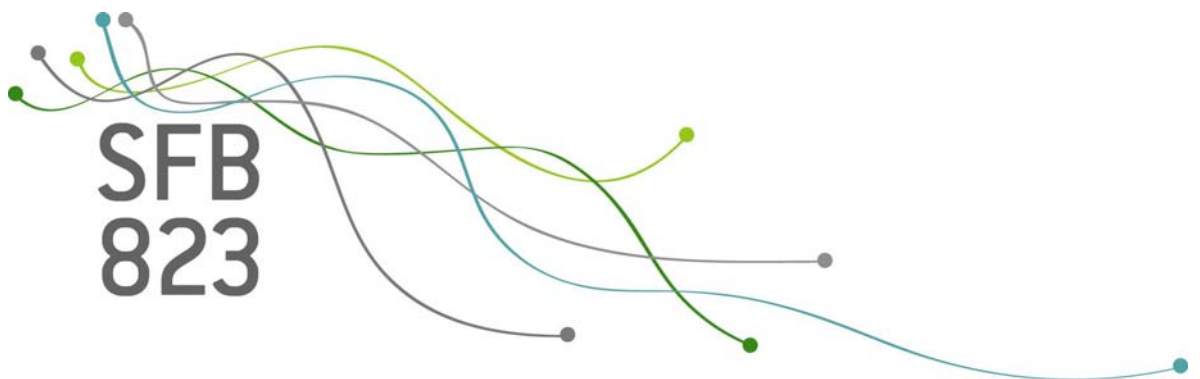SFB 823

Estimating nonlinear effects of fiscal policy using quantile regression methods

Discussion Paper

Ludger Linnemann, Roland Winkler

Nr. 3/2015

SFB 823

Estimating nonlinear effects of fiscal policy using quantile regression methods

Ludger Linnemann[*]

Roland Winkler[†]

TU Dortmund University

This version: January 21, 2015

Abstract: We use quantile regression methods to estimate the effects of government spending shocks on output and unemployment rates. This allows to uncover nonlinear effects of fiscal policy by letting the parameters of either vector autoregressive models or local projection regressions vary across the distribution of macroeconomic activity. In quarterly US data, we find that fiscal output multipliers are notably larger if GDP is predicted to be below trend. Conversely, higher government spending appears to significantly lower unemployment only if the unemployment rate is in the largest deciles of its conditional distribution.

---

[*]*Corresponding author*: TU Dortmund University, Department of Economics, Vogelpothsweg 87, D-44221 Dortmund, Germany, email: ludger.linnemann@tu-dortmund.de, tel: +49/231/7553102.
[†]roland.winkler@tu-dortmund.de

1 Introduction

The size of the effect of government spending shocks on output is the subject of a large empirical literature. Early contributions, starting from Blanchard and Perotti (2002), used linear time series models, in particular structural vector autoregressions (VAR) to estimate the magnitude of fiscal multipliers. More recently, the debate has turned to the question whether there are nonlinear effects of fiscal shocks that cannot be adequately captured by linear statistical methods. Blanchard and Leigh (2013), for example, suggest that overoptimistic forecasts of the consequences of fiscal consolidation in the aftermath of the Great Recession may be attributable to the failure of recognizing that fiscal multipliers were state-dependent and particularly large during this period. Auerbach and Gorodnichenko (2012, 2013) provide evidence based on nonlinear smooth transition vector autoregressions and local projection methods that the output effects of fiscal policy changes are systematically larger in recessions than in booms.

In this paper, we present new evidence on the nonlinear effects of government spending shocks using an alternative estimation technique that is particularly useful for this task. Specifically, we apply quantile regression methods (see Koenker and Basset, 1978) to the same basic frameworks that have been used in the previous literature. First, based on the methodology first proposed in Cecchetti and Li (2008) who study nonlinear effects of asset price booms or crashes, we use quantile methods to estimate vector autoregressions and the associated quantile-specific impulse responses. Second, we adapt the local projection method proposed by Jordà (2005) to estimate the effects of fiscal policy on the forecasts of various quantiles of the distribution of macroeconomic activity. Using quarterly post-WWII US data, we measure economic activity either by trend deviations of real GDP (henceforth referred to as output) or by the unemployment rate. In both cases, we find notable nonlinearities in the estimated effects of fiscal policy. In particular, fiscal expansions appear to have a considerably larger effect on output if the latter is depressed, in the sense of being predicted to be in the lower parts of its conditional distribution. Similarly, a strongly negative effect on unemployment only results if the latter is elevated.

The quantile methods that we use have important advantages both relatively to more traditional linear methods and relatively to the smooth transition estimates of Auerbach and Gorodnichenko (2012, 2013).

Relatively to estimates based on linear models, quantile methods have the appealing feature that they are able to estimate the impact of explanatory variables, like fiscal spending shocks, with different coefficients at the different quantiles of the conditional distribution of the outcome variable. Conventional linear regression methods can only describe the effects of changes in explanatory variables on the conditional mean of the dependent variable, thus capturing only the central tendency of effects. In the present context, this amounts to the assumption that a fiscal policy change shifts the whole conditional distribution of output or unemployment rates uniformly. Whether this is true or not is of course an empirical question. The quantile regressions that we employ allow for the possibility that policy measures may have different effects at the tails of the conditional distribution of the outcome variable than at the center or at the mean. The method is thus well suited to investigate whether fiscal policy effects are nonlinear, in the sense that there are quantile-specific parameters that lead to different responses if, for example, output is depressed and thus in a lower quantile of its conditional distribution.

Relatively to the nonlinear estimation methods used in Auerbach and Gorodnichenko (2012, 2013), quantile methods arguably have considerable advantages as well. In particular, the smooth transition autoregressions (either in the form of vector autoregression or local projections) used in these papers require that the estimates are explicitly conditioned on a classification of different regimes (that are interpreted as recessions and expansions) and on a transition function that governs in which regime the economy resides. In contrast, we do not need to characterize different parts of the data as belonging to pre-specified boom or recession episodes, because the quantile methods used here allow us to estimate the impact of shocks on the whole distribution of output or unemployment, since parameters can vary in an unconstrained way across different quantiles. The obvious advantage is that we can analyze the effects of shocks at extreme values of the conditional distribution of economic activity. Thereby, our results are not vulnerable to potentially contentious issues pertaining to the timing of business cycles, or the precise definition of recessions and booms. This is particularly useful, for example, with respect to an outcome variable like unemployment, whose variation may not only reflect business cycles, but also other influences. The results that we present always pertain to different parameter estimates at specific portions of the conditional distribution of the data, specified by the particular quantile at which the estimates are carried

out. Thus, we can make statements about the effects of fiscal policy on relatively low or relatively high predicted unemployment rates, irrespective of whether these belong to booms or recessions (however defined).

The main results are as follows. We find that the reaction of output to government spending is notably different across the quantiles of output. In particular, the fiscal multiplier is relatively large and significantly positive if GDP is predicted to be considerably below its trend. Contrarily, for the highest quantiles of the conditional output distribution, the maximum fiscal multiplier effect is still slightly positive, but the cumulative multiplier (measured as the cumulative effect over 12 periods after a fiscal shock) may even become negative. Interestingly, we find that the impact effect of a government spending expansion is essentially the same across all output quantiles. This implies that the initial effect of a fiscal shock is well captured even with standard linear regression techniques that constrain the effect to be the same at each point of the distribution and measure only changes in the conditional mean. However, the same is not true for the dynamics following fiscal shocks. At low quantiles, there is a persistent hump-shaped rise in output. In contrast, at output quantiles closer to the median or higher, the effects of spending innovations are short-lived and do not display notable humps.

Thus, the main thrust of these results is consistent with the finding of strongly nonlinear effects of fiscal policy in Auerbach and Gorodnichenko (2012, 2013). The first of these papers uses a smooth transition vector autoregression and estimates higher fiscal multipliers in recessions using US data, while the latter uses state-dependent local projections for a panel of OECD countries with the same basic result. The finding of higher fiscal multipliers in recessions is, however, disputed by Ramey and Zubairy (2014) who use state-dependent local projections with US data. While we cannot enter this debate directly since our results are not conditioned on an explicit measure of recessions or booms, the general finding of pronounced nonlinearities in fiscal policy effects nevertheless tends to corroborate the result in Auerbach and Gorodnichenko (2012).

The second outcome variable that we consider is the unemployment rate. Here, we find the effects of government spending shocks to be negligible for large parts of the conditional unemployment distribution. Only at the highest decile of unemployment rates we find consistently large and statistically significant negative effects. Again, the disparity found at different quantiles is not

4

due to differences in impact effects, which are small, but is almost entirely due to the much larger persistence and hump-shaped nature of the dynamics following a shock that hits the economy when the labor market is characterized by high unemployment.

All results seem robust to allowing for possible anticipation effects due to fiscal foresight on the part of the private sector that have been found important by Ramey (2011). Moreover, while the quantitative details certainly differ between vector autoregressive estimates and those based on local projections, the main results with respect to the nonlinear nature of the effects of fiscal shocks are qualitatively similar between both methods.

The rest of the paper is organized as follows. Section 2 presents the quantile regression methods that we employ throughout. Section 3.1 shows the results concerning GDP, and 3.2 concerning unemployment rates. Sections 3.3 and 3.4 check for robustness. Section 4 concludes. Data and details of the algorithms used are discussed in an appendix.

2   Method

Quantile regression, as originally developed by Koenker and Bassett (1978), can be seen as a flexible generalization of standard regression equations (see Koenker, 2005, for a more recent treatment). A standard least squares regression postulates a model for the mean of a variable y_t conditional on the values of a vector of explanatory variables x_t, such that $E(y_t|x_t) = x_t\beta$, where the parameter vector β is estimated so as to minimize the sum of squared residuals. The estimated model explains changes in the conditional mean of y_t resulting from changes in the x_t variables. In the context of quantile regressions, in contrast, the focus is not on the conditional mean, but on the whole conditional distribution of y_t. This distribution can be characterized by its quantiles $q \in (0, 1)$, and the quantile model seeks to explain changes in the q-th quantile of the y_t - distribution by changes in the explanatory variables. Formally, letting $F(y_t)$ be the probability distribution function of the random variable y_t, the q-th quantile is defined by the quantile function $Q_q(.)$ with $Q_q(y_t) = F^{-1}(q)$. The quantile regression model explains the q-the quantile of y_t given the values of some vector of explanatory variables x_t as

$$Q_q(y_t|x_t) = x_t\beta(q),$$

where the elements of the parameter vector β give the marginal effect of the corresponding explanatory variable. The notation $\beta(q)$ highlights that there is a potentially different parameter vector at each respective quantile q of the distribution. The estimate for the parameter vector at given q is, following Koenker and Bassett (1978), obtained as

$$\widehat{\beta}(q) = \arg\min_{\beta(q)} \sum_t \rho_q \left[y_t - x_t(\beta(q)) \right], \tag{1}$$

where the check function ρ_q is defined, for any z_t, as $\rho_q[z_t] = (q - I_{z_t<0})z_t$ with $I_{z_t<0} = 1$ if $z_t < 0$ and $I_{z_t<0} = 0$ otherwise.

The resulting estimate $\widehat{\beta}(q)$ answers the question: if x_t changes by one unit, how much does the q-th quantile of the conditional distribution of y_t change? In this way, by evaluating (1) for various values of q, one can characterize the impact of changes in the x_t variables on the whole conditional distribution of y_t, measured at any of its quantiles that is of interest to the researcher. This allows for a much richer set of results than classical regression, which only characterizes the effect of the x_t on the conditional mean of y_t, and which assumes (tacitly) that changes in x_t shift the whole distribution of y_t in the sense of constraining the responses to be the same at each quantile.

Note that for $q = 0.5$, i.e. at the median, one obtains a regression based on minimizing the mean of absolute deviations, which for many distributions should be close to a standard least squares regression. The distinct advantage of quantile methods is to be able to focus attention on quantiles that are farther away from the median or mean, and thus gather information about the nature of the relation between the variables of interest that is not captured by the relation between their means. In a business cycle context, where the variables denote the cyclical components of macroeconomic aggregates, this allows us to discern whether the parameters of behavioral equations are different when evaluated in the lower or upper quantiles of the conditional distributions of the variables, which corresponds to evaluating the model equations in distributional areas associated with booms or recessions.

The exploration of parameter differences across the distribution of variables characterizing the state of the business cycle, like the cyclical component of GDP, or the unemployment rate, is the central object of interest in this paper. In our application, we use quantile techniques in two

different specifications that we discuss now. Both of these specifications have been used by previous authors, who however relied on standard linear regression methods, namely vector autoregressions and direct local projection regressions, or on nonlinear methods that require, unlike in our case, the a priori definition of different regimes.

In Section 3.1 to 3.3, we apply quantile regression estimation to a p-th order vector autoregression (VAR). Here, we build on the insights of Cecchetti and Li (2008), who to our knowledge have been the first to use quantile estimation methods in the context of VAR models. Let $z_t = (z_{1t}, z_{2t}, ..., z_{kt})'$ be a vector of k variables measured at time t, and let $q = (q_1, ..., q_k)'$ be a vector of quantiles at which the conditional distributions of the variables that constitute the elements of z_t are evaluated. We postulate a linear vector autoregressive model for these quantiles, for a given constant $p > 0$, such that

$$Q_q(z_t|z_{t-1}, \ldots, z_{t-p}) = c(q) + \sum_{i=1}^{p} B_i(q) z_{t-i} + \varepsilon_t(q), \qquad (2)$$

where

$$B_i(q) = \begin{pmatrix} \beta_{i,11}(q_1) \ldots \beta_{i,1k}(q_1) \\ \beta_{i,21}(q_2) \ldots \beta_{i,2k}(q_2) \\ \vdots \qquad\qquad \vdots \\ \beta_{i,k1}(q_k) \ldots \beta_{i,kk}(q_k) \end{pmatrix}, \quad c(q) = \begin{pmatrix} c_1(q_1) \\ c_2(q_2) \\ \vdots \\ c_k(q_k) \end{pmatrix}, \quad \varepsilon_t(q) = \begin{pmatrix} \varepsilon_{1t}(q_1) \\ \varepsilon_{2t}(q_2) \\ \vdots \\ \varepsilon_{kt}(q_k) \end{pmatrix}.$$

Here, the coefficients $\beta_{i,jn}(q_j)$ give the effect of lag i of variable n, i.e. z_{nt-i}, on the q_j-quantile of the conditional distribution of variable z_{jt}. The constants c and the i.i.d. elements of the disturbance vector ε_t are also indexed by the respective quantile. The notation emphasizes the important fact that the parameters in each equation may pertain to different quantiles of the conditional distribution of the respective left hand side variable. This allows us to use the estimated model to answer a variety of interesting empirical questions. For example, we can ask how a change in the median of z_{1t} ($q_1 = 0.5$) affects the lowest decile of the conditional distribution of the second variable z_{2t} ($q_2 = 0.1$).

Since each equation of (2) has the same right hand side, we estimate the model equation by equation for each vector q using quantile regression. Here and in all applications below, we use Roger Koenker's Matlab function rq_fnm.m for estimation.[1] We then orthogonalize the covariance

---

[1] This software is available at http://www.econ.uiuc.edu/~roger/research/rq/rq.html

matrix of the residuals by a Cholesky decomposition, and compute impulse responses. The implied identification restrictions are discussed below in Section 3. Standard errors of impulse responses are bootstrapped by resampling from the estimated residuals in a way described in further detail in the appendix.

In Section 3.4, we check the robustness of the results with respect to using a different method to compute impulse responses, namely the local projection method proposed by Jordà (2005). We implement it as a single equation estimate of the quantiles of an outcome variable x_{t+h} on a policy shock measure s_t and a vector of k control variables z_t. That is, for a given forecast horizon h, we estimate the direct forecast quantile regression

$$Q_q(x_{t+h}|s_t,\ldots,s_{t-p},z_t,\ldots,z_{t-p}) = c_h(q) + \alpha_h(q)s_t + \sum_{i=1}^{p}\beta_{h,i}(q)s_{t-i} + \sum_{n=1}^{k}\sum_{i=1}^{p}\gamma_{h,ni}(q)z_{n,t-i} + \varepsilon_{t+h}(q),$$

where again q is a vector of quantiles of interest, c is a constant, and α, β, and γ are quantile-specific and horizon-specific parameters to be estimated. This equation is estimated via quantile regression for various forecast horizons $h = 0, ..., H$, which allows to construct impulse responses as local projections as proposed by Jordà (2005). To compute standard errors, we use the blocks-of-blocks bootstrap as further described in the appendix to preserve the dependencies among the variables while resampling.

3 Results

This section presents our estimation results. First, we estimate vector autoregressions at the quantiles of the distribution of detrended GDP and of the distribution of the unemployment rate. We start with a presentation of results of our baseline specification and discuss robustness checks subsequently. Second, we estimate local projections and discuss the results from this alternative estimation strategy.

3.1 Quantile VAR-GDP

In this section, we employ quantile regression to estimate vector autoregressive estimates of government spending shock effects on aggregate output. We use quarterly US data from 1954q3 to 2013q4, which is the longest sample over which all variables are available. Our baseline set of variables consists of government spending (real government consumption and gross investment)

g_t, real GDP y_t, and real net taxes τ_t (measured as the real value of government current tax and social security receipts, net of current transfers, deflated with the GDP deflator), as well as the short-run real interest rate r_t, constructed as the annualized difference between the Federal Funds Rate and the log-change in the GDP deflator. The data appendix contains details on the precise definitions and sources of the data. The level variables g_t, y_t, and τ_t are measured as log-deviations from quadratic time trends. In keeping with previous literature, we use a constant and four lags of all variables in the VARs.

The parameters of the VARs are estimated by applying equation-by-equation quantile regressions, as described above. On the left hand sides of the estimating equations, we measure government spending, taxes, and the real interest rate at the medians of their conditional distributions, but evaluate the effects on detrended real GDP (output) at various quantiles of its conditional distribution. In this way, we estimate if the effects of fiscal spending shocks differ with respect to whether they impact output in a situation where it is relatively low or where it is relatively high. While, in principle, the estimation can be performed at any quantile of output, for conciseness we present the baseline results for just three quantiles, namely for the $q = \{0.1,\ 0.5,\ 0.9\}$ quantiles of output. These provide information about how policy shocks impact output in three benchmark cases. In the case $q = 0.1$, the cyclical component of real GDP is in the lowest 10% of its conditional distribution, which we interpret as a recessionary period. If $q = 0.5$, we evaluate the policy effects at the median of output, which is close to its mean, and if $q = 0.9$ we estimate the effects of a fiscal shock occurring when cyclical output is booming such that it is in the highest 10% of its conditional distribution. This choice of quantiles is made to ensure the readability of the presentations and to focus the discussion on the most interesting aspects. However, it is important to note that these quantile values are just examples for various parts of the output distribution, and the general results are robust with respect to the choice of other possible quantiles. We demonstrate this by providing a more detailed analysis with a finer grid of output quantiles further below.

For any given quantile q, we estimate the reduced form VAR equations and compute impulse responses by orthogonalizing the residual covariance matrix through a Cholesky decomposition. Thus we impose a standard recursive identification scheme. We order the government spending variable first, following Blanchard and Perotti (2002) and a large literature thereafter, implying

Figure 1: Impulse responses at different quantiles of the output distribution.

(a) Output

(b) Government spending
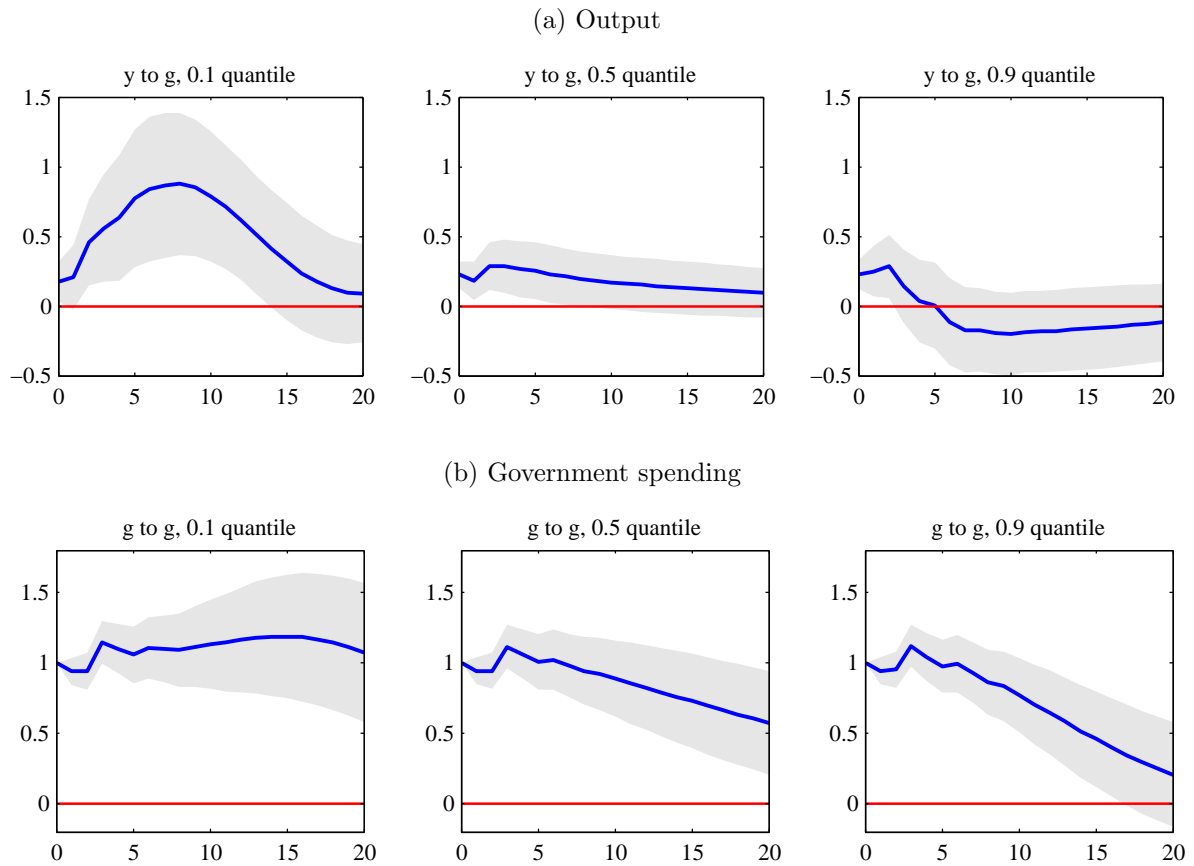
Notes: The solid lines show the responses of output and government spending to a one percent government spending shock, when the parameters are estimated at the 0.1, 0.5, and 0.9 output quantiles. The shaded areas are the 90% confidence intervals.

that a fiscal spending shock is identified by the assumption that government spending is exogenous within the quarter. This identification method has come under critique from Ramey (2011), who argues that it is vulnerable to confounding unanticipated spending shocks with responses to anticipated changes in spending. Below, we therefore discuss the extent to which our results are robust to the possible presence of anticipation effects.

Figure 1 shows the orthogonalized impulse responses of output and government spending, based on the VAR equations estimated at the various quantiles of output, for our baseline model (g_t, y_t, τ_t, r_t) (in this order) to a positive one percent shock in (the median of) government spending.[2] The shaded areas depict bootstrapped 90% confidence intervals.

The figure shows that output generally responds positively to a fiscal expansion, and the

---

[2] For brevity, we display only the responses that are most interesting for the question at hand. The full set of impulse response figures for all variables can be obtained from the authors upon request.

response is statistically significantly different from zero for at least some quarters following a shock. The most interesting aspects, however, can be seen by comparing the responses at different quantiles, i.e. by comparing responses across the columns of each row of the figure. The impact response in the shock period is approximately identical at all quantiles of output. We verified that this is not specific to the quantiles on display here, but holds more generally at all quantiles between $q = 0.025$ and $q = 0.975$. However, there is a strong difference with respect to the maximum output response. When output is at its lowest decile ($q = 0.1$), a fiscal shock engenders a strong hump-shaped recovery that peaks about 6 to 9 quarters after the shock. At the median or at the highest decile of output ($q = 0.5$ and $q = 0.9$), however, the response rises only weakly and very briefly in the quarters after the impact. The maximum response is reached two quarters after the shock, and is only about one third as strong as at the 0.1 quantile. At $q = 0.9$, there is even an indication of a negative response after some quarters, although this seems insignificantly different from zero. Hence, we find a strong nonlinearity of the dynamic reaction of the aggregate economy: if output is relatively low to begin with, a fiscal expansion triggers a protracted and relatively strong increase in economic activity. If output is average or high, a fiscal shock of the same size produces initially the same reaction, but is hardly followed by any further output increases.

The persistently strong output response at the low quantile could conceivably be due to higher persistence of the shocks occurring at this part of the output distribution. That this is at least partly true is conveyed by the second row of Figure 1, which shows the own response of government spending to the fiscal shock, again with all parameters estimated at different quantiles of the conditional output distribution. The graphs clearly show that fiscal expansions that occur in times of relatively low output are much more persistent. To capture the extent to which the differences in the size and persistence of the output responses are driven by this, we also show fiscal multipliers.

Specifically, we compute the point-to-point fiscal multiplier and the cumulative fiscal multiplier as follows: let $\widehat{y}_h(q)$ be the impulse response of output h periods after the shock impact, when the parameters are estimated at the q-quantile of the output distribution, and let $\widehat{g}_h(q)$ denote the same for the government spending response. Then the point-to-point multiplier $M_h(q)$ is defined as $M_h(q) = \widehat{y}_h(q)/\widehat{g}_h(q)$, that is the ratio of the impulse response of y at h periods after a shock

11

Figure 2: Output multipliers at different quantiles of the output distribution.

(a) Maximum point-to-point multiplier (b) Cumulative multiplier
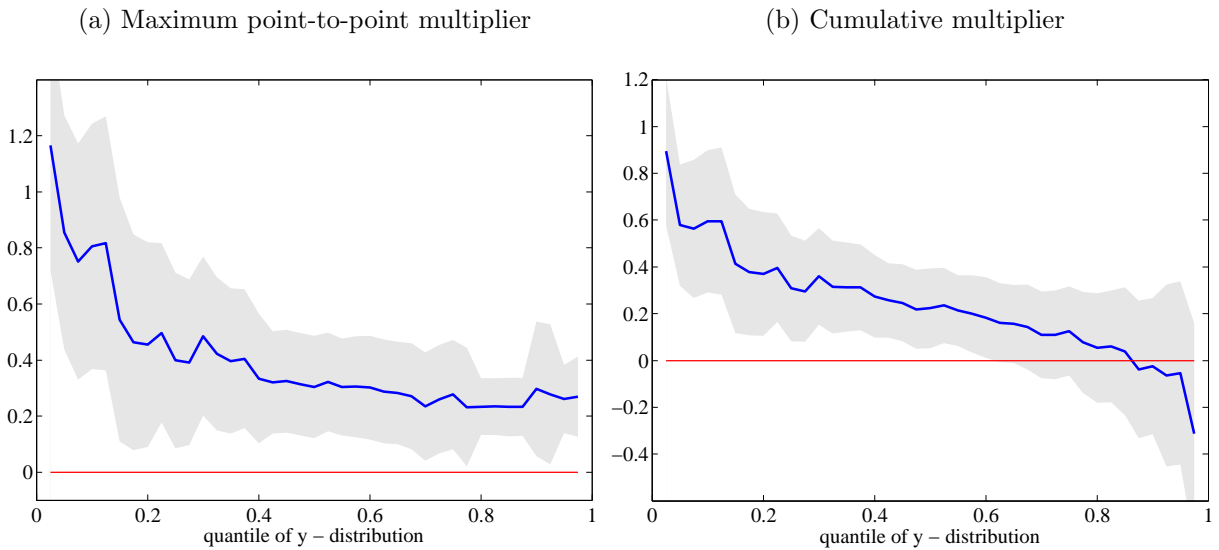
Notes: Panel (a) shows the maximum of the point-to-point multipliers, computed as $\widehat{y}_h(q)/\widehat{g}_h(q)$ for $h = 0, \ldots 12$, where $\widehat{y}_h(q)$ and $\widehat{g}_h(q)$ are the impulse responses of output and government spending h periods after the shock impact, when the parameters are estimated at the q-quantile of the conditional output distribution. Panel (b) shows the 12-quarter cumulative output multiplier, computed as $\sum_{h=0}^{12} \widehat{y}_h(q)/\sum_{h=0}^{12} \widehat{g}_h(q)$. The shaded areas are the 90% confidence intervals.

over the impulse response of g at the same time. We estimate this at all quantiles of y between $q = 0.025$ and $q = 0.975$ in steps of 2.5%, and record those values which are the largest over all horizons between zero and 12 periods ahead, $\max_{h \in \{0, \ldots, 12\}} M_h(q)$. The left panel of Figure 2 shows $\max_{h \in \{0, \ldots, 12\}} M_h(q)$ plotted against the output quantile at which it is estimated.

The cumulative multiplier, on the other hand, is computed as the ratio of the sum of output responses over the sum of government spending responses for the first 12 quarters following a shock, $CM_h(q) = \left[\sum_{h=0, \ldots, 12} \widehat{y}_h(q) \right] \Big/ \left[\sum_{h=0, \ldots, 12} \widehat{g}_h(q) \right]$. The right hand side of Figure 2 shows $CM_h(q)$ against the quantiles at which the coefficients are estimated.

One remark is in order. We compute multipliers in elasticity form. Multiplying our multipliers with the sample average of the GDP to government spending ratio, 4.0127 in our sample, yields multipliers in terms of absolute changes that are the objects often discussed in the literature.

The figure shows that the maximum pointwise response of output, normalized by the own response of government spending, varies strongly with respect to the quantile at which it is estimated. In particular, fiscal policy is effective at the lower output deciles, with a maximum multiplier peaking slightly above 1 if output is in the lowest 2.5% to 5% of its conditional distribution. By contrast, the multiplier levels out at medium to high quantiles, indicating that the

12

maximum of the output response relative to the government spending response is in the range between 0.4 and 0.2 for the middle and right parts of the output distribution.

These findings can be interpreted to imply that fiscal policy shocks are estimated to be rather strongly and persistently expansionary if they occur in phases where output is extraordinarily low. If output is close to its median, or if it is unusually high, the effects of fiscal shocks are still positive, but much weaker. The maximum pointwise fiscal multiplier may be about three times higher in a recession than it would be in an output boom. This nonlinearity in the responses appears quantitatively important and is measured rather precisely.

The finding of a larger fiscal multiplier in low output states is in line with the results in Auerbach and Gorodnichenko (2012, 2013). The results based on the method presented here have the advantage that they are not conditioned on a notoriously contentious delineation of regimes or business cycle phase dating issues. Instead, when analyzing the whole output distribution we find that fiscal spending not only changes its position, but also its shape, by strongly compressing the left tail which pertains to recessionary outcomes.

While the maximum multiplier only looks at one point of each set of responses, it might be interesting to have a more comprehensive summary measure. One such measure that is frequently discussed in the empirical literature is the cumulative fiscal multiplier. Looking at the right panel in Figure 2, we find a similar picture in that the cumulative multiplier is relatively large in low output states. It is significantly larger than zero for the lower roughly 60 per cent of output values. The difference with respect to the maximum pointwise multipliers is that the cumulative multipliers shown in the right part of the figure may tend to negative values when evaluated at very high output, even though this is not statistically significant. The reason is that, for the highest output quantiles, the output response tends to become weakly negative some quarters after a shock, as already indicated above in the discussion of Figure 1.

3.2 Quantile VAR-Unemployment

In this section, we provide evidence that the estimated nonlinearities are not specific to the output variable, but rather seem to characterize the business cycle effects of government spending shocks more generally. To do so, we check the robustness of our results concerning the estimated fiscal

Figure 3: Impulse responses at different quantiles of the unemployment distribution.
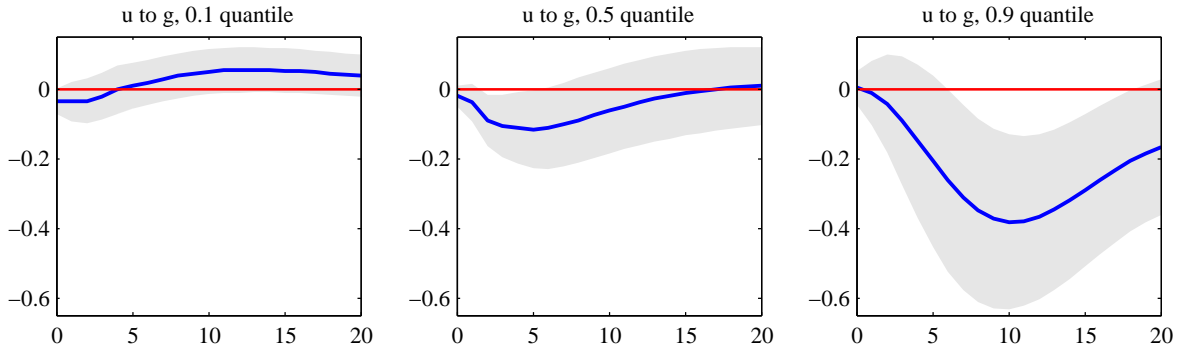
Notes: The solid lines show the response of unemployment to a one percent government spending shock, when the parameters are estimated at the 0.1, 0.5, and 0.9 unemployment quantiles. The shaded areas are the 90% confidence intervals.

policy effects on unemployment. In the estimating model, we replace output by the unemployment rate u_t, such that the vector of variables considered now reads (g_t, u_t, τ_t, r_t). We evaluate the responses to an orthogonalized shock in the median of g_t at different unemployment quantiles.

Figure 3 shows impulse responses to a one percent government spending shock.[3] The unemployment rate shows markedly different responses across the quantiles of its distribution: at $q = 0.1$, when the labor market is predicted to be booming, there is either no or even a very slight delayed positive response of unemployment. At the median, depicting an average labor market performance, the unemployment rate responds by a slight decline that reaches a trough at around -0.1 percentage points a few quarters after a shock. In contrast, in a depressed labor market, exemplified by $q = 0.9$ when unemployment is in its highest decile, a fiscal shock engenders a strong and protracted decline reaching almost -0.4 points after about two years. The finding with respect to unemployment are thus the mirror image of those with respect to GDP.

Figure 4 shows the results for a broader range of unemployment quantiles. Again, for easy readability we condense the information in the impulse responses by computing the extreme values of the pointwise multipliers and the cumulative multipliers. In the case of unemployment being the outcome variable of interest, we record those values of the point-to-point multipliers that occur at the horizon where the point-to-point multiplier reaches a minimum. The cumulative 12-quarter multiplier of fiscal spending with respect to the unemployment rate, on the other hand, shown in the right panel of the figure, is again obtained by summing over the first 12 entries of the impulse response of unemployment, and dividing by the sum of the impulse response values of government

---

[3]The responses of government spending, not shown for brevity, are very similar to the ones shown in Figure 1.

Figure 4: Unemployment multipliers at different quantiles of the unemployment distribution.

(a) Minimum point-to-point multiplier   (b) Cumulative multiplier
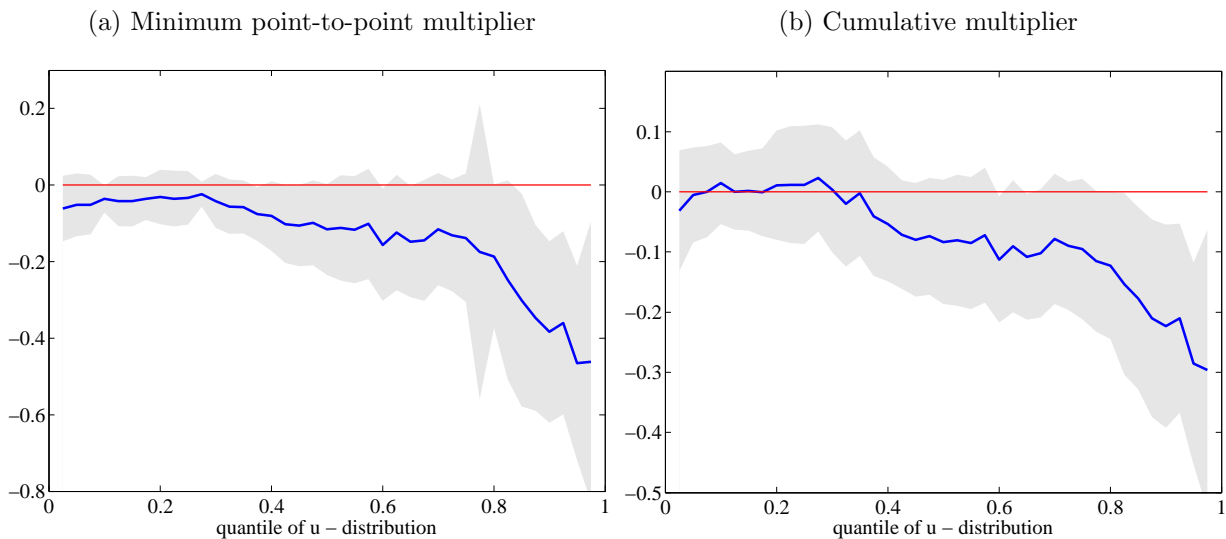
Notes: Panel (a) shows the minimum of the point-to-point multipliers, computed as $\widehat{u}_h(q)/\widehat{g}_h(q)$ for $h = 0, \ldots 12$, where $\widehat{u}_h(q)$ and $\widehat{g}_h(q)$ are the impulse responses of unemployment and government spending h periods after the shock impact, when the parameters are estimated at the q-quantile of the conditional unemployment distribution. Panel (b) shows the cumulative unemployment multiplier, computed as $\sum_{h=0}^{12} \widehat{u}_h(q) / \sum_{h=0}^{12} \widehat{g}_h(q)$. The shaded areas are the 90% confidence intervals.

spending.

Both measures allow very similar interpretations. For the largest part of the conditional unemployment distribution, the fiscal multiplier, however measured, appears small and insignificantly different from zero. In a booming labor market, as characterized by the left third of the distribution of unemployment rates, say, the effects of fiscal spending expansions on unemployment are practically zero. At unemployment rates around their median, the effect is slightly negative, but still insignificant. Only in a depressed labor market, at unemployment rates as high as the 10 to 15% highest values on record, does an increase in fiscal spending have a precisely measurable effect on unemployment. However, the effect is strongly negative in this case, suggesting that fiscal policy would be rather effective in lowering unemployment when applied in a situation where labor market performance is poor.

This, again, points out the importance of allowing for nonlinear effects of fiscal policy: an estimation approach that, in contrast to the one proposed here, relied on estimating average effects, would yield results close to those that we find at median unemployment rates and conclude that government shocks do not have much influence on the unemployment rate. As our quantile-based estimates show, this conclusion would be true for situations where the labor market is in good or

average shape, but erroneous for phases of the business cycle characterized by high unemployment.

3.3 Robustness

In this section, we present results of two robustness checks. First, we take into account anticipation effects of fiscal policy. Second, we re-estimate our VARs on a sample that leaves out the Great Recession and the recent period of the zero lower bound on nominal interest rates.

Fiscal foresight. The preceding results have been obtained under the assumption that fiscal shocks can be identified by imposing a recursive causal ordering on the VAR with government spending ordered first. This amounts, as has first been pointed out by Blanchard and Perotti (2002), to the assumption that government spending is exogenous within the quarter, because all feedback from innovations in output and other variables to government spending is presumed to take at least one quarter before becoming effective, for example due to institutional delays in the political and administrative process. This identification assumption, though applied extensively in the empirical literature, has been criticized by Ramey (2011). She points out that this identifying assumption is invalid if government spending shocks are anticipated by private agents. In this case, the timing of fiscal shocks as estimated by the method by Blanchard and Perotti (2002) is incorrect, since private responses do not react to actual spending increases, but already to news about impending future spending plans. Ramey (2011) finds that anticipation effects indeed seem to be important.

For this reason, we check the robustness of the results reported above with respect to the possibility of fiscal foresight. We use Ramey's (2011) news variable, which is constructed as the present discounted value of anticipated changes to military government spending, normalized by previous quarter's GDP.[4] We order the news variable, $news_t$, as the first one in the vector autoregressions, such that the variables are $(news_t, g_t, x_t, \tau_t, r_t)$, in this order, where x_t is either detrended GDP y_t or the unemployment rate u_t. By computing orthogonalized impulse responses from a Cholesky decomposition of the residual covariance matrix, this plausibly imposes that news are contemporaneously exogenous with respect to all other variables. Crucially, this setup allows that all other variables may be contemporaneously caused by fiscal news. Consequently, while innovations to the news variable capture anticipated fiscal shocks, we take the remaining

---

[4] The news variable can be downloaded from Valery Ramey's web site http://econweb.ucsd.edu/~vramey.

Table 1: Fiscal multipliers at different quantiles of the distribution of output or unemployment, robustness checks.

	Maximum point-to-point multipliers at the q-th quantile					Cumulative multiplier at the $q-$th quantile				
	0.025	0.1	0.5	0.9	0.975	0.025	0.1	0.5	0.9	0.975
Controlling for Ramey news										
Output	1.16***	0.63***	0.29***	0.31***	0.32***	0.74***	0.47***	0.17**	0.07	−0.22
	(0.18)	(0.24)	(0.11)	(0.15)	(0.09)	(0.20)	(0.17)	(0.10)	(0.18)	(0.28)
Unemployment	−0.04**	−0.04*	−0.11*	−0.23*	−0.60***	0.06*	0.01	−0.07*	−0.14*	−0.37***
	(0.02)	(0.03)	(0.07)	(0.17)	(0.10)	(0.05)	(0.04)	(0.06)	(0.10)	(0.07)
Leaving out the Great Recession										
Output	1.21***	0.81***	0.27***	0.27**	0.22***	0.90***	0.63***	0.17*	−0.10	−0.52*
	(0.28)	(0.19)	(0.12)	(0.15)	(0.09)	(0.20)	(0.13)	(0.12)	(0.26)	(0.45)
Unemployment	−0.03	−0.02	−0.13**	−0.30	−0.43**	−0.01	0.03	−0.10**	−0.23***	−0.33***
	(0.05)	(0.02)	(0.06)	(0.35)	(0.22)	(0.04)	(0.04)	(0.06)	(0.10)	(0.11)

Notes: For output, columns 2-6 report the maximum of the point-to-point multipliers, computed as $\widehat{y}_h(q)/\widehat{g}_h(q)$ for $h = 0, \ldots, 12$, where $\widehat{y}_h(q)$ and $\widehat{g}_h(q)$ are the impulse responses of output and government spending h periods after the shock impact, when the parameters are estimated at the q-quantile of the conditional output distribution. For unemployment, the minimal point-to-point multiplier is shown. Columns 7-11 report cumulative multipliers, computed as $\sum_{h=0}^{12} \widehat{y}_{qh} / \sum_{h=0}^{12} \widehat{g}_{qh}$ or $\sum_{h=0}^{12} \widehat{u}_h(q) / \sum_{h=0}^{12} \widehat{g}_h(q)$. Standard errors are reported in parentheses, '*', '**', '***' indicate statistical significance at 16, 5, and 2.5 percent levels.

orthogonal innovations in g_t after controlling in this way for anticipation effects as measuring unanticipated spending shocks.

We thus compute the impulse responses to unanticipated spending shocks, controlling for fiscal news, on each of the activity variables (either output or the unemployment rate), and again compare the results at various quantiles of the activity variable. The upper block of Table 1 gives an overview of the results.[5] Overall, the responses of output or unemployment rates are broadly similar to the ones presented in the previous section. In the case of output being the activity variable, the fiscal multiplier is large and significant at low output quantiles, while it tends to get smaller when evaluated at the higher output quantiles. The maximum pointwise multiplier remains positive throughout, while the cumulative multiplier, again measured over 12 quarters, may become negative in the highest ranges of output realizations, though this seems statistically indistinguishable from zero.

In the case of the unemployment rate as the indicator of macroeconomic activity, the picture is also broadly similar to the one shown above. The most compelling finding is that unemployment responses, and the associated fiscal multipliers, are numerically small (be it weakly positive or

---

[5]A graphical exposition of the impulse response figures can be obtained from the authors upon request.

negative, borderline significant or not at all) for the largest part of the conditional distribution of unemployment rates. However, fiscal shocks do seem to affect unemployment rates in their upper quantiles rather strongly negatively.

Thus, we conclude that the results reported above for the baseline model do not seem to be the artefact of an incorrect identification assumption, in the sense that controlling for the possibility of anticipated fiscal spending, or news shocks as constructed by Ramey (2011), leaves the salient features of the empirical findings largely unchanged.

Leaving out the Great Recession. The results presented so far are based on a sample that includes the recent period of the Great Recession in which output was considerably below its trend and monetary policy was constrained by the zero lower bound on nominal interest rates. Recent theoretical contributions by, for example, Eggertsson (2010) and Christiano, Eichenbaum and Rebelo (2011) suggest that government spending multipliers might be large at the zero lower bound. The question thus arises whether our findings are mainly driven by the observations pertaining to the period of the Great Recession and the associated zero lower bound on nominal interest rates. To check whether this is the case or whether our findings are robust when we leave out the Great Recession, we re-estimate our quantile VARs on a sample that ends in 2007q4.

The lower block of Table 1 shows that our results remain qualitatively intact when we leave out the Great Recession. At low quantiles of the output distribution, the multipliers are large and highly statistically significant. Likewise, at high quantiles of the unemployment distribution, unemployment multipliers are large and significant. This has an interesting theoretical implication because it suggests that the nonlinear effects of government spending shocks cannot, or at least not entirely, be explained by changes in the transmission mechanism of fiscal policy that occur only at the zero lower bound constraint on monetary policy.

3.4   Local projections

The impulse responses presented so far, and the multiplier values derived from them, relied on estimating the whole VAR system of dynamic equations. The results thus depend on inverting the autoregressive lag polynomial, and are affected by all parameters in all equations. An alternative that does not have this property has been proposed by Jordà (2005), who points out that impulse

responses can be directly estimated from local projections, i.e. by estimating the change in the forecast of an outcome variable that is brought about by a change in an impulse variable. Local projections have recently also been used in the empirical fiscal policy literature by Auerbach and Gorodnichenko (2013) and Ramey and Zubairy (2014). It is thus interesting to uncover whether our results are robust with respect to using this alternative technique of estimating impulse responses.

In this section, we thus present robustness checks in the form of estimates of quantile local projections. Specifically, let the activity variable $x_t \in \{y_t, u_t\}$, again be either detrended GDP y_t or the unemployment rate u_t. We then estimate the h-period ahead forecast equation

$$Q_q(x_{t+h}|g_t, Z) = \alpha_h(q)g_t + Z'\gamma(q) + \varepsilon_{t+h}(q), \quad h = 0, ..., 12, \tag{3}$$

where $Q_q(x_{t+h}|g_t, Z)$ is the q-quantile of x_{t+h}, α_h is a parameter and γ a parameter vector, g_t is detrended government spending as above, and Z' is a vector that contains other variables that are used to forecast future macro activity. The variables in Z' are, apart from a constant, four lags of government spending and of the respective activity variable (output or unemployment), and four lags of our measures of real net taxes and the short-run real interest rate as defined above, and ε is a disturbance term. For each quantile q at which the distribution of outcomes is to be evaluated, we estimate (3) for each h and record the estimate of $\alpha_h(q)$. The sequence of the estimates of $\alpha_h(q)|_{h=1,..,12}$ is then the local projection impulse response. It gives the change in the respective predicted q-quantile of the distribution of x_{t+h} in h periods if there is contemporaneous change in fiscal policy that leads to an unexpected movement in g_t.

Note that for $h = 0$, (3) is just identical to the first equation in the VAR used above. Therefore, $\alpha_0(q)$ is the immediate effect of a one unit change in current g_t, and thus can be interpreted as the impact multiplier. As such, it is subject to the same identification assumption as used above, namely that disturbances other than fiscal shocks (hence any innovation in ε) that affect output or unemployment do not feed back into changes in government spending during the same quarter. The plausibility of this assumption is, again, doubtful if large parts of fiscal spending changes are anticipated by private agents in advance. Therefore, we also check the robustness of the results by including Ramey's (2011) fiscal news variable, and four lags of it, in the vector Z'.

Table 2 presents the results. As is usual with local projection methods, the individual impulse

Table 2: Impulse responses at different quantiles of the output or unemployment distribution according to Jordà's local projection method.

	Maximum response at the q−th quantile					Cumulative response at the $q-th$ quantile				
	0.025	0.1	0.5	0.9	0.975	0.025	0.1	0.5	0.9	0.975
Baseline specification										
Output	0.65**	0.57**	0.40*	0.27*	0.36**	4.02***	3.59***	3.48***	1.68***	1.32*
	(0.36)	(0.32)	(0.27)	(0.16)	(0.22)	(1.23)	(1.04)	(0.93)	(0.79)	(0.83)
Unemployment	−0.14*	−0.10***	−0.16*	−0.49***	−0.47***	−0.50**	−0.48**	−0.95***	−3.22***	−2.67***
	(0.11)	(0.05)	(0.11)	(0.21)	(0.21)	(0.29)	(0.24)	(0.35)	(0.61)	(0.74)
Controlling for Ramey news										
Output	0.81***	0.74***	0.38*	0.29*	0.35**	3.87***	3.85***	3.15***	0.73	−0.39
	(0.37)	(0.34)	(0.24)	(0.19)	(0.20)	(1.23)	(1.07)	(0.94)	(0.83)	(0.94)
Unemployment	−0.15*	−0.11*	−0.08	−0.45***	−0.46***	−0.75***	−0.73***	−0.35	−2.97***	−3.13***
	(0.11)	(0.07)	(0.12)	(0.22)	(0.22)	(0.30)	(0.29)	(0.38)	(0.68)	(0.78)

Notes: The table report responses to a one percent government spending shock. Columns 2-6 report the maximum response at the q-th quantile of the conditional output or unemployment distribution, computed as $\max_{h=0,\dots 12} \widehat{\alpha}_h(q)$. For unemployment, columns 2-6 show the minimal response. Columns 7-11 report the cumulative responses, computed as $\sum_{h=0}^{12} \widehat{\alpha}_h(q)$. Standard errors are reported in parentheses, '*', '**', '***' indicate statistical significance at 16, 5, and 2.5 percent levels.

responses are much more volatile than those from inverted vector autoregressions. However, the summary statistics that are presented in Table 2 convey a rather clear picture. The left hand side of the table contains the maximum (minimum) responses of output (unemployment) over the 12 period horizon that we consider, and the right hand side shows cumulative responses that are obtained by summing over the responses of all 12 horizons. Note, for interpretation, that these responses are based on a single equation approach. We therefore do not compute multipliers in the same sense as above, where we reported the output response as a ratio over the own response of government spending following a fiscal shock. The values shown in Table 2 cannot, strictly speaking, be interpreted as fiscal multipliers, but are simply the responses, or the cumulated responses, of the respective left-hand variable. This also explains why the cumulative responses in the right hand part of the table are much larger than the cumulative multipliers computed earlier, since the former implicitly attribute all current and future changes in the activity variable to a one-time innovation in g_t without normalizing by relating these to the future changes in government spending itself.

These considerations explain, why the individual response values that are shown in Table 2 look somewhat different to the ones presented earlier for the VAR models. However, in regard of the

main question pursued here, namely the existence of nonlinear effects of fiscal shocks, the results in Table 2 convey a very similar message to the one arrived at in previous sections. In particular, the predicted output response is positive throughout and noticeably larger at low quantiles of its distribution. These low quantile responses, which correspond to the effects of fiscal shocks on GDP when the latter is predicted to be noticeably below its long-run trend, are strongly statistically significant. The same is not true for fiscal shock effects at the higher quantiles corresponding to above-trend predictions of GDP, where the responses are weaker and it is less clear whether they differ significantly from zero in a statistical sense. Conversely, the unemployment rate does not seem to be much affected by fiscal shocks, unless it is forecasted to be in the rightmost parts of its distribution. The latter case corresponds to a government spending innovation taking place in a situation where the baseline forecast is a depressed labor market h-periods ahead. In this situation, the unemployment rate is predicted to decrease strongly and statistically significantly. As the lower part of Table 2 shows, the results are quantitatively somewhat different if one controls for fiscal anticipation effects, but the qualitative conclusion to be derived remains largely the same.

To sum up, we conclude that our main empirical findings – that the effects of fiscal shocks are markedly different depending on the question at which particular quantile of the conditional distribution of the outcome variables they are estimated – are rather robust to the alternative estimation technique employed in this subsection.

4   Conclusion

This paper has documented nonlinear effects of government spending shocks in the US using a quantile regression approach. Estimating the effects of fiscal expansions at different quantiles of the conditional distribution of economic activity reveals that fiscal policy is more effective in a depressed economy than in normal times. If output is low relative to trend and unemployment is relatively high, fiscal multipliers are large. On the contrary, at high quantiles of the conditional distribution of detrended output and low quantiles of the conditional unemployment distribution, fiscal multipliers are small and most often insignificant. These findings are robust to different empirical strategies and to controlling for anticipation effects. Our results also go through if we exclude the years 2008-2013 from our sample. This suggests that our findings are not driven by

21

the Great Recession, during which monetary policy was constrained by the zero lower bound on nominal interest rates.

Our results have interesting policy and theoretical implications. With regard to stabilization policy, our results suggest that government spending expansions are an effective tool to stimulate an economy out of a recession. With regard to economic modeling, our findings strongly support theories that can explain nonlinear effects of fiscal policy. However, we have shown that the zero lower bound, which is one of the most prominent theories of nonlinearities in the effectiveness of fiscal policy, cannot, or at least not entirely, explain our results. Developing and testing alternative theories of countercyclical fiscal multipliers is thus a fruitful area for future research.

References

Auerbach, Alan J. and Yuriy Gorodnichenko, 2012. Measuring the Output Responses to Fiscal Policy. American Economic Journal: Economic Policy 4, 1-27.

Auerbach, Alan J. and Yuriy Gorodnichenko, 2013. Fiscal Multipliers in Recession and Expansion. In: Fiscal Policy after the Financial Crisis, Alberto Alesina and Francesco Giavazzi, eds., University of Chicago Press.

Berkowitz, Jeremy, Ionel Birgean and Lutz Kilian, 1999. On the Finite-Sample Accuracy of Nonparametric Resampling Algorithms for Economic Time Series. In: Advances in Econometrics: Applying Kernel and Nonparametric Estimation to Economic Topics 14, T. B. Fomby and C. Hill, eds., JAI Press.

Blanchard, Olivier and Daniel Leigh, 2013. Growth Forecast Errors and Fiscal Multipliers. American Economic Review 103, 117-120.

Blanchard, Olivier and Roberto Perotti, 2002. An Empirical Investigation of the Dynamic Effects of Shocks to Government Spending and Taxes on Output. Quarterly Journal of Economics 117, 1329-1368.

Christiano, Lawrence J., Eichenbaum, Martin and Sergio Rebelo, 2011. When is the Government Spending Multiplier Large? Journal of Political Economy 119, 78-121.

Cecchetti, Stephen G. and Hong Li, 2008. Measuring the Impact of Asset Price Booms Using Quantile Vector Autoregressions. Brandeis University, mimeo.

Eggertsson, Gauti B., 2011. What fiscal policy is effective at zero interest rates? In: NBER Macroeconomics Annual 2010, 59-112. University of Chicago Press.

Jordà, Òscar, 2005. Estimation and Inference of Impulse Responses by Local Projections. American Economic Review 95, 161-182.

Kilian, Lutz and Yun Jung Kim, 2011. How reliable are local projection estimators of impulse responses? Review of Economics and Statistics 93, 1460-1466.

Koenker, Roger, 2005. Quantile Regression, Econometric Society Monograph Series, Cambridge University Press.

Koenker, Roger and Gilbert Bassett, Jr., 1978. Regression Quantiles. Econometrica 46, 33-51.

Ramey, Valerie A., 2011. Identifying Government Spending Shocks: It's all in the Timing. The Quarterly Journal of Economics 126, 1-50.

Ramey, Valerie A. and Sarah Zubairy, 2014. Government Spending Multipliers in Good Times and in Bad: Evidence from U.S. Historical Data. University of California at San Diego and Texas A&M University, mimeo.

Appendix A: Data sources and definitions

Table A: Data sources

	Series Title	Series ID	Source
(1)	Real Gross Domestic Product	GDPC1	BEA
(2)	Real Government Consumption Expenditures and Gross Investment	GCEC96	BEA
(3)	Gross Domestic Product: Implicit Price Deflator	GDPDEF	BEA
(4)	Government Current Tax Receipts	W054RC1Q027SBEA	BEA
(5)	Government Current Transfer Payments	A084RC1Q027SBEA	BEA
(6)	Government Current Receipts: Contributions for Government Social Insurance	W782RC1Q027SBEA	BEA
(7)	Effective Federal Funds Rate	FEDFUNDS	BFED
(8)	Civilian Unemployment Rate	UNRATE	BLS
(9)	Defense News,Percent of Lagged Nominal GDP	RN	Ramey

Notes: BLS: U.S. Bureau of Labor Statistics, BEA: U.S. Bureau of Economic Analysis, BFED: Board of Governors of the Federal Reserve System. Variables (1)-(8) were downloaded from the FRED database. Variable (9) was downloaded from Valerie Ramey's website.

Table B: Definition of data variables

Variable	Construction	Description
Output	qd $(GDPC1)$	cyclical component of log of real GDP
Government Spending	qd $(GCEC96))$	cyclical component of log of real government spending
Net Taxes	qd $\left(\frac{(W054...)+(W782...)-(A084...)}{GDPDEF} \right)$	cyclical component of log of real net taxes
Real interest rate	$\frac{FEDFUNDS}{100} - \log\left(\frac{GDPDEF}{GDPDEF(-1)} \right) \cdot 4$	annualized ex-post real interest rate
Unemployment	$UNRATE/100$	unemployment rate
Ramey News	$RN/100$	spending news variable

Notes: The function qd stands for log-quadratic detrending, (-1) indicates a one-quarter lag.

Appendix B: Bootstrap algorithms

This appendix describes the bootstrap algorithms used to estimate standard error bands for the impulse response functions presented in the paper. For the quantile VAR models, we generate standard errors in the conventional way as follows: after estimation of a model, we resample with replacement from the residuals of all equations contained in the VAR. From these, a new sample

dataset is created using the estimated parameters. This sample data set is used to generate a new set of impulse responses. This procedure is repeated to produce 5000 sets of bootstrap impulse responses. The point-wise standard deviation of these is taken as the estimate of the standard error of the model's impulse response functions.

For the quantile local projection estimates, this simple bootstrap algorithm is not available, because the approach requires to estimate h-period ahead forecast equations which do not have a first-order autoregressive representation. We thus use the blocks-of-blocks bootstrap technique advocated by Kilian and Kim (2011) and discussed in more depth in Berkowitz et al. (1999). Specifically, from the left-hand and the right-hand side data in equation (3) we form all possible blocks of data containing $\mathfrak{L}$ consecutive observations. This is intended to preserve the dependencies in the data. Then we randomly draw blocks of data and form a new sample of the same size of the original data sample. From this bootstrap sample, a new parameter estimate of the quantile local projection equation is produced and stored. After repeating this procedure 5000 times, we construct 90% confidence intervals by omitting the largest and smallest five per cent of the estimated bootstrap parameters. The width of the confidence interval divided by 1.65 is used as an approximate standard error of the parameter in question. Following the suggestion in Berkowitz et al. (1999), we set the block length $\mathfrak{L}$ at $T^{1/3}$, where T is the sample size.

To construct the confidence bands for the fiscal multipliers and the cumulative multipliers described in the paper, we proceed similarly by directly applying the bootstrap algorithm to the multiplier estimate. In particular, we generate 5000 bootstrap samples with the blocks-of-blocks method described above, and estimate the respective multiplier (or cumulative multiplier) measure from each of these. From the distribution of bootstrapped sample multipliers or cumulative multipliers, we construct the 90% confidence intervals shown in the figures.