Macroeconomic Dynamics, **23**, 2019, 240–283. Printed in the United States of America.
doi:10.1017/S1365100516001176

FISCAL POLICY MULTIPLIERS IN AN RBC MODEL WITH LEARNING

KAUSHIK MITRA
University of Birmingham

GEORGE W. EVANS
University of Oregon
and
University of St Andrews

SEPPO HONKAPOHJA
Bank of Finland

Using the standard real business cycle model with lump-sum taxes, we analyze the impact of fiscal policy when agents form expectations using adaptive learning rather than rational expectations (RE). The output multipliers for government purchases are significantly higher under learning, and fall within empirical bounds reported in the literature, which is in sharp contrast to the implausibly low values under RE. Positive effects of fiscal policy are demonstrated during times of economic stress like the recent Great Recession. Finally, it is shown how learning can lead to consumption and investment dynamics empirically documented during some episodes of "fiscal consolidations."

Keywords: Government Purchases, Expectations, Output Multiplier, Fiscal Consolidation, Taxation

1. INTRODUCTION

There has been a recent revival of interest in the effects of fiscal policy in the wake of policy measures enacted by governments all over the world to combat the damaging effects of the "Great Recession." [1] This recent literature includes Hall (2009), Barro and Redlick (2011), Ramey (2011b), Ramey (2011a), Leeper et al. (2011), Coenen et al. (2012), and Ravn et al. (2012).

One thread running through this literature is measuring the effects of fiscal policy through examinations of government purchases multipliers in the context of exogenous changes in defense spending. An example often used in these studies is that of a war that leads to temporary increases in military expenditures. This

An earlier version of the paper was circulated as "Fiscal Policy and Learning." We are very grateful for financial support from ESRC Grant RES-062-23-2617 and from National Science Foundation Grant no. SES-1025011 without which this work could not have been carried out. Useful comments from referees and an associate editor are gratefully acknowledged. Any views expressed are those of the authors and do not necessarily reflect the views of the Bank of Finland. Address correspondence to: Seppo Honkapohja, Bank of Finland, P.O. Box 160, FI 00101 Helsinki, Finland; e-mail: seppo.honkapohja@bof.fi.

© 2017 Cambridge University Press 1365-1005/17

https://doi.org/10.1017/S1365100516001176 Published online by Cambridge University Press

interpretation is modeled by a surprise temporary increase in government purchases as emphasized in the earlier studies of Barro and King (1984) and Baxter and King (1993).

A common perception in the literature is that the standard neoclassical (Real Business Cycle or RBC) model is an inadequate model for the study of this particular policy experiment. For example, Hall (2009, p. 185) argues that the basic mechanism through which a temporary increase in government purchases works its way in the RBC model leads to the inescapable conclusion of very low output multipliers that are well outside the range found in empirical studies.[2] He argues that Keynesian or New Keynesian models with an aggregate demand channel are needed to deliver sizable government spending multipliers. It is also well known in New Keynesian models that the presence of non-Ricardian agents increases the multiplier and that the size of the multiplier is strongest when nominal interest rates are at its lower bound of zero.

The recent analyses are almost invariably developed under the "rational expectations" (RE) hypothesis. While not denying the potential importance of aggregate demand channels for changes in government spending, a question of considerable interest is the extent to which the generally small size of multipliers in the RBC model depends on RE. This question is of importance regardless of one's views concerning the role of aggregate demand channels since most dynamic macroeconomic general equilibrium models incorporate the neoclassical mechanisms that are central to the RBC model.[3] Our focus in this paper is purely on this positive issue of the size of the impact of government spending on aggregate output when agents are adaptive learners.

Thus, in this paper, we study the impact of government purchases in the standard RBC model with the sole modification that we replace RE with agents who have incomplete information about the effects of policy changes and are learning adaptively over time about these changes.[4] To make the comparison most cleanly, we use the basic RBC model with lump-sum taxes and a standard calibration.

As we have argued in Evans et al. (2009) and Mitra et al. (2013), the assumption of RE is very strong and usually unrealistic when analyzing policy changes. Economic agents need to have complete knowledge of the underlying structure, both before and after the policy change. They must also rationally and fully incorporate this knowledge in their decision making, and do so under the assumption that other agents are equally knowledgeable and equally rational. Our approach, in contrast, uses an adaptive learning model in which agents have partial structural knowledge. At each date, agents' consumption and labor supply choices depend on the time path of expected future wages, interest rates, and taxes. In line with the standard literature of adaptive learning, we assume agents' forecasts of wages and interest rates are based on a statistical model, with coefficients updated over time using least squares. However, to forecast the present value of future taxes, agents use the value implied by the announced path of future government spending under the assumption that this is announced credibly by policymakers.

https://doi.org/10.1017/S1365100516001176 Published online by Cambridge University Press

This approach seems very natural to us. The essence of the adaptive learning approach is that agents do not understand the general equilibrium considerations that govern the evolution of the central endogenous variables, so that aggregate capital, aggregate labor, and factor prices are assumed to be forecasted statistically. On the other hand, agents can be expected to immediately incorporate into their decisions the direct implications of credible announcements of the path of future government spending and taxes on their future net incomes. Under adaptive learning output dynamics within the RBC model can be quite different than under RE, as shown by Mitra et al. (2013) for cases of permanent increases in government spending.[5] Remarkably, substantial differences arise despite the fact that Ricardian equivalence can hold under both RE and learning, see Evans et al. (2012).

In this paper, our principal focus is on increases in government purchases that are known to be temporary, and in particular on the sizes of the multipliers for such policies. We find that, for the standard RBC model, output multipliers for a temporary change in government purchases can be much higher under learning than under RE, and indeed are in line with the range provided by the empirical literature.

Using this approach, the impact of fiscal policy undertaken during times of economic stress (negative shocks as in the Great Recession) is analyzed next. We model a scenario designed to capture important features of fiscal policy changes by governments to combat the Great Recession. We find that output multipliers for changes in government purchases continue to be high under adaptive learning in contrast to the values found under RE. This indicates that fiscal policy can raise output and employment in deep recessions.[6]

As a final contribution, we consider the episodes of so-called "expansionary fiscal consolidations" that have been widely studied since the contribution of Giavazzi and Pagano (1990). In the basic RBC model without distortions, a permanent reduction in government spending leads to steady-state reductions in output, so our focus here is on private sector expansion. It is known that the RBC model under RE is unable to deliver dynamics of consumption, and especially investment, matching the empirical evidence during these fiscal episodes. However, the introduction of adaptive learning can lead to short-run behavior of consumption and investment, consistent with the evidence of these episodes. Thus, we are able to provide a simple theory that can explain private sector expansion during these episodes without the need for "special theories" for large versus small changes in fiscal policy. The need for simple theories to explain these episodes has been strongly argued in Alesina et al. (2002).

Section 2 below gives a quick overview of the basic RBC model in the presence of learning by agents, and Section 3 elaborates on the learning mechanism used by agents. Section 4 analyzes the implications for multipliers of changes in government purchases. Section 5 explores robustness of these results to alternative econometric specifications of the perceived law of motion. Section 6 analyzes the effects of fiscal stimulus of the type conducted in the United States during the Great Recession. Section 7 describes how the introduction of learning in the RBC

https://doi.org/10.1017/S1365100516001176 Published online by Cambridge University Press

model can give a better match to some features of the data observed during the "fiscal consolidations." The final section concludes.

2. THE MODEL

There is a representative household who has preferences over nonnegative streams of a single consumption good c_t and leisure $1 - n_t$ given by

$$\hat{E}_t \left[\sum_{s=t}^{\infty} \beta^{s-t} U(c_s, 1 - n_s) \right], \text{ where } U(c_s, 1 - n_s) = \ln c_s + \zeta \ln(1 - n_s). \quad (1)$$

Here, $\hat{E}_t$ denotes potentially subjective expectations at time t for the future, which agents hold in the absence of RE. The analysis of the model under RE is standard. When RE is assumed, we indicate this by writing E_t for $\hat{E}_t$. Our presentation of the model is general in the sense that it applies under learning as well as under RE. The form of the utility function in (1) has been used frequently, e.g., Long and Plosser (1983).[7]

The household flow budget constraint is

$$a_{t+1} = w_t n_t + r_t a_t - c_t - \tau_{h,t}, \text{ where} \quad (2)$$

$$r_t = 1 - \delta + r_{k,t}. \quad (3)$$

Here, a_t is per capita household wealth at the beginning of time t, which equals holdings of capital k_t owned by the household less their debt (to other households), b_{pt}, i.e., $a_t \equiv k_t - b_{pt}$. r_t is the gross interest rate for loans made to other households, w_t is the wage rate, c_t is consumption, n_t is labor supply, and $\tau_{h,t}$ is per capita lump sum taxes. Equation (3) arises due to the absence of arbitrage from loans and capital being perfect substitutes as stores of value; $r_{k,t}$ is the rental rate on capital goods, and δ is the depreciation rate.

Households maximize utility (1) subject to the budget constraint (2) that yields the Euler equation for consumption:

$$c_t^{-1} = \beta \hat{E}_t r_{t+1} c_{t+1}^{-1}. \quad (4)$$

From the flow budget constraint (2), we can get the intertemporal budget constraint (in realized terms) assuming the relevant transversality condition holds:

$$0 = r_t a_t + \sum_{j=1}^{\infty} [D_{t,t+j}(t)]^{-1} \chi_{t+j} + \chi_t, \quad (5)$$

where $D_{t,t+j} = \prod_{i=1}^{j} r_{t+i}$, $j \geq 1$ and $\chi_t \equiv w_t n_t - c_t - \tau_{h,t}$.

Note that (5) involves future choices of labor supply by the household which can be eliminated to derive the linearized consumption function. For this, we make

https://doi.org/10.1017/S1365100516001176 Published online by Cambridge University Press

use of the static household first-order condition:

$$\zeta(1-n_t)^{-1} = w_t c_t^{-1}. \tag{6}$$

This relationship can be used to substitute out n_{t+j} in (5), and we can then obtain an expected value intertemporal budget constraint:

$$0 = r_t a_t + \chi_t + \sum_{j=1}^{\infty} \hat{E}_t (D_{t,t+j})^{-1} [w_{t+j} - (1+\zeta) c_{t+j} - \tau_{h,t+j}].$$

To obtain its optimal choice of consumption c_t, the household is assumed to use a consumption function based on a linearization around steady-state values. In particular, we assume agents linearize the expected value intertemporal budget constraint and the Euler equation around the initial steady-state values $\bar{c}$, $\bar{a}$, $\bar{w}$, $\bar{\tau}_h$, and $\bar{r} = \beta^{-1}$. This linearization point is natural since agents can be assumed to have estimated precisely the steady-state values before the policy change that takes place.

As shown in Mitra et al. (2013), substituting the linearized Euler equation (4) into the intertemporal budget constraint, one obtains the consumption function:

$$\begin{aligned}(c_t - \bar{c})\frac{(1+\zeta)}{(1-\beta)} &= \bar{a}(r_t - \bar{r}) + \beta^{-1}(a_t - \bar{a}) - (\tau_{h,t} - \bar{\tau}_h) \\ &\quad + (w_t - \bar{w}) - (\bar{w} - \bar{\tau}_h) S r_t^e - S\tau_{h,t}^e + S w_t^e,\end{aligned} \tag{7}$$

where

$$S r_t^e \equiv \sum_{j=1}^{\infty} \beta^{j+1} \sum_{i=1}^{j} (r_{t+i}^e - \bar{r}), \tag{8}$$

$$S\tau_{h,t}^e \equiv \sum_{j=1}^{\infty} \beta^j (\tau_{h,t+j}^e - \bar{\tau}_h), \tag{9}$$

$$S w_t^e \equiv \sum_{j=1}^{\infty} \beta^j (w_{t+j}^e - \bar{w}), \tag{10}$$

denote "present value" type expressions.

Equation (7) specifies a behavioral rule for the household's choice of current consumption based on pre-determined values of initial assets, real interest rates, wage rates, current values of lump-sum taxes, and (subjective) expectations of future values of wages, interest rates, and lump-sum taxes. Expectations are assumed to be formed at the beginning of period t and, for simplicity, we assume these to be identical across agents (though agents themselves do not know this to be the case). Equation (7) can then be viewed as the behavioral rule for per capita

https://doi.org/10.1017/S1365100516001176 Published online by Cambridge University Press

consumption in the economy with c_t, n_t, and w_t simultaneously determined given expectations.

To implement its behavioral rule, the household requires forecasts for r^e_{t+i}, w^e_{t+j}, and $\tau^e_{h,t+j}$. For taxes $\tau^e_{h,t+j}$ (and $\bar{\tau}_h$), we assume that agents use "structural" knowledge based on announced government spending rules. For convenience, we assume balanced budgets, so that $\tau_{h,t+j} = g_{t+j}$. For r^e_{t+i} and w^e_{t+j}, we assume that households estimate future values using a VAR-type model in k_t, w_t, $r_{k,t}$, and v_t, with coefficients updated over time by recursive least squares (RLS). The detailed procedure is described in Section 3.

Alternative assumptions could be made. For example, agents might forecast future taxes adaptively in the same way that they forecast wages and interest rates. We will be focusing below on announced temporary increases in government spending. If agents were to forecast future taxes purely adaptively, they would be ignoring the information given by the announced government policy. For the reasons given in Section 1, we think it is implausible that agents ignore this information. In addition, since this information is treated as central to the RE analysis of announced policy changes this would cut out a major channel by which the path of future government spending is usually assumed to affect current economic activity. Nevertheless, for completeness, we briefly consider this variation in Section 5. Other possible assumptions would be to assume that agents are unsure about the date at which the government spending increase will end, or even allow for the possibility that the increase in government spending may never end. For reasons of space, we do not pursue these latter extensions in this paper.

To complete the model, we describe the evolution of the other state variables, namely w_t, $r_{k,t}$, r_t, y_t, and k_{t+1}. Households own capital and labor services that they rent to firms. The firm uses these inputs to produce output y_t using the Cobb–Douglas production technology:

$$y_t = v_t k_t^{\alpha} n_t^{1-\alpha}, \tag{11}$$

where v_t is the technology shock that follows an AR(1) process:

$$\hat{v}_t = \rho \hat{v}_{t-1} + \tilde{u}_t \tag{12}$$

with $\hat{v}_t = (v_t - \bar{v})$. Here, $\bar{v}$ is the mean of the process, and $\tilde{u}_t$ is an independent and identically distributed (i.i.d.) zero-mean process following a normal distribution with constant variance σ_u^2.

Profit maximization by firms implies the standard first-order conditions involving wages and rental rates:

$$w_t = (1-\alpha) v_t \left(\frac{k_t}{n_t}\right)^{\alpha} \text{ and } r_{k,t} = \alpha v_t \left(\frac{n_t}{k_t}\right)^{1-\alpha}. \tag{13}$$

https://doi.org/10.1017/S1365100516001176 Published online by Cambridge University Press

In equilibrium, aggregate private debt b_{pt} is zero given our assumption of balanced government budgets. Thus, $a_t = k_t$, and market clearing determines k_{t+1} from

$$k_{t+1} = v_t k_t^{\alpha} n_t^{1-\alpha} + (1-\delta)k_t - c_t - g_t, \tag{14}$$

where g_t is per capita government spending.

For simulations of the model, we follow standard procedures and approximate the path using a linearization around the steady-state values. To analyze the impact of policy in the model, we compare the dynamics under learning to those under RE. At this stage, we remark that, as is well known, under RE and in the absence of a policy change the endogenous variables, $k_{t+1}, c_t, n_t, w_t, r_{k,t}$, and r_t can be written as an (approximate) linear function of k_t and v_t, e.g., Campbell (1994). The RE solution can be written in the form of a stationary VAR(1), in the state $\hat{x}_t' \equiv (\hat{k}_t, \hat{v}_t)$,

$$\begin{pmatrix} \hat{k}_{t+1} \\ \hat{v}_{t+1} \end{pmatrix} = B \begin{pmatrix} \hat{k}_t \\ \hat{v}_t \end{pmatrix} + \begin{pmatrix} 0 \\ 1 \end{pmatrix} \tilde{u}_{t+1}, \text{ where } B = \begin{pmatrix} \lambda_2 & f_{kv} \\ 0 & \rho \end{pmatrix}, \tag{15}$$

with the other variables given by linear combinations of the state; the hatted values are deviations from the RE deterministic steady state, i.e., $\hat{k}_t = k_t - \bar{k}$, etc. Note also that under RE forecasts of future $\hat{w}_{t+j}$ and $\hat{r}_{k,t+j}$ are given by linear combinations of the forecasted future state $\hat{x}^e_{t+j} = B^j \hat{x}_t$.

The focus of this paper is on policy changes. The method for obtaining the impact of policy changes under RE is standard, e.g., see Ljungqvist and Sargent (2012, Chap. 11) or Mitra et al. (2013) for the details. We now turn to obtaining the dynamics under learning when there is a policy change.

3. LEARNING DYNAMICS

In the standard adaptive learning approach, private agents formulate an econometric model to forecast future taxes as well as interest rates and wage rates, since these are required in order for agents to solve for their optimal level of consumption. We continue to follow this approach with respect to interest rates and wage rates, but for forecasting taxes agents are assumed to understand the future course of taxes implied by the announced policy. Agents in effect are given structural knowledge of the fiscal implications of the announced change in government purchases.[8]

As argued in Section 1, we think this is a natural way to proceed since changes in agents' own future taxes have a quantifiable direct effect, whereas future wages and interest rates are determined through dynamic general equilibrium effects. According to the adaptive learning perspective, it is unrealistic to assume that agents understand the economic structure sufficiently well to improve on reduced form econometric forecasts of aggregate variables like wages and interest rates. Thus, we assume that when a policy change is announced, agents calculate $S\tau^e_{h,t}$ using the announced changes. To keep things simple, we assume that the government operates and is known to operate under a balanced-budget rule. The assumption

https://doi.org/10.1017/S1365100516001176 Published online by Cambridge University Press

of balanced budget with lump-sum taxes is often the maintained assumption in the cited works in Section 1 for analyzing the effects of changes in government purchases on output. Additionally, with lump-sum taxes, exogenous spending, and appropriate additional assumptions, Ricardian Equivalence holds under both RE and learning, so that our results hold more generally; see, e.g., Evans et al. (2012).

The main policy change we examine in Section 4 is that of a temporary increase in (per capita) government purchases, g from $\bar{g}$ to $\bar{g}'$ for $T_g - 1$ periods, announced to take place immediately at $t = 1$, i.e.,

$$g_t = \tau_t = \begin{cases} \bar{g}', t = 1, \ldots, T_g - 1 \\ \bar{g}, t \geq T_g, \end{cases} \tag{16}$$

so that government purchases and taxes are changed in period $t = 1$, and this change is reversed at a later period T_g (this is often termed a *surprise* change in g in the literature). In our example in Section 4, we set $T_g = 9$ quarters so that we are considering a two-year increase in g.

Given their structural knowledge of the government budget constraint and the announced path of government purchases, the agents can thus compute the present value of the increase in their future taxes as

$$S\tau^e_{h,t} = \sum_{j=1}^{\infty} \beta^j (g_{t+j} - \bar{g}) = \begin{cases} \frac{\beta}{1-\beta}(\bar{g}' - \bar{g})(1 - \beta^{T_g - t - 1}), 1 \leq t \leq T_g - 2 \\ 0, \text{ for } t \geq T_g - 1. \end{cases}$$

Under learning, agents also need to form forecasts of future wages and interest rates since these are needed for their individual consumption choice in (7). Moreover, they need to form forecasts of these variables without full knowledge of the underlying model parameters. Wage and interest rate forecasts under learning depend on the perceived laws of motion (PLMs) of agents, with parameters updated over time in response to the data. We consider PLMs where, as in the stationary RE solution, future capital, wages, and rental rates depend on the current capital stock and technological shock, k_t and v_t.[9] That is, we consider PLMs of the form

$$k_{t+1} = b_k + a_{kk} k_t + a_{kv} \hat{v}_t + \text{noise}, \tag{17}$$

$$w_t = b_w + a_{wk} k_t + a_{wv} \hat{v}_t + \text{noise}, \tag{18}$$

$$r_{k,t} = b_r + a_{rk} k_t + a_{rv} \hat{v}_t + \text{noise}, \tag{19}$$

$$\hat{v}_t = \rho \hat{v}_{t-1} + \tilde{u}_t, \tag{20}$$

where the PLM parameters b_k, a_{kk}, a_{kv}, etc. will be estimated on the basis of actual data. The final line is the stochastic process for evolution of the (de-meaned) technological shock, which for simplicity is assumed known to the agents. In real-time learning, the parameters in (17)–(19) are estimated, and therefore time-dependent, and are updated using recursive least squares (RLS); see, e.g., Evans and Honkapohja (2001, p. 233). We also assume agents allow for structural change,

https://doi.org/10.1017/S1365100516001176 Published online by Cambridge University Press

which includes policy changes as well as other potential structural breaks, by discounting older data as discussed below.

In postulating that agents forecast using the PLM (17)–(20), we are implicitly assuming that they do not have useful information from previous policy changes. We think this is generally plausible since policy changes are relatively infrequent and since the qualitative and quantitative details of previous policy changes are unlikely to be the same. In particular, previous fiscal policy changes (if any), of the type considered in this paper, are likely to have varied in terms of the magnitude and duration of the change in government spending, and the state of the economy in which it was announced and implemented. Since older information of this type would probably have limited value, we assume that agents respond to policy change by updating the parameters of the PLM (17)–(19) as new data become available.[10]

Before discussing how the PLM coefficients are updated over time using least-squares learning, we describe how (17)–(19) are used by agents to make forecasts. Given coefficient estimates and the observed state $(k_t, \hat{v}_t)$, equations (17) and (20) can be iterated forward to obtain forecasts k^e_{t+j} and $\hat{v}_{t+j}$ for $j = 1, 2, \ldots$. Wage and rental rate forecasts w^e_{t+j} and $r^e_{k,t+j}$ are then obtained using the relationships (18) and (19), respectively, whereas interest-rate forecasts are given by $r^e_{t+j} = 1 - \delta + r^e_{k,t+j}$. Given these forecasts, Sw^e_t and Sr^e_t are computed from (10) and (8), respectively, which in turn are used in (7) in determining consumption in the temporary equilibrium. See the appendix of Mitra et al. (2013) for further details.

Parameter updating by agents using RLS learning is as follows. We define the time t parameter estimates as

$$\phi_{k,t} = \begin{pmatrix} b_{k,t} \\ a_{kk,t} \\ a_{kv,t} \end{pmatrix}, \phi_{w,t} = \begin{pmatrix} b_{w,t} \\ a_{wk,t} \\ a_{wv,t} \end{pmatrix}, \phi_{rk,t} = \begin{pmatrix} b_{r,t} \\ a_{rk,t} \\ a_{rv,t} \end{pmatrix}, z_t = \begin{pmatrix} 1 \\ k_t \\ \hat{v}_t \end{pmatrix}.$$

The RLS formulas corresponding to estimates of equations (17)–(19) are

$$\begin{aligned} \phi_{k,t} &= \phi_{k,t-1} + \gamma R_t^{-1} z_{t-1}(k_t - \phi'_{k,t-1} z_{t-1}), \\ \phi_{w,t} &= \phi_{w,t-1} + \gamma R_t^{-1} z_{t-1}(w_{t-1} - \phi'_{w,t-1} z_{t-1}), \\ \phi_{rk,t} &= \phi_{rk,t-1} + \gamma R_t^{-1} z_{t-1}(r_{k,t-1} - \phi'_{rk,t-1} z_{t-1}). \\ R_t &= R_{t-1} + \gamma(z_{t-1} z'_{t-1} - R_{t-1}). \end{aligned}$$

Note that $\phi_{k,t}$, $\phi_{w,t}$, and $\phi_{rk,t}$, which are used in (17)–(19) to make time t forecasts, are computed using data available at the end of time $t - 1$. The initial values of all parameter estimates ϕ and R are set to the initial steady-state values under RE.

Here, it is assumed that agents update parameter estimates using "discounted least squares," i.e., they discount past data geometrically at rate $1 - \gamma$, where $0 < \gamma < 1$ is typically a small positive number. In the learning literature, the

https://doi.org/10.1017/S1365100516001176 Published online by Cambridge University Press

parameter γ is known as the "gain," and discounted least squares is also called "constant-gain" least squares. For simplicity, the gain is assumed to be the same in all the regressions.

Constant-gain least squares is widely used in the adaptive learning literature because at t it weights recent data more heavily than older data. See, for example, Sargent (1999), Orphanides and Williams (2007), Carceles-Poveda and Giannitsarou (2008), and Eusepi and Preston (2011). In the current context, constant gain is particularly appealing since agents will be aware that policy changes will induce changes in forecast-rule parameter values taking a possibly complex and time-varying form. The use of a constant-gain rule allows parameter estimates to track changes in parameter values more quickly than does "decreasing-gain" least squares.

4. MULTIPLIERS FOR GOVERNMENT PURCHASES

In this section, we examine the effects of a temporary change in g. Our general aim is to compare the dynamics obtained under RE and adaptive learning, focusing on the multiplier for output to see the effects of such a policy. We assume that the economy is initially in the steady state corresponding to $g = \bar{g}$, and the temporary increase in g is assumed to be fully credible and announced at the start of period 1, taking the particular form given in equation (16). An example that is often used is a war that leads to a temporary increase in military expenditures, e.g., see Hall (2009), Barro and Redlick (2011), and Ramey (2011a,b).

Figure 1 compares the dynamics under RE and learning for key variables. The variables plotted are capital (k_t), gross investment ($i_t = k_{t+1} - (1 - \delta)k_t$), consumption ($c_t$), labor ($n_t$), output ($y_t$), and wages ($w_t$). All variables are measured in percentage deviations from the (unchanged) steady state. In period $t = 0$, all variables are in the steady state. We assume the following parametric form for the figures: $\zeta = 4$, $\delta = 0.025$, $\alpha = 1/3$, $\beta = 0.985$, $\rho = 0.95$, $\bar{v} = 1.359$, $\bar{g} = 0.20$, and $\gamma = 0.04$ in the learning rule. These parameter values conform to the ones used in the RBC literature, see, e.g., King and Rebelo (1999) or Heijdra (2009). To aid interpretation, $\bar{v} = 1.359$ is chosen to normalize output to (approximately) one, specifically $\bar{y} = 1.00057$. The government spending/output ratio is 21%, that of investment/output ratio is 20%, and that of consumption/output ratio is 59%. $\tilde{u}_t$ is assumed to be distributed normally with zero mean and standard deviation $\sigma_u = 0.007$, which is in line with the value used in this literature.[11]

Our choice of the gain parameter $\gamma = 0.04$ is in line with most of the literature, e.g., Branch and Evans (2006), Orphanides and Williams (2007), and Milani (2007). Eusepi and Preston (2011) use a much smaller value for the gain, but they do not consider changes in policy, for which a larger value of γ is more appropriate.[12]

For the policy exercises, there is an increase in government purchases from $\bar{g} = 0.20$ to $\bar{g}' = 0.21$ (a 5% increase) that takes place at $t = 1$, and lasts until $T_g = 9$, i.e., for eight quarters (e.g., a two-year war) in equation (16). We plot

https://doi.org/10.1017/S1365100516001176 Published online by Cambridge University Press

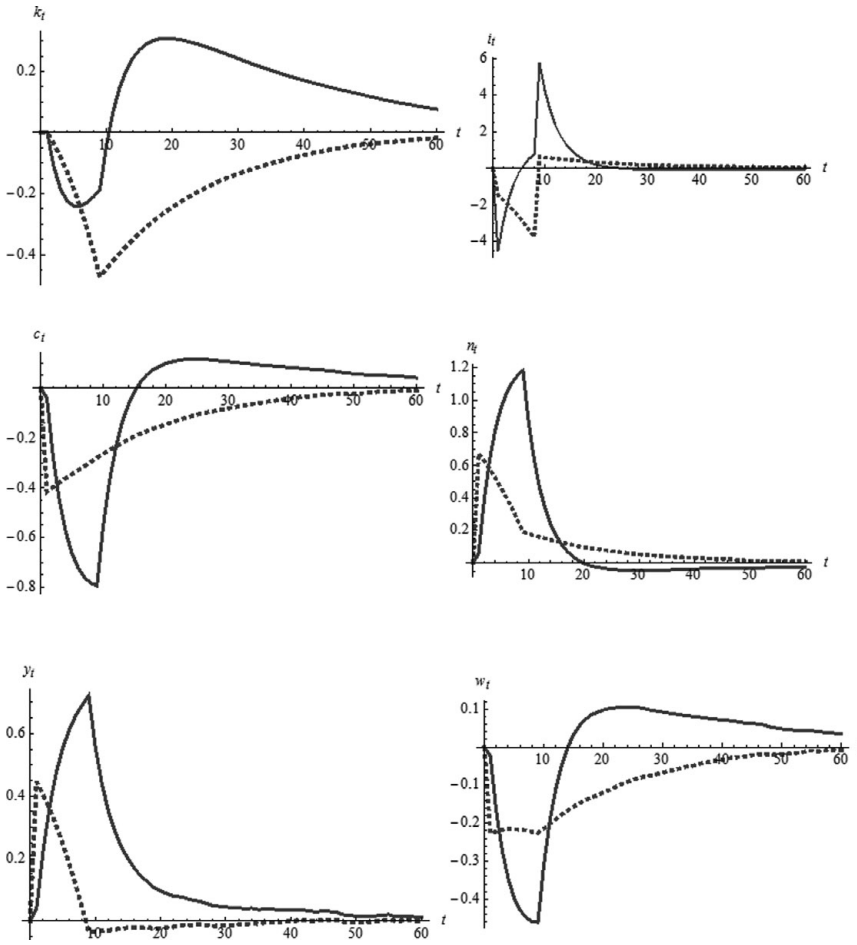


FIGURE 1. Dynamic paths for a two-year (eight periods) increase in government purchases. The solid lines are the learning paths, whereas the dashed lines are the RE paths. All variables are measured in percentage deviations from steady-state values. Mean paths over 100,000 simulations.

the mean time paths for each endogenous variable over 100,000 replications in Figure 1.

Under RE, the dynamics are well understood, see Baxter and King (1993) and Mitra et al. (2013) for details. k_t falls as long as the policy change is in effect and then increases toward the (unchanged) steady state. c_t falls on impact and then increases monotonically toward the steady state. An important feature of a temporary increase in g is that consumption smoothing by agents is achieved by a *reduction* in investment i_t. The small wealth effect due to a temporary change, as opposed to a permanent change in g, leads to small impact effects on c_t, n_t,

https://doi.org/10.1017/S1365100516001176 Published online by Cambridge University Press

and y_t. The k_t/n_t ratio falls on impact, which raises r_t and lowers w_t on impact. w_t continues to be low during the period of high g, and this reduces n_t over time. People maintain a rising path of c_t by reducing i_t as long as the period of increased g lasts, which also results in a falling path of y_t over time. Once the period of high g is over, a rising path of c_t can be maintained without the need to reduce capital, and there is an investment boom at this point and k_t starts increasing toward the steady state. The k_t/n_t ratio starts rising, which lowers r_t (raises w_t), leading to further declines in n_t as it converges toward the steady state.

Consider now the impacts of the policy under learning. The most marked difference under learning compared to RE is the sharper fall in investment i_t on impact. Under RE, agents foresee the path of low wages (high interest rates) in the future which reduces initial consumption more on impact compared to learning. With expectations of future wages and interest rates pre-determined, and only a small rise in $S\tau^e_{h,t}$ (due to the temporary change), the reduction in consumption at $t = 1$ is much smaller under learning than under RE. (The impact effects on other variables are also muted under learning for the same reason.) Consequently, there is a sharp fall in i_t with k_t run down rapidly. The sizable negative impact effect of c_t under RE, followed by a steady return to steady state, is sometimes viewed as implausible. In contrast under learning the response over the first five years is hump-shaped, followed by some overshooting and eventual convergence. This hump-shaped response is also seen in y_t and n_t.

Under learning, although agents correctly foresee the period of higher taxes, they fail to appreciate the precise form of the wage and price dynamics that result from the policy change. The reduction in k_t over $t = 1, \ldots, T_g - 1$ leads to lower wages and expected wages, Sw^e_t, and higher interest rates and expected interest rates, Sr^e_t, resulting in a period of excessive pessimism during the period of high g. The resulting reduction in c_t and increase in n_t during this period reverse the fall in i_t and stabilize k_t in excess of RE levels. When the period of high g ends at $T_g = 9$, the planned reduction in g leads to a sharp spike in i_t and build-up of k_t. This leads to a period of higher wages and expected wages, and lower actual and expected interest rates, and thus to an extended period of correction to the earlier period of overpessimism, before eventual convergence back to the steady state.

One way to view these results is that agents fail to foresee the full impacts of the crowding out or crowding in of capital from government purchases. In the present case, agents tend to extrapolate the low wages during the period of increased purchases, which result from the run-down of capital. Although agents understand that their future taxes will fall when the war ends, they fail to recognize the improvement in wages that will occur after the crowding in of capital after the war. This is the source of the excessive pessimism during the war, with a resulting correction after the war ends.

We turn now to a comparison of the government purchase multipliers under RE and learning. As argued by several authors, e.g., Hall (2009), the multipliers obtained in RBC models under RE are too small to be consistent with the data. Hall notes that the US evidence from World War II (WWII) and the Korean wars

https://doi.org/10.1017/S1365100516001176 Published online by Cambridge University Press

suggests multipliers for gross domestic product (GDP) in the 0.7–1.0 range. Ramey (2011a) concludes that for deficit-financed increases in purchases a range of 0.8–1.5 is likely; similarly, Ramey (2011b) finds the "implied government spending multipliers range from 0.6 to 1.2." [13]

The general view is that output multipliers in RBC models are very small, and unlikely to be consistent with these values. As emphasized, e.g., by Leeper et al. (2011), Keynesian elements need to be included in the model to obtain an aggregate demand channel and realistic multipliers. An issue that has not received attention is the potential role for adaptive learning to provide an additional channel for the multiplier within the standard RBC model. We now take up this issue.

Figure 2 shows the results for the output, investment, and consumption multipliers for the policy experiment displayed in Figure 1. In each case, we show both the multipliers viewed as a distributed lag response and the cumulative multiplier over time. For each graph within Figure 2, the RE and learning responses are shown. The cumulative multipliers are computed as a discounted sum using the discount factor β. Specifically, for the output multipliers, we compute

$$ym_t = \frac{y_t - \bar{y}}{\bar{g}' - \bar{g}} \text{ and } ycm_t = \frac{\sum_{i=1}^{t} \beta^{i-1}(y_i - \bar{y})}{(\bar{g}' - \bar{g}) \sum_{i=1}^{T_g - 1} \beta^{i-1}}, \text{ for } t = 1, 2, 3, \ldots,$$

with analogous formulas for the investment and consumption multipliers. We use discounting to ensure that, e.g., small persistent values of $y_i - \bar{y}$ do not receive undue weight. Note that for $t \geq T_g - 1$, the (discounted) cumulative output multiplier equals one plus the cumulative consumption multiplier plus the cumulative investment multiplier.

The output multipliers are particularly striking. Although the impact multiplier is larger under RE than under learning, by quarter 5 the learning multiplier is larger than the RE multiplier and by quarter 8 the RE multiplier is near zero, where it remains, whereas the learning multiplier has increased substantially, reaching a peak of over 0.7 in quarter 10. The difference in multiplier effects is captured well by the (discounted) cumulative multiplier, which over five years is more than 0.8 under learning but less than 0.25 under RE. In fact, in the final period of the figure (year 15), the cumulative output multiplier is 0.94 under learning and only 0.22 under RE. Strikingly, the output multipliers obtained under learning are in line with the empirical evidence cited above.

What accounts for the much larger output multiplier under learning compared to RE? This can be seen from the consumption and investment multipliers. Under both RE and learning, the higher g crowds out consumption, but there is a hump-shaped response under learning, which declines until quarter 10. In fact the consumption multiplier eventually (from $t = 16$) turns positive, and the long-run cumulative consumption multiplier is substantially less negative under learning than under RE. In the final period of the figure, the cumulative consumption multiplier

https://doi.org/10.1017/S1365100516001176 Published online by Cambridge University Press

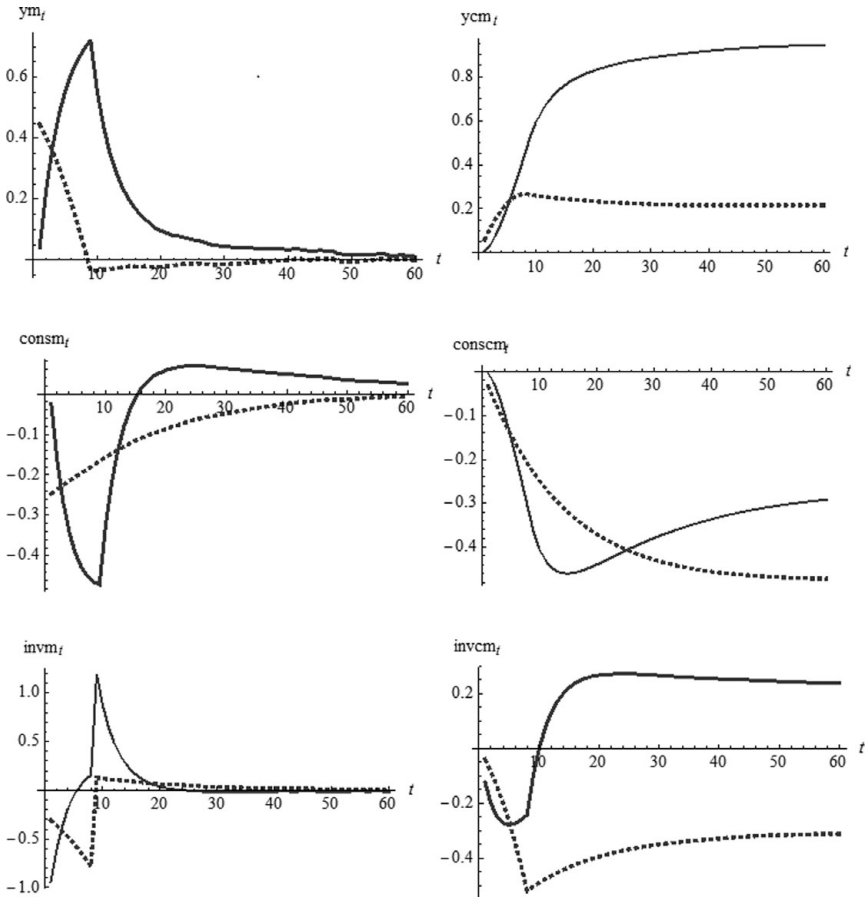


FIGURE 2. Multipliers as a distributed lag response (left-hand side) and cumulative multipliers (right-hand side), for output, consumption, and investment, for the increase in government purchases considered in Figure 1. The solid lines are the learning paths, whereas the dashed lines are the RE paths.

is −0.29 under learning and −0.47 under RE. That is, overall there is significantly less crowding out of consumption under learning than under RE.[14]

The biggest difference is, however, in the behavior of the investment multipliers. As discussed earlier, the negative impact effect on investment is larger under learning than under RE, but this quickly reverses and by quarter 6 the impact on investment is positive under learning and substantially negative under RE. The cumulative investment multipliers after five years are over 0.25 under learning and about −0.4 under RE. Thus, under RE, the overall small cumulative output multiplier reflects crowding out of investment as well as consumption, whereas the

https://doi.org/10.1017/S1365100516001176 Published online by Cambridge University Press

longer-run cumulative output multipliers under learning of over 0.94 reflect much less crowding out of consumption and substantial *crowding in* of investment.

We briefly discuss the robustness of our results to different choices of the gain parameter. Use of a higher constant gain parameter seems to result in higher output multipliers; for example, in year 15, this multiplier is 0.97 with a gain of 0.1, whereas it is 0.85 and 0.70 with gains of 0.02 and 0.01, respectively. On the other hand, if agents use a constant gain during the policy change and then switch to a decreasing gain, the cumulative output multiplier can even exceed one. For example, use of the baseline gain of $1/25$ for $t \leq T_g - 1$ and $(25 + t - T_g + 1)^{-1}$ for $t \geq T_g$ results in a cumulative output multiplier of 1.05 in year 15, whereas the corresponding multiplier rises to 1.10 if these gains are replaced by $1/10$ and $(10 + t - T_g + 1)^{-1}$, respectively.

We remark that adaptive learning can shed some light on the controversial issue of the qualitative response of consumption to a rise in government purchases. As noted by Ramey (2011b), some empirical studies find negative responses of private consumption, in the short to medium term, whereas others find positive responses. Under RE, it is well known that the consumption multiplier is quite negative in the RBC model as in Figure 2. Hall (2009, p. 198) puts it forcefully: "The model is fundamentally inconsistent with increasing and constant consumption when government purchases rise." Our study indicates that under learning the distributed lag response of consumption in the RBC model can eventually become positive (from quarter 16 onwards in Figure 2). Under learning, we have both a negative consumption response in the short to medium term and a positive response thereafter.

Many authors have demonstrated that the purely neoclassical (RBC) model has no potential to produce realistic output multipliers, because of the significant crowding out of consumption and investment, and that in order to get acceptable output multipliers consistent with the empirical evidence, one has to turn to models that blend neoclassical and Keynesian elements. Even if one accepts that New Keynesian features are part of a realistic mechanism by which government purchases affect output, it is useful to understand how large the multiplier can potentially be in RBC models as some of the micro-foundations are common in neoclassical and New Keynesian models. Our principal finding is that the introduction of adaptive learning to the RBC model can by itself rectify the apparent inability of this model to fit the evidence on output multipliers. RBC models with learning are capable of delivering higher multipliers and indeed are even within the range found in empirical studies.

5. VARIATIONS

In this section, we consider three variations of our technique that give alternative estimates of multipliers under learning. This allows us to asses the robustness of our earlier results.

https://doi.org/10.1017/S1365100516001176 Published online by Cambridge University Press

5.1. Adaptive Learning of Taxes

In the results from the preceding section, it was assumed that agents understood the tax path that would be followed during and after the stimulus. Although we feel this is plausible in the context of a credible fiscal policy, it may be of interest to see how the results will change if expectations about future taxes are formed using adaptive learning. Because taxes do not depend upon the capital stock or productivity levels, we assume a simple learning rule in which the agents update their estimate of the expected future taxes based on the observed actual taxes. Specifically, $\tau^e_{h,t+j} = b_{\tau,t}$ for $j \geq 1$, where

$$b_{\tau,t} = b_{\tau,t-1} + \gamma(g_t - b_{\tau,t-1}).$$

It follows that

$$S\tau^e_{h,t} = \frac{\beta}{1-\beta}(b_{\tau,t} - \bar{g}),$$

where $\bar{g} = \bar{\tau}_h$ is the initial level of steady-state taxes.

We have re-run the simulations under this assumption, and the results are quite similar both qualitatively, in terms of the paths of the variables, and quantitatively, in terms of the sizes of the responses. In particular, the peak of the distributed lag output multipliers occurs in period 9 and is equal to 0.770. The cumulative multiplier in period 10 is 0.770, in period 20 is 0.888, in period 40 is 1.007, and in period 60 is 1.034.[15] This compares to the baseline results in which the peak of the distributed lag output multipliers again occurs in period 9 and is 0.722 and the cumulative multiplier in period 10 is 0.593, in period 20 is 0.827, in period 40 is 0.917, and in period 60 is 0.945.[16] Thus, with adaptive learning of taxes, there is a slight strengthening of our earlier results in which the path of taxes is correctly foreseen.

5.2. Reversion to RE Coefficients at End of Policy

The results in Section 4 assume that agents continue to learn, i.e., to update the parameters of their forecast rules, after the end of the fiscal stimulus. An alternative assumption might be that agents revert to the original RE parameter values corresponding to the initial steady state to which the economy will again converge. The argument for this would be that since the stimulus is now over, there is no reason for the dynamics of the economy and hence the parameters of the forecast rule to be different from what they were before the policy was implemented. Against this agents might wonder if the stimulus will have a persistent effect on the dynamics of the economy that outlasts the stimulus. This might seem plausible given that the level of the capital stock is quite far from its steady-state value at the end of the stimulus. Continuing to update parameter estimates also seems to us likely if the agents have had no prior experience of a policy stimulus of comparable size and duration.

https://doi.org/10.1017/S1365100516001176 Published online by Cambridge University Press

Since both arguments have plausibility, we consider the implications of agents using forecast rules in the poststimulus period that are a weighted average of the original RE parameters and those that result from continued constant gain learning. Thus, we assume that in the poststimulus period agents use forecast rule coefficients:

$$\tilde{\phi}_{j,t} = \eta\bar{\phi}_{j,RE} + (1-\eta)\phi_{j,t},$$

where $j = k, w, rk$ and $0 \leq \eta \leq 1$. Here, $\bar{\phi}_{j,RE}$ denotes the RE parameter values at the steady state, and $\phi_{j,t}$ denotes the parameter values obtained from constant gain learning when agents are using these forecast rules.

When $\eta = 0$, we get the results of Section 4. These baseline results for output multipliers were summarized immediately above. In the polar opposite case of $\eta = 1$, the peak of the distributed lag output multiplier occurs in period 8 and takes the value 0.683.[17] The cumulative output multipliers in periods 10, 20, 40, and 60 are respectively 0.431, 0.421, 0.436, and 0.440. It can be seen that the effect of reverting to RE coefficients at the end of the stimulus is to largely eliminate additional multiplier effects in the postpolicy period. Even in this case the cumulative multipliers under learning are about twice the values that arise in the rational expectations equilibrium (REE).

For values $0 < \eta < 1$, we naturally get intermediate values of the multipliers. For $\eta = 0.8$, the peak of the distributed lag output multiplier occurs in period 8 and takes the value 0.677. The cumulative output multipliers in periods 10, 20, 40, and 60 are respectively 0.462, 0.494, 0.527, and 0.535. For $\eta = 0.5$, the peak of the distributed lag output multiplier occurs in period 8 and takes the value 0.683. The cumulative output multipliers in periods 10, 20, 40, and 60 are respectively 0.503, 0.611, 0.670, and 0.686.

5.3. Dummy Variables for Policy Regime

An alternative method by which agents could allow for an abrupt change in parameter estimates at the end of the policy period would be to include a suitable dummy variable in the forecasting equations. In particular, suppose that agents model the coefficients of the estimated law of motion for capital, wages, and rental rates as shifting discretely during the period of the policy change. To do this, they include a dummy variable in the equations (17)–(19) that takes the value $\bar{g}' - \bar{g} > 0$ during the period of high government spending and zero once the policy change is removed.

For this procedure, the crucial issue is how to initialize the coefficient of the dummy variable in each of the forecasting equations. One possibility would be to assume that there are also temporary government spending shocks that follow an AR(1) zero mean process. As a result, agents have experience with fluctuations in government spending before the policy change takes place. Earlier we argued against the use of this information by agents on the grounds that such a policy would be sufficiently different from past experience to make earlier fluctuations

https://doi.org/10.1017/S1365100516001176 Published online by Cambridge University Press

in g_t of limited information value. However, it is possible that agents make at least some use of the past effects in forecasting the effects of the new policy.

Thus, we now assume that agents have the following PLM, in which $\hat{g}_t$ is included as an additional state variable:

$$\begin{aligned} k_{t+1} &= b_k + a_{kk}k_t + a_{kv}\hat{v}_t + a_{kg}\hat{g}_t + a_{kd}D\bar{g}_t + \text{noise}, \\ w_t &= b_w + a_{wk}k_t + a_{wv}\hat{v}_t + a_{wg}\hat{g}_t + a_{wd}D\bar{g}_t + \text{noise}, \\ r_{k,t} &= b_r + a_{rk}k_t + a_{rv}\hat{v}_t + a_{rg}\hat{g}_t + a_{rd}D\bar{g}_t + \text{noise}, \\ \hat{v}_t &= \rho\hat{v}_{t-1} + \tilde{u}_t, \text{ and } \hat{g}_t = \mu\hat{g}_{t-1} + \varepsilon_t, \text{ where } 0 < \rho, \mu < 1. \end{aligned}$$

Here, the government spending process is given by

$$g_t = \bar{g} + D\bar{g}_t + \hat{g}_t.$$

The variable $D\bar{g}_t$ is a discrete variable taking the value $D\bar{g}_t = \bar{g}' - \bar{g} \equiv \Delta\bar{g}$ for $1 \leq t \leq T_g - 1$ and zero thereafter. The RLS formulas need to be modified suitably. See the appendix for details.

We now discuss how to initialize the parameter estimates on the dummy variable at $t = 1$.[18] One possibility is that agents believe they have no information on these values and set their initial estimates at zero. However, if we assume that coming into $t = 1$, agents' estimates of the PLM have converged to the RE values $(\bar{a}_{kg}, \bar{a}_{wg}, \bar{a}_{rg})$, corresponding to the stationary environment in which $g_t = \bar{g} + \hat{g}_t$; another possibility is that they treat the RE coefficients on $\hat{g}_t$ as useful information. More generally, and more plausibly, agents might use

$$a_{kd} = (1-\omega)\bar{a}_{kg},\ a_{wd} = (1-\omega)\bar{a}_{wg},\ \text{and } a_{rd} = (1-\omega)\bar{a}_{rg}$$

for some shrinkage parameter $0 \leq \omega \leq 1$. The parameter ω measures the distrust agents place on the relevance of temporary spending shock for the policy change.

In contrast to Section 4, the inclusion of the dummy variable $D\bar{g}_t$, with initial nonzero parameter estimates, implies that agents immediately project effects on capital, wages, and interest rates during the period of the fiscal stimulus. Furthermore, the inclusion of $D\bar{g}_t$ means that agents anticipate from $t = 1$ a discrete change in expected future capital, wages, and interest rates when the policy stimulus is ended. Against this, however, to the extent that initial estimates of the parameters on $D\bar{g}_t$ do use the experience from temporary changes $\hat{g}_t$, this may provide poor guidance for the effects of a persistent fiscal stimulus of the type considered here.

Table 1 presents the results for the cumulative output multipliers in period 40 under learning for a range of μ and ω. The corresponding cumulative multipliers under RE are around 0.18 for all the cases reported in Table 1. Although it is clear that the results depend on both μ and ω, unless both μ and ω are low, the multipliers under learning are substantially higher than under RE. Consider first the case when $\omega = 1$. This corresponds to prior coefficients of zero on $D\bar{g}_t$ in the

https://doi.org/10.1017/S1365100516001176 Published online by Cambridge University Press

TABLE 1. Cumulative output multipliers under learning for policy stimulus using dummy variable specification

Cumulative multiplier	ω			
μ	0.1	0.5	0.9	1.0
0.3	0.09	0.37	0.65	0.72
0.7	0.16	0.41	0.66	0.73
0.9	0.32	0.51	0.70	0.74
0.95	0.41	0.55	0.69	0.73
0.99	0.60	0.65	0.71	0.72

μ is the AR(1) coefficient for $\hat{g}_t$, and ω is the shrinkage parameter for the prior coefficients of $D\bar{g}_t$. The cumulative multipliers reported are at the end of period 40 on the basis of averages over 3,000 simulations. The corresponding cumulative multipliers under RE in all cases are around 0.18.

dummy variable specification. In this case the cumulative multipliers are always above 0.7. This is somewhat smaller than in Section 4 but remains within the empirical range. The somewhat smaller cumulative multipliers result from agents over period $t = 1, \ldots, T_g - 1$, anticipating a discrete change in the dynamics starting at $t = T_g$ when the level of government spending is reduced back to the old level. For high values of ω, like $\omega = 0.9$, the results are qualitatively the same, with cumulative multipliers for $\omega = 0.9$ above 0.65.

For smaller values of $\omega < 1$, the results depend on μ, with cumulative multipliers increasing in μ. The contrast is sharpest for small values of μ and ω. When $\omega = 0.1$ and $\mu = 0.3$, the cumulative multiplier is small, as it is under RE, though the paths of the variables are quite different. Under RE, there is a large impact effect at $t = 1$ on y, n, c, w, and r from the policy change. However, from $t = 2$ onwards, these variables all begin to return toward the steady state. For the case $\omega = 0.1$ and $\mu = 0.3$, the impact effect in $T = 1$ is in the same direction but smaller than under RE. Then, because agents forecast using coefficients for $D\bar{g}_t$ based on the RE values for $\hat{g}_t$ with $\mu = 0.3$, in which households view the impacts of government spending as very temporary, expected future wages are higher than under RE and expected interest rates are lower. Thus, over the next few periods, households are overly optimistic compared to RE, so that employment and output are lower than under RE. These expectations are eventually reversed near T_g, but on aggregate the cumulative output multipliers are small.

In contrast, if $\omega = 0.9$, so that households have priors close to zero on the coefficients of $D\bar{g}_t$, then the situation is similar to the original analysis, with an erosion of wages and growing pessimism that leads to declining consumption, higher employment, and output, and an investment recovery after the initial negative impact. As in Section 4, there is also a surge in investment after the government spending is reduced to previous levels and thus the cumulative multiplier is large.

Overall, it can be seen that the basic result from Section 4, that multipliers can be much higher under learning than under RE, remains when the dummy variable

https://doi.org/10.1017/S1365100516001176 Published online by Cambridge University Press

specification is used. If either the $\hat{g}_t$ process is strongly persistent or the agents treat as relatively uninformative the previously estimated coefficients on $\hat{g}_t$, then cumulative multipliers will be much higher under learning, in particular more than 0.6, which is over three times the values under RE and is within the empirically relevant range.

One other point that should be noted is that the fiscal stimulus considered here is, in fact, strongly persistent. The policy sets $D\bar{g}_t = \bar{g}' - \bar{g} > 0$ for $t = 1, \ldots, T_g - 1$ and $D\bar{g}_t = 0$ for $t \geq T_g$. The sample AR(1) coefficient for $D\bar{g}_t$ for large t can be shown to be $r_{\bar{g}} = (T_g - 2)/(T_g - 1)$, which for $T_g = 9$ gives $r_{\bar{g}} = 7/8$. Thus, if μ is small, agents should substantially discount the parameters on $\hat{g}_t$ when forming their prior, which suggests that in this case larger values of ω would be more plausible. From Table 1, it can be seen that for $\omega \geq 0.5$, the multiplier under adaptive learning is more than twice the value obtained under RE for all values of μ shown.

6. ROBUSTNESS OF RESULTS

Before turning to the remaining applications, we briefly consider the robustness of our multiplier results to several alternative specifications. Since there is evidence that labor supply is considerably less elastic than assumed in standard RBC models, we look at the impact of a temporary fiscal stimulus for a specification with greatly reduced Frisch elasticity of labor supply. Second, we have assumed that forecast learning parameters are updated before the start of the current period. It is of interest to study how results are altered if we instead assume that these forecast parameters are updated using contemporaneous data. Finally, we have made the commonly used assumption that adequate approximations are made by linearizing all the model equations. We therefore also check the robustness of our results to using exact temporary equilibrium relations.

6.1. Elasticity of Labor Supply

Our baseline model uses the utility function $U(c, 1 - n) = \ln c + \zeta \ln(1 - n)$. For this specification, the Frisch elasticity of labor supply depends only on the steady-state ratio of leisure time to work time, and hence will be relatively high. Greater flexibility for the labor supply elasticity can be obtained by using a more general class of utility functions, such as

$$U(c, 1 - n) = \ln c + \zeta(1 - \varepsilon)^{-1}(1 - n)^{1-\varepsilon},$$

for which our baseline model corresponds to $\varepsilon = 1$.[19] It can be shown that the Frisch elasticity is given by

$$E = \varepsilon^{-1}\frac{1 - n}{n}.$$

In our baseline calibration, steady-state employment is around $\bar{n} = 0.2194$, which gives a Frisch elasticity just over 3.55.

https://doi.org/10.1017/S1365100516001176 Published online by Cambridge University Press

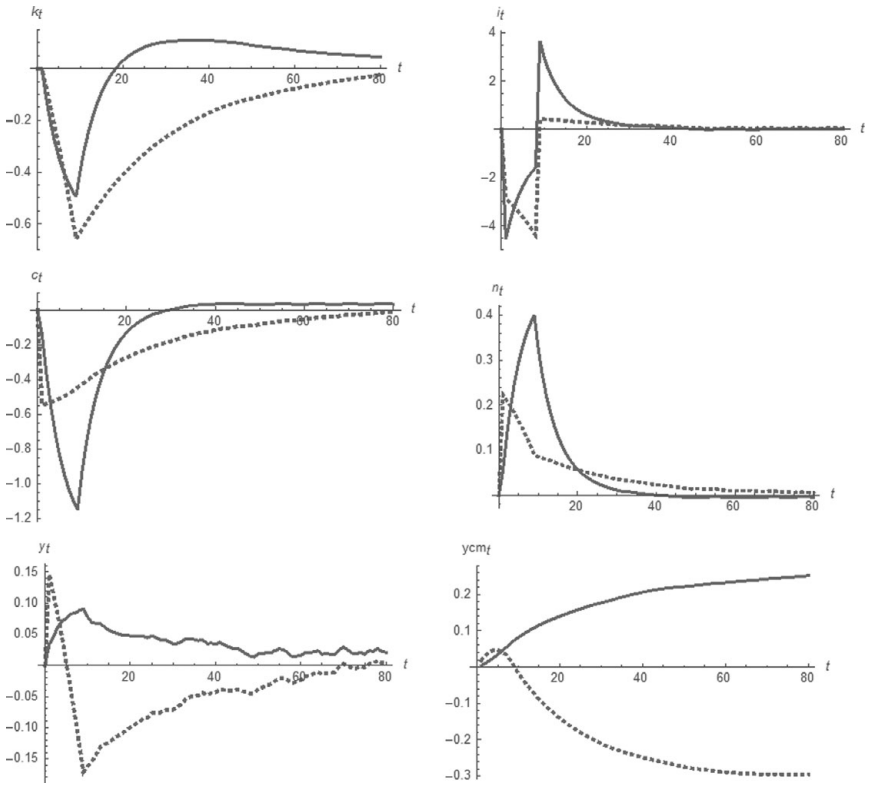


FIGURE 3. Dynamic paths for a two-year (eight periods) increase in government purchases corresponding to the simulation with Frisch elasticity of labor supply 0.475. The solid lines are the learning paths, whereas the dashed lines are the RE paths. All variables are measured in percentage deviations from steady-state values. Mean paths over 40,000 simulations. All variables are as explained before with the bottom panel displaying output (left panel) and the cumulative output multiplier (right panel).

There has been a large literature on the appropriate value of E, with controversy often centering on differences between the relatively high values typically used in the RBC literature and the much smaller values typically estimated in empirical microstudies. See, for example, the discussion in Chetty et al. (2011). Our baseline value $E = 3.55$ lies in the range used in most of the RBC literature, which requires high values to match the ratio of employment variance to output variance. The empirical microliterature typically estimates much smaller values, e.g., Frisch elasticities in the range 0.2–0.5.[20] It is therefore of interest to investigate the robustness of our results to labor supply elasticities in this range.

Figures 1 and 2 show the results for a two-year increase in government spending in our baseline case in which the Frisch elasticity of labor supply is 3.55. Figure 3 shows the results for the same policy experiment using a lower Frisch elasticity

https://doi.org/10.1017/S1365100516001176 Published online by Cambridge University Press

in line with empirical microstudies. We choose $\varepsilon = 8$, $\zeta = 0.85$, with all other parameters as in the baseline case. This corresponds to an elasticity $E = 0.475$. Figure 3 illustrates the dynamic paths in this case based on averages over 40,000 replications. There are many similarities of the paths to those in Figure 1, for both RE and learning, but also important differences in each case.

The RE paths in Figure 3 for k, i, and c are qualitatively similar to those in Figure 1. However, as would be expected with a lower E, n increases much less during the policy implementation period. This, together with the crowding out of investment and the reduction of k during the implementation period, now leads to a *reduction* in y, which is strongest near the end of the policy and continues well into the post-policy period. Consequently, under RE, there are negative output multipliers that emerge during the policy period and that continue long after the policy period. The cumulative output multiplier at period 40 is around -0.25 and at period 80 is -0.29. The sensitivity to labor elasticities of fiscal multipliers in the RBC model under RE is well known. In the extreme case of completely inelastic labor supply, there would be no increase in output at any point, and a persistent but temporary reduction in output due to the temporary crowding out of capital.

With the lower labor supply elasticity $E = 0.475$, as with the baseline higher calibration for E, the learning paths for consumption, investment, and output are strikingly different from the RE paths. Comparing the learning paths for employment and output in Figures 1 and 3, we see that although the dynamic shapes remain similar, the quantitative effects under learning are much smaller. In particular, under learning, the lower labor elasticity results in a smaller impact of policy on output and hence smaller multipliers.[21]

At the same time, the large difference in output multipliers, between the RE and adaptive learning cases, remains present in Figure 3. In particular, the cumulative output multiplier under learning rises gradually, reaching a value of 0.25 in period 80, whereas the corresponding path for RE shows an initial rise before falling continuously, reaching a value of -0.29 in period 80. The large positive difference between the learning and RE cumulative output multipliers is thus present with the current calibration, just as it was in the baseline case. The same mechanism is at work: Under learning, agents take into account the tax implications of the temporary increase in government spending, but they do not understand the precise pattern of the wage and price dynamics that result from the policy change.

We can summarize the results of our alternative calibration as follows. For the baseline calibration, with the higher labor supply elasticity calibration often used in the RBC literature, we showed in Section 4 that although output multipliers under RE are small, they are much larger under learning, and indeed are consistent with the range 0.6–1.2 often estimated in the empirical literature. When we instead use a much smaller labor supply elasticity in line with much of the empirical microliterature, the multipliers under learning are no longer large enough to match the empirical range. However, the *difference* between learning and RE remains almost as striking: With relatively inelastic labor supplies, under RE the cumulative output multipliers turn negative before the end of the policy period, whereas under

https://doi.org/10.1017/S1365100516001176 Published online by Cambridge University Press

learning output multipliers and cumulative output multipliers are consistently positive both during and following the policy period.

6.2. Learning from Contemporaneous Data

In our baseline results, it was assumed that agents used only past data in updating their estimates of parameters over time; see the parameter updating equations in Section 3. This is the typical formulation in learning models and is easy to implement. This corresponds to the main (lagged-information) timing convention examined in Marcet and Sargent (1989), and for reasons given below we think this is the most natural way to view the way agents update parameter estimates. However, as a check on the robustness of our results, we now assume that agents use contemporaneous values of wages and rental rates to update their estimates in the wage and rental rate equations. In particular, the RLS formulas are changed to "contemporaneous-data" timing as follows:

$$\begin{aligned}
\phi_{k,t} &= \phi_{k,t-1} + \gamma R_{t-1}^{-1} z_{t-1}(k_t - \phi'_{k,t-1} z_{t-1}),\\
\phi_{w,t} &= \phi_{w,t-1} + \gamma R_t^{-1} z_t(w_t - \phi'_{w,t-1} z_t),\\
\phi_{rk,t} &= \phi_{rk,t-1} + \gamma R_t^{-1} z_t(r_{k,t} - \phi'_{rk,t-1} z_t).\\
R_t &= R_{t-1} + \gamma(z_t z'_t - R_{t-1}).
\end{aligned}$$

In this formulation, the $\phi_{k,t}$ equation remains unchanged, but $\phi_{w,t}$, $\phi_{rk,t}$, and the R_t equations change so that $\phi_{w,t}$ and $\phi_{rk,t}$ reflect current observations of w_t and $r_{k,t}$, respectively. Before giving the results for this alternative timing convention, we discuss the reasons why we adopted the formulation given in Section 3.

Under adaptive learning, the economy is viewed as a sequence of temporary equilibrium outcomes, in which, at date t, the endogenous variables are jointly determined by current exogenous and pre-determined variables and by the expectations held at t of relevant future variables. The expectations at t are formed based on the agent's information set at t and on the parameter estimates ϕ_t that govern how forecasts are made given the information set. For the baseline case presented in given in Section 3, we adopt a recursive timing protocol within the temporary equilibrium, as follows. At the end of period $t-1$, parameter estimates ϕ_{t-1} are updated to ϕ_t using RLS, based on information (i.e., observed variables) in $t-1$. The variables assumed observed at the end of $t-1$ include w_{t-1}, $r_{k,t-1}$, and also k_t, which results from the time $t-1$ aggregate saving decisions. At the start of t, the productivity shock $\hat{v}_t$ is generated. Using observations at t of the predetermined variable k_t and the exogenous variables $\hat{v}_t$ and g_t, together with the updated parameter estimates ϕ_t, forecasts are made, using the estimated PLMs, for the future values of capital, wages, and rental rates. The temporary equilibrium values for the endogenous variables n_t, y_t, w_t, and $r_{k,t}$ are then jointly determined by the household decision rules for c_t, n_t, and market clearing. We think this timing protocol is attractive. The recursive timing for expectations is

https://doi.org/10.1017/S1365100516001176 Published online by Cambridge University Press

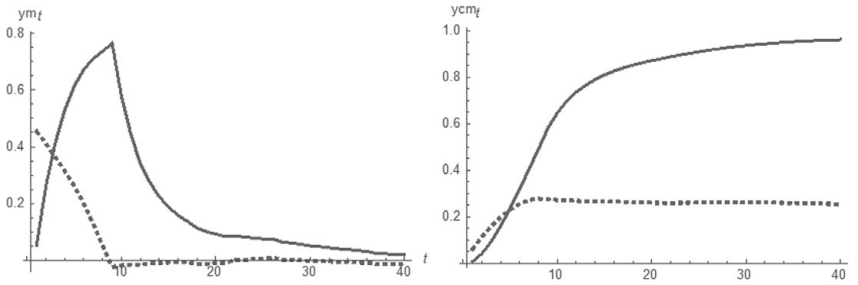


FIGURE 4. Output multiplier (distributed lag in left-hand panel and cumulative in right-hand panel) for the case when parameter estimates are updated based on contemporaneous values of wages and rental rates. The solid lines are the learning paths, whereas the dashed lines are the RE paths. Mean paths over 50,000 replications.

natural, for example, if forecasts are obtained from an econometric forecasting firm before agents go to the marketplace. It is also convenient computationally because it confines simultaneity within the temporary equilibrium to joint determination of the endogenous variables, given exogenous and predetermined variables and household and firm expectations about future variables.

Nonetheless it is useful to verify that our results are not overly sensitive to our assumed timing. We therefore also simulate the model under the alternative contemporaneous-data timing in which w_t and $r_{k,t}$ are used to update the parameter estimates $\phi_{w,t}$ and $\phi_{rk,t}$, respectively, that are in turn employed during t to forecast future variables using the PLM. We have left the $\phi_{k,t}$ equation unchanged since time t knowledge of k_{t+1} in this equation is particularly implausible. k_{t+1} is the aggregate capital stock carried into period $t + 1$ and results from the sum of the saving decisions of the individual agents in t. To assume an agent knows k_{t+1} during t amounts to assuming they have knowledge during t of the aggregate consumption and savings decisions of the other agents. We view this as both implausible and not in the spirit of the learning literature.[22]

Under our contemporaneous-data timing, there is now simultaneity under learning between current parameter estimates $\phi_{w,t}$ and $\phi_{rk,t}$ and the endogenous variables w_t, $r_{k,t}$, c_t, n_t, and y_t. This leads to additional computational complexity. Figure 4 displays the distributed lag multipliers and the cumulative output multipliers for the contemporaneous-data timing with parameters as in the baseline case of Figures 1 and 2. The multipliers in Figure 4 look very similar to the top panel of Figure 2. Thus, both qualitatively and quantitatively the multiplier results are robust to the altered timing. In addition to computational complexity and longer computer time required for simulations using contemporaneous-data timing, an additional issue with this timing is that occasionally, when the number of time periods is large, there are unstable paths that result from the simultaneity. Although we found this happens rarely with 50,000 simulations, it does make the contemporaneous-data timing assumption less attractive computationally. In any

https://doi.org/10.1017/S1365100516001176 Published online by Cambridge University Press

event, Figure 4 shows that our baseline results are not significantly affected by our timing convention.

6.3. Nonlinear Temporary Equilibrium Relations

Our analysis so far has assumed that linear approximations are satisfactory for studying the impact of fiscal policy both under RE and when agents are making decisions based on adaptive learning. Given the complexity of the agents' problem, we continue to assume that agents make their consumption decisions based on the consumption function (7). We also continue to assume that agents use linear forecasting rules based on PLMs of the form (17)–(19). However, for the remaining relationships, we use the exact nonlinear equations. In particular, we use the exact static first-order condition between consumption and labor supply (6), the production function (11), the wage rate, and the rental rate equations (13) as well as the market clearing condition that gives the capital accumulation equation (14).

Simulations using these nonlinear temporary equilibrium equations under learning are considerably slower: Using Mathematica 9, we have found this is approximately 150 times slower than working with the linear approximations. We briefly describe the differences in the adaptive learning results for the baseline case described in Section 4.

Qualitatively, the results are quite similar both in terms of the paths of the variables and the shapes of the distributed lag and cumulative multipliers. Quantitatively, we find that the results are also similar. To make the comparison as close as possible, we simulate both the linear and nonlinear systems over 40 periods for 2,000 replications using the same seed of random numbers. In the linear case, the peak of the distributed lag output multipliers occurs in period 9 and is equal to 0.677 and the cumulative multiplier in period 10 is 0.543, in period 20 is 0.752, and in period 40 is 0.890. In the nonlinear case, the peak of the distributed lag output multipliers again occurs in period 9 and is equal to 0.684 and the cumulative multiplier in period 10 is 0.548, in period 20 is 0.762, and in period 40 is 0.917.

We see that the output multipliers based on the nonlinear system are essentially the same as those based on the linear approximations used in Section 4. We have also checked for the simulations in Sections 7 and 8 that under learning the qualitative dynamics are unaffected and that there are only minor quantitative differences between the nonlinear system and the linearized system. Because of the heavy computational burden of the nonlinear setup, we primarily rely on using the linearized system in which a large number of replications are feasible.

7. FISCAL STIMULUS IN RECESSIONS

In this and the next section, we take up two applications of our analysis. For simplicity, we restrict attention to the baseline formulation that was developed in Section 4.

https://doi.org/10.1017/S1365100516001176 Published online by Cambridge University Press

Temporary increases in government spending are often motivated as policies to expand output and employment during recessions. A growing literature is reconsidering their effects owing to the large fiscal stimuli adopted in various countries in the aftermath of the Great Recession. For example, Christiano et al. (2011), Corsetti et al. (2010), and Woodford (2011) demonstrate the effectiveness of fiscal policy in models with monetary policy when the zero lower bound on nominal interest rate is reached. [For a contrary view, see Mertens and Ravn (2014).] Although the main argument for such policies relies on a demand channel, it is clearly of interest to examine the impact of a fiscal stimulus in the RBC model. We are particularly interested to know if such a policy has positive effects under learning when implemented during a severe recession.

With this in mind, we consider a situation motivated by events during the Great Recession in the United States. The NBER Business Cycle Dating Committee estimates December 2007 as the start of the recession and June 2009 as the trough, after which the economy again began to expand. Thus, the US economy was in recession during the whole of 2008 and the first half of 2009. It is widely agreed that the recession was the most severe in the United States since the Great Depression of the 1930s.

We model the above situation by assuming that the economy is initially in a steady state (corresponding to say the last quarter of 2007). We capture the main features of the Great Recession by a sequence of negative two-standard-deviation shocks to the innovation ($\tilde{u}_t$) that hit the economy for four periods in the technology equation (12), i.e., $\tilde{u}_t = -2\sigma_u$ in periods $t = 1, 2, 3, 4$. This captures the severity of the recession in 2008. This is followed by the economy being hit by negative one-standard-deviation shocks to the innovation $\tilde{u}_t$ in the next two periods (i.e., $\tilde{u}_t = -\sigma_u$ in periods $t = 5, 6$), i.e., the first half of 2009. Thereafter, from period $t \geq 7$ onwards, the evolution of the economy is governed by equation (12) with $\tilde{u}_t$ drawn from a zero mean normal distribution with variance σ_u^2 with $\sigma_u = 0.007$ as before.

Upon looking at changes in fiscal policy in this setting, it is natural to take a broad interpretation of the v_t shocks. In the traditional RBC model, the equilibrium is efficient, and the behavior of consumption and investment is an optimal response to productivity shocks. However, as emphasized in Chari et al. (2007), suitable input financing frictions are observationally equivalent to negative productivity shocks. One can thus view the sequence of negative innovations $\tilde{u}_t$ as a convenient short-cut for modeling distortions during the financial crisis that led to reduced productivity.

Features of the policy change motivated by the American Recovery and Reinvestment Act (ARRA) of February 2009[23] are captured in the model by an increase in g announced in period $t = 5$. We assume that at $t = 5$ it is announced credibly that there will be an increase in g two quarters hence from $\bar{g} = 0.2$ to $\bar{g}' = 0.21$ (a 5% hike in g, approximately 1% of GDP) for a period of two and half years, i.e., from periods $t = 7, \ldots, 16$. It is also announced that g will return to its original level of $\bar{g}$ in period $t = 17$.

https://doi.org/10.1017/S1365100516001176 Published online by Cambridge University Press

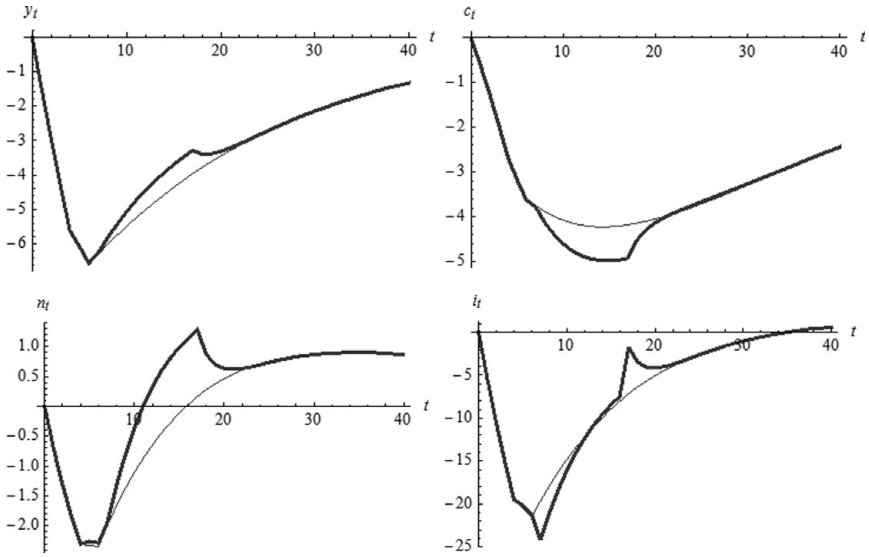


FIGURE 5. Dynamic paths showing the impact on major variables of a fiscal stimulus announced in the midst of the Great Recession. Mean paths over 20,000 simulations. The solid black line illustrates the learning paths with the policy change and the lighter shaded line the learning paths without the policy change. All variables are measured in percentage deviations from steady-state values.

The dynamics under learning are shown in Figure 5 for y_t, c_t, n_t, and i_t (the mean paths over 20,000 replications are reported).[24] The solid black line illustrates the learning paths with the policy change. We also depict the learning paths without any policy change with the lighter shaded line. Of course, there are no differences in the dynamics of the two economies for the first year until the policy change is announced at $t = 5$. The severity of the recession during the first year means that y_t has fallen by -5.61% as of $t = 4$. Once the policy change is announced at $t = 5$, the dynamics of the two economies start to differ, though the effect on y_t and c_t for the first few periods is small.

The impact of the policy builds up steadily after the policy change comes into effect at $t = 7$. y_t rises over time and is approximately 0.68% points higher at $t = 17$. The differences in dynamics start getting smaller from $t = 25$ onwards but y_t continues to be significantly higher with the policy change for 5 years and stays above the no-policy path throughout the 10-year period plotted in Figure 5. Employment n_t also gets a substantial boost during the time of higher g and in fact is above the steady state from period 11 onwards. The boost in n_t and the lower levels of c_t during the time of higher g help explain the significant expansionary effects of the fiscal policy under learning.[25]

We also plot the corresponding output multipliers for this policy experiment in Figure 6. The left-hand panel shows the distributed lag multiplier and the

https://doi.org/10.1017/S1365100516001176 Published online by Cambridge University Press

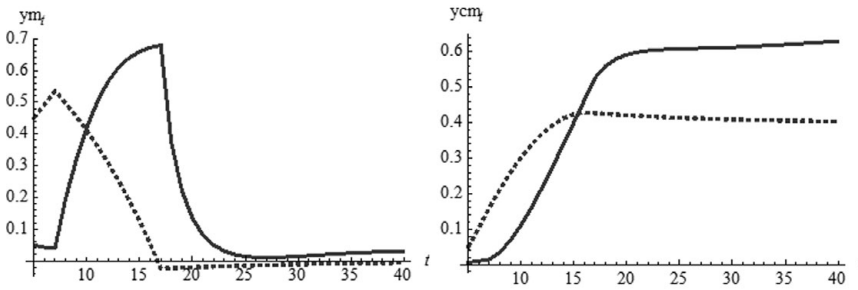


FIGURE 6. Output multiplier (distributed lag in left-hand panel and cumulative in right-hand panel) for the policy experiment illustrated in Figure 5. The solid lines are the learning paths, whereas the dashed lines are the RE paths.

right-hand panel the (discounted) cumulative output multipliers. In the figure, the solid black line illustrates the multipliers under learning, whereas the dashed line illustrates the multipliers under the assumption of RE. The output multipliers are higher under RE compared to learning until $t = 9$. However, the onset of the higher g from $t = 7$ gives a significant boost to the output multiplier under learning that goes above RE levels soon after the policy change and stays higher than RE for the entire period plotted in Figure 6. At $t = 40$, the cumulative output multiplier under learning is 0.63, whereas that under RE is only 0.4.[26] Interestingly, the size of the multiplier depends on the severity of the shocks hitting the economy in the first six quarters; if the size of these shocks is reduced by half, the cumulative output multiplier under learning increases to 0.8 (whereas the RE multiplier is unchanged).

We stress that our model lacks many elements thought to be necessary for generating large multipliers in a recession. Our focus in this section has been to check whether the mechanisms, we have identified in RBC models with learning, continue to deliver significant positive output multipliers when the economy has been subject to a sequence of large adverse "productivity" shocks. Although the multipliers under learning are somewhat smaller than in Section 4, a fiscal stimulus does raise output and employment during the recession. We again see that the assumption of RE underestimates the effects of fiscal policy when agents are learning adaptively over time. Fiscal policy can be effective in the standard RBC model not only when adopted during normal times but also when undertaken during recessionary times, even though our model does not include price or wage rigidities or liquidity constrained households.[27]

8. FISCAL CONSOLIDATION

Since the 1990s, there has been significant interest in the so-called "non-Keynesian" effects of fiscal policy spurred on by the seminal contribution of Giavazzi and Pagano (1990) who studied the two largest fiscal consolidations of

https://doi.org/10.1017/S1365100516001176 Published online by Cambridge University Press

the 1980s, Denmark in 1983–1986 and Ireland in 1987–1989. A striking feature of these contractionary fiscal policies was that the private sector boomed rather than fell into the deep recession that many economists and policy makers had predicted. A huge literature arose pointing to examples of fiscal consolidations (i.e., permanent reductions in government spending) displaying similar "non-Keynesian" effects.[28]

Although the empirical literature is vast, there have been some attempts to explain these effects at a theoretical level, including discussion of whether special theories were needed to explain the effects of large fiscal consolidations. Most of the focus of this literature has been on the explanation of the effects of fiscal policy on private *consumption*.[29] More recently, Alesina et al. (2002) have argued that descriptive evidence suggests that increases in private investment (rather than private consumption) explain a greater share of the response of private-sector GDP growth in large fiscal consolidations.[30] They find very little evidence that private investment reacts differently during these large fiscal adjustments than in the "normal" circumstances. As they remark, "This result questions the need for 'special theories' for large versus small changes in fiscal policy." (p. 586)

Episodes of large fiscal consolidations are good examples of situations that economic agents are unlikely to have experienced earlier in their lifetimes. As argued in Section 1, in such situations it is plausible to replace RE by the assumption that agents gradually learn about effects of these policy changes. Fiscal consolidation has a negative effect on aggregate output and employment under both RE and adaptive learning. However, we will see that the standard RBC model with adaptive learning is able to explain key features in the behavior of private consumption *and* investment in line with the fiscal episodes cited above. The consolidation has a negative effect on aggregate output and employment.

Fiscal consolidations are typically modeled as a surprise permanent reduction in government purchases, starting from steady state at $t = 0$. We consider the following scenario. At the beginning of period $t = 1$, a policy announcement is made that the level of government purchases will fall permanently from $\bar{g} = 0.22$ to $\bar{g}' = 0.20$ (i.e., an almost 10% drop in g). The policy announcement is assumed to be credible and known to the agents with certainty. We believe this is a realistic assumption; drastic cuts in purchases are typically implemented when things turn very bad and the public accepts that permanent adjustments are required.[31]

The long-run effects on the steady state of a decrease in government consumption are well known: higher consumption and lower levels of investment, output, labor, and capital. See, e.g., Baxter and King (1993).

The dynamics under RE are also standard; see, for instance, Baxter and King (1993, pp. 321–322), Heijdra (2009, Chap. 15), or Mitra et al. (2013). The qualitative dynamics are confirmed by the behavior of variables under RE in Figure 7. For our purposes, the most relevant issue is the behavior of c_t and i_t. Under RE, there is a big rise in c_t on impact overshooting the new (higher) steady state followed by a gradual fall toward this steady state. i_t, on the other hand, falls dramatically below the new (lower) steady state on impact followed by a gradual rise over time.

https://doi.org/10.1017/S1365100516001176 Published online by Cambridge University Press

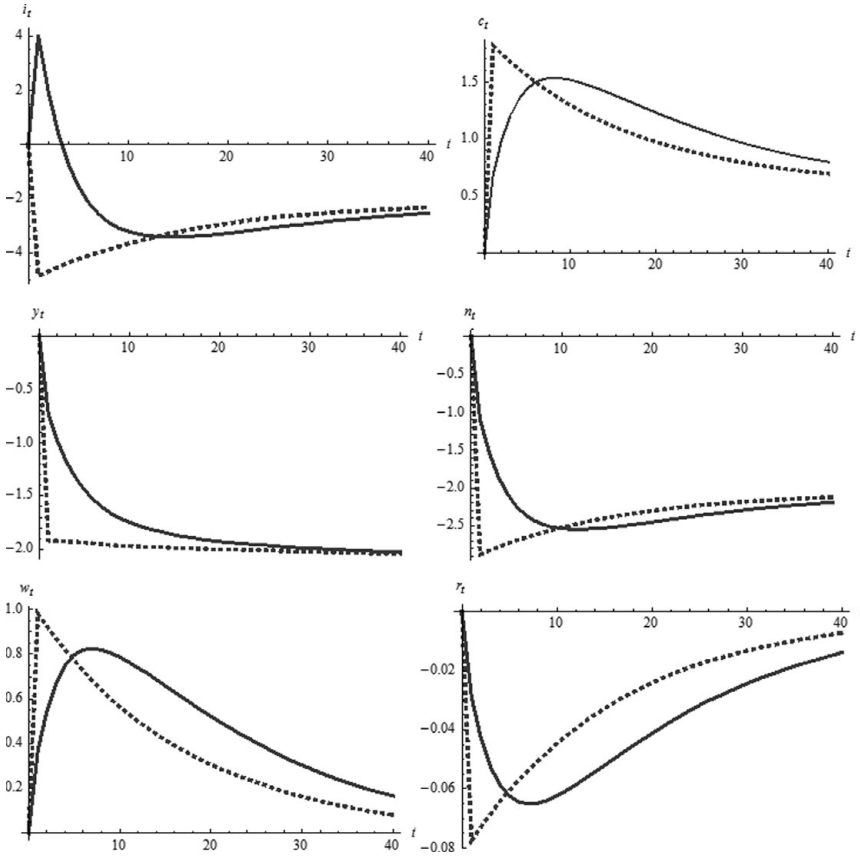


FIGURE 7. Dynamic paths for the fiscal consolidation. The solid lines are the learning paths, whereas the dashed lines are the RE paths. All variables are measured in percentage deviations from steady-state values. Mean paths over 100,000 simulations.

Although the behavior of c_t is consistent with the fiscal episodes mentioned above, the behavior of i_t is at odds with the empirical literature documented above.

Under learning, c_t rises on impact followed by a gradual hump-shaped increase in its level eventually going above the RE level before monotonically falling toward the steady state. The most striking difference from RE is, however, in the behavior of investment. Instead of the big drop in investment under RE, the opposite case of a large boom in investment and hence a rising path of capital occurs under learning in the initial periods after the policy change. Strikingly, this qualitative behavior of i_t under learning is consistent with the empirical evidence cited above.

Why is the behavior of i_t different under learning compared to RE? At $t = 1$, consumption rises because of the decrease in the present value of taxes $S\tau^e_{h,t}$. As in the case of a temporary change in g, discussed in Section 4, the impact effects

https://doi.org/10.1017/S1365100516001176 Published online by Cambridge University Press

are less under learning than under RE because the paths of future w_t and r_t are not fully anticipated. Under learning, w^e_{t+s} and r^e_{t+s} gradually respond to the data, leading initially to a gradual rise in w^e_{t+s} (and fall in r^e_{t+s}) before eventually falling toward the steady state.

As a consequence of the smaller sizes of the impacts on output and consumption at $t = 1$, the decrease in g necessarily leads to a higher level of i_t under learning than under RE, and in fact a sharp increase in investment follows. In the periods immediately following the policy change, expectations of wages and interest rates begin to adjust. Two factors are at work. The higher capital stock in the periods soon after the policy change leads to higher forecasts of future wages and lower forecasts of future interest rates and thus higher Sw^e_t and lower Sr^e_t. This leads to a further increase in c_t, and decreases in n_t and y_t, which results in decreases in i_t from its high level at $t = 1$. After several periods, this process moves k_t to a downward path, accompanied by a rise in n_t, and a decrease in k_t/n_t, driving w_t downwards and r_t upwards to their steady-state values. The other factor at work is that over time coefficient estimates under RLS learning gradually adjust in response to the shock and the evolution of the data. Eventually the coefficients converge to the values that correspond to the REE values at the new steady state, so that in the long run there is convergence to the new REE.

Under adaptive learning, the behavior of i_t and c_t is both in line with the episodes of fiscal retrenchment cited above. Investment increases sharply under learning: In period 1 it is more than 4% points higher than the initial steady value and continues to stay higher than RE levels for three years. c_t grows less rapidly under learning compared to RE levels for six quarters but is then significantly above RE levels for a sustained period. These results for i_t and c_t are obtained in the conventional RBC model under learning, without the need to introduce real frictions or distortionary taxes.

Table 2 summarizes the impact of learning on the behavior of investment, consumption, and output. For each variable, the table gives, over different horizons, the difference between the cumulative impact under learning and under RE (the table is based on the same data as in Figure 7). This difference is particularly striking for investment. For example, over five years, the cumulative difference between the level of investment under learning and under RE amounts to 6.61% of steady-state output or 31.91% of steady-state investment.[32] Over 5- and 10-year horizons, the cumulative effect on consumption is also greater under learning than under RE. It follows that the cumulative difference between the level of output under learning and under RE, which is equal to the sums of the differences for investment and consumption, is also large over all three horizons. Over 10 years, this difference amounts to over 7.57% of steady-state output.

It should be emphasized, however, that fiscal consolidation leads to a fall in aggregate output and employment under both RE and learning.[33] This is an unavoidable consequence of the lower steady-state level that necessarily accompanies a permanent reduction in g in the basic RBC model that we are using. However, n_t falls less rapidly under learning and is around 0.7% points higher than RE

https://doi.org/10.1017/S1365100516001176 Published online by Cambridge University Press

TABLE 2. Cumulative effects on key variables of a fiscal consolidation on the basis of averages over 100,000 simulations

Cumulative effects, %	3 years RLS−RE	5 years RLS−RE	10 years RLS−RE
$\left(\sum i_t\right)/\bar{y}$	6.91	6.61	5.28
$\left(\sum c_t\right)/\bar{y}$	−1.09	0.19	2.29
$\left(\sum y_t\right)/\bar{y}$	5.82	6.80	7.57
$\left(\sum i_t\right)/\bar{i}$	33.35	31.91	25.48
$\left(\sum c_t\right)/\bar{c}$	−1.84	0.32	3.87

Cumulative difference between effects under learning (RLS) and under rational expectations (RE) is in percentage relative to the new steady state.

levels after year 1. This feature explains the higher levels of output under learning compared to RE levels for the entire 10-year period depicted in Figure 5 and summarized in Table 2.

To summarize, the literature on fiscal consolidation emphasizes the possibility of positive effects on both private consumption and, especially, private investment resulting from permanent decreases in government spending. Adaptive learning can provide a natural mechanism, operating through expectations, for a surge in investment immediately following a fiscal consolidation, as well as a sustained period of higher consumption.

Of course, the detailed results will also depend on the specific econometric forecasting model used by agents. The dummy variable specification used in Section 5 is an alternative that could be used in the current context. Using this specification, we find that our fiscal consolidation results are essentially unchanged if ω is large, so that agents distrust estimates based on past temporary government spending shocks when forming priors on the dummy variable coefficients. However, for smaller ω, especially if μ is low, the surge in investment does not take place. This suggests that the behavior of private sector spending in response to fiscal consolidation may vary across countries, depending on specific circumstances.

9. CONCLUSIONS

In this paper, we have studied the impact of changes in government purchases in a standard RBC model with adaptive learning. Methodologically, our approach has been to assume that households understand the direct effects of announced changes in government purchases on their after-tax income, but have imperfect knowledge of the implications of the policy for the future paths of wages and interest rates. Expectations of these latter variables follow the adaptive learning approach in which agents estimate and update their forecasts using statistical learning rules.

Using this approach, we study the implications for three inter-related questions that have been a major focus of recent research. Our main finding is that the multiplier effects of government purchases in RBC models under learning can be

https://doi.org/10.1017/S1365100516001176 Published online by Cambridge University Press

much larger than under the standard RE assumption, and in particular they are compatible with the range found in empirical studies. In our baseline formulation of adaptive learning, there is less crowding out of consumption and there is substantial crowding in of investment. For a wide range of parameters, this result extends to alternative econometric specifications of the agents' perceived law of motion and to the extent of structural knowledge they have about future taxes. We also find that fiscal policy, taking the form of temporary increases in government purchases, can increase output and employment during severe recessions. Finally, we have seen that the behavior of both consumption and investment under fiscal consolidations can better match some stylized empirical facts when adaptive learning is incorporated into the RBC model.

In future work, we aim to study these issues in extended models that allow for more realistic forms of government financing, incorporating distortionary taxes and government debt, and in models that include aggregate demand channels.

NOTES

1. Active fiscal strategies adopted in the United States and United Kingdom include temporary tax cuts and credits and large public works projects; see, for instance, Auerbach et al. (2010).

2. We briefly discuss empirical estimates in Section 4.

3. Leeper et al. (2011) report simulated multipliers for a series of nested models in which New Keynesian models are specified as generalizations of the RBC model.

4. For discussion of the adaptive learning approach and extensive references, see, for example, Evans and Honkapohja (2001), Sargent (2008) and Evans and Honkapohja (2013). Policy change under learning has also been studied in Evans et al. (2001), Marcet and Nicolini (2003), Giannitsarou (2006), Nunes (2009), and Berardi (2013).

5. Both surprise and preannounced increases were considered in Mitra et al. (2013) using the methodology introduced in Evans et al. (2009) for learning dynamics.

6. We note that, as shown in Chari et al. (2007), productivity shocks in the RBC model are observationally equivalent to changes in the efficiency wedge resulting from changes in relative financing distortions.

7. King et al. (1988) emphasize that log utility for consumption is needed for steady-state labor supply along a balanced growth path.

8. A related approach is followed in Preston (2006) and Eusepi and Preston (2010) in connection with monetary policy: In some cases agents are assumed to incorporate the announced interest-rate rule in their forecasts.

9. The assumption that current productivity is observed is standard under RE and also typically assumed in the adaptive learning literature. Hamilton and Neville (2014) consider imperfect observability of technology shocks in an RBC model using Kalman filter learning.

10. See Evans et al. (2009) for an example of learning from repeated policy changes.

11. We use standard RE values for calibration since we are considering an economy initially in the RE equilibrium before the policy change.

12. As argued in Part I of Benveniste et al. (1990), the size of the gain should reflect the trade-off between tracking and filtering. A policy change is in effect a structural change that requires a higher weight on tracking, and hence a relatively large gain. We discuss later in this section the sensitivity of our results to different choices of the gain parameter.

13. Ilzetzki et al. (2013) look at cross-country estimates of the fiscal multipliers and find that there are substantial differences in government consumption multipliers depending, for example, on the level of development, exchange rate regime, and openness to trade. For high income countries, they find

https://doi.org/10.1017/S1365100516001176 Published online by Cambridge University Press

that the cumulative multiplier rises to a long-run value of 0.66 and for closed economies, the long-run cumulative multiplier is 1.1. For public investment, the long-run cumulative multipliers are well above one.

14. At the same time, we acknowledge that the consumption decline during the first 10 or so quarters in Figure 1 is inconsistent with most of the empirical literature. Thus, adaptive learning does not fully reconcile the RBC model with this feature of the data.

15. The numerical results with adaptive learning of taxes are obtained using 20,000 replications.

16. These baseline numerical results corresponding to Figures 1 and 2 used 100,000 replications.

17. The numerical results for $\eta > 0$ are obtained using 20,000 replications.

18. The RLS moment matrix estimate R_t is now 5×5 and must also be initialized. For the 4×4 submatrix corresponding to the second-moment matrix of $(1, k_t, \hat{v}_t, \hat{g}_t)$, we use the RE value. To this we append $R(i, 5) = R(5, i) = 0$ for $i = 1, 2, 3, 4$ and $R(5, 5) = (\bar{g}' - \bar{g})^2$. This implies that in the first period a_{kd}, a_{wd}, and a_{rd} are simply adjusted by the forecast error times the gain.

19. As is common, we have specified utility in terms of leisure, $1-n$, where the unit of time available for work or leisure is set to 1. An alternative specification sometimes used is in terms of the disutility of labor $\chi(1+\varphi)^{-1}n^{1+\varphi}$.

20. For discussion on how and whether macro-RBC elasticities and micro-type calibrations can be reconciled, see Chetty et al. (2011) and the references therein.

21. We have also simulated the results for an even lower Frisch labor supply elasticity of 0.31, and the conclusions remain roughly unchanged.

22. However, this could be implemented at the cost of even greater computation complexity.

23. For a summary of the features of the ARRA, see Romer and Bernstein (2009) and Cogan et al. (2010).

24. The policy we consider now is an announced anticipated change in g that takes place in the near future. See the appendix for details.

25. As discussed in Section 4, investment is to some extent crowded out during the first part of the implementation, followed by a recovery during the later part of the implementation and a surge as the policy ends.

26. When Figures 3 and 4 are recalculated for 1,000 replications using the nonlinear system under learning, there are only minor differences in the results. For example, for both the linear and nonlinear systems, the peak distributed lag multiplier occurs at $t = 17$ and takes the mean value 0.680 in the linear case and 0.696 in the nonlinear case. Similarly, the mean cumulative multipliers at $t = 20$ and $t = 40$ are 0.592 and 0.630, respectively in the linear case and the corresponding numbers for the nonlinear system are 0.601 and 0.642.

27. It should be noted, however, that there is empirical evidence, e.g., Blanchard and Leigh (2013), that output multipliers for fiscal policy have been substantially higher during the recession beginning in 2007 and its aftermath. This is consistent with a strong aggregate demand channel at the zero lower bound.

28. For recent discussion and references, see Hemming et al. (2002), Alesina et al. (1998), Briotti (2005), and Alesina and Ardagna (2010).

29. These attempts include Blanchard (1990), Bertola and Drazen (1993), and Perotti (1999).

30. See also Alesina et al. (1998). Perotti (1999, footnote 31) concedes that these episodes were characterized by big increases in investment (and net exports).

31. As noted by Corsetti et al. (2012), fiscal consolidations may also arise as reversals of earlier increases in government spending. It would be interesting to examine such policies under learning in future work.

32. When Table 2 is recalculated using the nonlinear system under learning, there are only minor differences in the results. For example, based on 3,000 replications the corresponding numbers are 6.71% and 32.39%.

33. Empirical evidence on aggregate effects is reviewed, e.g., in Briotti (2005) and IMF (2010, Chap. 3).

34. In effect, agents assume that the economy has a single steady state.

https://doi.org/10.1017/S1365100516001176 Published online by Cambridge University Press

35. The formulas for Sw_t^e and Sr_t^e for $t \geq T_g - 1$ are analogous to that in the appendix to Mitra et al. (2013) because of the absence of the dummy variable after this period.

REFERENCES

Alesina, A. and S. Ardagna (2010) Large changes in fiscal policy: Taxes versus spending. *Tax Policy and the Economy* 24, 35–68.

Alesina, A., S. Ardagna, R. Perotti and F. Schiantarelli (2002) Fiscal policy, profits, and investment. *American Economic Review* 92, 571–589.

Alesina, A., R. Perotti, J. Tavares, M. Obstfeld and B. Eichengreen (1998) The political economy of Fiscal adjustments. *Brookings Papers on Economic Activity* 1, 197–248.

Auerbach, A., W. G. Gale and B. H. Harris (2010) Activist fiscal policy. *Journal of Economic Perspectives* 24, 141–164.

Barro, R. J. and R. G. King (1984) Macroeconomic effects from government purchases and taxes. *Quarterly Journal of Economics* 99, 817–839.

Barro, R. J. and C. J. Redlick (2011) Macroeconomic effects from government purchases and taxes. *Quarterly Journal of Economics* 126, 51–102.

Baxter, M. and R. G. King (1993) Fiscal policy in general equilibrium. *American Economic Review* 83, 315–334.

Benveniste, A., M. Metivier and P. Priouret (1990) *Adaptive Algorithms and Stochastic Approximations*. Berlin: Springer-Verlag.

Berardi, M. (2013) Escape dynamics and policy specification. *Macroeconomic Dynamics* 17, 123–142.

Bertola, G. and A. Drazen (1993) Trigger points and budget cuts: Explaining the effects of fiscal austerity. *American Economic Review* 83, 11–26.

Blanchard, O. J. (1990) Can severe fiscal contractions be expansionary? A tale of two small european countries: Comment. *NBER Macroeconomics Annual* 1, 75–111.

Blanchard, O. J. and D. Leigh (2013) Growth Forecast Errors and Fiscal Multipliers. Working paper WP/13/1, International Monetary Fund.

Branch, W. A. and G. W. Evans (2006) A simple recursive forecasting model. *Economic Letters* 91, 158–166.

Briotti, M. G. (2005) Economic Reactions to Public Finance Consolidation: A Survey of the Literature. ECB Occasional paper no. 38.

Campbell, J. Y. (1994) Inspecting the mechanism: An analytical approach to the stochastic growth model. *Journal of Monetary Economics* 33, 463–506.

Carceles-Poveda, E. and C. Giannitsarou (2008) Asset pricing with adaptive learning. *Review of Economic Dynamics* 11, 629–651.

Chari, V., P. J. Kehoe and E. R. McGrattan (2007) Business cycle accounting. *Econometrica* 75, 781–836.

Chetty, R., A. Guren, M. Day and A. Weber (2011) Are micro and macro labor supply elasticities consistent? A review of evidence on the intensive and extensive margins. *American Economic Review: Papers and Proceedings* 101, 471–475.

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

Coenen, G., C. J. Erceg, C. Freedman, D. Furceri, M. Kumhof, R. Lalonde, D. Laxton, J. Lindè, A. Mouragane, D. Muir, S. Mursula, C. De Resende, J. Roberts, W. Roeger, S. Snudden, M. Trabandt and J. in't Veld (2012) Effects of fiscal stimulus in structural models. *American Economic Journal: Macroeconomics* 4, 22–68.

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

Corsetti, G., K. Kuester, A. Meier and G. J. Muller (2010) Debt consolidation and fiscal stabilization of deep recessions. *American Economic Review: Papers and Proceedings* 100, 41–45.

https://doi.org/10.1017/S1365100516001176 Published online by Cambridge University Press

Corsetti, G., A. Meier and G. J. Muller (2012) Fiscal stimulus with spending reversals. *Review of Economics and Statistics* 94, 878–895.
Eusepi, S. and B. Preston (2010) Central bank communication and expectations stabilization. *American Economic Journal: Macroeconomics* 2, 235–271.
Eusepi, S. and B. Preston (2011) Expectations, learning and business cycle fluctuations. *American Economic Review* 101, 2844–2872.
Evans, G. W. and S. Honkapohja (2001) *Learning and Expectations in Macroeconomics*. Princeton, NJ: Princeton University Press.
Evans, G. W. and S. Honkapohja (2013) Learning as a rational foundation for macroeconomics and finance. In Frydman and Phelps (2013) chap. 2, pp. 68–111.
Evans, G. W., S. Honkapohja and R. Marimon (2001) Convergence in monetary inflation models with heterogeneous learning rules. *Macroeconomic Dynamics* 5, 1–31.
Evans, G. W., S. Honkapohja and K. Mitra (2009) Anticipated fiscal policy and learning. *Journal of Monetary Economics* 56, 930–953.
Evans, G. W., S. Honkapohja and K. Mitra (2012) Does Ricardian equivalence hold when expectations are not rational? *Journal of Money, Credit and Banking* 44, 1259–1283.
Frydman, R. and E. Phelps (eds.) (2013) *Rethinking Expectations: The Way Forward for Macroeconomics*. Princeton, NJ: Princeton University Press.
Giannitsarou, C. (2006) Supply-side reforms and learning dynamics. *Journal of Monetary Economics* 53, 291–309.
Giavazzi, F. and M. Pagano (1990) Can severe fiscal contractions be expansionary? A tale of two small European countries. *NBER Macroeconomics Annual* 1, 75–111.
Hall, R. E. (2009) By how much does GDP rise If the government buys more output? *Brookings Papers on Economic Activity* 2009, 183–231.
Hamilton, B. F. and R. F. Neville (2014) Imperfect transmission of technology shocks and the business cycle consequences. *Macroeconomic Dynamics* 18, 418–437.
Heijdra, B. J. (2009) *Foundations of Modern Macroeconomics*. Oxford: Oxford University Press.
Hemming, R., M. Kell and S. Mahfouz (2002) The Effectiveness of Fiscal Policy in Stimulating Economic Activity – A Review of the Literature. Working paper 02/208, International Monetary Fund.
Ilzetzki, E., E. G. Mendoza and C. A. Vegh (2013) How big (small?) are fiscal multipliers? *Journal of Monetary Economics* 60, 239–254.
IMF (2010) *World Economic Outlook, October 2010*. Washington, DC: IMF Publication Services.
King, R. G., C. I. Plosser and S. T. Rebelo (1988) Production, growth and business cycles. *Journal of Monetary Economics* 21, 195–232.
King, R. G. and S. T. Rebelo (1999) Resuscitating real business cycles. In Taylor and Woodford (1999) pp. 927–1007.
Leeper, E. M., N. Traum and T. B. Walker (2011) Clearing Up the Fiscal Multiplier Morass. NBER working paper 17444.
Ljungqvist, L. and T. J. Sargent (2012) *Recursive Macroeconomic Theory,* 3rd ed. Cambridge, MA: MIT Press.
Long, Jr., J. B. and C. I. Plosser (1983) Real business cycles. *Journal of Political Economy* 91, 39–69.
Marcet, A. and J. P. Nicolini (2003) Recurrent hyperinflations and learning. *American Economic Review* 93, 1476–1498.
Marcet, A. and T. J. Sargent (1989) Convergence of least-squares learning mechanisms in self-referential linear stochastic models. *Journal of Economic Theory* 48, 337–368.
Mertens, K. and M. O. Ravn (2014) Fiscal Policy in an Expectation Driven Liquidity Trap. *Review of Economic Studies* 81, 1637–1667.
Milani, F. (2007) Expectations, learning and macroeconomic persistence. *Journal of Monetary Economics* 54, 2065–2082.
Mitra, K., G. W. Evans and S. Honkapohja (2013) Policy change and learning in the RBC model. *Journal of Economic Dynamics and Control* 37, 1947–1971.

https://doi.org/10.1017/S1365100516001176 Published online by Cambridge University Press

Nunes, R. (2009) Learning the inflation target. *Macroeconomic Dynamics* 13, 167–188.
Orphanides, A. and J. C. Williams (2007) Robust monetary policy with imperfect knowledge. *Journal of Monetary Economics* 54, 1406–1435.
Perotti, R. (1999) Fiscal policy in good times and bad. *Quarterly Journal of Economics* 114, 1399–1436.
Preston, B. (2006) Adaptive learning, forecast-based instrument rules and monetary policy. *Journal of Monetary Economics* 53, 507–535.
Ramey, V. A. (2011a) Can government purchases stimulate the economy? *Journal of Economic Literature* 49, 673–685.
Ramey, V. A. (2011b) Identifying government spending shocks: It's all in the timing. *Quarterly Journal of Economics* 126, 1–50.
Ravn, Morten, O., S. Schmitt-Grohe and M. Uribe (2012) Consumption, government spending, and the real exchange rate. *Journal of Monetary Economics* 59, 215–234.
Romer, C. D. and J. Bernstein (2009) The Job Impact of the American Recovery and Reinvestment Plan. Mimeo.
Sargent, T. J. (1999) *The Conquest of American Inflation*. Princeton, NJ: Princeton University Press.
Sargent, T. J. (2008) Evolution and intelligent design. *American Economic Review* 98, 5–37.
Taylor, J. and M. Woodford (eds.) (1999) *Handbook of Macroeconomics*, vol. 1, Amsterdam: Elsevier.
Woodford, M. (2011) Simple analytics of the government expenditure multiplier. *American Economic Journal: Macroeconomics* 3, 1–35.

APPENDIX A: LINEARIZATIONS

The linearized wage rate, rental rate, real interest rate, output, and capital accumulation equations are

$$w_t - \bar{w} = \bar{w}\left[\left(\frac{v_t}{\bar{v}} - 1\right) + \alpha\left(\frac{k_t}{\bar{k}} - 1\right) - \alpha\left(\frac{n_t}{\bar{n}} - 1\right)\right],$$

$$r_{k,t} - \bar{r}_k = \bar{r}_k\left[\left(\frac{v_t}{\bar{v}} - 1\right) - (1-\alpha)\left(\frac{k_t}{\bar{k}} - 1\right) + (1-\alpha)\left(\frac{n_t}{\bar{n}} - 1\right)\right],$$

$$r_t - \bar{r} = r_{k,t} - \bar{r}_k,$$

$$y_t - \bar{y} = \bar{y}\left[\left(\frac{v_t}{\bar{v}} - 1\right) + \alpha\left(\frac{k_t}{\bar{k}} - 1\right) + (1-\alpha)\left(\frac{n_t}{\bar{n}} - 1\right)\right],$$

$$k_{t+1} - \bar{k} = (y_t - \bar{y}) - (c_t - \bar{c}) - (g_t - \bar{g}) + (1-\delta)(k_t - \bar{k}).$$

The equations giving the steady state are $\bar{r} = 1 - \delta + \bar{r}_k = \beta^{-1}$, $\bar{c} = \bar{v}\bar{k}^{\alpha}\bar{n}^{1-\alpha} - \delta\bar{k} - \bar{g}$, $\zeta\bar{c}^{\sigma} = \bar{w}(1-\bar{n})^{\epsilon}$, $\bar{w} = (1-\alpha)\bar{v}(\frac{\bar{k}}{\bar{n}})^{\alpha}$, and $\bar{r}_k = \alpha\bar{v}(\frac{\bar{k}}{\bar{n}})^{\alpha-1}$.

APPENDIX B: FISCAL POLICY UNDER LEARNING DYNAMICS

Under learning, we define variables as deviations from estimated steady states, i.e., $\tilde{k}_t = k_t - \bar{k}_t^e$, $\tilde{r}_{k,t} = r_{k,t} - \bar{r}_{k,t}^e$, and $\tilde{w}_t = w_t - \bar{w}_t^e$. Using the PLMs (17)–(20), we can obtain

https://doi.org/10.1017/S1365100516001176 Published online by Cambridge University Press

estimates of the steady state (omitting the time subscripts on $\bar{k}_t^e$, etc.):

$$\bar{k}^e = \frac{b_k}{1-a_{kk}}, \bar{r}_k^e = b_r + a_{rk}\frac{b_k}{1-a_{kk}}, \bar{w}^e = b_w + a_{wk}\frac{b_k}{1-a_{kk}}.$$

Here, we can write

$$\tilde{k}_{t+1} = a_{kk}\tilde{k}_t + a_{kv}\hat{v}_t, \tag{B.1}$$

$$\tilde{w}_t = a_{wk}\tilde{k}_t + a_{wv}\hat{v}_t, \tag{B.2}$$

$$\tilde{r}_{k,t} = a_{rk}\tilde{k}_t + a_{rv}\hat{v}_t. \tag{B.3}$$

Here, under learning the coefficients a_{kk} etc. denote the estimated values at a moment in time. Then, under learning, forecasts of capital and productivity are made using

$$\begin{pmatrix} \tilde{k}_{t+1} \\ \hat{v}_{t+1} \end{pmatrix} = \tilde{B}\begin{pmatrix} \tilde{k}_t \\ \hat{v}_t \end{pmatrix} + \begin{pmatrix} 0 \\ \tilde{u}_{t+1} \end{pmatrix},$$

$$\tilde{B} = \begin{pmatrix} a_{kk} & a_{kv} \\ 0 & \rho \end{pmatrix}.$$

Defining $\tilde{x}_t \equiv \begin{pmatrix} \tilde{k}_t \\ \hat{v}_t \end{pmatrix}$, we have for $j \geq 1$,

$$\tilde{x}_{t+j}^e = \tilde{B}^j \tilde{x}_t. \tag{B.4}$$

Using the future forecasts of capital stocks from (B.4), we can in turn obtain the future forecasts of wages and rental rates from (B.2) and (B.3) as

$$\tilde{w}_{t+j}^e = \begin{pmatrix} a_{wk} & a_{wv} \end{pmatrix} \tilde{B}^j \tilde{x}_t \text{ and } \tilde{r}_{k,t+j}^e = \begin{pmatrix} a_{rk} & a_{rv} \end{pmatrix} \tilde{B}^j \tilde{x}_t.$$

These are then used to form Sw_t^e and Sr_t^e. It can be shown, e.g., that

$$Sw_t^e = \frac{\beta}{1-\beta}(\bar{w}_t^e - \bar{w}) + \begin{pmatrix} a_{wk} & a_{wv} \end{pmatrix} \beta\tilde{B}(I - \beta\tilde{B})^{-1}\tilde{x}_t,$$

and

$$Sr_t^e = \frac{\beta^2}{(1-\beta)^2}(\bar{r}_{k,t}^e - \bar{r}_k) + \beta^2 \begin{pmatrix} a_{rk} & a_{rv} \end{pmatrix} (I - \tilde{B})^{-1}\tilde{B}[I(1-\beta)^{-1} - \tilde{B}(I - \beta\tilde{B})^{-1}]\tilde{x}_t.$$

For details, see Mitra et al. (2013).

For Section 4 , the other needed item is $S\tau_{h,t}^e$, which is given in the text. For Section 7, we need to calculate this for an announced future temporary government spending change.

We consider a temporary anticipated policy change that is known to be temporary by the agents. We assume $T_{g2} > T_{g1} + 1$ as in our policy change. We have

$$g_t - \bar{g} = \begin{cases} 0, t = 0, \ldots, T_{g1} - 1, \\ \Delta\bar{g} \equiv (\bar{g}' - \bar{g}), t = T_{g1}, \ldots, T_{g2} - 1, \\ 0, t \geq T_{g2}. \end{cases}$$

https://doi.org/10.1017/S1365100516001176 Published online by Cambridge University Press

Under learning, we have when $1 \leq t \leq T_{g1} - 1$,

$$S\tau^e_{h,t} = \sum_{j=1}^{\infty} \beta^j (g_{t+j} - \bar{g}) = \Delta\bar{g} \sum_{j=T_{g1}-t}^{T_{g2}-t-1} \beta^j = \frac{\beta^{T_{g1}-t}(1 - \beta^{T_{g2}-T_{g1}})}{1-\beta} \Delta\bar{g}.$$

For $T_{g1} \leq t \leq T_{g2} - 1$,

$$S\tau^e_{h,t} = \sum_{j=1}^{\infty} \beta^j (g_{t+j} - \bar{g}) = \frac{\beta(1 - \beta^{T_{g2}-t-1})}{1-\beta} \Delta\bar{g}$$

and $S\tau^e_{h,t} = 0$ for $t \geq T_{g2}$. Thus,

$$S\tau^e_{h,t} = \begin{cases} \frac{\beta^{T_{g1}-t}(1-\beta^{T_{g2}-T_{g1}})}{1-\beta} \Delta\bar{g}, t = 1, \ldots, T_{g1} - 1, \\ \frac{\beta(1-\beta^{T_{g2}-t-1})}{1-\beta} \Delta\bar{g}, t = T_{g1}, \ldots, T_{g2} - 1, \\ 0, t \geq T_{g2}. \end{cases}$$

APPENDIX C: DETAILS FOR DUMMY VARIABLES SPECIFICATION IN SECTION 5

We now compute Sw^e_t and Sr^e_t when agents include a dummy variable in their regression equations as in Section 5. First, we can write the PLMs in Section 5 in deviation form as shown in Appendix A; with deviations under learning taken from the estimated steady-state values of capital, wage rate, and rental rate.[34] Using this notation, we have the PLMs with dummy variables included as follows:

$$\begin{aligned} \tilde{k}_{t+1} &= a_{kk}\tilde{k}_t + a_{kv}\hat{v}_t + a_{kg}\hat{g}_t + a_{kd} D\bar{g}_t + \text{noise}, \\ \tilde{w}_t &= a_{wk}\tilde{k}_t + a_{wv}\hat{v}_t + a_{wg}\hat{g}_t + a_{wd} D\bar{g}_t + \text{noise}, \\ \tilde{r}_{k,t} &= a_{rk}\tilde{k}_t + a_{rv}\hat{v}_t + a_{rg}\hat{g}_t + a_{rd} D\bar{g}_t + \text{noise}. \end{aligned}$$

The variable $D\bar{g}_t$ is a discrete variable taking the value $D\bar{g}_t = \Delta\bar{g} \equiv (\bar{g}' - \bar{g})$ for all $t \leq T_g - 1$ and is zero when $t \geq T_g$ .Forecasts use $\tilde{x}'_t \equiv (\tilde{k}_t, \hat{v}_t, \hat{g}_t)$, and we have for periods $1 \leq t \leq T_g - 1$,

$$\begin{aligned} \tilde{x}_{t+1} &= B\tilde{x}_t + d + \upsilon_{t+1}, \\ d &= \begin{pmatrix} a_{kd}\Delta\bar{g} \\ 0 \\ 0 \end{pmatrix}, \upsilon_{t+1} = \begin{pmatrix} 0 \\ \tilde{u}_{t+1} \\ \varepsilon_{t+1} \end{pmatrix}, \end{aligned}$$

and for periods $t \geq T_g$, we have

$$\tilde{x}_{t+1} = B\tilde{x}_t + \upsilon_{t+1}.$$

https://doi.org/10.1017/S1365100516001176 Published online by Cambridge University Press

First, consider periods $1 \leq t \leq T_g - 1$. We have

$$\tilde{x}^e_{t+i} = B^i \tilde{x}_t + (I - B^i)(I - B)^{-1} d \text{ for } i \leq T_g - t - 1,$$
$$\tilde{x}^e_{t+i} = B^{t+i-T_g+1} \tilde{x}^e_{T_g-1} \text{ for } i \geq T_g - t.$$

For $t \geq T_g - 1$, we have $\tilde{x}^e_{t+i} = B^i \tilde{x}_t$. We use these to compute Sw^e_t and Sr^e_t when the dummy variable is included in agents' regression equations.

Sw^e_t can be obtained from

$$Sw^e_t = \sum_{j=1}^{\infty} \beta^j (\bar{w}^e_t - \bar{w}) + \sum_{j=1}^{\infty} \beta^j \tilde{w}^e_{t+j}$$
$$= \frac{\beta}{1-\beta}(\bar{w}^e_t - \bar{w}) + \sum_{j=1}^{\infty} \beta^j \left[\begin{pmatrix} a_{wk} & a_{wv} & a_{wg} \end{pmatrix} \tilde{x}^e_{t+j} + a_{wd} Dg_{t+j} \right]$$

since $\tilde{w}^e_{t+j} = (a_{wk} \quad a_{wv} \quad a_{wg}) \tilde{x}^e_{t+j} + a_{wd} Dg_{t+j}$. The infinite sum above for $t \leq T_g - 2$ can be computed as

$$\sum_{j=1}^{\infty} \beta^j \left[\begin{pmatrix} a_{wk} & a_{wv} & a_{wg} \end{pmatrix} \tilde{x}^e_{t+j} + a_{wd} Dg_{t+j} \right]$$
$$= \sum_{j=1}^{\infty} \beta^j a_{wd} Dg_{t+j} + \sum_{j=1}^{\infty} \beta^j \begin{pmatrix} a_{wk} & a_{wv} & a_{wg} \end{pmatrix} \tilde{x}^e_{t+j}$$
$$= \sum_{j=1}^{T_g-t-1} a_{wd} \Delta \bar{g} \beta^j + \begin{pmatrix} a_{wk} & a_{wv} & a_{wg} \end{pmatrix} \sum_{j=1}^{\infty} \beta^j \tilde{x}^e_{t+j}$$
$$= a_{wd} \Delta \bar{g} \beta \frac{1 - \beta^{T_g-t-1}}{1-\beta} + \beta B (I - \beta B)^{-1} \tilde{x}_t$$
$$+ \begin{bmatrix} \frac{1-\beta^{T_g-t}}{1-\beta} I - (I - \beta B)^{-1}(I - (\beta B)^{T_g-t}) + \\ \beta^{T_g-t-1} \beta B (I - \beta B)^{-1} - (\beta B)^{T_g-t} (I - \beta B)^{-1} \end{bmatrix} (I - B)^{-1} d,$$

so that finally for $t \leq T_g - 2$,

$$Sw^e_t = \frac{\beta}{1-\beta}(\bar{w}^e_t - \bar{w}) + a_{wd} \Delta \bar{g} \beta \frac{1 - \beta^{T_g-t-1}}{1-\beta} + \beta B (I - \beta B)^{-1} \tilde{x}_t$$
$$+ \begin{bmatrix} \frac{1-\beta^{T_g-t}}{1-\beta} I - (I - \beta B)^{-1}(I - (\beta B)^{T_g-t}) + \\ \beta^{T_g-t-1} \beta B (I - \beta B)^{-1} - (\beta B)^{T_g-t} (I - \beta B)^{-1} \end{bmatrix} (I - B)^{-1} d.$$

For $t \geq T_g - 1$, we have[35]

$$Sw^e_t = \frac{\beta}{1-\beta}(\bar{w}^e_t - \bar{w}) + \beta \begin{pmatrix} a_{wk} & a_{wv} & a_{wg} \end{pmatrix} B (I - \beta B)^{-1} \tilde{x}_t$$

since the dummy variable is no longer present in this sum.

https://doi.org/10.1017/S1365100516001176 Published online by Cambridge University Press

For Sr_t^e, when $t \geq T_g - 1$, we have

$$Sr_t^e = \sum_{j=1}^{\infty} \beta^{j+1} \sum_{i=1}^{j} (r_{t+i}^e - \bar{r}) = \frac{\beta^2}{(1-\beta)^2} (\bar{r}_{k,t}^e - \bar{r}_k)$$
$$+ \beta^2 \begin{pmatrix} a_{wk} & a_{wv} & a_{wg} \end{pmatrix} (I - \tilde{B})^{-1} \tilde{B} [I(1-\beta)^{-1} - \tilde{B}(I - \beta\tilde{B})^{-1}] \tilde{x}_t.$$

When $t \leq T_g - 2$, we have

$$Sr_t^e = \frac{\beta^2}{(1-\beta)^2} (\bar{r}_{k,t}^e - \bar{r}_k) + S_{rtilda_t},$$
$$S_{rtilda_t} \equiv \sum_{j=1}^{\infty} \beta^{j+1} \sum_{i=1}^{j} \tilde{r}_{k,t+i}^e; \tilde{r}_{k,t+i}^e = \begin{pmatrix} a_{wk} & a_{wv} & a_{wg} \end{pmatrix} \tilde{x}_{t+i}^e + a_{wd} Dg_{t+i}.$$

Hence,

$$S_{rtilda_t} = \sum_{j=1}^{\infty} \beta^{j+1} \sum_{i=1}^{j} \left[\begin{pmatrix} a_{wk} & a_{wv} & a_{wg} \end{pmatrix} \tilde{x}_{t+i}^e + a_{rd} Dg_{t+i} \right]$$
$$= \sum_{j=1}^{\infty} \beta^{j+1} \sum_{i=1}^{j} \left[\begin{pmatrix} a_{wk} & a_{wv} & a_{wg} \end{pmatrix} \tilde{x}_{t+i}^e \right] + \sum_{j=1}^{\infty} \beta^{j+1} \sum_{i=1}^{j} a_{rd} Dg_{t+i}$$
$$\equiv S_{rtildaX_t} + S_{rtildaD_t}. \tag{C.1}$$

Note that

$$\sum_{i=1}^{j} a_{rd} Dg_{t+i} = \begin{cases} j a_{rd} \Delta\bar{g}, \; j \leq T_g - 1 \\ (T_g - 1) a_{rd} \Delta\bar{g}, \; j \geq T_g. \end{cases}$$

So,

$$S_{rtildaD_t} = \sum_{j=1}^{T_g - 2} \beta^{j+1} j a_{rd} \Delta\bar{g} + \sum_{j=T_g-1}^{\infty} \beta^{j+1} (T_g - 1) a_{rd} \Delta\bar{g}$$
$$= a_{rd} \Delta\bar{g} \frac{\beta^2 \left[1 - (T_g - 1)\beta^{T_g - 2} + (T_g - 2)\beta^{T_g - 1}\right]}{(1-\beta)^2} + a_{rd} \Delta\bar{g} (T_g - 1) \frac{\beta^{T_g}}{1-\beta}.$$

Now, we compute $S_{rtildaX_t}$ when $t \leq T_g - 2$. For this, we first note that when $j \leq T_g - t - 1$,

$$\sum_{i=1}^{j} \tilde{x}_{t+i}^e = \sum_{i=1}^{j} B^i \tilde{x}_t + j(I - B)^{-1} d - \sum_{i=1}^{j} B^i (I - B)^{-1} d$$
$$= [(I - B)^{-1} B - (I - B)^{-1} B^{j+1}] \tilde{x}_t + j(I - B)^{-1} d$$
$$- [(I - B)^{-1} B - (I - B)^{-1} B^{j+1}] (I - B)^{-1} d,$$

https://doi.org/10.1017/S1365100516001176 Published online by Cambridge University Press

and when $j \geq T_g - t$, we have

$$\sum_{i=1}^{j} \tilde{x}_{t+i}^e = [(I-B)^{-1}B - (I-B)^{-1}B^{T_g-t}][\tilde{x}_t - (I-B)^{-1}d] + (T_g - t - 1)(I-B)^{-1}d + \left[(I-B)^{-1}B - (I-B)^{-1}B^{t+j+2-T_g}\right]\tilde{x}_{T_g-1}^e,$$

where $\tilde{x}_{T_g-1}^e = B^{T_g-1-t}\tilde{x}_t + (I - B^{T_g-t-1})(I-B)^{-1}d$.

Using this information, we compute $S_{rtildaX_t}$ when $t \leq T_g - 2$. For this, we need to find

$$\sum_{j=1}^{\infty}\beta^{j+1}\sum_{i=1}^{j}\tilde{x}_{t+i}^e = \sum_{j=1}^{T_g-t-1}\beta^{j+1}\sum_{i=1}^{j}\tilde{x}_{t+i}^e + \sum_{j=T_g-t}^{\infty}\beta^{j+1}\sum_{i=1}^{j}\tilde{x}_{t+i}^e \equiv Z_1 + Z_2.$$

We simplify Z_1 now:

$$Z_1 = \sum_{j=1}^{T_g-t-1}\beta^{j+1}\left[(I-B)^{-1}B\tilde{x}_t - (I-B)^{-1}B(I-B)^{-1}d\right] + (I-B)^{-1}\left(\sum_{j=1}^{T_g-t-1}(\beta B)^{j+1}\right)\left[(I-B)^{-1}d - \tilde{x}_t\right] + \left(\sum_{j=1}^{T_g-t-1}\beta^{j+1}j\right)\left[(I-B)^{-1}d\right].$$

Evaluating the relevant sums above yields Z_1 finally

$$Z_1 = \frac{\beta^2\left[1-\beta^{T_g-t-1}\right]}{(1-\beta)}\left[(I-B)^{-1}B\tilde{x}_t - (I-B)^{-1}B(I-B)^{-1}d\right] + \beta^2(I-B)^{-1}B^2(I-\beta B)^{-1}[I-(\beta B)^{T_g-t-1}]\left[(I-B)^{-1}d - \tilde{x}_t\right] + \left\{\frac{\beta^2\left[1-(T_g-t)\beta^{T_g-t-1}+(T_g-t-1)\beta^{T_g-t}\right]}{(1-\beta)^2}\right\}(I-B)^{-1}d.$$

Next, we compute Z_2 :

$$Z_2 = \frac{\beta^{T_g-t+1}}{(1-\beta)}\left[(I-B)^{-1}B - (I-B)^{-1}B^{T_g-t}\right][\tilde{x}_t - (I-B)^{-1}d] + (T_g-t-1)(I-B)^{-1}d + \sum_{j=T_g-t}^{\infty}\beta^{j+1}\sum_{i=T_g-t}^{j}B^{t+i-(T_g-1)}\tilde{x}_{T_g-1}^e.$$

https://doi.org/10.1017/S1365100516001176 Published online by Cambridge University Press

The final term can be shown to be

$$\sum_{j=T_g-t}^{\infty} \beta^{j+1} \sum_{i=T_g-t}^{j} B^{t+i-(T_g-1)} \tilde{x}^e_{T_g-1}$$

$$= \frac{\beta^{T_g-t+1}}{(1-\beta)} (I-B)^{-1} B \tilde{x}^e_{T_g-1}$$

$$-(I-B)^{-1} (\beta B)^{T_g-t+1} (I-\beta B)^{-1} B^{t+1-T_g} \tilde{x}^e_{T_g-1},$$

which provides a simplified expression for Z_2. This finally completes the process of finding $S_{rtildaX_t}$ when $t \leq T_g - 2$.

When $t \geq T_g - 1$, the dummy variable is absent from agents' regression equations so that $S_{rtildaD_t} = 0$ and we have $S_{rtilda_t} = S_{rtildaX_t}$ in (C.1). The latter can be obtained using analogous techniques used in Mitra et al. (2013) to yield the following:

$$S_{rtilda_t} = \sum_{j=1}^{\infty} \beta^{j+1} \sum_{i=1}^{j} \left[\begin{pmatrix} a_{rk} & a_{rv} & a_{rg} \end{pmatrix} \tilde{x}^e_{t+i} \right] =$$

$$\beta^2 \begin{pmatrix} a_{rk} & a_{rv} & a_{rg} \end{pmatrix} (I-B)^{-1} B[(1-\beta)^{-1} I - B(1-\beta B)^{-1}] \tilde{x}_t.$$

This gives the formulas for Sw^e_t and Sr^e_t when the dummy variable is included in agents' regression equations which are required in Section 5.

APPENDIX D: TEMPORARY POLICY CHANGE UNDER RE

The basic techniques are described in Mitra et al. (2013). The key is to obtain the evolution of capital since the other variables can then be obtained using the above linearizations. Under RE, the capital sequence can be shown to be given by

$$\hat{k}_{t+1} = \lambda_2 \hat{k}_t - \lambda_2 A_{k0}^{-1} [S_g(t) + S_v(t)]. \tag{D.1}$$

See Mitra et al. (2013, Appendix B) for the derivation and the values of $S_v(t)$, λ_2, and A_{k0}. The formula for $S_g(t)$ depends on the details of the government spending policy. For the surprise temporary change considered in Section 4, we have

$$S_g(t) = \begin{cases} \left((A_{g0} - 1) \frac{(1-\lambda_1^{-(T_g-t-1)})}{1-\lambda_1^{-1}} + A_{g0} \lambda_1^{-(T_g-t-1)} \right) \Delta \bar{g}, t = 1, \ldots, T_g - 2, \\ A_{g0} \Delta \bar{g}, t = T_g - 1, \\ 0, t \geq T_g. \end{cases}$$

For the announced temporary policy change considered in Section 7, we have

$$g_t - \bar{g} = \begin{cases} 0, t < T_{g1}, \\ \Delta \bar{g}, T_{g1} \leq t < T_{g2}, \\ 0, t \geq T_{g2}, \end{cases}$$

https://doi.org/10.1017/S1365100516001176 Published online by Cambridge University Press

and

$$g_{t+j} - \bar{g} = \begin{cases} 0, t+j < T_{g1}, \\ \Delta\bar{g}, T_{g1} \le t+j < T_{g2}, \\ 0, t+j \ge T_{g2}. \end{cases}$$

Assume $T_{g2} > T_{g1} + 1$. Then,

$$A_{g0}(g_{t+j} - \bar{g}) + A_{g1}(g_{t+j+1} - \bar{g}) = \begin{cases} 0, t+j < T_{g1} - 2, \\ -\Delta\bar{g}, t+j = T_{g1} - 1, \\ (A_{g0} - 1)\Delta\bar{g}, T_{g1} \le t+j < T_{g2} - 1, \\ A_{g0}\Delta\bar{g}, t+j = T_{g2} - 1, \\ 0, t+j \ge T_{g2}. \end{cases}$$

We now compute $S_g(t)$.

First, for $1 \le t < T_{g1}$,

$$\begin{aligned} S_g(t) &= -\Delta\bar{g}\lambda_1^{t-(T_{g1}-1)} + (A_{g0} - 1)\Delta\bar{g} \sum_{j=T_{g1}-t}^{T_{g2}-t-2} \lambda_1^{-j} + A_{g0}\Delta\bar{g}\lambda_1^{-(T_{g2}-t-1)} \\ &= \left[-\lambda_1^{t-(T_{g1}-1)} + (A_{g0} - 1)\frac{\lambda_1^{-(T_{g1}-t)}(1 - \lambda_1^{-(T_{g2}-T_{g1}-1)})}{1 - \lambda_1^{-1}} \right. \\ &\quad \left. + A_{g0}\lambda_1^{-(T_{g2}-t-1)} \right] \Delta\bar{g}. \end{aligned}$$

For $T_{g1} \le t \le T_{g2} - 1$,

$$\begin{aligned} S_g(t) &= (A_{g0} - 1)\Delta\bar{g} \sum_{j=0}^{T_{g2}-t-2} \lambda_1^{-j} + A_{g0}\Delta\bar{g}\lambda_1^{-(T_{g2}-t-1)} \\ &= \left[(A_{g0} - 1)\frac{(1 - \lambda_1^{-(T_{g2}-t-1)})}{1 - \lambda_1^{-1}} + A_{g0}\lambda_1^{-(T_{g2}-t-1)} \right] \Delta\bar{g} \end{aligned}$$

and $S_g(t) = 0$ for all $t \ge T_{g2}$.

https://doi.org/10.1017/S1365100516001176 Published online by Cambridge University Press