EFCE
7/19/06
10:15 AM

Identifying Government Spending Shocks: It's All in the Timing

By

Valerie A. Ramey

University of California, San Diego
National Bureau of Economic Research

July 2006

Preliminary and Incomplete
Comments Welcome

I wish to thank Garey Ramey for helpful comments.

I. Introduction

How does the economy respond to a rise in government purchases? Do consumption and real wages rise or fall? The literature remains divided on this issue. VAR techniques in which identification is achieved by assuming that government spending is predetermined within the quarter typically find that a positive government spending shock raises not only GDP and hours, but also consumption and the real wage (or labor productivity) (e.g. Rotemberg and Woodford (1992), Blanchard and Perotti (2002), Fatás and Mihov (2001), Perotti (2004), Montford and Uhlig (2005), and Galí, López-Salido, and Vallés (2006)). In contrast, analyses using the Ramey-Shapiro (1998) "war dates" typically find that while government spending raises GDP and hours, it lowers consumption and the real wage (e.g. Ramey and Shapiro (1998), Edelberg, Eichenbaum, and Fisher (1999), Burnside, Eichenbaum, and Fisher (2004), and Cavallos (2005)).

Whether government spending raises or lowers consumption and the real wage is crucial for our understanding of how government spending affects GDP and hours. It is also important for distinguishing macroeconomic models. Consider first the neoclassical approach, as represented by papers such as Aiyagari, Christiano and Eichenbaum (1992) and Baxter and King (1993). A *permanent* increase in government spending financed by nondistortionary means creates a negative wealth effect for the representative household. The household optimally responds by decreasing its consumption and increasing its labor supply. Output rises as a result. The increased labor supply lowers the real wage and raises the marginal product of capital in the short run. The rise in the marginal product of capital leads to more investment and capital accumulation, which eventually brings the real wage back to its starting value. In the new steady-state, consumption is lower and hours are higher. A *temporary* increase in government

spending in the neoclassical model has less impact on output because of the smaller wealth effect. Depending on the persistence of the shock, investment can rise or fall. In the short run, hours should still rise and consumption should still fall.[1]

The new Keynesian approach seeks to explain a rise in consumption, the real wage, and productivity found in most VAR analyses. For example Rotemberg and Woodford (1992) and Devereux, Head and Lapham (1996) propose models with oligopolistic (or monopolistic) competition and increasing returns in order to explain the rise in real wages and productivity. More recently, Galí, López-Salido, and Vallés (2006) develop a model with sticky prices, noncompetitive behavior in labor markets, and "rule-of-thumb" consumers to explain how consumption and real wages can rise when government spending increases. Their paper makes clear how many special features the model must contain to explain the rise in consumption.

This paper reexamines the empirical evidence by comparing the two main empirical approaches to estimating the effects of government spending: the VAR approach and the Ramey-Shapiro narrative approach. After reviewing the set-up of both approaches and the basic results, I show that a key difference appears to be in the timing. In particular, I show that the Ramey-Shapiro dates Granger-cause the VAR shocks, but not vice versa. Thus, big increases in military spending are anticipated several quarters before they actually occur. I show this also true for several notable cases of non-defense government spending changes. I then use simple neoclassical model to demonstrate that failing to account for the anticipation effect can explain all of the differences in the empirical results of the two approaches.

Both approaches have weaknesses when applied to quarterly post-WWII data. While the VAR approach suffers from the anticipation effect, the Ramey-Shapiro narrative approach relies

[1] Adding distortionary taxes or government spending that substitutes for private consumption or capital adds additional complications. See Baxter and King (1993) and Burnside, Eichenbaum, and Fisher (2004) for discussions of these complications.

heavily on only a few episodes. After showing how small most variations in government spending are in the post-WWII period, I move on to study annual data from 1889 to the present. Moving to a long-span of annual data solves problems in both approaches. First, the annual frequency may attenuate the anticipation effect when identifying shocks in the VAR. Second, the longer time span includes more military dates. In these data I find that both methods give similar results: GDP and hours rise, but consumption falls after a government spending shock. While aggregate productivity rises, I find that manufacturing productivity falls. I used arguments from Ramey and Shapiro (1998) and Basu and Fernald (1997) to explain how these results are consistent with a multi-sector neoclassical model.

II. A Comparison of Results from the VAR Approach and the Ramey-Shapiro Narrative Approach

A. The VAR Approach

Blanchard and Perotti (2002) have perhaps the most careful and comprehensive approach to estimating fiscal shocks using VARs. To identify shocks, they first incorporate institutional information on taxes, transfers, and spending to set parameters, and then estimate the VAR. Their basic framework is as follows:

$$Y_t = A(L,q)Y_{t-1} + U_t$$

where Y_t consists of quarterly real per capita taxes, government spending, and GDP. Although the contemporaneous relationship between taxes and GDP turns out to be complicated, they find that government spending does not respond to GDP or taxes contemporaneously. Thus, their identification of government spending shocks is identical to a Choleski decomposition in which government spending is ordered before the other variables. When they augment the system to

include consumption, they find that consumption rises in response to a positive government spending shock. Galí et al (2006) use this basic identification method in their study which focuses only on government spending shocks and not taxes. They estimate a VAR with additional variables of interest, such as real wages, and order government spending first.[2]

B. The Ramey-Shapiro Narrative Approach

In contrast, Ramey and Shapiro (1998) use a narrative approach to identify shocks to government spending. Because of their concern that many shocks identified from a VAR are simply anticipated changes in government spending, they focus only on episodes where *Business Week* suddenly began to forecast large rises in defense spending induced by major political events that were unrelated to the state of the U.S. economy. The three episodes identified by Ramey and Shapiro were as follows:

Korean War

On June 25, 1950 the North Korean army launched a surprise invasion of South Korea, and on June 30, 1950 the U.S. Joint Chiefs of Staff unilaterally directed General MacArthur to commit ground, air, and naval forces. In the July 1, 1950 issue, *Business Week* wrote: "We are no longer in a peacetime economy. Even if the Communists should back down in Korea, we have had a warning of what can happen any time at all or in any of the Asiatic nations bordering on the USSR. The answer will be more money for arms." (p. 9) After early UN victories, *Business Week* in October 1950 predicted a quick end to hostilities in Korea, but a continuing defense spending increase. It predicted a somewhat faster pace of spending after China entered the war on November 9, 1950, but pointed out that it would take at least six months to translate

[2] See the references listed in the introduction to see the various permutations on this basic set-up.

defense programs into men and material. Many articles in late 1950 and in 1951 discussed Pentagon bottlenecks in awarding defense contracts. In January 1952, Washington decided to stretch the timetable of spending through 1954 instead of tapering off in 1953.

The Vietnam War

Despite the military coup that overthrew Diem on November 1, 1963, *Business Week* was still taking about defense cuts for the next year (November 2, 1963, p. 38; July 11, 1964, p. 86). After the Gulf of Tonkin incident on August 2, 1964, *Business Week* was still not forecasting increases in defense spending. However, after the February 7, 1965 attack on the U.S. Army barracks, *Business Week* began forecasting increases in defense spending (March 6, 1965, p. 41, and April 3, 1965, p. 29). (More to come on articles after 1965 about anticipations.)

The Carter-Reagan Buildup

The Soviet invasion of Afghanistan on December 24, 1979 led to a turnaround in U.S. defense policy. The event was particularly worrisome because some believed it was a possible precursor to actions against Persian Gulf oil countries. The January 21, 1980 *Business Week* (p.78) printed an article entitled "A New Cold War Economy" in which it forecasted a significant and prolonged increase in defense spending. (More to come on articles after 1980.)

These dates were based on data up through 1998. Owing to recent events, I now add the following date to these war dates:

9/11

On September 11, 2001, terrorists struck the World Trade Center and the Pentagon. On October 1, 2001, *Business Week* forecasted that the balance between private and public sectors would shift, and that spending restraints were going "out the window." In this case, though, it was clear that some of the increased spending they were discussing was not defense, but rather industry bailouts and the like. Also, as Eichenbaum and Fisher (2005) point out, the 9/11 episode was unusual in that it was not followed within a few years by an increase in labor and capital tax rates. To recall the timing of key subsequent events, the U.S. invaded Afghanistan soon after 9/11. It invaded Iraq on March 20, 2003.

Figure 1 shows the paths of log per capita real defense spending and government spending in the post-WWII era, along with lines indicating the military dates.[3] The major movements in defense spending all come following one of the four military dates. Korea is obviously the most important, but the other three are also quite noticeable. There are two minor blips in second half of the 1950s and early 1960. Both of these are probably linked to political events, such as Soviet tanks crushing the revolt in Hungary in 1956 and the tensions between the Soviet Union and the U.S. in 1960. (More to come on Business Week articles.)

Looking at the bottom graph in Figure 1, we see that total government spending shows the same distinguishable blips after the four dates. The impact of the Soviet invasion of Afghanistan has a delayed effect on total government spending, though. It is clear that other programs were cut in the years immediately after the invasion at the same time defense spending

[3] Per capita variables are created using the entire population, not just the civilian noninstitutional population ages 16 and over. Ramey and Francis (2006) present empirical and theoretical arguments for why "per capita" should encompass the entire population.

was starting to increase. Over the entire sample, total government spending differs from defense spending in that it has a significant upward trend over time.

C. Comparison of Empirical Results

Consider now a comparison of the effects of government spending increases based on the two identification methods. In both instances, I use a VAR which consists of the log real per capita quantities of total government spending, GDP, total hours worked, nondurable plus services consumption, and gross private investment, as well as the log of nominal compensation in private business divided by the deflator in private business. I use total hours worked instead of private hours worked based on Cavallo's (2005) work showing that a significant portion of rises in government spending consists of increases in the government payroll. Also, note that I use a product wage rather than a consumption wage. Ramey and Shapiro (1998) show both theoretically and empirically why it is the product wage that should be used when trying to distinguish models of government spending. Defense spending tends to be concentrated in a few industries, such as manufactured goods. Ramey and Shapiro show that the relative price of manufactured goods rise significantly during a defense buildup. Thus, product wages in the expanding industries can fall at the same time that the consumption wage is unchanged or rising.[4]

Both VARs are specified in levels, with a time trend and six lags included. Quarterly data from 1947:1 to 2006:1 are used. In the VAR identification, the government spending shock is identified by a Choleski decomposition in which government spending is ordered first. In the war dates identification, the current value and six lags of a dummy variable with the military date is also included. There is no feedback allowed from the other variables to the military date. The

[4] The main reason that Rotemberg and Woodford (1992) find that real wages increase is that they construct their real wage by dividing the wage in manufacturing by the implicit price deflator. Ramey and Shapiro show that the wage in manufacturing divided by the price index for manufacturing falls during a defense buildup.

military date takes a value of unity in 1950:3, 1965:1, 1980:1, and 2001:3.[5] I compare the effects of a unit shock to the military date to a .06 shock to the government spending equation error term, since this produces the same size rise in government spending.

Figures 2A and 2B show the impulse response functions. Standard error bands are 90% bands, based on monte carlo simulations. The first column shows the results from the VAR identification and the second column shows the results from the war dates identification. Figure 2A shows the effects on government spending, GDP, and hours. The results are qualitatively consistent across the two identification schemes for these three variables. By construction, total government spending rises by the same amount, although the peak occurs several quarters earlier in the VAR identification. GDP rises in both cases, but its rise is much greater in the case of the war dates identification. Hours rise slightly in the VAR identification, but much more strongly in the war dates identification. A comparison of the output and hours response shows that productivity rises slightly in both specifications. I will discuss later why this is likely a consequence of sectoral reallocation.

Figure 2B shows the cases in which the two identification schemes differ in their implications. The VAR identification scheme implies that government spending shocks raise consumption, lower investment for two years, and raise the real wage. In contrast, the war dates identification scheme implies that government spending shocks lower consumption, raise investment for a few quarters before lowering it, and lower the real wage.

One might wonder why there is so much controversy, though, because the standard error bands are rather large. More parsimonious specifications with fewer variables or fewer lags or ones that use restricted sample periods usually yield point estimates that are similar, but that are

[5] Burnside, Eichenbaum and Fisher (2004) allow the value of the dummy variable to differ across episodes according to the amount that government spending increase. They obtain very similar results.

estimated more precisely. For example, Ramey and Shapiro (1998), who use a bivariate system with current and lagged values of the war dates and the variable of interest, find significant negative effects on nondurable plus services consumption, durable consumption expenditures, residential investment and real product wages. Using the VAR identification framework, Perotti (2004) and Galí et al (2006) find significant positive effects on consumption.

Overall, these two approaches give diametrically opposed answers with regard to some key variables. The next section presents empirical evidence and a theoretical argument that can explain the differences.

III. The Importance of Timing

A concern with the VAR identification scheme is that some of what it classifies as "shocks" to government spending may well be anticipated. Indeed, Blanchard and Perotti (2002) worried about this very issue, and devoted Section VIII of their paper to analyzing it. To test for the problem of anticipated policy, they included future values of the estimated shocks to determine whether they affected the results. They found that the response of output was greater once they allowed for anticipation effects (see Figure VII). Unfortunately, they did not show how the responses of consumption or real wages were affected. Perotti (2004) approached the anticipation problem by testing whether OECD forecasts of government spending predicted his estimated government spending shocks. For the most part, he found that they did not predict the shocks.

In the next section, I show that the war dates do predict the VAR government spending shocks. I also show how in each war episode, the VAR shocks are positive several quarters after *Business Week* started forecasting increases in defense spending. In the following section, I

present a simple theoretical model that show how the difference in timing can explain most of the differences in results.

A. Empirical Evidence on Timing Lags

To compare the timing of war dates versus VAR-identified shocks, I estimate shocks using the VAR discussed above except with defense spending rather than total government spending as the key variable. I then plot those shocks around the war dates.

Figure 3 shows the path of log per capita real defense spending and the series of identified shocks. Consider first the Korean War. The first vertical line shows the date when the Korea War started. The second vertical line indicates when the armistice was signed in July 1953. According to the VAR estimates, there were a series of three large positive shocks to defense in 1950:4, 1950:1, and 1950:2. However, as *Business Week* made clear, the path of defense spending during these three quarters was anticipated as of July 1950. Moreover, the negative shock to defense estimated in the first quarter of 1952 corresponds to the decision to spread out defense spending over a longer period. It is also interesting to note that while *Business Week* was predicting a future decline in defense spending as early as April 1953 when a truce seemed imminent, the VAR records a negative defense spending shock late in the first quarter of 1954.

Vietnam shows a similar pattern in that the war date predates a series of positive shocks from the VAR. (More to come on Business Week forecasts after initial shocks for these later episodes.) After 9/11 the VAR implies a large negative shock in the first quarter of 2003 and a large positive shock in the second quarter of 2003. Most people would have been anticipating an increase in military spending at least by the first quarter of 2003.

Thus, it appears that the VAR might be labeling as "shocks" changes in defense spending that were forecastable. To test this hypothesis, I ran Granger causality tests between the war dates and VAR-based defense and total government spending shocks. Table 1 shows the results. The evidence is very clear: the war dates Granger-cause the VAR shocks but the VAR shocks do not Granger-cause the war dates.

One should be clear that this is not an issue only with defense spending. Consider the interstate highway program. In early 1956, *Business Week* was predicting that the "fight over highway building will be drawn out." By May 5, 1956, *Business Week* thought that the highway construction bill was a sure bet. It fact it passed in June 1956. However, the multi-billion dollar program was intended to stretch out over 13 years. It is difficult to see how a VAR could accurately reflect this program. Another example is schools for the Baby Boom children. Obviously, the demand for schools is known several years in advance. Between 1949 and 1969, real per capita spending on public elementary and secondary education increased 300%.[6] Thus, a significant portion of non-defense spending is known months, if not years, in advance.

B. The Importance of Timing in a Theoretical Model

To see how important anticipation effects can be, consider the following calibrated neoclassical model:

[6] The nominal figures on expenditures are from the *Digest of Education Statistics*. I used the GDP deflator to convert to real.

$$Y_t = N_t^{0.67} K_t^{0.33}$$

$$U = \log(C_t) + \log(500 - N_t)$$

$$Y_t = C_t + I_t + G_t$$

$$K_{t+1} = I_t + (1 - 0.025)K_t$$

Y is output, N is labor, K is capital, C is consumption, I is investment, and G is government purchases. Government purchases are financed with nondistortionary taxes. Households maximize the present discounted value of utility U with discount factor $\beta = 0.984$.

Consider an increase in government spending similar to that which occurred during the Korea War. Assume that the change in government spending is announced one period in advance and that the entire path is perfectly foreseen.

Figure 4 shows how important anticipation is. Although government spending does not begin to rise until the second quarter, output, hours, and investment immediately jump up in the first quarter. Furthermore, consumption and the real wage both fall in the first quarter before beginning to rise slowly in the second quarter.

Now consider Figure 5 which overlays these responses on ones that we would infer if we were off one period on the timing. In particular, suppose that a VAR identifies the government spending shock as occurring in quarter 2 rather than in quarter 1, since government spending doesn't start rising until then. The red dotted lines show the kinds of impulse responses we would obtain if our timing were off by a quarter: they are just the blue solid lines led a quarter. According to the delayed impulse responses, an increase in government spending leads consumption and the real wage to rise and investment to fall. The reason that they give these

effects is that they missed the impact effect of the announcement. Thus, being one quarter off changes the signs of the implied impulse responses.

Note how similar the theoretical patterns in Figure 5 are to the empirical patterns in Figure 2. The delayed theoretical responses give much smaller responses of output and hours than the true responses. Similarly, the VAR shocks give much smaller responses of output and hours than the military dates. The delayed theoretical responses indicate that consumption and real wages rise, while the true responses indicate that consumption and the real wage fall. Similarly, the VAR shocks imply that consumption and real wages rise, whereas the military dates indicate that consumption and real wages fall. Thus, the theoretical model can explain all of the differences across identification methods as a difference in timing caused by announcement effects.

One could modify the theoretical model to incorporate elements such as habit persistence in consumption and/or adjustment cost in investment. In this case, the true responses would be more dragged out and missing the timing by one quarter would have a somewhat smaller effect. However, missing the timing but a few quarters would have change the impulse responses just as they did in the previous example.

IV. New Annual Historical Evidence

We have seen that VAR methods for identifying government spending shocks may suffer from anticipation effects. Increases in government spending are typically known at least a quarter before they occur, and this can potentially have a significant effect on the estimated impulse response functions. Unfortunately, many of the estimated quarterly shocks to

government spending are probably blips in the timing of programs such as defense contracts, school construction, or highway construction. It is just not clear how much of the estimated shocks are truly shocks and how much is just noise.

On the other hand, the military dates appear to be much better in terms of timing. However, their paucity leads to imprecise estimates and the possibility that other major events around the episodes could account for the estimates.

Given the short-comings of both methods, we might learn more by studying the effects of government spending in annual historical data. Annual historical data has several advantages. First, moving from a quarterly to an annual frequency might mitigate some of the timing issues in the VAR methodology. If most changes in government spending are anticipated by only a few quarter or two, then VARs might be able to identify true unanticipated shocks better in annual data. Unfortunately, moving to the annual frequency and using post-WWII data results in too small a sample size. For this reason, I extend the sample back as far as possible. Extending the sample back has the added advantage that it includes a period with much bigger government spending shocks. Thus, we might obtain more precise estimates from these data. Recently, McGrattan and Ohanian (2006) have studied the effects of spending government spending during World War II.

Figure 6 shows the path of real per capita government spending from 1889 to 2004. The lines indicate military dates, including the Spanish-American War (dated as starting in1898), World War I (starting in 1917 for the U.S.), World War II (starting in 1941 for the U.S., although 1942 would be equally valid), and the four war dates of the post-WWII period. The biggest surges in government spending clearly occur in the early sample. While the Korean War dwarfs the other military build-ups in the post-WWII period, Figure 6 shows that the Korean War is

dwarfed by the two world wars. Furthermore, annual government spending seems to be more volatile in the early period even in peace time. Thus, these historical data may have more useful information, both with respect to military dates and the VAR methodology.

To analyze the effects of military spending for 1889 to 2004, I estimate a VAR with government spending, GDP, nondurable plus services consumption, gross private investment, total hours, and manufacturing labor productivity. Since there is no consistent real wage series extending back to 1889, I focus on labor productivity instead. The path of total productivity is implied from the impulse responses of GDP and total hours. I include manufacturing labor productivity for reasons I will discuss below. Three lags of each variable is included.

Figure 7 shows the estimated shocks to government spending based on the VAR. In annual data, the VAR methodology seems better at picking up the onset and resolution of wars. Also, note how much bigger the shocks are even in peacetime during the early period. The 1920s and 1930s had a number of big government spending shocks.

Figure 8 shows the impulse response function from shocks identified using the standard VAR. The results imply that an increase in government spending raises GDP and hours, but lowers consumption and investment. These results are perfectly consistent with the neoclassical model. While investment should increase in response to a permanent increase in government spending, it will decrease if the increase in government spending is temporary and not too persistent. Thus, the investment response is also consistent with the neoclassical model.

The only result at odds with the neoclassical model is the implied path of total productivity. Comparing the responses of GDP and hours, it is clear that total productivity rises in response to the increase in government spending. However, this effect could be due to differing sectoral responses. As Ramey and Shapiro (1998) document for the post-WWII period,

increases in government spending during war time tend to be directed toward manufactured goods, and durable manufacturing in particular. Basu and Fernald (1997) show that manufacturing industries, and durables in particular, have larger returns to scale than other industries. Thus, a sectoral shift in the composition of output to higher returns to scale industries appears as an increase in productivity in the aggregate, even though each industry has diminishing returns to scale. To test this, the last graph in Figure 8 shows labor productivity in manufacturing. Consistent with the neoclassical explanation, this variable falls in response to a government spending shock.

One might worry that rationing during World War II is accounting for the decline in consumption. When I re-estimate the model omitting 1940 to 1946, I find similar results to those of the full sample. In particular, consumption falls significantly.

Thus, in annual historical data government spending shocks as identified by a VAR give results that are consistent with the neoclassical model. It appears that temporal aggregation helps mitigate the anticipation effects.

How does the narrative approach do in the annual historical data? Timing of military dates is more difficult in annual data. Also, *Business Week* did not exist before 1929. As a first pass, for military dates, I use 1898 (the U.S. declared war on Spain in April after the sinking of the USS Maine in Havana Harbor in February), 1917 (the U.S. declared war on Germany in April and on Austria and Hungary in December), 1941 (the U.S. was supplying Britain during the year and Pearl Harbor occurred in December), and 1950, 1965, 1980 and 2001. I estimate the VAR with current and three lags of the military dates included and followed the same procedure as with the post-WWII data.

Figure 9 shows the results using the war date identification. Interestingly, the results are similar to those from the VAR procedure, although the standard error bands are larger. Thus, both procedures give qualitatively similar results in the annual historical data. It appears that moving to the lower frequency makes both methods consistent.

V. Conclusions

This paper has reexamined the post-WWII empirical evidence on the effects of government spending and has provided new evidence based on annual data back to 1889. I have shown that the main reason for the differing results across the VAR and narrative approaches is timing and the consequences of anticipation effects. I show that the Ramey-Shapiro war dates Granger-cause the VAR shocks, but not vice versa. I then show using a simple theoretical model how getting omitting anticipation effects can switch the sign of key impulse responses.

In the final part of the paper, I turn to historical annual data in the hopes that lower frequency data mitigates the problem of anticipation. It appears to, as both the VAR approach and the war date approach give the same results. In the historical annual data, an increase in government spending raises GDP and hours, but lowers consumption and investment. While aggregate productivity rises, manufacturing productivity falls. I used arguments from Ramey and Shapiro (1998) and Basu and Fernald (1997) to explain how these results are consistent with a multi-sector neoclassical model.

Table 1. Granger Causality Tests

Hypothesis Tests	p-value
Defense shocks	
Do War dates Granger cause VAR shocks?	0.002
Do VAR shocks Granger cause War dates?	0.316
Government spending shocks	
Do War dates Granger cause VAR shocks?	0.004
Do VAR shocks Granger cause War dates?	0.174

VAR shocks were estimated by regressing the log of the variable of interest on 6 lags of itself, log real per capita GDP, log real per capita nondurable plus services consumption, log real per capita gross private investment, log real per capita total hours worked, and log compensation in private business divided by the deflator for private business.

The test regressed the second variable on 6 lags of the first variable.

Figure 1

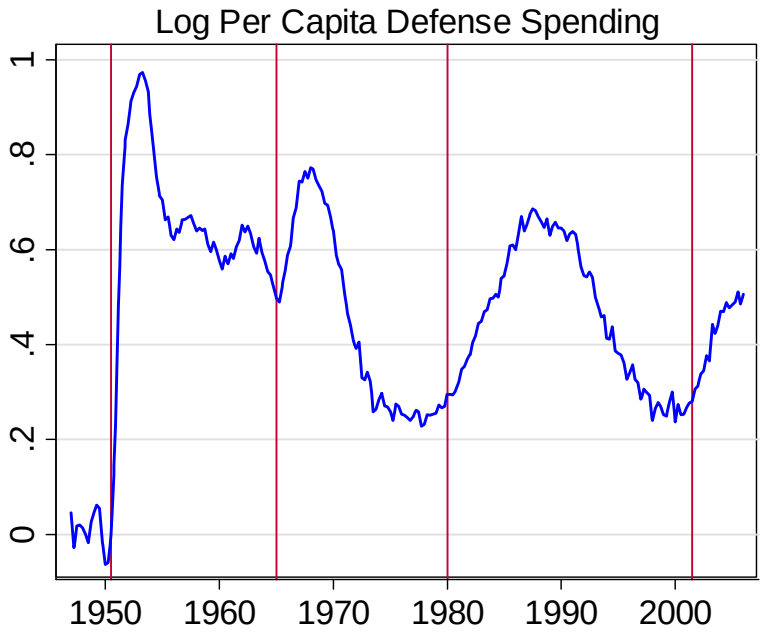
Log Per Capita Defense Spending
0
.2
.4
.6
.8
1
1950
1960
1970
1980
1990
2000

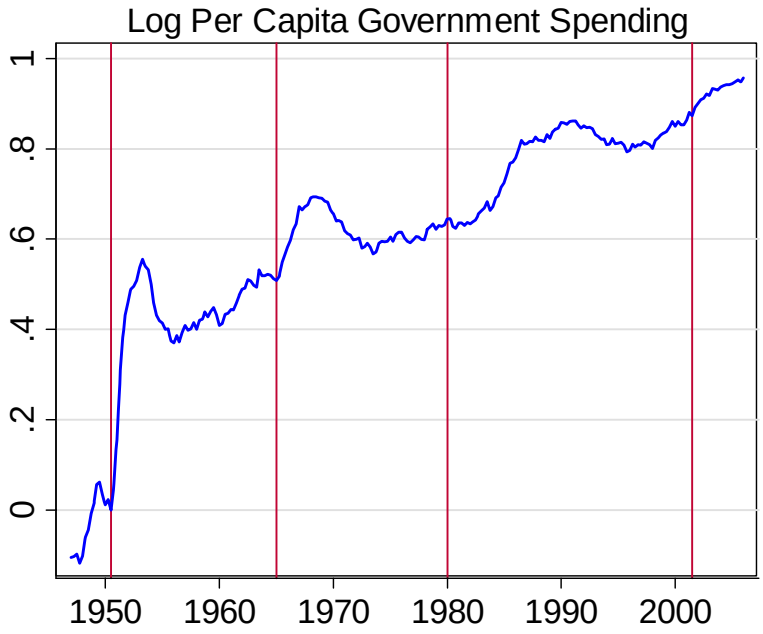
Log Per Capita Government Spending
0
.2
.4
.6
.8
1
1950
1960
1970
1980
1990
2000

Figure 2A. Comparison of Identification Methods
Response to a government spending shock
(Standard error bands are 90% confidence intervals)

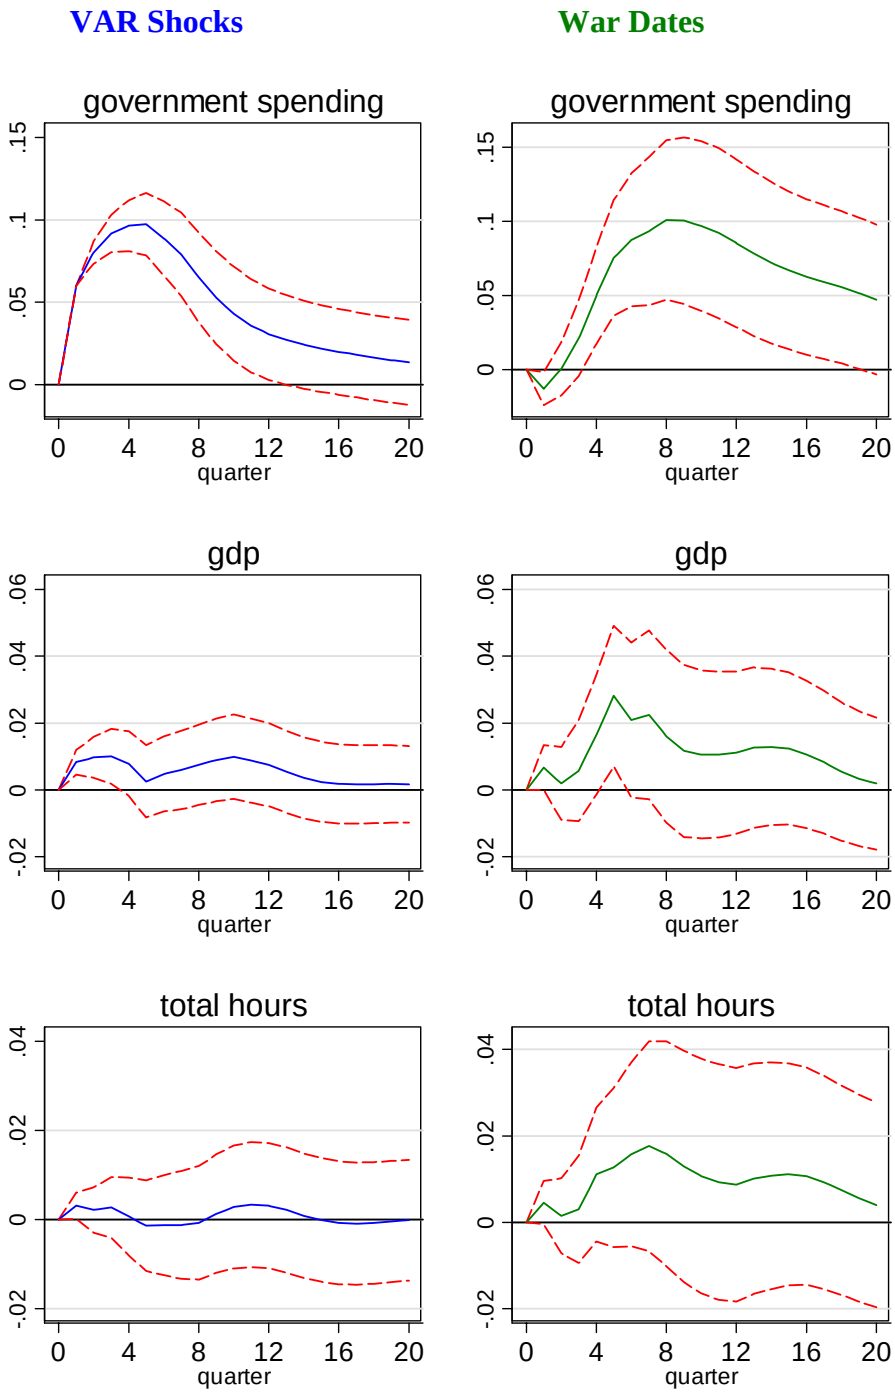

Figure 2B. Comparison of Identification Methods (continued)
Response to a government spending shock
(Standard error bands are 90% confidence intervals)

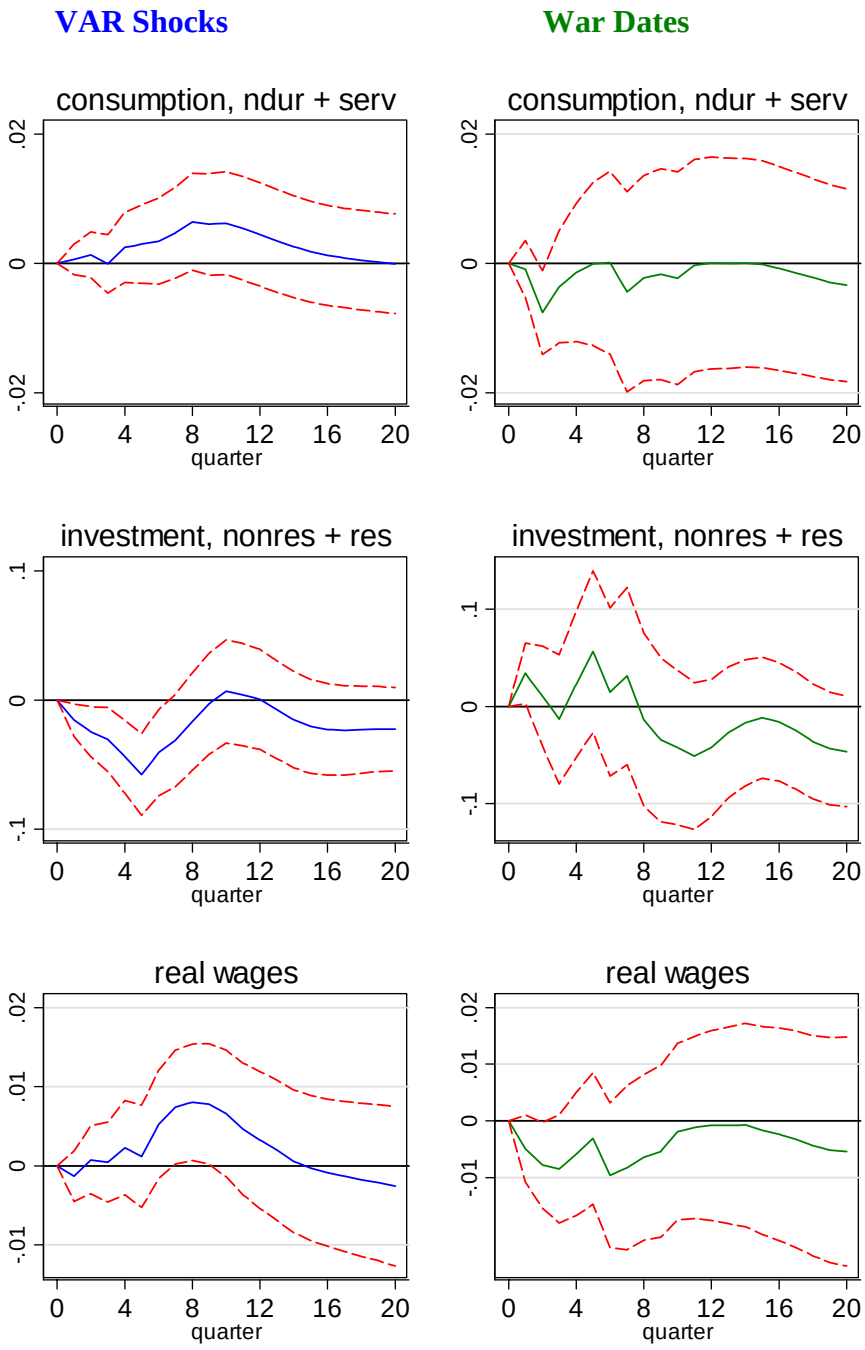

Figure 3: Episodes of Defense Spending and VAR Defense Shocks

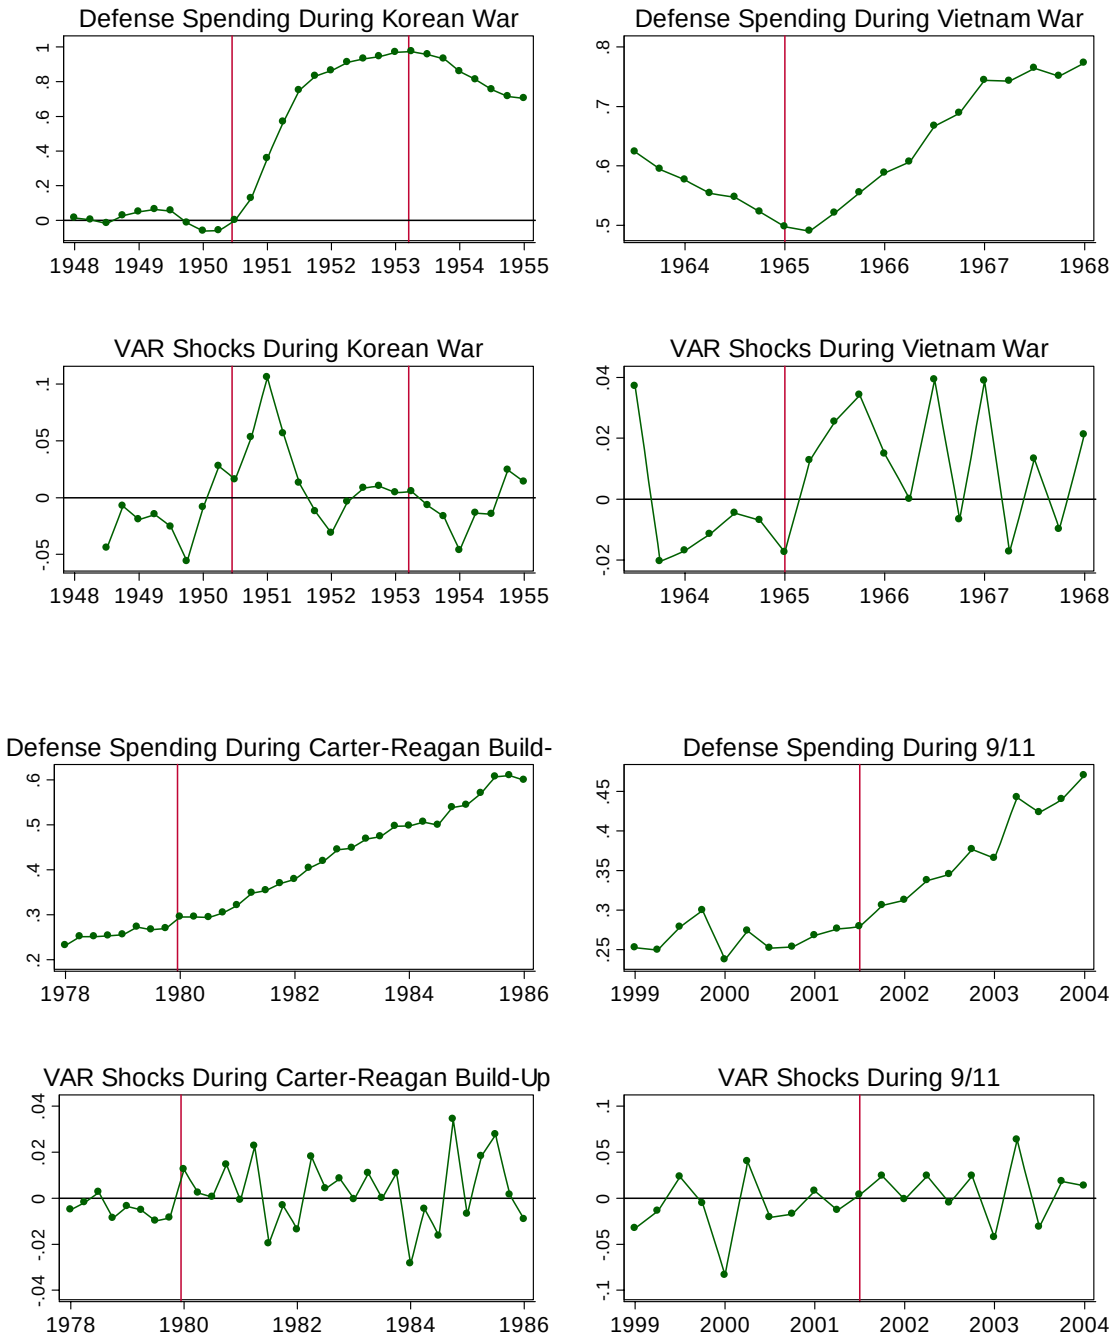

Figure 4. The Effect of an Increase in Government Spending

Announced One Quarter in Advance

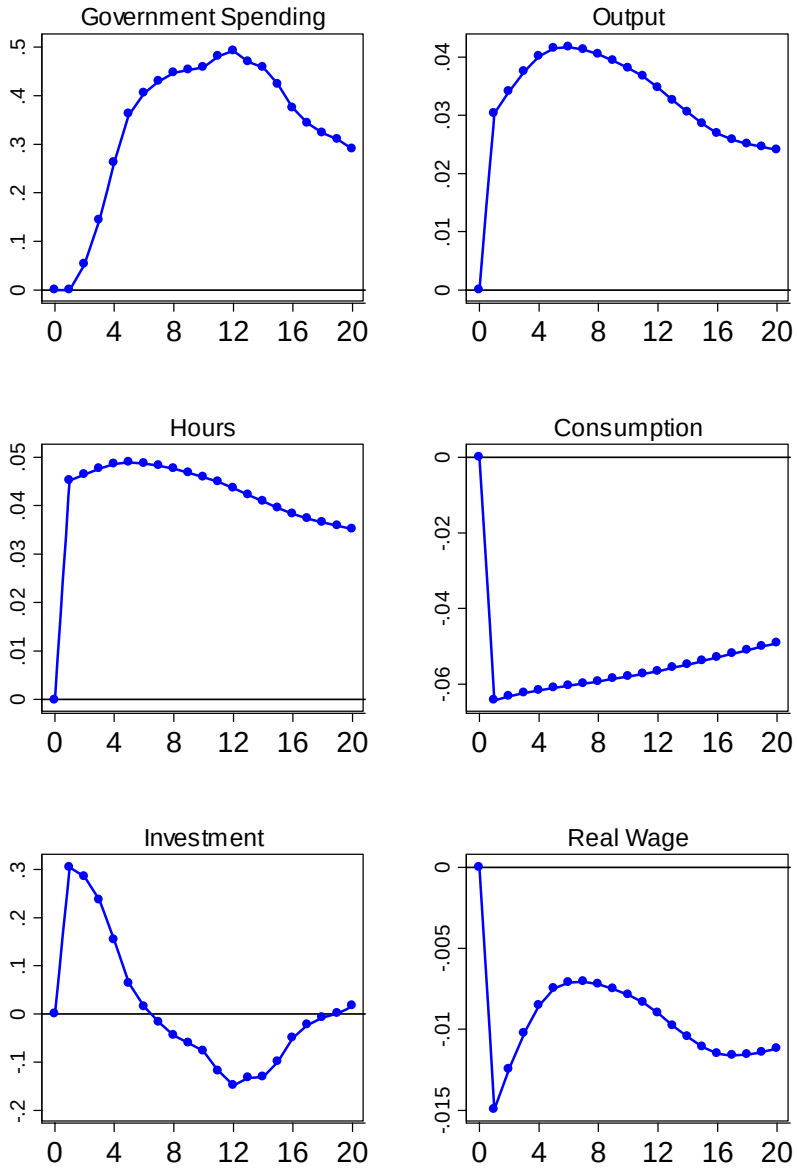

Figure 5. The Effect of Missing the Timing

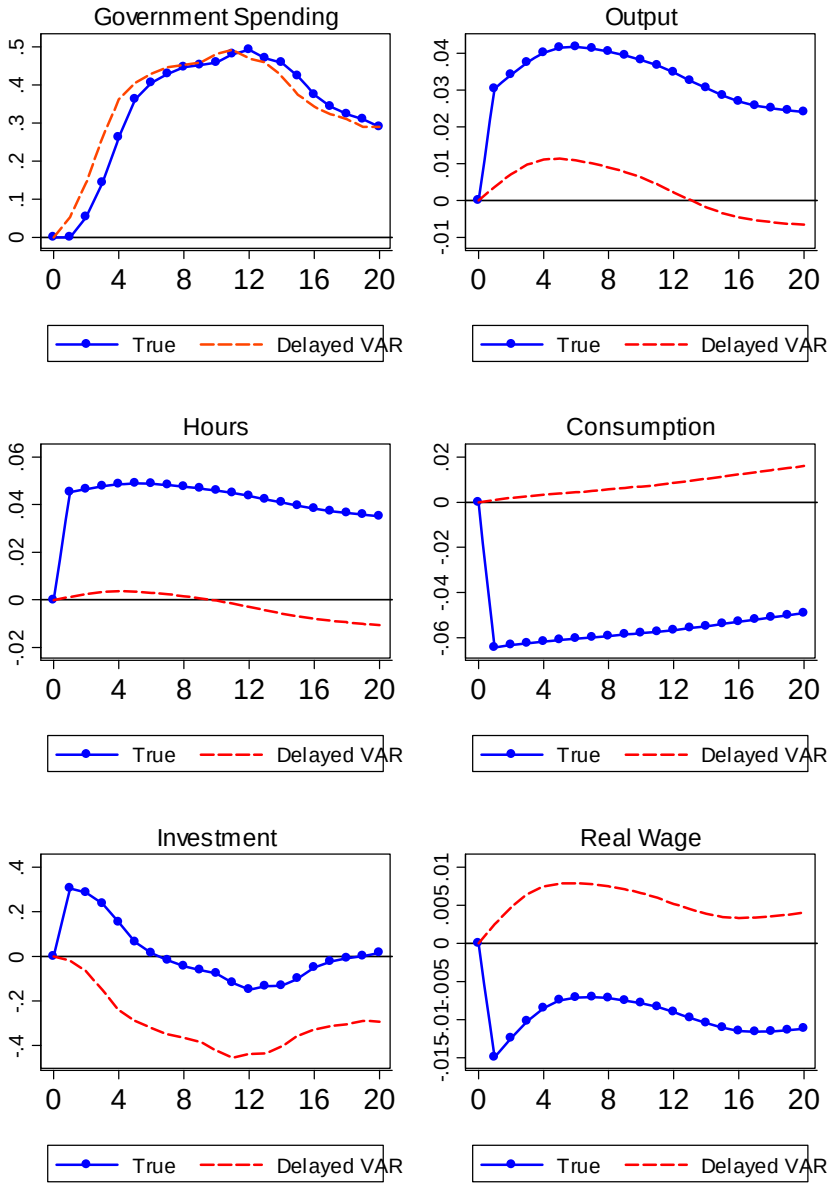

Figure 6

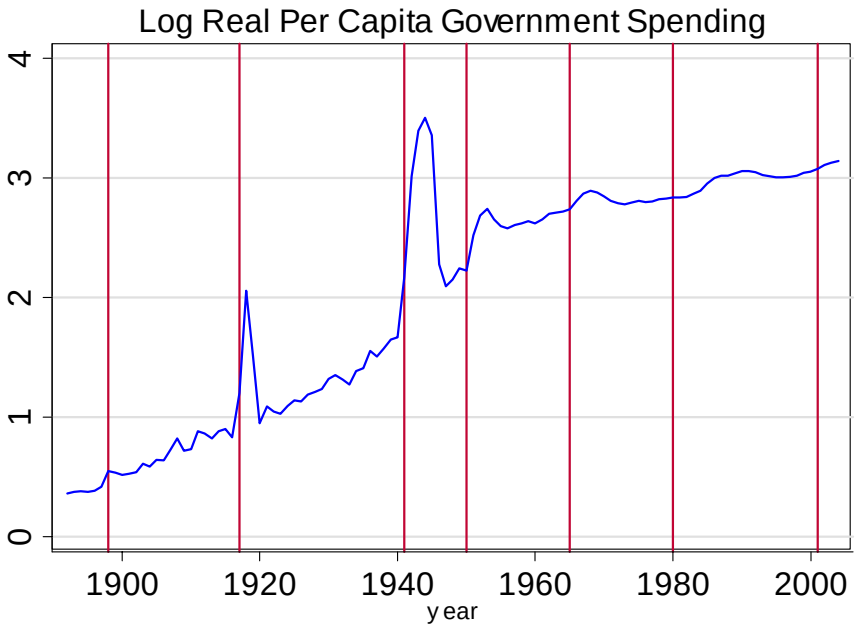
Log Real Per Capita Government Spending
0
1
2
3
4
1900
1920
1940
1960
1980
2000
year

Figure 7. Government Spending Shocks as Identified by the VAR

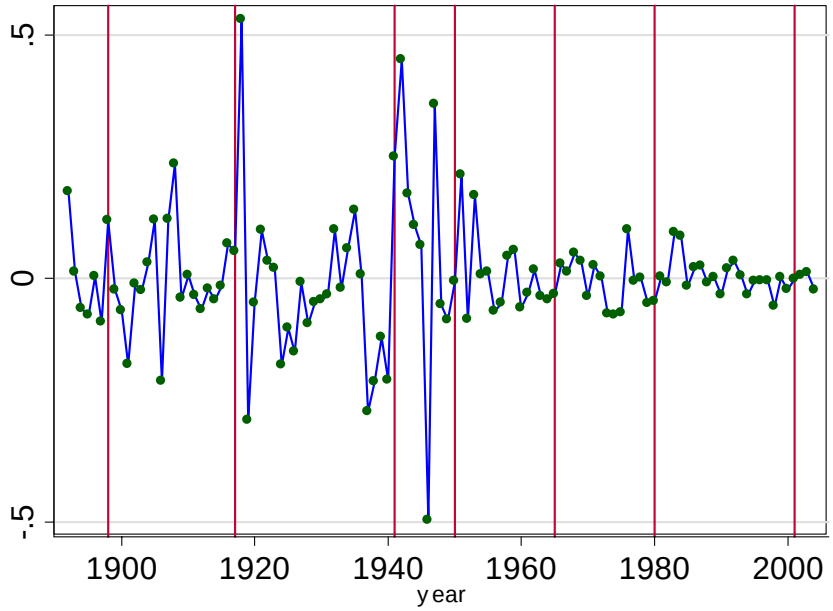

Figure 8. The Effect of a Government Spending Shock
VAR Identification, Annual, 1889-2004
(Standard error bands are 90% confidence intervals)

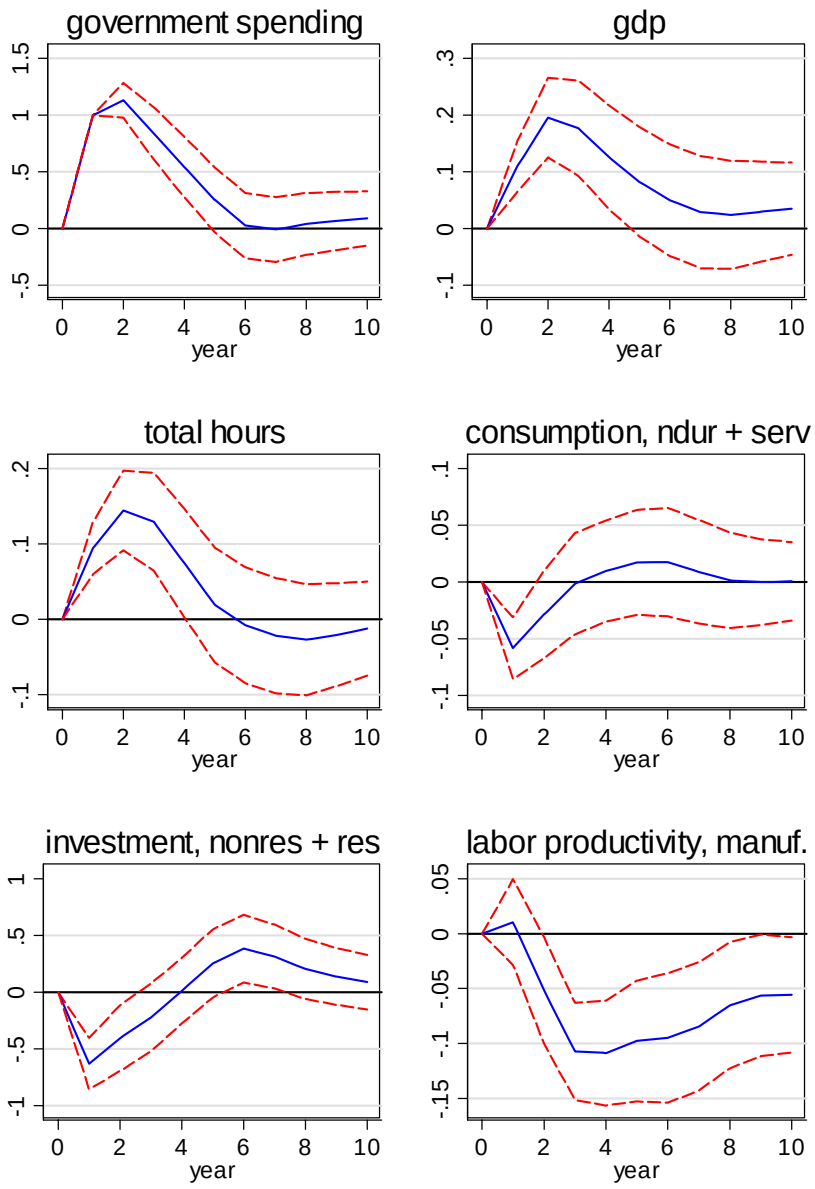

Figure 9. The Effect of a Government Spending Shock
War Dates Identification, Annual, 1889-2004
(Standard error bands are 90% confidence intervals)