Federal Reserve Bank of New York
Staff Reports

A Bayesian Approach to Estimating Tax and Spending Multipliers

Matthew Denes
Gauti B. Eggertsson

Staff Report no. 403
November 2009

This paper presents preliminary findings and is being distributed to economists and other interested readers solely to stimulate discussion and elicit comments. The views expressed in the paper are those of the authors and are not necessarily reflective of views at the Federal Reserve Bank of New York or the Federal Reserve System. Any errors or omissions are the responsibility of the authors.

A Bayesian Approach to Estimating Tax and Spending Multipliers
Matthew Denes and Gauti B. Eggertsson
Federal Reserve Bank of New York Staff Reports, no. 403
November 2009
JEL classification: E52

Abstract

This paper outlines a simple Bayesian methodology for estimating tax and spending multipliers in a dynamic stochastic general equilibrium (DSGE) model. After forming priors about the parameters of the model and the relevant shock, we used the model to exactly match only one data point: the trough of the Great Depression, that is, an output collapse of 30 percent, deflation of 10 percent, and a zero short-term nominal interest rate. Because we form our priors as distributions, the key economic inference of our analysis—the multipliers of tax and spending—are well-defined probability distributions derived from the posterior of the model. While the Bayesian methods used are standard, the application is slightly unusual. We conjecture that this methodology can be applied in several different settings with severe data limitations and where more informal calibrations have been the norm. The main advantage over usual calibration exercises is that the posterior of the model offers an interesting way to think about sensitivity analysis and gives researchers a useful way to describe model-based inference. We apply our simple estimation method to the American Recovery and Reinvestment Act (ARRA), passed by Congress as part of the 2009 stimulus plan. The mean of our estimate indicates that ARRA increased output by 3.6 percent in 2009 and 2010. The standard deviation of this estimate is 1 percent.

Key words: tax and spending multipliers, zero interest rates, deflation

---

Denes: Wharton, University of Pennsylvania (e-mail: matthew.denes@gmail.com). Eggertsson: Federal Reserve Bank of New York (e-mail: gauti.eggertsson@ny.frb.org). We thank Vasco Curdia and Christopher Sims for guidance and comments. The views expressed in this paper are those of the authors and do not necessarily reflect the position of the Federal Reserve Bank of New York or the Federal Reserve System.

1 Introduction

What is the effect of cutting taxes or increasing government spending on aggregate output? This is a key question in macroeconomics that is of central importance in policy making. The crisis of 2008 made this question urgent, because the economic stimulus passed in January 2009 was directly aimed at increasing output and employment. The numbers involved were staggering. The stimulus bill passed by Congress (the American Recovery and Reinvestment Plan) was about 775 billions dollars, over two years, or close to 3 percent of GDP. Meanwhile, until quite recently, there has been a bit less research than one might expect on the stimulatory effect of government spending (with many important exceptions, some of which that are cited in the paper), at least if compared to the research on the effect on output and prices of cutting the short-term nominal interest rate. To a large extent this has been driven by the fact that up until now, most economists have assumed that cutting nominal interest rates should be enough to increase output and unemployment at business cycle frequencies. With the Fed Funds rate at zero since December 2008, it was only natural to look to fiscal policy.

This paper proposes a Bayesian method to estimate fiscal multipliers that answer the question, by how much does output increase for every dollar in government spending or tax cuts? In recent years Bayesian methods have become increasingly popular to estimate dynamic stochastic equilibrium (DSGE) models, see e.g. An and Schorfeide (2007) for an overview. The advantage of Bayesian methods is that it allows the researcher to formally incorporate outside information into the estimation. Typically these methods have been used to estimate models with decades of data. In this application, however, there is limited direct data so the application is somewhat unusual.

The key feature of the policy environment in 2008 was that nominal interest rates collapsed to zero. Given a reasonable specification of monetary policy, Eggertsson (2009a) shows that the effect of fiscal policy is fundamentally different in this environment than when interest rates are positive. This makes any estimate based upon past data highly suspect, since the only period in which we have observed zero interest rates in US economic history is the Great Depression. This aspect of the current environment, however, makes Bayesian methods particularly suitable because it allows us to formally incorporate outside information into the estimation framework in a way we hope is transparent. We are able, therefore, to make statistical inference even if we have very little data at zero interest rates. Obviously, then, our predictions are heavily reliant on the specific structure of the model and our conjecture of how the crisis should be interpreted in the context of our model. But conditional on the strong assumption that the model is correct (and our theory of the crisis of 2008), we can still formally incorporate the uncertainty about the parameters and the underlying shock, which is what gives rise to the uncertainty of our estimates.

The approach we propose here is to rely heavily on the priors – i.e. outside information – to parameterize the structural parameters of the model. One can think of the parameterization of the structural parameters as stemming from a standard Bayesian estimation of post war data outside of the zero bound. This is the first step. The second step is to parameterize a shock

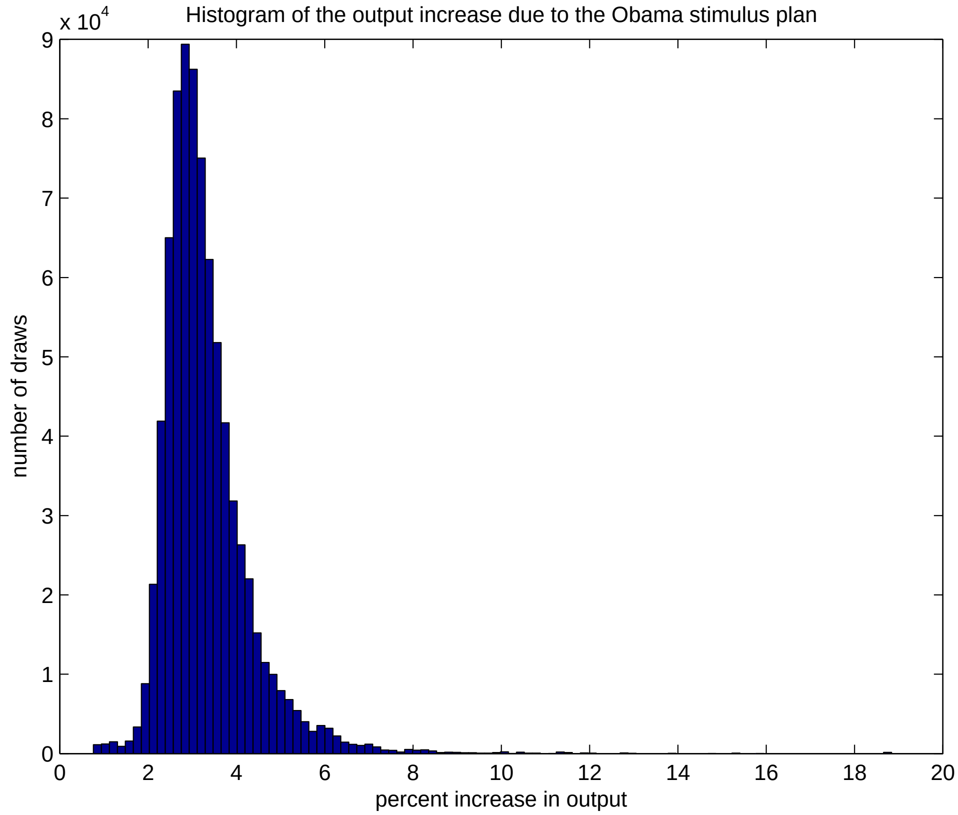

Figure 1: The percentage increase in output that results from the Obama fiscal stimulus plan according to our Bayesian estimate.

process which gives rise to the "crisis of 2008" in our model. Our prior of the shock is that it is "big" – in a way that we make precise. The third step is to ask the model to match a disaster scenario: The trough of the Great Depression. The reason why we choose this benchmark, is that we believe that policymakers conducted policy in order to prevent a Great Depression scenario (or in any event an insurance against the risk of a Great Depression scenario). It is also a natural to use as benchmark the only episode in US economic history in which nominal interest rates collapsed to zero in the US.

Using our approach we compute tax and spending multipliers at zero interest rate. The mode for the government spending multiplier – the statistic that has generated most recent discussion – is 2.28. This means that one should expect that a one percent increase in government spending (as a fraction of GDP) should increase output by 2.28 percent. The 5 and 95 percentiles of this estimate are 1.4 and 3.2, respectively. Because there has been so much discussion of this particular statistic recently, we focus our discussion on it, although we also report three tax multipliers.

Bernstein and Romer (2009) summarized a model simulation in preparation of "the American Recovery and Reinvestment Plan," that was ultimately passed by Congress, presumably informed by this estimate. They found that a 775 billion dollar stimulus increases output in 2009 and 2010 by 3.7 percent. However, they only provided a point estimate. Figure 1 shows the probability distribution of the percentage output increase in response to a 775 billion dollar stimulus plan according to our Bayesian estimation of the multipliers. The estimate assumes that 2/3 of the stimulus is implemented by temporarily increasing government spending, while the rest is used to temporarily cut labor tax rates. The mean of the estimate is 3.3 with a standard deviation of 1. This is – at least to us – surprisingly similar numbers to those found by Bernstein and Romer. The calculation underlying the estimate, however, is quite different. Most importantly – in our case – the labor tax rate cut is *contractionary* and reduces the effectiveness of the plan (the mean of output growth is 3.9 in the absence of labor tax cuts). The reason for this result is quite specific to our model (and we do not take the numbers literally) although it reflects a relatively general principle of DSGE models at zero interest rates, i.e., that policy should be aimed at increasing aggregate *spending incentives* rather than *aggregate supply incentives* (this point is explained in detail in Eggertsson (2009a) and briefly reviewed in section 4). The paper also documents an estimate of a tax cut that is quite effective in increasing output, i.e. a temporary reduction in sales taxes.

The numerical example shown figure 1 is not meant to be a comprehensive estimate of the stimulus plan passed by Congress in 2009. Instead, we report it to show the kind of results the methodology can deliver. The main objective of this paper is to show how Bayesian methods can be used to answer the type of questions posed by Bernstein and Romer (2009). As this is our objective, we have chosen to keep our model as simple as possible, which allows us to illustrate the method with almost exclusively closed form solutions. We are hopeful that this strategy will help future researchers in formulating more detailed models. It is our conjecture that more complex DSGE models would deliver results that are quantitatively similar to the ones we report here, altough this – of course – remains to be seen.

The first paper on the effect of government spending at zero interest rates in a New Keynesian DSGE model is Eggertsson (2001).The main point of that paper was analytical: Government spending is a natural way of solving the problem of the zero bound, and the paper shows that this is particularly true if the government cannot commit to future policy. The paper also analyses the Ramsey solution, i.e., optimal policy under commitment, and shows that in this case government spending is not as important to stabilize demand. Most recent paper have tried to make quantitative statements, with more reduced form representation of policy than Eggertsson (2001), and our paper is more in that spirit (i.e. we study the effect of policy on the margin). Some papers find quite different results than reported here. For example, Cogan, Coenen, Wieland and Taylor (2009), find a very small effect of government spending, even if their model is similar (altough a bit more detailed). What is the reason for the difference? Our simple framework, allows us to see explicitly why this is the case. In short, the main reason is that while they assume a permanent increase in government spending, we assume a well targeted stimulus package that is *temporary,* a point developed in a bit more detail in Eggertsson (2009b). Another recent study by Christiano, Eichenbaum and Rebelo (2009) also find a large effect of government spending, consistent with the previous literature on fiscal policy at the zero bound such as Eggertsson (2004,6) and Christiano (2004). In contrast to this study, however, they find that the multiplier is extremely sensitive to parameter values and report a large range for plausible values for the multiplier. In contrast, our analysis reveals relatively tight estimates of the multipliers (at least in comparison to that study).

What is the reason? Our simple example, and closed form solutions, allows us to show explicitly why a Bayesian estimation gives a tighter estimate of the multipliers. The reason is that the estimation strategy constrains the model parameters to replicate a "benchmark scenario", i.e. a the Great Depression, while Christiano et al (2009) consider parameter configurations that can lead to unbounded output and deflation collapses, in which case the multipliers become very large (what they term as the "divine coincidence").

Eggertsson (2009a) details the theoretical foundations of the model we use in this paper. Accordingly, we only briefly review the model. The main contribution of this paper is the estimation strategy. The tax cuts and spending are modelled by reduced form policy rules, while Eggertsson (2004,6) and Eggertsson and Woodford (2004) studies the optimal commitment and/or the Markov Perfect equilibrium of a subset of the policy instruments.

2 A Microfounded Model and Key Analytical Results

This section summarizes a standard New Keynesian model described in a bit more detail in Eggertsson (2009a). A consumer maximizes $E_t \sum_{T=t}^{\infty} \beta^{T-t} \xi_T \left[u(C_T) + g(G_T) - \int_0^1 v(l_T(j))dj \right]$ over time, where E_t is an expectation operator, $C_t \equiv \left[\int_0^1 c_t(i)^{\frac{\theta}{\theta-1}} di \right]^{\frac{\theta-1}{\theta}}$ is a Dixit-Stiglitz aggregate of consumption of variety $c_t(i)$, $u(.)$ is utility of consumption, $v(.)$ disutility of labor of each of labor variety $l_T(j)$, β is a discount factor between 0 and 1 and ξ_T is a shock to preferences. The budget constraint of the households is $(1+\tau_t^s)P_t C_t + B_t = (1-\tau_{t-1}^A)(1+i_{t-1})B_{t-1} + (1-\tau_t^P)\int_0^1 Z_t(i)di + (1-\tau_t^w)\int_0^1 W_t(j)l_t(j)dj - T_t$ where $Z_t(i)$ is profits that are distributed lump sum to the households. There are five types of taxes in the baseline model: a sales tax τ_t^s on consumption purchases, a payroll tax τ_t^w, a tax on financial assets τ_t^A, a tax on profits τ_t^p, and a lump-sum tax T_t. On the firm side, there are monopolistically competitive firms (each producing each variety of the consumption good) with a production function that is linear in labor (we abstract from capital) and can only adjust their price with a probability $(1-\alpha)$ in each period. For the details on the non-linear model, see Eggertsson (2009a). It is convenient to summarize the model by "aggregate demand" and "aggregate supply".Aggregate demand (AD) is the equilibrium condition derived from the optimal consumption decisions of the household where we have used the aggregate resource constraint to substitute out for consumption. In its log-linearized form it can be written as

$$\hat{Y}_t = E_t \hat{Y}_{t+1} - \sigma(i_t - E_t \pi_{t+1} - r_t^e) + (\hat{G}_t - E_t \hat{G}_{t+1}) + \sigma E_t(\hat{\tau}_{t+1}^s - \hat{\tau}_t^s) + \sigma \hat{\tau}_t^A, \tag{1}$$

where i_t is the one-period risk-free nominal interest rate[1], π_t is inflation, r_t^e is an exogenous shock, and E_t is an expectation operator and the coefficient is $\sigma > 0$.[2] $\hat{Y}_t$ is output in log

[1] In terms of our previous notation, i_t now actually refers to $log(1+i_t)$ in the log-linear model. Observe also that this variable, unlike the others, is not defined in deviations from steady state. I do this so that we can still express the zero bound simply as the requirement that i_t is nonnegative.

[2] The coefficients of the model are defined as $\sigma \equiv -\frac{\bar{u}_{cc}}{\bar{u}_c \bar{Y}}$, $\omega \equiv \frac{\bar{v}_y \bar{Y}}{\bar{v}_{yy}}$, $\psi \equiv \frac{1}{\sigma^{-1}+\omega}$, $\kappa \equiv \frac{(1-\alpha)(1-\alpha\beta)}{\alpha} \frac{\sigma^{-1}+\omega}{1+\omega\theta}$, where bar denotes that the variable is defined in steady state.

deviation from steady state, $\hat{G}_t$ is government spending in log deviation from steady state, $\hat{\tau}_t^s$ is sales taxes in log deviation from steady state, $\hat{\tau}_t^A$ is log deviation from steady state,[3] and r_t^e is an exogenous disturbance.[4] Aggregate supply (AS), is the equilibrium condition derived by the optimal production and pricing decisions of the firms. It's log-linearized form can be written as

$$\pi_t = \kappa \hat{Y}_t + \frac{\kappa}{\sigma^{-1} + \omega}(\hat{\tau}_t^w + \hat{\tau}_t^s) - \frac{\kappa \sigma^{-1}}{\sigma^{-1} + \omega}\hat{G}_t + \beta E_t \pi_{t+1}, \tag{2}$$

where the coefficient $\kappa \equiv \frac{(1-\alpha)(1-\alpha\beta)}{\alpha}\frac{\sigma^{-1}+\omega}{1+\omega\theta}$, β is the rate of time preference and $0 < \beta < 1$. Without going into details about how the central bank implements a desired path for the nominal interest rates, it is assumed that it cannot be negative so that

$$i_t \geq 0 \tag{3}$$

Monetary policy follows a Taylor rule, with a time-varying intercept, that takes the zero bound into account

$$i_t = \max(0, r_t^e + \phi_\pi \pi_t + \phi_y \hat{Y}_t) \tag{4}$$

where the coefficients $\phi_\pi > 1$ and $\phi_y > 0$. For given policy rules for taxes and spending, equations (1)-(4) close the model. Note that we assume, as the model in Eggertsson (2009a), lump-sum taxes. This means that we do not need to keep track of the government budget constraint in our experiments; for a given evolution of government spending and the tax instruments, we can assume the lump-sum taxes make up for any budgetary shortfall (Eggertsson (2009a)).

3 An output collapse at the Zero Bound

We study the response of the model to a one time shock to r_t^e which reverses back to steady state with probability $1 - \mu$ in each period.

A1 – Structural shocks: $r_t^e = r_L^e < 0$ *unexpectedly at date* $t = 0$. *It returns back to steady state* $r_H^e = \bar{r}$ *with probability 1*$-\mu$ *in each period. The stochastic date the shock returns back to steady state is denoted* T^e. *To ensure a bounded solution, the probability* μ *is such that* $L(\mu) = (1-\mu)(1-\beta\mu) - \mu\sigma\kappa > 0$.

One can interpret a negative r_t^e as a preference shock. Everyone suddenly wants to save more so that the real interest rate has to decline for output to stay constant. There are better interpretations, however. Curdia and Eggertsson (2009), show that a model with financial frictions can also be reduced to equations (1)-(2). The shock r_t^e corresponds to an exogenous increase in the probability of default by borrowers. What is nice about this interpretation is that r_t^e can now

[3] Here, $\hat{G}_t$ is the percentage deviation of government spending from steady-state over steady-state aggregate output. In the numerical examples, $\hat{\tau}_t^A$ is scaled to be comparable to percent deviation in annual capital income taxes in steady state so that it corresponds to $\hat{\tau}_t^A \equiv 4 * (1-\beta)\log\{\tau_t^A/(1-\bar{\tau}^A)\}$.

[4] It is defined as $r_t^e \equiv \log \beta^{-1} + E_t(\hat{\xi}_t - \hat{\xi}_{t+1})$, where $\hat{\xi}_t \equiv \log \xi_t/\bar{\xi}$.

be mapped into the wedge between a risk free interest rate rate and a interest rate paid on risky loans. Both rates are observed in the data. The wedge implied by these interest rates exploded in the US economy during the crisis of 2008, giving empirical evidence for a large negative shock to r_t^e. A banking crisis – characterized by an increase in probability of default by banks and borrowers – is our story for the model's recession.

Eggertsson (2009a) describes the model graphically and gives intuition for condition A1 on μ. It is replicated in figure 2. The figure plots the AS and the AD curve, conditional on the shock given in A1, i.e.

$$\mu\hat{Y}_L = +\sigma(1-\mu)\pi_L + \sigma r_L^e + \mu\hat{G}_L - \sigma\mu\hat{\tau}_L^s + \sigma\hat{\tau}_L^A$$

$$(1-\beta\mu)\pi_L = \kappa\hat{Y}_t + \frac{\kappa}{\sigma^{-1}+\omega}(\hat{\tau}_L^w + \hat{\tau}_L^s) - \frac{\kappa\sigma^{-1}}{\sigma^{-1}+\omega}\hat{G}_L$$

where we assume that policy moves perfectly in sync with the shock (as further discussed below) and thus "shifts" the curves.

The solid lines in figure 2.show the equilibrium in period 0 if we know for sure that the shock is over in period 1. It is completely determined by the AD equation and the AS equation just pins down inflation. Consider now an increase in μ so there is a probability the economy can stay in the depressed state longer than one period. Then the AD curve becomes flatter and the AS curve steeper, and the equilibrium is characterized by greater deflation and an output collapse. As we increase μ, the collapse increases without a bound so that in the limit, the two curves become parallel and the deflation and the collapse approach infinity, a *deflationary blackhole.* We assume that μ is never large enough to approach this limit. Hence the collapse in output and prices is bounded by assumption A1.

Solving the AD and AS equations with respect to π_L and $\hat{Y}_L$, we obtain

$$\pi_t = \frac{1}{(1-\mu)(1-\beta\mu)-\mu\sigma\kappa}\kappa\sigma r_L^e < 0 \text{ if } t < T^e \text{ and } \pi_t = 0 \text{ if } t \geq T^e \tag{5}$$

$$\hat{Y}_t = \frac{1-\beta\mu}{(1-\mu)(1-\beta\mu)-\mu\sigma\kappa}\sigma r_L^e < 0 \text{ if } t < T^e \text{ and } \hat{Y}_t = 0 \text{ if } t \geq T^e \tag{6}$$

$$i_t = 0 \text{ if } t < T^e \quad \text{and} \quad i_t = r_t^e \quad if \quad t \geq T^e$$

We consider both tax and spending multipliers at zero and at positive interest rates. Our discussion will focus mostly on the multiplier of government spending, as this has been the focus of much of the literature, but we report all other multipliers as well. Consider first the multiplier at positive interest rates. Imagine an increase in government spending at time 0, $\hat{G}_L > 0$, that is reversed back with probability $1-\rho$ in each period. Because there are no shocks, the interest rate remains positive. Also note that because the model is purely forward looking then $\hat{Y}_t = \hat{Y}_L$ for all $t < T$ where T is the stochastic data at which government spending reverts back to steady state. Solving (1)-(4) together this yields the multiplier of government spending at positive interest rates:

$$\frac{\Delta\hat{Y}_L}{\Delta\hat{G}_L^N} = \frac{(1-\rho)(1-\rho\beta)+(\phi_\pi-\rho)\frac{\kappa}{\sigma^{-1}+\omega}}{(1-\rho+\sigma\phi_y)(1-\rho\beta)+(\phi_\pi-\rho)\sigma\kappa} > 0$$

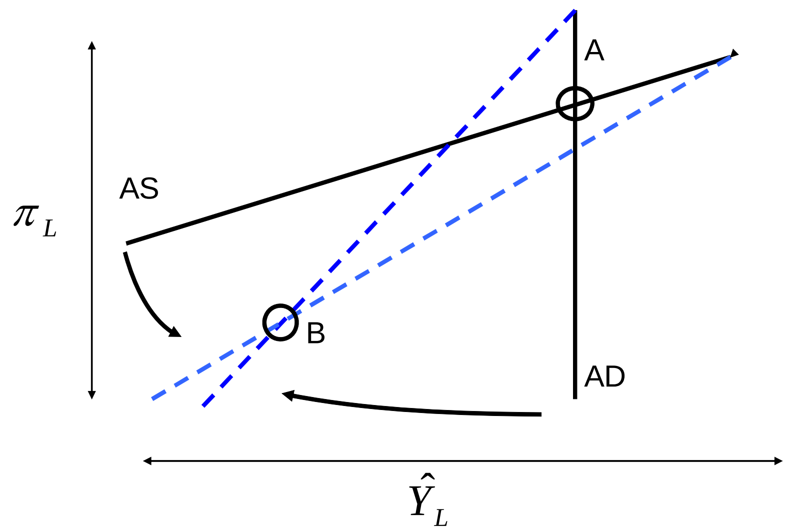

Figure 2: The effect of multiperiod recession.

This multiplier answers the question: By how much does a one dollar increase in government spending (which is reversed with probability $1-\rho$) increase output? Let us now compute the multiplier of government spending at zero interest rates. We now assume the shock in A1 that reverts back to steady state with probability $1-\mu$ in each period. Our thought experiment is to compute the effect of increasing government spending in all states of world in which the shock perturbs the economy, i.e., we are considering an increase in $\hat{G}_L > 0$ that is reversed with probability $1-\mu$ in each period (and is perfectly correlated with the shock). Hence we are interested in measuring the effect of increasing government spending at zero interest rates, where the zero interest rate bound is binding because of the shock. At zero interest rates, the multiplier of government spending is:

$$\frac{\Delta \hat{Y}_L}{\Delta \hat{G}_L^N} = \frac{(1-\mu)(1-\beta\mu) - \frac{\mu\kappa}{\sigma^{-1}+\omega}}{(1-\mu)(1-\beta\mu) - \sigma\mu\kappa} > 1 \tag{7}$$

The objective of the estimation is to find some reasonable values for these multipliers. The multiplier for the other tax instruments can also be expressed analytically, see Eggertsson (2009a) for details. For taxes, we also consider the same thought experiment as above, i.e. taxes are cut through the duration of the shock.

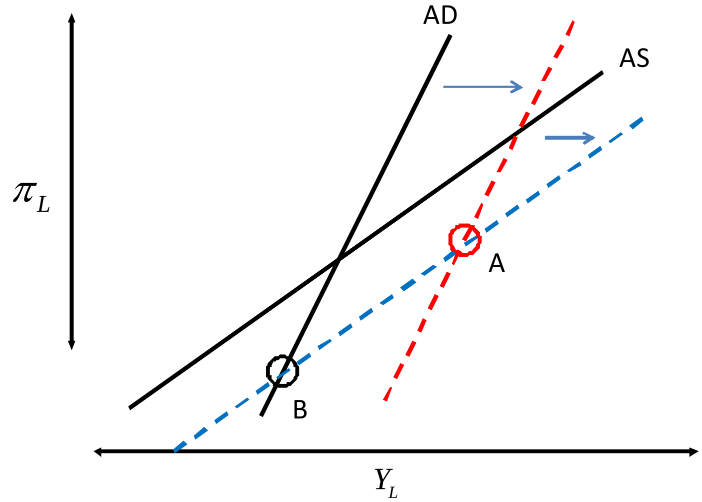

Figure 3: The effect of increasing aggregate demand vs aggregate supply.

4 The basic mechanism: A graphical illustration

Before getting into the estimation, it is helpful to give the reader a quick feel for what to expect in terms of signs of the multipliers of tax and spending and build up further understanding of the main mechanism of the model. Consider the two equations (1)-(2) from previous section. They are plotted up in figure 3. What is the effect of increasing government spending? This shifts the AD curve out because aggregate demand is the sum of government and consumption spending. But it also increases aggregate supply. Why? Because higher government spending takes away resources for consumption for a given level of output, thus increasing people willingness to work. This is the standard "wealth effect" documented in the RBC literature. A new equilibrium is found at point A.

Consider now the effect of cutting labor taxes. In our setting this has no effect on aggregate demand. It does, however, increase aggregate supply. Because people now get more money in their pockets for working, they want to do more of it. A new equilibrium is found at point B. Note, however, that output is lower at point B than at the original intersection of the two curves! Because everybody want to work more, then somehow, output is lower. How can this be? This is peculiar to the economic environment of zero nominal interest rates. A tax cuts get everybody to work more. At zero interest rate, however the problem is insufficient demand. Because everybody wants to works more, this reduces labor costs which creates current and expected deflation. Because AD depends on the real interest rate, the difference between the nominal interest rate and expected inflation, this reduces demand, which is precisely what happens at point B. *Because* *everybody wants to work more, then everybody works less in equilibrium!* Again, it is important to stress that this result is special to zero interest rates. Under usual circumstances a central bank would offset the deflationary pressure by cutting the nominal interest rate (by more than one-to-one). This means that the AD curve would be downwards sloping in (π_L, Y_L) space . Then a shift out in the AS curve increases output. The current result is created by the fact that the central bank cannot cut the nominal interest rate to accommodate the increase in aggregate supply.

The general principle we can grasp from the discussion in the past two paragraphs is that at zero interest rates the problem is insufficient spending, non insufficient supply. Anything that increases aggregate spending is stimulative, while anything that increases aggregate supply can be counterproductive because it may increase deflationary pressures. Armed with this insight, it is straight forward to consider the effect of other tax variations. Consider first the effect of cutting sales taxes temporarily. This has the effect of shifting out aggregate demand because people want to spend today in order to take advantage of the lower taxes. Meanwhile, the lower taxes today, means that the utility of each additional dollar spend is higher, so people want to work more. This shifts out aggregate supply. Hence the analytics work out the same way as an increase in government spending (as can be observed by looking at equation (1)-(2)) and the equilibrium is found at again found in point A. [5] Consider now the consequence of cutting taxes on savings. This gives people the incentive to save, instead of spend, thus shifting the AD curve backwards, leading to a new equilibrium at lower inflation and output. Clearly this is a contractionary tax cut in our environment.

5 A simple calibration and an informal explanation for why a Bayesian approach is helpful

Our objective is to study the value of the policy multipliers in the model, i.e. what is the quantitative effect of cutting taxes or increasing spending when the interest rate is zero? A natural approach is to estimate the model over some data range and base the analysis on past episodes. The main problem with this approach is that until now US economic history documents only one period in which the zero bound was binding, that is, the Great Depression. This period was also associated with a host of other policy initiatives, and eventually a war. Essentially, this period thus only reflects "one" datapoint.

The approach we take here is to ask the model to choose the shocks and the model parameters to match the trough of the Great Depression, i.e. the value of output and deflation in 1932 assuming that no policy was in place at that time. This is a reasonable benchmark because it was only in the spring of 1933 that FDR takes office and implements several expansionary policies (see e.g. Eggertsson (2008)). At that time output was around -30 percent in deviation from its peak value in 1929 and deflation around -10 percent, and the short-term nominal interest rate, as for example measured by 3-month Treasuries, was close to zero.

[5] There are some subtleties involved which is discussed in some more detail in Eggertsson (2009a).

How can the model match the numbers from the Great Depression cited above? Inflation and output are given by (5) and (6). These numbers depend on the three parameters (κ, σ, β) and the shock r^e_L and its persistence μ. One approach is to pick a number for (κ, σ, β) and then "back out" the value for (r^e_L, μ) that generates the data. Consider the following values $(\kappa, \sigma, \beta) = (0.02, 1, 0.99)$. What is the implied value for (r^e_L, μ)? Using (2) we obtain

$$\mu = \frac{\pi_{L-}\kappa Y_L}{\beta \pi_L} = \frac{-0.1/4 + 0.02 * 0.3}{-0.99 * 0.1/4} = 0.7677 \tag{8}$$

Then using (1) we obtain

$$r_L = \sigma^{-1} \frac{(1-\mu)(1-\beta\mu) - \mu\sigma\kappa}{1-\beta\mu} \hat{Y}_t = -0.0505 \tag{9}$$

In order to evaluate the multiplier of government spending, we also need to parameterize $\frac{1}{\sigma^{-1}+\omega}$. For illustration let us assume this parameter is $\frac{1}{2}$, but we go into a more detailed calibration in the next section. Using these values we can now compute the multiplier of government spending at zero interest rates as

$$\frac{\Delta \hat{Y}_L}{\Delta \hat{G}_L} = 1.19 \tag{10}$$

How sensitive is the conclusion to the calibration? One approach is to increase the value of one of the parameters, keeping all the others fixed. Consider for example varying κ while maintaining $(\sigma, \beta, r^e_L, \mu, \omega)$ fixed. Figure 4 shows the value of the multiplier as a function of κ keeping the other parameters constant. As can be seen in the figure the value of the multiplier increases without a bound, approaching infinity at the boundary of the inequality in A1, a similar experiment as in Christiano et al (2009). In terms of figure 2, this is when the AD and AS curves are parallel. At this point the multiplier approaches infinity. Observe that the variations in κ that gives rise to this behavior are relatively "reasonable" in the sense that one can justify the level of price rigidities needed to increase κ to this extent.

Is then all we can say about the multiplier of government spending that it is between 1 and ∞? The answer is no in our view. We can exclude the extreme values of the multiplier and come up with a much more reasonable range if we are ready to make some assumptions about what the shocks and the parameters are supposed to generate in terms of the observable variables, i.e. output and inflation.

Consider the parameter value for $\kappa = 0.0699$ shown by a circle in figure 4. This value gives rise to a multiplier of 13.8. To compute this number, however, we have kept all the other parameters of the model constant at $(r^e_L, \mu, \sigma, \beta) = (-0.0505, 0.7677, 1, 0.99)$ while increasing κ. As a consequence the model no longer matches the data from the Great Depression which guided our value for the shocks in the first place. Consider the implied value for output and inflation as we increase κ keeping the other parameters fixed. It is shown in figure 5. The figure shows that as the multiplier increases, output decreases without a bound, and so does deflation. In the case of $\kappa = 0.0699$ output "collapses" by -578 percent and deflation by -47 percent. Clearly, this is not a feasible solution, since output cannot be negative. In terms of figure 2, the way this is working is

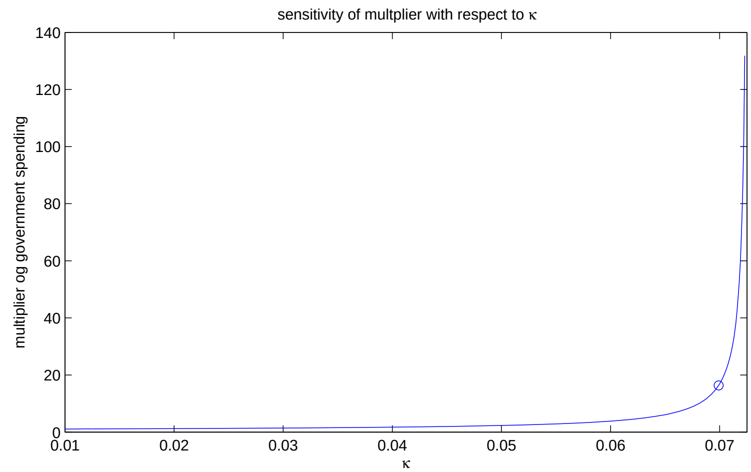

Figure 4: For given value for the other parameters and a given value for the shock, the multiplier of government spending is extreemely sensitive to increasing κ.

that the increase in κ makes the AS curve steeper, which pushes deflation and output into more and more negative territory and thus approaching the boundary condition excluded by A1.

Recall that the shocks are chosen in (8) and (9) so as to generate a particular scenario, i.e. that of -30% drop in output and -10% inflation. Hence a sensitivity analysis of the kind just performed may not be very informative if we want the model to generate the Great Depression. A different value of κ should lead one to consider a different value of the shocks. The implied value of the shocks using (8) and (9) and $(\kappa, \sigma, \beta) = (0.0699, 1, 0.99)$ is $(r_L, \mu) = (-0.2471, 0.16)$. Using these values, the multiplier is no longer 13.8 but instead 1.0082! The sensitivity analysis above does give some valuable insights, however. What we found was that the multiplier becomes very large in the event that the output collapse is very severe. This is an intuitive result, and one that may be important for policymakers, as emphasized by Christiano, Eichenbaum and Rebello (2009).

The approach of backing the shocks mechanically out of (8) and (9) when changing the other parameters has obvious weaknesses. Consider the example we just constructed. In this case, the value of μ is 0.16, which suggests that the public puts only a 16 percent chance on that the recession will last more than one quarter! Because of this small duration of the disturbance, it also has to take on an extremely large negative value, measured in annual terms, in order to generate a Great Depression. The shock corresponds to the "first best" real interest rate – i.e. the real interest rate needed for no contraction in the model – of -98.84 percent. These values are not reasonable, at least if one believes that the data on real rates in post war date have been

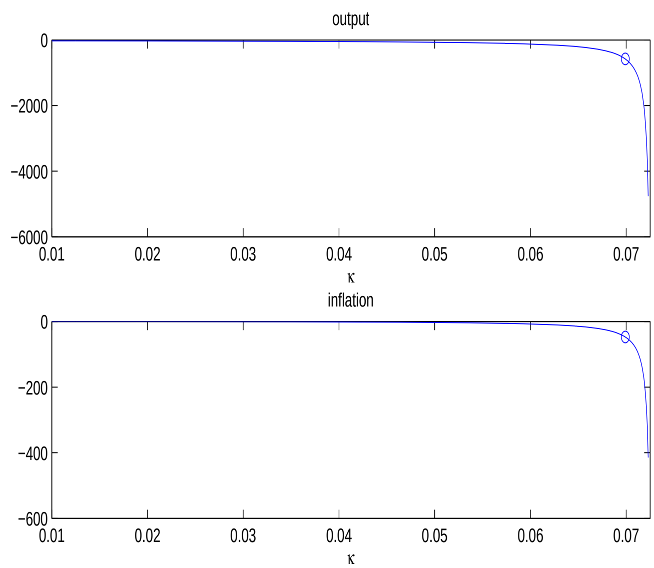

Figure 5: High values of the multiplier tend to be associated with explosive collapse in output and deflation.

within some reasonable range from a first best allocation (more on this in the next section). One can even also argue that our original example was "unreasonable" because the value of the shock was then (in annual terms) -20%. Another obvious question is, even if we choose $\kappa = 0.0699$, can there still be values for the shock that are "reasonable" as long as we are willing to change the values assumed for σ,β? But how should one decide between the different values of σ, β, κ and the shocks? And what should we consider "reasonable"?

To address the issues sketched out above, we suggest a Bayesian approach. We specify a prior distribution for all of the structural parameters of the model and the shocks. We then choose those parameters so that the model matches a 30 percent drop in output and 10 percent deflation as closely as possible, while simultaneously matching the priors as well as we can. In other words, we choose the parameters and the shocks to maximize the posterior distribution of the model. Once this is done, we argue, one can construct a reasonably informative sensitivity analysis, and one that has some advantages over the simple comparative statics we did above. Moreover, one can construct with some confidence an estimate for the multipliers that are much tighter than an analysis based on figure 4. In the next section we discuss our priors. We then move onto deriving the posterior of the model.

6 Choosing the priors

Let us denote the priors by Ω. The priors, shown in Table 1, reflect our ex ante "believes" about the parameters and shock in the model, i.e. what we think is "reasonable". The priors are chosen so that θ has a mean of 8 (consistent with markup of 14 percent) but a standard deviation of 3. We choose a loose prior on this parameter, since the literature documents very broad range for it, and thus we wish to take this uncertainty into account. We assume that price rigidities are consistent with prices being adjusted on average once every nine months, with a standard deviation of about one and a half month. The value of β is consistent with a 1.33 percent average annual interest rate (extracted from post WWII data) with a 0.4 percent standard deviation. For σ^{-1} we assume a mean of 2 with a standard deviation of 1. For ω we assume a mean of 1 with a standard deviation of 0.75. Both values are relatively standard in the literature and encompass some reasonable degree of uncertainty.

Our conjecture is that the crisis is created by an intertemporal shock. Hence choosing our prior for the shock is important. To choose this prior for we do the following. First, imagine an economy which is always at its first best. In this fictional economy, the process for r_t^e is equivalent to the ex-ante real interest rate, i.e. $r_t = i_t - E_t \pi_{t+1}$. Let us now suppose that the economy has been at its first best since WWII (i.e. since 1952 onwards), because the zero bound has not been binding during that period. While we do not believe this literally, it seems "reasonable" to believe that the first best allocation has not been radically different from observed outcomes in this period, given that unemployment has never been close to the levels seen during the Great Depression, and that we have furthermore not observed zero interest rate in this period. This means that the actual real rate is informative to construct our prior for stochastic process underlying the first best real interest rate, i.e. our fundamental shock. This prior is what helps inform our approach because we will then impose the prior that the shock causing the Great Depression was "big" where big will be defined in reference to the stochastic process for actual observed ex ante real interest rates.

We first construct a times series for the ex-ante real interest rate. We then fit it by an AR(1) process. We construct inflation expectations by a three variable reduced form VAR with inflation, output and the short-term nominal interest rate. Our estimated AR(1) process is

$$r_t = 0.058 + 0.83 r_{t-1} + \epsilon_t \tag{11}$$

Suppose the shock r_L^e in the low state is equivalent to a 3 standard deviation shock given by this process. This corresponds to a shock that would only occur about 0.3% of the time. This is a very extreme shock. This results in a value for $r_L^e = -0.0102$ or -4.1% in annual frequencies. We will see soon, that this assumption is important in estimating the "size" of the multipliers. In particular if we assume that a *smaller shock* would give rise to the Great Depression, then the value of the multiplier would be bigger. Our prior thus stacks the cards against finding a big multiplier as we will make clear later on.

While ex-ante real rates would be an accurate measure of the efficient rate of interest only in the event output was at its efficient rate at all times, this gives at least some sense of a "large" shock as a source of the Great Depression. A key element of this assumption, is that we assume that the economy is not "too far" from its first best post WWII. In particular, extreme shocks like we considered in section 5 get low weight in the posterior of the model. The prior on the persistence of the shock is that it is expected to reach steady state in 12 quarters, or 3 years. It seems reasonable to suppose that in the midst of the Great Depression people expected it to last about 3 years. Because we don't have a strong prior on either the exact value of the shock, nor its persistence, we allow for relatively large variance for the priors on both the size of the shock and of its persistence. A one standard deviation of our prior for r_L from it's mean leads the shock to being -6% (or -2% if smaller) or the expected duration of 30.5 quarters (or 7.5 if smaller).

Table 1: Priors

	distribution	mean	standard deviation
α	beta	0.66	0.05
β	beta	0.99669	0.001
$1-\mu$	beta	1/12	0.05
σ^{-1}	gamma	2	0.5
ω	gamma	1	0.75
θ	gamma	8	3
r_L	gamma	-0.010247	0.005

7 Analytical Derivation of the Posterior

We only try to match the model to one datapoint, the trough of the Great Depression, i.e. the beginning of the year 1933. At that time the short-term nominal interest rate was close to zero, output was near -30% from its 1929 value and deflation was on the order of -10 percent (Eggertsson (2008)). A posterior distribution is the probability distribution for observing a particular value for each of the parameters and shocks, given the data observed. Let us denote the posterior $p(\Omega/X)$ where X is the data and π_{1933} and $\hat{Y}_{1933}$ are the elements of X. In order to construct the posterior, it is necessary, for computational reasons, to assume that there is a random discrepancy between the model and the data.[6] We assume

$$\pi_{1933}^{data} = \pi_L + \epsilon^{\pi}$$

$$\hat{Y}_{1933}^{data} = \hat{Y}_L + \epsilon^{Y}$$

where ϵ is a normally distributed measurement error. The posterior of the model can now be constructed using Bayes rule. By Bayes rule, we have

$$p(\Omega/X) = \frac{p(X/\Omega)p(\Omega)}{p(\Omega)}$$

We assume that the priors are independent so that we can write $\ln p(\Omega) = \sum_{\psi_s \in \Omega} f(\psi_s)$. Because we assumed that the measurement error is normally distributed we can write $\ln p(X/Y) = -\frac{(\epsilon^{\pi})^2}{2\sigma_{\pi}^2} - \frac{(\epsilon^{Y})^2}{2\sigma_{Y}^2} + constant$. Combining, and ignoring variables that are not a functions of Ω, the natural logarithm of the posterior likelihood can be written as

$$\log p(\Omega/X) = -\frac{(\pi_L(\Omega) - (-\frac{0.1}{4}))^2}{2\sigma_{\pi}^2} - \frac{(\hat{Y}_L(\Omega) - (-0.3))^2}{2\sigma_{Y}^2} + \sum_{\psi_s \in \Omega} f(\psi_s), \tag{12}$$

where $\pi_L(\Omega)$ and $\hat{Y}_L(\Omega)$ are given by (5) and (6) and we have substituted for the data points (note that the time unit is a quarter of a year). The likelihood is formed conditional on the hypothesis that the shock r_L is in the "low state." The only data we match is a decline in output of 30 percent and a drop in inflation of 10 percent. As implied by the discussion above, the functions $f(\psi_s)$ measure the distance of the variables in Ω from the priors, where the parameters and shocks are denoted $\psi_s \in \Omega$. The distance functions $f(\psi_s)$ are given by the statistical distribution of the priors listed in Table 1 below. Note that we use a gamma distribution for parameters that are constrained to be positive and a beta distribution for parameters that must be between 0 and 1. In the case of θ, for example, we have

$$\log f(\theta; a, b) = (a-1)\log\theta - \theta * b$$

[6] This may not be obvious from the context here. We are currently working on a version of the estimation the does away with the measurement error, deriving the posterior of the model without assuming it. This version of the paper will replace this draft shortly. In this application we assume that the measurement error has standard deviation of 10^(-6).

where a is called the shape parameter and b the rate parameter. They are set to match the mean and the standard deviation we specified for the prior for gamma.

What is the interpretation of the posterior likelihood function (12)? This probability measure indicates the probability of a particular parameter/shock configuration Ω, given the data point from the Great Depression, and conditional on that the model solution given by (5) and (6) is a correct up to a normally distributed measurement error. The mode of this posterior will then tell us the "most likely" value for Ω given that we ask the model to match this data.

There is much more information in the posterior than the mode of this distribution, however. We can, for example, compute probability distributions for any function of the parameters/shocks Ω, e.g., the various multipliers, given the criteria suggested by the posterior. As we will discuss below, the posterior also gives us a formal way of thinking about the "sensitivity" of each of the various economic objects, such as the spending and tax multipliers, to variations in some of the structural parameters.

8 Numerical Approximation of Posterior

This section summarizes the numerical methods used to characterize the posterior. The next section discusses the results and their interpretation. The mode of the posterior is characterized numerically and is reported in Table 1. To compute the mode of the posterior we use a Matlab maximization routine developed by Christopher Sims. This code is available on our website. In order to characterize the entire posterior distribution, we use a Metropolis algorithm, briefly summarized below, using the mode as a starting point. Let y^T denote the set of available data and Ω the vector of coefficients and shocks. Further, let Ω^j denote the jth draw from the posterior of Ω. A new draw is obtained by drawing a candidate value, $\widetilde{\Omega}$, from a Gaussian proposal distribution with mean Ω^j and variance sV, where s is a scaling factor. We then set $\Omega^{(j+1)} = \widetilde{\Omega}$ with probability equal to

$$\min\{1, \frac{p(\Omega/y^T)}{p(\Omega^j/y^T)}\}. \tag{13}$$

If the proposal is not accepted, we set $\Omega^{(j+1)} = \Omega^j$.

We set V equal to the inverse Hessian of the posterior evaluated at the mode, while s is chosen in order to achieve an acceptance rate approximately equal to 23 percent, as suggested by Gelman et al (2003). We run two four chains of 1,000,000 draws and discard the first 200,000 to allow convergence to the ergodic distribution. We conducted a variety of checks to make sure that the sample distribution has converged to the ergotic distribution.[7]

[7] To ensure that the sample distribution has converged to the ergodic distribution, we follow the approach of Gelman et al (2003). We checked that the potential scale reduction of the distribution is near one, calculated the effective number of independent draws, and inspected the plots of the draws from the posterior.

9 Results

	Prior 5%	Prior 50%	Prior 95%	Posterior 5%	Posterior 50%	Posterior 95%	Mode
alpha	0.5757	0.6612	0.7402	0.7026	0.7633	0.8164	0.7747
beta	0.9949	0.9968	0.9981	0.9948	0.9967	0.9981	0.9970
1-mu	0.0198	0.0740	0.1788	0.0705	0.1045	0.1523	0.0971
omega	0.1519	0.8200	2.4631	0.7756	1.9033	3.8332	1.5692
reL	-0.0196	-0.0094	-0.0036	-0.0266	-0.0146	-0.0066	-0.0104
sigma^-1	1.2545	1.9585	2.8871	0.8107	1.2161	1.7779	1.1599
theta	3.7817	7.6283	13.4871	8.3394	13.2496	19.7410	12.7721

Table 2: Priors and Posteriors of the parameters and shocks.

Table 2 shows priors for the parameters, the shock and the posterior. We see that the prior and posterior distributions for each of the parameters and shocks are relatively similar for most of the parameters, in most cases slightly tighter. In the case of θ, the posterior is broader. Does this mean that the data and the model puts no structure on the parameters and the shocks, so that there is no updating of the priors? No. Even if the unconditional distributions for each parameter do not change much, the correlations between each parameters is significantly different according to the posterior than according to the priors. The priors of the parameters are not correlated. Table 3 illustrates, however, that the correlation between the parameters in the posterior of the model are different from zero.

What is the intuition for the significant correlation between the parameters in the posterior of the model even if the priors are independent? The logic is straightforward. Consider, for example, the equation for output in 1933 given by (6). It can be written as

$$\hat{Y}_{1933} = \frac{1-\beta\mu}{\sigma^{-1}(1-\mu)(1-\beta\mu)-\mu\kappa} r_L^e < 0$$

Our estimation procedure makes $\hat{Y}_{1933}$ as close to the data – i.e. -0.3 – as possible, subject to the priors. Consider variations in the parameters around the mode in the posterior, where, for example, $\sigma^{-1} = 1.16$. Consider increasing σ^{-1} a bit. We see that if all other parameters are kept "constant", this will result in the coefficient in front of r_L^e to be smaller so that then $\hat{Y}_{1933}$ is less negative and thus moves away from -0.3.[8] However, our assumption is that the "measurement error" is small. This means that this parameter configuration gets a small probability weight according to the posterior. In other words, this parameter configuration is very improbable. More probable is that as σ^{-1} gets higher, then this increase is offset by some of the other parameters moving. For example, as σ^{-1} increases, then in order to keep $\hat{Y}_{1933}$ constant this can be offset by an increase in r_L^e. That should create a negative correlation between r_L^e and σ^{-1}. This conjecture is confirmed in Table 3. One can make a similar argument to explain the correlation between the other parameters.

[8] I put quotation mark around "constant" because the parameter κ is also a function of σ^{-1} and this works in the other direction. However, that effect is of order of magnitude smaller.

	alpha	beta	gamma	sigmainv	omega	theta	reL
alpha	1	-0.0115	-0.5595	0.4633	-0.1913	-0.4403	0.1770
beta	-0.0115	1	0.0319	-0.0157	-0.0091	0.0202	-0.0266
gamma	-0.5595	0.0319	1	-0.7517	-0.2389	-0.2506	-0.4121
sigmainv	0.4633	-0.0157	-0.7517	1	0.2803	0.2033	-0.2119
omega	-0.1913	-0.0091	-0.2389	0.2803	1	-0.0844	-0.0313
theta	-0.4403	0.0202	-0.2506	0.2033	-0.0844	1	0.0650
reL	0.1770	-0.0266	-0.4121	-0.2119	-0.0313	0.0650	1

Table 3: Correlation between the parameters in the posterior distribution.

Here is another way to look at the strong correlations between the parameters in the posterior: It means that if we move one of the parameters while keeping all other fixed, we are likely to obtain a parameter configuration that is very unlikely according to the posterior. A simple way of seeing this is to explore the sensitivity analysis we did in section 5. As we increased the value of κ, this had a very large effect on the value of the multiplier of government spending, *holding the other parameters and the shock fixed*. At the same time, however, this implied that that model predicted an extreme output collapse and exploding deflation. Because we parameterize the model to match almost exactly 30 percent collapse in output and 10 percent deflation, the posterior says that this parameter configuration is unlikely. Hence, the large range of multipliers discussed in section 5 are unlikely, if the posterior is used as a judge.

Table 4 uses the posterior to create probability distributions for the multipliers. We see that the mode for the multiplier for government spending is 2.27 which is a bit higher than many other estimates. The 5-95 percent posterior confidence bands for this number is between 1.42 and 3.2. While this is a relatively large posterior band, it still gives a better sense for the relevant magnitudes than the sensitivity analysis in section 5 might lead one to believe, where reasonable parameter variations appeared to give a multiplier pretty much anywhere between 1 and ∞. Moreover, because we now have a well-defined posterior distribution, we can do a quite informative sensitivity analysis.

	Posterior 5%	Posterior 50%	Posterior 95%	Mode
tax cut multiplier i>0	0.0476	0.0800	0.1434	0.0962
tax cut multiplier i=0	-1.3890	-0.4994	-0.2132	-0.8153
gov spending multiplier i>0	0.2911	0.3428	0.4038	0.3247
gov spending multiplier i=0	1.4295	1.9258	3.2064	2.2793
sales tax cut multipler i>0	0.2541	0.4180	0.6578	0.3766
sales tax cut multipler i=0	1.4883	2.3851	4.1760	2.6438
capital tax cut multipler i>0	-0.0049	-0.0034	-0.0024	-0.0033
capital tax cut multipler i=0	-0.6748	-0.3031	-0.1605	-0.4048

Table 4: The posterior distibution of the fiscal multipliers.

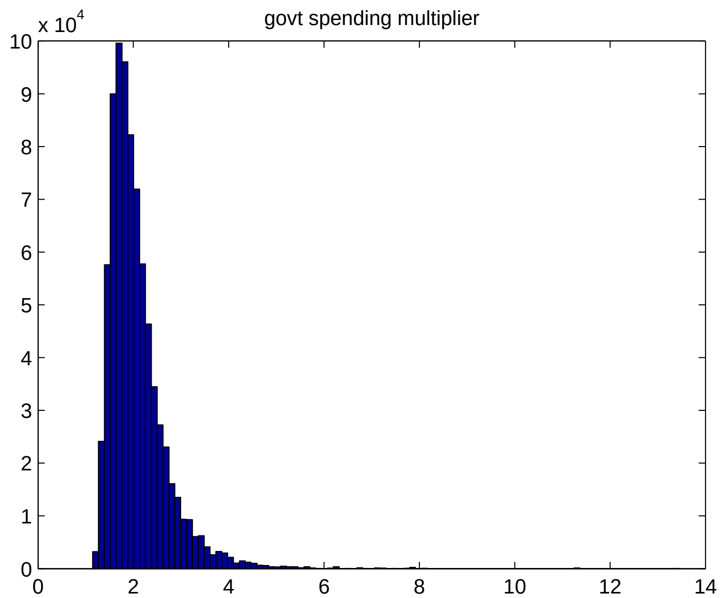

Figure 6: The posterior of the multiplier of government spending.

10 Sensitivity analysis

Figure 6 shows the posterior distribution for the government spending multiplier. The figure reveals what we already hinted at in Table 4. The multiplier is relatively narrowly centered in the vicinity of 2. We see, however, that there are some quite large values for the multiplier, but these values get a low posterior probability.

	alpha	beta	gamma	sigmainv	omega	theta	reL	kappa
tax cut multiplier	-0.1474	0.0090	0.0532	0.3617	0.4343	-0.0556	-0.6506	0.0285
gov spending multiplier	0.0083	-0.0240	-0.3092	-0.1551	0.2399	-0.0044	0.7456	-0.2662
sales tax cut multipler	0.3167	-0.0273	-0.7536	0.5276	0.3962	0.1292	0.5002	-0.6781
capital tax cut multipler	-0.3538	0.0258	0.6580	-0.2763	-0.1050	-0.1482	-0.6952	0.6081

Table 5: The correlation between the multipliers and the variation in the parameters and the shocks.

What parameter values give rise to a large multiplier? In section 5 we explored the sensitivity of the multiplier of government spending with respect to the parameter κ, varying this parameter, while keeping the other parameters fixed. This gave us the extreme sensitivity shown in figure 4. But this figure, as we have already stressed, shows the value of the multiplier when we increase κ holding all other parameters constant, thus leading simultaneously to an explosive collapse in output and deflation. Another way of doing sensitivity analysis is to ask: How does the value of the multiplier change, if we increase κ, but still try to match the data from the Great Depression?

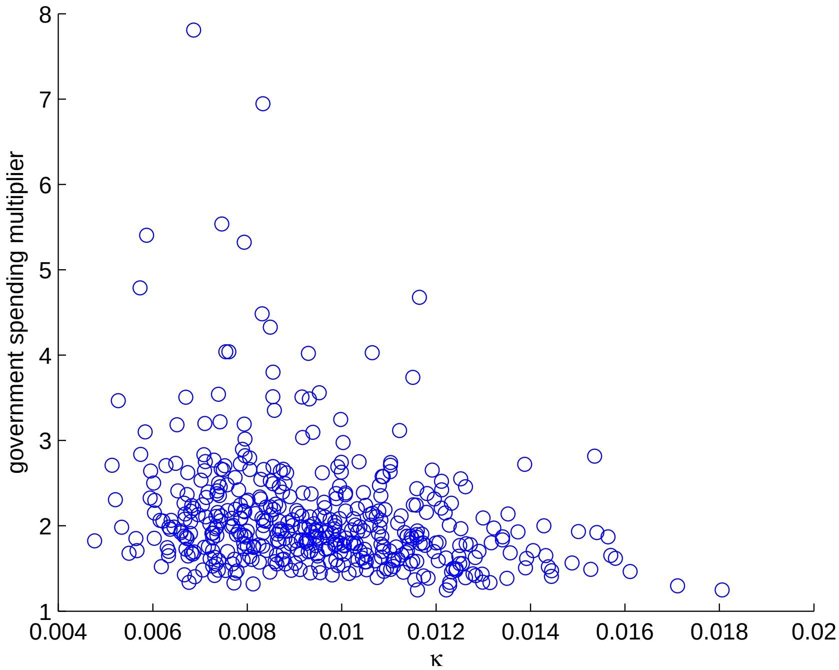

Figure 7: The value of the multiplier declines with and increase in κ.

A natural way to answer that question is to explore the correlations between κ and the multiplier of government spending in the posterior of the model. Figure 7 shows a scatter plot, plotting different values of the multiplier for different values of κ (plotting a random draw of 500 sample points from the posterior). As this figure reveals, in contrast to figure 4 in section 5, the value of the multiplier tends to decrease, the higher values of κ. This is summarized by a correlation coefficient of the full sample of the posterior in Table 5 showing a negative correlation of -0.26. The reason for this is that a different value of κ results in different values for the other parameters, and there is not any strong reason to expect this to be associated with a higher value for the multiplier. In fact, as the figure and the table shows, the opposite is true.

The table summarizes the correlation coefficient for the multipliers for the other parameters. Figure 8 shows a scatter plot for the multiplier of government spending for the parameters. Figure 9 draws out one interesting scatter plot, i.e. the correlation between r_L^e and the spending multiplier. We see that this relationship is very strong, as r_L^e become smaller, the multiplier becomes bigger. The reason for this is that as r_L gets smaller, the estimation still insists upon that the estimated parameter configuration generates a collapse in output and prices corresponding to -30 and -10 percent. The only way this can be achieved is by a parameter configuration in which the model is very "unstable", i.e. the two curves in figure 2 are very close to being parallel. In this case a "small shocks" generates a very large contraction. But it is also true that in this case a small increase in spending also has big effects, for precisely the same reason.

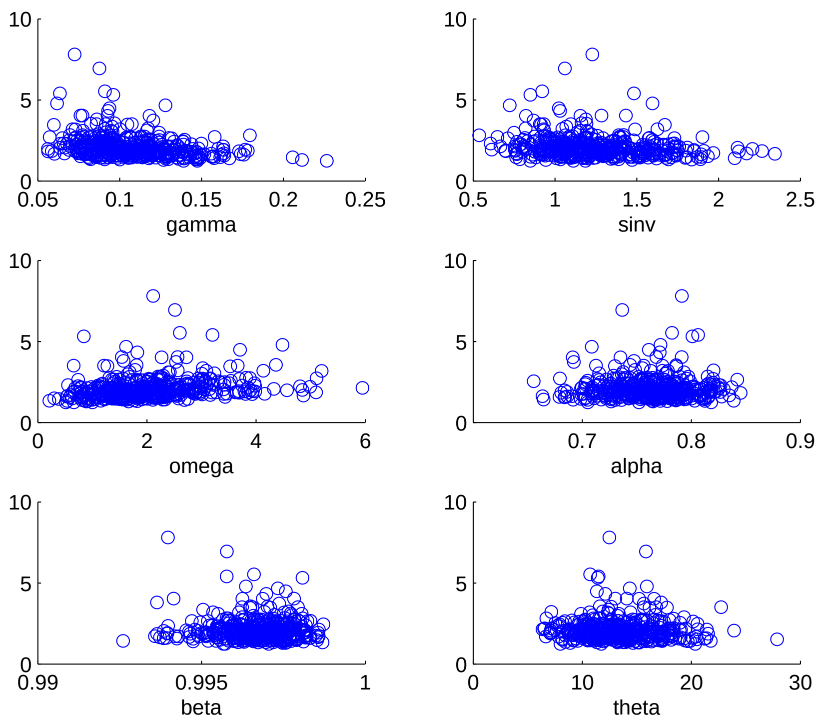

Figure 8: Scatter plot of the government spending multiplier aganist variations in the various parameters.

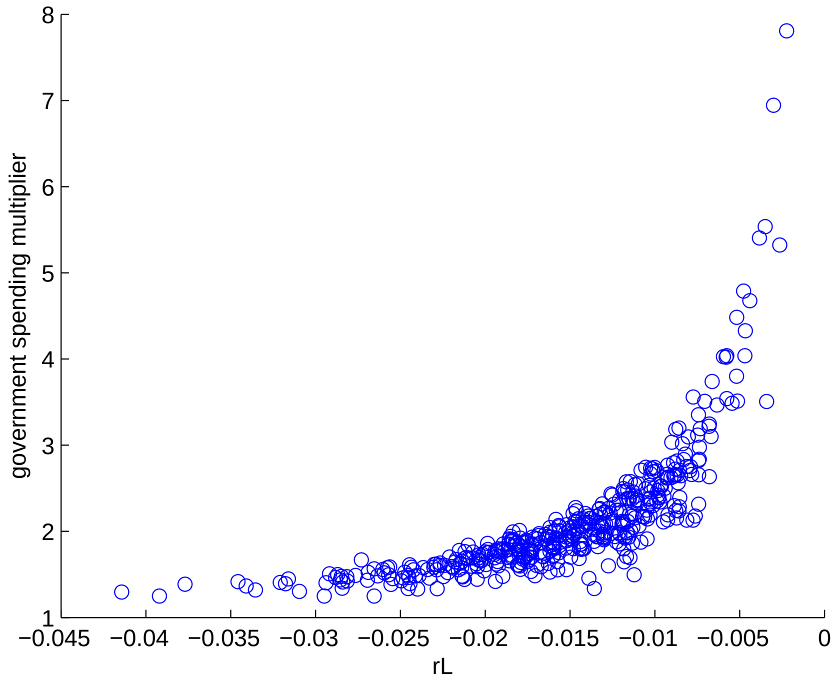

Figure 9: The lower the value of the shock the bigger the multipliier in the posterior of the model.

11 Conclusions

This paper asks a classic question: What is the effect of spending and taxes on output? The paper answers this question from a Bayesian perspective. Conditional on a matching the trough of the Great Depression we can calibrate a relatively standard model to find a quite precise estimate of the multipliers of tax and spending. We hope that the simple methodology we suggest here will be helpful to study this question in richer models.

Will the economic performance of the US economy in 2009 and 2010 possibly prove the multipliers suggested here wrong? It will be almost impossible to judge. The reason for this is that the model only has a prediction on what happens as a consequence of the fiscal stimulus relative to a counterfactual evolution of an economy in the absence of a stimulus. As it happens, this is also what makes the evaluation of the effect of past fiscal spending so hard.

12 References

An, Sungbae and Frank Schorfeide (2007), "Bayesian Analysis of DSGE Models," Econometric Reviews, 26(2-4), 113-172.

Bernstein, Jared and Christina Romer, "The Job Impact of the American Recovery and Reinvestment Plan," mimeo, Council of Economic Advisers.

Availble at http://otrans.3cdn.net/45593e8ecbd339d074_l3m6bt1te.pdf.

Christiano, Lawrence, Martin Eichenbaum, and Sergio Rebelo (2009), "When Is the Government Spending Multiplier Large," Northerwestern, mimeo.

Cogan, John, Tobias Cwik, John Taylor, and Volker Wieland (2009), “New Keynesian versus Old Keynesian Government Spending Multipliers,” mimeo, Stanford University.

Curdia, V., and Gauti Eggertsson (2009). "What Caused the Great Depression?" mimeo, New York Federal Reserve.

Eggertsson, Gauti (2004). "Monetary and Fiscal Coordination in a Liquidity Trap," Chapter 3 of Optimal Monetary and Fiscal Policy in the Liquidity Trap, Ph.D. dissertation, Princeton University, June.

Eggertsson, Gauti (2006). "Fiscal Multipliers and Policy Coordination," Federal Reserve Bank of New York Staff Report No. 241.

Eggertsson, Gauti, (2008a), "Great Expectations and the End of the Depression," American Economic Review.

Eggertsson, Gauti (2008b). "Was the New Deal Contractionary?" mimeo, New York Federal Reserve.

Eggertsson, Gauti (2009a), "What fiscal policy is effective at zero interest rates?", mimeo, Federal Reserve Bank of New York.

Eggertsson, Gauti (2009b), "The simple analytics of the government spending multiplier", mimeo, Federal Reserve Bank of New York.

Gelman, Andrew, John B. Charlin, Hal S Stern and Donald B. Rubin (2003), "Bayesian Data Analysis, Second Edition.