AEA Papers and Proceedings 2019, 109: 433–437
https://doi.org/10.1257/pandp.20191084

Policy Options at the Zero Lower Bound When Foresight is Limited†

By Michael Woodford and Yinxi Xie*

The recent global financial crisis and the slow recovery from it have resulted in considerable experimentation with less-conventional approaches to stabilization policy, including both renewed interest in countercyclical fiscal policy and more aggressive use of "forward guidance" with regard to future monetary policy. They have also led to increased debate about more exotic proposals, such as price-level targeting and "helicopter money." All of these policies have been the subject of a good deal of theoretical analysis, but generally under the assumption of rational expectations (RE) equilibrium—that if a novel policy is announced, people's actions should adjust in a way that would be optimal under model-consistent predictions regarding the economy's subsequent dynamics.

The RE assumption is always something of an idealization, but it seems particularly heroic in the case of very novel policies, like many of those recently considered, since one cannot suppose that people should know from previous experience how things should unfold under such a policy. Yet at the same time, conventional analyses of these policies lean quite heavily on the assumptions that are made about what people anticipate about the future.

Here we re-examine a number of policies that have been discussed under a more modest assumption about the degree to which people should be able to correctly foresee the future consequences of a novel policy. The approach that we take is the one proposed in Woodford (2018), based on the architecture of state-of-the-art programs to play games of strategy such as chess or go. Our analysis assumes that in any period t, both households and firms look forward from their current situations some finite distance into the future, to the possible situations that they can reach in period $t+h$ through some finite-horizon action plan; they use structural knowledge (including any announcements about novel government policies) to deduce the consequences of their intended actions over this horizon. (The planning horizon h is here taken to be exogenously fixed.)

Interim situations that someone imagines reaching in period $t+h$ are evaluated using a *value function* that has been learned from past experience. Crucially, we suppose that the value functions cannot be adjusted to take account of an unusual shock or a change in policy, if neither the shock nor the new policy are ones with which people have had much prior experience, though their value functions may be well-adapted to the prior environment. Under some circumstances, this kind of analysis leads to conclusions very similar to conventional RE analysis (at least, under a suitable equilibrium selection criterion), as discussed in Woodford (2018). However, a situation in which monetary policy is constrained by the zero lower bound (ZLB) for a period that may last longer than the length of many people's planning horizons is one in which the finiteness of planning horizons can make a significant difference for the predicted macroeconomic dynamics.

I. A Simple Crisis Model

We re-examine several policies treated in the recent literature using the simple New Keynesian model with finite-horizon planning developed in Woodford (2018). We assume fundamentals described by a two-state Markov chain, as in Eggertsson and Woodford (2003). For a long time prior to date $t = 0$, the economy is assumed to have been in a stationary state in which (i) financial frictions have been

*Woodford: Columbia University, New York, NY 10027 (email: mw2230@columbia.edu); Xie: Columbia University, New York, NY 10027 (email: yx2288@columbia.edu).

†Go to https://doi.org/10.1257/pandp.20191084 to visit the article page for additional materials and author disclosure statement(s).

negligible; (ii) both government purchases and the real public debt have been constant; and (iii) the inflation target π^* has always been achieved, and that as a consequence both households and firms have learned the value functions appropriate to that simple, stationary environment. At $t = 0$, a financial disturbance occurs that increases the demand for safe assets, and causes the nominal interest rate on safe assets that would be required to maintain the inflation target to fall below the ZLB.

Once the financial wedge has increased, we assume that each period there is a probability $0 < \delta < 1$ of it continuing to have the same large value in the following period (the "crisis state" continues), and a probability $1 - \delta$ that instead it reverts permanently to its previous small value (the "normal state"). Let us consider first the case in which both monetary and fiscal policy remain unchanged, and monetary policy is specified by a strict inflation targeting (IT) rule, under which the central bank adjusts the interest rate as needed in order to keep inflation equal to π^*, if this is consistent with the ZLB, and if not sets the lowest interest rate possible. In this case, as in the RE analysis of Eggertsson and Woodford (2003),[1] the solution is Markovian, with output $y_t = \underline{y}$ and inflation $\pi_t = \underline{\pi}$ each period as long as the crisis state persists, but with both variables immediately returning to their normal steady-state values once fundamentals revert to the normal state.

When planning horizons are finite, the contractionary effects of such a shock are less dramatic than in the RE analysis. In our calibrated model,[2] the value of $\underline{y}$ for different horizons h is shown in Figure 1 below, as the value of y when $g = 0$. Even if we assume a ten-year planning horizon for all households and firms (the case $h = 40$), the contraction is only a bit more than half as severe as under the RE analysis. Nonetheless, if there is some degree of foresight, even a relatively modest financial wedge can substantially impact stabilization goals, raising the question whether tools besides conventional interest-rate policy are available to mitigate such effects.

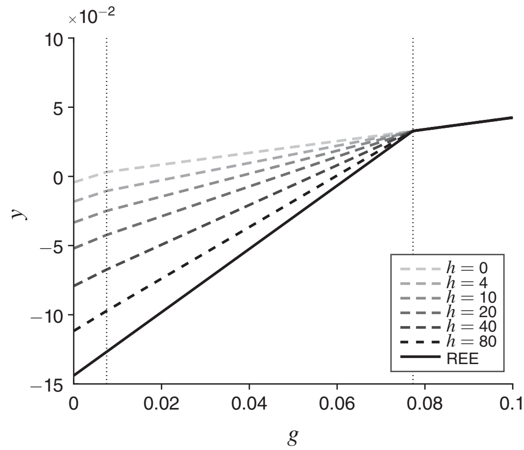

FIGURE 1. OUTPUT EFFECT OF INCREASING GOVERNMENT PURCHASES DURING THE CRISIS

Note: Output relative to the steady-state level as a function of g, for alternative horizons h (in quarters).

II. Efficacy of Fiscal Stimulus

As an example of countercyclical fiscal policy, consider a policy that increases government purchases by fraction g of steady-state GDP, and maintains the higher level as long as the financial wedge remains large, returning to the normal level of purchases when the financial wedge reverts to its normal level. (In order to separate the effects of government purchases from those of government deficits, we assume for simplicity that the path of the real public debt remains unchanged.) There is again a Markovian solution, in which $y_t = \underline{y}$, $\pi_t = \underline{\pi}$ in the crisis state, and both variables return to their normal values in the normal state. We can then investigate the effect of the value of g on outcomes in the crisis state.

Figure 1 shows output (relative to the normal steady state) in the crisis state as a function of g. In the RE case (shown by the lowest line), the contraction would be severe with no increase in g; but output is sharply increasing in g—the "multiplier" is 2.3 in this calibrated example.[3]

[1] Note, however, that in the finite-horizon analysis we avoid the multiplicity of solutions that bedevils RE analysis. See further discussion in Woodford (2018).

[2] We follow the calibration proposed by Eggertsson (2010), in which this kind of shock creates a "Great Depression" under the RE analysis, except that π_* equals 2 percent per year for us. An online Appendix explains our numerical calculations.

[3] Similar conclusions are obtained under the RE assumption by Eggertsson (2010); Christiano, Eichenbaum, and Rebelo (2011); and Woodford (2011).

Moreover, the multiplier remains this high up until the point (g near 8 percent of GDP, in the figure) where the output gap is closed and the inflation target is achieved.

But when planning horizons are finite, the multiplier is not as large. The large multipliers in RE analyses depend on the additional stimulus that comes from anticipation of higher output and inflation in the future (in states in which the financial wedge remains large); a shorter planning horizon reduces these effects. In fact, if horizons are relatively short (for example, $h = 10$ quarters), the multiplier is less than 1 for additional purchases beyond the leftmost vertical dotted line (reached when g is increased by less than 1 percent of GDP). Not only do people not anticipate continuing fiscal stimulus too far into the future in this case, but they anticipate the output gap closing within their planning horizon, and so anticipate monetary tightening in response to the fiscal stimulus, offsetting some of its effect.

Finite planning horizons do increase the effectiveness of fiscal stimulus policies in one respect: they break Ricardian equivalence, and allow government transfers (or deficit spending) to stimulate aggregate demand. In the online Appendix, we consider the effects of a policy that permanently increases the real public debt when a financial crisis shock occurs. A debt-financed transfer increases both the output level $\underline{y}$ and inflation rate $\underline{\pi}$ during the crisis state; for while the present value of future taxes increases by the amount of the transfer, increased taxes beyond households' planning horizon are not taken into account. However, in our model the effects of such transfers remain weak, because of consumption smoothing. A transfer equal to a full quarter's GDP only increases $\underline{y}$ by 0.3 percent, even when $h = 0$, and by less if planning horizons are longer.

Nor is the solution simply to make larger transfers, if monetary policy continues to follow the IT rule. For any increase in public debt above the bound $b^{\max}$ (only 1.4 quarters' GDP if a quarterly model is used in forward planning, and even less if "periods" are shorter), no further increase in $\underline{y}$ or $\underline{\pi}$ is achieved. The reason is that with this degree of fiscal stimulus, the output gap is anticipated to close within the planning horizon, so that monetary policy is expected to be tightened in order to prevent inflation from over-shooting the target; further increases in b simply increase anticipated interest rates later in the planning horizon, preventing any further increase in desired spending.

One can instead achieve a larger stimulus to output and inflation during the crisis by combining the fiscal transfer with a commitment by the central bank to keep the interest rate at its lower bound, *even if* this meant over-shooting the inflation target, as long as the financial wedge remains large. (We continue to suppose, for now, that monetary policy returns to the IT regime as soon as fundamentals revert to normal.) Thus a coordinated change in both fiscal and monetary policy can achieve more than the sum of the effects of the two policies considered individually. It is also notable that under this combined policy, the commitment to allow inflation to over-shoot increases aggregate demand even though, in our model simulations, inflation does not ever over-shoot the target; the commitment matters only because of its effects on the (incorrect) calculations of myopic decision-makers.

III. Forward Guidance

Even in the absence of any change in fiscal policy, the contractionary effects of the financial shock can be reduced if it is possible to make credible commitments about the conduct of monetary policy *after* the ZLB ceases to prevent attainment of the inflation target. Here we consider the effects of alternative commitments regarding monetary policy after reversion to the normal state. These may be thought of as ad hoc commitments, announced only once the financial shock has occurred; in particular, we assume that policy prior to date zero has been the simple IT regime, and that the value functions of households and firms are the ones appropriate to that regime, regardless of the new policy that may be announced during the crisis.

The effects of such "forward guidance" on aggregate demand during the crisis obviously depend on people's being able to reason about future conditions using the information provided about the new policy. Hence an assumption of finite planning horizons will weaken the stimulative effects of such a policy, relative to the predictions of the RE analysis of Eggertsson and Woodford (2003), as discussed in Woodford (2018). Nonetheless, as long as horizons are not too short, appropriate forward guidance should help to mitigate the effects of a financial shock despite the binding ZLB. Moreover, the fact that the policy should stimulate demand and price increases more in the case of households and firms with longer horizons is desirable, insofar as these are precisely the parties whose behavior would otherwise be most affected by the financial shock; thus the heterogeneous effects of forward guidance reduces the distortions that would otherwise result from heterogeneous responses to the financial shock.

This can be illustrated by considering the effects of a commitment to keep the nominal interest rate at its lower bound until the price level is restored to a *target path* that grows deterministically at the target inflation rate π^*. If the ZLB prevents the inflation target from being achieved during the crisis period, the price level will fall below this target path; the commitment then requires the interest rate to remain at the lower bound for a time even after the financial wedge becomes small again, even though this causes inflation above the rate π^* for a time. Once the "price-level gap" has been closed, we suppose that monetary policy is used to achieve the target inflation rate π^* thereafter. Since policy after the gap has been closed is the same as under the IT policy, this commitment is equivalent to a "temporary price-level target" (TPLT) of the kind proposed by Bernanke (2017).

Figure 2 compares the solution under this policy to that under the simple IT policy, in the case that the financial shock lasts for ten quarters; the two simulations are compared for each of three possible lengths of planning horizons. The solid line in each panel shows the response of the price-level gap $\tilde{p}_t \equiv \log(P_t/P_t^*)$ under IT; the price-level gap becomes steadily more negative while the shock persists, and is then stabilized at a permanent negative level once fundamentals revert to normal. The dashed line in each panel shows the alternative dynamics under the TPLT commitment; the price level falls behind the target path to a much lower extent during the crisis, in addition to being eventually returned to the target path within a few quarters of reversion to the normal state. The smaller price-level gap during the crisis corresponds to more successful output stabilization as well, as shown in the online Appendix. Notably, the predicted dynamics are also more similar across different values of h in the case of this policy, implying less distortions from heterogeneous responses in the case of heterogeneous planning horizons.

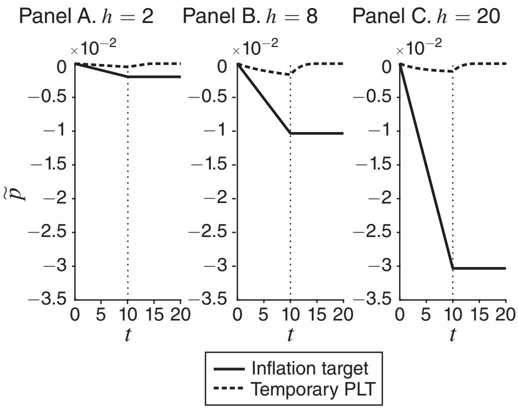

Figure 2. Price-Gap Dynamics under Inflation Targeting versus Temporary PLT

Note: Response of the price-level gap when shock lasts for ten quarters, under two alternative policy commitments when the shock occurs.

IV. Ad Hoc Commitments versus Policy Rules

In order to obtain the advantages of commitment to a PLT when the ZLB binds, is it necessary for a central bank to *always* conduct its policy in accordance with a PLT, even when financial constraints are minimal, or does it suffice to announce a TPLT policy on an ad hoc basis only when a financial shock that causes the ZLB to bind occurs? Under the RE analysis, the two kinds of policy should achieve *identical* outcomes during a crisis period and the immediately following period in which the price-level gap is being closed. Since the advantages of a PLT over an IT regime are most compelling in the case of such a crisis, it might be thought that a TPLT policy offers the more prudent approach.

However, in the limited-foresight analysis there is an advantage to a systematic PLT rule over an ad hoc commitment: this is that a different systematic approach to policy in the period before a financial crisis occurs can change the value functions that households and firms learn, and then apply in their forward planning during the crisis period. In the case of commitment to a PLT rule, we continue to assume that the value functions are not updated during a temporary period in which the ZLB binds, but we suppose that the value functions that have been learned are the ones that would be optimal under a PLT

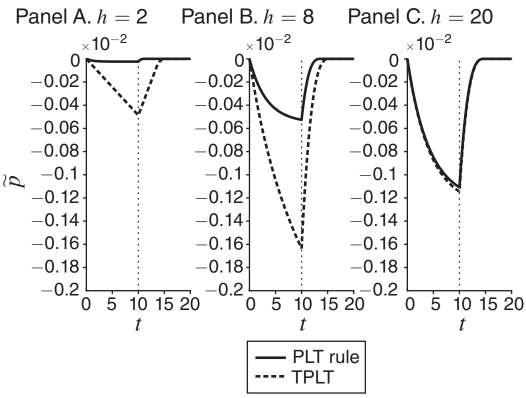

FIGURE 3. PRICE-GAP DYNAMICS UNDER TEMPORARY PLT VERSUS PLT RULE

Note: Response of the price-level gap when shock lasts for ten quarters, under two alternative assumptions about policy prior to the shock.

regime in which the ZLB never binds. In these latter value functions, the price-level gap is recognized as a crucial state variable (affecting the expected marginal value of wealth).

Figure 3 shows simulated responses to the same history of fundamentals as in Figure 2, but now comparing outcomes under the TPLT to the outcomes under a systematic PLT rule; the difference in the two cases reflects not any difference in the expected conduct of monetary policy during and after the crisis period, but a difference in the value functions used in the forward planning exercises of households and firms in the two cases.

Allowing households and firms to learn the value functions appropriate to a PLT regime (the solid lines in each panel) results in less of a decline in prices (and associated with this, less contraction of output) when the ZLB binds. The difference is particularly great when planning horizons are relatively short. If we assume planning horizons of five years (the third panel) or more, the difference made by allowing different value functions to be learned is minimal, but the difference is considerable if h is only a few quarters.

Some might assume that recognizing limitations on people's ability to correctly deduce the future consequences of a new policy should reduce the theoretical benefits of commitment to policy rules, and hence favor a purely discretionary approach to policy. The comparison in Figure 3 shows that this need not be the case. In our analysis, recognizing that planning horizons may be short reduces the predicted efficacy of ad hoc commitments in response to a special situation—and so increases the case for seeking to ensure that the "default" expectations implicit in people's value functions are ones that help to stabilize the economy during a crisis. These "default" expectations are best shaped by systematic action in accordance with a relatively simple policy rule, since they are learned by induction from past experience, rather than being derived through deductive reasoning about a concrete current situation.

REFERENCES

Bernanke, Ben S. 2017. "Monetary Policy in a New Era." Presented at Rethinking Macroeconomic Policy Conference at the Peterson Institute, Washington, DC, October 12–13.

Christiano, Lawrence, Martin Eichenbaum, and Sergio Rebelo. 2011. "When Is the Government Spending Multiplier Large?" *Journal of Political Economy* 119 (1): 78–121.

Eggertsson, Gauti B. 2010. "What Fiscal Policy Is Effective at Zero Interest Rates?" *NBER Macroeconomics Annual 2010*, Vol. 25, edited by Daron Acemoglu and Michael Woodford, 59–122. Chicago: University of Chicago Press.

Eggertsson, Gauti B., and Michael Woodford. 2003. "The Zero Bound on Interest Rates and Optimal Monetary Policy." *Brookings Papers on Economic Activity* 34 (1): 139–211.

Woodford, Michael. 2011. "Simple Analytics of the Government Expenditure Multiplier." *American Economic Journal: Macroeconomics* 3 (1): 1–35.

Woodford, Michael. 2018. "Monetary Policy Analysis when Planning Horizons are Finite." National Bureau of Economic Research Working Paper 24692.