Estimation of Impulse response functions with term structure local projections

by

James McNeil
Dalhousie University

Working Paper No. 2020-05

August 06, 2020

DEPARTMENT OF ECONOMICS

DALHOUSIE UNIVERSITY
6214 University
Avenue PO Box 15000
Halifax, Nova Scotia, CANADA
B3H 4R2

Estimation of impulse response functions with term structure local projections

James McNeil*

August 4, 2020

Abstract

I propose a new method to calculate impulse response functions of the term structure using term structure local projections. The method imposes a set of testable, linear, cross-equation restrictions that the term premium does not respond to a given economic shock. This substantially reduces the number of parameters to estimate and leads to notable efficiency gains compared to the standard local projection approach. Using this method, I show that contractionary monetary policy shocks increases the term premium and the effect lasts for several months. By contrast, tax and total factor productivity shocks have much smaller effects on the term premium.

Keywords: Local projections, monetary policy, term premium, term structure.

JEL classification: C32, E43, E52.

*Department of Economics, Dalhousie University, Halifax, Nova Scotia, Canada. Email: `mcneilj@dal.ca`. This research was supported by the Social Sciences and Humanities Research Council of Canada. I am grateful to Gregor Smith, Allan Gregory, James MacKinnon, Fransisco Ruge-Murcia, and seminar participants at Queen's University for helpful comments.

1 Introduction

Macroeconomists are often interested in the response of many highly correlated variables to a policy change. An example is the response of the term structure of interest rates—the joint response of yields on bonds of different maturities—to a monetary policy change. While central banks generally conduct monetary policy by adjusting short-run nominal interest rates, borrowing and lending typically occur at longer time horizons and hence it is through the response of long-term interest rates that monetary policy is likely to affect the real economy. This motivates the desire for a better understanding of the links between short- and long-run interest rates and in particular how a change to the short-run nominal interest rate affects the entire term structure of interest rates. Because interest rates at different maturities are highly correlated, researchers may anticipate that their responses after a policy move will also be highly correlated, which could be exploited to achieve more precise estimates of the policy effect.

The challenge researchers face in this context is how to specify the relationship between the potentially numerous response variables, while remaining as agnostic as possible about the precise nature of this relationship. In this paper I propose a solution to this issue by estimation of impulse response functions using term structure local projections. Proposed by Jordà (2005), local projections are a method to calculate impulse response functions by directly projecting future values of the response variables onto current values of the causal variable, akin to the method of direct multi-step forecasting. Plagborg-Møller and Wolf (2020) show that, in population, local projections and vector autoregressions estimate the same impulse response functions, although they may have different finite sample properties.

The term structure local projections I introduce in this paper use the law of iterated expectations to derive a set of linear cross-equation restrictions for the local projections regressions which provide impulse responses of the entire term structure. The method can be seen as an extension of the standard local projection framework, using a set of restrictions to link together estimates across the term structure without making any assumptions about the shape of the term structure itself. Further, the restrictions substantially reduce the number of parameters to estimate, which can provide significant efficiency gains, and rejection of the restrictions may also have an interesting economic interpretation.

In the main application of this paper I use the term structure local projections to estimate the response of the term structure of interest rates to a monetary policy shock. In this case rejection of the restrictions can be interpreted as evidence of a significant response of the term premium to a monetary policy shock. I show that there is some evidence that the term premium increases in response to a contractionary monetary policy shock, but the effect is short-lived. The restrictions offer a substantial efficiency gain over the usual local projections of the short-term interest rate, in some cases as large as 30%. I then apply these restrictions to estimate the response of the term structure of interest rates to tax and total factor productivity shocks. Both these shocks have a quantitatively smaller impact than monetary policy shocks and only the tax shocks have a significant effect on the term premium overall.

2 Research context

The restrictions I propose in this paper are similar to those proposed by Sargent (1979) for vector autoregressions. However, because the restrictions I derive here are always linear, estimation is far simpler and not susceptible to the multiplicity of solutions documented by Kurmann (2007). In addition, these estimates retain the traditional benefits of local projections compared to vector autoregressions: estimation errors do not amplify through the impulse responses and no assumptions about invertibility are required. They also offer benefits over the standard local projection method. As noted by Kilian and Kim (2011), local projections typically require estimates of many parameters and, as a result, precision can sometimes be low, leading to large confidence intervals. Barnichon and Brownlees (2019) propose a method to smooth local projections by approximation with B-splines, which results in narrower confidence bands at the price of biased coefficients. Miranda-Agrippino and Ricco (2017) propose a Bayesian local projection framework to reduce the estimation variance of standard local projections by shrinking them towards estimates from a vector autoregression, optimally balancing the bias-variance trade-off between the two methods. Because the restrictions I propose in this paper substantially reduce the number of parameters to estimate, they may also result in narrower confidence intervals, but they do not introduce any bias under the null hypothesis that the restrictions are correct.

An alternative approach to estimate the response of highly correlated variables is to first summarize the dynamics of these variables using a term structure model, which achieves dimensionality reduction by relating the cross-section of observables to a few common latent components, and then estimates the responsiveness of these latent components to a policy change. The response of the variables of interest can then be reconstructed from the response of the latent components and the structure of the underlying term structure model. This is the approach taken by Diebold *et al.* (2006) who estimate the response of the term structure of interest rates after several different macroeconomic shocks. A drawback to this approach is that results are then contingent on the particular term structure model employed. Furthermore, estimation of these models by maximum likelihood methods is often challenging and estimates can be sensitive to the choice of optimization algorithm and starting values. Estimation of the policy impact on each variable independently is another alternative but this ignores the common features of the data, which should be helpful to achieve more precise estimates, and may also preclude the study of interesting research questions.

One such question is the ability of the central bank to influence long-run interest rates, and whether this occurs through expectations of future short-run interest rates or changes to the term premium. When short-run interest rates are constrained at their effective lower bound, monetary policy may turn to focus on how to lower long-term interest rates, as seen during the recent recessions beginning in 2008 and 2020. Because long-term interest rates may be decomposed into the sum of expected future short-term interest rates and the term premium, monetary policy may influence long-term interest rates by changing expectations of future short-run interest rates or adjusting the term premium. For example, forward guidance is designed to influence long-term interest rates by changing expectations of future short-run interest rates. However, the efficacy of these policies depends on the response of the term premium, which may amplify or mitigate the response of long-run interest rates depending on the sign of response. Hence, an understanding of how the components of long-term interest rates respond to monetary policy adjustments is an important question.

The expectations hypothesis states that the term premium is constant, which implies that the response of long-term interest rates after a policy adjustment is entirely determined by the sum of the response of expected future short-term interest rates. Many empirical papers, however, reject the expectations hypothesis (See Fama and Bliss (1987), Campbell and Shiller (1991), Backus *et al.* (2001), Cochrane and Piazzesi (2005) and Sarno *et al.* (2007), for example). The most common interpretation of this finding is that the term premium is not constant, a view supported by Crump *et al.* (2018) and Bauer and Rudebusch (2020) who document time variation in the term premium. Given this time variation in the term premium and the importance of the term premium to explain long-term interest rates, of interest is what can explain these changes in the term premium, and in particular the role of monetary policy in this context.

Complicating this issue is that the term premium is not directly observable, and estimates can vary widely depending on assumptions about the stationarity of interest rates and choice of term structure model, as demonstrated by Bauer and Rudebusch (2020). The advantage of the method I propose in this paper is that it requires only the restriction that bond yields of different maturities respond in the same way to a monetary policy shock. This restriction is testable and a rejection of the restriction can be interpreted as imposing the condition that the term premium does not respond to changes in monetary policy. I show that over the full impulse response horizon this restriction is not rejected by the data, indicating no evidence that monetary policy actions cause long-lasting variations in the term premium.

However, there is a significant response of the term premium at short horizons.

Existing empirical evidence on the behaviour of the term premium has been mixed: Crump *et al.* (2018) and Tillmann (2020) find that a contractionary monetary policy shock lowers the term premium, which is the opposite finding of Gertler and Karadi (2015) and Abrahams *et al.* (2016) who find that the term premium responds in the same direction as the monetary policy shock. One potential explanation for this discrepancy is the different treatments of the expectations of future short-run interest rates, which are necessary to uncover the term premium. As noted by Gertler and Karadi (2015), their finding could be because they calculate the response of the term premium indirectly using impulse responses of short-run interest rates from their vector autoregression. These expectations may not accurately reflect the change in expected interest rates of financial market participants and, because the expectations are generated from a vector autoregression, this misspecification will amplify at more distant impulse response horizons. Because local projections are generally less sensitive model to misspecification they should be less sensitive to this issue. They also do not rely on survey forecasts of future interest rates, which may not coincide with the expectations of financial market participants and hence could misrepresent the term premium.

A better understanding of the behaviour of the term structure is also of interest from a theoretical perspective. As Gertler and Karadi (2015) note, standard macroeconomic theory predicts that the response of long-term interest rates occurs only through adjustments to expected future short-term interest rates with the term premium remaining constant. Boivin *et al.* (2010) provide an extensive review of the transmission of monetary policy including neo-classical models, models with financial frictions, and a New Keynsian DSGE model, all of which share the common feature that the term premium is not responsive to monetary policy. Without some modification, the term premium is constant in standard macroeconomic models, which would appear to contradict empirical evidence against the expectations hypothesis.

More recent theoretical macroeconomic models have been able to generate a time-varying term premium by introducing habit formation, as in Rudebusch and Swanson (2008) and Christoffel *et al.* (2011); the recursive preferences proposed by Epstein and Zin (1989), as in Rudebusch and Swanson (2012) and van Binsbergen *et al.* (2012); or financial market frictions, as in Chen *et al.* (2012), Gertler and Karadi (2013), and Carlstrom *et al.* (2017). Habit models typically cannot match the size of the term premium with sacrificing model fit to other macroeconomic variables, in particular inflation. Epstein-Zin preferences have been more successful at matching both macro and financial data, but only with a very large value for the coefficient of relative risk aversion. These models also find that, while the term premium should decline after a monetary policy tightening, the magnitude of the effect is very small so that only implausibly large monetary policy shocks could induce the variability apparent in the term premium. By contrast, total factor productivity and government spending shocks have much larger effects. Models with financial frictions can produce sizable movements in the term premium but have typically focused on the effects of unconventional monetary policy—quantitative easing in particular—rather than conventional monetary policy.

3 Term structure local projections

In this section I introduce the term structure local projections. I first show how the standard local projection method can be applied to the set of interest rates individually, disregarding any relationship across the term structure, and then derive a set of cross-equation parameter restrictions which give the term structure local projections.

Let i_t denote the policy interest rate, ϵ_t a series of exogenous monetary policy shocks, $i_{3,t}$ the interest rate on a three-month treasury bill, and X_t a vector of control variables. The monetary policy shocks are unobserved but we do observe an instrumental variable, z_t, which is correlated with the monetary policy shocks, uncorrelated with all other shocks driving the interest rate, and uncorrelated with both itself and other shocks at all leads and lags, as outlined by Stock and Watson (2018). Then the response of the three-month treasury bill h months after a 100 basis point increase to the policy interest rate, defined as:

$$E[i_{3,t+h}|X_{t-1}, \epsilon_t = 1] - E[i_{3,t+h}|X_{t-1}, \epsilon_t = 0] = \gamma_{3,h}, \tag{1}$$

can be estimated from the local projections:

$$i_{3,t+h} = \beta_{3,h} X'_{t-1} + \gamma_{3,h} i_t + u_{3,h,t}, \tag{2}$$

by instrumenting for i_t with z_t.

In addition, we also observe many yields on long-term treasury securities, each of which matures at some multiple n of the maturity on the short-run interest rate, in this case the three-month treasury bill. Local projections estimating the response of these yields after the same monetary policy shock can be written as:

$$i_{n\times 3,t+h} = \beta_{n\times 3,h} X'_{t-1} + \gamma_{n\times 3,h} i_t + u_{n\times 3,h,t}, \tag{3}$$

where again instrumenting for i_t with z_t gives consistent estimates of the coefficients $\gamma_{n\times 3,h}$.

Equations (2) and (3) can be estimated independently, giving impulse response functions for each of the bond yields, denoted $\gamma_{3,h}$ and $\gamma_{n\times 3,h}$. However, the responses of yields on securities of different maturities after a monetary policy adjustment may be highly correlated, which suggests that more efficient estimates may be achieved.

Following Kozicki and Tinsley (2005), the yield on a bond maturing $n \times 3$ months into the future can be expressed as the sum of expected future values of the short yield and a term premium, $\phi_{n,t+h}$,

$$i_{n\times 3,t+h} = \frac{1}{n}\sum_{j=0}^{n-1} E[i_{3,t+j\times 3+h}|x_{t+h}] + \phi_{n,t+h}, \tag{4}$$

where x_{t+h} is the information set of financial market participants. I assume that $\{X_t, z_t\} \in x_t$ because the econometrician observes only a subset of the information sets of market participants. Taking conditional expectations of (4) and applying the law of iterated expectations gives:

$$E[i_{n\times 3,t+h}|X_{t-1}, \epsilon_t] = \frac{1}{n}\sum_{j=0}^{n-1} E[i_{3,t+j\times 3+h}|X_{t-1}, \epsilon_t] + E[\phi_{n,t+h}|X_{t-1}, \epsilon_t]. \tag{5}$$

Differencing this conditional expectation in the same manner as the impulse response function defined in equation (1) gives:

$$\begin{aligned} &E[i_{n\times 3,t+h}|X_{t-1}, \epsilon_t = 1] - E[i_{n\times 3,t+h}|X_{t-1}, \epsilon_t = 0] = \\ &\frac{1}{n}\sum_{j=0}^{n-1} (E[i_{3,t+j\times 3+h}|X_{t-1}, \epsilon_t = 1] - E[i_{3,t+j\times 3+h}|X_{t-1}, \epsilon_t = 0]) + \\ &E[\phi_{n,t+h}|X_{t-1}, \epsilon_t = 1] - E[\phi_{n,t+h}|X_{t-1}, \epsilon_t = 0], \end{aligned} \tag{6}$$

which says that the response of the long-term bond yield after a monetary policy shock equals the sum of the current and future responses of the short-term bond yield plus the response of the term premium. Imposing the restriction that the term premium does not respond to a monetary policy adjustment, $E[\phi_{n,t+h}|X_{t-1}, \epsilon_t = 1] = E[\phi_{n,t+h}|X_{t-1}, \epsilon_t = 0]$, gives the following set of linear cross-equation restrictions:

$$\gamma_{n\times 3,h} = \frac{1}{n}\sum_{j=0}^{n-1} \gamma_{3,h+j\times 3}. \tag{7}$$

These restrictions allow the impulse responses of yields on all long-term bonds to be written as an average of the impulse responses of the yield on the short-term bond, which substantially reduces the number of parameters to estimate. By incorporating data for yields on bonds of different maturities the term structure local projections should also provide more accurate estimates of the response of interest rates after a monetary policy shock. Coibion

(2012) shows that much of the discrepancy between the existing findings on the effects of monetary policy can be explained by different responses of interest rates so that accurate estimates of this response should be of interest to researchers. The restrictions are also testable using a Wald statistic and rejection can be interpreted as a test of the hypothesis that the term premium does not respond to monetary policy shocks.

There are several advantages to the term structure local projections over a vector autoregression approach. First, imposing these restrictions on a vector autoregression would require non-linear cross-equation restrictions, making estimation much more difficult. The term structure local projection restrictions are always linear so estimation is simple. Of course, the VAR could instead be estimated without any restrictions, as in the unrestricted local projections. But that may preclude the inclusion of all observable bond yields without running up against curse of dimensionality issues. Second, the term structure local projections produce distinct restrictions for each impulse response horizon and bond maturity, so testing can be conducted with more precision. Restrictions on the vector autoregression will hold at all impulse response horizons, which may mask changes in the term premium. Finally, Kurmann (2007) shows that the standard vector autoregression restrictions can take the form of a high-order polynomial for even moderately-sized vector autoregressions, resulting in multiple solutions.

4 Results

I use the following panel of annualized yields of constant-maturity treasury securities, taken from the Federal Reserve's H.15 report, to make up the left-hand-side variables of equations (2) and (3): the three-month, six-month, and one-year treasury bills, and the two-year, three-year, and five-year treasury notes. As control variables I include a constant and twelve lags each of the log of industrial production, the log of the consumer price index, the Gilchrist and Zakrajšek (2012) excess bond premium, and the first three principal components of all interest rates. I use principal components, rather than including lags of all yields as controls, because of the high degree of multicollinearity in the set of yields.

This set of control variables is chosen in order to follow Gertler and Karadi (2015) as closely as possible. For the same reason I take the one-year treasury bill rate as the policy rate and the change in three-month ahead federal funds futures in a thirty minute window around a monetary policy announcement as the instrumental variable. As demonstrated by Miranda-Agrippino and Ricco (2017), because of differences between the information sets of policy makers and financial market participants, this instrument contains an information component; monetary policy announcements convey information about the true state of the economy. Accounting for this component is especially important in this context because this information may affect expectations of future interest rates, which make up a substantial component of long-run interest rates and, by extension, the term premium. To account for this I use the modified instrument proposed by Miranda-Agrippino and Ricco (2017), which first projects changes in federal funds futures onto its own lags and Greenbook forecasts. This instrument has a first-stage F-statistic of 10.93, above the threshold value of 10 proposed by Stock *et al.* (2002) so that weak instruments do not appear to be a concern. This means that the policy rate i_t corresponds with a point along the yield curve, $i_{12,t}$ (that is, $n = 4$ in the notation of equation (3)). The restrictions given by (7) will then impose the condition that, on average, all other interest rates follow the response of the policy variable. The sample period is 1990–2014 and data is at a monthly frequency.

The term structure local projections can be interpreted as imposing the restriction that the term premia do not respond after a monetary policy shock. This is testable with the Wald statistic, which follows a χ^2 distribution with degrees of freedom equal to the number of parameter restrictions. The joint test of all 182 restrictions given by (7) has a test statistic of 181.32 with an associated p-value of 0.5, indicating no significant evidence against the null hypothesis that the term premium does not respond to changes in monetary policy. This can be interpreted as evidence that, overall, the term premium does not respond to changes in monetary policy.

Figure 1 shows the response of the term structure of interest rates after a contractionary monetary policy shock normalized to increase the unrestricted response of the one-year treasury bill rate by 100 basis points. The term structure local projections use all interest rates described above, along with the restrictions given by equations (7), to estimate the response of the entire term structure after a monetary policy shock. The figure can be interpreted as the response of the short-run interest rate some number of months h after a monetary policy change. Because the restrictions express the response of long-term interest rates as the average response of the current and future short-term interest rates, the figure can equivalently be interpreted as the response of the yield curve to a monetary policy change, under the restriction that the term premium does not respond to the policy change.

The response of interest rates closely matches the response of the one-year interest rate documented by Gertler and Karadi (2015) and the responses of the federal funds rate documented by Coibion (2012) based on the Romer and Romer (2004) monetary policy shocks. The monetary policy contraction increases interest rates on impact and demonstrates a hump-shaped response in the first few months after the policy change. The effect remains significant for approximately one year, after which interest rates respond negatively before the effect eventually wears off after approximately three years. Thus, contractionary monetary policy twists the yield curve, pushing up interest rates in the short term while pulling them down in the medium term.

The term structure local projections offer substantial efficiency gains compared with standard local projections. Table 1 shows the ratio of the standard errors for the unrestricted and restricted estimates of $\gamma_{3,h}$, where the restrictions are defined by (7). Standard errors for both estimators are calculated using a Newey-West HAC with 12 lags. Table 1 shows that in nearly all cases the ratio of standard errors is greater than one, indicating that the restricted estimates have a lower standard error than the unrestricted estimates, and in many cases the efficiency gain is quite large. The standard errors for the unrestricted local projections are always larger than those for term structure local projections I propose in this paper. The new estimates have an efficiency gain as large as 30%, demonstrating a notable improvement over the unrestricted estimates. The final column of Table 7 shows that the average ratio of standard errors at all horizons is 1.21, once again indicating a substantial efficiency gain for the term structure local projections.

As described above, the joint test of no response of the term premium at all bond maturities and impulse response horizons fails to reject. However, one of the great advantages of the local projections is that they allow this test to be conducted with much greater precision. This can be done using the responses of the term premia for each of the long-term bonds after a monetary policy shock, calculated from the unrestricted local projections as:

$$\Delta tp(m,h) = \hat{\gamma}_{m,h} - \frac{1}{n}\sum_{j=0}^{n-1} \hat{\gamma}_{3,h+3\times j}, \tag{8}$$

where $m = 3 \times n$ denotes the maturity in months. Notice that, because (7) relates the response of a long-term bond to n responses of the short-term bond, the term structure local projections will truncate the maximum impulse response horizon for long-term bonds, which is emphasized in the figure.

Figure 2, which shows these responses, sheds greater light on how the term premia respond to monetary policy shocks. At most horizons and maturities the hypothesis of no response cannot be rejected, which matches the finding of the Wald statistic. However, there is evidence that the term premium significantly increases in response to a contractionary monetary policy shock on impact and in the first few periods after the shock. Intuitively, the increase in short rates does not last long enough to match the increase in long rates, which results in a temporary but notable increase in the term premium. This effect is apparent for treasury securities at all maturities but is especially apparent for bonds maturing between two and five years ahead. The average impact effect on the term premium for these four maturities is 70 basis points, or 70% of the impact effect of the monetary policy shock on the term structure. This indicates that the response of long-term bond yields is, on average, 70 basis points larger than the average response of the short-term bond yield over the maturity of the long-term bond, so that the term premium has an amplifying effect on long-term bonds in the periods following a monetary policy shock. For bonds maturing two, three, and five years ahead the term premium responds by more than 100 basis points on impact.

I note that my finding that the term premium responds in the same direction as the monetary policy shock is the opposite prediction from the theoretical model proposed by Rudebusch and Swanson (2012), which finds that contractionary monetary policy lowers the term premium, but consistent with empirical evidence by Gertler and Karadi (2015) and Abrahams *et al.* (2016).

5 Tax and total factor productivity shocks

I now extend this analysis to other potential explanations for the variability of the term premium. In particular, I consider the effects of exogenous changes to total factor productivity and tax revenue. The local projections can now be written as:

$$i_{n\times 3,t+h} = \beta_{n\times 3,h} X'_{t-1} + \gamma_{n\times 3,h} y_t + u_{n\times 3,h,t}, \tag{9}$$

where y_t is an endogenous variable, specifically the growth rate of one of GDP or federal tax revenue, both in real per capita terms. Because each of these is an endogenous variable I again require instrumental variables in order to identify $\gamma_{n\times 3,h}$. As an instrument for exogenous variation in GDP I use the utilization-adjusted total factor productivity series proposed by Fernald (2014). The tax instrument is the narrative series of unanticipated exogenous tax changes proposed by Mertens and Ravn (2013), and extended by Liu and Williams (2019). The first-stage regressions have F-statistics of 13.48 for total factor productivity shocks and 13.79 for tax revenue shocks tests, indicating that both instruments are sufficiently correlated with the relevant endogenous variable that weak instruments is not a concern.

The data are now at a quarterly frequency and I modify the set of control variables X_t to include a constant and four lags each of the growth rates of GDP, federal government spending and federal tax revenue, the log of the consumer price index, the Gilchrist and Zakrajšek (2012) excess bond premium, and the first three principal components of all interest rates. The set of interest rates making up the left-hand-side variables is the same as in the previous section. Because the instruments for tax and total factor productivity shocks are available for a longer period than the monetary policy instrument used in the previous section, I extend the sample to cover 1982–2016.

Figures 3 and 4 show the responses of the term premium after total factor productivity and tax shocks along with 68% and 90% confidence intervals calculated using a Newey-West HAC with four lags. The scales of the shocks are normalized to increase the growth rates of GDP or tax revenue by one percentage point. The effect of a positive total factor productivity shock is to decrease the term premium on bonds maturing two and three years into the future, after a lag of several quarters. The effect is temporary, wearing off after approximately two years, and has a maximum effect of about 20 basis points. There is no significant response of the term permia for bonds maturing in six months, one year, or five years, which explains why the joint test of all restrictions fails to reject: the Wald statistic for the joint test of no response of the term premium to a total factor productivity shock is 45.51 with 64 degrees of freedom and a p-value of 0.96.

The initial effect of a tax shock is to significantly decrease the term premium on bonds with maturity of one year or more. Although the effect is statistically significant, it is quantitatively small, affecting the term premium by less than five basis points. There is also some evidence of a delayed positive effect on the term premium: the response of the term premia on both two- and three-year bonds is significantly positive five quarters after the shock. Compared with the effects of monetary policy and total factor productivity shocks, there is now more evidence against the joint restrictions: the Wald statistic for the joint test of no response of the term premium to a tax shock is 78.16 with 64 degrees of freedom and a

p-value of 0.11. From Figure 4 this appears to be primarily driven by the strong significance of the impact responses of the term premia on bonds with maturity one year or more.

The theoretical model proposed by Rudebusch and Swanson (2012) suggests that the quantitative effects of monetary policy shocks on the term premium are much smaller than fiscal policy shocks and especially total factor productivity shocks. While the normalization of the impact effect on the local projections makes this comparison somewhat challenging, the local projections do not appear to support this conclusion. The maximal effect of a monetary policy shock on the term premium is an order of magnitude larger than the maximal effect of a total factor productivity shock. To match the scale of the response of the term premium to a 100 basis point monetary policy shock would require a total factor productivity shock that increases the quarterly GDP growth rate by approximately ten percentage points. Likewise for tax shocks, which have an even smaller impact on the term premium. Despite this discrepancy in magnitude, the sign of the response of the term premium to total factor productivity shocks does match the prediction of Rudebusch and Swanson (2012).

6 Conclusion

In this paper I propose a set of linear cross-equation restrictions that can be used to calculate impulse responses of the term structure using local projections. In addition to the usual advantages of local projections, the term structure local projections substantially reduce the number of parameters to estimate, which can lead to notable efficiency gains, and require no underlying model of the term structure. Estimation of term structure models via maximum likelihood methods is a notoriously challenging exercise and estimates are often sensitive to choice of starting values and optimization algorithms. In addition, impulse response functions are conditional on the underlying term structure model and may be sensitive to estimation and misspecification errors. By contrast, the restrictions implied by the term structure local projections are always linear and hence estimation is straightforward.

I use the term structure local projections to estimate the response the term premium to monetary policy, fiscal policy, and total factor productivity shocks. The estimates demonstrate a substantial efficiency gain compared with traditional local projections of the short-run interest rate, in some cases as much as 30%. A contractionary monetary policy shock increases the term premium on impact but the effect is short-lived, lasting only a few periods. Tax shocks have a statistically significant effect on the term structure of interest rates, first decreasing and then increasing the term premium, and total factor productivity shocks decrease the term premium on bonds maturing in two and three years. The quantitative effects of total factor productivity and tax shocks also appear to be much smaller than that of monetary policy shocks, a useful finding for researchers interested in theoretical models with time-varying term premia.

Many other applications of the term structure local projections are possible. This method could be used to estimate the term structures of expectations and may also have applications outside of the context of term structures. One possible extension would be to draw upon theoretical relationships between variables or even accounting identities to derive a similar set of cross-equation restrictions introduced in this paper.

References

Abrahams, M., Adrian, T., Crump, R. K., Moench, E., and Yu, R. (2016). Decomposing real and nominal yield curves. *Journal of Monetary Economics*, 84:182–200.

Backus, D., Foresi, S., Mozumdar, A., and Wu, L. (2001). Predictable changes in yields and forward rates. *Journal of Financial Economics*, 59(3):281–311.

Barnichon, R. and Brownlees, C. (2019). Impulse response estimation by smooth local projections. *Review of Economics and Statistics*, 101(3):522–530.

Bauer, M. D. and Rudebusch, G. D. (2020). Interest rates under falling stars. *American Economic Review*, 110(5):1316–1354.

Boivin, J., Kiley, M. T., and Mishkin, F. S. (2010). How has the monetary transmission mechanism evolved over time? volume 3 of *Handbook of Monetary Economics*, 369–422.

Campbell, J. Y. and Shiller, R. J. (1991). Yield spreads and interest rate movements: A bird's eye view. *The Review of Economic Studies*, 58(3):495.

Carlstrom, C. T., Fuerst, T. S., and Paustian, M. (2017). Targeting long rates in a model with segmented markets. *American Economic Journal: Macroeconomics*, 9(1):205–242.

Chen, H., Cúrdia, V., and Ferrero, A. (2012). The macroeconomic effects of large-scale asset purchase programmes. *The Economic Journal*, 122(564):F289–F315.

Christoffel, K., Jaccard, I., and Kilponen, J. (2011). Government bond risk premia and the cyclicality of fiscal policy. (ECB working paper, No. 1411).

Cochrane, J. H. and Piazzesi, M. (2005). Bond risk premia. *American Economic Review*, 95(1):138–160.

Coibion, O. (2012). Are the effects of monetary policy shocks big or small? *American Economic Journal: Macroeconomics*, 4(2):1–32.

Crump, R. K., Eusepi, S., and Moench, E. (2018). The term structure of expectations and bond yields. (Federal Reserve Bank of New York Staff Reports, No. 775).

Diebold, F. X., Rudebusch, G. D., and Aruoba, S. B. (2006). The macroeconomy and the yield curve: A dynamic latent factor approach. *Journal of Econometrics*, 131(1-2):309–338.

Epstein, L. G. and Zin, S. E. (1989). Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework. *Econometrica*, 57(4):937–969.

Fama, E. F. and Bliss, R. R. (1987). The information in long-maturity forward rates. *The American Economic Review*, 77(4):680–692.

Fernald, J. (2014). A quarterly, utilization-adjusted series on total factor productivity. (Federal Reserve Bank of San Francisco, Working Paper No. 2012-19).

Gertler, M. and Karadi, P. (2013). QE 1 vs. 2 vs. 3. . . : A framework for analyzing large-scale asset purchases as a monetary policy tool. *International Journal of Central Banking*, 9:49.

Gertler, M. and Karadi, P. (2015). Monetary policy surprises, credit costs, and economic activity. *American Economic Journal: Macroeconomics*, 7(1):44–76.

Gilchrist, S. and Zakrajšek, E. (2012). Credit spreads and business cycle fluctuations. *American Economic Review*, 102(4):1692–1720.

Jordà, Ò. (2005). Estimation and inference of impulse responses by local projections. *American Economic Review*, 95(1):161–182.

Kilian, L. and Kim, Y. J. (2011). How reliable are local projection estimators of impulse responses? *The Review of Economics and Statistics*, 94(4):1460–1466.

Kozicki, S. and Tinsley, P. (2005). What do you expect? Imperfect policy credibility and tests of the expectations hypothesis. *Journal of Monetary Economics*, 52(2):421–447.

Kurmann, A. (2007). VAR-based estimation of Euler equations with an application to New Keynesian pricing. *Journal of Economic Dynamics and Control*, 31(3):767–796.

Liu, C. and Williams, N. (2019). State-level implications of federal tax policies. *Journal of Monetary Economics*, 105:74–90.

Mertens, K. and Ravn, M. O. (2013). The dynamic effects of personal and corporate income tax changes in the United States. *American Economic Review*, 103(4):1212–47.

Miranda-Agrippino, S. and Ricco, G. (2017). The transmission of monetary policy shocks. (Sciences Po OFCE Working Paper No. 15, 2017-05-05).

Plagborg-Møller, M. and Wolf, C. K. (2020). *Local projections and VARs estimate the same impulse responses*. Unpublished manuscript, Princeton University.

Romer, C. D. and Romer, D. H. (2004). A new measure of monetary shocks: Derivation and implications. *American Economic Review*, 94(4):1055–1084.

Rudebusch, G. D. and Swanson, E. T. (2008). Examining the bond premium puzzle with a DSGE model. *Journal of Monetary Economics*, 55:S111–S126.

Rudebusch, G. D. and Swanson, E. T. (2012). The bond premium in a DSGE model with long-run real and nominal risks. *American Economic Journal: Macroeconomics*, 4(1):105–143.

Sargent, T. J. (1979). A note on maximum likelihood estimation of the rational expectations model of the term structure. *Journal of Monetary Economics*, 5(1):133–143.

Sarno, L., Thornton, D. L., and Valente, G. (2007). The empirical failure of the expectations hypothesis of the term structure of bond yields. *Journal of Financial and Quantitative Analysis*, 42(1):81–100.

Stock, J. H. and Watson, M. W. (2018). Identification and estimation of dynamic causal effects in macroeconomics using external instruments. *The Economic Journal*, 128(610):917–948.

Stock, J. H., Wright, J. H., and Yogo, M. (2002). A survey of weak instruments and weak identification in generalized method of moments. *Journal of Business & Economic Statistics*, 20(4):518–529.

Tillmann, P. (2020). Monetary policy uncertainty and the response of the yield curve to policy shocks. *Journal of Money, Credit and Banking*, 52(4):803–833.

van Binsbergen, J. H., Fernndez-Villaverde, J., Koijen, R. S., and Rubio-Ramrez, J. (2012). The term structure of interest rates in a DSGE model with recursive preferences. *Journal of Monetary Economics*, 59(7):634–648.

A Tables and figures

Table 1: Ratio of standard error estimates

h	0	1	2	3	6	12	18	24	36	48	60	all
$\frac{std.err.(\hat{\gamma}_{3,h})}{std.err.(\tilde{\gamma}_{3,h})}$	1.13	1.10	1.13	1.28	1.23	1.26	1.28	1.30	1.25	1.26	1.13	1.21

Note: Ratio of standard errors for the restricted and unrestricted estimates of $\gamma_{3,h}$. Values larger than one indicate that the standard errors from the term structure local projections are lower than standard errors from the usual local projection method. The restricted estimates, $\tilde{\gamma}_{3,h}$, are estimated under the term structure local projections, which impose the restrictions (7). The unrestricted estimates, $\hat{\gamma}_{3,h}$, are estimates of the response of the short-term interest rate. The final column shows the average ratio of standard errors at all horizons. Standard errors for both estimators are calculated using a Newey-West HAC with 12 lags.

Figure 1: Response of the term structure to a monetary policy shock

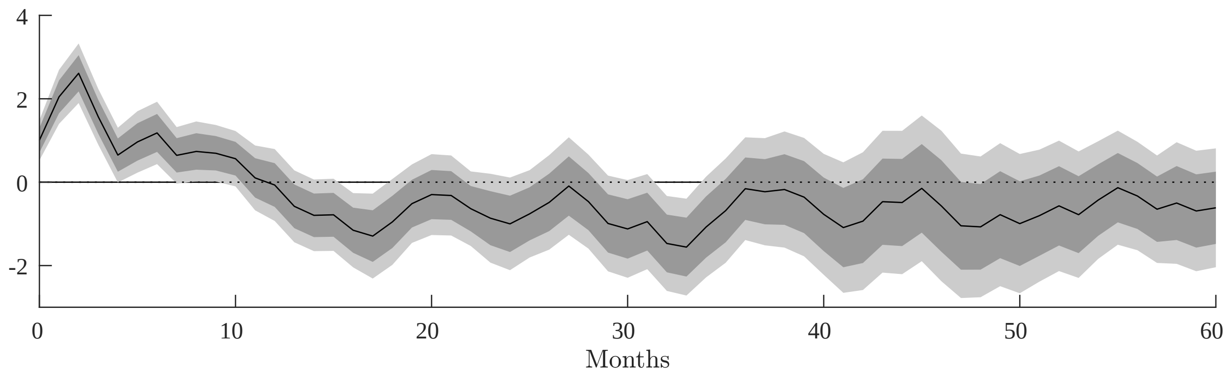

Note: Response of the term structure of interest rates to a monetary policy shock normalized to increase the one-year treasury yield by 100 basis points upon impact calculated by imposing the restrictions (7) to equations (2) and (3). The left-hand-side variables are yields on the three-month, six-month, and one-year treasury bills and the two-year, three-year, and five-year treasury notes. Control variables include a constant and twelve lags each of the log of industrial production, the log of the consumer price index, the Gilchrist and Zakrajšek (2012) excess bond premium, and the first three principal components of all interest rates. The instrumental variable is the change in three-month ahead federal funds futures in a thirty minute window around a monetary policy announcement, after adjusting for information asymmetries as in Miranda-Agrippino and Ricco (2017), which has a first-stage F-statistic of 10.93. The dark and light shaded areas are the 68% and 90% confidence intervals, calculated using a Newey-West HAC with 12 lags.

Figure 2: Response of the term premium to a monetary policy shock

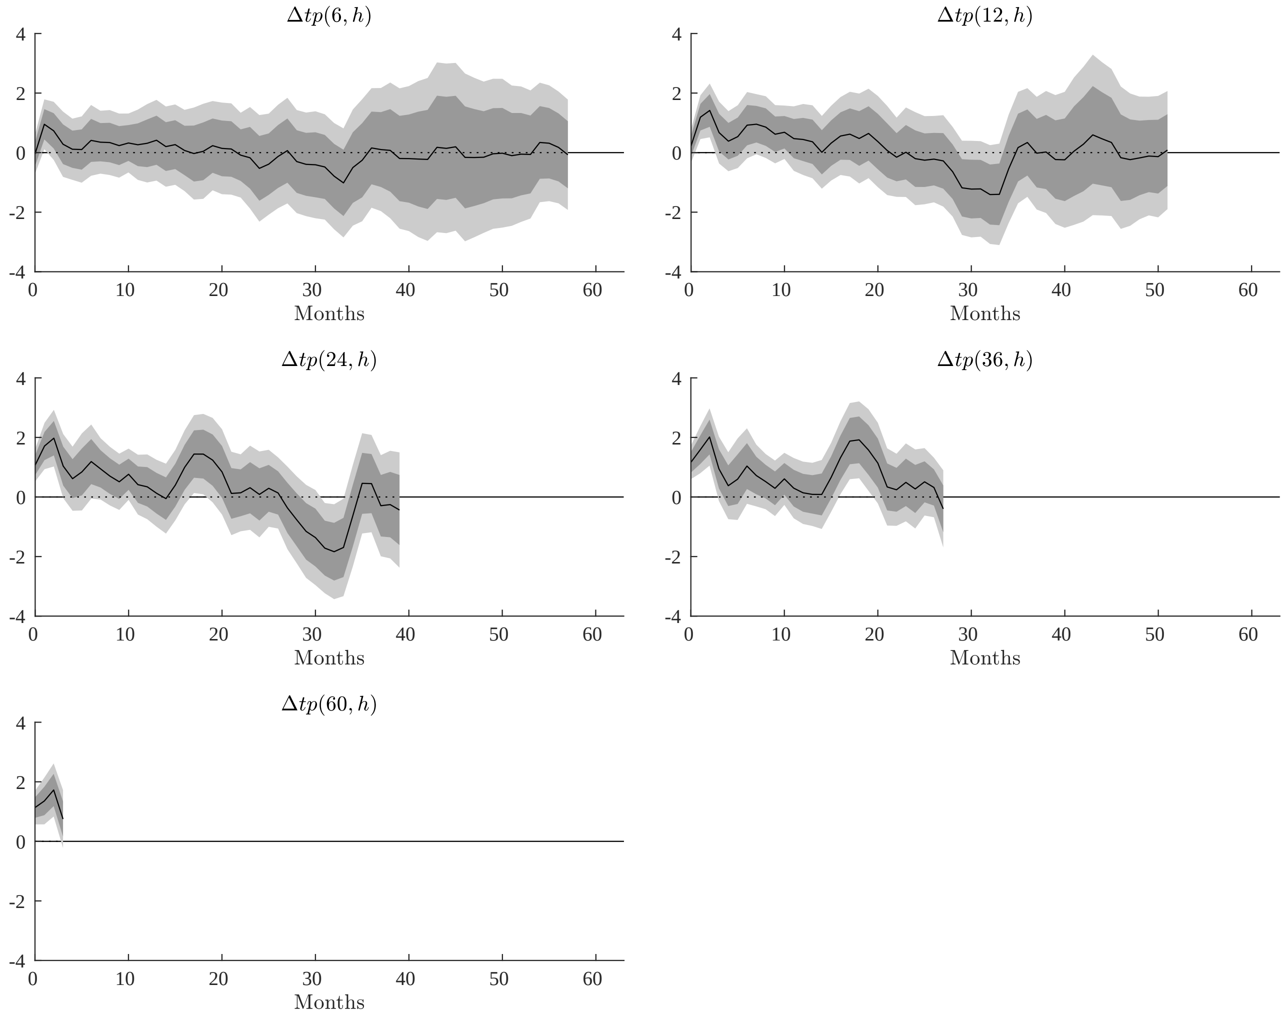

Note: Response of the term premium for a bond with a given maturity, defined by equation (8), to a monetary policy shock normalized to increase the one-year treasury yield by 100 basis points upon impact. $\Delta tp(m, h)$ denotes the response of the term premium for a bond maturing m months in the future h months after a monetary policy shock. The dark and light shaded areas are the 68% and 90% confidence intervals, calculated using a Newey-West HAC with 12 lags.

Figure 3: Response of the term premium to a total factor productivity shock

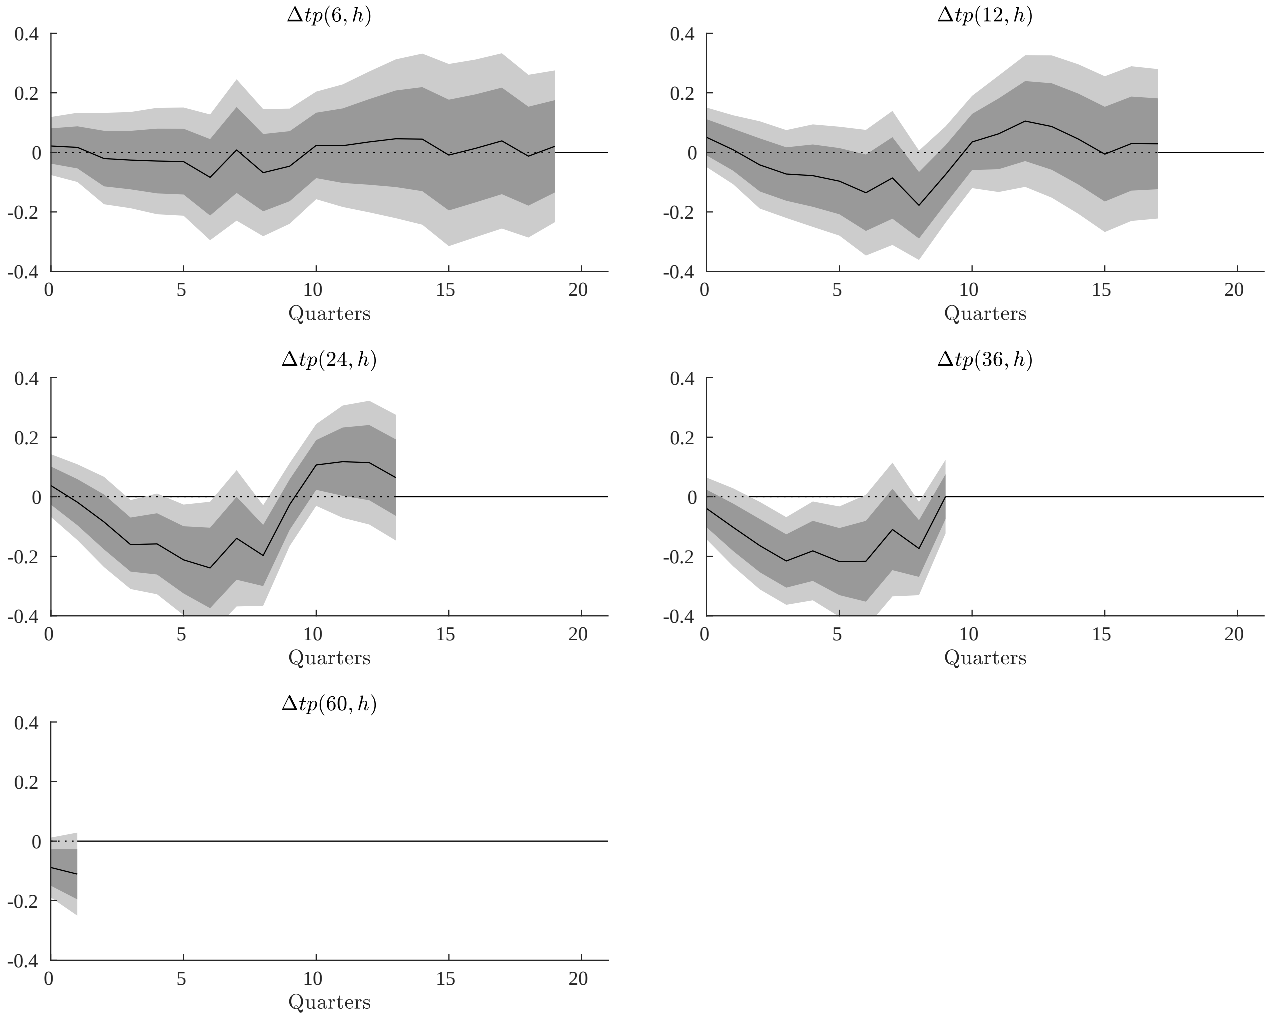

Note: Response of the term premium for a bond with a given maturity, defined by equation (8), to a total factor productivity shock normalized to increase the growth rate of GDP by one percentage point upon impact. $\Delta tp(m,h)$ denotes the response of the term premium for a bond maturing m months in the future h quarters after a total factor productivity shock. The dark and light shaded areas are the 68% and 90% confidence intervals, calculated using a Newey-West HAC with four lags.

Figure 4: Response of the term premium to a tax shock

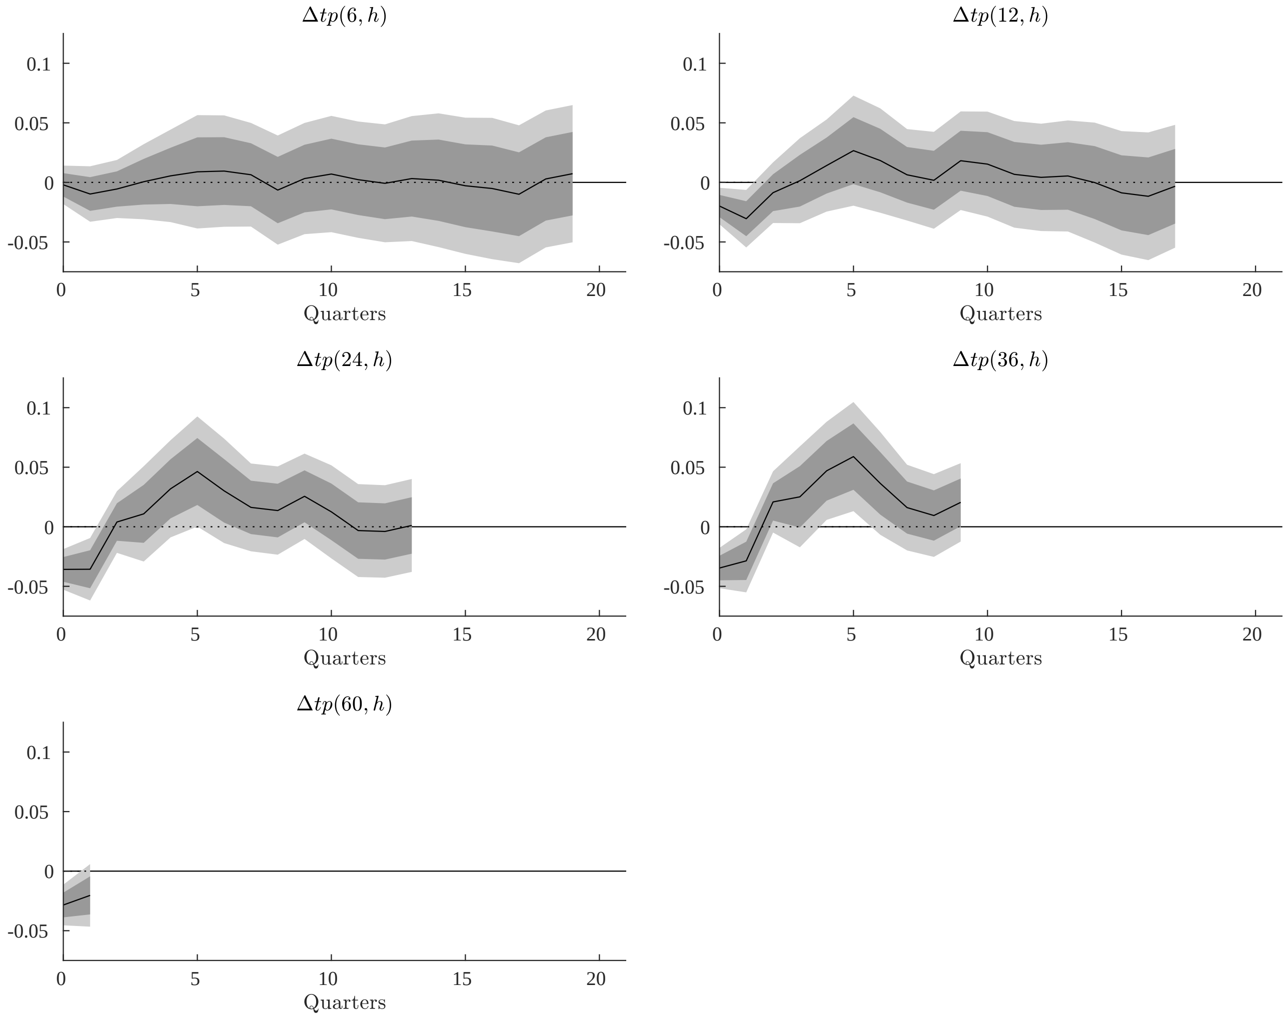

Note: Response of the term premium for a bond with a given maturity, defined by equation (8), to a tax revenue shock normalized to increase the growth rate of GDP by one percentage point upon impact. $\Delta tp(m,h)$ denotes the response of the term premium for a bond maturing m months in the future h quarters after a tax revenue shock. The dark and light shaded areas are the 68% and 90% confidence intervals, calculated using a Newey-West HAC with four lags.