Working Paper Series
Congressional Budget Office
Washington, DC

The Short-Term Effects of Tax Changes—Evidence for State Dependence

U. Devrim Demirel
Congressional Budget Office
devrim.demirel@cbo.gov

August 2016
Working Paper 2016-03

To enhance the transparency of the work of the Congressional Budget Office (CBO) and to encourage external review of that work, CBO's working paper series includes both papers that provide technical descriptions of official CBO analyses and papers that represent original, independent research by CBO analysts. The information in this paper is preliminary and is being circulated to stimulate discussion and critical comment as developmental work for analysis for the Congress. The views expressed here should not be interpreted as CBO's. Papers in this series are available at http://go.usa.gov/ULE.

The author wishes to thank Paul Burnham, Wendy Edelberg, Jeffrey Kling, Sylvain Leduc, Benjamin Page, Jonathan Parker, Jeanine Rees, Jeffrey Werling, and the session participants at the 2015 World Congress of the Econometric Society, 2015 Biennial Conference of Georgetown Center for Economic Research, and 2014 Midwest Macroeconomics Conference for valuable comments and suggestions.

Abstract

This study evaluates how the state of the economy, as measured by the rate of unemployment, influences the short-term effects of tax changes on output and employment. I examine the effects of narratively identified tax changes by employing two alternative approaches to estimating state-dependent impulse response functions that have been widely used in the recent literature. Both approaches suggest that short-term effects of tax changes on output and employment become smaller during times of higher unemployment. That may be because changes in incentives governing the supply of productive inputs may have a lesser effect on resource utilization when there is slack in the economy.

Keywords: Tax Policy, Fiscal Policy, State Dependence, Regime Switching

JEL classification: E2, E6, H2

1 Introduction

The effect of taxes on economic activity has long been a central area of interest to economists and policymakers. Although the research literature has identified the major channels through which tax changes affect the economy and provided theoretical insights into how those channels operate, no broad consensus yet exists regarding the size of the effects from tax changes. Romer and Romer (2010) estimate that a reduction in tax revenues by 1 percent of gross domestic product (GDP) can boost GDP by roughly 3 percent. Blanchard and Perotti (2002), however, estimate a much smaller figure—about 1 percent, and Favero and Giavazzi (2012) estimate a still smaller figure of roughly 0.5 percent.

To better understand some of the factors that might underlie those diverse estimates, this paper investigates how the effects from a tax change might vary across different states of the economy, in particular, across states of high and low unemployment. The main conclusion to emerge from the analysis is that tax changes tend to have less of an effect on employment and output when the amount of slack in the economy (as measured by the rate of unemployment) is greater.

Macroeconomic effects of fiscal policy changes may vary depending on the state of the economy for various reasons.[1] For example, the effect of fiscal expansions on output may become larger when unemployment is high because a larger segment of households would become liquidity-constrained and thus spend a greater fraction of their additional income (see Tagkalakis, 2008). Moreover, in the absence of inflationary concerns, monetary policy is unlikely to offset the boost in aggregate demand that a fiscal expansion would induce. Indeed, a number of recent empirical studies (including Fazzari and others, 2014; Auerbach and Gorodnichenko, 2013, 2012; Bachman and Sims, 2012; and Mittnik and Semmler, 2012) find that fiscal expansions boost output by a larger amount in recession than they do in expansion.[2] (On the contrary, Ramey and Zubairy, 2014; and Owyang and others, 2013; find no evidence of such relationship.) Those studies, however, only examine the effects of government purchases. Because government purchases effect the economy mainly by changing aggregate demand, their results primarily reflect how the effects on aggregate demand from fiscal expansions vary across different states of the economy. The results of this study most likely capture a wider range of effects than those reflected in the previous liter-

[1] See Parker (2011) for a detailed discussion. The Congressional Budget Office's (CBO) analyses of fiscal policies also take into account the contemporaneous economic conditions under which such policies are implemented. For example, CBO's range of estimates for the demand multiplier (defined as the total change in gross domestic product for each dollar of direct effect on aggregate demand) varies with the degree to which the economy's resources are utilized (that is, the amount of economic slack) and monetary policy responses during the periods when the changes in fiscal policies take place. For a detailed discussion, see Reichling and Whalen (2015) and Congressional Budget Office (2014a).

[2] A related branch of the literature examines how the effects of government purchases vary with the state of monetary policy. Cogan and others (2010), Christiano and others (2011), and Woodford (2011) argue that government purchases have larger effects on output when the nominal interest rate is near zero.

ature because the tax changes I examine influence the economy in many ways, only one of which is by changing aggregate demand.

A broad tax change can have many effects because it can simultaneously alter individuals' decisions regarding work, consumption, and saving, and business' decisions regarding hiring, investment, and capacity utilization. While some of those effects (such as those on individuals' consumption and firms' investment decisions) primarily influence aggregate demand, others (such as those on individuals' decisions regarding work) mainly alter the incentives that govern the supply of labor and other productive inputs. Intuitively, when the economy is weak and there is substantial slack, altering the supply of productive inputs would, for the most part, add to or subtract from the existing pool of underused resources without significantly affecting the extent to which resources are utilized in the short term. For example, in a state of high unemployment, a boost in the supply of labor driven by a tax cut is unlikely to raise employment and output as much as it does when there is heightened competition among employers (that is, when the economy is in a state of low unemployment). As a result, the effect of a tax cut on employment and output would be smaller than what it would be in the absence of slack.

Estimating the macroeconomic effects of tax changes presents a series of methodological challenges. I confront the difficulties related to the identification of exogenous tax changes and accurate dynamic representation of data by building on recently developed methods. In an influential study, Romer and Romer (2010) address the identification problem by carefully examining the narrative record of legislated tax changes (historical documentation pertaining to pieces of past legislation) in the post–World War II period. I draw heavily on that narrative record to identify exogenous changes in taxes. I then evaluate the effects of those changes using two methods, one that is based on local projections, and another that uses regime-switching vector autoregressions (VARs).

The local-projection approach is based on the method proposed by Jorda (2005) and follows Auerbach and Gorodnichenko (2013) and Ramey and Zubairy (2014) in adapting that method to a state-dependent framework. I then directly estimate the responses of output and employment to narratively identified tax changes at different time horizons by estimating a separate state-dependent model for each horizon. The result is a set of impulse response estimates that show how the responses of output and employment to a tax change vary with the state of the economy.

The VAR approach accounts for the possibility that Romer and Romer's (2010) exogenous tax changes may be measured with error; following Mertens and Ravn (2013), I treat narrative tax changes as noisy observations of latent tax shocks. In the linear VAR model of Mertens and Ravn, however, effects of tax shocks do not depend on the state of the economy. To introduce state dependence, I embed the regime-switching structure proposed in Auerbach and Gorodnichenko (2012) into an otherwise linear VAR model, thereby allowing the propagation of tax shocks to vary across different states. I then develop a strategy to extend the identification approach of Mertens and Ravn to cases in which the relationship between VAR residuals and latent tax shocks are state dependent, thereby allowing the contemporaneous responses to tax shocks to vary across states as well.

I use the local-projection and VAR-based methods in tandem because, for the purposes of this analysis, neither method is clearly superior to the other, and also because the two methods together provide a more complete characterization of state dependence. Crucially, however, both methods suggest that the effects of tax changes on employment and output are smaller in a state of high unemployment and larger when unemployment is low. Moreover, that pattern remains intact under a series of checks for robustness.

Section 2 of this paper discusses the proposed identification strategy and explains how it relates to Romer and Romer's (2010) narrative approach. Section 3 describes a local-projection-based approach to estimating impulse responses and presents the estimated response profiles. Section 3 also explores the implications of some alternative specifications of the state-dependent model on which these local projections are based. Section 4 lays out the VAR-based approach, discusses how it captures state dependence, and describes the identification and estimation procedures. Section 4 also presents VAR-based impulse response estimates, and evaluates their sensitivity to alternative specifications of the VAR model. Section 5 concludes.

2 Identifying Tax Shocks

One of the most challenging problems of tax policy analysis is that of identification. Tax changes are correlated with myriad economic factors; some legislated tax changes occur in response to budgetary needs while others are primarily motivated by policymakers' views on the economy's current and future prospects of growth. Taxes also fluctuate without any active changes in tax policy; revenues and average rates change automatically as the tax base expands and contracts over the business cycle or as the distribution of income across tax brackets shifts (through a process known as bracket creep). For those and other reasons, identifying the tax changes that are truly exogenous with respect to macroeconomic variables and disentangling their effects from those of other sources of economic fluctuations can be extremely difficult, which often compels researchers to rely on questionable assumptions to achieve identification.

I address the identification problem by drawing heavily on the narrative record of legislated changes to tax policy documented in Romer and Romer (2010). Carefully examining various sources of information, Romer and Romer determine the pieces of legislation that aim to enhance the long-term performance of the economy or that were enacted in response to inherited budget deficits.[3] Those pieces of legislation are considered to be exogenous because, unlike the tax

[3]Those sources of information include the Economic Report of the President, presidential speeches, the Annual

changes that evidently respond to business cycles or to changes in government spending, they do not aim to stabilize macroeconomic fluctuations, nor were they enacted because spending was changing. Therefore, rather than being direct consequences of prevailing economic and budgetary conditions, they are considered drivers of those conditions themselves.

In the local-projection approach, I treat Romer and Romer's narrative tax changes as structural shocks—exogenous and unexpected changes in revenues that are uncorrelated with past realizations of the variables included in the analysis. That accords with Romer and Romer's original treatment and yields a straightforward interpretation of the impulse response estimates. Narrative tax changes, however, may not perfectly overlap with structural tax shocks. One reason is that the narrative record identifies the exogenous changes in tax *liabilities* which, as Barro and Redlick (2011) note, correspond to intended (or targeted) changes in revenues rather than actual changes. In addition, as Romer and Romer (2010) acknowledge, determining the underlying motivation for certain tax changes and deciding if that motivation suggests exogeneity can be very difficult. It is, therefore, quite possible—if not highly probable—that the exogenous tax changes in Romer and Romer's narrative record are measured with error. The presence of measurement error, however, breaks the presumed one-to-one correspondence between the narrative tax changes and the structural tax shocks that I seek to identify.

The nonlinear VAR approach addresses the problem of measurement error by treating narrative tax changes as a series of noisy observations that convey information about latent tax shocks instead of treating them as tax shocks themselves. To recover that information, I follow Mertens and Ravn (2013) in using the narrative tax changes as a proxy for tax shocks, adopting a procedure that resembles instrumental variable estimation. My approach extends the method of Mertens and Ravn to regime-switching VAR settings by allowing the relationship between VAR residuals and narrative tax changes to vary with the state of the economy.

3 Estimating the Effects of Tax Changes Using Local Projections

This section evaluates the effects of tax changes using a version of Jorda's (2005) local projections (LP) approach that allows for state dependence.[4] Following Ramey and Zubairy (2014) and

Report of the Secretary of the Treasury, and the reports prepared by the House Ways and Means Committee and the Senate Finance Committee.

[4]See Stock and Watson (2007) and (2003) for applications of this method to economic forecasting.

Auerbach and Gorodnichenko (2013), the benchmark model is specified as

$$\begin{aligned} y_{t+k} &= \mu_k + \zeta_t + [1 - h(d_{t-1})] \left[\sum_{i=1}^{q} T^i_{H,k} X_{t-i} + \phi_{H,k} z_t \right] \\ &\quad + h(d_{t-1}) \left[\sum_{i=1}^{q} T^i_{L,k} X_{t-i} + \phi_{L,k} z_t \right] + \xi_{t+k}, \end{aligned} \tag{1}$$

with

$$h(d_{t-1}) = \frac{\exp(-\eta d_{t-1})}{1 + \exp(-\eta d_{t-1})}, \tag{2}$$

where $k = 0, 1, 2...$ denotes the impulse-response horizon and y_{t+k} is the variable of interest (discussed below) in time $t+k$. The terms μ_k and ζ_t, respectively, denote a constant and a deterministic time trend; X_{t-i} is a vector of control variables (which may include lags of y); and z_t denotes the tax change that the narrative record of Romer and Romer (2010) shows for time t. The term ξ_{t+k} is a potentially autocorrelated and heteroskedastic error term. (That is, ξ_{t+k} can be correlated with its past values, and its variance can depend on X_{t-i} and z_t.)

Variable d_{t-1} is the unemployment rate gap series constructed by the Congressional Budget Office (CBO). That variable measures the state of the economy when the tax change occurs. Unemployment rate gap is the difference between the actual unemployment rate and CBO's estimate of the underlying long-term rate of unemployment.[5] The underlying long-term rate is defined as the rate of unemployment caused by reasons other than cyclical downturns and short-term structural factors. Therefore, a large positive value for d_{t-1} indicates a state of high cyclical unemployment.

The function $h(d_{t-1})$ (henceforth, the transition function) takes values between 0 and 1 and determines the weights attached to high- and low-unemployment regimes for different values of d_{t-1}. The parameter $\eta > 0$ controls the speed of transition from one regime to another. In the benchmark case, it is set so that the economy experiences high unemployment—which I, following Ramey and Zubairy (2014), define as a state in which the rate of unemployment exceeds 6.5 percent—roughly 26% of time. That figure matches the frequency of episodes with high unemployment since early 1950s.

The LP method involves estimating a separate equation of the form (1) for each impulse-response horizon (k) and for each variable of interest (such as real GDP and tax revenues). Once the coefficients $\phi_{H,k}$ and $\phi_{L,k}$ in (1) are estimated for a given k, the response of variable y in time $t+k$ to a tax change that occurs in time t and state $d_{t-1} = d$ can be found as

$$[1 - h(d)]\phi_{H,k} + h(d)\phi_{L,k}. \tag{3}$$

[5] CBO assesses the cyclical component of various macroeconomic variables by estimating the relationship between those variables and the unemployment rate gap. For additional information, see Congressional Budget Office (2014b).

State-dependent impulse response functions of length N are then constructed by estimating (1) for $k = 1, 2...N$ and computing (3) for each k. Because the function $h(d_t)$ has the properties $\lim_{d_t \to \infty} h(d_{t-1}) = 0$ and $\lim_{d_t \to -\infty} h(d_{t-1}) = 1$, the coefficients $\phi_{H,k}$ and $\phi_{L,k}$ represent the responses of y to a tax change in two limiting states with very high and very low unemployment, respectively.

To express the changes in output and tax revenues in same units, I follow Ramey and Zubairy (2014) and Owyang and others (2013) in defining the changes in the variables of interest as shares of real GDP. For example, when estimating (1) for output, I define y_{t+k} as the change in real GDP from quarter $t-1$ to quarter $t+k$ divided by real GDP in quarter $t-1$. Similarly, when estimating (1) for real tax revenues, I define y_{t+k} as the change in real revenues from $t-1$ to $t+k$ divided by real GDP in $t-1$.

A major advantage of the LP approach over standard VAR estimation is that it accounts for the economy's state-to-state transition without requiring additional assumptions about the manner in which such transition occurs. VAR-based estimates of state-dependent impulse responses (including the benchmark estimates of Auerbach and Gorodnichenko, 2012; and Bachman and Sims, 2012) are often computed assuming that the state of the economy remains fixed throughout the evaluated response period. Under that assumption, however, tax changes do not influence the economy's transition across states. That feature can be problematic, especially when the average amount of time in which the economy remains in a particular state is significantly less than the amount of time in which the state is assumed to remain fixed in the estimated model. LP-based impulse responses incorporate the economy's transition across states because the estimated coefficients $\phi_{H,k}$ and $\phi_{L,k}$ reflect, among other things, how the state of the economy tends to evolve in the sample within k periods after a tax change. Thus, if, on average, tax cuts push the economy toward a state of lower unemployment in the sample (or if tax increases push toward a state of higher unemployment), the estimated coefficients $\phi_{H,k}$ and $\phi_{L,k}$ will reflect that tendency.

Another advantage of the LP approach is that it is less sensitive to specification error than the VAR-based method of estimating impulse responses. In the standard VAR analysis of how tax changes affect the economy, impulse responses of length N are constructed by estimating the contemporaneous responses to a tax change and then iterating the VAR model forward to trace out the variables' evolution over N periods after the tax change. If the estimated VAR is incorrectly specified, however, that iterative procedure tends to amplify the specification error. The LP method reduces the potential impact of specification error because it estimates the responses directly by projecting the variables of interest on tax changes for each horizon $k = 1, 2, ..N$ without imposing additional restrictions on how variables evolve over time after a tax change.

However, neither approach is clearly superior to the other in all situations. For example, although the LP approach can be advantageous if the impulse response analysis is based on an incorrectly specified model, VAR-based estimates are asymptotically more efficient if the underlying model is correctly specified. Moreover, my sensitivity analysis and the findings of Ramey (2012) suggest that, over longer horizons, LP-based responses can exhibit statistically significant yet erratic oscillations. Using the VAR- and LP-based methods in tandem, therefore, almost certainly helps to obtain a more complete picture.

3.1 Impulse Responses

To assess the variation in the effects of tax changes across different states, response functions are evaluated under alternative values of the unemployment rate gap d. Those values are determined by balancing two important considerations. First, the values must be sufficiently differentiated to allow for a well-defined pattern of state dependence, provided that one exists, to emerge from the analysis. Second, because inference can become highly inaccurate if the number of observations in the neighborhood of a particular state is too small or if the amount of time the economy remains in that neighborhood is too short, the determined values must fall well within historical experience. Balancing those two factors, I evaluate the response functions for the values $d = 0.5\%$ (lower-unemployment state), and $d = 1.5\%$ (higher-unemployment state).

In all cases, I estimate (1) and (3) for 12 quarters (that is, for $k = 1, 2...12$) and, using the procedure described above, compute the responses of real GDP and employment to an exogenous change that reduces tax liabilities by 1 percent of GDP. All estimated responses are presented within plus and minus 1-standard-error bands, which are based on heteroskedasticity and autocorrelation consistent (HAC) standard errors (see Newey and West, 1987). The lag length for each control variable under each impulse response horizon (k) is decided using the Akaike and Schwartz information criteria. Following Francis and Ramey (2009), I include a quadratic trend in all specifications to control for slow-moving components of the variables.

3.1.1 Benchmark Specification

In the benchmark case, the vector of control variables X_{t-i} includes real GDP, real cyclically-adjusted tax revenues, and real federal deficits, all of which are in 2009 dollars and entered in quarterly levels. CBO constructs the series for cyclically adjusted revenues by removing the effects of business cycle fluctuations so that the adjusted series do not incorporate automatic movements generated during business ups and downs. The sample period runs from the first quarter of 1951 to the fourth quarter of 2006.

Figure 1 suggests a negative relationship between the size of the responses of real GDP and employment to a tax change and the value of the unemployment rate gap at the time of the tax change; the responses are larger in the lower-unemployment state. Table 1 shows the estimated values of the average, maximum, and relative responses of GDP (denoted $\frac{1}{N}\Sigma_{i=1}^{N} y_i$, $\max_{i=1..N}\{y_i\}$, and $\Sigma_{i=1}^{N} y_i / \Sigma_{i=1}^{N} \tau_i$, respectively) over $N = 8$ quarters. The relative response is defined as the average response of GDP divided by that of revenues over the same time period. All three measures suggest that a tax cut's effect on output is stronger during times of lower unemployment. In particular, the maximum response of real GDP to a reduction in revenues by 1 percent of GDP is roughly 2.5 percent in the lower-unemployment state and around 1.2 percent in the higher-unemployment state.

The responses of employment are estimated to be more gradual than those of output. As Figure 1 shows, real GDP and employment gradually increase after a tax cut, with the response of real GDP peaking in about 8 quarters in both states. The response of employment peaks in about 11 quarters in the lower-unemployment state and in 9 quarters in the higher-unemployment state.

3.1.2 Alternative Specifications

Additional Variables. To evaluate the sensitivity of the benchmark results, I first extend the vector of control variables by including real government spending and the interest rate on 3-month U.S. Treasury bills. That specification can be regarded as a natural extension of the benchmark case. Including government purchases alongside deficits helps to control for the interaction of spending and revenues, and including the short-term interest rate helps to account for the role of monetary policy and attributes movements in output and employment to the correct policy variable (see Rossi and Zubairy, 2011). Estimated responses do not change qualitatively when those two variables are included in the vector of controls. Figure 2 and Table 1 show that the increases in real GDP and employment are stronger in the lower-unemployment state. Both sets of responses are, however, noticeably smaller than those of the benchmark case in both unemployment states.

I also reestimate (1) by including the unemployment rate gap, d, in the set of control variables. With d introduced as a separate independent variable, (1) accounts for potential direct effects the state of the economy may have on output and employment responses in addition to the effects of interactions between d and tax changes. Estimated responses again do not change qualitatively when d is included into the benchmark control vector. The increases in real GDP and employment are larger in the lower-unemployment state regardless of whether d is the only addition to the benchmark control vector or it is included alongside real government spending and the interest rate.

Controlling for Anticipation. The next step is to assess the potential impact of anticipation effects on the results. The responses of output and employment to anticipated changes in fiscal policy can be quite different from responses to unanticipated changes (see, for example, Leeper and others, 2012; Mertens and Ravn, 2012; Ramey, 2011; and Yang, 2005) because forwardlooking agents can react to future policy changes as soon as news about such changes arrives.[6] To check the sensitivity of the estimates to anticipation effects, I follow Mertens and Ravn (2013) in excluding from the sample the tax changes with implementation lags exceeding 90 days and retaining only those with shorter lags. That eliminates roughly half of the tax changes that Romer and Romer (2010) classify as exogenous. Figure 3 and Table 1 show that the main results are not sensitive to excluding anticipated tax changes; except for the first two quarters, the increases in real GDP and employment are stronger in the lower-unemployment state.

A Threshold Specification. I also examine if the results depend on how the transition function $h(d_{t-1})$ is specified by considering the threshold specification adopted by Ramey and Zubairy (2014). Specifically, I replace the transition function $h(d_{t-1})$ in (1) and (3) with an indicator function $I_{t-1} \in \{0, 1\}$ that takes the value 1 if the unemployment rate gap in quarter $t-1$ is greater than a threshold level $\overline{d}$ and the value 0 if otherwise. That is, the economy is considered to be in a state of low unemployment in quarter t if the unemployment rate gap in $t-1$ is smaller than $\overline{d}$ and in a high-unemployment state if otherwise. Unlike the continuous function $h(d_{t-1})$, which suggests smooth transition between states, the indicator function I_{t-1} implies that the state of the economy shifts abruptly at the threshold $\overline{d}$. I set the threshold unemployment rate gap, $\overline{d}$, to 1.1 percent, which roughly corresponds to a threshold unemployment rate of 6.5 percent (given CBO's estimate of the average natural rate of unemployment in the sample period). As Figure 4 and Table 1 show, the increases in real GDP and employment from a tax cut are again estimated to be larger in the lower-unemployment state (that is, when $\acute{d}_t < 1.1$) than in the higher-unemployment state (when $d_t > 1.1$).

4 Estimating the Effects of Tax Changes Using Regime-Switching VARs

The basic VAR approach is based on the regime-switching model proposed in Auerbach and Gorodnichenko (2012). The benchmark specification is

$$X_t = \mu + [1 - h(d_{t-1})] \sum_{i=1}^{q} A_i^H X_{t-i} + h(d_{t-1}) \sum_{i=1}^{q} A_i^L X_{t-i} + u_t, \tag{4}$$

where

$$u_t \sim N[0, \Omega(d_{t-1})], \text{ and } E[u_t u_j'] = 0 \text{ for } t \neq j,$$

[6]Contrary to that argument, Perotti (2012) finds no significant evidence for anticipation effects.

with

$$\Omega(d_{t-1}) = [1 - h(d_{t-1})]\Omega_H + h(d_{t-1})\Omega_L. \tag{5}$$

The vector $X_t = [\tau_t \, W_t']'$ is $n \times 1$ and includes taxes (τ_t) and a subvector W_t that contains additional variables of interest. The objects $\{A_i^H\}_{i=1}^q$ and $\{A_i^L\}_{i=1}^q$ are $n \times n$ coefficient matrices, and μ is a $n \times 1$ vector of constants. The transition function $h(d_{t-1})$ is as specified in (2).

The model captures state dependence in two ways. First, differences between autoregressive coefficients $\{A_i^H\}_{i=1}^q$ and $\{A_i^L\}_{i=1}^q$ allow the propagation of shocks to vary with the state of the economy. Second, differences between Ω_H and Ω_L in (5) allow the contemporaneous responses of the VAR variables to various shocks to vary with the state of the economy. The core autoregressive part of (4) and the residual covariance matrix $\Omega(d_{t-1})$ approach $\{A_i^H\}_{i=1}^q$ and Ω_H, respectively, as unemployment increases and approach $\{A_i^L\}_{i=1}^q$ and Ω_L as unemployment decreases. Therefore, the pairs $\{A_i^H\}_{i=1}^q$, Ω_H and $\{A_i^L\}_{i=1}^q$, Ω_L represent two polar regimes that correspond to very high and very low unemployment states. The VAR described in (4) continually switches regimes because, at any given point in time, its dynamics are governed by a weighted average of two regimes and the weights assigned to each regime vary with the state of the economy.

4.1 VAR Residuals, Narrative Tax Changes, and Tax Shocks

The LP-based impulse responses are computed assuming that exogenous tax changes identified in Romer and Romer's (2010) narrative account are the structural tax shocks that I seek to identify. Because narrative tax changes may be subject to measurement error, however, they may not perfectly overlap with tax shocks. A more explicit examination of the relationship between tax shocks and narrative tax changes helps to account for that imperfect overlap.

The relationship between VAR residuals and the vector of all structural shocks that buffet the economy is described by

$$u_t = R(d_{t-1})v_t, \tag{6}$$

where $v_t \sim N(0, I_{n\times n})$ and $I_{n\times n}$ denotes the identity matrix. The matrix $R(d_{t-1})$ describes how the mapping from structural shocks to the residuals varies with the state of the economy. Given $u_t \sim N[0, \Omega(d_{t-1})]$, equation (6) suggests that $R(d_{t-1})$ is linked to the covariance matrix $\Omega(d_{t-1})$ through the relationship

$$\Omega(d_{t-1}) = R(d_{t-1})R(d_{t-1})'. \tag{7}$$

Without loss of generality, the first element of the vector v_t can be assumed to be the tax shock in question. The vector of structural shocks can then be written as $v_t = [v_{1,t} \, v_{2,t}']'$, where $v_{1,t}$ is the tax shock and $v_{2,t}$ is a vector that collects the remaining structural shocks. Accordingly, the matrix

$R(d_{t-1})$ can be partitioned in the form

$$R(d_{t-1}) = \begin{bmatrix} \underset{1\times 1}{R(d_{t-1})_{11}} & \underset{1\times (n-1)}{R(d_{t-1})_{12}} \\ \underset{(n-1)\times 1}{R(d_{t-1})_{21}} & \underset{(n-1)\times (n-1)}{R(d_{t-1})_{22}} \end{bmatrix}, \tag{8}$$

where the scalar $R(d_{t-1})_{11}$ gives the contemporaneous effect of a tax shock on the first variable in X_t and the column vector $R(d_{t-1})_{21}$ gives the same for the remaining variables in X_t. Because the tax shock is ranked first in v_t, the first column of $R(d_{t-1})$ (that is, the *impact vector* of the tax shock) describes the contemporaneous effect of tax shocks on X_t.

Because of potential measurement error, narrative tax changes may be best viewed as imperfect observations of the true structural tax shocks. In other words, they can be considered as a series of noisy signals that convey information about underlying shocks. To make the relationship between narrative tax changes and latent tax shocks more precise, I next define the random variable z_t as the tax change that the narrative record shows for quarter t. Specifically, z_t denotes the measured exogenous change in federal tax liabilities in quarter t expressed as a percentage of GDP and demeaned by subtracting the average from nonzero observations. As in Mertens and Ravn (2013), the relationship between z_t and the tax shock $v_{1,t}$ can be thought to take the form

$$z_t = \theta_t(\psi v_{1,t} + \alpha_t), \tag{9}$$

where $\alpha_t \sim N(0, \sigma_\alpha^2)$ represents measurement error and ψ is a positive constant. The variable $\theta_t \in \{0, 1\}$ is a random indicator function describing a censoring process that governs the observations on narrative tax changes. That variable takes the value 1 in a particular quarter if the narrative account shows a tax change for that quarter. Otherwise, it takes the value 0.

I assume that α_t is orthogonal to $v_{1,t}$ and the censoring process is independent of α_t and $v_{1,t}$. It then follows from (6) and (9) that

$$E[z_t u_t'] = (1-\delta)\psi R(d_{t-1})_1', \tag{10}$$

where $1-\delta$ is the probability of the event $\theta_t = 1$ and $R(d_{t-1})_1 = [R(d_{t-1})_{11}\ R(d_{t-1})_{21}']'$ denotes the first column vector of the matrix $R(d_{t-1})$. Although the relationship between narrative tax changes and structural tax shocks does not depend on the state of the economy, the covariance of z_t with u_t' is state dependent, as 10 suggests. That follows directly from the state-dependent residual covariance specification described in (5).

Next, observe that, given (6), equation (9) can be rewritten as

$$\begin{aligned} z_t &= \theta_t \left[\psi \gamma R(d_{t-1})^{-1} u_t + \alpha_t\right] \\ &= \theta_t \left[\psi \kappa(d_{t-1}) u_t + \alpha_t\right], \end{aligned} \tag{11}$$

where γ is a $1 \times n$ vector of all zeros but the first element is replaced with unity and $\kappa(d_{t-1}) = \gamma R(d_{t-1})^{-1}$ is a $1 \times n$ vector-valued function that describes how the relationship between narrative tax changes and VAR residuals varies with the state of the economy. Provided that tax shocks, measurement error, and the censoring process are uncorrelated with each other, it follows from (11) that

$$E[z_t u_t'] = (1-\delta)\psi \kappa(d_{t-1}) \Omega(d_{t-1}). \tag{12}$$

Therefore, combining (10) and (12), the first column of (8) can be identified as

$$R(d_{t-1})_1 = \Omega(d_{t-1})' \kappa(d_{t-1})'. \tag{13}$$

Equation 13 shows how the impact vector of tax shocks can be identified from estimates of the state-dependent objects $\kappa(d_{t-1})$ and $\Omega(d_{t-1})$. Appendix A discusses $R(d_{t-1})_1$ in more detail and highlights its relation to Merten and Ravn's (2013) analogous construct.

4.2 Estimation

Estimating VAR-based impulse responses involves computing variables' contemporaneous responses to a tax shock and then iterating (4) forward to construct their responses over time. To implement that procedure, estimates for the objects $\kappa(d_{t-1})$ and $\Omega(d_{t-1})$ are needed to recover contemporaneous responses via (13), and estimates for the autoregressive coefficients $\{A_i^H\}_{i=1}^q$ and $\{A_i^L\}_{i=1}^q$ are needed to iterate (4) forward. In principle, all of those objects can be estimated jointly by using a maximum likelihood method. To construct an appropriate likelihood function, however, the functional form of $\kappa(d_{t-1})$ needs to be determined. Appendix B shows that, combining (5) and (7) with the definition $\kappa(d_{t-1}) = \gamma R(d_{t-1})^{-1}$, that form can be identified up to a sign convention as

$$\kappa(d_{t-1}) = g(d_{t-1})\overline{\kappa}, \tag{14}$$

where $\overline{\kappa}$ is a $1 \times n$ vector of constants and

$$g(d_{t-1}) = \left[\frac{1+\exp(-\eta d_{t-1})}{a_1 + (a_1 + a_2)\exp(-\eta d_{t-1})}\right]^{\frac{1}{2}}. \tag{15}$$

The constants a_1 and a_2 are defined as

$$a_1 = \overline{\kappa}\Omega_H\overline{\kappa}' \text{ and } a_2 = \overline{\kappa}(\Omega_L - \Omega_H)\overline{\kappa}'. \tag{16}$$

As shown in Appendix C, given (14)–(16), the relevant log-likelihood function can be constructed as

$$\begin{aligned} \log \mathcal{L}(\Phi) \;\; = \;\; & \text{constant } + \frac{1}{2}\sum_{t=1}^{T}\left[\log\left|\Omega(d_{t-1})^{-1}\right| - (X_t - H_t'\beta)'\Omega(d_{t-1})^{-1}(X_t - H_t'\beta)\right] \quad (17) \\ & + \sum_{z_t=0}\log\delta + \sum_{z_t\neq 0}\left[\log(1-\delta) - \log\sigma_\alpha - \frac{1}{2}\left(\frac{z_t - \psi\kappa(d_{t-1})(X_t - H_t'\beta)}{\sigma_\alpha}\right)^2\right]. \end{aligned}$$

where

$$H_t = I_{n\times n} \otimes [1\ X_{t-1}' ... X_{t-q}'\ \ h(d_{t-1})X_{t-1}' ... h(d_{t-q})X_{t-q}']'$$

and

$$\beta = [\mu_1\ A_{11}^H ... A_{q1}^H ... D_{1n} ... D_{qn}]'.$$

The expression A_{ij}^H denotes the jth row vector of the matrix A_i^H, the objects $\Omega(d_{t-1})$ and $\kappa(d_{t-1})$ are defined as in (5) and (14)–(16), and $D_{ij} = A_{ij}^L - A_{ij}^H$. The set $\Phi = \{\beta, \delta, \sigma_\alpha, \Omega_H, \Omega_L, \overline{\kappa}, \psi\}$ collects all parameters of the log-likelihood. Given estimates for the elements of Φ, the impact vector of the tax shock can be easily recovered using (2), (13), and (14)–(16), and impulse response functions can be constructed for any desired horizon by iterating (4) forward.

The first-order conditions of the maximum likelihood problem yield closed-form expressions that describe $\{\beta, \delta, \sigma_\alpha\}$ in terms of $\{\Omega_H, \Omega_L, \overline{\kappa}, \psi\}$ (see Appendix C). Thus, maximizing (17) over $\{\Omega_H, \Omega_L, \overline{\kappa}, \psi\}$ while using those expressions to link $\{\beta, \delta, \sigma_\alpha\}$ to $\{\Omega_H, \Omega_L, \overline{\kappa}, \psi\}$ greatly simplifies the computation. However, the estimation procedure is still complicated by the possibility that the log-likelihood function (17) might have several local maxima, flat segments, and/or nonconvexities in certain regions of the parameter space. To address those issues, I use the Markov chain Monte Carlo (MCMC) method proposed in Chernozhukov and Hong (2003) and also adopted by Auerbach and Gorodnichenko (2012) to estimate the elements of the parameter set Φ and to simultaneously compute confidence intervals for the estimated parameters. In addition to yielding consistent estimates for the parameters, the method makes it easier to incorporate additional constraints that might reflect prior information or stem from supplementary identifying restrictions, although I do not impose any restrictions other than those implied by (7) and (13). Appendix C discusses the computational algorithm in greater detail.

4.3 Impulse Responses

The next steps are to evaluate the effects of tax shocks using the described identification scheme and to investigate how impulse responses vary with the state of the economy.

4.3.1 Benchmark Specification

The benchmark VAR involves a three-variable specification, $X_t = [\tau_t \; y_t \; l_t]'$, which includes a tax measure (τ_t), real GDP (y_t, chained in 2009 dollars), and total nonfarm business employment (l_t), all of which are entered as percentage changes from the previous period. That specification includes all core variables of the preceding LP-based analysis.

Because tax shocks simultaneously effect various components of the tax policy (such as revenues, marginal and average rates, and brackets), considering alternative tax measures can help to capture a broader range of dependencies. For that reason, I estimate two versions of the benchmark VAR: In the first version, the tax measure τ_t corresponds to real federal tax revenues in quarter t. In the second version, the tax measure is the average personal income tax rate (APTR) in quarter t. In the two versions, real revenues and the APTR are included sequentially (one at a time) into the estimated VAR, which helps to reduce the number of estimated parameters in each experiment.

Tax Revenues. To facilitate direct comparison with the literature that uses VARs to evaluate the effects of tax changes, I first estimate the change in GDP resulting from a shock to revenues. That is, I take the tax measure τ_t (in the three-variable vector $X_t = [\tau_t \; y_t \; l_t]'$) to correspond to cyclically adjusted real revenues defined as a percentage of real GDP.

Figure 5 displays the impulse responses of real GDP and employment to a tax shock that reduces revenues by 1% of GDP in the first quarter. Following Auerbach and Gorodnichenko (2012) and Bachman and Sims (2012), I compute impulse responses assuming that the state of the economy, as measured by the prevailing unemployment rate gap at the time of the tax change, remains fixed throughout the evaluated response period.[7] That is equivalent to ascribing a high degree of persistence to the unemployment rate gap so that it remains at its initial value for a certain number of periods following a tax change. Because that assumption rules out any feedback from impulse responses to the state of the economy and the probability of remaining at the initial state diminishes over time, impulse response estimates become increasingly less informative as the response horizon extends. For that reason, I compute the VAR-based impulse responses for a relatively small number of periods (8 quarters).

[7] Auerbach and Gorodnichenko (2012) also consider a case that allows for feedback from the state of the economy to the macroeconomic variables. Allowing for that feedback does not have a significant effect on their results.

To gauge the amount of uncertainty surrounding the central estimates, I also compute confidence intervals for the impulse response functions using the MCMC approach discussed earlier.[8] As in the LP-based analysis, confidence intervals mark bands of one standard error around each impulse response point estimate. Unlike the standard VAR-based and narrative-based approaches, however, the joint log-likelihood specification in (17) takes into account the uncertainty arising from the measurement and identification of exogenous tax changes (through the terms in the second row). That uncertainty is, therefore, also reflected in the estimated confidence bands.

Like the LP-based responses, the VAR-based responses show larger boosts in real GDP and employment in the lower-unemployment state, although the increases in both variables are estimated to occur faster than those suggested by the LP-based analysis. As shown in Table 2, VAR-based estimates of the average, maximum, and relative responses of GDP over 8 quarters suggest that the effect on output from a tax cut is stronger during times of lower unemployment. In particular, the maximum response of real GDP to a reduction in tax liabilities by 1 percent of GDP is roughly 1.8 percent in the lower-unemployment state (reached within 4 quarters after the shock) and around 1.2 percent in the higher-unemployment state (reached within 2 quarters after the shock).

The responses of employment are also larger and more persistent in the lower-unemployment state and more gradual than those of output in both states. The maximum increase in employment occurs after 6 quarters in the lower-unemployment state and after 4 quarters in the higher-unemployment state.

Average Personal Income Tax Rates. Next, I incorporate the average personal income tax rate (APTR) into the benchmark VAR. The average tax rate on personal income is the ratio of the sum of federal personal current taxes and contributions to government social insurance to the personal income tax base. The tax base for personal income is calculated as the national income and product accounts (NIPA) personal income (excluding government transfers) plus contributions for government social insurance.[9]

Figure 6 displays the responses of the APTR, real GDP, and nonfarm business employment to a tax cut that lowers the APTR by one percentage point. The tax cut generates an immediate positive response in real GDP and employment in both states. The boost in both variables is strong in the lower-unemployment state and becomes weaker in the higher-unemployment state. The observed pattern again suggests that the effects of a tax cut on real GDP and employment significantly vary across high- and low-unemployment states, becoming stronger when the unemployment rate gap is larger (that is, when there is increased slack in the economy at the time of the tax cut).

Although the size of the tax shock is set so that, in both states, the APTR is reduced by the

[8] See Appendix C for details.

[9] See the online data appendix for Mertens and Ravn (2013).

same amount (one percentage point) in the first quarter, the responses of the APTR in subsequent quarters differ considerably across states. The reduction in the APTR is estimated to be greater in the lower-unemployment state in the quarters following the initial shock. That outcome most likely reflects the stronger response of taxable incomes in the lower-unemployment state, which tends to make revenues a smaller percentage of the tax base and thus results in a more sizable decline in the APTR.

4.3.2 Alternative Specifications

To assess the sensitivity of the VAR-based results, I next reestimate the effects of tax changes under a series of alternative specifications.

Additional Variables. As in the LP-based analysis, I first incorporate additional variables into the estimated VAR. Increasing the number of variables in the VAR can be crucial for differentiating across different sources of variation more accurately and can help capture additional transmission channels. It is, however, also important that the number of variables remains small so that the estimation procedure is computationally manageable.

The extended VAR takes the form $X_t = [\tau_t \; y_t \; l_t \; d_t \; i_t]'$, where the additional variables d_t and i_t, respectively, denote real government purchases and the interest rate on 3-month U.S. Treasury bills. Including those variables into the VAR helps control for the interaction between taxes and spending and the role of monetary policy. Figure 7 displays the responses of the variables in the extended VAR to a tax cut that lowers revenues by 1 percent of GDP in higher- and lower-unemployment states. The estimated responses of employment in the lower-unemployment state are very similar to those of the benchmark case. The responses of output and employment, however, fade out faster than benchmark responses in the higher-unemployment state. As the estimated average, maximum, and relative responses of real GDP (shown in Table 2) suggest, the benchmark pattern remains unaltered; the boost in real GDP and employment generated by the tax cut is stronger when unemployment is lower.

Controlling for Anticipation. I now take into account the potential problem that may stem from conflating anticipated and unanticipated tax changes by eliminating from the sample the tax changes that took longer than 90 days to implement. As Figure 8 shows, there is no evidence that the larger output and employment effects of tax cuts in the lower-unemployment state are sensitive to excluding the tax changes with long implementation lags. Because the sample includes considerably fewer observations than the benchmark sample, however, estimated error bands (especially those surrounding the responses of employment) are markedly wider. In addition, the responses of output are substantially larger than the benchmark responses (see Table 2). That result is consistent with Mertens and Ravn's (2012) finding that anticipated tax cuts lead to an initial decline in output. (Thus, eliminating those tax cuts from the sample should result in larger output responses.)

A Level Specification. In the preceding analyses, variables are entered into the estimated VAR as percentage differences from the previous quarter. If the true process that governs the dynamics of the variables is a VAR in first differences, then differencing can enhance the small-sample performance of estimates (see Hamilton, 1994). Differencing can also yield misleading results if the true data-generating process does not admit a representation in first differences or if the variables are cointegrated.

To ensure that potential cointegrating relationships in the data are preserved, I estimate a version of the benchmark case by entering all variables in levels. Specifically, the estimated VAR takes the form $X_t = [\tau_t \; y_t \; l_t]'$, where τ_t denotes cyclically-adjusted revenues as a share of GDP and y_t and l_t, respectively, denote the log levels of real GDP and employment. I also include a deterministic time trend into the estimated VAR.

Figure 9 illustrates the responses to a tax change that reduces revenues by one percent of GDP in the first quarter. There are same noticeable differences between the levels and the benchmark cases. In the levels case, the estimated responses (especially those of employment) are quite similar across the two unemployment states in the first few quarters following the shock. The differences between the average and maximum output responses across the two states are also smaller than those estimated using the benchmark specification, as Table 2 shows. The overall pattern of state dependence, however, remains similar to that of the benchmark case: the boost in real GDP and employment is more pronounced in the lower-unemployment state. Although not reported here, using the APTR instead of revenues in the levels VAR yields similar results.

Jointly Estimating the Transition Function. The results discussed thus far are based on a particular choice for the parameter η that characterizes the transition function $h(d_t)$ defined in (2). Choosing a value for η, however, is difficult because of a lack of a clear-cut calibration target. A natural alternative to calibration is to estimate η simultaneously with other VAR parameters. With η added to the list of estimated parameters, however, the estimation problem becomes highly nonlinear. As a result, impulse responses may become highly sensitive to a few observations. Nevertheless, estimation can help to compensate for the lack of an obvious calibration target.

My estimation procedure returns the value 0.441 for the parameter η, which is less than a third of the calibration value (1.5). As mentioned in Auerbach and Gorodnichenko (2012), the smaller the value for η, the smoother the VAR system transitions between states. Thus, the estimated value of η suggests considerably smoother transition between unemployment states than what is implied by that parameter's calibration value.

Figure 10 displays the impulse responses that are computed using the estimated value of η. Despite the sizable difference between the estimated and calibrated values of η, the characteristics of the estimated responses are quite similar to those of the benchmark case. The estimates for the average, maximum, and relative responses of output in lower- and higher-unemployment states (listed in Table 2) are slightly larger than their benchmark counterparts. Like the benchmark values, however, they indicate larger effects on output and employment from a tax cut in the lower-unemployment state.

5 Conclusion

Effects of tax changes on output and employment can be different in different states of the economy. A common pattern that emerges from a variety of empirical specifications is that the effects on output and employment from a tax change tend to become smaller when unemployment is high; that is, in the presence of greater slack in the economy. Under the benchmark LP specification, for example, the estimated maximum boost in real GDP from a cut in revenues by 1% of GDP is about twice as large in the low-unemployment state as the boost a similar cut would induce in the high-unemployment state. That pattern remains largely intact under several extensions of the benchmark LP model and is also insensitive to using a VAR-based method to estimate impulse responses.

The size of the estimated responses to a tax change in a particular state varies significantly across different model specifications ranging from modest to strikingly large. But, in the majority of the examined cases, the estimated responses are quite sizable in a state with low unemployment. Although large output responses are consistent with the findings of a number of narrative-based studies (including Romer and Romer, 2010; and Mertens and Ravn, 2013), large employment responses are difficult to reconcile with the evidence provided by the research literature that examines how pre-tax incomes react to changes in marginal tax rates. (See Saez and others, 2012, for a review of that literature.) With some exceptions, those studies conclude that labor supply (as measured by the reported pre-tax incomes) responds very little to changes in marginal tax rates—a result that runs contrary to the large employment effects found in this analysis. However, the main conclusions of this study depend on how responses vary across states rather than how large those responses are in a particular state; in that regard, the evaluated cases invariably suggest that effects are smaller when unemployment is higher.

The current analysis does not offer a full theoretical framework to examine why the short-term effects of tax changes may become smaller during times of slack—I investigate the theoretical underpinnings of that finding in ongoing work. Intuitively, however, those effects may occur because changes in the supply of productive inputs (in particular, of labor) are likely to have a smaller effect on resource utilization when there is substantial slack in the economy. For example, in times of high unemployment and severe job rationing, stronger work incentives brought about by lower taxes (or, conversely, weaker work incentives because of higher taxes) are unlikely to effect actual employment as much as they would when there is more intense competition among employers for workers.

One limitation of this analysis is that when the number of observations for a particular unemployment state is small and the length of time the economy remains in that states is short, the precision of estimates and accuracy of inference can be significantly reduced. As discussed earlier, our benchmark state specifications are chosen with that consideration in mind. Adopting a smooth-transition framework (as opposed to conjecturing definite thresholds at which the state of the economy shifts abruptly) mitigates that problem by facilitating a continuous transition between states. Another issue stems from the lack of a universal measure for the state of the economy. Because this study concentrates on the implications of tax changes on output and employment, the adopted state measure mainly reflects the prevailing conditions in the labor market. Given those limitations, a potential direction for future research involves adopting more targeted measures of state and evaluating the degree of being in a particular state using those measures.

Appendix A

This appendix discusses the relationship between my identification approach and that of Mertens and Ravn (2013). First, consider the partition $E_t[z_t u_t'] = [\Sigma(d_{t-1})_{zu_1'} \ \Sigma(d_{t-1})_{zu_2'}]$. The scalar $\Sigma(d_{t-1})_{zu_1'}$ gives the covariance of z_t with the first element of u_t'. The row vector $\Sigma(d_{t-1})_{zu_2'}$ collects the covariances of z_t with the remaining elements of u_t'. Thus, Equation 10 can be written as

$$[\Sigma(d_{t-1})_{zu_1'} \Sigma(d_{t-1})_{zu_2'}] = (1-\delta)\psi \left[R(d_{t-1})_{11}' \ R(d_{t-1})_{21}'\right]. \tag{18}$$

Equation 12 suggests that

$$[\Sigma(d_{t-1})_{zu_1'} \Sigma(d_{t-1})_{zu_2'}] = (1-\delta)\psi\kappa(d_{t-1})\Omega(d_{t-1}). \tag{19}$$

Thus, given estimates for the state-dependent objects $\kappa(d_{t-1})$ and $\Omega(d_{t-1})$, and provided that $\Sigma(d_{t-1})_{zu_1'}$ (which, in this case, is a scalar) is invertible, Equations 8 and 18 can be combined to obtain

$$R(d_{t-1})_{21} = \left[\Sigma(d_{t-1})^{-1}_{zu_1'} \Sigma(d_{t-1})_{zu_2'}\right]' R(d_{t-1})_{11}. \tag{20}$$

Equation 20 provides a set of state-dependent covariance restrictions on the elements of the first column of $R(d_{t-1})$. Those restrictions can be combined with the ones that follow from (7) to derive the first element of the impact vector as

$$\begin{aligned} R(d_{t-1})_{11} \ = \ & \{\Omega(d_{t-1})_{11} - [\Omega(d_{t-1})_{12} - \Omega(d_{t-1})_{11}\Lambda(d_{t-1})'] \times [\Lambda(d_{t-1})\Omega(d_{t-1})_{11}\Lambda(d_{t-1})' \\ & -\Omega(d_{t-1})_{12}'\Lambda(d_{t-1})' - \Lambda(d_{t-1})\Omega(d_{t-1})_{12} + \Omega(d_{t-1})_{22}]^{-1} \\ & \times[\Omega(d_{t-1})_{12}' - \Lambda(d_{t-1})\Omega(d_{t-1})_{11}]\}^{\frac{1}{2}}, \end{aligned} \tag{21}$$

where

$$\Lambda(d_{t-1}) = \left[\Sigma(d_{t-1})^{-1}_{zu_1'} \Sigma(d_{t-1})_{zu_2'}\right]' \tag{22}$$

and $\Omega(d_{t-1})_{ij}$ denotes an appropriate partitioning of (5) that conforms with (8). Once an estimate for the first element of the impact vector is obtained from (21), the remaining elements can be subsequently recovered using (20).

Mertens and Ravn (2013) derive an analytical solution for the impact vector of tax shocks for a general case in which z_t is a k dimensional vector. That solution, however, is based on a linear VAR. Equation 21 gives a state-dependent reformulation of that solution for the case $k = 1$.

Appendix B

This appendix describes the derivation of the functional form of $\kappa(d_{t-1})$. First, the VAR residual covariance (5) can be reexpressed in the form

$$\Omega(d_{t-1}) = \Omega_H + h(d_{t-1})\Omega_D, \tag{23}$$

where $\Omega_D = \Omega_L - \Omega_H$. Also, post-multiplying both sides of the definition $\kappa(d_{t-1}) = \gamma R(d_{t-1})^{-1}$ with $R(d_{t-1})$ gives $\kappa(d_{t-1})R(d_{t-1}) = \gamma$, which, together with (7), implies that

$$\kappa(d_{t-1})\Omega(d_{t-1})\kappa(d_{t-1})' = 1. \tag{24}$$

Next, the function $\kappa(d_{t-1})$ is guessed to be of the form

$$\kappa(d_{t-1}) = g(d_{t-1})\overline{\kappa},$$

where $\overline{\kappa}$ is a $(1 \times n)$ vector of constants. Incorporating that guess and (23) into (24) produces the relationship

$$g(d_{t-1})^2[\overline{\kappa}\Omega_H\overline{\kappa}' + h(d_{t-1})\overline{\kappa}\Omega_D\overline{\kappa}'] = 1. \tag{25}$$

Plugging the definition (2) into (25) and reorganizing yields

$$g(d_{t-1}) = \left[\frac{1 + \exp(-\eta d_{t-1})}{\overline{\kappa}\Omega_H\overline{\kappa}' + (\overline{\kappa}\Omega_H\overline{\kappa}' + \overline{\kappa}\Omega_D\overline{\kappa}')\exp(-\eta d_{t-1})}\right]^{\frac{1}{2}},$$

which verifies the initial guess and leads to (15) and (16) in the main text. Observe that incorporating the guess $\kappa(d_{t-1}) = -g(d_{t-1})\overline{\kappa}$ into (24) also results in (25). Therefore, the form of $\kappa(d_{t-1})$ is identified only up to a sign convention because (24) is equally consistent with the guesses $\kappa(d_{t-1}) = g(d_{t-1})\overline{\kappa}$ and $\kappa(d_{t-1}) = -g(d_{t-1})\overline{\kappa}$.

Appendix C

This appendix provides the derivation of the log-likelihood function (17), drives the analytical first-order conditions, and discusses the computational algorithm followed to implement the estimation approach.

Equation 4 can be written more compactly in the form

$$X_t = H_t'\beta + u_t, \tag{26}$$

with $H_t = I_{n\times n} \otimes [1\ X_{t-1}' ... X_{t-q}'\ \ h(d_{t-1})X_{t-1}' ... h(d_{t-q})X_{t-q}']'$ and $\beta = [\mu_1\ A_{11}^H ... A_{q1}^H ... D_{1n} ... D_{qn}]'$, where A_{ij} denotes the jth row vector of the matrix A_i and $D_{ij} = A_{ij}^L - A_{ij}^H$. Incorporating (26) into (11) yields

$$z_t = \theta_t[\psi\kappa(d_{t-1})(X_t - H_t'\beta) + \alpha_t]. \tag{27}$$

Given (26) and (27), the initial value H_1, and a series for the unemployment rate gap index $(d_{T-1}, ..d_0)$, the joint conditional density of the sample $(X_T, z_T, ..X_1, z_1)$, expressed as a product of conditional densities, can be written as

$$f(X_T, z_T, ..X_1, z_1|d_{T-1}, ..d_0, H_1) = \prod_{t=1}^{T} f(z_t|X_t, H_t, d_{t-1})f(X_t|H_t, d_{t-1}).$$

Therefore, the log-likelihood of the sample is

$$\log \mathcal{L}(\Phi) = \sum_{t=1}^{T}[\log f(z_t|X_t, H_t, d_{t-1}; \Phi) + \log f(X_t|H_t, d_{t-1}; \Phi)],$$

where $\Phi = \{\beta, \delta, \sigma_\alpha, \Omega_H, \Omega_L, \overline{\kappa}, \psi\}$ denotes the set of parameters. Provided that the censoring process is independent of the measurement error and the structural shocks, and measurement errors are distributed independently over time, the conditional density function for z_t can be written as

$$f(z_t|X_t, H_t, d_{t-1}; \Phi) = \delta^{1-\theta_t}\left\{(1-\delta)\left(\frac{1}{\sigma_a\sqrt{2\pi}}\right)\exp\left[-\frac{z_t - \psi\kappa(d_{t-1})(X_t - H_t'\beta)}{2\sigma_a}\right]^2\right\}^{\theta_t}, \tag{28}$$

where δ is the probability of the event $\theta_t = 0$. Likewise, the conditional density for X_t can be written as

$$f(X_t|H_t, d_{t-1}; \Phi) = (2\pi)^{-n/2}\left|\Omega(d_{t-1})^{-1}\right|^{1/2}\exp\{-1/2(X_t - H_t'\beta)'\Omega(d_{t-1})^{-1}(X_t - H_t'\beta)\} \tag{29}$$

Taking the logs of (28) and (29) and summing over all observations produces

$$\log \mathcal{L}(\Phi) = \text{constant} + \frac{1}{2}\sum_{t=1}^{T}\left[\log\left|\Omega(d_{t-1})^{-1}\right| - (X_t - H_t'\beta)'\Omega(d_{t-1})^{-1}(X_t - H_t'\beta)\right] \quad (30)$$

$$+\sum_{z_t=0}\log\delta + \sum_{z_t\neq 0}\left[\log(1-\delta) - \log\sigma_\alpha - \frac{1}{2}\left(\frac{z_t - \psi\kappa(d_{t-1})(X_t - H_t'\beta)}{\sigma_\alpha}\right)^2\right],$$

which corresponds to the sample log-likelihood (17). Given that the narrative tax changes and VAR residuals are orthogonal to H_t' (that is, $E[z_t H_t'] = E[u_t H_t'] = 0$), the first-order conditions with respect to β, σ_a, and δ yield

$$\beta = \left[\sum_{t=1}^{T} H_t\Omega(d_{t-1})^{-1}H_t'\right]^{-1}\left[\sum_{t=1}^{T} H_t\Omega(d_{t-1})^{-1}X_t\right], \quad (31)$$

$$\sigma_\alpha^2 = (1/K)\sum_{z_t\neq 0}[z_t - \kappa(d_{t-1})(X_t - H_t'\beta)]^2, \text{ and} \quad (32)$$

$$\delta = (T-K)/T, \quad (33)$$

where $K > 0$ denotes the number of non-zero observations for z_t.

Given a set of values for $\{\Omega_H, \Omega_L, \overline{\kappa}, \psi\}$, estimates for $\{\beta, \sigma_\alpha, \delta\}$ can be obtained from (31)–(33). Therefore, parameters can be jointly estimated by iterating on $\{\Omega_H, \Omega_L, \overline{\kappa}, \psi\}$ to maximize (30) while using (31)–(33) to recover $\{\beta, \sigma_\alpha, \delta\}$. Because the log-likelihood function (30) is highly nonlinear in $\{\Omega_H, \Omega_L, \overline{\kappa}, \psi\}$, however, several local maxima might exist. To avoid being stuck at a local maximum, I use the Markov chain Monte Carlo (MCMC) procedure proposed in Chernozhukov and Hong (2003). I follow Auerbach and Gorodnichenko (2012) in adopting the Hastings-Metropolis algorithm to implement the procedure, which involves three steps:

1. Iterate on $\{\Omega_H, \Omega_L, \overline{\kappa}, \psi\}$ to maximize the log-likelihood function (30) subject to (31)–(33), using a standard numerical optimization routine. Denote the vector of parameters that achieves that maximum $\Phi^{(0)}$.

2. Adopting $\Phi^{(0)}$ as the initial value, draw a vector of shocks $\Lambda^{(i)}$ from the distribution $N(0, \Sigma_\Phi)$, where Σ_Φ is a diagonal matrix and calculate a candidate vector of parameters as $\Gamma^{(i)} = \Phi^{(i)} + \Lambda^{(i)}$.

3. Determine the $n+1$ state of the Markov chain by following the rule

$$\Phi^{(i+1)} = \begin{cases} \Gamma^{(i)} \text{ with probability } \min\{1, e^{\log\mathcal{L}(\Gamma^{(i)}) - \log\mathcal{L}(\Phi^{(i)})}\} \\ \Phi^{(i)} \text{ with probability } 1 - \min\{1, e^{\log\mathcal{L}(\Gamma^{(i)}) - \log\mathcal{L}(\Phi^{(i)})}\}, \end{cases}$$

where $\log \mathcal{L}(\Gamma^{(i)})$ and $\log \mathcal{L}(\Phi^{(i)})$, respectively, denote the log-likelihood functions evaluated at the candidate and current states of the chain.

Following Auerbach and Gorodnichenko (2012), I set Σ_Φ so that the acceptance rate for candidate states is roughly 30% and, following the practice in Gelman and Rubin (1992), I drop the first half of the draws to exclude the "burn-in" period.

Chernozhukov and Hong (2003) show that $(1/I)\sum_{q=1}^{I} \Phi^{(q)}$ provides a consistent estimate for the parameter vector Φ under a set of regularity conditions. Following Auerbach and Gorodnichenko (2012), I compute error bands by generating random draws (with replacement) from the chain $\{\Phi^{(q)}\}_{q=1}^{I}$ and computing impulse response functions for each draw. The reported error bands are constructed using the standard deviations of point responses in a sample of $1,000$ independent draws from the simulated chain.

References

[1] Auerbach, A., Gorodnichenko, Y., 2013. Fiscal multipliers in recession and expansion. In: Alesina, A., Giavazzi, F. (Eds.), Fiscal Policy after the Financial Crisis 2013, 19–55. University of Chicago Press.

[2] Auerbach, A., Gorodnichenko, Y., 2012. Measuring the output responses to fiscal policy. American Economic Journal: Economic Policy 4 (2), 1–27.

[3] Bachmann, R., Sims, E., 2012. Confidence and the transmission of government spending shocks. Journal of Monetary Economics 59, 235–249.

[4] Barro, R. J., Redlick, C. J., 2011. Macroeconomic effects from government purchases and taxes. Quarterly Journal of Economics 126, 51–102.

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

[6] Chernozhukov, V., Hong, H., 2003. An MCMC approach to classical estimation. Journal of Econometrics 115, 293–346.

[7] Christiano, L., Eichenbaum, M., Rebelo, S., 2011. When is the government spending multiplier large? Journal of Political Economy 119 (1), 78–121.

[8] Cogan, J. F., Cwik, T., Taylor, J. B., Wieland, V., 2010. New Keynesian versus old Keynesian government spending multipliers. Journal of Economic Dynamics and Control 34 (3), 281–295.

[9] Congressional Budget Office, 2014a. How CBO Analyzes the Effects of Changes in Federal Fiscal Policies on the Economy. www.cbo.gov/publication/49494.

[10] Congressional Budget Office, 2014b. Revisions to CBO's Projection of Potential Output Since 2007. www.cbo.gov/publication/45150.

[11] Favero, C., Giavazzi, F., 2012. Measuring tax multipliers: the narrative method in fiscal VARs. American Economic Journal: Economic Policy 4 (2), 69–94.

[12] Fazzari, S. M., Morley, J., Panovska, I., 2013. State dependent effects of fiscal policy. Unpublished paper, Australian School of Business.

[13] Francis, N., Ramey, V. A., 2009. Measures of per capita hours and their implications for the technology-hours debate. Journal of Money, Credit and Banking 41, 1071–1097.

[14] Gelman, A., Rubin, D. B. 1992. Inference from iterative simulation using multiple sequences (with discussion). Statistical Science 7, 457–511.

[15] Hamilton, J. D., 1994. Time Series Analysis. Princeton University Press, Princeton.

[16] Jorda, O., 2005. Estimation and inference of impulse responses by local projections. American Economic Review 95, 161–182.

[17] Leeper, E. M., Walker, T. B., Yang., S. S., 2013. Fiscal foresight and information flows. Econometrica 81 (3), 1115–1145.

[18] Mertens, K., Ravn, M., 2013. The dynamic effects of personal and corporate income tax changes in the United States. American Economic Review 103 (4), 1212–1247.

[19] Mertens, K., Ravn, M., 2012. Empirical evidence on the aggregate effects of anticipated and unanticipated US tax policy shocks. American Economic Journal: Economic Policy 4 (2), 145–181.

[20] Mittnik, S., Semmler, W., 2012. Regime dependence of the fiscal multiplier. Journal of Economic Behavior & Organization 83, 502–522.

[21] Newey, W. K., West, K. D., 1987. A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. Econometrica 55, 703–708.

[22] Owyang, M. T., Ramey V. A., Zubairy, S., 2013. Are government spending multipliers greater during periods of slack? Evidence from twentieth-century historical data. American Economic Review: Papers & Proceedings 103 (3), 129–134.

[23] Parker, J., A., 2011. On measuring the effects of fiscal policy in recessions. Journal of Economic Literature 49 (3), 703–718.

[24] Perotti, R., 2012. The effects of tax shocks on output: not so large, but not small either. American Economic Journal: Economic Policy 4 (2), 214–237

[25] Ramey, V. A., 2012. Comment on Roads to Prosperity or Bridges to Nowhere? Theory and evidence on the impact of public infrastructure investment. In NBER Macroeconomics Annual 2012, Volume 27, 147–153. National Bureau of Economic Research.

[26] Ramey, V. A., 2011. Identifying government spending shocks: it's all in the timing. Quarterly Journal of Economics 126 (1), 1–50.

[27] Ramey, V. A., Zubairy, S., 2014. Government spending multipliers in good times and in bad: evidence from U.S. historical data. NBER Working Paper No. 20719, National Bureau of Economic Research.

[28] Reichling, F., Whalen, C., 2015. The Fiscal Multiplier and Economic Policy Analysis in the United States. Working Paper 2015-02, Congressional Budget Office.

[29] Romer, C., Romer, D., 2010. The macroeconomic effects of tax changes: estimates based on a new measure of fiscal shocks. American Economic Review 100 (3), 763–801.

[30] Rossi, B., Zubairy, S., 2011. What is the importance of monetary and fiscal shocks in explaining US macroeconomic fluctuations? Journal of Money, Credit, and Banking 43 (6), 1247–1270.

[31] Saez, E., Slemrod, J., Giertz, S., 2012. The elasticity of taxable income with respect to marginal tax rates: a critical review. Journal of Economic Literature 50 (1), 3–50.

[32] Stock, J., Watson, M., 2007. Why has U.S. inflation become harder to forecast? Journal of Money, Credit, and Banking 39, 3–33.

[33] Stock, J., Watson, M., 2003. Forecasting output and inflation: the role of asset prices. Journal of Economic Literature 41, 788–829.

[34] Tagkalakis, A., 2008. The effect of fiscal policy on consumption in recessions and expansions. Journal of Public Economics 92, 1486–1508.

[35] Woodford, M., 2011. Simple analytics of the government expenditure multiplier. American Economic Journal: Macroeconomics 3 (1), 1–35.

[36] Yang, S. S., 2005. Quantifying tax effects under policy foresight. Journal of Monetary Economics 52, 1557–1568.

Tables and Figures

Table 1. Responses of Real GDP to a Reduction in Tax Liabilities by 1% of GDP Estimated Using the Local Projection (LP) Method.

	Average: $\frac{1}{N}\Sigma_{i=1}^{N} y_i$	Maximum: $\max_{i=1..N}\{y_i\}$	Relative: $\frac{\Sigma_{i=1}^{N} y_i}{\Sigma_{i=1}^{N} \tau_i}$
Benchmark			
Lower Unemployment	1.36	2.52	−3.50
Higher Unemployment	0.32	1.25	−0.87
Additional Controls			
Lower Unemployment	0.72	1.19	−1.86
Higher Unemployment	0.16	0.56	−0.43
No Anticipation			
Lower Unemployment	0.94	1.67	−2.93
Higher Unemployment	0.29	1.14	−0.56
Threshold			
Lower Unemployment	1.59	2.89	−3.29
Higher Unemployment	0.50	1.67	−1.40

Note: Variables y_i and τ_i denote the percent changes (in the i^{th} quarter) of real GDP and revenues, respectively, from what would occur without the reduction in tax liabilities.

Table 2. Responses of Real GDP to a Reduction in Tax Liabilities by 1% of GDP Estimated Using the Regime-Switching Vector Autoregression (VAR) Method.

	Average: $\frac{1}{N}\Sigma_{i=1}^{N} y_i$	Maximum: $\max_{i=1..N}\{y_i\}$	Relative: $\frac{\Sigma_{i=1}^{N} y_i}{\Sigma_{i=1}^{N} \tau_i}$
Benchmark			
Lower Unemployment	1.57	1.79	-2.57
Higher Unemployment	0.88	1.22	-1.50
Additional Controls			
Lower Unemployment	1.65	1.98	-2.87
Higher Unemployment	0.85	1.60	-1.35
No Anticipation			
Lower Unemployment	3.34	3.57	-3.27
Higher Unemployment	2.47	2.74	-2.43
Levels Specification			
Lower Unemployment	1.44	1.62	-2.67
Higher Unemployment	1.10	1.43	-1.89
Estimating η			
Lower Unemployment	1.74	1.97	-2.66
Higher Unemployment	1.18	1.41	-1.94

Note: Variables y_i and τ_i denote the percent changes (in the i^{th} quarter) of real GDP and revenues, respectively, from what would occur without the reduction in tax liabilities.

Figure 1. LP-Based Responses of Real GDP and Employment to a Reduction in Tax Liabilities by 1% of GDP Evaluated Under Alternative Values of the Unemployment Rate Gap.

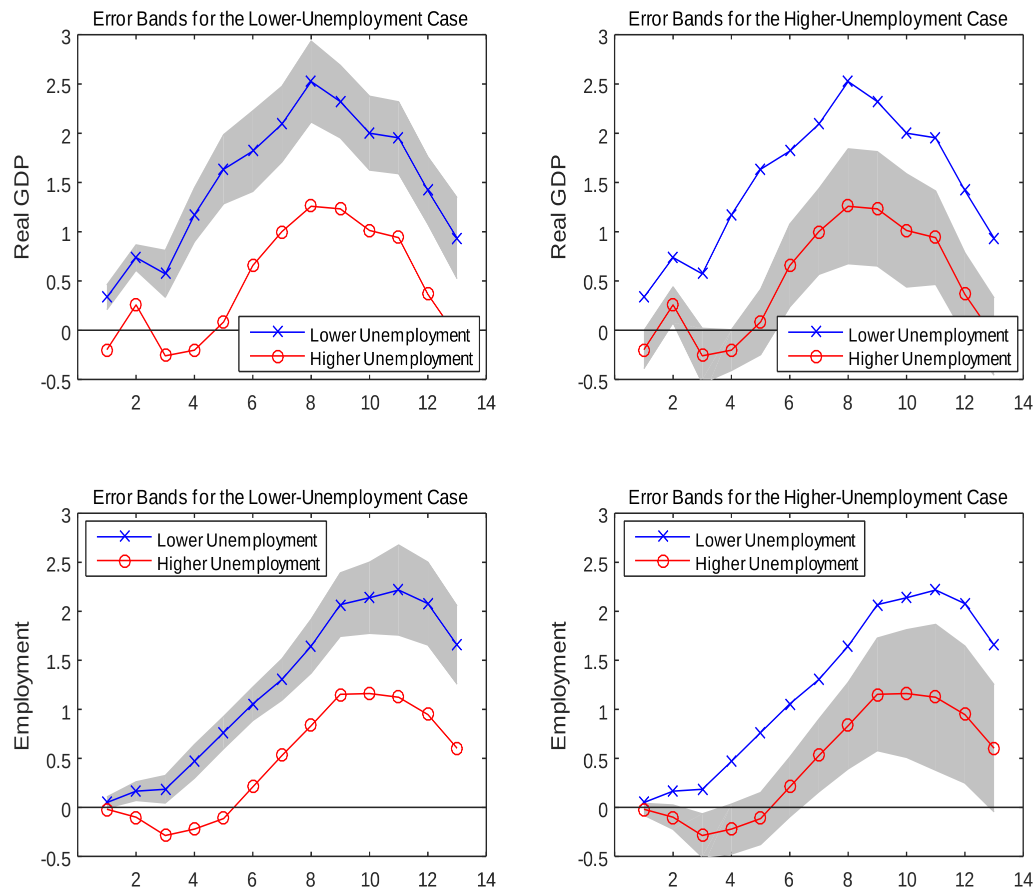

Note: In each panel, the vertical axis shows the estimated percentage change in each variable in response to a tax shock and the horizontal axis denotes the quarter after the shock. "Lower Unemployment" and "Higher Unemployment" indicate states in which the unemployment rate gap equals 0.5 and 1.5 percent, respectively. Shaded areas mark plus and minus one-standard-error bands.

Figure 2. LP-Based Responses of Real GDP and Employment to a Reduction in Tax Liabilities by 1% of GDP Estimated Using Additional Control Variables.

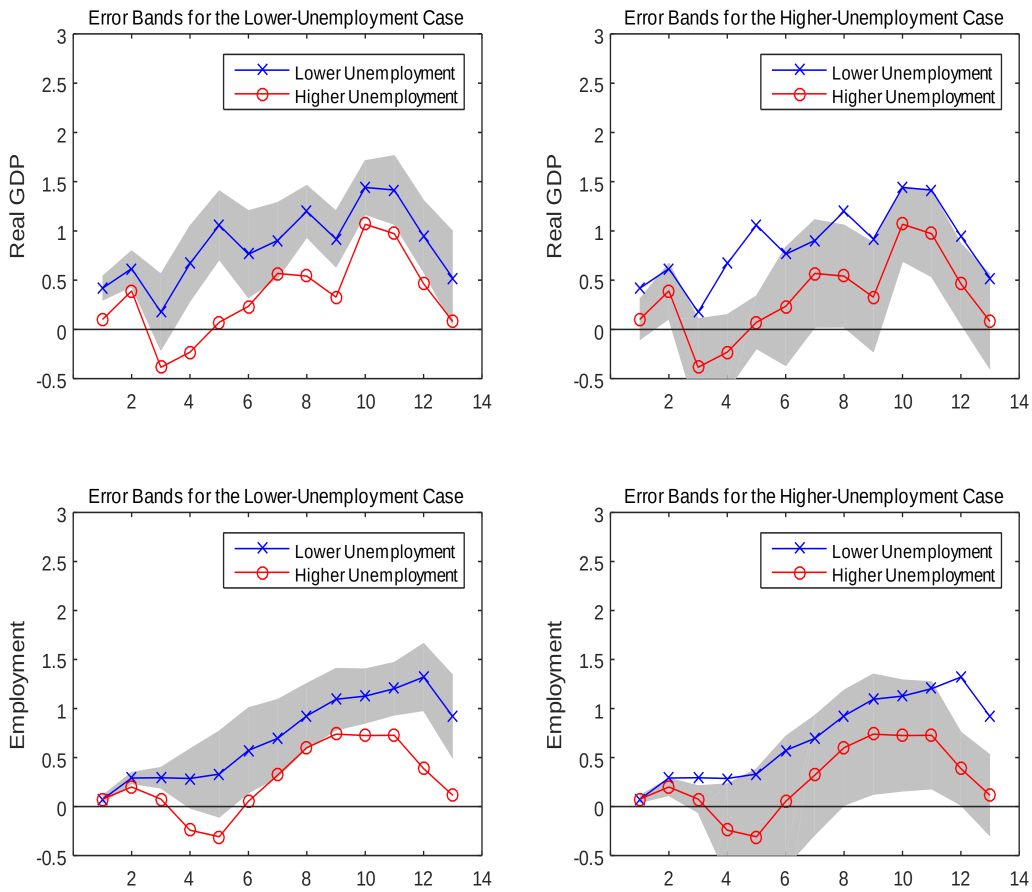

Note: In each panel, the vertical axis shows the estimated percentage change in each variable in response to a tax shock and the horizontal axis denotes the quarter after the shock. "Lower Unemployment" and "Higher Unemployment" indicate states in which the unemployment rate gap equals 0.5 and 1.5 percent, respectively. Shaded areas mark plus and minus one-standard-error bands.

Figure 3. LP-Based Responses Estimated by Excluding From the Sample the Tax Changes with Implementation Lags That Exceed 90 Days.

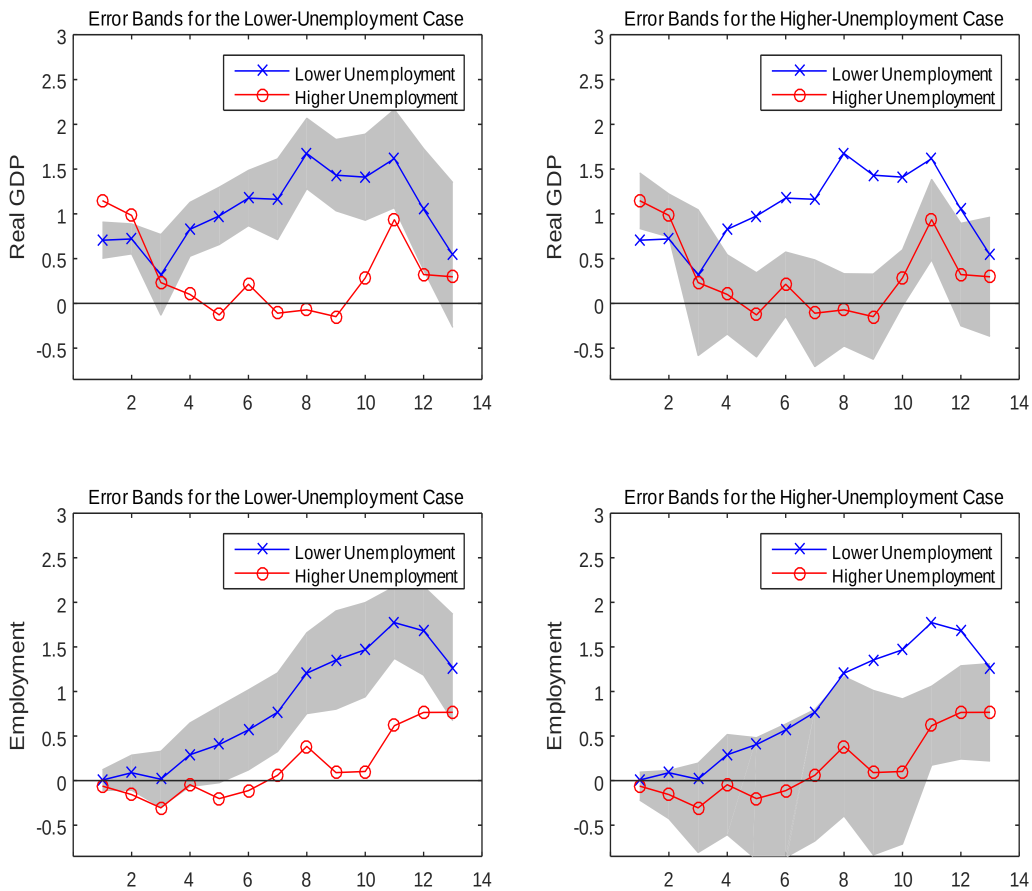

Note: In each panel, the vertical axis shows the estimated percentage change in each variable in response to a tax shock and the horizontal axis denotes the quarter after the shock. "Lower Unemployment" and "Higher Unemployment" indicate states in which the unemployment rate gap equals 0.5 and 1.5 percent, respectively. Shaded areas mark plus and minus one-standard-error bands.

Figure 4. LP-Based Responses Estimated Using the Threshold Specification.

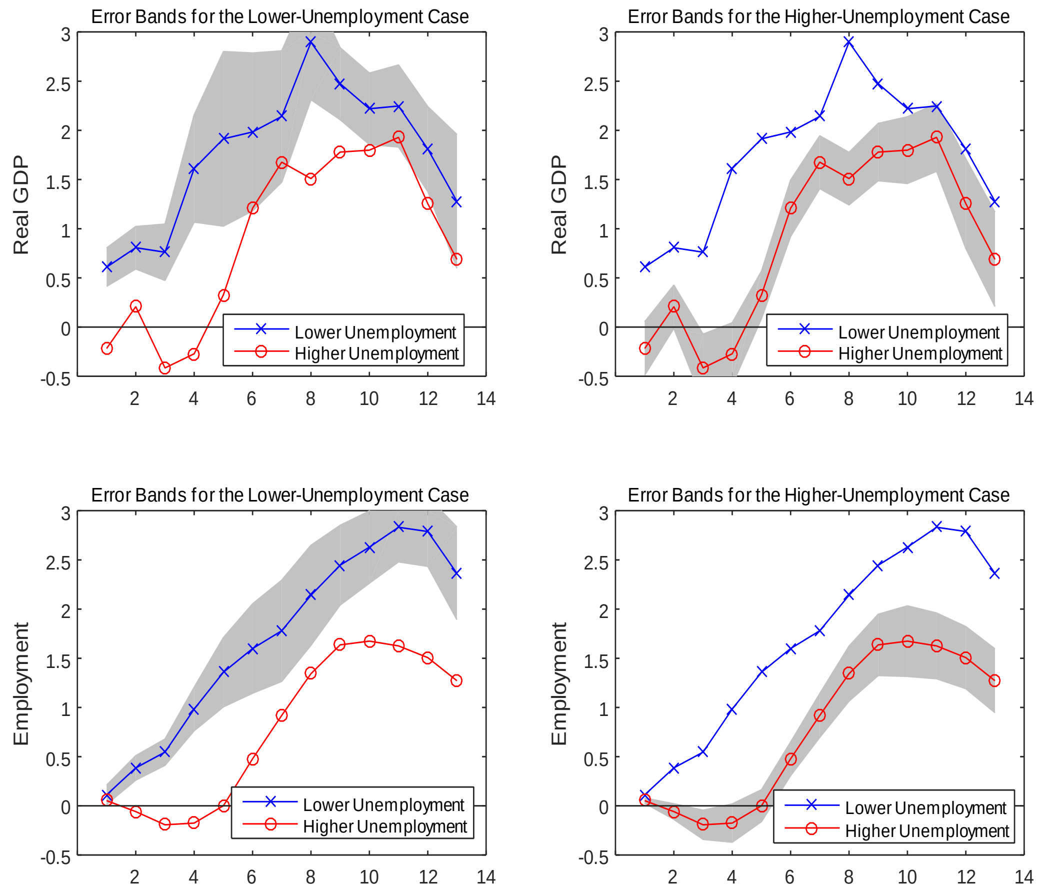

Note: In each panel, the vertical axis shows the estimated percentage change in each variable in response to a tax shock and the horizontal axis denotes the quarter after the shock. "Lower Unemployment" and "Higher Unemployment" indicate states in which the unemployment rate gap equals 0.5 and 1.5 percent, respectively. Shaded areas mark plus and minus one-standard-error bands.

Figure 5. VAR-Based Responses of Real GDP and Employment to a Reduction in Tax Revenues by 1% of GDP Evaluated Under Alternative Values of the Unemployment Rate Gap.

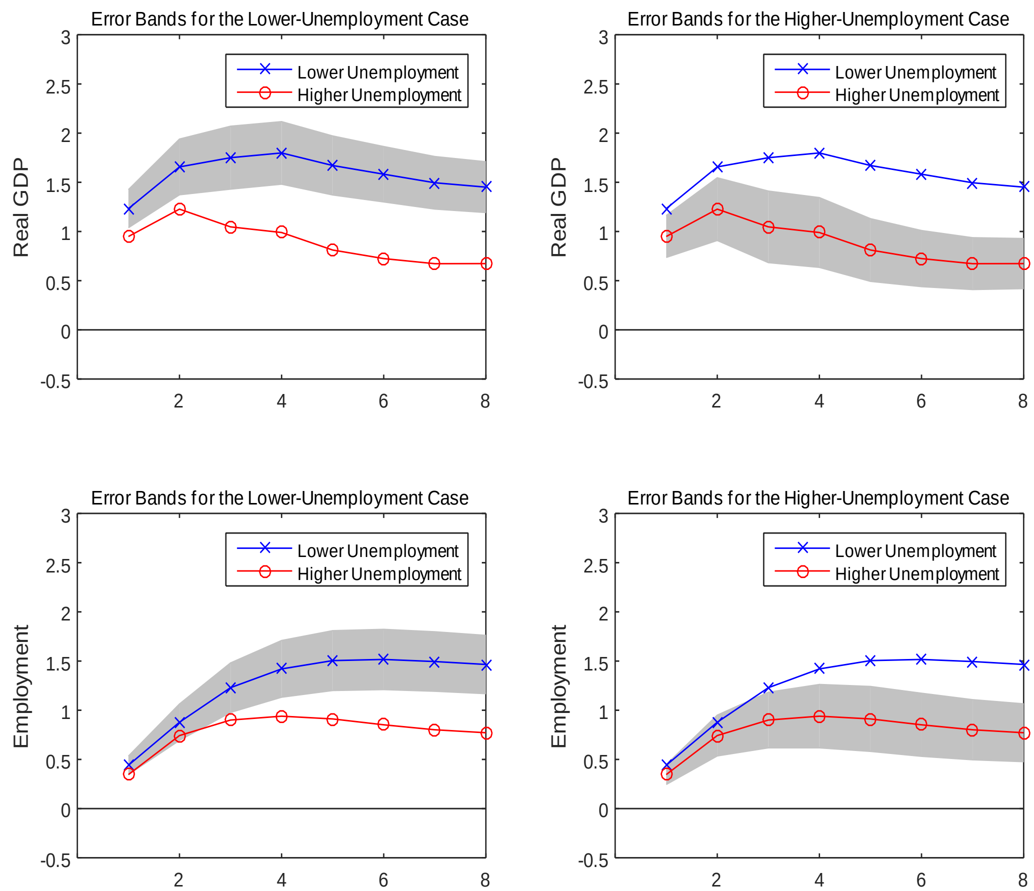

Note: In each panel, the vertical axis shows the estimated percentage change in each variable in response to a tax shock and the horizontal axis denotes the quarter after the shock. "Lower Unemployment" and "Higher Unemployment" indicate states in which the unemployment rate gap equals 0.5 and 1.5 percent, respectively. Shaded areas mark plus and minus one-standard-error bands.

Figure 6. VAR-based Responses of Real GDP and Employment to a Shock That Reduces the Average Personal Income Tax Rate (APTR) by 1 Percentage Point in the First Quarter.

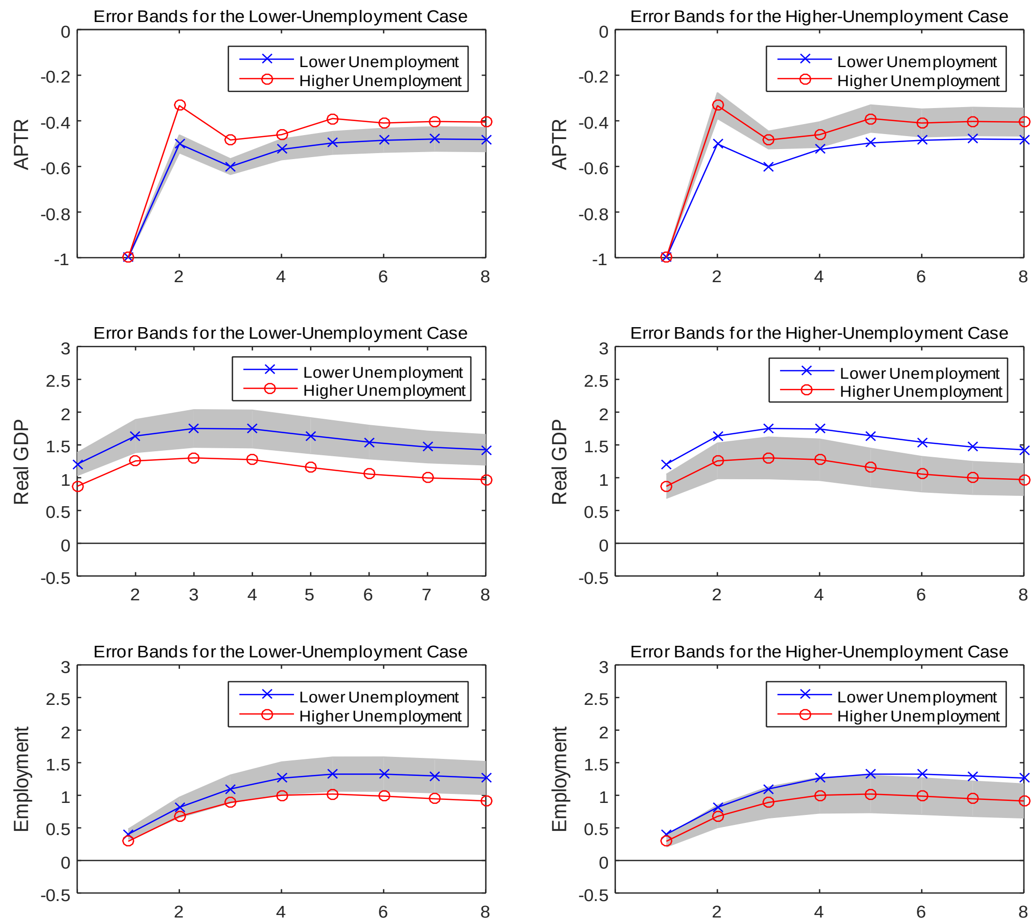

Note: In each panel, the vertical axis shows the estimated percentage change in each variable in response to a tax shock and the horizontal axis denotes the quarter after the shock. "Lower Unemployment" and "Higher Unemployment" indicate states in which the unemployment rate gap equals 0.5 and 1.5 percent, respectively. Shaded areas mark plus and minus one-standard-error bands.

Figure 7. VAR-Based Responses Estimated Using Additional Control Variables.

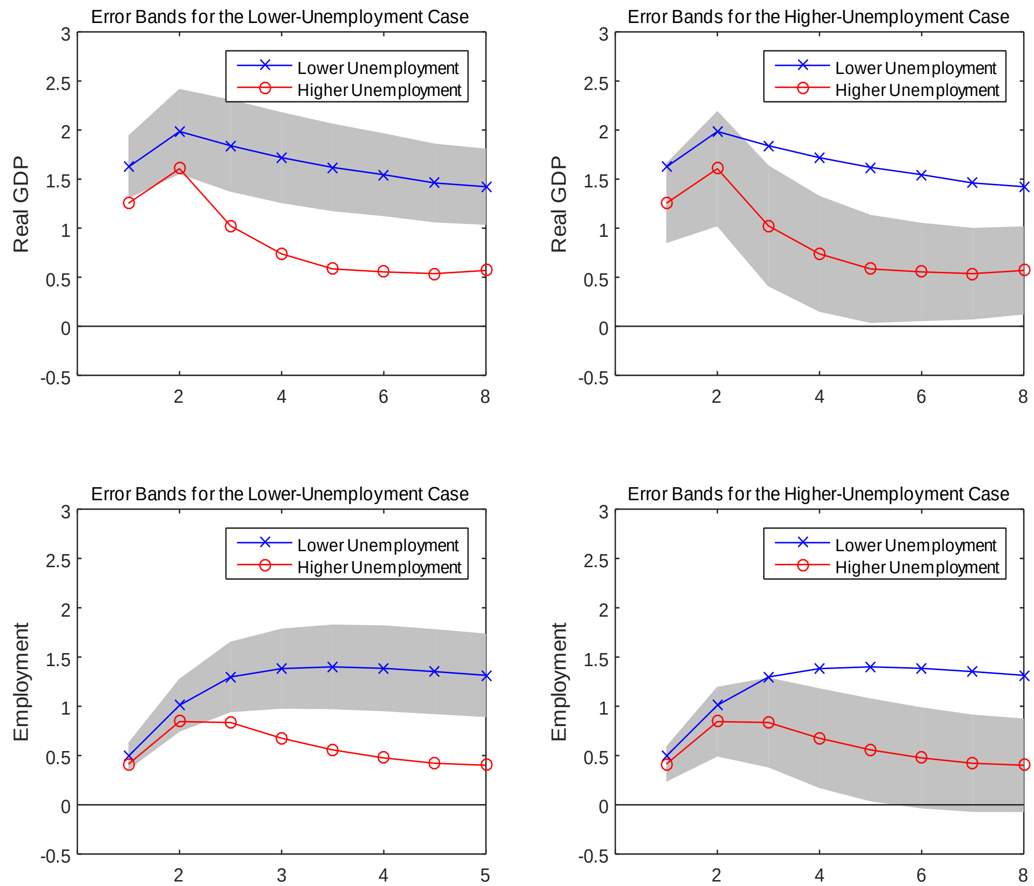

Note: In each panel, the vertical axis shows the estimated percentage change in each variable in response to a tax shock and the horizontal axis denotes the quarter after the shock. "Lower Unemployment" and "Higher Unemployment" indicate states in which the unemployment rate gap equals 0.5 and 1.5 percent, respectively. Shaded areas mark plus and minus one-standard-error bands.

Figure 8. VAR-Based Responses Estimated by Excluding From the Sample the Tax Changes with Implementation Lags That Exceed 90 Days.

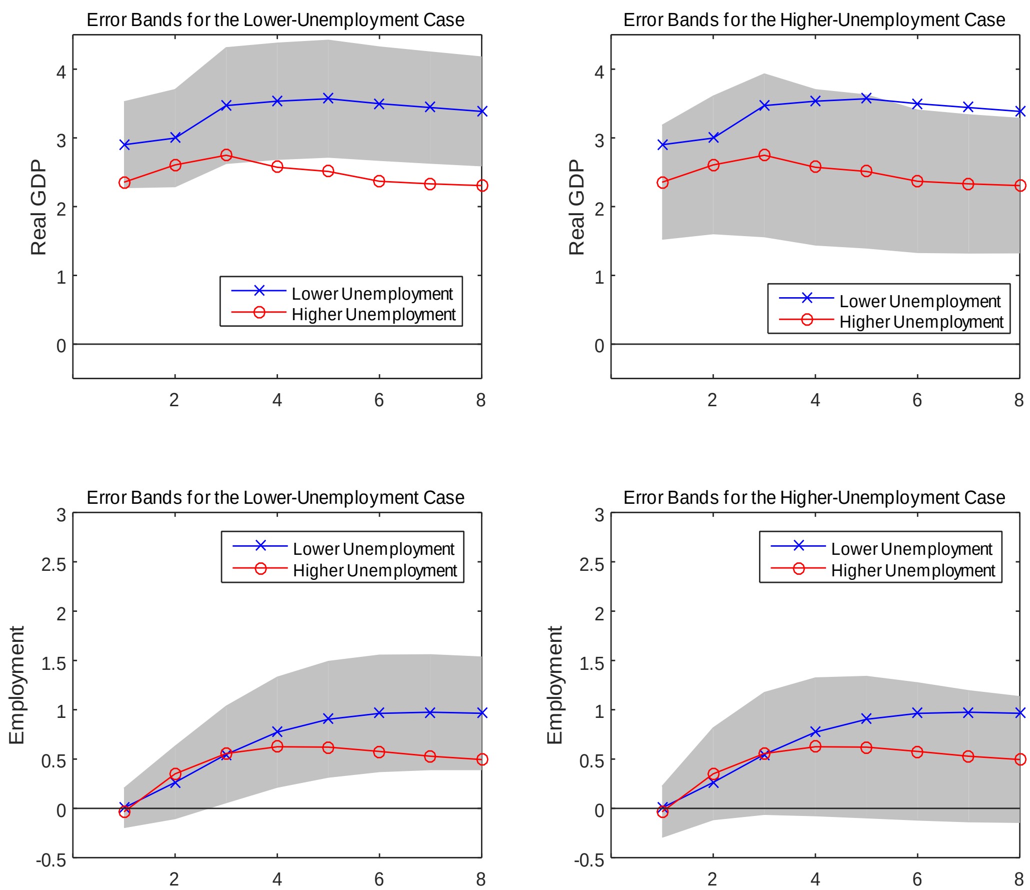

Note: In each panel, the vertical axis shows the estimated percentage change in each variable in response to a tax shock and the horizontal axis denotes the quarter after the shock. "Lower Unemployment" and "Higher Unemployment" indicate states in which the unemployment rate gap equals 0.5 and 1.5 percent, respectively. Shaded areas mark plus and minus one-standard-error bands.

Figure 9. VAR-Based Responses Estimated Using the Levels Specification.

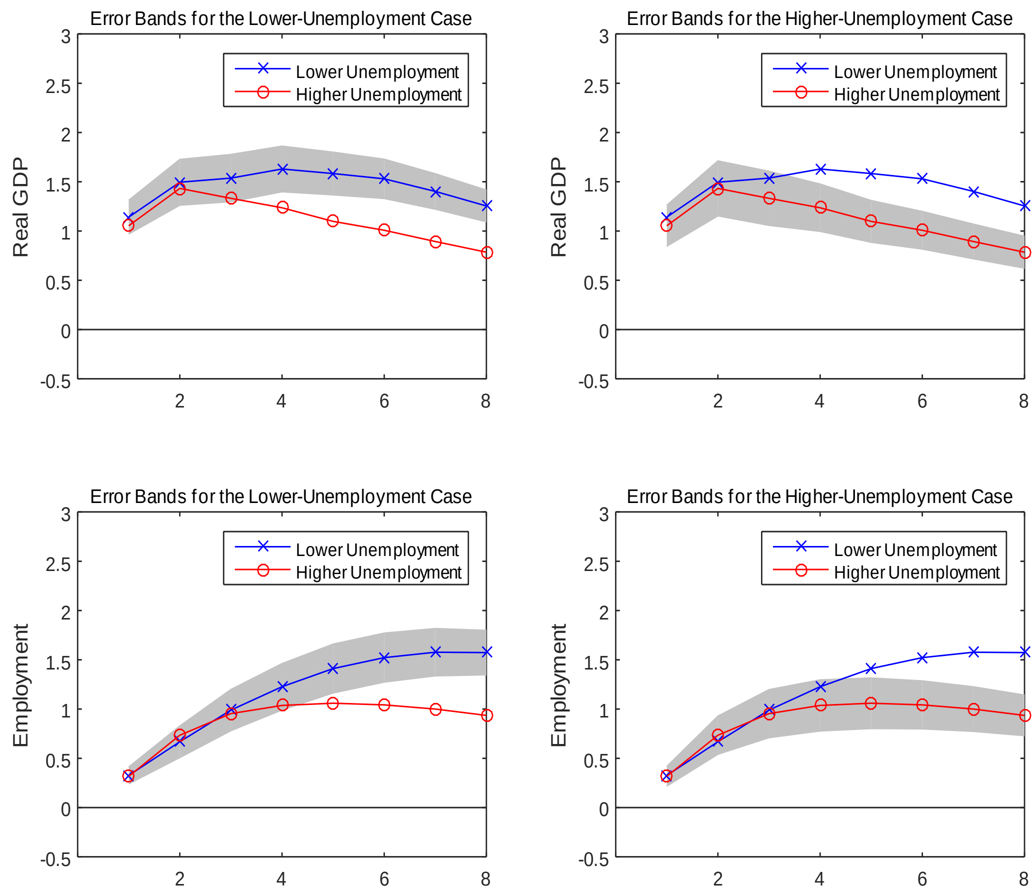

Note: In each panel, the vertical axis shows the estimated percentage change in each variable in response to a tax shock and the horizontal axis denotes the quarter after the shock. "Lower Unemployment" and "Higher Unemployment" indicate states in which the unemployment rate gap equals 0.5 and 1.5 percent, respectively. Shaded areas mark plus and minus one-standard-error bands.

Figure 10. VAR-Based Responses Computed Using the Estimated Value of the Parameter η.

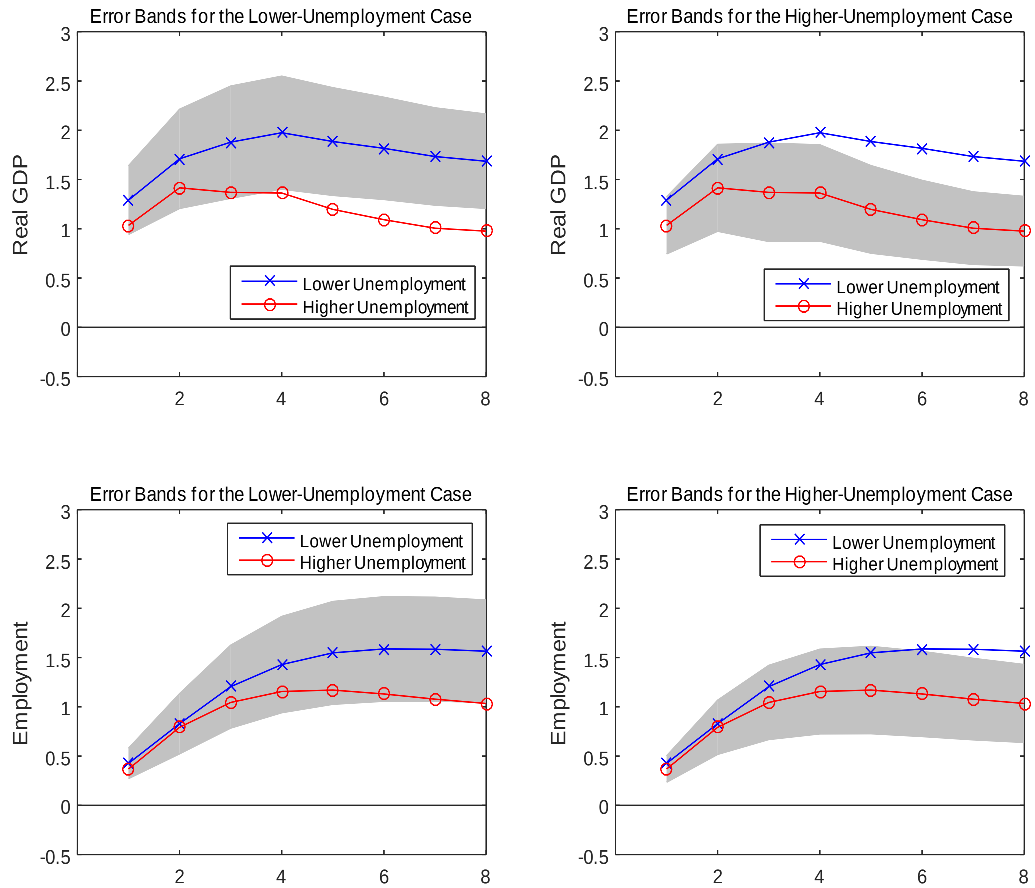

Note: In each panel, the vertical axis shows the estimated percentage change in each variable in response to a tax shock and the horizontal axis denotes the quarter after the shock. "Lower Unemployment" and "Higher Unemployment" indicate states in which the unemployment rate gap equals 0.5 and 1.5 percent, respectively. Shaded areas mark plus and minus one-standard-error bands.