multiplier when applied to simulated model data. Section 4 conducts robustness checks. We then conclude and discuss some policy implications.

2. Statistical design

This section introduces the statistical methods that will allow us to do three things: (1) decompose the impulse response into direct, indirect, and composition effects as discussed in the introduction; (2) calculate how other state variables affect the response to a policy intervention; and (3) investigate time-variation in the response according to variation in the conditioning information set.

We will be able to accomplish all three goals with a simple extension of the usual local projection framework that keeps the model linear in parameters and therefore can be estimated by simple regression methods. This extension is based on the well-known KBO decomposition (Kitagawa, 1955; Blinder, 1973; Oaxaca, 1973), a tool that is used often in applied microeconomics. In our context, the decomposition provides a general framework for decomposing heterogeneity around the average effect of a policy intervention in a multi-variate and time-varying manner.

2.1. Preliminary statistical discussion and intuition

This section introduces the main ideas with simplified examples. Formal statements of the usual assumptions can be found in, e.g., Wooldridge (2001) and Fortin, Lemieux, and Firpo (2011). Later on we provide assumptions for typical macroeconomics applications as we expand on our examples. Here we focus on the intuition. Note that when describing the behavior of random variables, we omit observation indexes, which are used once the discussion moves on to a finite sample.

A review of the main features of the KBO decomposition Suppose that, in a static setting for now, we are interested in the response of an outcome variable, y, to a *randomly assigned* intervention, f. We make this assumption to simplify the exposition, without getting into the details of how this is achieved. In general, we will assume only that $f \in \{0, 1\}$ is randomly assigned, at least conditional on controls $\boldsymbol{x}$, and the observed data are generated by the following mixture of unobservable latent variables, y_1 and y_0,

$$y = (1 - f)\, y_0 + f\, y_1 = y_0 + f\, (y_1 - y_0)\,. \tag{1}$$

That is, the observed random variable y is either the random variable y_0, which is observed when $f = 0$, or it is y_1 when $f = 1$. Note that the observed data belong to one state or the other. One cannot simultaneously observe both states—the observed data come from the latent mixture process described by Equation 1. As is standard, we refer to y_0 and y_1 as *potential outcomes* in the terminology of the Rubin causal model (Rubin, 1974).

These potential outcomes are random variables y_j with $j \in \{0, 1\}$. Suppose that they have unconditional means $E(y_j) = \mu_j$. A natural statistic of interest is $E(y_1 - y_0) = \mu_1 - \mu_0$, that is, the average difference in the unconditional mean between the treated and the control subpopulations. Although the potential outcomes y_1 and y_0 cannot be simultaneously observed, their moments (under random assignment), can be easily calculated.

We note that the potential outcomes approach and its notation can be somewhat new to applied macroeconomists. A few examples can help clarify basic notions. In a randomized controlled trial, a common (strong) ignorability assumption is that $y_j \perp f$ for $j = 0, 1$. This assumption does not imply that y and f are unrelated. Rather, the assumption means that the choice of intervention f is unrelated to the potential outcomes that may happen for a given choice of $f \in \{0, 1\}$. Hence a quantity such as $E(y_1 \,|\, f = 0)$ is well defined. It refers to the expected value that the random variable y_1 — referring to units in the treated subpopulation — would *counterfactually* take were it not exposed to treatment and instead had been placed in the control group (i.e. with $f = 0$). We will use such counterfactual expectations below.

We might reflect on the strong ignorability condition. It bears noting that even when this fails in practice, a milder condition of selection on observables, that is, $y_j \perp f \,|\, \boldsymbol{x}$ for $j = 0, 1$ would allow most of the results here to carry through and is akin to common identification based on exclusion restrictions (depending on what is included in $\boldsymbol{x}$). We expand on this point below.

2.2. The static Kitagawa-Blinder-Oaxaca decomposition

Without loss of generality, we can write $y_j = \mu_j + v_j$ where $E(v_j) = 0$, since $E(y_j) = \mu_j$ by definition, with $j \in \{0, 1\}$. Any heterogeneity in the treated and control subpopulations is therefore relegated to the terms v_j. Whenever covariates (explanatory variables or, simply, controls) $\boldsymbol{x}$ are available, they are useful to characterize heterogeneity across units (and later for us, across time) and we may assume additivity so that $v_j = g(\boldsymbol{x}) + \epsilon_j$.

As a starting point it is natural to further assume that these covariates enter linearly, so that $v_j = (\boldsymbol{x} - E(\boldsymbol{x}))\gamma_j + \epsilon_j$. The assumption of linearity simplifies the assumptions and the analysis considerably with little loss of generality in economics applications. We include the covariates in deviations from their unconditional mean to ensure that $E[(\boldsymbol{x}-E(\boldsymbol{x}))\gamma_j] = 0$, in which case unobserved heterogeneity is such that $E(\epsilon_j) = 0$. If observed heterogeneity is well captured by the vector of explanatory variables and the linearity assumption is correct, then it is also the case that $E(\epsilon_j \mid \boldsymbol{x}_j) = 0$. That is, the projection of y_j onto $\boldsymbol{x}_j$ is properly specified. Failure of identification results in this last condition being violated, as is well known. Later on, failure of this condition will lead us to propose an instrumental variable approach for controls that are simultaneously determined with the intervention f.

Researchers are often interested in understanding the overall effect of the intervention on outcomes. The KBO decomposition (Kitagawa, 1955; Blinder, 1973; Oaxaca, 1973) is used often in applied microeconomics for this purpose. It is worth going through its derivation here before later using similar arguments when analyzing a dynamic setting and using local projections. These derivations borrow heavily from Wooldridge (2001) and Fortin, Lemieux, and Firpo (2011), but are otherwise fairly standard in the literature.

A researcher's first line of inquiry is likely to involve calculating the average treatment effect defined earlier as $E(y_1 - y_0)$. However, we can think of this overall average treatment effect, based on our earlier assumption that heterogeneity in the latent potential outcomes enters linearly, by writing $E(y_1 - y_0)$ equivalently as:

$$\begin{aligned} E[E(y_1 \mid f = 1) - E(y_0 \mid f = 0)] &= E[E(y_1 \mid \boldsymbol{x}, f = 1) \; - \; E(y_0 \mid \boldsymbol{x}, f = 0)] \\ &= \{\, \mu_1 + E[\boldsymbol{x} - E(\boldsymbol{x}) \mid f = 1]\gamma_1 \; + \; \underbrace{E(\epsilon_1 \mid f = 1)}_{=0} \,\} \\ &- \{\, \mu_0 + E[\boldsymbol{x} - E(\boldsymbol{x}) \mid f = 0]\gamma_0 \; + \; \underbrace{E(\epsilon_0 \mid f = 0)}_{=0} \,\}\,. \end{aligned} \tag{2}$$

Of course, in population, the second equality is $\mu_1 - \mu_0$, the usual average treatment effect. However, in finite samples, the required balance condition may not be met exactly, as we shall see. To go further, by adding and subtracting $E[\boldsymbol{x} - E(\boldsymbol{x}) \mid f = 1]\gamma_0$ to the previous expression, we can begin to see how to think about heterogeneity more directly since:

$$\begin{aligned} E[E(y_1 \mid f = 1) - E(y_0 \mid f = 0)] &= (\mu_1 - \mu_0) \\ &+ E[\boldsymbol{x} - E(\boldsymbol{x}) \mid f = 1](\gamma_1 - \gamma_0) \\ &+ \{\, E[\boldsymbol{x} - E(\boldsymbol{x}) \mid f = 1] \; - \; E[\boldsymbol{x} - E(\boldsymbol{x}) \mid f = 0] \,\}\gamma_0\,. \end{aligned} \tag{3}$$

Equation 3 is central for us, and contains three interesting terms. The first, $\mu_1 - \mu_0$, is the difference in the unconditional means of the treated and control subpopulations. We refer to it as the *direct* effect of an intervention rather than simply the average treatment effect to draw a contrast with the *indirect* effect that we now explain.

The second term, $E[\boldsymbol{x} - E(\boldsymbol{x}) \mid f = 1](\gamma_1 - \gamma_0)$, reflects changes in how the covariates affect the outcome due to the intervention. We will refer to this term as the *indirect* effect of intervention. This term would capture the idea that, e.g, mathematics knowledge may translate into a higher salary for workers assigned to take additional training in computer science, but may not be otherwise helpful if there is no complementarity between both knowing mathematics and computer science. Notice that $E[\boldsymbol{x} - E(\boldsymbol{x}) \,|\, f = 1]\boldsymbol{\gamma}_0$ measures the salary of workers with a given background in mathematics, had they been counterfactually assigned *not* to take the additional training in computer science.

A natural hypothesis we will be interested in testing is $H_0 : \boldsymbol{\gamma}_1 - \boldsymbol{\gamma}_0 = \mathbf{0}$. Failure to reject the null suggests that the effect of the covariates on the outcome is not affected by the intervention. This would rule out any interaction between the intervention and the controls. Crucially, traditional estimates of impulse responses have implicitly assumed this to be the case up to now. Later on, we will see that such a hypothesis plays a critical role in evaluating impulse response state-dependence. Note that, in a properly designed randomized control trial, covariate balance would hold and would imply that the expectation term $E[\boldsymbol{x} - E(\boldsymbol{x}) \,|\, f = 1] = 0$ since $E(\boldsymbol{x}) = E[\boldsymbol{x}|f = 1] = E[\boldsymbol{x}|f = 0]$, so the *average* indirect effect should be zero. However, this does *not* mean that $\boldsymbol{\gamma}_1 - \boldsymbol{\gamma}_0 = \mathbf{0}$. The covariates $\boldsymbol{x}$ will still influence the way in which treatment affects the outcomes for individual realizations of the random variable $\boldsymbol{x}$. As we shall see this turns out to have important applications in understanding how responses vary along different dimensions of heterogeneity.

The final term $\{E[\boldsymbol{x} - E(\boldsymbol{x}) \,|\, f = 1] - E[\boldsymbol{x} - E(\boldsymbol{x}) \,|\, f = 0]\}\boldsymbol{\gamma}_0$ reflects how, all else equal, the effect of the intervention may be driven simply by differences in the average value of the explanatory variables between the treated and control subpopulations. We will call this term the *composition* effect. In a typical application, where observables are appropriately controlled for, an implication of this effect is that the residuals should not be predictable by covariates (a violation of selection-on-observables).

In practice, a natural way to obtain each term in the decomposition of Equation 3 in a finite sample with $t = 1, \ldots, T$, would be to estimate the following regression, using Equation 1 as the springboard (and under our maintained assumption of linearity),

$$y_t \; = \; \mu_0 \; + \; (\boldsymbol{x}_t - \overline{\boldsymbol{x}})\boldsymbol{\gamma}_0 \; + f_t\beta + f_t(\boldsymbol{x}_t - \overline{\boldsymbol{x}})\boldsymbol{\theta} \; + \; \omega_t \, , \tag{4}$$

where $\hat{\beta} = \hat{\mu}_1 - \hat{\mu}_0$ is an estimate of the *direct* (or average treatment) effect; and $\hat{\boldsymbol{\theta}} = \hat{\boldsymbol{\gamma}}_1 - \hat{\boldsymbol{\gamma}}_0$ so that $(\overline{\boldsymbol{x}}_1 - \overline{\boldsymbol{x}})\hat{\boldsymbol{\theta}}$ is an estimate of the *indirect* effect, where the notation $\overline{\boldsymbol{x}}_1$ refers to the sample mean of the covariates for the treated units, while $\overline{\boldsymbol{x}}_0$, used later, refers to the sample mean for control units.

From this expression, an important lesson is that $\boldsymbol{x}_t$ should be expressed relative to the mean $\overline{\boldsymbol{x}}$ to ensure the direct effect captures the average impact and the indirect effect captures heterogeneity around the average. For expositional purposes, we do not discuss here straightforward panel data extensions with fixed effects, but will consider these later in our application.

Finally, some important diagnostics emerge here, and they will reappear later. Rejection of the null $H_0 : \beta = 0$ says the intervention has a direct effect on the outcome. Even if one cannot reject this null, the intervention can affect the outcome *indirectly* as long as the null $H_0 : \boldsymbol{\theta} = \mathbf{0}$ is rejected. Of course, the reverse is plausible as well: i.e., that there is a direct effect, but not an indirect effect (in which case there is no heterogeneity due to state variables). This means that the effect of covariates on outcomes does not depend on whether a unit is treated or not. Finally, both nulls could be rejected in which case the intervention has no effect on the outcome, directly or indirectly.[8]

The literature has, of course, expanded on Equation 4 into more flexible semi- and non-parametric settings that we do not discuss here as our next goal is to adapt these lessons to local projections. By far, applications of local projections are implemented in linear settings and this the natural starting point of the next Section.

2.3. Understanding dynamic impulse response heterogeneity with LPs

The static methods discussed in Sections 2.1 and 2.2, while common in applied microeconomics research, have not permeated macroeconomics as much. In this section we show that local projections offer a natural bridge between literatures and hence offer a more detailed understanding of dynamic impulse responses, the workhorse of applied macroeconomics research.

In order to move from the preliminary statistical discussion to a time series setting in which to investigate impulse responses, we define the outcome random variable observed at a horizon h periods after the intervention as $y(h)$, where a typical single observation from a finite sample of T observations is denoted y_{t+h}. We focus on the pure time series case for expositional purposes though our application later uses panel data.

As before, we begin with a binary policy intervention (i.e., the treatment) denoted $f \in \{0, 1\}$ where a typical single observation from a finite sample is denoted f_t. Later we allow for a continuous f. A vector of observable variables is denoted $\boldsymbol{x}$, where a typical single observation from a finite sample is denoted $\boldsymbol{x}_t$. Note that $\boldsymbol{x}$ may include contemporaneous values and lags of a vector of variables including the intervention, as well as lags of the (possibly transformed) outcome variable, among others. Moreover, define the $(H + 1) \times 1$ vector $\boldsymbol{y} = (y(0), y(1), \ldots, y(H))$, or when denoting an observation from a finite sample, $\boldsymbol{y}_t = (y_t, y_{t+1}, \ldots y_{t+H})$.

Assumptions As a natural starting point regarding the assignment of the policy intervention we follow Angrist, Jordà, and Kuersteiner (2018), whose selection-on-observables assumption we restate here for convenience.

Assumption 1. Conditional ignorability or selection on observables. *Let $\boldsymbol{y}_f$ denote the potential outcome that the vector $\boldsymbol{y}$ can take on impact and up to H periods after intervention $f \in \{0, 1\}$. Then we say f is randomly assigned conditional on $\boldsymbol{x}$ relative to $\boldsymbol{y}$ if*

$$\boldsymbol{y}_f \perp f \,|\, \boldsymbol{x} \qquad \textit{for } f = f(\boldsymbol{x}, \eta; \boldsymbol{\phi}) \in \{0, 1\}\,; \;\; \boldsymbol{\phi} \in \Phi\,.$$

[8] Note that the error term is $\omega_t = \epsilon_{0,t} + f_t\,(\epsilon_{1,t} - \epsilon_{0,t})$. Under the maintained assumptions, it has mean zero conditional on covariates.

The conditional ignorability assumption makes explicit that the policy intervention f is itself a function the observables $\boldsymbol{x}$, unobservables η, and a parameter vector $\boldsymbol{\phi}$. It means that $\boldsymbol{y}_f \perp \eta$, that is, the unobservables are random noise. Moreover, we assume that $\boldsymbol{\phi}$ is constant for the given sample considered. In other words, we rule out variation in the rule assigning intervention. In more general settings, if this assumption fails, it would be natural to use an instrumental variable approach, as we do in our application. It is worth remarking that statistical independence is a stricter condition than simple absence of correlation as the former rules out any dependence on any potential nonlinear transformation as well.

The literature has carefully considered identification issues using assumptions such as the conditional ignorability assumption just introduced. However, considering Equation 4, is this assumption enough to explore the interaction $f_t(\boldsymbol{x}_t - \bar{\boldsymbol{x}})$ to interpret $\boldsymbol{\theta}$ causally? For example, denote the first variable in $\boldsymbol{x}_t$, as r_t, where we think of this r_t as being the single stratification variable of interest, for now. As in our application later, think of f_t as the fiscal intervention, and r_t as a monetary intervention.

In practice, we would expect f_t and r_t to be jointly determined so that analyzing the response of the outcome due to an intervention in f_t and stratifying on a particular value of r_t would require some further assumptions to attach causal meaning to responses to f_t allowed to vary depending on r_t. This is, of course, well understood in applied microeconomics (e.g. Fortin, Lemieux, and Firpo, 2011).[9] The implicit assumption adopted by much of the literature is to impose a hierarchical causal structure. Decisions on f_t are contingent on r_t, but not the other way around. This is also implicit in Assumption 1. However, in our application below, we will view f_t and r_t as simultaneously determined, hence we will need an identification strategy.

In some cases, stratification is based on pre-determined variables, such as the lagged output gap. Does this mean that we can directly attach causal meaning to a given pre-determined variable of interest? The answer is no. Pre-determined variables could be simultaneously correlated with each other. For example, ZLB states also tend to exhibit negative output gaps. Negative output gaps are also correlated with low inflation, and so forth. Past values of output and inflation are also likely to respond to previous policy interventions over the same period. One would need to invoke further assumptions (e.g., short-run exclusion restrictions) or look for additional instrumental variables to make precise causal statements.

Although the general statement of conditional ignorability in Assumption 1 provides a great deal of flexibility (see Angrist, Jordà, and Kuersteiner, 2018), a simpler assumption can be made when considering a linear framework, as we do in the analysis that follows. In particular, for our purposes, the following assumption will suffice:

Assumption 2. Conditional mean independence. *Let $E(\boldsymbol{y}_f) = \boldsymbol{\mu}_f$ for $f \in \{0, 1\}$ so that, without loss of generality, $\boldsymbol{y}_f = \boldsymbol{\mu}_f + \boldsymbol{v}_f$. As before, we now assume linearity so that $\boldsymbol{v}_f = (\boldsymbol{x} - E(\boldsymbol{x}))\Gamma_f + \boldsymbol{\epsilon}_f$. Because of the dimensions of $\boldsymbol{y}_f$, we use the notation Γ_f instead of γ_f since Γ_f is now a matrix of coefficients with row*

[9] Awareness of the issue in macroeconomics resurfaced in Gonçalves, Herrera, Kilian, and Pesavento (2022).

dimension $H+1$, each row containing the corresponding γ_f^h vector for $h=0,\ldots,H$. Then, if $E(\boldsymbol{\epsilon}_f \mid \boldsymbol{x}) = \mathbf{0}$ for $f \in \{0,1\}$, we have:

$$E(\boldsymbol{y}_f) = E[E(\boldsymbol{\mu}_f + \boldsymbol{v}_f|\boldsymbol{x})] \quad = \quad \boldsymbol{\mu}_f + E\{E[(\boldsymbol{x}-E(\boldsymbol{x}))\Gamma_f + \boldsymbol{\epsilon}_f|\boldsymbol{x}]\} \tag{5}$$

$$= \quad \boldsymbol{\mu}_f + \underbrace{E\{(\boldsymbol{x}-E(\boldsymbol{x}))\Gamma_f\}}_{=0} + E[\underbrace{E(\boldsymbol{\epsilon}_f|\boldsymbol{x})}_{=0}] \qquad \textit{for } f \in \{0,1\}\,. \tag{6}$$

Based on Assumption 2, local projections can be easily extended to have the same format as the expression at Equation 4 when dealing with a finite sample of $t = 1,\ldots,T$ observations:

$$y_{t+h} \quad = \quad \begin{cases} \mu_0^h + (\boldsymbol{x}_t - \bar{\boldsymbol{x}})\gamma_0^h + \epsilon_{t+h}^0 & \text{if} \quad f_t = 0\,, \\ \mu_1^h + (\boldsymbol{x}_t - \bar{\boldsymbol{x}})\gamma_1^h + \epsilon_{t+h}^1 & \text{if} \quad f_t = 1\,, \end{cases}$$

which we can write as

$$= \quad \underbrace{\mu_0^h \;+\; (\boldsymbol{x}_t - \bar{\boldsymbol{x}})\gamma_0^h \;+\; f_t\,\beta^h}_{\text{usual local projection}} \;+\; \underbrace{f_t\,(\boldsymbol{x}_t - \bar{\boldsymbol{x}})\,\boldsymbol{\theta}^h}_{\text{KBO extension}} \;+\; \omega_{t+h}\,; \qquad \text{for} \;\; h = 0,1,\ldots,H\,;\; t = h,\ldots,T\,. \tag{7}$$

where clearly $\beta_h = (\mu_1^h - \mu_0^h)$; $\boldsymbol{\theta}^h = (\gamma_1^h - \gamma_0^h)$ and $\omega_{t+h} = \epsilon_{t+h}^0 + f_t(\epsilon_{t+h}^1 - \epsilon_{t+h}^0)$. Again, we maintain the time series notation and refrain from adding more parameters, such as fixed effects as is common in panel data applications, for simplicity and without loss of generality. Thus, relative to the usual specification of a local projection, the only difference is the additional KBO term, $f_t\,(\boldsymbol{x}_t - \bar{\boldsymbol{x}})\,\boldsymbol{\theta}^h$.[10]

Finally, note that, in a time series context, we require an additional technical assumption about the stationarity of the covariate vector $\boldsymbol{x}$. Without it, calculating means for the treated and control subpopulations would not be a well-defined exercise. In a typical local projection, it is not necessary to make such an assumption because the parameter of interest is β^h and all that is required for inference is for the projection to have a sufficiently rich lag structure to ensure that the residuals are stationary. Consequently, we make an additional assumption here, as follows:

Assumption 3. Ergodicity. *The vector of covariates $\boldsymbol{x}_t$ — which can potentially include lagged values of the (possibly transformed) outcome variable and the treatment, as well as current and lagged values of other variables — is assumed to be a covariance-stationary vector process ergodic for the mean (Hamilton, 1994).*

Ergodicity ensures that the sample mean converges to the population mean. Assuming covariance-stationarity is a relatively standard way to ensure that this is the case. More general assumptions could be made to accommodate less standard stochastic processes. However, covariance-stationarity and ergodicity are sufficiently general to include commonly observed processes in practice.

[10] Though not central to our discussion, it is worth noting that the residual ω_{t+h} will have the well known moving average structure common to local projections. Several simple inferential solutions are available (see, e.g. Jordà, 2005; Montiel Olea and Plagborg-Møller, 2021).

The average LP As a result of this simple extension, estimates of the *average* impulse response at any horizon h can be calculated in parallel fashion to Section 2.2.

It is important to note that the counterfactual experiment here is as follows: what is the difference between the expected average outcome path when the intervention is $f = 1$ and the expected average outcome path when the intervention is $f = 0$, that is, $E_x[E[y_{t+h}|f_t = 1, \boldsymbol{x}_t] - E[y_{t+h}|f_t = 0, \boldsymbol{x}_t]]$. Note that we use the law of iterated expectations to average over all values of $\boldsymbol{x}$. The *average* components of this impulse response at horizon h are given by

$$\begin{aligned}
\text{Direct effect:} \quad & \hat{\mu}_1^h - \hat{\mu}_0^h = \hat{\beta}^h\,, \\
\text{Indirect effect:} \quad & (\bar{\boldsymbol{x}}_1 - \bar{\boldsymbol{x}})(\hat{\boldsymbol{\gamma}}_1^h - \hat{\boldsymbol{\gamma}}_0^h) = (\bar{\boldsymbol{x}}_1 - \bar{\boldsymbol{x}})\hat{\boldsymbol{\theta}}^h\,, \\
\text{Composition effect:} \quad & (\bar{\boldsymbol{x}}_1 - \bar{\boldsymbol{x}}_0)\hat{\boldsymbol{\gamma}}_0^h\,,
\end{aligned}$$

where, recall, $\bar{\boldsymbol{x}}_f$ refers to the sample mean of the controls in each of the subpopulations $f \in \{0, 1\}$.

Note that, under a design with true randomization, we would have *balance* hold with $\bar{\boldsymbol{x}}_1 = \bar{\boldsymbol{x}}_0 = \bar{\boldsymbol{x}}$. In that case, the indirect and composition effects would vanish, on average, and the direct effect β^h alone would be the average treatment effect at horizon h. However, without true randomization, omitting the covariates in the regression would clearly lead to a biased estimate of the treatment effect. This is easy to see from the composition effect, which would not be zero. This highlights current best practice when researchers estimate local projections by including a rich set of covariates.

The time-varying LP However, we may also be interested in the time-varying impulse response at time t. The counterfactual here is different: What is the difference between the expected average outcome path when $f = 1$ relative to when $f = 0$, holding *fixed* $\boldsymbol{x} = \boldsymbol{x}_t$. That is, we assume the policy intervention which changes f from 0 to 1 is implemented with covariates held at their realized time t level. More specifically and in contrast to the static LP, we are interested in calculating $E[y_{t+h}|f_t = 1, \boldsymbol{x}_t] - E[y_{t+h}|f_t = 0, \boldsymbol{x}_t]$.

The components of this time-varying impulse response at horizon h are different and given by

$$\begin{aligned}
\text{Direct effect:} \quad & \hat{\mu}_1^h - \hat{\mu}_0^h = \hat{\beta}^h\,, \\
\text{Indirect effect, time } t\text{:} \quad & (\boldsymbol{x}_t - \bar{\boldsymbol{x}})(\hat{\boldsymbol{\gamma}}_1^h - \hat{\boldsymbol{\gamma}}_0^h) = (\boldsymbol{x}_t - \bar{\boldsymbol{x}})\hat{\boldsymbol{\theta}}^h\,, \\
\text{Composition effect:} \quad & (\boldsymbol{x}_t - \boldsymbol{x}_t)\hat{\boldsymbol{\gamma}}_0^h = \boldsymbol{0}\,.
\end{aligned}$$

The direct effect is time invariant by construction given linearity, but the indirect effect is now time-varying. The composition effect is now zero by construction. Two advantages of our approach stand out. First, time-variation in the impulse response can be calculated with simple regression techniques, in fact, with models linear in parameters where estimation and inference and standard. Second, time-variation is due to observable economic characteristics rather than being driven by unobservable latent processes, as is done in, e.g., Cogley and Sargent (2005) and Primiceri (2005). This makes interpretation of the impulse response easier as we can refer back to the conditions that generated the impulse response at any point in the sample rather than having to correlate latent processes with observables and then trying to deduce the relationship.

The counterfactual LP Finally, we consider the case where there might be an interaction between the original policy intervention, f, and perhaps and alternative intervention $r \in \boldsymbol{x}$ (or, more generally, another state variable). In our application below, we consider the interaction of fiscal and monetary policy. Because interest is now in counterfactual experimentation with f when we *manipulate* r to attain different levels, we can not take r as given: the policy variable r is likely chosen in response to f, among other factors. This will require further identification assumptions on r. One option is to use instrumental variables to isolate exogenous variation in r. In our application, as mentioned earlier, we exploit cross country variation in the response of r to the intervention f instead given the panel structure of our data.

Thus, we are interested in calculating $E_{\tilde{x}}[E[y_{t+h}|f_t = 1, r_t = r^*, \tilde{x}] - E[y_{t+h}|f_t = 1, r_t = \bar{r}, \tilde{x}]]$, where as before $\tilde{x}$ refers to the vector of controls excluding r. The notation r^* simply refers to a value of interest for the policy variable r_t that the analyst chooses to compare to, say $\bar{r}$. Note also that the results are averaged across all possible values of the remaining controls, $\tilde{x}$. In terms of the parameters from our original regression, we now have:

$$
\begin{aligned}
\text{Direct effect, } r_t = \bar{r}, \text{ time } t: &\qquad \hat{\mu}_1^h - \hat{\mu}_0^h = \hat{\beta}^h\,, \\
\text{Direct effect, } r_t = r^*, \text{ time } t: &\qquad \hat{\mu}_1^h - \hat{\mu}_0^h + (r^* - \bar{r})(\hat{\gamma}_{1r}^h - \hat{\gamma}_{0r}^h) = \hat{\beta}^h + (r^* - \bar{r})\hat{\theta}_r^h\,.
\end{aligned}
$$

In other words, $(r^* - \bar{r})\hat{\theta}_r^h$ modulates the direct effect of the intervention f depending on r.

Continuous policy interventions Equation 7 is in regression form, which naturally accommodates situations when the policy intervention varies continuously instead of discretely, as we have entertained thus far. However, more generally, one has to acknowledge the restrictions that such a setting imposes. In particular, because of the assumption of linearity, doubling f means that the effect on the outcome is doubled as well. In some situations this restriction may be undesirable. In pharmacology, for example, doubling the dose does not make the treatment twice as effective, for example, since higher doses can have diminishing effects, or even become lethal. However, as is true with regression in general, interpreting estimated coefficients within the training sample is considered to be a reasonable approximation. That is, we should interpret the results with caution once we venture far from steady state. This is, however, no different to the usual assumptions made in the empirical macroeconomics literature.

2.4. Putting the tools to use

The core assumption of impulse responses estimated with linear methods is that the effect of the controls on the outcome is independent of intervention. The previous sections showed that a simple extension of standard local projections, Equation 7, opens up numerous possibilities that we now detail for the empirical user.

Hypothesis tests Equation 7 allows the decomposition of the impulse response into two main components, as discussed, the direct, and indirect effects, though it also highlights the importance of appropriate controls via the composition effect. The natural first step is to test the null hypothesis of the direct effect, $H_0 : \mu_1^h = \mu_0^h$ or in terms of the LP in Equation 7, $H_0 : \beta^h = 0$ for $h = 0, 1, \ldots, H$. This can be easily done with standard Wald tests for individual coefficients.

Next, in order for there to be an indirect effect, it must be the case that $H_0 : \gamma_1^h = \gamma_0^h$ can be rejected, or in terms of the coefficients of the extended LPs, $H_0 : \boldsymbol{\theta}^h = \mathbf{0}$ for $h = 0, 1, \ldots, H$ is rejected. Again, for individual coefficients, this can be tested with standard Wald statistics. Thus, note that even if there is no direct effect, the intervention could still affect the outcome indirectly through its differential effect on the controls.

In the unlikely event that $\gamma_0^h = \mathbf{0}$ for $h = 0, 1, \ldots, H$, one might be tempted to omit the covariates. However, one should probably check the balance condition, $H_0 : E(\boldsymbol{x}|f = 1) = E(\boldsymbol{x}|f = 0)$. Failure of this hypothesis, *and* rejection of $H_0 : \boldsymbol{\theta}^h = \mathbf{0}$ for $h = 0, 1, \ldots, H$ would indicate that $\gamma_1^h \neq 0$ and that identification has not been achieved—the distribution of the covariates across treated and control subpopulations would be different (in mean) and they would be relevant in measuring the indirect effect. For this reason, it safest to include $\boldsymbol{x}$ even if insignificant since the interaction terms $f_t(\boldsymbol{x}_t - \overline{\boldsymbol{x}})$ could be significant.

The baseline impulse response Next, a natural next step is to report the direct effect, β^h for $h = 0, 1, \ldots, H$. This is the usual impulse response reported in the literature except that the LP now includes the terms resulting from the interaction of the intervention and the covariates. Omitting these terms, specially when $H_0 : \boldsymbol{\theta}_h = 0$ is rejected, could result in bias for $\hat{\beta}^h$. If the interaction terms are correlated with the policy variable, their omission would generate omitted variable bias.

Time-variation in the impulse response Before investigating how policy variables might interact with one another, which requires further identification assumptions, it makes sense to examine the time-variation of the impulse response. This can help the researcher identify specific periods of interest that can then be related to significant economic events.

State dependence One of the most popular extensions of LPs is in the estimation of state-dependent impulse responses. Auerbach and Gorodnichenko (2012) find asymmetric effects of government spending changes based on whether the economy is in a boom or a bust. Jordà and Taylor (2016) find similar asymmetries using the Guajardo, Leigh, and Pescatori (2014) dataset. In the monetary policy literature, for example, Angrist, Jordà, and Kuersteiner (2018) show that monetary policy loosening is less effective at stimulating the economy than tightening is at dampening it. Tenreyro and Thwaites (2016) find asymmetric effects based on whether the economy is in a boom or a bust. Jordà, Schularick, and Taylor (2020), using a different approach, find that low inflation environments and large output gaps seem to dull stimulative policy.

The previous discussion, however, showed that there are two issues often overlooked in these studies (including some of our own!). One is that state-dependence is likely to be multidimensional. Thus, omitting potential sources of state heterogeneity and simply focusing on the one dimension of interest, could result in omitted variable bias. Second is that state-variables are likely to respond to policy interventions themselves and be correlated with one another for this and other economic mechanisms. This requires further identification assumptions, as we have shown.

3. Decomposing the fiscal multiplier

We now illustrate our approach by applying the KBO decomposition to the fiscal multiplier. Paying close attention to identification of the indirect effect, we show how the dynamic causal effect of changes in fiscal policy interacts with the response of monetary policy. The first part of this section explains how we identify the indirect effects coming from monetary-fiscal interactions. We then present our main results. Finally, given our framework is multi-dimensional, we decompose the fiscal multiplier along a number of other popular dimensions and illustrate the challenges of considering uni-dimensional state dependence alone.

3.1. Data and approach

Identifying the causal effect of a change in fiscal policy requires some exogenous variation in policy, even if we are interested in the average effect. As a result, there is a large literature on the identification of exogenous changes in fiscal policy and we rely on an off-the-shelf and well-established dataset of exogenous fiscal interventions. Guajardo, Leigh, and Pescatori (2014) construct a cross-country panel dataset of plausibly exogenous movements in government spending and taxes that were introduced for the purpose of fiscal consolidation. The identification approach follows Romer and Romer (2010) and focuses on consolidations that were designed to tackle an inherited historical budget deficit, but were not responding to current or prospective business cycle fluctuations.

Although Guajardo, Leigh, and Pescatori (2014) use a mix of distributed lag models and vector autoregressions for estimation, in the next section we will follow Jordà and Taylor (2016) and employ local projections to show how the KBO decomposition can be tractably applied to the estimation of impulse response functions in that framework. Following Equation 7 in Section 2, we estimate the following sequence of panel local projections,

$$y_{i,t+h} - y_{i,t-1} = \mu_i^h + (\boldsymbol{x}_{i,t} - \overline{\boldsymbol{x}}_i)\,\gamma_0^h + f_{i,t}\,\beta^h + f_{i,t}\,(\boldsymbol{x}_{i,t} - \overline{\boldsymbol{x}}_i)\,\boldsymbol{\theta}_x^h + \omega_{i,t+h}\,, \qquad h = 0, 1, \ldots, H\,, \tag{8}$$

where y is a particular variable of interest, for example log GDP, the deficit to GDP ratio or the real interest rate; t refers to the time period and i refers to the country; μ_i^h is a country fixed effect; $f_{i,t}$ is the policy intervention or treatment, in this case the country-specific fiscal consolidation shock. $\boldsymbol{x}_{i,t}$ is the vector of additional covariates, with mean $\bar{\boldsymbol{x}}_i$.

In typical empirical fiscal multiplier papers, β^h is the key object of interest: the percent effect on, e.g., GDP, following a 1% of GDP fiscal consolidation.[11] The KBO interaction terms are typically ignored in Equation 8 or a particular specification is employed using a single state variable. In our framework, $\boldsymbol{x}_{i,t}$ are both control variables *and* characteristics of the treated subpopulation that may influence the way in which treatment affects the outcome. In principle, a broad range of $\boldsymbol{x}_{i,t}$ variables could modulate the effects of fiscal treatment and many $\boldsymbol{x}_{i,t}$ variables in macroeconomics will be highly correlated, making causal statements difficult.

In our baseline specification $\boldsymbol{x}_{i,t}$ includes the typical controls used in other studies. Specifically, we include two lags of real GDP growth, the deficit to GDP ratio, the change in the real interest rate and, following Jordà and Taylor (2016), the output gap to control for the state of the cycle.[12] In terms of the dependent variables, the response of the deficit to GDP ratio is not available in the Guajardo, Leigh, and Pescatori (2014) dataset, but is useful for computing certain definitions of the fiscal multiplier. We therefore merge the Guajardo, Leigh, and Pescatori (2014) fiscal consolidation shocks with the Jordà, Schularick, and Taylor (2017) Macrohistory Database (`http://www.macrohistory.net/database/`), which contains a wider array of variables that we can employ as outcomes in our LP analysis.

As an empirical starting point, we show in Figure 2 the impulse response functions estimated from Equation 8 under the assumption that $\theta_x^h = 0$. This approach was chosen first as it is close to specifications typically seen in the existing literature and provides a baseline average effect. The figure shows that a 1% of GDP improvement in the government fiscal balance leads to a peak fall in GDP of around 1% over 4 years. Despite some differences in sample and specification, Panel (a) of Figure 2 is very similar to the original results in Guajardo, Leigh, and Pescatori (2014). The most comparable IRF is shown in Figure 2 of the working paper (Guajardo, Leigh, and Pescatori, 2011), and is very similar to Figure 2, with a peak effect on GDP occurring 2–3 years after the shock, and between 0.5% and 1% in magnitude.

Panel (b) of Figure 2 show how monetary conditions change on average. Real short term interest rates fall following a fiscal consolidation. To the extent that monetary policy can support the economy when GDP falls, the decline in the real short rate implies that the average fiscal consolidation is associated with monetary accommodation, which is perhaps not unexpected. The exact effect on GDP, however, will depend in the precise degree of accommodation by the monetary authority, in other words the strength of the "monetary offset" at the time. What we see in Figure 2 is only the effect on average. If the fall in the real rate were smaller, for example, we might expect to see a more severe contraction in GDP. Decomposing this average, and characterizing the heterogeneity

[11]This could be interpreted as one measure of a fiscal multiplier. But later we compute cumulative multipliers from the IRFs to explicitly take account of the full dynamic path of GDP and the fiscal variables.

[12]Including time fixed effects extends standard error bands without materially affecting the point estimates. Thus, to improve the precision of the estimates in the Kitagawa-Blinder-Oaxaca decomposition discussed below, we capture a time-varying global factor by including world real GDP growth.

Figure 2: *Effects of a 1 percentage point of GDP fiscal consolidation*

(a) *Response of GDP (%)* **(b)** *Response of short term real interest rate (% points)*

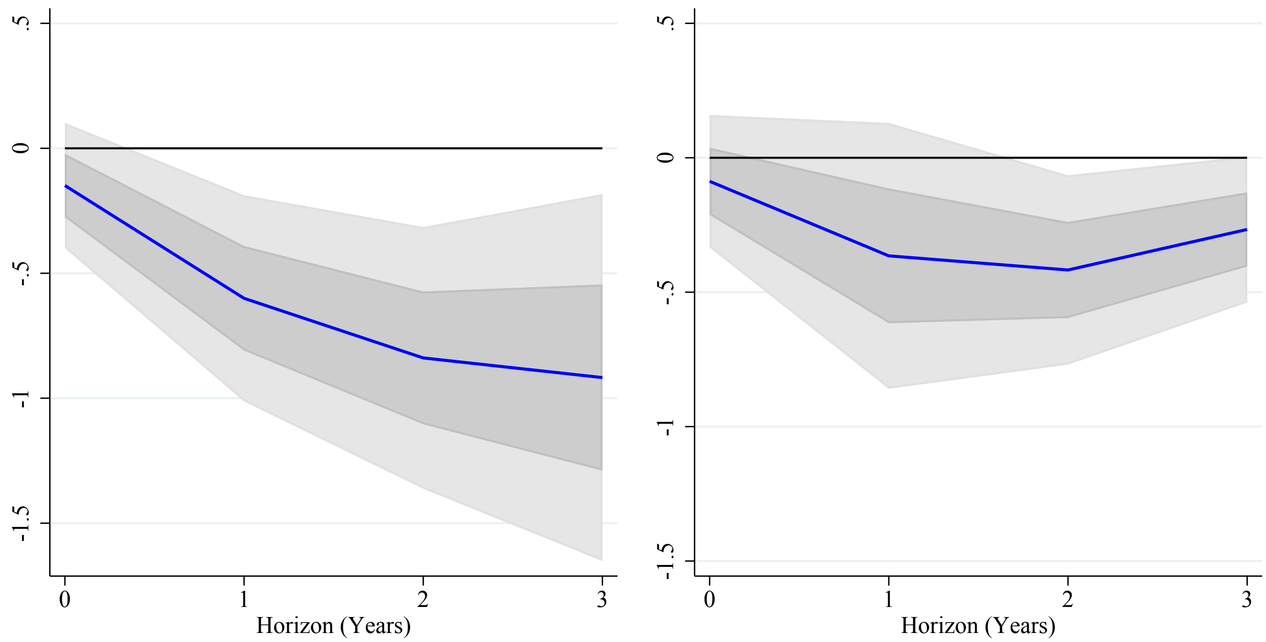

Notes: Vertical axes reported in percent change with respect to the origin. One and two standard deviation confidence bands for each coefficient estimate shown as grey areas. Local projections as specified in equation (8) excluding indirect effects and using two lags of each control described therein. Sample 1978:1–2009:4. See text.

around it, is, of course, the goal of this section.

To decompose the average effect into the part related to the monetary offset, we start from the general KBO equation above, Equation 8. Let $\boldsymbol{x}_{i,t}$ denote the vector of conventional covariates (discussed above). In addition, we will now include a further covariate that captures the sensitivity of monetary policy to the fiscal treatment. Denote this variable $\Theta_{i,t}$. We will discuss below how we go about measuring this. Then we have

$$y_{i,t+h} - y_{i,t-1} \;=\; \mu_i^h \;+\; (\boldsymbol{x}_{i,t} - \bar{\boldsymbol{x}}_i)\,\gamma_0^h \;+\; f_{i,t}\,\beta^h \;+\; \underbrace{f_{i,t}\,(\boldsymbol{x}_{i,t} - \bar{\boldsymbol{x}}_i)\,\boldsymbol{\theta}_x^h}_{\substack{\text{more traditional} \\ \text{indirect effects}}} \;+\; \underbrace{f_{i,t}\,\Theta_{i,t}\,\theta_f^h}_{\substack{\text{monetary-fiscal} \\ \text{indirect effects}}} \;+\; \omega_{i,t+h}\,. \tag{9}$$

The term $f_{i,t}\,\Theta_{i,t}\,\theta_f^h$ explores how the effects of the fiscal intervention are modulated by monetary policy. From the Kitagawa-Blinder-Oaxaca decomposition, the indirect effect from coming from the interaction of monetary and fiscal policy is given by $\Theta_{i,t}\,\theta_f^h$.

As discussed in Section 2, this empirical specification follows from a *decomposition* approach. As such, without further assumptions, we cannot necessarily interpret the coefficients on the $\boldsymbol{x}$ variables as causal. We therefore need a measure of $\Theta_{i,t}$ that can be used for identification of the monetary offset.[13]

To address this identification question, we employ a logic akin to Nakamura and Steinsson (2014) and Guren, McKay, Nakamura, and Steinsson (2020), who use the differential sensitivity of regions to more aggregate fluctuations as an identification strategy. In our approach, we use the differential sensitivity of interest rates to identified fiscal treatment across countries.

Our identifying variation relies on the idea that there is some cross-country variation in monetary policy. This could be related to different preferences of policymakers, different historical attitudes towards the degree of stabilization, different compositions of the policy committees etc. In the illustrative model used to produce Figure 1, this idea is captured by variation in the coefficient in the Taylor Rule (see Appendix J for full details of the model). For our empirical approach, we will exploit the panel nature of the fiscal dataset and focus on average cross-country differences in monetary policy behavior. As such, $\Theta_{i,t} = \Theta_i$. We will now explain how this works in more detail.

3.2. A simple policy sensitivity-based proxy

In this section, we propose a simple proxy for Θ_i based on the cross-country sensitivity of monetary policy rates to exogenous fiscal treatment. There are two steps. First, we estimate Θ_i by regressing the change in the nominal policy interest rate from $t-1$ to $t+h$ on the fiscal consolidation variable $f_{i,t}$, allowing the coefficient to vary by country. More specifically, we estimate the following sequence of local projections,

$$R_{i,t+h} - R_{i,t-1} \;=\; \mu_i^h \;+\; (x_{i,t} - \bar{x}_i)\,\gamma_0^h \;+\; \sum_{j=1}^{N} f_{i,t} \cdot \mathrm{I}\left[i = j\right] \cdot \tilde{\Theta}_i^h \;+\; \omega_{i,t+h}\,, \tag{10}$$

where $R_{i,t}$ is the short-run interest rate under the control of the monetary authority in country i in time t. Our key variable of interest is $\tilde{\Theta}_i^h$ which is allowed to vary by country. $\tilde{\Theta}_i^h$ captures the differential sensitivity of policy rates in country i to fiscal treatment.[14] We therefore use this as the basis of our proxy for the monetary response in a second stage regression that explores heterogeneity in the fiscal multiplier.

In the second stage we estimate our main empirical specification, Equation 9 using the estimated coefficients $\tilde{\Theta}_i^h$ from the first step above as our proxy. Note that in Equation 9 all covariates should be expressed relative to the mean. Θ_i in Equation 9 therefore refers to the de-meaned $\tilde{\Theta}_i^h$ from the first stage above.[15] Since we are interested in the dynamic causal effect via impulse response

[13] As noted earlier, monetary-fiscal interactions reflect the *subsequent* response of monetary policy rates. This environment is different to typical exercises studying state dependence, which tend to focus on how the effect of policy varies with some lagged state variable, e.g., the lagged output gap.

[14] In terms of other controls, $x_{i,t}$ includes the lagged change in the policy rate to capture persistence in the policy rate and the regression also includes country fixed effects. For the baseline results we keep this specification parsimonious, which helps improve the precision of the estimates. We have also reproduced the main results considering more elaborate specifications with further controls. The results for how the multiplier varies with monetary policy are very similar, so we maintain the parsimonious specification here.

[15] This ensures that our experiments below are relative to the "typical" average response of policy rates

functions, all these steps are run for each h. The right hand side of Equation 9 therefore contains $f_{i,t}\,\Theta_i^h\,\theta_f^h$. This is like interacting fiscal treatment at time t with the predicted subsequent response of the interest rate: $f_{i,t}\,\Theta_i^h$ is the fitted value for the *future interest rate response* from the first stage regression. As a result, this approach is like instrumental variables estimation where $f_{i,t}\,\Theta_i^h$ is the fitted value taken from the first stage. Our approach is therefore somewhat related to the sensitivity instrument approach of Guren, McKay, Nakamura, and Steinsson (2020).[16]

The identification assumption underlying this approach is that there is variation in the average response of monetary policy to shocks across countries but that this variation is not, on average, correlated with other factors that make the economy more sensitive to fiscal policy. Note that average differences across country are not a threat to identification, as these factors are captured by the country fixed effects. It is also not a problem if policy rates are more sensitive to *other* macroeconomic shocks, this is because our fiscal shocks are identified as orthogonal to other macroeconomic disturbances.[17] The strategy *is* therefore identifying how the fiscal multiplier varies with the cross-country sensitivity of policy rates to fiscal treatment.

The remaining issue is more specific. The concern is that this cross-country sensitivity of monetary policy to exogenous fiscal treatment might occur for reasons other than the monetary offset. This could occur if there is heterogeneity in the multiplier on average across countries for non-monetary reasons, and if monetary policy responds to fiscal treatment *indirectly*. Such concerns would tend to imply our estimates are a lower bound on the strength of the monetary offset, however. The simple proxy-based approach we study here is therefore a natural starting point and in Section 4 we will show that relaxing these assumptions broadly delivers similar results.

Inspecting the approach Can the reduced form KBO approach outlined above recover the true theoretical variation in the fiscal multiplier shown earlier in Figure 1? To do this, we simulate data from the New Keynesian model used to produce Figure 1 (see the appendix for the model details) for a hypothetical set of "countries" where each country differs in how monetary policy responds to inflation.[18] This environment theoretically captures the identification assumptions made above.[19]

---

in the sample. The average of the coefficients from the first stage regression also accords with conventional wisdom about how monetary policy tends to loosen following a fiscal consolidation on average. For example, in years 0, 1, 2 and 3 the average value of the country level coefficients is $-0.3, -0.6, -0.7$ and -0.5 respectively.

[16] We are therefore using cross-country heterogeneity in the country interest rate response to identified country-level shocks. This differs from Nakamura and Steinsson (2014) and Guren, McKay, Nakamura, and Steinsson (2020) who exploit differential local sensitivity to *common* aggregate or regional fluctuations.

[17] The appendix provides an illustration of how our approach a works in a simple two equation static system.

[18] We use the term "country" loosely here. In this simple example these are simply cross sectional units with different degrees of monetary accommodation, which is captured by a different coefficient on inflation in the monetary policy rule.

[19] As explained earlier, the rest of the model is the standard 3-equation New Keynesian model including hand to mouth households. Full details of the model are shown in Appendix J. Our goal is to illustrate how the Kitagawa-Blinder-Oaxaca approach identifies the importance of monetary-fiscal interactions for the size of the empirical fiscal multiplier. This section does not develop a theoretical framework to quantitatively rationalize the empirical magnitudes we will estimate below.

Figure 3: *Theoretical state dependence versus Kitagawa-Blinder-Oaxaca estimates*

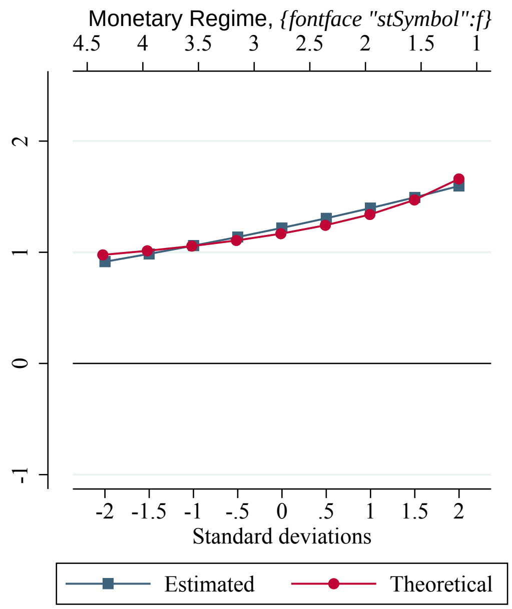

Notes: This chart shows how the cumulative fiscal multiplier (the cumulative effect on GDP divided by the cumulative increase in government spending) at 2 years varies with the monetary policy response both in our theoretical model and when the effect is estimated on simulated data. The red circles show the true theoretical variation by inflation coefficient ϕ (top horizontal axis). The blue squares show the empirical estimates obtained by using our Kitagawa-Blinder-Oaxaca decomposition estimates on data simulated from the model. The estimated monetary sensitivity parameter Θ_i is varied by 2 standard deviations (bottom horizontal axis). We simulate the model for 2000 periods, discarding the first 10%. For presentation, the theoretical results shown in the red line have been collapsed into nine bins to facilitate comparison with the nine points in the blue line. Moving from left to right on the horizontal axis implies a "less active" monetary policy which results in a larger multiplier.

Figure 3 repeats the red line with circles from Figure 1, which shows how the fiscal multiplier varies with the coefficient on inflation in the monetary policy rule (ϕ) in the model (now the top x-axis). The blue line with squares shows the Kitagawa-Blinder-Oaxaca decomposition-implied fiscal multiplier estimated on simulated data. The response of interest rates to the fiscal shock is estimated by "country" and this coefficient is used as a state variable. Thus the bottom horizontal axis refers to standard deviations of this object.[20] The figure shows that the empirical estimation with the KBO decomposition captures the monetary interaction in the simulated data very well.

3.3. Results

In this section we show the results from estimating the KBO specification Equation 9. To calculate the fiscal multiplier, two outcome variables are of interest: the cumulative percentage change in GDP, i.e. $y_{i,t+h} = {}^{(GDP_{i,t+h} - GDP_{i,t-1})}/_{GDP_{i,t-1}}$ and the cumulative change in the deficit (D) relative to initial

[20]The figure uses 2 standard deviations which captures well the range of different ϕ values in the simulated data. Note that, the KBO indirect effect estimates a non-linear function of the model's parameters so there is a monotonic but not one-to-one mapping between ϕ and the indirect effect estimated in the data.

Figure 4: *Policy experiments varying the response of monetary policy*

(a) *GDP Response (%)* **(b)** *Indirect Effect for GDP $h = 3$*

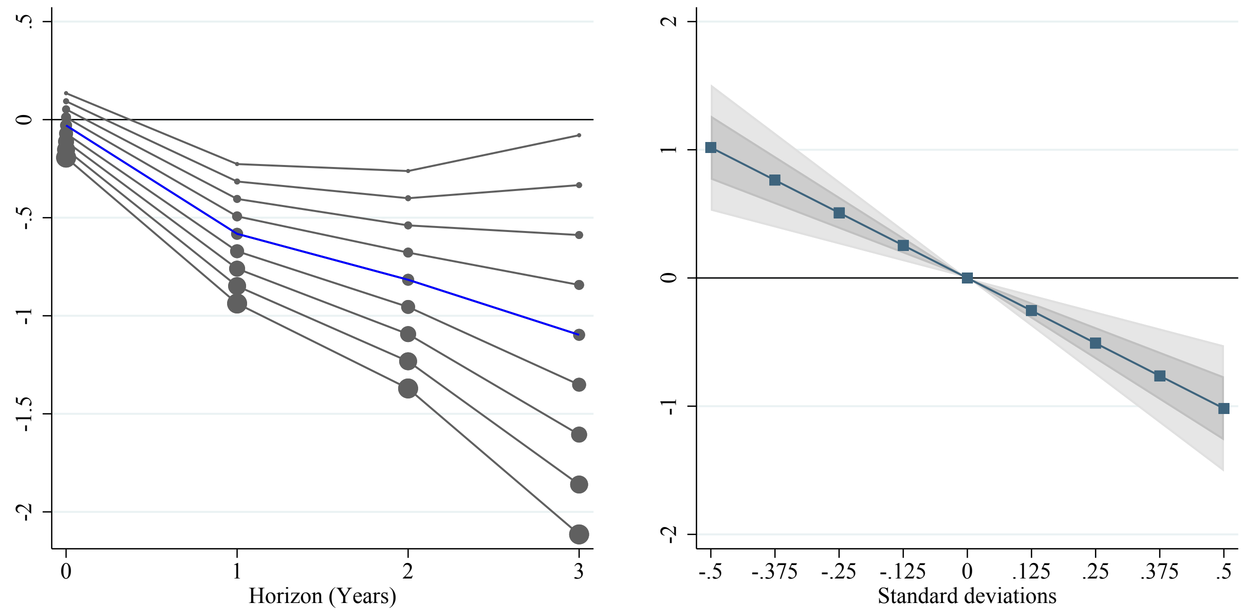

Notes: Panel (a) shows how the response of GDP varies with the degree of monetary policy accommodation. The blue lines report the direct effect, which should be compared to the average effect in Figure 2. The gray lines consider experiments which vary the degree of monetary accommodation. A larger marker indicates a tighter monetary policy scenario. Panel (b) plots the indirect effect on GDP for the peak effect at $h = 3$. The figure illustrates the effect of the monetary-fiscal interaction relative to the average multiplier in the full sample. This also allows us to formally test whether the indirect effect is statistically significant. The light and dark gray areas refers to a confidence interval of two and one standard deviations.

GDP, $y_{i,t+h} = {}^{(D_{i,t+h} - D_{i,t-1})}/_{GDP_{i,t-1}}$.[21] The β^h coefficients estimate the conventional impulse response function for the percentage change in the level of GDP or the deficit relative to GDP.

Figure 4 reports the main results from this exercise. Panel (a) shows the percentage response of GDP following a 1% of GDP fiscal consolidation (as measured by Guajardo, Leigh, and Pescatori, 2014). The central blue line in the fan reports the direct effect which, roughly, should be compared to the results from the linear model in Figure 2. As in Figure 2, GDP falls by around 1% over the course of 2–3 years.

To examine how the effect varies with monetary policy, the gray lines then conduct experiments where we vary the indirect effect estimated using the KBO decomposition from Equation 9. In particular, Figure 4 shows a range of scenarios where we vary Θ_i^h, the sensitivity of interest rates to fiscal policy. In keeping with the KBO formulation Θ_i^h — like all state variables — is expressed relative to its mean. The results therefore consider how the multiplier varies as we change the degree of monetary offset relative to the average degree of accommodation in the sample (captured in the direct effect).

In Figure 4 the size of the circular marker indicates a less accommodative (i.e. more contractionary) monetary policy. We vary Θ_i^h by one standard deviation, which produces real interest rate

[21]This ensures the ratio of these two IRFs can be interpreted as a multiplier, as discussed below.

variation of the order of 100bps on average over the period (see Appendix Figure A.4). In the face of a fiscal consolidation (a negative shock to GDP), a more muted real rate response is consistent with less monetary accommodation and a larger fiscal multiplier. This is indeed what is shown in Figure 4. As monetary policy becomes less accommodative, the multiplier becomes larger. Appendix Figure A.4 reports the associated figure for the response of the real interest rate. As expected tighter monetary policy is associated with less accommodation in terms of the real interest rate.

Panel (b) of Figure 4 shows the indirect effect on GDP for the peak effect at $h = 3$ and the standard errors.[22] This figure therefore shows the effect of the monetary-fiscal interaction relative to the average multiplier in the full sample (the central blue line in Panel (a)). This also allows us to formally test whether the indirect effect is statistically significant. The light and dark gray areas refer to a confidence interval of two and one standard deviations. As shown in the figure, for a less accommodative monetary policy, the negative effect on GDP is nearly 1% larger than in the baseline and this effect is statistically significant. In Table A.1 we report the precise coefficient estimates for β^h (the direct effect), θ_f^h (the strength of the indirect effect) and the standard errors at all horizons.[23]

The fiscal multiplier is typically defined as the \$ movement in GDP for a one \$ change in fiscal policy. Following Ramey (2016), this object can be computed empirically by estimating the effect on GDP and dividing by the associated change in the deficit relative to GDP. It is therefore instructive to also consider what happens to the deficit to GDP ratio to get a sense of the magnitude of the fiscal intervention in the data. The response of the deficit may also vary with the behavior of monetary policy, for example higher interest rates and lower demand could make it harder to reduce the deficit. The response of the deficit relative to GDP is shown in Figure A.2. A 1% fiscal consolidation (as measured by Guajardo, Leigh, and Pescatori (2014)) takes some time to have its full effect. The deficit to GDP ratio moves by around 0.5% in the current year, and is around 1% lower from the following year onwards. This path also depends on monetary policy, although in these experiments, there is not much state-dependence in the deficit to GDP ratio until the later years.

Figure 5 converts the state-dependent IRFs from Figure 4 and Figure A.2 into cumulative multipliers at different horizons. The cumulative multiplier is reported on the y-axis. As before, the x-axis varies Θ_i^h from -0.5 to $+0.5$ standard deviations. The three lines report the multiplier at different horizons. Note that the 0 point on the x-axis corresponds to the average treatment effect usually estimated in linear models. Interestingly, this is around or below 1 at *all* horizons. As monetary policy becomes more inert (rates are cut less aggressively in the face of falling demand), the multiplier rises. In these experiments the multiplier varies from around zero to nearly 2. Thus, in any fiscal intervention, the fiscal multiplier crucially depends on the monetary response. Interestingly, magnitudes around 2 are close to Keynes' original prediction of 2.5 (Keynes, 1936).

To highlight the importance of identification, Figure A.5 shows that a very different conclusion

[22]Standard errors computed using Monte Carlo, applied to both estimation of the proxy and second stage KBO decomposition. Standard errors capture the fact that Θ_i is a generated regressor.

[23]For presentational reasons, panel (a) of Figure 4 does not display the standard errors but Figure A.3 also visually shows that the direct effect (the blue line) is statistically significant.

Figure 5: *Cumulative fiscal multiplier by monetary response*

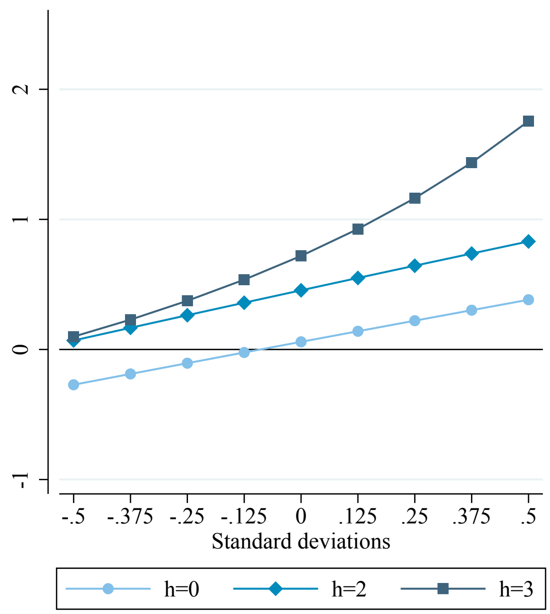

Notes: This chart shows the cumulative fiscal multiplier from each scenario in Figure 4. This is computed as the cumulative sum of the GDP response relative to the cumulative deficit to GDP response (based on Figure A.2). Each line refers to a different horizon, h. As in Figure 4, h goes from the current year $h = 0$ to the third year after the shock $h = 3$. $h = 1$ is omitted to avoid overcrowding the figure. Moving from left to right on the horizontal axis implies a "less active" monetary policy which results in a larger multiplier.

would be reached using the actual future change in the policy in x instead of using our sensitivity-based proxy. Without identification, the importance of the monetary offset is greatly reduced.

3.4. Decomposition along other dimensions

Our regressions contain a number of other state variables and the KBO decomposition allows us to consider how the fiscal multiplier varies along each of these dimensions while controlling for the other states.[24] Figure 6 shows how the multiplier varies according to each of the other macro controls in our regressions, holding the other variables constant. The variables are the output gap, the change in the fiscal deficit to GDP ratio, World GDP growth and the size of the fiscal consolidation.[25]

Figure 6 shows that along each of these dimensions, once we control for the other variables simultaneously, there is only sizable state dependence by the size of the output gap. This echos

[24] This is a type of selection-on-observables strategy and, of course, there still remains a question about whether these variables can be moved around assuming "all else equal". Still, these dimensions of heterogeneity are important and common in the literature. It is therefore interesting to revisit these in our framework. In the literature, state dependence is often investigated by considering one dimension at a time and with a variety of empirical specifications although, as discussed above, typical macro variables that are often used to define the state are likely to be highly correlated.

[25] Our main regression also includes GDP growth. The results are similar to those using the output gap and are thus omitted for brevity.

Figure 6: *Other forms of state dependence in the fiscal multiplier*

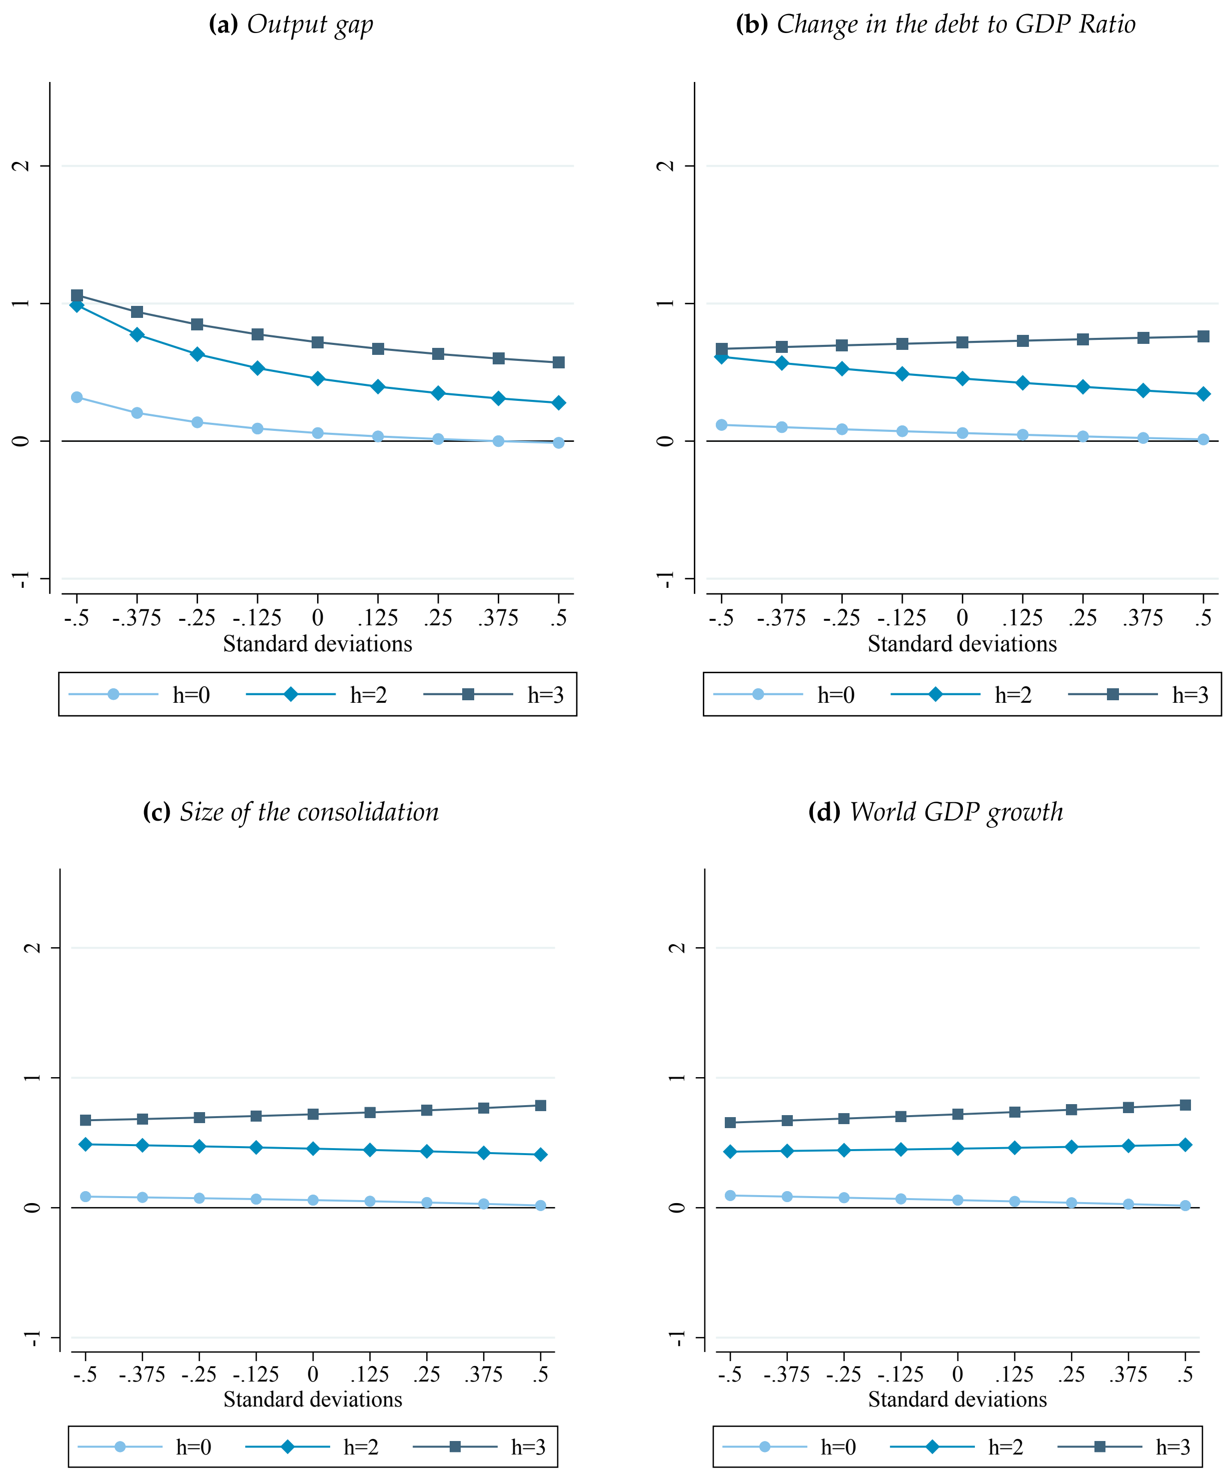

Notes: This chart shows how the cumulative fiscal multiplier varies with the other state variables in our regressions. As before, the multiplier is computed as the cumulative sum of the GDP response relative to the cumulative deficit to GDP response. Each line refers to a different horizon, h. As in Figure 4, h goes from the current year $h = 0$ to the third year after the shock $h = 3$. $h = 1$ is omitted to avoid overcrowding the figure. Panel (a) shows variation in the multiplier depending on the size of the (lagged) output gap, Panel (b) is for difference changes in the (lagged) deficit to GDP ratio, Panel (c) varies the size of the fiscal consolidation and Panel (d) varies World GDP growth.

findings in some other papers, such as Auerbach and Gorodnichenko (2012) and Jordà and Taylor (2016), that fiscal multipliers may be larger in periods of above-average slack. To the extent that a large change in the deficit to GDP ratio is associated with fiscal stress, our results do not suggest a smaller multipliers in these states.

In principle the lagged state variables in $\boldsymbol{x}$ are correlated with each other. How important is this concern for typical state dependence applications? To examine this question, Figure 7 shows the difference between our baseline specification (left column, which repeats Figure 6) and an alternative specification where a single x variable is included in the interaction term. The second column therefore comes from a specification that is closer to typical ways of estimating state dependence. The figure illustrates the challenges of a "single cut" approach for lagged endogenous variables. For example, the first row shows that if we only consider heterogeneity according to the change in the deficit to GDP ratio, we could erroneously conclude that fiscal multipliers are much larger in when deficits are large. In fact, once we allow all the control variables to enter as interactions, there is little heterogeneity by the deficit to GDP ratio. In reality, fiscal treatment states that see large increases in the deficit are also states in which the output gap is more negative. The flip side of this is that state dependence by the output gap tends to *increase* when we condition on all other dimensions of heterogeneity, as shown in the third row.

3.5. Time variation in the multiplier

Previous sections have shown how the KBO decomposition can be used to stratify impulse responses with state variables. This section shows how to calculate variation in the impulse response throughout the sample as a function of variation in the conditioning information set. The three-dimensional figure showing the variation of the impulse response over time is shown in Figure 8.

The figure is constructed as follows. We use a GDP-weighted average of all the controls in $\boldsymbol{x}$ and then simulate a 1% of GDP fiscal consolidation. Thus the figure traces the fiscal multiplier based on average conditions across the countries in our sample.

Several remarks are worth noting. First, the figure illustrates the complexity embodied in the dependence of the fiscal multiplier on a large set of controls. Though we have examined how the fiscal multiplier varies when monetary policy is more or less accommodating, when the output gap is positive or negative and so on, it would be impossible to combine that information as succinctly as Figure 8 does. Second, the figure shows that a policymaker interested in evaluating the costs and benefits of a fiscal consolidation would be well-advised to consider the current macroeconomic outlook. Figure 8 shows that the multiplier has fluctuated between as high as 3.5 and as low as 0. Third, the effect of the fiscal consolidation across years 1-3 is fairly stable, that is, the same forces inducing variation in the multiplier in one year, induce the same variation in subsequent years. Fourth, on average, the fiscal multiplier is approximately 1 across the entire sample.

Figure 7: *Baseline Results vs. Single Dimension State Dependence*

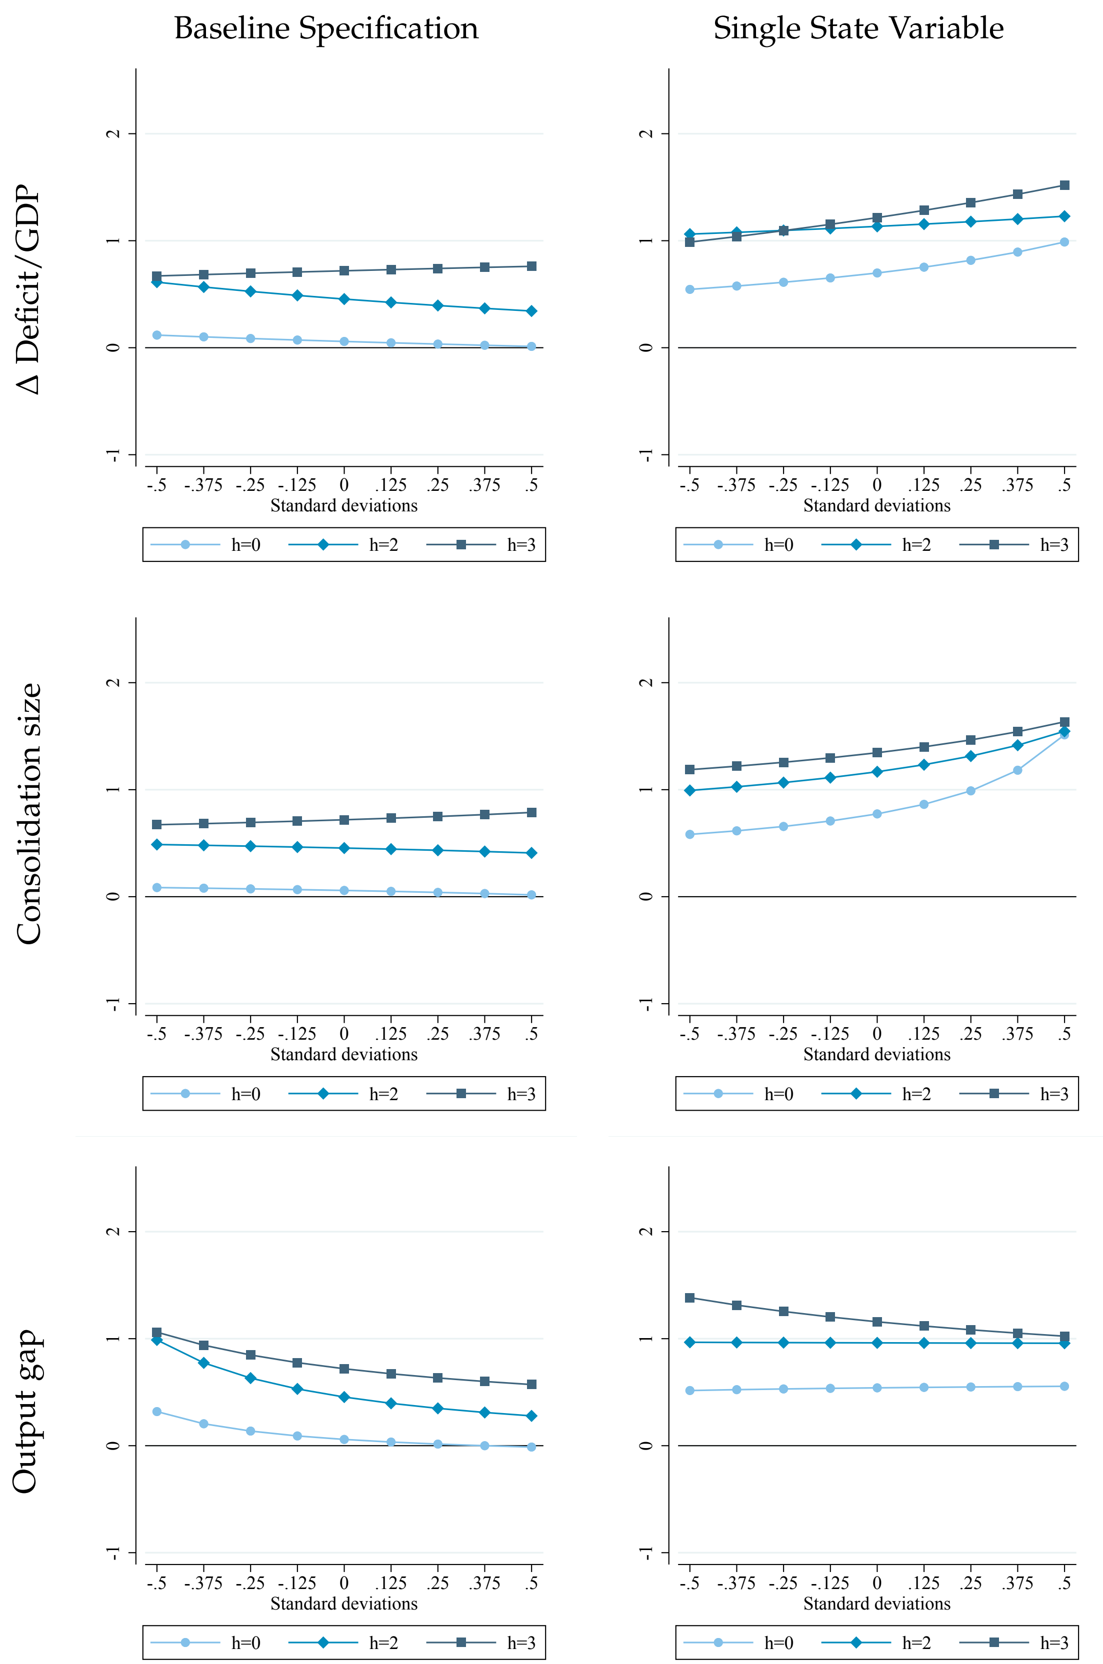

Notes: This chart compares state dependence in the fiscal multiplier from the baseline multivariate specification with a specification that only includes each state variable in isolation. The first column repeats the baseline results from Figure 6. The second column re-estimates the main regression but only includes the state variable of interest in the interaction term.

Figure 8: *Time variation in the fiscal multiplier*

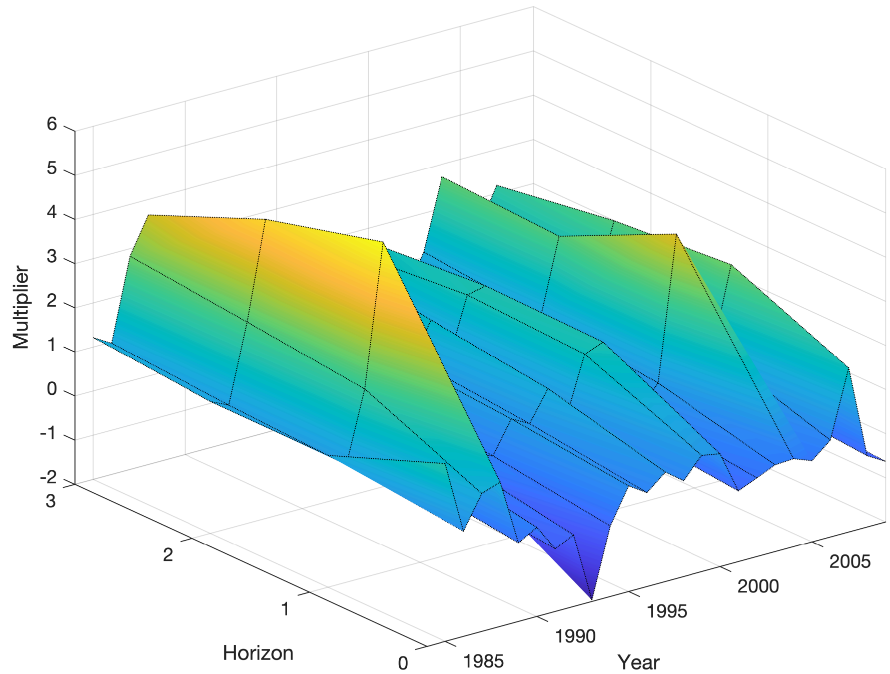

Notes: This chart shows time variation in the cumulative fiscal multiplier at each horizon of the IRF. This chart is constructed by evaluating the multiplier at horizon h given the state of the economy in each year of the sample. Given the lag structure, this figure shows time variation from 1982 to 2009.

4. Robustness and extensions

In this section we subject our approach to several robustness checks and extensions.

4.1. Relaxing the identification assumptions

The previous section considered a straightforward and easy to implement approach to constructing a proxy for the nature of the monetary response. This was based on the sensitivity of policy rates to fiscal treatment. In this section, we consider a more general approach that relaxes the identification assumptions made in the simpler approach. Rather than constructing a proxy by regressing policy rates on the fiscal shock directly, this section makes use of a more conventional Taylor Rule-type approach.

To motivate this further approach, and fix ideas, consider a policy rule for the interest rate R in terms of lagged interest rates and inflation π_t.

$$R_t = \rho_i R_{t-1} + (1 - \rho_i)\, \phi^i \pi_t\,, \tag{11}$$

The variation we are interested in capturing is the idea that the policy rule may vary across country

Figure 9: *The cumulative fiscal multiplier by monetary response: more general approach*

(a) *True theoretical versus estimates*

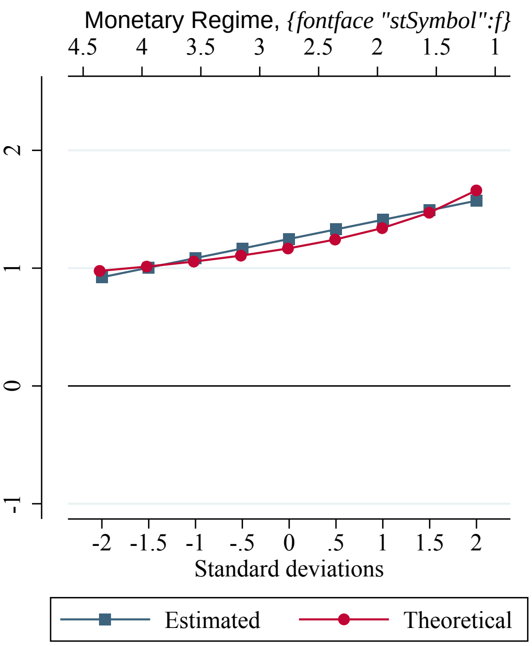

(b) *Empirical fiscal multiplier by monetary response*

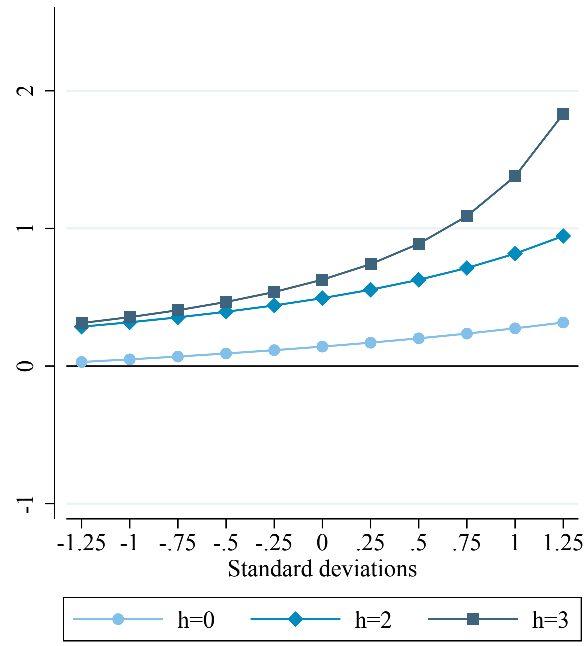

Notes: Panel (a) shows how the cumulative fiscal multiplier at 2 years varies with the monetary policy response both in theory and when the effect is estimated on simulated data. The red circles show the true theoretical variation by inflation coefficient (top x-axis). The blue squares show the empirical estimates obtained by using our KBO decomposition estimates on data simulated from the model. The estimated monetary sensitivity parameter Θ_i is varied by 2 standard deviations (bottom x-axis). For presentation, the theoretical results shown in the red line have been collapsed into nine bins to facilitate comparison with the nine points in the blue line. This figure is produced using an extended version of the model where the other structural parameters of the model also vary by "country". See main text for details. Panel (b) shows the cumulative fiscal multiplier varying the degree of monetary offset. This is computed as the cumulative sum of the GDP response relative to the cumulative deficit to GDP response. Each line refers to a different horizon, h. As in Figure 4, h goes from the current year $h = 0$ to the third year after the shock $h = 3$. $h = 1$ is omitted to avoid overcrowding the figure. Relative to the baseline figure in the main text, this is produced using the alternative proxy as discussed in Section 4. Note that the standard deviation of the two proxies (the new proxy and the baseline method) are different. The experiment here is therefore calibrated so the real interest rate varies on impact by a similar amount to the baseline figure.

i through variation in ϕ^i.

Note that, as discussed in the previous section, we are not trying to identify all the specific individual parameters of a Taylor Rule. Instead, all we need is to obtain a proxy for the sensitivity of interest rates to the economy where the ranking across countries is correctly captured.

We can therefore estimate a more reduced form expression,

$$R_t = \alpha_l^i R_{t-1} + \alpha_\pi^i \pi_t . \tag{12}$$

In particular, we estimate a variant of Equation 10 but where the $f_{i,t}$ variable is replaced with inflation. In keeping with the Taylor Rule estimation literature, we then instrument inflation with its lag. The coefficient Θ_i is then a proxy for the cross-country sensitivity of interest rates to inflation.

To illustrate how this works, the first panel of Figure 9 implements this approach with simulated data from the theoretical model. In the model we allow for heterogeneity in ϕ^i, but also allow for heterogeneity in the other structural parameters of the model. In particular, the model now features "cross-country" heterogeneity in the share of non-savers, the degree of price stickiness, the persistence of the fiscal shock and the Frisch elasticity. Panel (a) in Figure 9 is the counterpart of Figure 3 and shows that even in this more general environment our KBO approach still correctly recovers how the fiscal multiplier varies with the monetary offset.

The second panel of Figure 9 then conducts the same experiment in the data. The results are similar to the baseline findings. The drawback of this more flexible method is that it requires making assumptions about the arguments of the Taylor Rule, so it is not as transparent or easy to implement.[26] As a result, we see both approaches as useful for examining the strength of the monetary offset in the data.

To conclude this section, it is worth noting the flexibility and general applicability of our approach. In principle, these methods could be used to study the role of other sources of heterogeneity in the fiscal multiplier. The challenge, of course, would be to find good proxies for those characteristics that could be used for identification. This seems an interesting avenue for future work.

4.2. Tax- versus spending-led consolidations

It is possible that countries differ in the composition of the fiscal consolidation. For example, some countries may rely more on tax increases than spending cuts. The fiscal multiplier literature has often noted differences in spending versus tax multipliers. Furthermore, Guajardo, Leigh, and Pescatori (2014) find that tax-based consolidations are more contractionary.

This could affect our results in the following way. Suppose, for example, that tax multipliers are larger than those for spending (for reasons unrelated to monetary policy, as is the case in some macro models). Different policies might then induce different relative movements in GDP and interest rates. If countries differ in their *average* reliance on tax increases versus spending cuts, this could conceivably be captured in the Θ_i in our simple sensitivity proxy approach. It should not, however, bias the Taylor-Rule approach provided that the degree of monetary activism is uncorrelated with the fiscal authorities preferences for adjusting taxes versus spending.

Still, the flexibility of the KBO specification offers a further way to control for this possibility. Specifically, we construct a country-specific measure of the average propensity to use tax increases versus spending cuts.[27] We then interact this cross-country characteristic with the fiscal treatment $f_{i,t}$, essentially adding it as an additional KBO state variable. The *residual* variation in Θ_i is then being used to examine the monetary offset.[28]

[26] As an extension we considered a factor approach where inflation above is replaced by the first principal component of inflation and the output gap. This is one way to incorporate more arguments in the rule but estimate a single parameter to act as the proxy for the nature of the monetary response. Again the results are similar.

[27] In particular, we calculate the share of consolidation episodes that are tax-led by country.

[28] Guren, McKay, Nakamura, and Steinsson (2020) follow a similar logic to focus on residual variation in their sensitivity instrument, albeit in a different setting where the variation of concern is a time-region effect. The strategy used above could also be applied to rule out other cross-country concerns, although these types

Figure 10: *Policy experiments: time fixed effects*

(a) *Response of GDP (%)* **(b)** *Cumulative fiscal multiplier by monetary response*

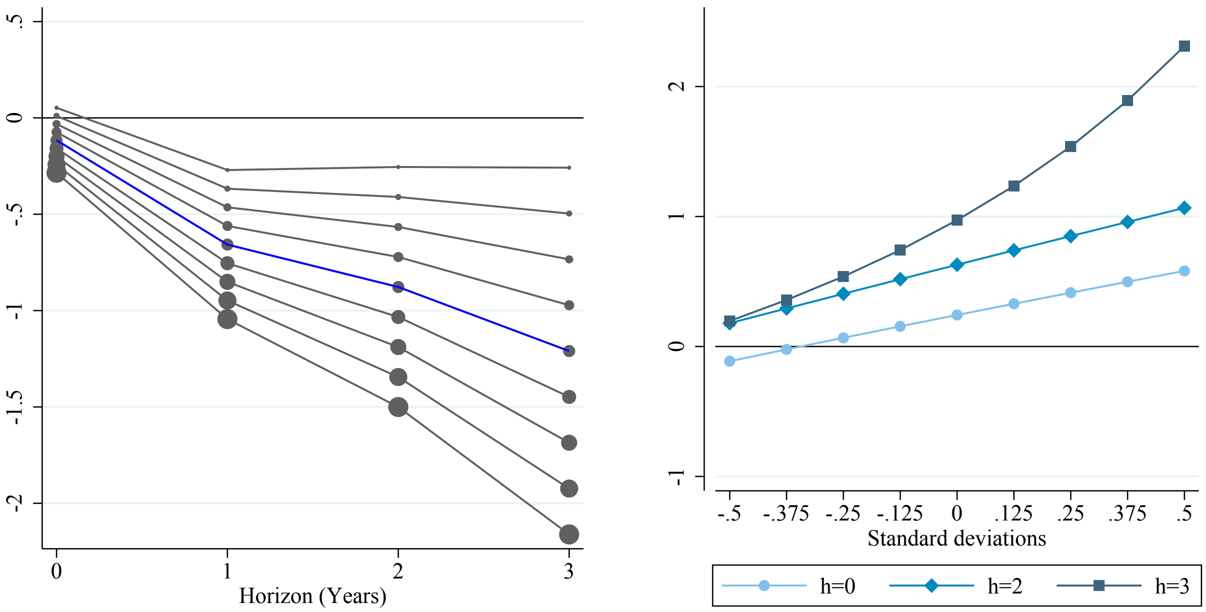

Notes: Panel (a) shows how the response of GDP varies with the degree of monetary policy accommodation. The blue lines report the direct effect. The gray lines consider experiments which vary the degree of monetary accommodation. A larger marker indicates a tighter monetary policy scenario. Panel (b) reports the associated cumulative fiscal multiplier. Relative to baseline Figure 4 and Figure 5, the specification in these figures include time fixed effects rather than world GDP growth.

The results of this exercise are shown in Figure A.6. The estimated monetary offset is very similar to the baseline case.

4.3. Global factors

In the baseline specification we included world GDP growth to capture time varying global factors. In the original Guajardo, Leigh, and Pescatori (2014) paper, the authors use time fixed effects as a more general way of capturing global factors. Earlier we noted that this seems to come at the cost of precision in our specification, but in this section we re-estimate our main results for the monetary-fiscal multiplier using time fixed effects rather than world GDP growth.

The results are shown in Figure 10. These figures are very similar to the baseline specification in Figure 4 and Figure 5. Our use of world GDP growth does not, therefore, affect our main findings.

4.4. Openness

A number of papers in the open economy macroeconomics literature have noted that the fiscal multiplier may vary with the degree of openness (e.g. Ilzetzki, Mendoza, and Végh (2013) find that

of issues are also dealt with in a more general sense by the approach in Section 4.1.

Figure 11: *Policy experiments: controlling for trade openness*

(a) *Response of GDP (%)* **(b)** *Cumulative fiscal multiplier by monetary response*

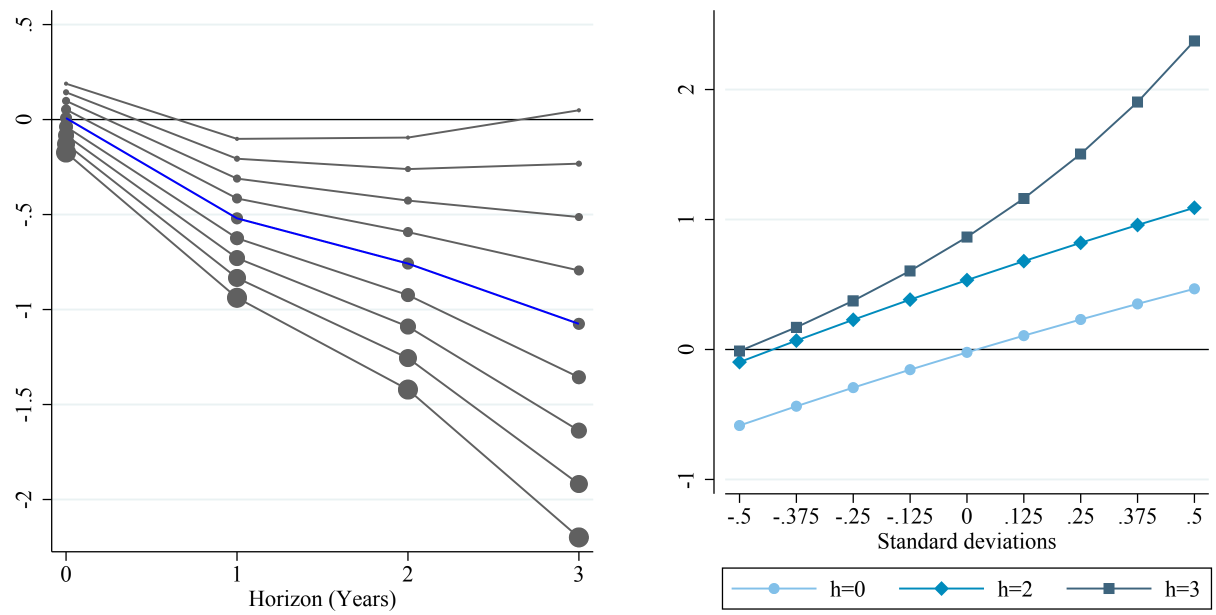

Notes: Panel (a) shows how the response of GDP varies with the degree of monetary policy accommodation. The blue lines report the direct effect. The gray lines consider experiments which vary the degree of monetary accommodation. A larger marker indicates a tighter monetary policy scenario. Panel (b) reports the associated cumulative fiscal multiplier. Relative to baseline Figure 4 and Figure 5, the specification in these figures controls for the influence of openness.

more open economies tend to have smaller fiscal multipliers). To what extent are our main results driven by variations in openness across countries? Note that, if our sensitivity proxy is correctly ranking countries based on the degree of monetary activism, variations in openness should not bias our findings. But one may be concerned that variation in openness is correlated with our monetary proxy, therefore confounding the conclusions.

The flexibility of our approach is that we can easily check this by including trade openness as an additional state variable. We measure openness as exports plus imports relative to GDP and include this in the vector of state variables x. The results are shown in Figure 11. The GDP and multiplier plots are very similar to the baseline case, and the multiplier varies from around 0 to 2.

4.5. Lag structure

If the fiscal shocks reflect purely random variation, the choice of additional controls should not affect the main set of estimates. In small samples however, serial correlation could potentially be an issue. As a further robustness check we show that the main results are not overturned by using a slightly longer lag structure for the controls. In the baseline results we chose two years of lags. Note that, relative to standard empirical papers using quarterly data, this already controls for a sizable degree of persistence. We also face a trade-off in that longer lag structures lead to loss of data and

Figure 12: *Policy experiments: longer lag structure*

(a) *Response of GDP (%)* **(b)** *Cumulative fiscal multiplier by monetary response*

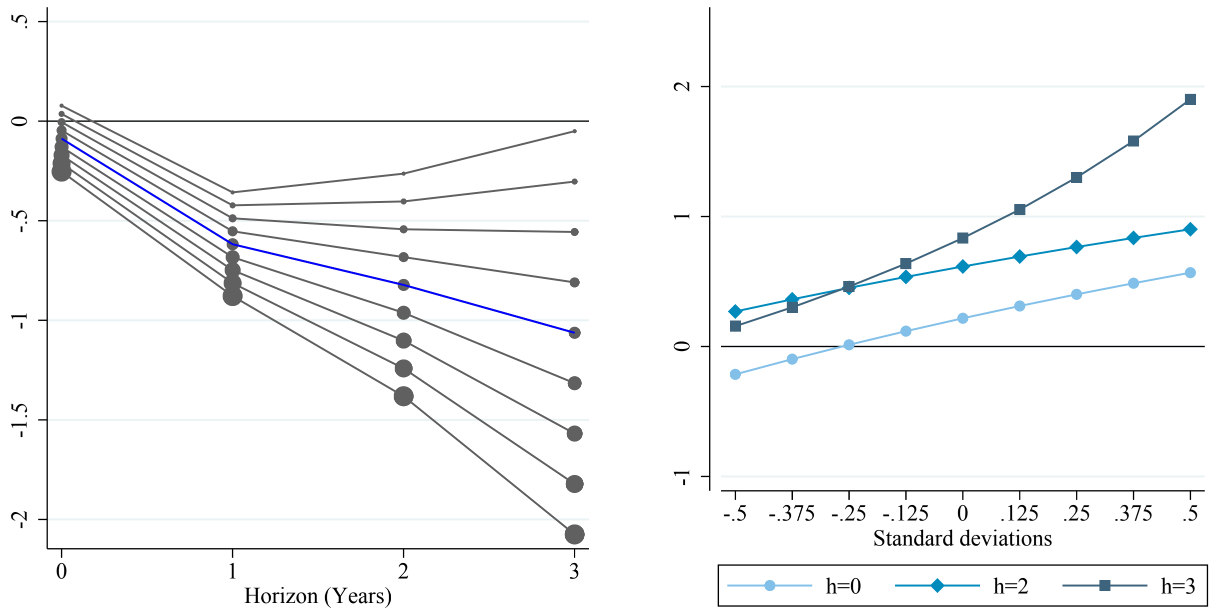

Notes: Panel (a) shows how the response of GDP varies with the degree of monetary policy accommodation. The blue lines report the direct effect. The gray lines consider experiments which vary the degree of monetary accommodation. A larger marker indicates a tighter monetary policy scenario. Panel (b) reports the associated cumulative fiscal multiplier. Relative to baseline Figure 4 and Figure 5, the specification in these figures include 3 lags of all controls in x.

more parameters to be estimated.

That said, we re-run our main results using three years of lags (equivalent, of course, to 12 quarters of lags in typical macro papers). Figure 12 shows that the results are very similar to the baseline findings in Figure 4 and Figure 5.

4.6. Monetary-fiscal interactions using shocks

To further corroborate the magnitudes found above, in this section we consider a different approach to studying monetary-fiscal interactions. Instead of relying on variation in the response of interest rates to fiscal policy across countries for identification, here we use an approach based on monetary policy shocks.

To motivate the approach consider the following simple motivating setup:[29]

$$y_{i,t} = \delta_f f_{i,t} + \delta_r r_{i,t} + u^y_{i,t}, \tag{13}$$

$$r_{i,t} = \Theta_y y_{i,t} + \Theta_f f_{i,t} y_{i,t} + u^r_{i,t}. \tag{14}$$

Here $y_{i,t}$ denotes real GDP growth and $r_{i,t}$ denotes the real rate. For simplicity, assume the fiscal

[29]A related setup is considered in the Appendix, which illustrates the baseline approach used earlier.

intervention is binary with $f_{i,t} \in [0, 1]$ and that the monetary authority sets the real interest rate directly. In this case, the sensitivity of interest rates to fiscal policy does not vary across country, but it does vary with the type of shock. During episodes of fiscal treatment, the monetary authority may respond to output fluctuations differently than in other periods. In the formulation above, this is captured by the Θ_f term, which is only relevant in periods of fiscal treatment. Note that, when there is no fiscal treatment, $f_{i,t} = 0$.

We can combine these expressions to create a reduced form equation. Given the binary nature of this example, we can then inspect the reduced form in the case of treatment, $f_{i,t} = 1$ and no treatment, $f_{i,t} = 0$. The resulting equation for estimation can be written as:

$$y_{i,t} = \beta_f f_{i,t} + \beta_r u^r_{i,t} + \beta_{rf} u^r_{i,t} f_{i,t} + u^y_{i,t}, \tag{15}$$

where

$$\beta_r = \frac{\delta_r}{1 - \delta_r \Theta_y}, \qquad \beta_f = \frac{\delta_f}{1 - \delta_r (\Theta_y + \Theta_f)}, \qquad \beta_{rf} = \frac{\delta_r}{1 - \delta_r (\Theta_y + \Theta_f)} - \beta_r .$$

The third term on the right hand side of Equation 15 captures the indirect effect from the interaction of monetary and fiscal policy. The degree of monetary policy accommodation is captured by the size of the monetary shocks $u^r_{i,t}$ (since these capture the policy stance relative to what would have been expected given the rule). The indirect effect captures the fact that, less accommodative monetary policy may translate into a larger recession during periods of fiscal treatment.

In estimating an equation of the form of equation 15 the technical challenge is that we do not observe $u^r_{i,t}$ directly and commonly constructed proxies for $u^r_{i,t}$ are usually only available for countries like the United States. To our knowledge, there is no consistent cross-country dataset of monetary policy shocks. In this section, as a robustness check, we therefore rely on a simple approach to validate the results in the previous section.

First, using a panel ordered probit, we predict the probability of observing an interest rate change based on two lags of GDP growth, inflation, the lagged change in the policy rate, and world GDP growth. We are implicitly assuming a common policy rule across countries in this exercise. Monetary policy shocks are then constructed as follows,

$$u^r_{i,t} = \Delta R_t - (p_{-1} \times -1 + p_0 \times 0 + p_{+1} \times 1),$$

where R_t is the nominal policy rate and the p terms are the predicted probabilities of a rate cut, no change or an increase. This approach therefore attempts to remove the predictable component of monetary policy. As in the previous section, the KBO decomposition is estimated from the following regression,

$$y_{i,t+h} = \mu^h_0 + (\boldsymbol{x}_{i,t} - \bar{\boldsymbol{x}}_i)\, \gamma^h_{0,x} + u^r_{i,t+h}\, \gamma^h_{0,r} + f_{i,t}\, \beta^h + f_{i,t}\, (\boldsymbol{x}_{i,t} - \bar{\boldsymbol{x}}_i)\, \boldsymbol{\theta}^h_x + f_{i,t}\, u^r_{i,t+h}\, \theta^h_r + \omega_{i,t+h} . \tag{16}$$

The main difference from the previous section is that the future stance of monetary policy during

Figure 13: *Policy experiments: an alternative approach*

(a) *Response of GDP (%)* **(b)** *Cumulative fiscal multiplier by monetary response*

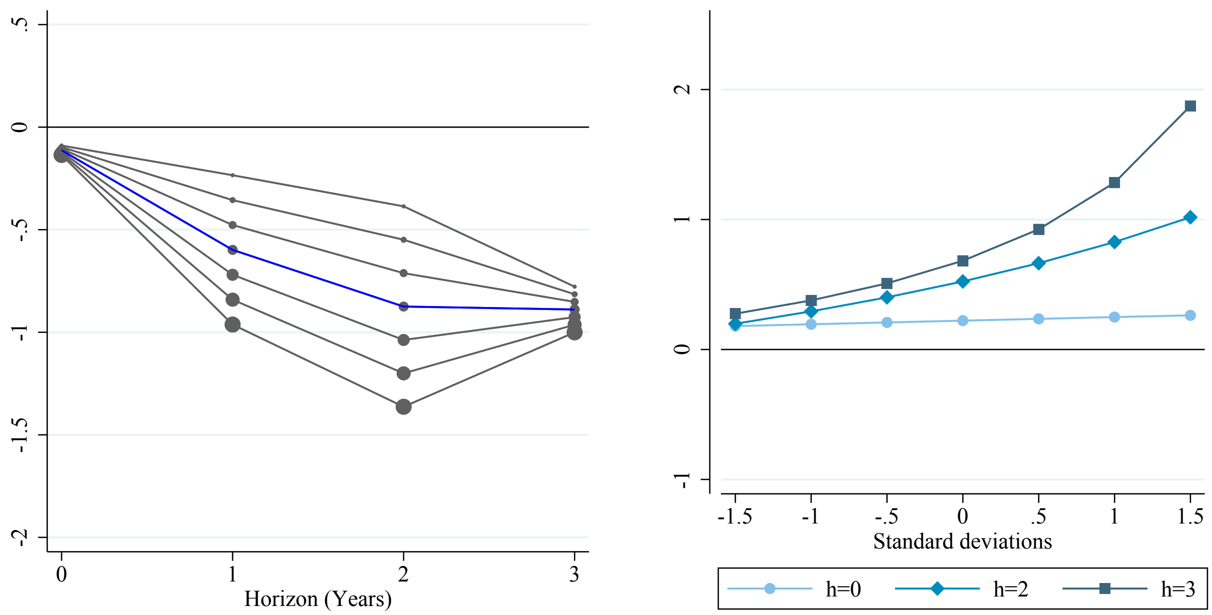

Notes: Panel (a) shows how the response of GDP varies with the degree of monetary policy accommodation. The blue lines report the direct effect. The gray lines consider experiments which vary the degree of monetary accommodation. A larger marker indicates a tighter monetary policy scenario. Panel (b) reports the associated cumulative fiscal multiplier. Relative to baseline Figure 4 and Figure 5, this figure is produced using an alternative approach to monetary-fiscal interactions, as discussed in the text.

the consolidation episode is captured by the deviation of the policy from what was expected, i.e., the shock term $u^r_{i,t+h}$.

Figure 13 shows the results. Once again, the first panel shows the effect on GDP on average (blue line), and for tighter and looser monetary policy during the consolidation episode (gray lines). We consider experiments from −1.5 standard deviation shocks to +1.5 standard deviation shocks. We use a wider range for this experiment as a one-standard deviation shock produces smaller variations in interest rates. Episodes with tighter monetary policy are associated with a much larger fall in GDP. In Figure A.7, the deficit to GDP ratio also improves by less in these more contractionary episodes. In Figure 13 Panel (b), we therefore report the cumulative fiscal multiplier. The multiplier rises to nearly 2 when monetary conditions are tight.

5. Conclusion and policy implications

This paper has shown that, using the Kitagawa-Blinder-Oaxaca decomposition from applied microeconomics in a local projections framework, the traditional macroeconomic impulse response can be decomposed into (1) the *direct* effect of the intervention on the outcome; (2) the *indirect* effect due to changes in how other covariates affect the outcome when there is an intervention; and (3) a *composition* effect due to differences in covariates between treated and control subpopulations. This provides a unified framework for decomposing the dynamic average treatment effects contained in impulse responses and, hence, for evaluating policy-dependence, state-dependence, and the balance conditions for identification in a multi-dimensional setting.

A natural application of these ideas is in the area of monetary-fiscal interactions. The fiscal multiplier is a key statistic for understanding how fiscal policy changes might stimulate or contract the macroeconomy. The size of the multiplier has been a subject of intensive debate since the Global Financial Crisis in 2008. But, despite the importance of this object, there is still much disagreement about existing empirical estimates. A large literature has focused on identifying the direct effect. Our paper tackles a more conceptual problem: there is no such thing as *the* fiscal multiplier in the data. One of the most obvious reasons is that monetary policy may not offset the effects of fiscal policy in the same way across time or across countries. We show that the Kitagawa-Blinder-Oaxaca decomposition provides a natural way to try to disentangle these effects and decompose the drivers of variation in the multiplier.

Our main result is that fiscal multipliers can be large when monetary policy is less activist. This accords with conventional wisdom and the mechanism can be found in many models. In our experiments, fiscal multipliers can be as low zero or as high as 2 and above, depending on the offsetting actions of the monetary authority. This has important implications for measuring "the multiplier" and for evaluating and predicting the likely effects of particular macro-policy interventions.

The Kitagawa-Blinder-Oaxaca decomposition we propose also has wider implications for measuring the effects of all kinds of policy treatments in macroeconomics, and can allow for many other possible dimensions of heterogeneity in a very flexible way. Using our decomposition approach, the tasks of estimation and inference can be easily undertaken by using standard linear regression methods while still being sufficiently general to allow for a great deal of unspecified state dependence and time-variation. We therefore hope these techniques will be of use to all researchers interested in the study of state-dependent, non-linear, and time-varying effects of policy interventions more generally.

References

Acconcia, Antonio, Giancarlo Corsetti, and Saverio Simonelli. 2014. Mafia and Public Spending: Evidence on the Fiscal Multiplier from a Quasi-experiment. *American Economic Review* 104(7): 2185–2209.

Alpanda, Sami, Eleonora Granziera, and Sarah Zubairy. 2021. State dependence of monetary policy across business, credit and interest rate cycles. *European Economic Review* 140: 103936.

Alpanda, Sami, and Sarah Zubairy. 2019. Household Debt Overhang and Transmission of Monetary Policy. *Journal of Money, Credit and Banking* 51(5): 1265–1307.

Angrist, Joshua D., Òscar Jordà, and Guido M. Kuersteiner. 2018. Semiparametric Estimates of Monetary Policy Effects: String Theory Revisited. *Journal of Business & Economic Statistics* 36(3): 371–387.

Auerbach, Alan J., and Yuriy Gorodnichenko. 2012. Measuring the Output Responses to Fiscal Policy. *American Economic Journal: Economic Policy* 4(2): 1–27.

Bachmann, Rüdiger, and Eric R. Sims. 2012. Confidence and the transmission of government spending shocks. *Journal of Monetary Economics* 59(3): 235–249.

Banerjee, Ryan, and Fabrizio Zampolli. 2019. What drives the short-run costs of fiscal consolidation? Evidence from OECD countries. *Economic Modelling* 82: 420–436.

Barnichon, Régis, Davide Debortoli, and Christian Matthes. 2020. Understanding the size of the government spending multiplier: It's in the sign. Economics Working Papers 1688, Department of Economics and Business, Universitat Pompeu Fabra.

Barro, Robert J., and Charles J. Redlick. 2011. Macroeconomic Effects From Government Purchases and Taxes. *Quarterly Journal of Economics* 126(1): 51–102.

Blanchard, Olivier, and Roberto Perotti. 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.

Blinder, Alan S. 1973. Wage Discrimination: Reduced Form and Structural Estimates. *Journal of Human Resources* 8(4): 436–455.

Born, Benjamin, Francesco D'Ascanio, Gernot J. Müller, and Johannes Pfeifer. 2021. Mr. Keynes meets the Classics: Government Spending and the Real Exchange Rate. ifo Working Paper 352.

Born, Benjamin, Falko Juessen, and Gernot J. Müller. 2013. Exchange rate regimes and fiscal multipliers. *Journal of Economic Dynamics and Control* 37(2): 446–465.

Canova, Fabio, and Evi Pappa. 2011. Fiscal Policy, Pricing Frictions and Monetary Accommodation. *Economic Policy* 26(1): 555–598.

Chodorow-Reich, Gabriel. 2019. Geographic Cross-Sectional Fiscal Spending Multipliers: What Have We Learned? *American Economic Journal: Economic Policy* 11(2): 1–34.

Christiano, Lawrence, Martin Eichenbaum, and Sergio Rebelo. 2011. When is the government spending multiplier large? *Journal of Political Economy* 119(1): 78–121.

Cloyne, James. 2013. Discretionary Tax Changes and the Macroeconomy: New Narrative Evidence from the United Kingdom. *American Economic Review* 103(4): 1507–1528.

Cloyne, James, Nicholas Dimsdale, and Natacha Postel-Vinay. 2018. Taxes and Growth: New Narrative Evidence from Interwar Britain. NBER Working Paper 24659.

Cloyne, James S., and Paolo Surico. 2017. Household Debt and the Dynamic Effects of Income Tax Changes. *Review of Economic Studies* 84(1): 45–81.

Cogley, Timothy, and Thomas J Sargent. 2005. Drifts and volatilities: monetary policies and outcomes in the post WWII US. *Review of Economic dynamics* 8(2): 262–302.

Corbi, Raphael, Elias Papaioannou, and Paolo Surico. 2019. Regional Transfer Multipliers. *Review of Economic Studies* 86(5): 1901–1934.

Corsetti, Giancarlo, André Meier, Gernot J. Müller, and Michael B. Devereux. 2012. What determines government spending multipliers? *Economic Policy* 27(72): 521–565.

Crafts, Nicholas, and Terence C. Mills. 2013. Rearmament to the Rescue? New Estimates of the Impact of "Keynesian" Policies in 1930s' Britain. *Journal of Economic History* 73(4): 1077–1104.

DeLong, J. Bradford, and Lawrence H. Summers. 2012. Fiscal Policy in a Depressed Economy. *Brookings Papers on Economic Activity* 2012(1): 233–297.

Eggertsson, Gauti B. 2011. What Fiscal Policy is Effective at Zero Interest Rates? In *NBER Macroeconomics Annual 2010*, edited by Acemoglu, Daron, and Michael Woodford, volume 25, 59–112. Chicago: University of Chicago Press.

Ferrière, Axelle, and Gaston Navarro. 2020. The Heterogeneous Effects of Government Spending : It's All About Taxes. International Finance Discussion Papers 1237, Board of Governors of the Federal Reserve System.

Fortin, Nicole, Thomas Lemieux, and Sergio Firpo. 2011. Decomposition Methods in Economics. In *Handbook of Labor Economics*, edited by Ashenfelter, O., and D. Card, volume 4, chapter 1, 1–102. Amsterdam: Elsevier.

Galí, Jordi. 2015. *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework and Its Applications*. Princeton, N.J.: Princeton University Press, 2nd edition.

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

Ghassibe, Mishel, and Francesco Zanetti. 2022. State dependence of fiscal multipliers: the source of fluctuations matters. *Journal of Monetary Economics* 132: 1–23.

Gonçalves, Sílvia, Ana María Herrera, Lutz Kilian, and Elena Pesavento. 2022. When do state-dependent local projections work? .

Guajardo, Jaime, Daniel Leigh, and Andrea Pescatori. 2011. Expansionary Austerity: New International Evidence. IMF Working Paper WP/11/158.

Guajardo, Jaime, Daniel Leigh, and Andrea Pescatori. 2014. Expansionary Austerity? International Evidence. *Journal of the European Economic Association* 12(4): 949–968.

Gunter, Samara, Daniel Riera-Crichton, Carlos Végh, and Guillermo Vuletin. 2018. Non-Linear Effects of Tax Changes on Output: The Role of the Initial Level of Taxation. World Bank Policy Research Working Paper 8668.

Guren, Adam M., Alisdair McKay, Emi Nakamura, and Jón Steinsson. 2020. Housing Wealth Effects: The Long View. *Review of Economic Studies* 88(2): 669–707.

Hamilton, James D. 1994. *Time Series Analysis*. Princeton, N.J.: Princeton University Press.

Hayo, Bernd, and Matthias Uhl. 2014. The Macroeconomic Effects of Legislated Tax Changes in Germany. *Oxford Economic Papers* 66(2): 397–418.

Hussain, Syed M., and Liu Lin. 2018. Comparing the Effects of Discretionary Tax Changes between the US and the UK. *The B.E. Journal of Macroeconomics* 18(1): 1–17.

Ilzetzki, Ethan, Enrique G. Mendoza, and Carlos A. Végh. 2013. How big (small?) are fiscal multipliers? *Journal of Monetary Economics* 60(2): 239–254.

Jo, Yoon J., and Sarah Zubairy. 2021. State dependent government spending multipliers: Downward nominal wage rigidity and sources of business cycle fluctuations. Working Papers 20210127-001, Texas A&M University, Department of Economics.

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

Jordà, Òscar, Moritz Schularick, and Alan M. Taylor. 2017. Macrofinancial History and the New Business Cycle Facts. *NBER Macroeconomics Annual* 31(1): 213–263.

Jordà, Òscar, Moritz Schularick, and Alan M. Taylor. 2020. The effects of quasi-random monetary experiments. *Journal of Monetary Economics* 112: 22–40.

Jordà, Òscar, and Alan M. Taylor. 2016. The Time for Austerity: Estimating the Average Treatment Effect of Fiscal Policy. *Economic Journal* 126(590): 219–255.

Kato, Akihisa, Wataru Miyamoto, Thuy Lan Nguyen, and Dmitriy Sergeyev. 2018. The Effects of Tax Changes at the Zero Lower Bound: Evidence from Japan. *AEA Papers and Proceedings* 108: 513–18.

Keynes, John Maynard. 1936. The Marginal Propensity to Consume and the Multiplier. In *The General Theory of Employment, Interest, and Money*, chapter 10. London: Macmillan.

Kitagawa, Evelyn M. 1955. Components of a difference between two rates. *Journal of the American Statistical Association* 50(272): 1168–1194.

Leeper, Eric M., Michael Plante, and Nora Traum. 2010. Dynamics of fiscal financing in the United States. *Journal of Econometrics* 156(2): 304–321.

Leeper, Eric M., Nora Traum, and Todd B. Walker. 2017. Clearing Up the Fiscal Multiplier Morass. *American Economic Review* 107(8): 2409–54.

Lo, Ming Chien, and Jeremy Piger. 2005. Is the Response of Output to Monetary Policy Asymmetric? Evidence from a Regime-Switching Coefficients Model. *Journal of Money, Credit and Banking* 37(5): 865–886.

Mertens, Karel, and Morten O. Ravn. 2013. The Dynamic Effects of Personal and Corporate Income Tax Changes in the United States. *American Economic Review* 103(4): 1212–1247.

Miyamoto, Wataru, Thuy Lan Nguyen, and Dmitriy Sergeyev. 2018. Government Spending Multipliers under the Zero Lower Bound: Evidence from Japan. *American Economic Journal: Macroeconomics* 10(3): 247–77.

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.

Mountford, Andrew, and Harald Uhlig. 2009. What are the Effects of Fiscal Policy Shocks. *Journal of Applied Econometrics* 24(6): 960–992.

Nakamura, Emi, and Jón Steinsson. 2014. Fiscal Stimulus in a Monetary Union: Evidence from US Regions. *American Economic Review* 104(3): 753–92.

Nguyen, Anh D. M., Luisanna Onnis, and Raffaele Rossi. 2021. The Macroeconomic Effects of Income and Consumption Tax Changes. *American Economic Journal: Economic Policy* 13(2): 439–66.

Oaxaca, Ronald. 1973. Male-Female Wage Differentials in Urban Labor Markets. *International Economic Review* 14(3): 693–709.

Peersman, Gert, and Frank Smets. 2001. Are the effects of monetary policy in the euro area greater in recessions than in booms? Working Paper Series 52, European Central Bank.

Primiceri, Giorgio E. 2005. Time varying structural vector autoregressions and monetary policy. *The Review of Economic Studies* 72(3): 821–852.

Ramey, Valerie A. 2016. Macroeconomic Shocks and Their Propagation. In *Handbook of Macroeconomics*, edited by Taylor, John B., and Harald Uhlig, volume 2, chapter 2, 71–162. Amsterdam: Elsevier.

Ramey, Valerie A. 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.

Riera-Crichton, Daniel, Carlos A. Vegh, and Guillermo Vuletin. 2015. Procyclical and countercyclical fiscal multipliers: Evidence from OECD countries. *Journal of International Money and Finance* 52: 15–31.

Romer, Christina D., and David H. Romer. 2010. The Macroeconomic Effects of Tax Changes: Estimates Based on a New Measure of Fiscal Shocks. *American Economic Review* 100(3): 763–801.

Rubin, Donald B. 1974. Estimating Causal Effects of Treatments in Randomized and Nonrandomized Studies. *Journal of Educational Psychology* 66(5): 688–701.

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

Wolf, Christian. 2023. Fiscal Stimulus and the Systematic Response of Monetary Policy. *AEA Papers and Proceedings* Forthcoming.

Woodford, Michael. 2011. Simple Analytics of the Government Expenditure Multiplier. *American Economic Journal: Macroeconomics* 3(1): 1–35.

Wooldridge, Jeffrey M. 2001. *Econometric Analysis of Cross Section and Panel Data*. Cambridge, Mass.: MIT Press.

ONLINE APPENDIX

A: Original Guajardo, Leigh, and Pescatori (2014) specification

Figure A.1: *Effects of a 1% of GDP fiscal consolidation: original IMF specification*

(a) *% response of GDP* **(b)** *Response of the short term real interest rate*

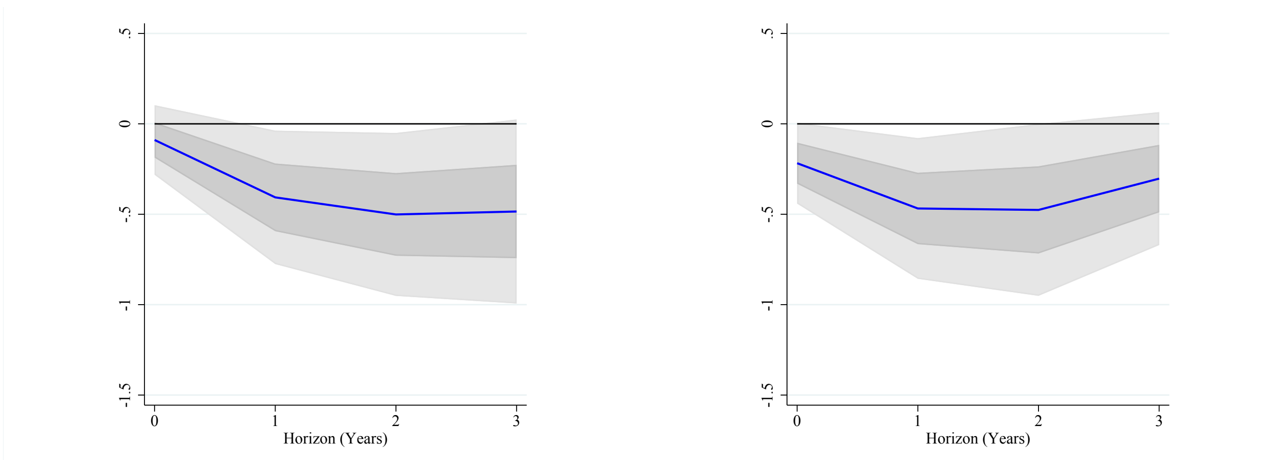

Notes: Vertical axes reported in percent change with respect to the origin. One and two standard deviation confidence bands for each coefficient estimate shown as grey areas. Local projections as specified in equation (8) without indirect effects and using two lags of each control described therein. Sample 1978:1–2009:4. This specification uses the original control set from Guajardo, Leigh, and Pescatori (2014). See text for details.

B. Sensitivity Proxy: A Simple Example

To formalize the interaction we have in mind consider the following, stylized, setup. In the main text we show that this approach works well using simulations from a standard New Keynesian DSGE model. Let some outcome $y_{i,t}$, e.g., GDP growth in country i at time t, depend on fiscal treatment $f_{i,t}$ and the choice of the real interest rate $r_{i,t}$. All variables as expressed relative to their means. Furthermore, suppose the interest rate is set by a monetary authority following a rule. For simplicity, assume that the monetary authority sets the real interest rate directly. Real interest rates are set to offset the negative effects of shocks to GDP, including changes in fiscal policy. Specifically,

$$y_{i,t} = \delta_f f_{i,t} + \delta_r r_{i,t} + u^y_{i,t}, \tag{17}$$

$$r_{i,t} = \overline{\Theta}^f f_{i,t} + \Theta^f_i f_{i,t} + \Theta^y u^y_{i,t} + u^r_{i,t}, \tag{18}$$

where δ_f measures the fiscal multiplier holding interest rates constant. In the data this cannot typically be estimated because interest rates are likely to endogenously respond to fiscal treatment, as is the case in Equation 18. This equation says that monetary policy responds to fiscal interventions but, in the way this rule is written, the degree of monetary accommodation could vary across countries. $\overline{\Theta}^f$ reflects the average response across all countries and Θ^f_i is the idiosyncratic component. Monetary policy also potentially responds to other economic shocks, captured by the term $u^y_{i,t}$. Combining Equation 17 and Equation 18 yields

$$y_{i,t} = (\delta_f + \delta_r \overline{\Theta}^f) f_{i,t} + \delta_r \Theta^f_i f_{i,t} + \delta_r u^r_{i,t} + (1 + \delta_r \Theta^y) u^y_{i,t}. \tag{19}$$

On the assumption that treatment $f_{i,t}$ is randomly assigned (as should be the case if the fiscal shocks are exogenous), the first term illustrates that the reduced-form estimate of the fiscal multiplier depends on the average monetary response in the data, $\delta_r \overline{\Theta}^f$. In other words, δ_f, δ_r and $\overline{\Theta}_f$ are not separately identified using the fiscal shock alone. The second term captures heterogeneity in the interest rate response. Note that Equation 19 has the form of the Kitagawa-Blinder-Oaxaca decomposition in Equation 7. In this simple case without any other controls, f (the policy treatment) in Section 2.2 corresponds to $f_{i,t}$ here and $(x_{i,t} - \overline{x}) = \Theta^f_i$. The indirect effect is then $\delta_r \Theta^f_i$. Since the total response (ignoring the composition effect) is simply the direct effect plus the indirect effect, we can consider experiments around the average effect by arbitrarily varying the indirect effect.

C. Variation in the deficit to GDP ratio by monetary policy response

Figure A.2: *Deficit/GDP ratio* $((D_{t+h} - D_{t-1})/Y_{t-1})$

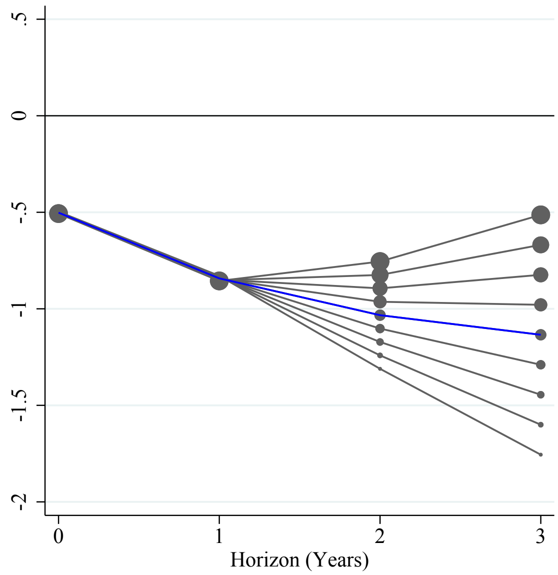

Notes: This figure shows how the response of the deficit to GDP ratio varies with the degree of monetary policy accommodation. The blue lines report the direct effect. The gray lines consider experiments which vary the degree of monetary accommodation. A larger marker indicates a tighter monetary policy scenario.

D. Significance of the direct effect

Figure A.3: *Direct effect: response of GDP (%) to a 1% of GDP fiscal consolidation*

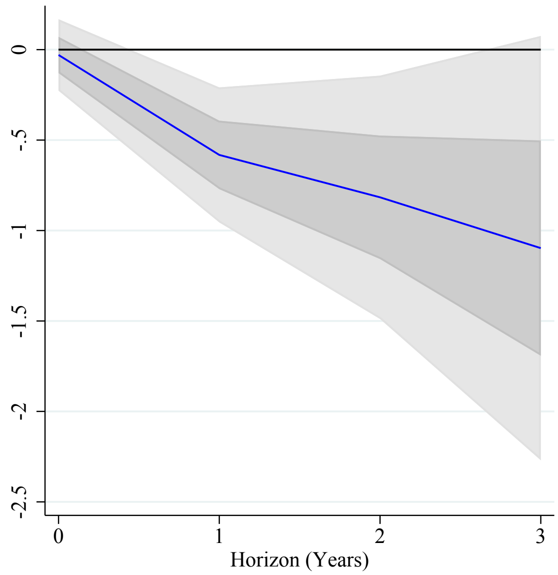

Notes: This figure shows how the response GDP (%) following a 1% of GDP fiscal consolidation. The blue lines report the direct effect estimated from the Kitagawa-Blinder-Oaxaca decomposition together with the one and two standard deviation error bands.

E. Coefficient estimates

Table A.1: *Coefficient estimates for the direct and indirect effects*

Horizon (Years)	β^h	θ_f^h
0	-0.03	-0.50
	(0.10)	(0.43)
1	-0.58	-0.60
	(0.19)	(0.29)
2	-0.82	-0.71
	(0.34)	(0.15)
3	-1.10	-1.25
	(0.60)	(0.16)

Notes: This table reports coefficient estimates based on equation 9. β^h is the impulse response function for the direct effect and therefore corresponds to the Figure A.3. θ_f^h governs the strength of the indirect effect. A negative value implies that, following a fiscal consolidation, real GDP is more negative when monetary policy is less accommodative. Standard errors are reported in parenthesis.

F. Response of the real rate

Figure A.4: *Response of the real interest rate*

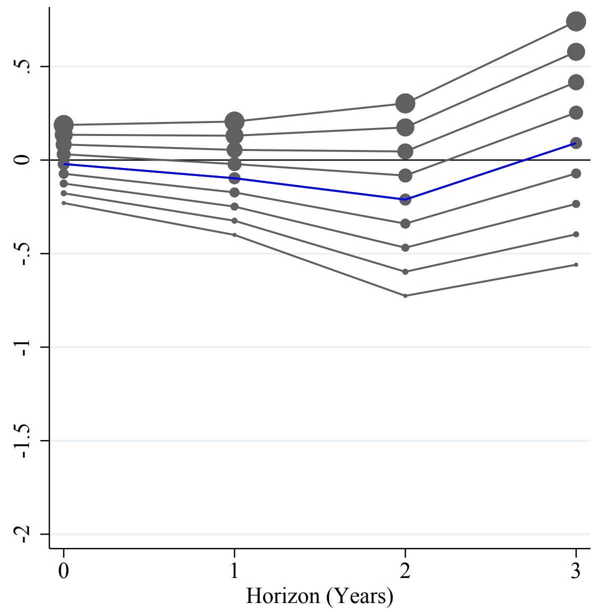

Notes: This figure shows the response of the real interest rate to a 1% of GDP fiscal consolidation. The blue lines report the direct effect. The gray lines consider experiments which vary the degree of monetary accommodation. A larger marker indicates a tighter monetary policy scenario.

G. No identification of the monetary offset

Figure A.5: *Cumulative fiscal multiplier by monetary response*

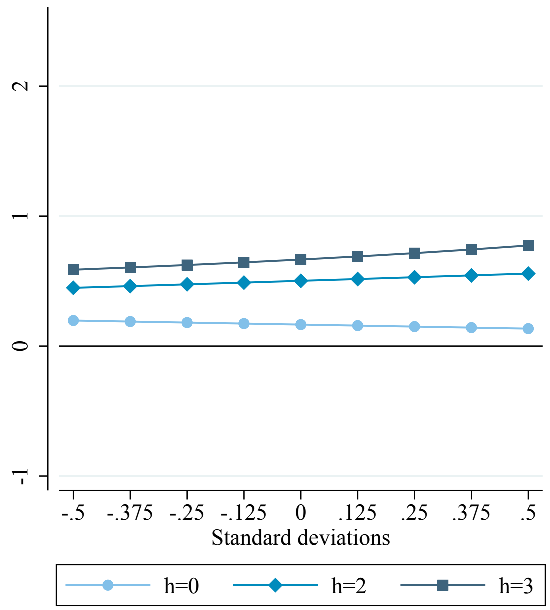

Notes: This chart shows the cumulative fiscal multiplier from each scenario in Figure 4. This is computed as the cumulative sum of the GDP response relative to the cumulative deficit to GDP response (based on Figure A.2). Each line refers to a different horizon, h. As in Figure 4, h goes from the current year $h = 0$ to the third year after the shock $h = 3$. $h = 1$ is omitted to avoid overcrowding the figure. Moving from left to right on the horizontal axis implies a "less active" monetary policy which results in a larger multiplier. Estimation uses the actual change in the future interest rate rather than our sensitivity-based proxy for the monetary offset.

H. Controlling for fiscal composition

Figure A.6: *Cumulative fiscal multiplier by monetary response*

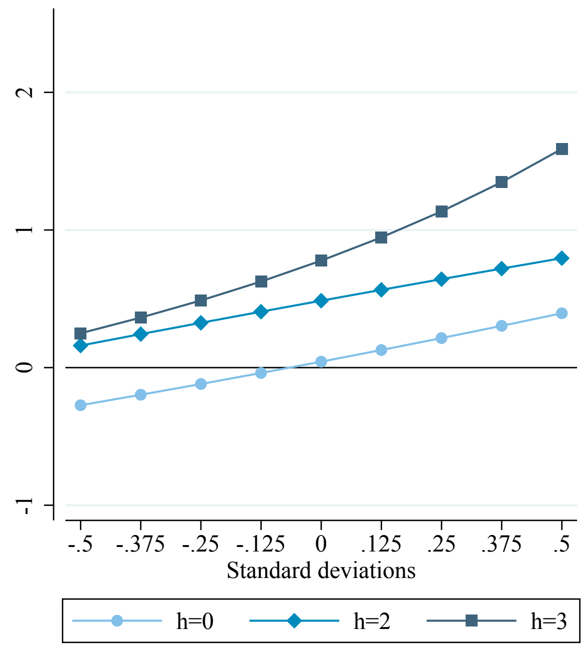

Notes: This chart shows the cumulative fiscal multiplier varying the degree of monetary offset. This is computed as the cumulative sum of the GDP response relative to the cumulative deficit to GDP response. Each line refers to a different horizon, h. As in Figure 4, h goes from the current year $h = 0$ to the third year after the shock $h = 3$. $h = 1$ is omitted to avoid overcrowding the figure. Relative to the baseline figure in the main text, this is produced controlling for the cross-country propensity to use taxes versus spending, as discussed in Section 4

I. Robustness exercises: deficit/GDP ratio

Figure A.7: *Robustness exercises: deficit/GDP ratio* $((D_{t+h} - D_{t-1})/Y_{t-1})$

(a) *Time fixed effects*

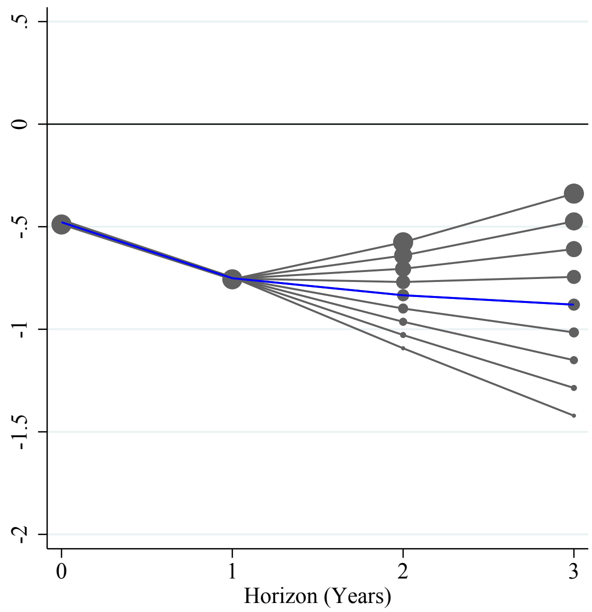

(b) *Longer lag structure*

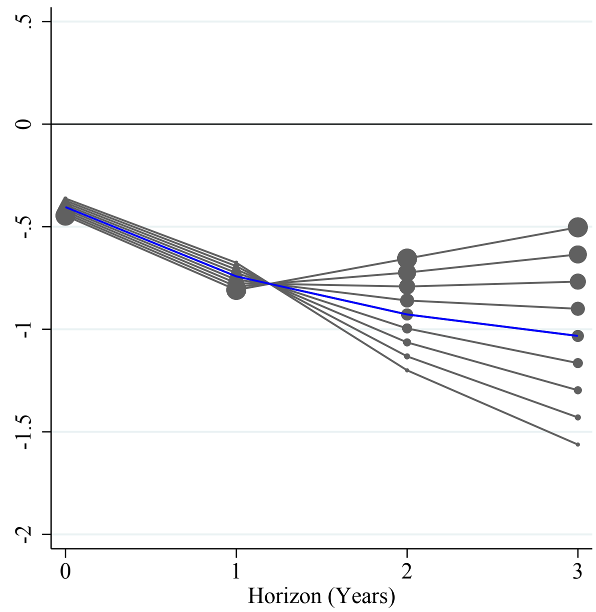

(c) *Alternative interactions approach*

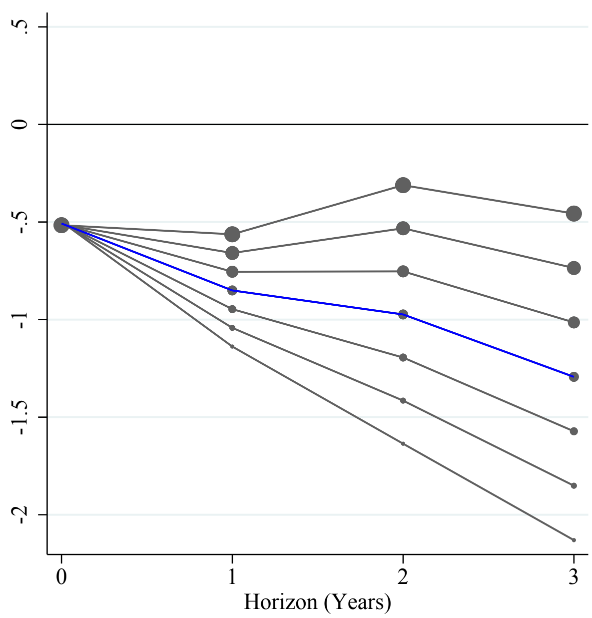

Notes: This figure shows how the deficit to GDP ratio varies with the degree of monetary policy accommodation in each of the robustness exercises covered in Section 4. The blue lines report the direct effect. The gray lines consider experiments which vary the degree of monetary accommodation. A larger marker indicates a tighter monetary policy scenario.

J. Overview of the theoretical model

To motivate our empirical application, Figure 1 in the main paper showed how the fiscal multiplier varies with monetary policy in a relatively standard New Keynesian model. This appendix outlines the model used to produce Figure 1 and to conduct the model simulation exercises in Figure 3 and Figure 9.

Many papers in the literature have, of course, noted that the fiscal multiplier depends on the model's monetary policy rule (see, for example, Woodford (2011) or Leeper, Traum, and Walker (2017)), and a more realistic quantitative model would also include a range of other features and mechanisms. Our goal here, however, is simply to illustrate the key monetary-fiscal interaction we have in mind, and to fix ideas about what we seek to quantify in the data.[30]

For monetary policy to affect the fiscal multiplier, the model needs some form of nominal rigidity. This motivates our focus on the New Keynesian class of models. To generate a wider range of multipliers the model needs to have some other rigidities beyond the simple textbook New Keynesian model. A range of mechanisms could be included but, for simplicity, we follow Galí, López-Salido, and Vallés (2007) and include two types of households, one group who fully optimize and another group who act in a rule-of-thumb manner.[31] In the presence of nominal rigidities, this allows the model to produce a range of different results for the multiplier, some of which are larger than 1 (see Leeper, Traum, and Walker, 2017). Fiscal policy is modeled as a persistent change in government spending. We will assume this policy experiment is financed by lump-sum taxes on saver households.[32] The model is therefore very standard. We first sketch the most important features for our purpose. The second part of this appendix provides full details of the model.

On the household side, the economy is populated by a continuum of households. A share $1-\mu$ of the households can save (or borrow) freely. They fully optimize their intertemporal choices: they choose consumption, saving, hours worked, and bond holdings to maximize expected lifetime utility subject to their budget constraint. We refer to these households as saver households, and their choices with a superscript S. The saver household's consumption plan follows the familiar Euler equation which relates consumption growth to the real interest rate. In linearized form this is

$$E_t\, \Delta \hat{c}^S_{t+1} = \frac{1}{\sigma}\left(\hat{R}_t - E_t\, \hat{\pi}_{t+1}\right)\ ,$$

where $\hat{c}^S_t$ is consumption of the saver household in log deviation from steady state. $\hat{\pi}_t$ is the log change in the price of the consumption good and $\hat{R}_t$ is the monetary policy interest rate, both in deviations from the deterministic steady state. With sticky prices monetary policy can exert control

[30]See Leeper, Traum, and Walker (2017) for a recent and extensive investigation of the effects of changes in government spending changes in a range of macroeconomic models with different frictions and assumptions.

[31]Again, this is purely expositional and, as discussed in Leeper, Traum, and Walker (2017), a number of modeling devices can be used to generate positive consumption effects that produce larger multipliers following a fiscal stimulus.

[32]Alternatively we could have assumed that government spending changes are financed with debt owned by the saver households but, because saver households finance the government, a form of Ricardian equivalence applies here and there is no need to model debt explicitly.

over the real interest rate ($\hat{R}_t - E_t\,\hat{\pi}_{t+1}$).

We assume that the remaining share μ of households are rule-of-thumb decision-makers in the sense that they have no access to bonds and consume all their labor income. We refer to these households as non-saver households, and denote their choices with a superscript N.[33] For them, consumption is therefore pinned down by their budget constraint

$$C_t^N = w_t\,N_t^N\,.$$

where C_t^N is the level of non-saver consumption, w_t is the real wage and N_t^N are hours worked by non-savers. Total consumption in this economy is equal to

$$C_t = \mu C_t^N + (1-\mu)C_t^S\,.$$

On the production side of the model, to rationalize price stickiness, there are two types of firms. Intermediate goods $y_t(j)$ are produced by a constant returns to scale technology $y_t(j) = An_t(j)$ under monopolistic competition. We normalize total factor productivity, A, to 1. Intermediate goods are turned into final goods Y_t by competitive final goods firms using a CES production function $Y_t = (\int_0^1 y_t(j)^{\frac{\epsilon-1}{\epsilon}} dj)^{\frac{\epsilon}{(\epsilon-1)}}$. Final goods are either purchased by households or government, i.e. $Y_t = C_t + G_t$ where G_t is government consumption expenditure. All varieties of intermediate good are substitutable with one another with an elasticity of demand ϵ and the demand curve for variety $y_t(j)$ is given by $y_t(j) = \left(\frac{p_t(j)}{P_t}\right)^{-\epsilon} Y_t$, which the intermediate goods firm takes as given.

Intermediate goods firms set prices and choose labor demand to minimize costs. Their decision problem is standard in the New Keynesian literature. With probability θ a firm is unable to change its price and keeps the same price as at $t-1$. With probability $1-\theta$ the firm is able to fully reset its price. The equilibrium conditions from the firm side lead to a standard dynamic pricing relationship. In linearized form this is the familiar New Keynesian Phillips Curve where inflation depends on expected future inflation and real marginal cost,

$$\hat{\pi}_t \;=\; \beta\, E_t\hat{\pi}_{t+1} \;+\; \tilde{\kappa}\,\widehat{mc}_t\,, \tag{20}$$

where $\tilde{\kappa} = \frac{1}{\theta}(1-\theta)(1-\beta\theta)$, θ is the probability of having a fixed price, and β is the household's discount factor. $\widehat{mc}_t$ is real marginal cost in log deviation from steady state. As usual in the New Keynesian model, the dynamics of real marginal cost are closely related to the output gap.

Fiscal policy is described by an exogenous, persistent stream of government purchases, G_t. We assume that the government finances government spending using lump sum taxes on saver households and the government budget constraint is simply $G_t = T_t$, but similar results would be obtained if we formally allowed for government debt owned by the savers. We will write $\hat{g}_t = \frac{G_t-G}{Y}$, the deviation of G_t from its steady state G relative to steady state output Y. Impulse response functions can then be expressed as fiscal multipliers under a standard definition $\frac{\Delta Y_t}{\Delta G_t}$. We will

[33] These households still make an intratemporal consumption and labor choice. The intratemporal labor supply equation is the same as for the saver household and, given the competitive nature of the labor market, both types of household face the same real wage.

assume that government spending deviations follow an AR(1) process,

$$\hat{g}_t = \rho_g \hat{g}_{t-1} + e_t .$$

Monetary policy follows a Taylor Rule. In terms of monetary-fiscal interactions, our point of departure is that the effects of a change in fiscal policy will be modulated by the nature of the monetary response. We therefore assume that the policy rule relates interest rate changes to inflation, with some persistence,

$$\hat{R}_t = \rho \hat{R}_{t-1} + (1-\rho)\phi^i \hat{\pi}_t , \tag{21}$$

and where ϕ^i can take a number of distinct values, indicating a different monetary policy regime i.

To illustrate how the fiscal multiplier in this model depends on the degree of the "monetary offset", we solve the model using standard linearization-based methods for a range of possible values of ϕ. In particular, we consider $\phi \in [1, 4.5]$.[34] Each time we solve the model we are assuming that agents in the model take ϕ as given. In the main empirical application in the paper, we exploit potential cross-country variation in ϕ for identification and we will interpret differences in ϕ as cross-country variation in average monetary policy behavior over the sample.

Figure 1 in the main text shows how the model-implied fiscal multiplier varies with the strength of the monetary offset, governed by ϕ. A more aggressive monetary policy rule with higher values of ϕ is associated with smaller fiscal multipliers. In contrast, a less activist monetary authority is associated with larger fiscal multipliers. Of course, this figure confirms what is already known in the theoretical literature: that the monetary policy rule has important implications for the size of the fiscal multiplier. Furthermore, note that the slope of the line in Figure 1 is only positive when nominal rigidities are present. Otherwise, monetary policy would have no effect on the fiscal multiplier.

Although our model is standard, Figure 1 in the main text is useful for two reasons. First, one goal of this paper is to produce the empirical counterpart of Figure 1. Second, in Sections 3 and 4 we use this model to validate the empirical approach. In particular, we show that our empirical approach correctly recovers Figure 1 when applied to simulated data from the model.

Further details on the model

The model is a simple variant of the textbook 3-equation New Keynesian model (e.g., as in Galí, 2015) with optimizing and hand-to-mouth households as in Galí, López-Salido, and Vallés (2007). The details below are therefore very standard.

[34] The rest of the model's parameters are calibrated as follows. We set $\psi = 1.7$, implying a Frisch elasticity of around 0.6. The probability of having a fixed price is set to $\theta = 0.85$. The household's discount factor $\beta = 0.99$. The persistence of government spending is set to $\rho_g = 2/3$ and interest rate smoothing is set to $\rho = 0.75$. We set the share of hand-to-mouth households to $\mu = 30\%$. Following Leeper, Plante, and Traum (2010), we set the government consumption share to 8%.

Households

Savers The economy is populated by $1-\mu$ saver/optimizing households:

$$\max_{C_t, N_t, B_t} E_0 \sum_{t=0}^{\infty} \beta^t \left(\log C_t - \frac{N_t^{1+\psi}}{1+\psi} \right) , \tag{22}$$

subject to

$$P_t C_t + Q_t B_t = B_{t-1} + W_t N_t + D_t - T_t \, . \tag{23}$$

Which leads to the following set of first order conditions:

$$(C_t) : \lambda_t = C_t^{-\sigma} \, , \tag{24}$$

$$(N_t) : \lambda_t \frac{W_t}{P_t} = N_t^{\psi} \, , \tag{25}$$

$$(B_t) : Q_t = 1/R_t = 1/(1+i_t) = \beta E_t \frac{\lambda_{t+1}}{\lambda_t} \frac{P_t}{P_{t+1}} \, , \tag{26}$$

where saver households own firms and receive any profits D_t lump sum. These households also finance government activities via a lump sum tax T_t.

In linearized form these equilibrium conditions can be written as:

$$\hat{w}_t = \hat{c}_t^S + \psi \hat{n}_t^S \, ,$$

$$E_t \Delta \hat{c}_{t+1}^S = \frac{1}{\sigma} \left(\hat{R}_t - E_t \hat{\pi}_{t+1} \right) \, .$$

Non-savers Non-saver rule of thumb households simply consume their entire labor income.

$$C_t^N = w_t N_t^N \, .$$

They also, in principle, have an intratemporal labor supply condition which comes from solving the same optimization problem as above for C and N but where $B = D = T = 0$.

$$\frac{W_t}{P_t} = C_t^N N_t^{N^{\psi}} \, . \tag{27}$$

Total consumption is given by:

$$C_t = \mu C_t^N + (1-\mu) C_t^S \, .$$

In linearized form these are:

$$\hat{w}_t = \hat{c}_t^N + \psi \hat{n}_t^N \, ,$$

$$\hat{c}^N_t = \hat{w}_t + \hat{n}^N_t\,,$$

$$\hat{c}_t = \mu\frac{C^N}{C}\hat{c}^N_t + (1-\mu)\frac{C^S}{C}\hat{c}^S_t\,.$$

Firms

Final goods firms Different varieties of goods $y(j)_t$ are aggregated by the final goods firm:

$$Y_t = \left[\int_0^1 y_t(j)^{\frac{\epsilon-1}{\epsilon}}dj\right]^{\frac{\epsilon}{(\epsilon-1)}}\,, \tag{28}$$

where ϵ is price elasticity of demand for good j. Final goods firms choose intermediate inputs to maximize profit:

$$\max_{y_t(j)}\left(P_t\left[\int_0^1 y_t(j)^{\frac{\epsilon-1}{\epsilon}}dj\right]^{\frac{\epsilon}{(\epsilon-1)}} - \int_0^1 p_t(j)y_t(j)dj\right)\,. \tag{29}$$

Which yields the following demand curve and aggregate price index

$$y_t(j) = \left(\frac{p_t(j)}{P_t}\right)^{-\epsilon} Y_t\,, \tag{30}$$

$$P_t = \left(\int_0^1 p_t(j)^{1-\epsilon}dj\right)^{\frac{1}{1-\epsilon}}\,. \tag{31}$$

Intermediate goods firms Intermediate goods firms solve a static labor demand problem and an intertemporal pricing problem subject to Calvo pricing frictions. Each period firms can re-optimize labor demand.

The firm minimizes labor costs by choosing $n(j)_t$ to minimize the following Lagrangian:

$$\min_{n_t(j)} \frac{W_t}{P_t}n_t(j) + mc_t(y_t(j) - An_t(j))\,, \tag{32}$$

where mc_t is real marginal cost. The first order condition is:

$$mc_t = (W_t/P_t)/A\,. \tag{33}$$

When the firm is able to re optimize, they choose $p_t(j)$ to maximize expected profits:

$$E_t\sum_{s=0}^{\infty}\theta^s\left(\beta^s\frac{\lambda_{t+s}}{\lambda_t}\right)\left[\frac{p_t(j)}{P_{t+s}}y_{t+s}(j) - mc_{t+s}y_{t+s}(j)\right]\,, \tag{34}$$

subject to

$$y_t(j) = \left(\frac{p_t(j)}{P_t}\right)^{-\epsilon} Y_t\,, \tag{35}$$

which yields:

$$\sum_{s=0}^{\infty} \theta^s E_t \left(\beta^s \frac{\lambda_{t+s}}{\lambda_t} \right) \left(\frac{p_t^*}{P_{t+s}} y_{t+s}(j) - \frac{\epsilon}{\epsilon - 1} mc_{t+s} y_{t+s}(j) \right) = 0 \,. \tag{36}$$

Linearization of equation 36 and the price index 29 yields the New Keynesian Phillips Curve shown above.

Policy

As mentioned above, government consumption is simply a persistent exogenous stream of purchases funded with lump sum taxes on savers. The budget constraint is therefore:

$$G_t = T_t \,.$$

In linearized form, government spending evolves as follows:

$$\hat{g}_t = \rho_g \hat{g}_{t-1} + e_t \,,$$

where e_t is a mean zero i.i.d. shock.

Monetary policy follows a simple Taylor Rule. The nominal interest rate $\hat{R}_t$, written in deviations from steady state, is set relative to inflation. Importantly, we will think of this rule as varying across country, c, but where each country operates as a closed economy. The policy rule is therefore:

$$\hat{R}_t = \rho_i \hat{R}_{t-1} + (1 - \rho_i) \phi^i \hat{\pi}_t \,. \tag{37}$$