

Fiscal and Monetary Stabilization Policy at the Zero Lower Bound: Consequences of Limited Foresight

by Michael Woodford¹ and Yinxi Xie²

¹ Columbia University

² Canadian Economic Analysis Department
Bank of Canada, Ottawa, Ontario K1A 0G9

mw2230@columbia.edu, YinxiXie@bankofcanada.ca



Bank of Canada staff working papers provide a forum for staff to publish work-in-progress research independently from the Bank's Governing Council. This research may support or challenge prevailing policy orthodoxy. Therefore, the views expressed in this paper are solely those of the authors and may differ from official Bank of Canada views. No responsibility for them should be attributed to the Bank.

Acknowledgements

Prepared for the Carnegie–Rochester–NYU Conference on Public Policy, “Central Banking in the 2020s and Beyond.” We thank Luigi Iovino, Narayana Kocherlakota, Chen Lian, Rosemarie Nagel, David López-Salido, an anonymous referee and participants at the Carnegie–Rochester–NYU Conference on Public Policy for helpful comments. The views in this paper are those of the authors and do not necessarily reflect those of the Bank of Canada.

Abstract

This paper reconsiders the degree to which macroeconomic stabilization is possible when the zero lower bound is a relevant constraint on the effectiveness of conventional monetary policy, under an assumption of bounded rationality. In particular, we reconsider the potential role of countercyclical fiscal transfers as a tool of stabilization policy. Because Ricardian equivalence no longer holds when planning horizons are finite (even when relatively long), we find that fiscal transfers can be a powerful tool to reduce the contractionary impact of an increased financial wedge during a crisis and can even make possible complete stabilization of both aggregate output and inflation under certain circumstances, despite the binding lower bound on interest rates. However, the power of such policies depends on the degree of monetary policy accommodation. We also show that a higher level of welfare is generally possible if both monetary and fiscal authorities commit themselves to history-dependent policies in the period after the financial disturbance that causes the lower bound to bind has dissipated.

Topics: Monetary policy; Fiscal policy; Central bank research; Business fluctuations and cycles

JEL codes: E52, E63, E7

1 Introduction

The events of the period since the financial crisis of 2008 have required a significant reappraisal of the previous conventional wisdom, according to which interest-rate policy alone — and more specifically, a policy of adjusting the central bank’s operating target for a short-term interest rate in response to contemporaneous economic conditions (as proposed, for example, by Taylor, 1993) — should suffice to maintain macroeconomic stability. It has become evident that conventional interest-rate policy will often be constrained by the zero lower bound (ZLB) on nominal interest rates.¹ One consequence has been a greater willingness on the part of central banks to make statements about likely future interest-rate policy, even years in advance, as a substitute for further immediate interest-rate reduction. But another has been a revival of interest in the use of counter-cyclical fiscal policy for macroeconomic stabilization.

An important research literature since the crisis has supported the view that fiscal stabilization policy can be especially valuable when interest-rate policy is constrained by the ZLB. This literature has mainly addressed the effects of countercyclical government purchases,² rather than government transfers, on the grounds that in simple representative-agent New Keynesian models, Ricardian equivalence holds in the case of lump-sum taxes and transfers; lump-sum transfers during a crisis, if expected to be financed by future lump-sum taxes, should have no effect at all on economic activity or inflation.³ On the other hand, the actual fiscal stimulus packages enacted in response to the crisis consisted, to an important extent, of increases in government transfers (Taylor, 2018). This makes it important to further consider the potential role of countercyclical government transfers as a tool of stabilization policy.

The Ricardian Equivalence result in standard treatments depends crucially on an assumption of rational expectations on the part of all decision makers. Yet the grounds for assuming rational expectations in such a case are especially weak. To the extent that a fiscal stimulus package is an *ad hoc* response to a single crisis, rather than an implication of a systematic policy of adjusting the government’s budget in

¹The effective lower bound need not be exactly zero; for purposes of our argument here, it is only important that there be some lower bound, and that in certain situations even reducing the policy rate to that lower bound will provide an insufficient stimulus to aggregate demand.

²See, for example, Eggertsson (2011), Christiano *et al.* (2011), or Woodford (2011).

³An exception is the paper of Ascari and Rankin (2013), who instead analyze the question under an assumption of rational expectations, but using an overlapping generations model. We show that the effectiveness of fiscal stimulus need not depend critically on the demographic structure (or the absence of bequest motives), once one allows for finite planning horizons.

response to the business cycle, one cannot expect that people should have rational expectations as a result of learning from experience; instead, one needs to ask what people should be able to deduce from reasoning about the predictable effects of a novel policy.

Moreover, in order for it to make sense to suppose that people should anticipate the future tax increases that must result from the increased public debt occasioned by the stimulus policy, one must assume not merely that people are capable of forward planning (taking into account the policy change), but that their forward planning extends over quite a long horizon — as long as is required for the increased public debt to be fully paid off. This means that in order for full (or nearly full) Ricardian Equivalence to obtain, one needs to assume that most people’s planning horizons extend far into the future.

In this paper, we reconsider the usefulness of government transfers as a tool of stabilization policy, and the issue of coordination between monetary and fiscal policies, 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 (2019), 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, 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 the end period of foresight 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.

Interim situations that someone imagines reaching in the end period of foresight 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 is the one 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 rational-expectations analysis (at least, under a suitable equilibrium selection criterion), as discussed in Woodford (2019). However, a situation in which monetary policy is constrained by the 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.

This suggests that in a model with finite planning horizons, countercyclical fiscal stimulus might be a powerful tool, and indeed one that might make it possible to stabilize the economy despite the lower bound on interest rates, without any need to resort to commitments about future monetary policy. Here we consider what can be achieved by state-contingent transfer policies⁴ when people's planning horizons are finite (and perhaps extend only a few quarters into the future). We focus on possible responses to a temporary increase in the size of a financial wedge (intended to capture a situation like the 2008 crisis) that is large enough to prevent complete stabilization using interest-rate policy alone, owing to the ZLB.⁵

We show that fiscal transfers can indeed reduce the contractionary impact of an increase in the financial wedge, and that, at least under some circumstances, a willingness to use fiscal policy with sufficient aggressiveness makes it possible to achieve complete stabilization of both aggregate economic activity and the overall rate of inflation, despite the zero lower bound and regardless of the size of the increase in the financial wedge. Thus the existence of state-contingent transfer policies expands the degree to which stabilization would be possible using interest-rate policy alone; and we obtain this result under conditions that would guarantee Ricardian Equivalence under an assumption of rational expectations.

At the same time, we show that it would be a mistake to conclude that countercyclical transfers are so effective a tool that there is no need for a central bank to ever indicate that it would allow inflation to overshoot the bank's long-run inflation target, nor any need for a commitment to conduct future interest-rate policy in any way different from what will best serve the bank's goals at that future date. We find that state-contingent transfers make possible equilibria that could not be achieved

⁴Woodford and Xie (2019) similarly show how the effects of government purchases depend on the length of decision makers' planning horizons.

⁵While we refer to the shock that creates the crisis in our model as an increase in the financial wedge, similar conclusions would be reached in the event of a reduction for other reasons in the real interest rate required to induce an efficient level of expenditure. For example, the ZLB has again become a binding constraint on US monetary policy in 2020 as a result of the COVID-19 pandemic, though in this case the disturbance did not originate in the financial sector. If a shock of this kind increases uncertainty about future economic conditions, the result can be an increase in precautionary saving that affects the equilibrium real interest rate in a way similar to the shock to the financial wedge in the model set out here.

using interest-rate policy alone, but that there is a limit to the stimulus that can be achieved even by massive fiscal transfers, in the absence of monetary accommodation — that is, a commitment *not* to raise interest rates, even if inflation overshoots its long-run target.⁶

We also find that there is a limit to what can be achieved, even by coordinated fiscal and monetary policy, if policy is expected to return to pursuit of its normal targets as soon as the wedge returns to a normal level. A higher level of welfare is possible, in general, if the monetary and fiscal authorities commit themselves to history-dependent policies in the period after the real disturbance has dissipated. Hence forward guidance — a commitment to conduct monetary policy differently in the future than would be the case in normal times — remains valuable even when fiscal transfers are also available as a tool of stabilization policy.

Our approach to introducing bounded rationality into a New Keynesian model is related to a number of contributions to an active recent literature that has offered reasons for central bank forward guidance about monetary policy years in the future to have weaker effects than those predicted by a New Keynesian model under the assumptions of rational expectations and full credibility of the policy announcement — a prediction of the standard model that is widely viewed as inconsistent with available evidence (the “forward guidance puzzle,” Del Negro *et al.*, 2015).⁷ Perhaps most obviously, Gabaix (2020) proposes a framework in which predictable future states are assumed to be down-weighted in their effects on agents’ decision rules, relative to the optimal decision rule under rational expectations, to an extent that is greater the greater the distance in the future of these states. This leads to predictions that are in many ways similar to those of our model with finite-horizon planning (especially if

⁶Our conclusions about the importance of monetary accommodation of fiscal stimulus are in line with those of Ascari and Rankin (2013), though the reason for Ricardian Equivalence to fail is different in our analysis.

⁷Not all of the proposed resolutions of the “forward guidance puzzle” have direct implications for the effects of fiscal policy considered here, however. For example, Andrade *et al.* (2019) discuss how forward guidance regarding interest-rate policy can fail to have the desired effect when the policy commitment is not correctly understood or believed by everyone in the private sector, and Gust *et al.* (2018) show how the effects of forward guidance can be attenuated by learning dynamics. These issues are also important for predictions of the effect of fiscal policy, but the relevant grounds for misunderstanding of the implications of a policy announcement are not the same for different kinds of policy.

we assume an exponential distribution of planning horizons); our framework can be viewed as providing an explicit model of cognitive processing that might underlie the kind of reduced-form effects proposed by Gabaix.

The implications of the two models are, however, not equivalent.⁸ A difference that matters for the current discussion is that while an exponential distribution of planning horizons in our model leads to similar predictions as the representative-agent model of Gabaix regarding the possibility of complete stabilization of both output and inflation using fiscal policy (as discussed in section 3), the welfare implications are not the same in the two models; thus the micro-foundations matter for policy design, even when the predicted behavior of macroeconomic aggregates is similar.

The effects of “level- k reasoning” (García-Schmidt and Woodford, 2019; Farhi and Werning, 2019; Iovino and Sergeyev, 2020) also have important similarities to the effects of finite-horizon planning. In the “level- k ” models, decision makers formulate infinite-horizon plans but assume that the future values of aggregate variables will be determined by the behavior of others who do not take announced policy changes into account to the same degree as they themselves do.⁹ This type of departure from rational-expectations analysis (motivated by work in experimental game theory by authors such as Nagel, 1995) is in the same spirit as our proposal of finite-horizon planning, in that both proposals truncate the length of a decision maker’s chain of deductive reason, to reduce the complexity of the required calculations; moreover, the two proposals are complementary, rather than alternatives. A more general model would allow both for truncation of forward planning after a finite number of steps and truncation of the number of levels of higher-order expectations that are taken into account.

We do not here present calculations with a limited level of reasoning, because this departure from standard methodology only matters much when planning horizons are long (as in the analyses of García-Schmidt and Woodford or Farhi and Werning);

⁸The differences are particularly notable in the case of long-lasting policy changes. For example, the Gabaix model yields “neo-Fisherian” long-run predictions that are quite different from those implied by finite-horizon planning, as discussed in Woodford (2019).

⁹In related work, Angeletos and Lian (2018) show how the aggregate effects of a policy announcement can be attenuated when economic conditions are perceived with noise by individual decision makers. While this analysis is based on imperfect information rather than a bound on the complexity of reasoning, it similarly reduces the extent to which higher-order expectations adjust in response to a policy change.

if planning horizons are short, this fact already limits the degree to which decisions depend on higher-order expectations. And while the two approaches lead to similar conclusions in some respects,¹⁰ the finite-horizon planning approach has the advantage of providing an integrated analysis of both short-run and long-run effects of a policy change¹¹ and offers the prospect of explaining lower-frequency macroeconomic dynamics in addition to the effects of unusual individual policy experiments.¹²

The paper proceeds as follows. Section 2 describes the New Keynesian DSGE model with finite planning horizon and the financial shocks considered in this paper. As a baseline to which more active stabilization policies can be compared, it analyzes the effects of such shocks, under alternative assumptions about the length of planning horizons, when monetary policy is specified by a purely forward-looking inflation target and strict budget balance is maintained. Section 3 then introduces fiscal transfer policies, while section 4 considers what can be achieved through coordinated fiscal and monetary stabilization policies. Section 5 concludes.

2 Output and Inflation Determination with Finite Planning Horizons

2.1 Forward Planning with a Finite Horizon

We study the consequences of limited foresight in a New Keynesian DSGE model with finite-horizon forward planning, building upon the approach developed in Woodford (2019). Households and firms make contingent plans for a finite distance into the future, and use a value function learned from past experiences to evaluate all possible terminal states in the last period of the planning horizon. Over this horizon, they

¹⁰As an example, note the similarity of the conclusions in García-Schmidt and Woodford (2019) and those in Woodford (2019) regarding the validity of “neo-Fisherian” predictions in the case of an interest-rate peg.

¹¹As García-Schmidt and Woodford (2019) discuss, the “level- k ” approach only offers an analysis of the immediate effects of the announcement of a novel policy and does not answer the question of how reasoning should change as people observe that outcomes under the new regime differ from those that they had expected.

¹²A version of the model proposed in Woodford (2019) is empirically estimated in Gust *et al.* (2019), who find that it is as good as or better than more ad hoc modifications of the rational-expectations New Keynesian model in accounting for US macroeconomic dynamics.

use structural knowledge (including any announcements about novel central bank or government policies) to deduce the consequences of their intended actions. For simplicity, we assume that the planning horizon is taken to be exogenously fixed.

2.1.1 The household decision problem

We illustrate the approach by briefly discussing here the problem of households in our model.¹³ As in standard New Keynesian models, we assume an economy made up of infinite-lived households, here assumed to be identical apart from possible differences in their planning horizons. But we suppose that at any date t , a state-contingent expenditure plan is selected only for dates between t and some date $t + h$, a finite distance in the future.

Letting C_τ^i be household i 's planned consumption in period τ of a composite good (a CES aggregate of the many differentiated goods produced in the economy), we suppose that at time t the household chooses state-contingent values $\{C_\tau^i\}$ for each of the dates $t \leq \tau \leq t + h$ (specifying real expenditure in each of the exogenous states that may arise at any of those dates, given the state of the world at the time of the planning) so as to maximize the expected value (according to the household's calculations at time t) of an objective of the form

$$\sum_{\tau=t}^{t+h} \beta^{\tau-t} u(C_\tau^i) + \beta^{h+1} v(B_{t+h+1}^i; s_{t+h}).$$

Here the first terms represent the discounted sum of flow utilities from consumption in periods t through $t + h$, while the final term represents the household's estimate of the value of the discounted sum of flow utilities that it can expect to receive in later periods, if the wealth that it holds at the end of the planning horizon is B_{t+h+1}^i . We allow in general for the possibility that the value assigned to the household's continuation problem after period $t + h$ may depend on the state of the world s_{t+h} that has been reached in period $t + h$.

In the household's planning exercise, it takes into account its budget constraint, and thus the way in which the value of B_{t+h+1}^i will depend on its planned level of expenditure. We assume that there exists only a single financial asset each period, a one-period riskless nominal debt instrument, the interest rate i_t on which is also

¹³The decision problem of price-setting firms is treated using similar methods in Woodford (2019).

the central bank's policy instrument. Because wealth can take this single form, the implications of the household's choices over the planning horizon for the value of its continuation problem can be summarized by a single quantity, B_{t+h+1}^i , indicating the wealth carried into period $t + h + 1$ in the form of this riskless nominal asset.

The evolution of this quantity is determined by a flow budget constraint of the form

$$B_{\tau+1}^i = (1 + i_\tau + \Delta_\tau) [B_\tau^i / \Pi_\tau + Y_\tau + T_\tau - C_\tau^i] - \Delta_\tau [B_\tau / \Pi_\tau + T_\tau] \quad (2.1)$$

for each period $t \leq \tau \leq t + h$. Here B_τ^i is the value of the nominal debt held by the household that matures at date τ , deflated by the period $\tau - 1$ price index $P_{\tau-1}$,¹⁴ so that it is a predetermined real variable.¹⁵ This quantity must be deflated by $\Pi_\tau \equiv P_\tau / P_{\tau-1}$, the gross inflation rate between $\tau - 1$ and τ , to obtain the real value of the maturing debt in units of the period τ composite good. The term Y_τ indicates production of the composite good in period τ , the value of which is received as income by the households (and treated as independent of any household decision, in the household's forward planning exercise); and T_τ is the value of lump-sum government transfers (the same to each household), also in units of the composite good.

End-of-period asset balances earn a nominal financial yield of i_τ between periods τ and $\tau + 1$. In addition, we assume that there is an additional benefit of holding riskless claims, which we represent in (2.1) as an additional dividend equal to Δ_τ per unit of savings held in this form. This additional dividend is intended to represent the existence of a (time-varying) safety premium, as in the models of Del Negro *et al.* (2017) and Caballero and Farhi (2017); increases in the size of such a premium are an important reason for the lower bound on the safe nominal interest rate to become a tighter constraint during financial crises.¹⁶

¹⁴As usual, this price index is the minimum cost at which a unit of the composite good can be purchased in period $\tau - 1$.

¹⁵In particular, B_{t+h+1}^i is a quantity that is determined in period $t + h$, and thus within the planning horizon, like the other terms in the objective function.

¹⁶The only consequence of a non-zero value of Δ_τ in our model is the introduction of a time-varying factor in the household Euler equations (2.4)–(2.5) below. The same kind of exogenous shift factor in the Euler equation could alternatively arise from exogenous variation in households' rate of time preference, as assumed in Eggertsson and Woodford (2003). While the latter assumption would allow for a simpler and more conventional model, we believe that variation in the size of the financial wedge represented by Δ_τ provides a more realistic picture of the kind of disturbance that is likely to give rise to the policy challenges that we address in this paper.

The final term in (2.1) is a lump-sum effective tax on households, equal in size to the safety dividend received by households in aggregate (using the notation B_τ for the aggregate supply of public debt carried into period τ). This indicates that the advantages to an individual household of holding more safe assets are at the expense of other households (as the “safety dividend” does not correspond to any additional resources created by the safe assets). We model the safety premium $\{\Delta_t\}$ as an exogenous process, satisfying $\Delta_t \geq 0$ at all times; we do not consider in this paper the possibility of government policies that can directly affect the size of this wedge.

2.1.2 Model-based expectations

The decision problem of a household at time t depends on the financial wealth B_t^i that it brings into the period and on the household’s expectations about the state-contingent evolution of the variables $\{\Pi_\tau, Y_\tau, T_\tau, i_\tau, \Delta_\tau\}$ over periods $t \leq \tau \leq t + h$, that is, the household’s planning horizon. We assume that in their forward planning exercises, households make use of correct structural information about how the economy works (including a correct understanding of monetary and fiscal policy, taking into account any new policies that may have been announced in response to an unexpected exogenous disturbance).

First, we assume a correct understanding of the state-contingent evolution of all exogenous state variables; this means that households correctly understand the current value of Δ_t (since they know the economy’s exogenous state before undertaking forward planning) and the conditional probability of different possible future paths $\{\Delta_\tau\}$.

Second, households are assumed to correctly understand the rules that will determine the policy variables $\{T_\tau, i_\tau\}$ over the planning horizon. For simplicity, we restrict attention in this paper to fiscal rules under which the path of the real public debt $\{B_\tau\}$ is exogenously specified;¹⁷ this allows us to consider both the case of no public debt (often assumed in analyses of alternative monetary policies) and various ways in which the level of public debt might depend on the path of the financial wedge $\{\Delta_\tau\}$. Aggregating over households, and assuming no government purchases,¹⁸

¹⁷See Xie (2020) for analysis of regimes in which there is instead feedback from endogenous variables to the path of real public debt, including “active” fiscal policy regimes according to the classification of Leeper (1991).

¹⁸The framework can easily be extended to allow for government purchases as well. See Woodford

it follows from (2.1) that the evolution of the real public debt must satisfy

$$B_{\tau+1} = (1 + i_\tau)[B_\tau/\Pi_\tau + T_\tau] \quad (2.2)$$

for each of the periods $t \leq \tau \leq t + h$. This, like other structural equations of our model, is assumed to be correctly understood by households. Then the assumption that fiscal policy is specified by an exogenous process $\{B_\tau\}$ implies that T_τ must endogenously adjust, to ensure that (2.2) is satisfied, in response to any changes in i_τ by the central bank, or changes in Π_τ as a result of firms' pricing decisions.

We similarly assume that households correctly understand the way in which i_τ will be determined under any contingency by the central bank's policy. For example, if the central bank follows a Taylor rule, then the state-contingent evolution assumed in a household's forward planning will necessarily satisfy that relation. Any feasible policy is assumed to be subject to a ZLB constraint

$$i_t \geq 0 \quad (2.3)$$

at all times.

Finally, households are also assumed to correctly understand how the variables Y_τ and Π_τ are determined by the decisions of households and price-setting firms respectively. However, in order not to have to model how the economy should evolve (or anyone else should be modeling it to evolve) beyond the horizon $t + h$, a household with horizon h at time t must model Y_τ and Π_τ as being determined by households and firms who do not look beyond the horizon $t + h$ while making their decisions at time τ . Just as the household, in its planning at time t , models its own behavior at some later date τ as the behavior that will appear optimal to someone with a planning horizon at that time of only $t + h - \tau$ periods, it similarly models the behavior of other households and firms at date τ under the assumption that they will all have planning horizons of $t + h - \tau$ periods. This means that the household will model all other households as spending the same amount at time τ as it plans itself to spend at that time. Hence the amount of income Y_τ that it expects to receive in any future state will be the same as the amount C_τ^i that it expects to spend in that state.

Let Y_t^j, Π_t^j, i_t^j be the (counterfactual) output, inflation, and nominal interest rate in the case that all economic units (households and firms) have a planning horizon

and Xie (2019) for analysis of how the government purchases multiplier is affected by finite planning horizons.

of $j \geq 0$ periods at time t . Then the Euler equation for optimal forward planning requires that for any $j \geq 1$,

$$u'(Y_t^j) = \beta(1 + i_t^j + \Delta_t) E_t[u'(Y_{t+1}^{j-1})/\Pi_{t+1}^{j-1}], \quad (2.4)$$

while for $j = 0$,

$$u'(Y_t^0) = \beta(1 + i_t^0 + \Delta_t) v'(B_{t+1}). \quad (2.5)$$

In (2.5) we use the fact that in equilibrium, a household with planning horizon zero must anticipate an interest rate i_t^0 that leads it to choose to hold wealth B_{t+1}^0 equal to the exogenously specified supply of public debt B_{t+1} (given that it expects other households to optimize over the same planning horizon as it does, and it expects the debt market to clear).

Thus we obtain a system of equations that can be recursively solved for the state-contingent evolution of the variables $\{Y_t^j\}$ for each possible horizon $j \geq 0$, given the state-contingent evolution of the endogenous variables $\{\Pi_t^j, i_t^j\}$ for all j , and the state-contingent evolution of the exogenous variables $\{\Delta_t, B_{t+1}\}$, along with any exogenous disturbances to the monetary policy rule.¹⁹ (Equation (2.5) can be solved for the value of Y_t^0 in any state of the world, given the values of the other variables; then, given a solution for the state-contingent evolution of $\{Y_\tau^0\}$, the $j = 1$ case of equation (2.4) can be solved for the value of Y_t^1 in any state of the world; and so on for progressively higher values of j .)

Modeling the optimizing decision of price-setting firms with finite planning horizons, we similarly obtain a system of equations that can be recursively solved for the state-contingent evolution of the variables $\{\Pi_t^j\}$ for each possible horizon $j \geq 0$, given the state-contingent evolution of the endogenous variables $\{Y_t^j, i_t^j\}$ and the state-contingent evolution of the exogenous variables. These equations, together with the monetary policy rule with which the endogenous variables must be consistent for each value of j , provide a system that can be jointly solved for the state-contingent evolution of the endogenous variables $\{Y_t^j, \Pi_t^j, i_t^j\}$ for each possible horizon $j \geq 0$, given the state-contingent evolution of the exogenous variables.

In writing the above equations, we take as given the value function $v(B)$ that households will use in their forward planning, and similarly the value function that

¹⁹The model can easily be extended to allow for exogenous disturbances to productivity, preferences, and government consumption, as treated in Woodford (2019), but in this paper, we are concerned only with possible policy responses to disturbances to the financial wedge Δ_t .

firms will use. In Woodford (2019), the endogenous evolution of these value functions in response to additional experience is also modeled; here, however, we abstract from this additional source of dynamics and assume fixed value functions that will be the same for the different policies that we consider. Our assumption is that the value functions are determined in a backward-looking way (as an inference from outcomes observed in the past) and not through a forward-looking deductive process; the whole point of the use of a value function to evaluate conditions that might be reached at the planning horizon $t+h$ is to avoid having to reason deductively about what should happen under various contingencies beyond that date.

Thus when an unusual shock hits, and unusual policies are announced in response, the value functions that households and firms use, at least initially, will continue to be ones that they learned from macroeconomic conditions prior to either the disturbance or the new policies.²⁰ Because our concern in this paper is solely with the effects of temporary policy changes in response to a transitory disturbance, we simplify the discussion by abstracting from the changes in the value functions that would eventually occur if the new conditions were to persist sufficiently long.²¹ Instead we assume that the value functions remain fixed over the scenarios that we consider below and are ones that represent an optimal adaptation to the stationary conditions assumed to have existed prior to the disturbance.

In the analyses below, the situation prior to the disturbance is assumed to have been the one in which the government debt has been zero ($B_t = 0$ at all times); the central bank has pursued a forward-looking inflation targeting policy, setting i_t each

²⁰If crises of a similar sort occur repeatedly and similar policies are adopted each time, one might expect that the value function used when such a crisis arrives should eventually adapt to this experience. We do not pursue this extension of the analysis here, but see the analysis in Woodford and Xie (2019) of the effects on equilibrium dynamics during a crisis of learning to expect compliance with a price-level targeting rule.

²¹Allowing the value functions to adapt is instead critical for certain other kinds of discussions. These include consideration of the eventual effects of commitment to an interest-rate peg for a long period of time, as in Woodford (2019); empirical modeling of US economic data over a period of decades that included significant shifts in both output and inflation trends, as in Gust *et al.* (2019); analysis of the conditions under which joint monetary-fiscal policy regimes imply sustainable long-run dynamics, as in Xie (2020); and consideration of the difference between commitment to a systematic price-level targeting rule and adoption of an *ad hoc* “temporary price-level target” when the ZLB binds, as in Woodford and Xie (2019).

period at the level required to ensure that $\Pi_t = \Pi^*$, the long-run inflation target,²² and the financial wedge Δ_t has at all times been small enough to make it possible for the central bank to achieve that target without violating the zero lower bound (2.3). In a stationary equilibrium in which these conditions always hold, the maximum attainable discounted utility for a household that enters period t with wealth B is given by

$$v(B) = \frac{1}{1 - \beta} u(\bar{Y} + (1 - \beta)B/\bar{\Pi}), \quad (2.6)$$

where \bar{Y} and $\bar{\Pi}$ are the stationary values of Y_t and Π_t .

This is the optimal value function for households in this stationary environment; its use in a finite-horizon planning exercise in the stationary environment would result in optimal behavior, regardless of the length of the planning horizon. It is also the value function to which the adaptive process described in Woodford (2019) would converge, if such an environment were maintained for a sufficiently long time. Thus we assume the value function (2.6) for households in our analyses below; we similarly assume for firms a value function that is optimally adapted to that same stationary environment.

2.1.3 Log-linear approximate dynamics

As in many rational-expectations analyses, it will be convenient to approximate the solution to the model structural equations using a log-linear approximation. We linearize the model's equations around a stationary equilibrium in which $\Delta_t = 0$ at all times, and the policy regime is the one discussed above for which the value functions of households and firms are adapted. We express the linearized equilibrium relations in terms of deviations from the stationary equilibrium values of the various state variables, using the following notation:

$$y_t^j \equiv \log(Y_t^j/\bar{Y}), \quad \pi_t \equiv \log(\Pi_t/\bar{\Pi}), \quad b_t \equiv B_t/(\bar{\Pi}\bar{Y}),$$

$$\hat{i}_t \equiv \log\left(\frac{1 + i_t}{1 + \bar{i}}\right), \quad \hat{\Delta}_t \equiv \frac{\Delta_t}{1 + \bar{i}}.$$

Here $\bar{i} \equiv \beta^{-1}\bar{\Pi} - 1 > 0$ is the stationary equilibrium value of the nominal interest rate.

²²This target is assumed to satisfy $\Pi^* > \beta$, so that a stationary equilibrium is possible in which this inflation rate is maintained at all times, and in this equilibrium, the ZLB constraint (2.3) is a strict inequality. Note that this will be satisfied in the case of any non-negative inflation target.

In terms of this notation, equilibrium conditions (2.4) and (2.5) can be log-linearized to yield

$$y_t^j = -\sigma(\hat{i}_t^j + \hat{\Delta}_t - \mathbb{E}_t \pi_{t+1}^{j-1}) + \mathbb{E}_t y_{t+1}^{j-1} \quad (2.7)$$

for each $j \geq 1$, and

$$y_t^0 = -\sigma(\hat{i}_t^0 + \hat{\Delta}_t) + (1 - \beta)b_{t+1}. \quad (2.8)$$

Note that except for the superscripts, (2.7) has the same form as the “New Keynesian IS equation” obtained under rational expectations (see, e.g., Woodford, 2003, chap. 4).

Similarly, the structural relations describing optimal price-setting behavior by firms can be log-linearized to yield

$$\pi_t^j = \kappa y_t^j + \beta \mathbb{E}_t \pi_{t+1}^{j-1} \quad (2.9)$$

for each $j \geq 1$, and

$$\pi_t^0 = \kappa y_t^0. \quad (2.10)$$

(See Woodford, 2019, for the derivation.) Here again, it will be observed that except for the superscripts, (2.9) has the same form as the “New Keynesian Phillips curve” obtained under rational expectations (Woodford, 2003, chap. 3).

Up to a log-linear approximation, the predicted evolution of aggregate variables is then given by

$$y_t \equiv \sum_h \omega_h y_t^h, \quad \pi_t \equiv \sum_h \omega_h \pi_t^h,$$

where ω_h is the fraction of both households and firms each period with planning horizon h (for all $h \geq 0$). We treat the frequency distribution $\{\omega_h\}$ as exogenously given in the exercises reported here.²³ In some of our numerical results, we assume an exponential distribution,

$$\omega_h = (1 - \rho)\rho^h, \quad (2.11)$$

for all $h \geq 0$, where $0 < \rho < 1$, though our methods do not depend on this.

Finally, in terms of the deviations variables, the zero lower bound constraint can be written as

$$\hat{i}_t \geq \underline{\hat{i}}, \quad (2.12)$$

²³There is no logical reason why this distribution needs to be the same for households and firms, but in the results presented here, we simplify the reporting of results by assuming the same planning horizons for both.

where $\hat{\imath} < 0$, meaning that the constraint does not bind when i_t is near its stationary equilibrium value \bar{i} .²⁴

2.1.4 How long are planning horizons?

The quantitative relevance of the departure from rational expectations that we propose depends, of course, on how long we assume planning horizons to be. Because our goal in this paper is to clarify how the assumed planning horizon matters, we do not take a stand on the most plausible numerical assumption about planning horizons.²⁵ There is, however, good reason to believe both that people are capable of some degree of forward planning and at the same time that planning horizons are often not too long.

Keramati *et al.* (2016) provide experimental evidence for finite-horizon planning by human subjects in a multi-stage decision problem.²⁶ In this experiment, a horizon $h = 2$ corresponds to full backward-induction solution of the problem; the authors find evidence of use of the decision strategy corresponding to $h = 1$ (some forward planning, but also not full backward induction), to an extent that varies with time pressure, among other factors. Johnson *et al.* (2002) study behavior in a three-round bargaining game, monitoring the information that subjects collect before making their first offer. They find that all subjects look at first-stage information; yet on 19% of trials, they fail to even look at information about the second-stage situation (evidence of a planning horizon $h = 0$), and on another 10% of all trials, they look at stage-two information but fail to look at information about the third-stage situation (consistent with a horizon $h = 1$). These experiments involve only relatively shallow decision trees, but nonetheless provide fairly clear evidence of finite-horizon planning.

There is also considerable evidence that laboratory subjects playing repeated games or multi-stage games with many stages often fail to “solve” these games by

²⁴If the lower bound is exactly zero, then we will have $\hat{\imath} = -(r^* + \pi^*) < 0$, where $r^* \equiv \beta^{-1} - 1$ is the stationary equilibrium real rate of return inclusive of the safety premium. This is assumed in our numerical calibration, but our qualitative results depend only on our assumption that $\hat{\imath} < 0$.

²⁵When possible, in our numerical results we illustrate how results differ under different assumptions about the planning horizon. Additional numerical results under alternative assumptions about the planning horizon are shown in the online appendix.

²⁶This study is of particular interest because it examines the predictions of a model much like ours that incorporates both finite-horizon forward planning and a value function learned from prior experience.

backward induction when there are more than a few stages. For example, in the case of a finitely repeated prisoner’s dilemma, with the number of repetitions known from the start, backward induction would require players to never cooperate at any stage; instead, players often cooperate initially²⁷ but then defect systematically for the last few rounds (e.g., Selten and Stoecker, 1986). Similarly, in the “centipede game,” where players whose moves alternate can continue the game or immediately terminate it at each of a sequence of stages, backward induction would dictate that a player should immediately terminate whenever given the opportunity; yet in laboratory experiments, players typically continue the game until three or four stages before the (known) final stage, but terminate at that point (McKelvey and Palfrey, 1992). Both of these patterns of behavior are consistent with a model in which people can plan three or four stages ahead, but not farther.

Coibion *et al.* (2020) conduct an experiment that is arguably more relevant to the decision situations considered in this paper by measuring the effect on household expectations of provision of information about the Fed’s projections for the future path of nominal interest rates. They find that this information affects households’ expectations (as reported in a survey) of variables such as US inflation most notably when it is information about the path of interest rates over the coming year; that there is some additional effect of information about the path over the year after that; but that there is little measurable effect on expectations of information provided about the path more than two years in the future. These findings are consistent with a model of the kind that we propose, on the assumption that most households’ planning horizons extend no more than two years into the future. At the same time, they suggest that forward planning with h equal to 3 or 4 quarters may well be within the capabilities of many households.

Perhaps most obviously relevant to the calibration of our model are the empirical

²⁷While this is not rational if the value function used to evaluate terminal positions is the one implied by backward induction — which would reflect an expectation that in subsequent play, one’s opponent will defect regardless of what has happened thus far — it can be rational if the value function used reflects an expectation that having cooperated to that point makes it more likely that the opponent will cooperate. The latter sort of value function could well be learned, given how subjects actually play, if subjects do not learn a value function that conditions on the number of remaining rounds. We leave further analysis of the application of our approach to such settings for future work. Other learning models are proposed by Selten and Stoecker (1986), Nagel and Tang (1998), and Ho and Su (2013).

estimates of Gust *et al.* (2019), who estimate a complete New Keynesian DSGE model with finite-horizon forward planning, assuming an exponential distribution of planning horizons (2.11). They find that quarterly US aggregate data are best fit by a parameterization in which ρ is approximately 0.5, meaning that 50 percent of decision makers only calculate what should happen in the coming quarter when deciding what to do ($h = 0$), another 25 percent calculate only what should happen in the coming quarter and the next one ($h = 1$), and only 6 percent think beyond the coming year ($h \geq 4$).

It is possible that people typically do not need to engage in much forward planning, because they expect the value function that they have learned to be fairly reliable under ordinary circumstances, but that forward planning extends farther when there is a reason to believe that it is worth the additional effort.²⁸ If so, the kind of crisis situations treated below are precisely the kind of situations in which it would be plausible to expect more forward planning than usual. On the one hand, people should realize that an unusual situation has arisen, so that the value function that has been correct on average in the past may be a less accurate guide than usual; and in addition, if the monetary and fiscal authorities announce novel policies, this is new information with consequences for people's decisions that can be determined only through forward planning. Thus we might well expect somewhat longer planning horizons in the situation of interest in this paper than are indicated by estimates like those of Gust *et al.* (2019). Nonetheless, in what follows, we primarily consider the consequences of planning horizons on the order of two years or less.

2.2 A Crisis Scenario

We consider the effects of alternative monetary and fiscal policies under the following scenario: prior to date $t = 0$, we suppose that the economy has for a long time been in the stationary equilibrium discussed above, in which the financial wedge has always been small, the government's budget has been balanced each period (so that government debt has remained equal to zero), and the inflation target π^* has been consistently achieved. As a result, households and firms have learned the value functions that are appropriate to a stationary environment of that kind. At time

²⁸Generalizing our model framework to endogenize the planning horizon in this way would be a valuable extension, but is beyond the scope of the current paper.

Table 1: Calibrated Parameter Values

	Value
Discount factor	$\beta = 0.997$
Price inertia (Calvo parameter)	$\alpha = 0.7747$
Elasticity of substitution between goods	$\theta = 12.7721$
Inverse of Frisch elasticity of labor supply	$\omega = 1.5692$
Phillips curve slope coefficient	$\kappa = 0.00859$
Intertemporal elasticity of substitution	$\sigma = 0.862$
Financial wedge in “crisis” state	$\hat{\Delta} = 0.013$
Persistence of “crisis” state	$\mu = 0.903$
Inflation target	$\pi^* = 0.005$

$t = 0$, however, an unexpected financial disturbance occurs, and the economy enters a “crisis” state, in which there is a substantial financial wedge $\hat{\Delta}_t > 0$ between the return on safe assets (balances held at the central bank) and other assets.

We assume that this crisis state persists, with the size of the financial wedge unchanged, until some date T at which the economy reverts back to its “normal” state, in which we suppose that the financial wedge $\hat{\Delta}_t$ will subsequently equal zero forever after. In some of the exercises below, we assume that the duration T of the crisis is known at the time that the shock occurs. In others, we assume that it is stochastic, but for simplicity we assume that there is a fixed probability $1 - \mu$ of reversion to the “normal” state each period, so that the exogenous fundamental $\{\Delta_t\}$ evolves according to a two-state Markov chain, as in Eggertsson and Woodford (2003). We write the constant financial wedge in the crisis state as $\hat{\Delta}_t = -\hat{r} + \underline{\Delta}$, where $\underline{\Delta} > 0$; the latter quantity measures the degree to which the financial wedge is too large to be offset through a contemporaneous interest-rate reduction.²⁹ The fact that $\underline{\Delta} > 0$ means that the inflation target can no longer be maintained at all times, using only conventional interest-rate policy and with a balanced government budget.

²⁹In the notation of Eggertsson (2011), this quantity corresponds to $\underline{\Delta} = -\underline{r} - \pi^*$, where $\underline{r} < 0$ is the natural rate of interest in the crisis state.

2.2.1 Numerical calibration

We illustrate a number of our conclusions about the effects of alternative policies under such a scenario for economic fundamentals using numerical computations. In these calculations, we calibrate the model — including our assumption about the size and persistence of the disturbance to fundamentals — largely in accordance with the parameter values proposed by Eggertsson (2011), who shows that under the assumption of rational expectations and a zero inflation target, these parameter values would imply a contraction of the size experienced by the US economy during the Great Depression, as shown by Eggertsson (2011). However, in this paper, we specify “normal” monetary policy as involving an inflation target π^* of two percent per year, rather than a target of zero inflation, as in Eggertsson’s model of the Great Depression. This makes the zero lower bound a less severe constraint in our scenario than in the one considered by Eggertsson, since we continue to assume the same size of increase in the financial wedge as in his Depression scenario.

In our numerical calculations, the periods of our discrete-time model are identified with quarters. We set the subjective discount factor $\beta = 0.997$, the slope of the Phillips curve $\kappa = 0.00859$, and the elasticity of intertemporal substitution $\sigma = 0.862$.³⁰ The shock required to account for the size of the contraction during the Great Depression is one in which $\hat{\Delta} = 0.013$,³¹ and the probability of remaining in the crisis state is $\mu = 0.903$, so that the expected length of a crisis is about 10 quarters. In addition, we assume a long-run inflation target of 2 percent per year; that is, $\pi^* = 0.005$ in quarterly terms, which implies that the part of financial wedge that cannot be offset by monetary policy owing to the ZLB is $\underline{\Delta} = 0.005$, or 2 percent per year.³² The calibrated parameter values are summarized in Table 1.

³⁰The table also reports values for several additional parameters (α, ω, θ) that matter only in section 4.3, when we consider welfare analysis in the presence of heterogeneous planning horizons. These parameters and their relevance for welfare calculations are discussed further in Appendix D.

³¹This is a quarterly rate; thus the assumed increase in the size of the financial wedge is a bit greater than 5 percent per annum. The natural rate of interest in the normal state is $r^* = \beta^{-1} - 1$, or slightly above 1 percent per annum; thus we assume that in the crisis state, the natural rate of interest falls to -4 percent per annum, as in Eggertsson (2011).

³²Note that this is only half the size of $\underline{\Delta}$ in the crisis state considered by Eggertsson (2011).

2.2.2 Contraction in the absence of a policy response

We first consider the consequences of a temporary large increase in the size of the financial wedge (the “crisis scenario” explained above), in the case of a two-state Markov disturbance, under an assumption that monetary and fiscal policy continue to be conducted as under normal conditions, which is to say as assumed above in our discussion of the stationary equilibrium prior to the occurrence of the shock. We assume that the government budget continues to be balanced each period, so that $B_{t+1} = 0$ at all times, and that the central bank continues to conduct monetary policy in accordance with a strict inflation target. The latter stipulation implies that in each period, \hat{i}_t will be set as necessary to ensure that inflation is equal to the target rate ($\pi_t = 0$, in our deviations notation), if this is consistent with the ZLB; if inflation undershoots the target in any period t even when the interest rate is at its lower bound, then \hat{i}_t will equal \hat{i} in that period (the policy as close as possible to achieving the inflation target in that period, taking as given the expected conduct of monetary policy in all future periods).

Let us first recall the analysis of such a situation under the assumption of rational expectations (RE) by Eggertsson and Woodford (2003) and Eggertsson (2011). The linearized equations of the RE model can be written in vector form as

$$z_t = A E_t z_{t+1} - \sigma a (\hat{i}_t + \hat{\Delta}_t), \quad (2.13)$$

where we define

$$z_t \equiv \begin{bmatrix} y_t \\ \pi_t \end{bmatrix}, \quad A \equiv \begin{bmatrix} 1 & \sigma \\ \kappa & \beta + \kappa\sigma \end{bmatrix}, \quad a \equiv \begin{bmatrix} 1 \\ \kappa \end{bmatrix}.$$

(Note that the path of public debt is irrelevant, owing to Ricardian Equivalence.) Under the assumption that $\hat{\Delta}_t$ evolves according to the two-state Markov process and that \hat{i}_t is chosen according to the inflation targeting policy, there exists a rational-expectations solution that is also Markovian, in the sense that the vector z_t takes only two possible values: a vector \underline{z} in any period t in which the crisis state persists, and the zero vector in each period after the reversion to the normal state (in which case the inflation target is achievable each period from then on).³³

³³Mertens and Williams (2018) call this the “target equilibrium”; it is not the only possible RE solution, even if one restricts attention to Markovian solutions. It is, however, the solution

In the case that

$$\kappa\sigma\mu < (1 - \mu)(1 - \beta\mu),$$

the matrix A has two positive real eigenvalues, both less than μ^{-1} , and the Markovian solution is also the unique bounded solution to the linear system (2.13). This condition holds if and only if

$$\mu < \bar{\mu}, \tag{2.14}$$

where $\bar{\mu}$ is a bound between zero and 1 that depends on the values of $\kappa\sigma$ and of β ; this is the case considered by Eggertsson and Woodford (2003).³⁴ In this case, the Markovian RE solution under which the crisis state persists is given by

$$z_t = \underline{z}^{RE} \equiv -\sigma(I - \mu A)^{-1} a \underline{\Delta} \lll 0. \tag{2.15}$$

In this equilibrium, both output and inflation remain persistently below their target values as long as the crisis state continues, but return immediately to their target values as soon as the financial wedge returns to its normal (negligible) value. As Eggertsson and Woodford (2003) show in a calibrated example, this equilibrium can involve quite a severe contraction as well as substantial deflation, in response to even a few percentage points' increase in the financial wedge. We now examine the robustness of these conclusions to allowing for finite planning horizons.

Assume again that the central bank adheres to a strict inflation targeting policy, and suppose also that there is no government debt (so that the fiscal authority maintains a balanced budget).³⁵ Equations (2.7) and (2.9) can then be written in vector form as

$$z_t^j = A E_t z_{t+1}^{j-1} - \sigma a (\hat{v}_t^j + \hat{\Delta}_t) \tag{2.16}$$

emphasized in the RE literature, following Eggertsson and Woodford (2003); we show below that restriction of attention to this RE solution can be justified as the limit of the unique solution associated with a model with finite planning horizons, when the length of the planning horizons is made arbitrarily long.

³⁴A Markovian rational-expectations solution can also be defined when μ exceeds the bound (2.14), but in this case it does not correspond to the limit of an equilibrium with finite-horizon planning, as planning horizons are made arbitrarily long.

³⁵This is a common assumption in New Keynesian models used for monetary policy analysis, though in models where Ricardian Equivalence would hold, it is without loss of generality. With finite planning horizons, the assumption is not innocuous, as we show in section 3.

for each $j \geq 1$, using the notation z_t^j for the vector $[y_t^j \ \pi_t^j]'$, while (2.8) and (2.10) can be written as

$$z_t^0 = -\sigma a (\hat{i}_t^0 + \hat{\Delta}_t) + (1 - \beta) a b_{t+1}. \quad (2.17)$$

Under the assumption of zero public debt, equation (2.17) implies that an expectation of strict inflation targeting requires that horizon-zero agents expect an interest rate

$$\hat{i}_t^0 = \max\{-\hat{\Delta}_t, \hat{\imath}\}.$$

Under the assumption that the financial wedge evolves as a two-state Markov chain, this implies that $z_t^0 = 0$ if t is any date after the reversion to the normal state, while

$$z_t^0 = \underline{z}^0 \equiv -\sigma a \underline{\Delta} \lll 0 \quad (2.18)$$

if t is any date at which the crisis state continues.

We can then use this result to solve recursively for the behavior of households and firms with progressively longer planning horizons. We find that z_t^h has a common value \underline{z}^h in each period t in which the crisis state continues, given by³⁶

$$\underline{z}^h = -\sigma \sum_{j=0}^h (\mu A)^j a \underline{\Delta} \lll 0 \quad (2.19)$$

for any planning horizon $h \geq 0$. Note that the solution is well defined for any finite h ; if in addition to our more general assumptions, μ satisfies the bound (2.14), the solution has a well-defined limit as h is made unboundedly large. In this latter case, we find that as $h \rightarrow \infty$, $\underline{z}^h \rightarrow \underline{z}^{RE}$, so that the unique equilibrium with finite-horizon planning approaches the Markovian rational-expectations equilibrium discussed above. It follows that any long enough finite planning horizon will lead to outcomes similar to those in the RE analysis.

If planning horizons are only of modest length, however, the quantitative predictions of the model with finite-horizon planning are different from those of the RE analysis. Since each of the terms in the sum (2.19) is a vector with both elements negative, it is evident that both \underline{y}^h and $\underline{\pi}^h$ are more negative the longer the planning horizon. This is illustrated in Figure 1, for the numerical parameter values listed in Table 1.

³⁶See Appendix A for details of the derivation.




Figure 1: Expenditure and rates of price increase during the crisis period, under different assumptions about the planning horizon h (in quarters) of households and firms, when the central bank follows a strict inflation targeting policy and there is no response of fiscal policy.

This solution tells us the value of \underline{y}^h and $\underline{\pi}^h$ for each possible planning horizon h . These calculations are the same regardless of the distribution of planning horizons in the economy. For a given distribution of planning horizons $\{\omega_h\}$, we can then compute the predicted state-contingent evolution of aggregate output and inflation by aggregating the individual decisions of the agents with different horizons. In the case of an exponential distribution of planning horizons (2.11), the condition required for the infinite sum $\sum_{h=0}^{\infty} \omega_h z^h$ to converge — and hence for there to be a well-defined equilibrium under the assumed policies — is

$$\rho\mu < \bar{\mu}, \quad (2.20)$$

where $\bar{\mu}$ is defined as in (2.14).

This is a weaker condition than (2.14), that requires only that the product $\rho\mu$ not be too large; it will be satisfied if either most planning horizons are not too long (ρ is well below 1) or the financial disturbance is not expected to last too long (μ is well

below 1), or both. In the case that it is satisfied, aggregate outcomes in the crisis state will be given by

$$\underline{z} = -\sigma [I - \rho\mu A]^{-1} a \underline{\Delta} \ll 0. \quad (2.21)$$

Note that if (2.20) is satisfied, (2.21) is the unique solution to our model, not simply one among multiple possible solutions, as in the rational-expectations analysis. In the case that (2.14) is satisfied, the solution (2.21) approaches the RE solution specified in (2.15) as ρ approaches 1; this provides a possible justification for selecting that solution in a rational-expectations analysis.

We see from Figure 1 that when households and firms have finite planning horizons, the contractionary and disinflationary effects of an increase in the financial wedge are less severe than in a rational-expectations analysis; the more short-sighted people are assumed to be, the milder the effects. Nevertheless, assuming some degree of foresight, the ZLB can pose a serious problem under these assumptions about policy. (A larger increase in the financial wedge would produce a correspondingly larger contraction than those shown in the figure.) Thus it is still desirable to explore whether alternative policies can mitigate this problem.

One possibility would be to consider what can be achieved by committing to a more expansionary monetary policy following the return of the financial wedge to its normal level, as proposed by Eggertsson and Woodford (2003). In their RE analysis, such a policy can greatly improve upon the outcomes associated with a purely forward-looking inflation targeting regime; however, the effects of such “forward guidance” depend entirely upon its being taken into account in the expectations of households and firms during the crisis period, which depends upon planning horizons being sufficiently long.³⁷ An alternative approach is to consider what can be achieved by increasing fiscal transfers in response to the financial disturbance. As we shall see, when planning horizons are finite, the use of this additional tool can achieve greater stabilization than even the best-designed forward guidance policy can on its own. However, the ideal policy response will involve both increased fiscal transfers and forward guidance regarding future interest-rate policy.

³⁷See Woodford and Xie (2019) for quantitative analysis of the degree to which shortening the assumed length of planning horizons reduces the predicted effects of such policies, even when clearly explained and fully credible.

3 Fiscal Transfers and Aggregate Demand

As explained in section 2, in this paper we consider only fiscal policies in which the real public debt B_{t+1} is a function of the exogenous state in period t (including the history of exogenous evolution of the financial wedge, through period t , and any information available at time t about future financial wedges), but in this section we no longer require that $B_{t+1} = 0$ at all times. The implied state-contingent level of net lump-sum transfers T_t is then given by equation (2.2). While we now allow the path of the debt to respond to shocks, we consider only policies under which the process $\{B_{t+1}\}$ remains within finite bounds with certainty for all time; this means that we consider only policies under which any increase in the public debt is eventually paid off, with certainty.³⁸ Given this — together with the facts that all taxes and transfers are lump sum and distributed equally to all households, and that there are no financial constraints (other than the “financial wedge” that allows riskless claims on the government to trade at a lower equilibrium rate of return than private debt) — our model is one in which Ricardian Equivalence would hold under an assumption of rational expectations.

Instead, if households have finite planning horizons — or even if a sufficient number of them do — a bounded increase in the path of the real public debt (resulting from an initial increase in lump-sum transfers, followed eventually — though possibly much later — by the tax increases required to keep the debt from exploding) will increase aggregate demand. Note that the household FOCs (2.4)–(2.5) imply that real expenditure Y_t^h by households with a planning horizon of h periods must satisfy

$$u'(Y_t^h) = \text{E}_t\left[\prod_{j=1}^h D_{t+j}^{h+1-j} \cdot \tilde{D}_{t+h}^0 v'(B_{t+h+1})\right], \quad (3.1)$$

where the stochastic discount factors are defined by

$$D_{t+1}^j \equiv \beta \frac{1 + i_t^j + \Delta_t}{\Pi_{t+1}^{j-1}} \quad \text{for any } j \geq 1, \quad \tilde{D}_t^0 \equiv \beta(1 + i_t^0 + \Delta_t).$$

Now consider the effect of a fiscal policy change that increases the planned level of B_{t+1} for at least some future dates (in at least some possible states of the world), while

³⁸This is true regardless of how prices, interest rates, and economic activity may evolve; thus we do not consider the effects of “non-Ricardian” fiscal policy rules of the kind discussed, for example, in Woodford (2001).

decreasing it at no dates. If the paths of neither goods prices nor asset prices change (as would be the case under Ricardian Equivalence), then (3.1) implies that Y_t^h must increase in any period t with the property that B_{t+h+1} is increased in at least some states that remain possible, conditional on the state at date t .³⁹ Aggregating across households with different planning horizons, one concludes that aggregate output Y_t must increase, in at least some periods; thus Ricardian Equivalence does not obtain.

The key to this result, of course, is our assumption that announcement of the policy change does not change the value function $v(B)$ used to evaluate terminal states. A household with rational expectations should instead understand that if a policy change results in a higher real public debt B_{t+h+1} , it must imply higher tax obligations in periods subsequent to $t + h$ (that is, beyond the planning horizon); and this should change the level of private wealth B_{t+h+1}^i needed in order to ensure a given level of continuation utility. Thus the correct value function $v(B_{t+h+1}^i)$ would have as another argument the aggregate supply of debt B_{t+h+1} .

Because the value function takes account only of a coarse description of the household's situation — and because the situation that gives rise to an unusually large public debt following a financial crisis may not be similar to situations that the household has frequently encountered in the past — we suppose that households have not already learned how to take this additional state variable into account in the way that they value terminal states.⁴⁰ Neglect of this state variable is what breaks Ricardian Equivalence. The degree to which this is quantitatively important will depend on the extent to which the time that it takes for the real public debt to return to its normal level following a shock exceeds the planning horizons of many households.

The failure of Ricardian Equivalence adds another dimension along which government policy can shift the equilibrium allocation of resources, possibly in ways that can improve stabilization outcomes. This is particularly easy to see in the case of an exponential distribution of planning horizons (2.11), where the parameter $0 < \rho < 1$

³⁹This follows from the fact that both $u'(Y)$ and $v'(B)$ are decreasing functions of their respective arguments.

⁴⁰One might wonder whether after crises have repeatedly occurred, to which fiscal policy always responds in the way discussed here, people should not learn a value function which takes account of the predictable increase in tax obligations following the crisis, undercutting the effects of fiscal transfers during the crisis. While this logic is consistent with our account of how the value function is learned, such learning is likely to be slow, in the absence of commitment by the fiscal authorities to a strict rule for the way in which taxes adjust to any increase in the public debt.

determines the mean planning horizon $\bar{h} \equiv \rho/(1 - \rho)$. In this case, the log-linearized aggregate demand relations (2.7)–(2.8) can be aggregated to yield

$$y_t = -\sigma(\hat{i}_t^e + \hat{\Delta}_t - \rho E_t \pi_{t+1}) + \rho E_t y_{t+1} + (1 - \rho)(1 - \beta)b_{t+1}, \quad (3.2)$$

where

$$\hat{i}_t^e \equiv (1 - \rho) \sum_{j=0}^{\infty} \rho^j \hat{i}_t^j$$

is an average of the interest rates expected by households with different planning horizons.⁴¹ The linearized aggregate supply relations (2.9)–(2.10) can similarly be aggregated to yield

$$\pi_t = \kappa y_t + \rho \beta E_t \pi_{t+1}. \quad (3.3)$$

Note that equations (3.2) and (3.3) relating the evolution of aggregate output and inflation reduce to the structural equations of the standard New Keynesian model under rational expectations in the limit as $\rho \rightarrow 1$.

Equation (3.2) shows that variation in the level of real public debt b_{t+1} (the debt issued in period t) shifts the aggregate-demand relation in exactly the same way as does variation in \hat{i}_t , the central bank's interest-rate target. It follows that, if one is concerned solely with stabilization of the aggregate variables y_t and π_t , there is no need to consider varying the path of the real public debt, as long as it is possible for the central bank to vary \hat{i}_t to the desired degree instead. However, when the zero lower bound is a binding constraint on interest-rate policy, the fact that the public debt can still be increased through transfer policy can effectively relax this constraint.

This allows stabilization of the aggregate economy in cases where this would not be possible under a policy that maintained $b_{t+1} = 0$ at all times. Note that the paths in which $y_t = \pi_t = 0$ at all times are consistent with both equations (3.2) and (3.3) holding at all times, if and only if

$$-\sigma(\hat{i}_t^e + \hat{\Delta}_t) + (1 - \rho)(1 - \beta)b_{t+1} = 0 \quad (3.4)$$

⁴¹In Woodford (2019), this equation involves \hat{i}_t , the actual interest-rate target of the central bank, rather than the variable \hat{i}_t^e defined here. The form (3.2) is more generally valid. In the earlier paper, monetary policy is assumed to be characterized by a linear relationship among \hat{i}_t and other aggregate variables, such as a Taylor rule $\hat{i}_t = \phi(\pi_t, y_t; s_t)$, where s_t is an exogenous state and ϕ is linear in the first two arguments. In such a case, the fact that the policy rule is understood by all households implies that $\hat{i}_t^j = \phi(\pi_t^j, y_t^j; s_t)$ for each horizon j ; aggregating over the different horizons then implies that $\hat{i}_t^e = \phi(\pi_t, y_t; s_t) = \hat{i}_t$, owing to the linearity of ϕ . When the zero lower bound sometimes constrains policy, as in the cases considered in this policy, \hat{i}_t^e will in general no longer equal \hat{i}_t .

at all times. Since everyone is assumed to understand that the central bank's policy must conform to the lower bound $\hat{i}_t \geq \underline{\hat{i}}$, the interest rates expected by households must satisfy $\hat{i}_t^e \geq \underline{\hat{i}}$ at all times. Hence if $\hat{\Delta}_t > -\underline{\hat{i}}$ at some time, it will not be possible to satisfy (3.4) with $b_{t+1} = 0$.

Instead, the condition can always be satisfied if we allow fiscal transfers. Let us suppose that the central bank's interest-rate target tracks variations in the financial wedge to the extent that this is consistent with the ZLB, i.e., that monetary policy ensures that

$$\hat{i}_t = \max\{-\hat{\Delta}_t, \underline{\hat{i}}\} \quad (3.5)$$

each period.⁴² Then (since the interest rate is specified as a function of the exogenous state) \hat{i}_t^e will equal \hat{i}_t , and condition (3.4) will be satisfied if and only if fiscal policy is given by

$$b_{t+1} = \frac{\sigma}{(1-\rho)(1-\beta)} \tilde{\Delta}_t, \quad (3.6)$$

where

$$\tilde{\Delta}_t \equiv \max\{\hat{\Delta}_t + \underline{\hat{i}}, 0\} \quad (3.7)$$

measures the part of the financial wedge that is not offset by interest-rate policy. If monetary policy is given by (3.5) and fiscal policy by (3.6), equilibrium will involve $y_t = \pi_t = 0$ at all times, regardless of the path of the financial wedge.⁴³

⁴²More precisely, we assume that this policy is followed during a relatively brief period in which there is a non-trivial financial wedge, but that after that period the central bank reverts to a policy rule that ensures achievement of its inflation target. The latter stipulation is required in order to ensure that there should not be any long-run drift in the value functions of households and firms, allowing us to abstract from modeling the endogenous adjustment of value functions, as discussed in section 2. If the rule (3.5) were followed forever, then the learning process for the value functions specified in Woodford (2019) would lead to unstable dynamics, as shown in that paper for the case of a permanent zero financial wedge.

⁴³It is immediately obvious from inspection of equations (3.2) and (3.3) that the asserted solution is consistent with both of these equations at all times. We show in Appendix B that this is indeed the unique equilibrium outcome, assuming a bound on the asymptotic growth rate of the excess financial wedge. The required condition holds, for example, in the case of the two-state Markov process for the financial wedge introduced in section 2.2 as long as (2.20) is satisfied.

4 Coordinated Monetary and Fiscal Stabilization Policy

The striking result of the previous section might make it seem that there is no need for a central bank to depart from its commitment to a strict inflation-targeting policy, given that fiscal transfers can be varied to offset any effects on aggregate demand of variations in the financial wedge. Can one simplify the tasks of both policy authorities and communication with the public by stating that the sole concern of the central bank should be to ensure that inflation remains equal to the target rate, while it is the responsibility of the fiscal authority to offset any excessive financial wedge (any positive value of $\tilde{\Delta}_t$) with fiscal transfers, so as to maintain a zero output gap?

We shall argue that this would be a mistake. Successful use of fiscal policy as a tool of stabilization policy requires that it be supported by an appropriate monetary policy; moreover, the ideal joint policy will involve a commitment that monetary policy will continue to depart from the central bank's usual inflation targeting policy, even after the financial wedge has returned to its normal size.

4.1 The Dependence of Fiscal Stimulus on Monetary Accommodation

It might seem from the analysis above that the central bank can commit itself to the inflation targeting policy considered in section 2.2.2 (setting \hat{i}_t as needed to achieve the inflation target, or as low as possible if the target cannot be achieved), and that as long as fiscal policy is given by (3.6), the outcome will be complete stabilization of both inflation and the output gap. This would, however, be incorrect. It is true that the equilibrium described at the end of the previous subsection is one in which the paths of \hat{i}_t and π_t conform to the proposed monetary policy rule, but it is *not* true that that equilibrium is consistent with everyone expecting that monetary policy will be conducted in accordance with that rule. In our model, because of people's finite planning horizons, it matters not only what happens in equilibrium, but what the central bank would be expected to do out of equilibrium; and the complete stabilization of macroeconomic aggregates actually depends on people's understanding that the central bank is *not* determined to prevent overshooting of the long-run inflation target under any circumstances.

In order to see this, we need to consider the forward plans of agents with differing planning horizons in the equilibrium in which $y_t = \pi_t = 0$ at all times. Substituting the monetary policy rule (3.5) into (2.16)–(2.17) yields

$$z_t^j = A E_t z_{t+1}^{j-1} - \sigma a \tilde{\Delta}_t \quad (4.1)$$

for each $j \geq 1$, and

$$z_t^0 = -\sigma a \tilde{\Delta}_t + (1 - \beta)a b_{t+1}.$$

In the case of an arbitrary process for the financial wedge and an arbitrary fiscal policy, this system of equations can be solved recursively to yield

$$z_t^h = -\sigma \cdot \sum_{j=0}^h [A^j a] E_t \tilde{\Delta}_{t+j} + (1 - \beta)[A^h a] E_t b_{t+h+1} \quad (4.2)$$

for any planning horizon $h \geq 0$. The implied solutions for the aggregates y_t and π_t are then obtained by averaging over the various planning horizons h . If fiscal policy is given by (3.6), these equations imply $y_t = \pi_t = 0$; however, they do not generally imply $y_t^h = \pi_t^h = 0$ for each individual planning horizon.

Consider, for example, the case in which the financial wedge evolves according to a two-state Markov chain of the kind proposed in section 2.2. In this case, the right-hand side of (4.2) depends only on whether the economy is still in the crisis state at date t or has already returned to normal. In any period t such that the economy remains in the crisis state, the solution is given by

$$z_t^h = \underline{z}^h \equiv \sigma \cdot \left\{ \frac{\mu^h}{1 - \rho} [A^h a] - \sum_{j=0}^h \mu^j [A^j a] \right\} \cdot \underline{\Delta},$$

where $\underline{\Delta} > 0$ is the excess financial wedge in this state. Instead, in any period after the return to the normal state, $z_t^h = 0$. Note that the solution for \underline{z}^h is well defined for any finite horizon h , regardless of parameter values.

One observes that in the crisis state, the elements of \underline{z}^h are different for different horizons h . For example, when $h = 0$,

$$\begin{bmatrix} \underline{y}^0 \\ \underline{\pi}^0 \end{bmatrix} = \sigma \frac{\rho}{1 - \rho} \begin{bmatrix} 1 \\ \kappa \end{bmatrix} \underline{\Delta} \gg 0.$$

Moreover, one can show that the largest of the two positive real eigenvalues of A is equal to $\bar{\mu}^{-1} > 1$, where $\bar{\mu}$ is the quantity introduced in (2.14). Then if $\mu < \bar{\mu}$, one

finds that

$$\underline{z}^h \rightarrow \underline{z}^{RE} \ll 0$$

as $h \rightarrow \infty$, where \underline{z}^{RE} is the Markovian rational-expectations solution defined in (2.15). Thus both \underline{y}^h and $\underline{\pi}^h$ are positive in the case of short enough planning horizons, while both are negative in the case of long enough horizons.⁴⁴

The situation would be quite different if, instead, the central bank were understood to be committed to setting the interest rate required to achieve its inflation target, unless constrained by the ZLB. In that case, there would be a maximum degree of aggregate demand stimulus that could be achieved through fiscal transfers, no matter how large the transfers might be. Under strict inflation targeting, (2.8) and (2.10) imply

$$\begin{aligned} \pi_t^0 &= -\kappa\sigma(\hat{i}_t^0 + \hat{\Delta}_t) + \kappa(1 - \beta)b_{t+1} \\ &= \min\{-\kappa\sigma(\hat{i}_t + \hat{\Delta}_t) + \kappa(1 - \beta)b_{t+1}, 0\}. \end{aligned}$$

In the case of the assumed two-state Markov chain for the financial wedge, this implies that as long as the crisis state persists, one will have

$$\pi_t^0 = \underline{\pi}^0 = \kappa \min\{(1 - \beta)b_{t+1} - \sigma\underline{\Delta}, 0\}.$$

(The corresponding value of \underline{y}^0 is simply this quantity without the prefactor κ .) Thus increases in the public debt are stimulative only up to the level

$$b^{max} \equiv \frac{\sigma\underline{\Delta}}{1 - \beta}.$$

(In our numerical calibration, this amounts to 0.36 of annual GDP.⁴⁵) For any level $b_{t+1} \geq b^{max}$, the model predicts that $\underline{y}^0 = \underline{\pi}^0 = 0$.

For any longer horizon h , we similarly will have $z_t^h = \underline{z}^h$ as long as the crisis state persists, where crisis values $\{\underline{z}^h\}$ can be computed recursively as follows. For any $j \geq 1$, (4.1) implies that in any crisis period,

$$\pi_t^j = [\kappa \beta + \kappa\sigma] \mu \underline{z}^{j-1} - \kappa\sigma\underline{\Delta}$$

⁴⁴See Figure 4 below for illustration of a similar result in the case that the duration of the crisis period is known with certainty from the beginning.

⁴⁵Note that this does not mean that there would be no effect of increasing the public debt beyond 36 percent of GDP — a level that the US is already well past. It means that, under the assumptions of our calibration, there would be no effect of an increase by more than 36 percent of GDP relative to the normal steady-state level of public debt.




Figure 2: Expenditure and rates of price increase during the crisis period, for households and firms with different planning horizons h (in quarters) when the central bank follows a strict inflation targeting policy. The two lines correspond to the minimal and maximal sizes of fiscal stimulus.

if \hat{v}_t^j is expected to be at the lower bound. If both elements of \underline{z}^{j-1} are non-positive, this implies inflation below target, even with the interest rate at the lower bound. Hence the ZLB will bind, and we must have

$$\underline{z}^j = \mu A \underline{z}^{j-1} - \sigma a \underline{\Delta} \ll 0. \quad (4.3)$$

Under the assumption that $b_{t+1} \geq b^{max}$ for as long as the crisis state persists (the most favorable assumption for a stimulative effect of fiscal policy), we have shown in the previous paragraph that $\underline{z}^0 = 0$; we can then show recursively using (4.3) that both elements of \underline{z}^j are non-positive for all $j \geq 0$. It follows that the assumption used to derive (4.3) is valid for all $j \geq 1$.

Thus under the most expansive possible fiscal policy, we will have $z_t^h = \underline{z}^h$ as long as the crisis state persists, where the sequence $\{\underline{z}^h\}$ can be computed recursively

using (4.3), starting from the initial condition $\underline{z}^0 = 0$. This yields the solution

$$\underline{z}^h = -\sigma \sum_{j=1}^h (\mu A)^{j-1} a \underline{\Delta} \lll 0$$

for each $h \geq 1$. Both y_t^h and π_t^h remain below their target values for all horizons $h > 0$, and more so the longer the horizon. (This is illustrated for our numerical example in Figure 2, which shows the values of \underline{z}^h both in the case of zero fiscal stimulus⁴⁶ and in the case of the maximum fiscal stimulus.) We see that fiscal stimulus can mitigate the contractionary and disinflationary effects of the financial disturbance, but both spending and the rate of price increase continue to fall, even with the maximum fiscal stimulus, for all horizons $h > 0$; if planning horizons extend years into the future, the fraction of the contractionary effect that can be offset using fiscal policy alone is quite modest.

Summing over the different planning horizons (again assuming an exponential distribution of horizons), the net effect on both aggregate output and aggregate inflation is necessarily contractionary. As long as (2.20) is satisfied,⁴⁷ the weighted average of the $\{\underline{z}^h\}$ is a convergent sum, and equal to

$$\underline{z} = (1 - \rho) \sum_{h=0}^{\infty} \rho^h \underline{z}^h = -\rho\sigma [I - \rho\mu A]^{-1} a \underline{\Delta} \lll 0.$$

It is not possible to fully stabilize either aggregate output or inflation; both necessarily fall in the crisis state. Indeed, the effects on aggregate output and inflation are similar to those obtained in the case of no fiscal response (see equation (2.21) above): they are simply both reduced by a factor of ρ . This means that the contractionary effects are reduced by less than half, in the case of any mean planning horizon \bar{h} greater than one quarter.

Instead, it is possible to completely eliminate the contractionary effects of the increased financial wedge on both output and inflation, if an expansionary fiscal policy (an increase in the real public debt through lump-sum transfers, in an amount proportional to the excess financial wedge $\underline{\Delta}$) is combined with monetary accommodation — a commitment to keep the nominal interest rate at its lower bound during

⁴⁶Note that the results for $b = 0$ repeat those shown in Figure 1 above.

⁴⁷This is the condition required for both eigenvalues of $\rho\mu A$ to be less than 1, the same condition required for convergence as in the case of zero fiscal stimulus considered in section 2.2.2.

the period in which the financial wedge is large, even if this causes inflation to overshoot its long-run target. A coordinated change in both monetary and fiscal policy in response to the financial disturbance can achieve more than either policy can on its own.

Thus while fiscal transfers have an important contribution to make, in the case that planning horizons are finite, the availability of this additional instrument does not make monetary stabilization policy irrelevant. Moreover, the important aspect of monetary policy is not what the central bank actually *does* during the period when the financial wedge is large (since the ZLB binds during this period); rather, it is what it leads people to believe that it *would* do, in the event that the ZLB were to cease to bind. In this sense, commitments about the determinants of future interest-rate policy remain a crucial dimension of policy, even when aggressive use of government transfers is possible.

4.2 The Continuing Relevance of Forward Guidance

The example considered above not only shows that the use of fiscal transfers can improve stabilization outcomes, relative to what monetary policy alone can accomplish; the results obtained might seem to make the details of monetary policy unimportant, given sufficient latitude in the way that fiscal policy can be used. If we assume a conventional objective for stabilization policy, in which the aim is to minimize the expected value of a discounted sum of squared target misses

$$E_0 \sum_{t=0}^{\infty} \beta^t [\pi_t^2 + \lambda y_t^2], \quad (4.4)$$

then it is easy to characterize an optimal joint monetary-fiscal policy in the case of an exponential distribution of planning horizons.

We have shown in this case that if (3.4) is satisfied at all times, we will have $y_t = \pi_t = 0$ at all times, which obviously achieves the minimum possible value of criterion (4.4). Moreover, it is possible to choose a state-contingent evolution $\{b_{t+1}\}$ that satisfies (3.4) at all times, in the case of *any* assumed state-contingent evolution for $\{\hat{i}_t\}$, as long as \hat{i}_t is a function only of the exogenous state, so that $\hat{i}_t^e = \hat{i}_t$.⁴⁸ For example, it is not necessary for interest-rate policy to respond at all to increases in

⁴⁸It is necessary, however, that we assume that interest rates do not adjust endogenously in response to changes in b_{t+1} in such a way as to keep $(1 - \rho)(1 - \beta)b_{t+1} - \sigma \hat{i}_t^e$ constant; this is

the financial wedge in order for complete aggregate stabilization to be possible; we could assume that $\hat{i}_t = 0$ at all times and make the fiscal authority solely responsible for responding to variations in financial conditions.

The example suggests another strong conclusion as well: it would seem that there is no need to contemplate any deviation from our baseline policy regime (strict inflation targeting and zero public debt) in periods when financial wedges are small (small enough so that $\tilde{\Delta}_t = 0$), simply because there are transitory periods in which the wedges are large. Thus there might seem to be no need for forward guidance, in the sense of a commitment to more stimulative than ordinary policy for a time even after the financial wedge is again small, for the sake of improved stabilization during the period when the wedge is large.

But even if we assume that the expected value of (4.4) is the sole objective of policy (an assumption that we challenge below), this conclusion is not generally true outside the special case of the exponential distribution of planning horizons. As a simple (but instructive) alternative case, suppose instead that all households and firms have a common planning horizon $h > 0$, and consider the general class of policies specified by exogenous state-contingent paths for both $\{\hat{i}_t, b_{t+1}\}$. Within this family of policies, our goal is to choose state-contingent paths $\{\hat{i}_t, b_{t+1}\}$ so as to minimize the expected value of (4.4), where $y_t = y_t^h, \pi_t = \pi_t^h$ at all times (because of the common planning horizon).

Let us begin by considering the optimal evolution of $\{b_{t+1}\}$, taking as given the state-contingent path of $\{\hat{i}_t\}$. In the case of an arbitrary interest-rate policy, we can use the same methods as above to show that z_t^h will be given by (4.2), except that in the more general case the variable $\tilde{\Delta}_t$ defined in (3.7) must be replaced by $\check{\Delta}_t \equiv \hat{\Delta}_t + \hat{i}_t$.

Since only the evolution of the variables $\{z_t^h\}$ for horizon h matters for the stabilization objective, it follows from this solution that the choice of b_{t+h+1} for any exogenous state s_{t+h} in period $t+h$ affects no variables relevant to the stabilization objective other than z_t^h in the state at period t in which it is possible to reach the particular state s_{t+h} at date $t+h$. We can then reduce the problem of choosing an optimal state-contingent evolution for $\{b_{t+1}\}$ to a sequence of independent static problems: for any state s_t in period t , choose the level of b_{t+h+1} in the states at date $t+h$ that are

the problem with an expectation that the central bank is committed to whatever interest-rate adjustments are needed to achieve a fixed inflation target.

possible conditional on being in state s_t so as to minimize $L(z_t^h) \equiv (\pi_t^h)^2 + \lambda(y_t^h)^2$, where z_t^h is given by the generalized version of (4.2).

This is a convex minimization problem, with a unique interior solution characterized by a first-order condition. If we introduce the notation

$$A^j a = \begin{bmatrix} \alpha_j \\ \gamma_j \end{bmatrix}$$

for each $j \geq 0$, then the first-order condition is given by

$$\gamma_h \pi_t^h + \lambda \alpha_h y_t^h = 0.$$

Substituting the generalization of (4.2) into this yields a linear equation (with a unique solution) for the expected public debt at the end of the planning horizon:

$$E_t b_{t+h+1} = \frac{\sigma}{1-\beta} \sum_{j=0}^h \frac{\lambda \alpha_h \alpha_j + \gamma_h \gamma_j}{\lambda \alpha_h^2 + \gamma_h^2} E_t \check{\Delta}_{t+j}. \quad (4.5)$$

Substituting this solution into the generalization of (4.2) then yields an equation for z_t^h in the case of an optimal transfer policy, but arbitrary interest-rate policy,

$$z_t^h = \theta_t \begin{bmatrix} -\gamma_h \\ \lambda \alpha_h \end{bmatrix}, \quad \text{where } \theta_t \equiv \sigma \sum_{j=0}^{h-1} \frac{\alpha_j \gamma_h - \alpha_h \gamma_j}{\lambda \alpha_h^2 + \gamma_h^2} E_t \check{\Delta}_{t+j}. \quad (4.6)$$

It follows that the minimum achievable value of $L(z_t^h)$, given interest-rate policy, will be given by

$$L_t = \lambda(\lambda \alpha_h^2 + \gamma_h^2) \theta_t^2,$$

where θ_t is the function of financial wedges and interest rates defined in (4.6).

Given that these results obtain regardless of the assumed interest-rate policy, the choice of an optimal interest-rate policy reduces to the choice of a state-contingent evolution $\{\hat{i}_t\}$ subject to the lower bound (2.12) holding at all times, so as to minimize the expected value of $\sum_{t=0}^{\infty} \beta^t \theta_t^2$, where θ_t is given by (4.6). If the ZLB never binds, this problem will be solved by choosing $\hat{i}_t = -\hat{\Delta}_t$ each period, so that $\check{\Delta}_t = 0$ at all times, implying that $\theta_t = 0$ at all times. However, this will be possible if and only if financial wedges are never large, i.e., $\tilde{\Delta}_t = 0$ at all times — the same condition as is required for complete stabilization to be possible under the constraint that $b_{t+1} = 0$ at all times. While the availability of countercyclical fiscal transfers as an additional

policy instrument reduces the losses associated with a given process $\{\check{\Delta}_t\}$ for the financial wedges not offset by contemporaneous interest-rate adjustments, it does not change the fact that complete stabilization requires (except in the case where $h = 0$) that one be able to ensure that $\check{\Delta}_t = 0$ at all times. As discussed in section 2.2, this is sometimes precluded by the ZLB.

We can also see that in general, when complete stabilization is not possible, the optimal second-best policy will involve committing to maintain $b_{t+1} > 0$ and/or $\hat{i}_t < 0$ (that is, deviation from the policies associated with the long-run steady state) even in some periods t after the financial wedge has again become small, so that an immediate return to the long-run steady state would be possible. Suppose instead that one were to have $\hat{i}_t = 0$ (and hence $\check{\Delta}_t = 0$) for all $t \geq T$, where T is the (possibly random) date at which reversion to the “normal” state occurs, while $\check{\Delta}_t$ is instead necessarily positive (because of the ZLB) at all dates $0 \leq t < T$. It would then follow that at any date t (and in any state of the world at that date) at which the financial wedge remains large, $\check{\Delta}_t > 0$ and $\check{\Delta}_{t+j}$ is also anticipated to be non-negative in all possible successor states with $j > 0$; hence the right-hand side of (4.5) will necessarily be positive.⁴⁹ We can thus conclude that optimal policy would require that $E_t b_{t+h+1} > 0$.

If we further suppose that t is a date such that the financial wedge remains large at date t , but it is foreseen that it will necessarily be small at date $t + h$, then this requires that $b_{t+h+1} > 0$ with positive probability even after reversion to the normal state. Furthermore, on the assumption that the economy is already in the normal state at date $t + h$, and hence that $\hat{i}_{t+h} = 0$ in all possible states at that date, a policy under which $b_{t+h+1} > 0$ in some state s_{t+h} will also have to involve $\pi_{t+h}^0 > 0$ in that state.⁵⁰ Thus the optimal joint fiscal-monetary policy must also involve an understanding that inflation would be allowed to overshoot its long-run target, even in some periods $t \geq T$. Stabilization outcomes during the period when the financial wedge remains large (and the ZLB consequently binds) are improved by committing to continue expansionary policies for a time beyond the date T at which it would be possible to again achieve complete stabilization using orthodox (and purely forward-looking) policies.

⁴⁹Here we use the fact that all elements of the vector a and the matrix A are positive, implying that $\alpha_j, \gamma_j > 0$ for all $j \geq 0$.

⁵⁰This follows from (4.2), given that $\gamma_h > 0$.




Figure 3: Equilibrium trajectories in the case of an elevated financial wedge for 10 quarters (panel (a)), under three alternative assumptions about policy: (i) $b_{t+1} = 0$ at all times, and $\hat{i}_t = \max\{-\hat{\Delta}_t, \hat{i}\}$; (ii) $b_{t+1} = 0$ at all times, but the path $\{\hat{i}_t\}$ is chosen optimally; or (iii) the paths of both $\{b_{t+1}\}$ and $\{\hat{i}_t\}$ are chosen optimally. Planning horizons extend 8 quarters into the future, and t measures quarters since the onset of the elevated financial wedge.

This is illustrated by a numerical example in Figure 3.⁵¹ In this example, all households and firms are assumed to have planning horizons extending eight quarters into the future ($h = 8$); we furthermore assume for simplicity that it is known from $t = 0$ onward that the financial wedge will be elevated for exactly ten quarters (rather than assuming stochastic exit from the crisis state, as in the two-state Markov case), so that $T = 10$ with certainty.

The several panels of the figure show the (deterministic) evolution of the financial

⁵¹We discuss further the calculations involved and show how the results depend on the assumed planning horizon, in Appendix C.

wedge, output, inflation, the nominal interest rate, and the real public debt in response to such a disturbance,⁵² under three possible assumptions about monetary and fiscal policy, which is to say about the paths of $\{\hat{i}_t\}$ and $\{b_{t+1}\}$. (Both of these evolve deterministically under all of the policies considered, since no further uncertainty is resolved after date $t = 0$.) In case (i), we assume that \hat{i}_t tracks the variation in the “natural rate of interest” (the interest rate required for stabilization of the output gap, as specified in (3.5)), and that $b_{t+1} = 0$ (no response of fiscal policy to the disturbance). These assumptions lead to the same outcomes as under the “orthodox” policy discussed in section 2.2.

In case (ii), we again assume that $b_{t+1} = 0$ at all times, but consider optimal forward guidance with respect to the future evolution of the central bank’s nominal interest-rate target $\{\hat{i}_t\}$. In case (iii), we instead allow the paths of both $\{\hat{i}_t\}$ and $\{b_{t+1}\}$ to be optimized. In the latter case, we see that the optimal joint fiscal-monetary commitment involves promising to maintain both $b_{t+1} > 0$ and $\hat{i}_t < 0$ for a time after the reversion to the normal state in quarter 10. The figure (panel (e)) shows that optimal policy requires an increase in the public debt (by an amount equal to nearly two years’ GDP) by at least quarter 8, which must then be maintained in quarter 9. In quarter 10, when the financial wedge has returned to zero (see panel (a)), it continues to be optimal to maintain a larger public debt than in the long-run steady state (though not as large as the debt in quarters 8 and 9); the optimal level of the debt continues to be somewhat positive in quarters 11 and later, though much smaller than the earlier levels of debt.

The optimal joint fiscal-monetary commitment also involves keeping the nominal interest rate lower than its long-run steady-state level for two quarters following the reversion of the financial wedge to zero. Panel (d) of the figure shows that under policy (iii), the nominal interest rate remains at the zero lower bound in quarter 10, even though it would be possible at this time to return immediately to the long-run steady state (and policy (i) would require $\hat{i}_t = 0$ from $t = 10$ onward). The nominal interest rate also remains well below its long-run steady-state level in quarter 11,

⁵²Here the financial wedge, inflation, and the nominal interest rate are reported in annualized terms: $\hat{\Delta}_t = 0.05$ means a safety premium of 5 percentage points per year (and is equivalent to the value $\hat{\Delta} = 0.013$ given in Table 1, where the value is for a quarterly model). Output is reported as a percentage deviation from the long-run steady-state level of output, and real public debt in units of years of long-run steady-state real GDP.

though no longer at the lower bound.

As in rational-expectations analyses of optimal forward guidance, a commitment to keep the interest rate “low for longer” following the reversion of the financial wedge to zero improves stabilization during earlier periods when the financial wedge is large (and the ZLB precludes complete stabilization as a result). Indeed, the degree to which it is optimal to commit to keep interest rates low beyond date $T = 10$ is similar in the case when fiscal transfers are used optimally (case (iii)) as in the case where fiscal transfers cannot be used (case (ii)). Despite the optimal use of counter-cyclical transfer policy, it remains valuable for the central bank to communicate that it will not quickly return to pursuit of its normal targets following a period in which the financial wedge has been so elevated as to cause the ZLB to bind.

4.3 Optimal Policy with a Welfare-theoretic Stabilization Objective

It might be thought that the conclusions about the role of forward guidance in the previous section depend on our having assumed (unrealistically) that all households and firms have planning horizons of exactly the same length, while forward guidance about policy after the financial wedge reverts to a normal level would be unnecessary in the case of a heterogeneous distribution of planning horizons of the kind assumed in section 3. But the results in section 3 show only that in the case of an exponential distribution of planning horizons, a relatively simple policy commitment suffices to completely stabilize both an overall price index (or inflation rate) and aggregate output. This implies that the loss function (4.4) can be minimized by such a policy; yet the proposed policy does not really eliminate all distortions in the allocation of resources.

In a representative-household model, a loss function of the form (4.4) can be justified as a quadratic approximation to the level of expected utility of the representative household, under conditions discussed by Woodford (2003) and Benigno and Woodford (2005). However, those derivations apply to a model in which the aggregate output measure Y_t (or rather, $Y_t - G_t$, where G_t is the quantity of the composite good consumed by the government) represents the quantity of the composite good consumed by each household, and in which all firms that reconsider their price in period t choose the same (optimal) new price, so that the only reason for the prices of

different goods to be misaligned is that different firms adjust their prices at different dates.

In the model considered here, instead, if the planning horizons of different households and firms are heterogeneous, then households with different planning horizons will generally consume different amounts at a given point in time, and firms with different planning horizons will generally set different prices even when they adjust their prices at the same point in time. This creates additional sources of inefficiency in the allocation of resources: non-uniform allocation of the goods produced at a given date to the different households reduces average utility (and hence reduces the representative household's *ex ante* expected utility, since households do not know *ex ante* which planning horizon they will have), and dispersion in the prices of the goods supplied by different firms (even though the aggregate inflation rate never varies) means that the composite good will be obtained in a way that uses more resources than necessary (because the quantities supplied of the different differentiated goods will not be uniform).

Using the same methods as are explained in Woodford (2003) for the representative-household case, a quadratic approximation to the average level of expected utility in the economy (averaging over the agents with different planning horizons) leads to a loss function for stabilization policy of a more general form,⁵³

$$E_0 \sum_{t=0}^{\infty} \beta^t [\pi_t^2 + \alpha^{-1} \text{var}(\pi_t^h) + \lambda_{agg} y_t^2 + \lambda_{disp} \text{var}(y_t^h)], \quad (4.7)$$

where $\text{var}(\pi_t^h)$ and $\text{var}(y_t^h)$ measure the dispersion of the values of π_t^h and y_t^h respectively, across decision makers with heterogeneous planning horizons at a given point in time (and in a particular state of the world).⁵⁴ Here $0 < \alpha < 1$ is the fraction of price-setters that do not reconsider their price from one period to the next (a measure of price stickiness in the Calvo model of price adjustment), and the coefficients $\lambda_{agg} > \lambda_{disp} > 0$ are functions of the model's underlying micro parameters. In the case of the parameter values specified in Table 1, the weight on $\text{var}(\pi_t^h)$ is 1.29 times

⁵³See Appendix D for details of the derivation.

⁵⁴For example, we define $\text{var}(\pi_t^h) \equiv \sum_h \omega_h (\pi_t^h)^2 - (\sum_h \omega_h \pi_t^h)^2$. Like the aggregate inflation rate $\pi_t \equiv \sum_h \omega_h \pi_t^h$, this is a random variable that takes a particular value in each possible state of the world at each date t . Thus the variance refers to the distribution of different values of h in the population, not to uncertainty about which of various future states of the world will be reached.

the weight on π_t^2 , while the weight on $\text{var}(y_t^h)$ is 0.43 times the weight on y_t^2 ; thus the dispersion terms are of non-trivial significance.

Note that in the case that all households and firms have a common planning horizon h , as assumed in section 4.2, $\text{var}(\pi_t^h) = \text{var}(y_t^h) = 0$ at all times, regardless of the policy chosen, and in this case (4.7) reduces to the simpler loss function (4.4) assumed above. Thus the characterization of optimal policy given in section 4.2 continues to be correct if the criterion is the welfare-based objective (4.7). However, the welfare-based objective has different consequences in the case of heterogeneous planning horizons.

For example, we can reconsider the question of optimal stabilization policy in the case of an exponential distribution of planning horizons, as assumed in section 3. This case is of particular interest, apart from the opportunity it provides to reconsider the results of section 3. One of the less appealing features of the optimal policy exercise reported in section 4.2 is that the criterion used to judge which policy is best is based on the projected evolution of the economy by the policymaker (i.e., by our model), who knows that each period's actual outcomes will be determined by households with planning horizon h , whereas people's subjective assessments of their welfare will be based on their own expectations of the economy's future evolution, which assume that outcomes at future dates will be determined by the decision rules of people who all have horizons shorter than h .

Thus there can be a systematic difference between the consequences of the policy projected in the welfare-evaluation exercise and those projected by the people in the economy. Indeed, it might be judged desirable to create expectations of outcomes that people won't like (based on their analysis of the future behavior of short-horizon decision makers) in order to get people to do things earlier that the policymaker likes; in the proposed welfare analysis, the policymaker need not be concerned by such expectations because *actually* the later outcomes will be different (because people will actually have longer planning horizons at the later date than they had earlier projected themselves to have). Such an analysis, while correct under its assumptions, raises questions about the degree of public support that can be expected for the "optimal" policy.

This problem is much less glaring in the case of an exponential distribution of planning horizons. If the distribution of planning horizons is given by (2.11), then (for example) the future inflation rate predicted at date t by the policymaker for some

future date $\tau > t$ will be given by (the state-contingent value of)

$$y_\tau = (1 - \rho) \sum_{h=0}^{\infty} \rho^h \pi_\tau^h. \quad (4.8)$$

Instead, a private decision maker with planning horizon h who projects an inflation rate for that state of the world at date τ as part of a forward planning exercise at time t will project an inflation rate of $\pi_\tau^{t+h-\tau}$. But an inflation rate for period τ will only be projected by decision makers with horizons $h \geq \tau - t$. Among these, fraction $(1 - \rho)$ have horizon $h - (\tau - t)$, fraction $(1 - \rho)\rho$ have horizons $h + 1 - (\tau - t)$, and so on; hence fraction $(1 - \rho)$ project the inflation rate π_τ^0 , fraction $(1 - \rho)\rho$ project the rate π_τ^1 , and so on. It follows that the average inflation rate projected for period τ by private decision makers (among those who calculate a projected future inflation rate at all) will be given by the right-hand side of (4.8).

There will thus be no systematic difference between the inflation rate projected by the policymaker under a given policy and the average inflation rate predicted by private decision makers. Of course, it is still the case that to the extent that decision makers with different planning horizons make different projections, they will not all agree with the policymaker's projection; they will only agree with it on average. But the stabilization objective (4.7) also directs the policymaker to prefer policies that reduce the extent to which people with different planning horizons expect different future outcomes. Thus a policy chosen to minimize (4.7) in the case of an exponential distribution of planning horizons does not result in such an uncomfortable degree of difference between the outcomes projected by the policymaker and those expected by people in the economy.

In the case of an exponential distribution of planning horizons, we have seen (in section 3) that the policy that minimizes the ad hoc loss function is one that completely stabilizes both output and inflation at all times. But such a policy is not optimal from the standpoint of the welfare-based criterion (4.7), because it leads to substantial dispersion in the forecasts of decision makers with different planning horizons, as discussed in section 4.1 for the case of a Markovian disturbance process. Here we consider (as in section 4.2) the simpler case of a disturbance that increases the financial wedge for T periods, with the value of T known with certainty at the time of the shock.

In Figure 4, we consider a financial disturbance expected to last for $T = 10$ quarters, and assume for purposes of our numerical illustration an exponential distribution




Figure 4: The paths of real spending (top panel) and inflation (bottom panel) at each date t (the horizontal axis) as determined by the decisions of agents with each of several different planning horizons h . As in Figure 3, the shock is assumed to result in an elevated financial wedge for 10 quarters. The policy is one that completely stabilizes aggregate output and inflation in all periods, in the case of an exponential distribution of planning horizons with mean horizon $\bar{h} = 8$ quarters.

of planning horizons with mean horizon $\bar{h} = 8$ quarters. If the fiscal-monetary policy discussed in section 3 is followed, it is possible to completely stabilize the paths of both aggregate output y_t and inflation π_t . The figure shows the projected time paths of π_t^h and y_t^h implied by (4.2) for the 20 quarters following the shock, for each of several different planning horizons h .

In this thought experiment, everyone understands that it is possible to return to the economy's long-run steady state from period T onward, and the policy authorities are expected to adopt policies consistent with this; thus everyone agrees on the expectations $y_t^h = \pi_t^h = 0$ for all $t \geq 10$. But in the periods with the elevated financial wedge, the behavior of decision makers with different planning horizons are not at all the same. All those with planning horizons that do not extend to period T spend *more* during this period than they would normally, as a result of their expectation that the public debt (which is increased in response to the shock) will not yet have




Figure 5: The paths of aggregate variables in the case of the disturbance considered in Figures 3 and 4, in an economy with the same distribution of planning horizons as in Figure 4. Two policies are compared: a policy that completely stabilizes aggregate output and inflation, as in Figure 4 (dashed lines), and the policy that maximizes average expected utility (solid lines). The variables shown are as in Figure 3.

been paid back at the end of their planning horizon. But all those whose planning horizons extend to date T or beyond spend *less* than normally, just as would be the case in the absence of expansionary fiscal policy (given that there is no monetary easing beyond date T), as they expect tax increases over their planning horizon that fully offset the effect of the initial fiscal transfers. Because the expected level of the public debt is very different before and after date T under this policy, the decisions made by people with horizons that do or do not extend to date T are notably different.

Because of this, not all terms in (4.7) are reduced to zero by such a policy, and one can do better under a policy that does not promise such an abrupt change in policy around date T . Figure 5 shows how the policy that minimizes (4.7) differs from the

policy derived in section 3. Here we optimize over policies specified by deterministic paths $\{\hat{i}_\tau, b_{\tau+1}\}$ for all $\tau \geq 0$, announced at time $t = 0$ and similarly understood by all decision makers that plan far enough ahead (at any date $t \geq 0$) to model economic conditions at date τ . (We can specify policy by deterministic paths for the policy variables because all uncertainty about fundamentals is resolved in period $t = 0$.) Under any policy of this kind, the paths $\{y_t^h, \pi_t^h\}$ for all $t \geq 0$ implied by the policy are given by (4.2). This allows us to evaluate (4.7) for any such policy; this objective is optimized over the set of possible sequences $\{\hat{i}_\tau, b_{\tau+1}\}$.⁵⁵

Panels (d) and (e) of the figure compare the paths $\{\hat{i}_\tau, b_{\tau+1}\}$ that minimize the welfare-based objective (4.7), shown by solid lines, to the ones that minimize the ad hoc objective (4.4), shown by dashed lines. (The latter policy is the one discussed in section 3.) The policy that minimizes the ad hoc objective reduces the nominal interest rate to its lower bound and increases the public debt by several multiples of GDP. It maintains both variables at these constant (very accommodative) levels until date T , but returns both variables immediately to their long-run steady-state levels from date T onward.⁵⁶ Instead, the welfare-optimal policy involves a less aggressive fiscal stimulus during the period of the elevated financial wedge, which is also withdrawn less abruptly after time T . Instead, forward guidance regarding future interest-rate policy is also used as an additional stimulus to demand during the period with the elevated financial wedge. Interest rates are not raised to their normal level until three quarters after the financial disturbance has dissipated (and indeed remain at the lower bound for the first two quarters after the reversion to normal fundamentals).

This alternative policy does not fully stabilize the economy from time T onward (even though this would be possible, as shown in Figure 4), and it is less successful than the other policy in stabilizing aggregate output and inflation in the period before time T as well, as shown in panels (b) and (c) of Figure 5. But it has the virtue of producing less disagreement about optimal action across decision makers with different planning horizons. This is shown in Figure 6, where the dashed and solid lines correspond to the same two policies as in Figure 5. This makes the policy

⁵⁵See Appendix E for details of our numerical method.

⁵⁶This is essentially the policy that is recommended by Gabaix (2020) to minimize an ad hoc objective of the form (4.4) in his model with cognitive discounting of future outcomes. Our model agrees with his conclusions as to the effects of such a policy on aggregate variables, under the assumption of an exponential distribution of planning horizons, but it does not imply that this is the policy that is best for private welfare.




Figure 6: The paths of the dispersion measures over time, in the case of the same disturbance and same two policies as are considered in Figure 5.

superior, from the standpoint of criterion (4.7), despite larger values for the terms in the loss function measuring departures of aggregate output and inflation from their target values.

Just as in the case where everyone is assumed to have a common (but finite) planning horizon (Figure 3), the optimal combined monetary-fiscal regime involves a commitment to maintain more accommodative policies than normal, using both monetary and fiscal instruments, beyond the date T at which it becomes possible to fully stabilize the economy (that is, to achieve zero values for *all* of the terms in (4.7), not just stabilization of aggregate output and inflation). It is optimal to commit to a more expansionary policy in the period immediately following the reversion of fundamentals, not because of any desirable effects from period T onward, but purely for the sake of better outcomes during the period of the elevated financial wedge, owing to the anticipation of more accommodative policies later. Thus forward guidance continues to be a valuable component of an optimal policy, despite the fact that planning horizons are finite (and indeed, the horizons of some decision makers are assumed to be quite short), and also despite the fact that there are now two instruments of stabilization policy.

And just as in the case of the common planning horizon, the optimal use of interest-rate forward guidance is fairly similar, even when fiscal transfer policy is also available as a tool of stabilization (and is optimally used), as it would be under the




Figure 7: The same three policies are compared as in Figure 3, but now for the case of an economy with the heterogeneous distribution of planning horizons assumed in Figure 5. The format and the assumed disturbance are as in Figures 3 and 5.

assumption of no use of fiscal transfers for stabilization purposes. Figure 7 compares the optimal coordinated monetary-fiscal policy with the optimal interest-rate path in the case that transfer policy cannot be used, as well as with the equilibrium outcome if neither transfer policy nor interest-rate forward guidance can be used (as in Figure 3, but now under the assumption of an exponential distribution of planning horizons). The format of the figure is the same as in Figure 3. One sees that the optimal interest-rate path is fairly similar regardless of whether fiscal policy can also be used for stabilization policy; however, use of the additional fiscal instrument considerably improves on the degree to which both output and inflation can be stabilized in the period of the financial disturbance.

While the numerical example presented is a special case, it illustrates a fairly general point. Once one recognizes that there are many different distortions each

period (associated with the separate pricing and expenditure decisions of agents with different planning horizons) that must all be set to zero in order to achieve the first-best allocation of resources, it is evident that a first-best outcome will not generally be achievable, even when both fiscal transfer policy and interest-rate policy are available as instruments (assuming that neither can be separately targeted to households or firms depending on their planning horizon).

And the fact that the first-best allocation is not achievable makes it not generally optimal to fully stabilize both inflation and output as soon as the financial wedge reverts to its normal level, even though this would be optimal if one only cared about outcomes from date T onward. The reason is that changing anticipated outcomes after date T can change the allocation of resources prior to date T (when stabilization is incomplete, even under the second-best optimal policy), and it will be optimal to choose at least some degree of distortion after date T for the sake of reducing the much larger distortions before date T , just as in the RE analysis of Eggertsson and Woodford (2003). Hence we should expect in general that forward guidance regarding policy after date T will be useful in mitigating the distortions created by the large financial wedge of the crisis period.

5 Concluding Remarks

In this paper, we reconsider the nature of effective stabilization policy when the zero lower bound is a relevant constraint on the effectiveness of conventional monetary policy by relaxing the unrealistic assumption that people should be able to deductively reason about the economy's future evolution under a novel policy regime arbitrarily far into the future. We examine the robustness of conclusions about the consequences of particular combined monetary-fiscal regimes to changes in the assumed degree of decision makers' foresight in the economy. We find that when planning horizons are finite, the contractionary effects of a financial disturbance are less dramatic than in the rational-expectations analysis. But, as long as there is some degree of foresight, even a relatively modest financial wedge can substantially impact stabilization goals, if additional tools of stabilization policy beyond those needed under normal circumstances are not available.

Given that Ricardian equivalence does not hold when people have finite horizons, we consider in particular the extent to which pure variation in the government's

budget balance, i.e., changes in the size of lump-sum transfers, can serve as a tool of stabilization policy. We show that fiscal transfers can be a powerful tool to reduce the contractionary impact of a financial disturbance and can even make possible a complete stabilization of both aggregate output and inflation, despite the binding ZLB constraint. But the power of fiscal transfers relies on the degree of monetary accommodation of such transfers.

Moreover, neither the availability of transfer policy nor the fact that the length of planning horizons is bounded makes commitments about interest-rate policy beyond the date at which the financial disturbance has dissipated, of the kind argued for in the rational-expectations analysis of Eggertsson and Woodford (2003), no longer relevant. We show that the use of such forward guidance along with fiscal policy achieves better stabilization outcomes than fiscal policy alone would achieve under an understanding that the central bank will return to pursuit of its usual inflation target once the financial wedge is again modest in size.

In a numerical example that we present, the degree to which it is optimal to commit to a continuation of looser monetary policy beyond the time at which fundamentals have reverted to normal is roughly the same in the case of an optimal state-contingent transfer policy, as in the case of no response of fiscal transfers to the disturbance at all. Thus, while the rational-expectations analysis exaggerates the quantitative effects of forward guidance policies as a response to the kind of financial disturbance considered here (as stressed in the literature on the “forward guidance puzzle” cited in the introduction), a commitment to keeping interest rates “lower for longer” following a crisis that causes the ZLB to become a binding constraint continues to be desirable in the framework for policy analysis proposed here.

References

- [1] Andrade, Philippe, Gaetano Gaballo, Eric Mengus, and Benoit Mojon, “Forward Guidance and Heterogeneous Beliefs,” *American Economic Journal: Macroeconomics* 11: 1-29 (2019).
- [2] Angeletos, George-Marios, and Chen Lian, “Forward Guidance Without Common Knowledge,” *American Economic Review* 108: 2477-2512 (2018).
- [3] Ascari, Guido, and Neil Rankin, “Effectiveness of Government Debt for Demand Management: Sensitivity to Monetary Policy Rules,” *Journal of Economic Dynamics and Control* 37: 1544-1566 (2013).
- [4] Benigno, Pierpaolo, and Michael Woodford, “Inflation Stabilization and Welfare: The Case of a Distorted Steady State,” *Journal of the European Economic Association* 3: 1185-1236 (2005).
- [5] Caballero, Ricardo J., and Emmanuel Farhi, “The Safety Trap,” *Review of Economic Studies* 85: 223-274 (2017).
- [6] Calvo, Guillermo A., “Staggered Prices in a Utility-Maximizing Framework,” *Journal of Monetary Economics* 12: 383-398 (1983).
- [7] Christiano, Lawrence, Martin Eichenbaum, and Sergio Rebelo, “When Is the Government Spending Multiplier Large?” *Journal of Political Economy* 119: 78-121 (2011).
- [8] Coibion, Olivier, Dimitris Georgarakos, Yuriy Gorodnichenko, and Michael Weber, “Forward Guidance and Household Expectations,” working paper, University of Texas, Austin, February 2020.
- [9] Del Negro, Marco, Domenico Giannone, Marc P. Giannoni, and Andrea Tambalotti, “Safety, Liquidity and the Natural Rate of Interest,” *Brookings Papers on Economic Activity*, Spring 2017, pp. 235-294.
- [10] Del Negro, Marco, Marc P. Giannoni, and Christina Patterson, “The Forward Guidance Puzzle,” Federal Reserve Bank of New York Staff Report no. 574, revised December 2015.

- [11] Eggertsson, Gauti B., “What Fiscal Policy Is Effective at Zero Interest Rates?” *NBER Macroeconomics Annual* 25: 59-122 (2011).
- [12] Eggertsson, Gauti B., and Michael Woodford, “The Zero Bound on Interest Rates and Optimal Monetary Policy,” *Brookings Papers on Economic Activity*, Spring 2003, pp. 139-211.
- [13] Farhi, Emmanuel, and Ivan Werning, “Monetary Policy, Bounded Rationality, and Incomplete Markets,” *American Economic Review* 109: 3887-3928 (2019).
- [14] Gabaix, Xavier, “A Behavioral New Keynesian Model,” *American Economic Review* 110: 2271-2327 (2020).
- [15] García-Schmidt, Mariana, and Michael Woodford, “Are Low Interest Rates Deflationary? A Paradox of Perfect-Foresight Analysis,” *American Economic Review* 109: 86-120 (2019).
- [16] Gust, Christopher, Edward Herbst, and David López-Salido, “Forward Guidance with Bayesian Learning and Estimation,” Finance and Economics Discussion Series no. 2018-072, Federal Reserve Board, September 2018.
- [17] Gust, Christopher, Edward Herbst, and David López-Salido, “Short-Term Planning, Monetary Policy, and Macroeconomic Persistence,” Finance and Economics Discussion Series no. 2020-003, Federal Reserve Board, December 2019.
- [18] Ho, Teck-Hua, and Xuanming Su, “A Dynamic Level- k Model in Sequential Games,” *Management Science* 59: 452-469 (2013).
- [19] Iovino, Luigi, and Dmitriy Sergeyev, “Central Bank Balance Sheet Policies Without Rational Expectations,” working paper, Bocconi University, June 2020.
- [20] Johnson, Eric J., Colin Camerer, Sankar Sen, and Talia Rymon, “Detecting Failures of Backward Induction: Monitoring Information Search in Sequential Bargaining,” *Journal of Economic Theory* 104: 16-47 (2002).
- [21] Keramati, Mehdi, Peter Smittenaar, Raymond J. Dolan, and Peter Dayan, “Adaptive Integration of Habits into Depth-Limited Planning Defines a Habitual-Goal-Directed Spectrum,” *Proceedings of the National Academy of Sciences* 113: 12868-12873 (2016).

- [22] Leeper, Eric, “Equilibria under ‘Active’ and ‘Passive’ Monetary and Fiscal Policies,” *Journal of Monetary Economics* 27: 129-147 (1991).
- [23] McKelvey, Richard D., and Thomas R. Palfrey, “An Experimental Study of the Centipede Game,” *Econometrica* 60: 803-836 (1992).
- [24] Mertens, Thomas, and John C. Williams, “What to Expect from the Lower Bound on Interest Rates: Evidence from Derivatives Prices,” Federal Reserve Bank of New York Staff Report no. 865, August 2018.
- [25] Nagel, Rosemarie, “Unraveling in Guessing Games: An Experimental Study,” *American Economic Review* 85: 1313-1326 (1995).
- [26] Nagel, Rosemarie, and Fang Fang Tang, “Experimental Results on the Centipede Game in Normal Form: An Investigation on Learning,” *Journal of Mathematical Psychology* 42: 356-384 (1998).
- [27] Selten, Reinhard, and Rolf Stoecker, “End Behavior in Sequences of Finite Prisoner’s Dilemma Supergames: A Learning Theory Approach,” *Journal of Economic Behavior and Organization* 7: 47-70 (1986).
- [28] Taylor, John B., “Discretion Versus Policy Rules in Practice,” *Carnegie-Rochester Conference Series on Public Policy* 39: 195-214 (1993).
- [29] Taylor, John B., “Fiscal Stimulus Programs During the Great Recession,” Economics Working Paper no. 18117, Hoover Institution, December 2018.
- [30] Woodford, Michael, “Fiscal Requirements for Price Stability,” *Journal of Money, Credit and Banking* 33: 669-728 (2001).
- [31] Woodford, Michael, *Interest and Prices: Foundations of a Theory of Monetary Policy*, Princeton: Princeton University Press, 2003.
- [32] Woodford, Michael, “Simple Analytics of the Government Expenditure Multiplier,” *American Economic Journal: Macroeconomics* 3: 1-35 (2011).
- [33] Woodford, Michael, “Monetary Policy Analysis when Planning Horizons are Finite.” *NBER Macroeconomics Annual* 33: 1-50 (2019).

- [34] Woodford, Michael, and Yinxi Xie, “Policy Options at the Zero Lower Bound when Foresight is Limited,” *AEA Papers and Proceedings* 109: 433-437 (2019).
- [35] Xie, Yinxi, “Fiscal and Monetary Policy Interaction under Limited Foresight,” working paper, Columbia University, January 2020. [Available at: econ.columbia.edu/e/yinxi-xie/]

A Output and Inflation Contraction in the Crisis State in the Absence of a Policy Response

Here under the assumption that the financial wedge evolves as a two-state Markov chain, we show the derivation for equation (2.19), i.e., the expressions of output and inflation in the crisis state in the absence of a policy response as in section 2.2.2. In the text, we have already shown that for the case of $h = 0$, $z_t^0 = 0$ if t is any date after the reversion to the normal state, while

$$z_t^0 = \underline{z}^0 \equiv -\sigma a \underline{\Delta} \lll 0 \tag{A.1}$$

if t is any date at which the crisis state continues.

We can then use this result to solve recursively for the behavior of households and firms with progressively longer planning horizons. First we observe that if t is any date after the reversion to the normal state, $z_t^h = 0$ for all h . This can be established recursively; we first show that if $z_t^h = 0$ for all dates after the reversion to normal for some horizon $h \geq 0$, (2.16) implies that inflation targeting will require $\hat{i}_t^{h+1} = 0$ at any date after the reversion to normal, and hence that $z_t^{h+1} = 0$ as well. Then the fact that we have already shown that $z_t^0 = 0$ after the reversion to normal implies that $z_t^h = 0$ for all h .

Next, consider instead dates t at which the crisis state continues, and suppose that it has already been established for some horizon h that in any crisis state, $z_t^h = \underline{z}^h$, where \underline{z}^h is a vector that is negative in both elements. Then it follows from (2.16) that even if $\hat{i}_t^{h+1} = \hat{i}$ (the most expansionary possible monetary policy that can be expected), in any crisis state the vector z_t^{h+1} will equal

$$z_t^{h+1} = \mu A \underline{z}^h - \sigma a \underline{\Delta} \lll 0.$$

Hence the ZLB will necessarily bind, and we will have $z_t^{h+1} = \underline{z}^{h+1}$ in any such state, where

$$\underline{z}^{h+1} = \mu A \underline{z}^h - \sigma a \underline{\Delta} \lll 0. \tag{A.2}$$

It follows that the equilibrium will be Markovian,⁵⁷ and that the sequence of vectors

⁵⁷Note that this is not an assumption (equilibrium selection criterion), as in the case of the rational-expectations analysis; we have shown that in the case of finite-horizon planning, there is necessarily a unique solution, and that it has this property.

$\{z^h\}$ characterizing the Markovian equilibrium can be computed recursively, using (A.2) together with the initial condition (A.1).

This system of equations can be recursively solved to yield

$$z^h = -\sigma \sum_{j=0}^h (\mu A)^j a \underline{\Delta} \ll 0$$

for any planning horizon $h \geq 0$.

B Output and Inflation Stabilization with an Exponential Distribution of Planning Horizons

Here we demonstrate that the combination of a monetary policy specified by (3.5) and a fiscal policy specified by (3.6) each period imply complete stabilization of aggregate output and inflation at all times, in the case of an exponential distribution of planning horizons. In the text, we have already shown that this monetary rule implies that the spending and price-increase decisions of households and firms with an arbitrary planning horizon h are given by equation (4.2). This is a well-defined, unique solution, independent of any assumption about the distribution of planning horizons in the economy. There will therefore exist a well-defined, unique solution in the case of an exponential distribution of planning horizons (2.11) if and only if the infinite sums

$$z_t = (1 - \rho) \sum_{h=0}^{\infty} \rho^h z_t^h \tag{B.3}$$

converge, where z_t^h is given by (4.2). This is the issue that remains to be addressed.

Let us first consider the partial sum that aggregates the decisions of only the part of the population with horizons less than or equal to k periods,

$$z_t^{(k)} = (1 - \rho) \sum_{h=0}^k \rho^h z_t^h$$

for some finite k . This finite sum is obviously well defined; it remains to be determined whether the sequence $\{z_t^{(k)}\}$ converges as k is made large.

Substituting (4.2) into this definition yields

$$\begin{aligned}
z_t^{(k)} &= -\sigma(1-\rho) \sum_{h=0}^k \rho^h \sum_{j=0}^h [A^j a] E_t \tilde{\Delta}_{t+j} + (1-\beta)(1-\rho) \sum_{h=0}^k \rho^h [A^h a] E_t b_{t+h+1} \\
&= -\sigma(1-\rho) \sum_{h=0}^k \sum_{j=h}^k \rho^j [A^h a] E_t \tilde{\Delta}_{t+h} + (1-\beta)(1-\rho) \sum_{h=0}^k \rho^h [A^h a] E_t b_{t+h+1} \\
&= \sum_{h=0}^k [A^h a] E_t [(1-\beta)(1-\rho)\rho^h b_{t+h+1} - \sigma(\rho^h - \rho^{k+1})\tilde{\Delta}_{t+h}].
\end{aligned}$$

If we further substitute the fiscal rule (3.6) into this, we obtain

$$z_t^{(k)} = \sigma \rho^{k+1} \sum_{h=0}^k [A^h a] E_t \tilde{\Delta}_{t+h}. \quad (\text{B.4})$$

Thus the condition required for a well-defined solution is convergence of the sequence defined by the right-hand side of (B.4) as k becomes large.

The existence of a well-defined limit depends on the asymptotic rate of growth of the expected future excess financial wedge $E_t \tilde{\Delta}_{t+h}$. A sufficient condition for the existence of a well-defined solution is that $\tilde{\Delta}_t = 0$ with probability one beyond some finite future date T . In this case, for all $k \geq T - t$, $z_t^{(k)}$ is a constant multiple of ρ^k , and hence converges to zero as k is made large (regardless of the value of $\rho < 1$).

But this is not necessary: a weaker sufficient condition is that there exists a finite constant $C > 0$ such that $E_t \tilde{\Delta}_{t+h} \leq C \cdot \mu^h$ for all h , where $\mu \geq 0$ is a growth factor satisfying (2.20). As discussed in the text, this bound implies that both eigenvalues of A are less than $\rho^{-1}\mu^{-1}$. The partial sum $\sum_{h=0}^k [A^h a]$ is therefore positive, increases in k , and grows asymptotically with a growth factor less than $\rho^{-1}\mu^{-1}$. Hence the right-hand side of (B.4) is necessarily non-negative (since (3.7) implies that $\tilde{\Delta} \geq 0$ at all times) and bounded above by a positive sequence that converges to zero at an exponential rate as k is made large.

Hence, under this condition, the infinite sum in (B.3) is well defined, and equal to

$$z_t = \lim_{k \rightarrow \infty} z_t^{(k)} = 0.$$

Thus the specified joint fiscal-monetary regime implies the existence of a well-defined unique equilibrium, in which $y_t = \pi_t = 0$ at all times, as stated in the text. Among the cases in which a well-defined equilibrium exists is the two-state Markov process

for the financial wedge introduced in section 2.2, under the assumption that the probability μ of continuation of the crisis state satisfies (2.20), the condition already discussed in section 2.2.2 for the existence of a well-defined equilibrium in the case of a balanced-budget policy and strict inflation targeting.

C Optimal Fiscal-Monetary Policy Coordination with a Common Planning Horizon: Numerical Methods

In this section, we propose a numerical method to compute the solutions for optimal exogenous state-contingent fiscal transfer policy and interest-rate policy. Assume all agents have the same planning horizon, and the path of financial wedge is perfectly predictable, i.e., $\hat{\Delta}_t = -\hat{i}_t + \underline{\Delta}$ for $0 \leq t < T - 1$, where $\underline{\Delta} > 0$ is the excess financial wedge that cannot be offset by a reduction in nominal interest rate, and $\hat{\Delta} = 0$ for all $t \geq T$. We consider the following class of policies: the fiscal policy specified by an exogenous path of the real public debt $\{b_{t+1}\}$ and the monetary policy specified by an exogenous path of the nominal interest rate $\{\hat{i}_t\}$ consistent with the ZLB constraint (2.12).

The structural equations (2.7) and (2.9) can be written as

$$z_t^j = Az_{t+1}^j - \sigma a(\hat{i}_t + \hat{\Delta}_t)$$

for all $j \geq 1$, while (2.8) and (2.10) can be written as

$$z_t^0 = -\sigma a(\hat{i}_t + \hat{\Delta}_t) + (1 - \beta)ab_{t+1},$$

where $z_t^j = [y_t^j \ \pi_t^j]'$ and the matrices A and a are defined as in (2.13).

Since the path of $\hat{\Delta}_t$ is exogenously given, policy variables $\{\hat{i}_t, b_{t+1}\}$ can be equivalently described by the sequences of $\{\check{\Delta}_t, b_{t+1}\}$, where $\check{\Delta}_t = \hat{i}_t + \hat{\Delta}_t$. The problem of solving optimal monetary and fiscal policy is then to choose $\{\check{\Delta}_t, b_{t+1}\}$ for all $t \geq 0$, subject to the constraints that $\check{\Delta}_t \geq \underline{\Delta} > 0$ for all $0 \leq t < T$ and $\check{\Delta}_t \geq \hat{i}$ for all $t \geq T$, so as to minimize the welfare loss (4.4).

Now we characterize the solution to such an optimal fiscal and monetary policy problem. Let us first take the sequence of $\{\check{\Delta}_t\}$ as given, and derive the optimal

choice of the sequence $\{b_{t+1}\}$, which is a sequence of independent static optimization problems. More specifically, for any period $t \geq 0$, we choose b_{t+h+1} to minimize $(\pi_t^h)^2 + \lambda(y_t^h)^2$, where $\{y_t^h, \pi_t^h\}$ are given by

$$z_t^h = -\sigma \sum_{j=0}^h [A^j a] \check{\Delta}_{t+j} + (1 - \beta) [A^h a] b_{t+h+1} \quad (\text{C.5})$$

for any horizon $h \geq 0$.

Denote $[A^j a] = [\alpha_j \ \gamma_j]'$ for each $j \geq 0$, and then the F.O.C.s of the problem for optimal fiscal transfer policy are given by

$$\begin{bmatrix} \lambda \alpha_h & \gamma_h \end{bmatrix} z_t^h = 0.$$

Substituting (C.5) into the above formula yields the unique solution of b_{t+h+1} as

$$b_{t+h+1} = \frac{\sigma \sum_{j=0}^h (\lambda \alpha_h \alpha_j + \gamma_h \gamma_j) \check{\Delta}_{t+j}}{(1 - \beta) (\lambda \alpha_h^2 + \gamma_h^2)}. \quad (\text{C.6})$$

By substituting the expression of b_{t+h+1} into (C.5), the output and inflation under the optimal fiscal transfer policy conditional on a given state-contingent interest rate policy are thus given by

$$\begin{aligned} z_t^h &= -\sigma \sum_{j=0}^{h-1} \left\{ [A^j a] - \frac{(\lambda \alpha_h \alpha_j + \gamma_h \gamma_j)}{\lambda \alpha_h^2 + \gamma_h^2} [A^h a] \right\} \check{\Delta}_{t+j} \\ &= [-\sigma \sum_{j=0}^{h-1} \theta_j \check{\Delta}_{t+j}] \begin{bmatrix} \gamma_h \\ -\lambda \alpha_h \end{bmatrix}, \end{aligned}$$

where $\theta_j = \frac{\alpha_j \gamma_h - \alpha_h \gamma_j}{\lambda \alpha_h^2 + \gamma_h^2}$ for each $0 \leq j \leq h - 1$. It follows that the minimized value of the objective function $L_t \equiv (\pi_t^h)^2 + \lambda(y_t^h)^2$ is equal to

$$L_t = [\sigma \sum_{j=0}^{h-1} \theta_j \check{\Delta}_{t+j}]^2 \lambda (\lambda \alpha_h^2 + \gamma_h^2).$$

We now consider the optimal monetary policy $\{\check{\Delta}_t\}$ so as to minimize $\sum_{t=0}^{\infty} \beta^t L_t$, subject to the constraint $\check{\Delta}_t \geq \underline{\Delta} > 0$ for all $0 \leq t < T$ and $\check{\Delta}_t \geq \hat{i}$ for all $t \geq T$. The F.O.C.s of the optimal choice of $\check{\Delta}_t$ for any $t \geq h - 1$ are given by

$$\sum_{j=0}^{h-1} \beta^{-j} [\sum_{l=0}^{h-1} \theta_l \check{\Delta}_{t-j+l}] \theta_j \geq 0, \quad \check{\Delta}_t \geq \underline{\Delta}_t, \quad (\text{C.7})$$

where at least one of these inequalities must hold with equality, and $\check{\Delta}_t = \underline{\Delta}$ for all $0 \leq t < T$ and $\check{\Delta}_t = \hat{i}$ for all $t \geq T$. Instead, for any $0 \leq t < h - 1$, the F.O.C.s are given by

$$\sum_{j=0}^t \beta^{-j} [\sum_{l=0}^{h-1} \theta_l \check{\Delta}_{t-j+l}] \theta_j \geq 0, \quad \check{\Delta}_t \geq \check{\Delta}_t. \quad (\text{C.8})$$

We conjecture that the solution of $\{\check{\Delta}_t\}$ to (C.7) and (C.8) has the following form: there exists a T^* such that the ZLB binds in every period up to some date $T^* \geq 0$, and then the ZLB never binds for any dates $t \geq T^*$, i.e., for any $0 \leq t < T^*$, $\check{\Delta}_t = \underline{\Delta}_t$, while $\sum_j \beta^{-j} [\sum_{l=0}^{h-1} \theta_l \check{\Delta}_{t-j+l}] \theta_j = 0$ for all $t \geq T^*$, and $\check{\Delta}_t \rightarrow 0$ as $t \rightarrow \infty$.

Under this conjecture, for numerical purposes, we assume that there exists a large enough T^{max} such that $\check{\Delta}_t = 0$ for any date $t > T^{max}$.⁵⁸ Then (C.7) and (C.8) give a total number of $T^{max} - T^* + 1$ linear equations for the periods $T^* \leq t \leq T^{max}$ in which the ZLB is not binding, with the unknown variables $\{\check{\Delta}_t\}_{t=T^*}^{T^{max}}$. This linear system yields a unique solution if one exists. Thus we can start with $T^* = 0$ and increase the possible values of T^* until we find a value of T^* so that the unique solution $\{\check{\Delta}_t\}$ satisfies all the inequality conditions. In other words, for a given guess of T^* , we have a system of linear equations to solve with a unique solution $\{\check{\Delta}_t\}$; then we check whether the solution satisfies a sequence of inequalities (in which case it is the desired solution).

Once we get the solution of T^* and the sequence of $\{\check{\Delta}_t\}$, the optimal fiscal transfer policy is accordingly pinned down by (C.6). With the optimal fiscal and monetary policy, the realized paths of output and inflation are then given by (C.5), while the expected paths of output and inflation are given by (2.7)-(2.10).

Figure C.8 and C.9 show the equilibrium dynamics of output, inflation, interest rate, and public debt with different common planning horizons under the optimal fiscal transfer policy and monetary policy. In these figures, the economy enters the crisis state at time $t = 0$ and reverts to the normal state after 10 quarters ($T = 10$). All the variables are reported in annualized terms.

So far, we have shown the numerical methods for solving optimal combined fiscal transfer policy and interest rate policy. In order to highlight the role of fiscal transfer policy, we now consider optimal exogenous state-contingent interest rate policy but with $b_{t+1} = 0$ at all times (as included in Figure 3). In this case, output and inflation (C.5) are instead given by

$$z_t^h = -\sigma \sum_{j=0}^h [A^j a] \check{\Delta}_{t+j}.$$

Then the minimized value of the objective function $L_t \equiv (\pi_t^h)^2 + \lambda (y_t^h)^2$ is equal to

$$L_t = [\sigma \sum_{j=0}^h \gamma_j \check{\Delta}_{t+j}]^2 + \lambda [\sigma \sum_{j=0}^h \alpha_j \check{\Delta}_{t+j}]^2.$$

⁵⁸In the numerical exercise, we take $T^{max} = 200$.




Figure C.8: Equilibrium trajectories in the case of an elevated financial wedge for 10 quarters, for households and firms with relatively short common planning horizons, under the paths of both $\{b_{t+1}\}$ and $\{\hat{i}_t\}$ being chosen optimally. The planning horizon h is in quarters, and t measures quarters since the onset of the elevated financial wedge.

The F.O.C.s for the optimal monetary policy with $b_{t+1} = 0$ for any t are thus given by

$$\sum_{j=0}^h \beta^{-j} [(\sum_{l=0}^h \gamma_l \check{\Delta}_{t-j+l}) \gamma_j + \lambda (\sum_{l=0}^h \alpha_l \check{\Delta}_{t-j+l}) \alpha_j] \geq 0, \quad \check{\Delta}_t \geq \underline{\Delta}_t,$$

for any $t \geq h$, while for any $0 \leq t < h$, the F.O.C.s are given by

$$\sum_{j=0}^t \beta^{-j} [(\sum_{l=0}^h \gamma_l \check{\Delta}_{t-j+l}) \gamma_j + \lambda (\sum_{l=0}^h \alpha_l \check{\Delta}_{t-j+l}) \alpha_j] \geq 0, \quad \check{\Delta}_t \geq \underline{\Delta}_t.$$

We similarly conjecture that there exists a T^* such that the ZLB binds in every period up to some date $T^* \geq 0$, and then the ZLB never binds for any dates $t \geq T^*$. With the same numerical method as in solving the optimal combined fiscal-monetary




Figure C.9: Equilibrium trajectories in the case of an elevated financial wedge for 10 quarters, for households and firms with relatively long common planning horizons, under the paths of both $\{b_{t+1}\}$ and $\{\hat{i}_t\}$ being chosen optimally. The planning horizon h is in quarters, and t measures quarters since the onset of the elevated financial wedge.

policy problem, we can similarly solve for the optimal monetary policy of $\{\check{\Delta}_t\}$ under the assumption of $b_{t+1} = 0$ at all times.

Besides the numerical methods to calculate the optimal fiscal and monetary policy, it is also important to recognize that the superior stabilization outcomes shown in case (iii) of Figure 3 in section 4.2 depend on people's understanding that under this policy, the central bank is committed to maintaining low interest rates even if inflation and/or output overshoot their long-run target values, and even if such overshooting occurs at date $T = 10$ or later (which is to say, after complete stabilization has again become feasible). We first note that panel (c) of Figure 3 shows that under the optimal fiscal-monetary commitment, inflation is allowed to overshoot its long-run target value in quarters 0 through 8; this is because over this period, this degree

of demand stimulus remains insufficient to raise output to its target level — some degree of overshooting of inflation is tolerated in order not to have even greater undershooting of output. But we further observe that this degree of stabilization depends on allowing people to believe that (under circumstances that are actually counter-factual) monetary policy would allow both output and inflation to overshoot their targets simultaneously.

Figure C.10 shows the paths of output and inflation that are anticipated in the forward plans of households and firms, in the case (iii) equilibrium of Figure 3.⁵⁹ Here the assumed path of the financial wedge is shown again in panel (a) — because this is an exogenous variable, actual and anticipated paths coincide. In the other two panels, the paths of output and inflation that are anticipated looking forward from each date are shown. If we let $y_{\tau|t}$ (respectively, $\pi_{\tau|t}$) denote the level of output (inflation) in period τ that is anticipated during forward planning in period t , then the figure shows the paths $\{y_{\tau|t}\}$ and $\{\pi_{\tau|t}\}$ for dates $t \leq \tau \leq t + 8$, looking forward from each of a succession of dates t . The solid line in each panel shows the actually realized path of the variable (the paths of $y_t = y_{t|t}$ and $\pi_t = \pi_{t|t}$ as functions of t); the dashed lines instead show the paths of $y_{\tau|t}$ and $\pi_{\tau|t}$ as functions of τ , with a separate dashed line for each value of t .

Note that these anticipated paths involve greater overshooting of long-run targets than the actually realized paths do. Furthermore, the overshooting is anticipated to extend into the period $t \geq 10$ in which complete stabilization would be possible. For example, both inflation and output are anticipated to overshoot their long-run targets simultaneously in quarter 10, in the forward plans of people looking forward from quarters $t = 2$ through 7; both are anticipated to overshoot simultaneously in quarter 11, by people looking forward from quarters $t = 3$ through 6. Thus not only does the optimal forward guidance express an intention to keep interest rates below their long-run level in these quarters, but it requires people to believe that the central bank will do so even though (in the calculations of people considering this period several quarters earlier) this is expected to lead to overshooting of the long-run targets for both inflation and output. Despite the optimal use of counter-cyclical transfer policy, it remains valuable for the central bank to communicate that it will not quickly return to pursuit of its normal targets following a period in which the

⁵⁹Figure C.10 shows the case with a common planning horizon $h = 8$ quarters, while Figure C.11 shows the case with a longer common planning horizon $h = 20$ quarters.




Figure C.10: Dashed lines show the expected paths of output ($y_{\tau|t}^h$) and inflation ($\pi_{\tau|t}^h$) for dates $t \leq \tau \leq t + h$, under the plans calculated by households and firms with horizon $h = 8$ at successive dates t , in the case that both monetary and fiscal policy commitments are optimal (case (iii) from Figure 3). The solid lines show the predicted actual paths of output ($y_{t|t}^h$) and inflation ($\pi_{t|t}^h$). Both t and τ indicate quarters since the onset of the disturbance (again shown in panel (a)).

financial wedge has been so elevated as to cause the ZLB to bind.

D A Utility-based Welfare Loss Function with Heterogeneous Planning Horizons

Here we derive a quadratic approximation to the objective function of policymakers who maximize the average level of expected utility of households (with different planning horizons) in the economy. We use the same methods as are explained in Woodford (2003) for the representative household case, and thus make the welfare loss function with heterogeneous planning horizons comparable to the familiar one for the standard representative-agent New Keynesian model. To facilitate this comparison, in this section we use the same notation as in Woodford (2003, chap. 6).




Figure C.11: Dashed lines show the expected paths of output ($y_{\tau|t}^h$) and inflation ($\pi_{\tau|t}^h$) for dates $t \leq \tau \leq t + h$, under the plans calculated by households and firms with horizon $h = 20$ at successive dates t , in the case that both monetary and fiscal policy commitments are optimal. The solid lines show the predicted actual paths of output ($y_{t|t}^h$) and inflation ($\pi_{t|t}^h$). Both t and τ indicate quarters since the onset of the disturbance.

For any given distribution of planning horizons $\{\omega_h\}$, the average period utility of households in period t is given by

$$U_t = \sum_h \omega_h u(C_t^h; \xi_t) - \int_0^1 v(l_t(i); \xi_t) di,$$

where ξ_t is a vector of random exogenous disturbances and $v(l_t(i); \xi_t)$ is the disutility of working $l_t(i)$ hours in producing differentiated good i .

We can equivalently define a function $\tilde{v}(y; \tilde{\xi})$ indicating the disutility of supplying quantity y as follows:

$$\tilde{v}(y; \tilde{\xi}) \equiv v(f^{-1}(y/A); \xi),$$

where $Af(l)$ is the output produced with labor input and $\tilde{\xi} \equiv (\xi, a)$ denotes the complete vector of exogenous disturbances, including both the exogenous shocks ξ and the technology shock $a \equiv \log A$.

By substituting the equilibrium condition $C_t^h = Y_t^h$, where Y_t^h is what agents with planning horizon h believe that aggregate demand will be, the period utility to a second-order approximation can be written as

$$\begin{aligned}
U_t &= \Sigma_h \omega_h u(Y_t^h; \xi_t) - \int_0^1 \tilde{v}(y_t(i); \tilde{\xi}_t) di \\
&= \bar{Y} u_c \{ \Sigma_h \omega_h [\hat{Y}_t^h + \frac{1}{2}(1 - \sigma^{-1})(\hat{Y}_t^h)^2 + \sigma^{-1} g_t \hat{Y}_t^h] \\
&\quad - \int_0^1 [\hat{y}_t(i) + \frac{1}{2}(1 + \omega)(\hat{y}_t(i))^2 - \omega q_t \hat{y}_t(i)] \} + \text{t.i.p.}, \tag{D.9}
\end{aligned}$$

where $\hat{Y}_t^h \equiv \log(Y_t/\bar{Y})$ and $\hat{y}_t(i) \equiv \log(y_t(i)/\bar{Y})$.⁶⁰ The parameter σ is the intertemporal elasticity of substitution and ω is the elasticity of real marginal cost with respect to its own output (i.e., the inverse of Frisch elasticity of labor supply). As in Woodford (2003, chap. 6), given the preference shock ξ_t , the notation g_t measures the percentage variation in output required to keep the marginal utility of expenditure u_c at its steady-state level and q_t measures the percentage variation in output required to keep the marginal disutility of supply \tilde{v}_y at its steady-state level.⁶¹ The term “t.i.p.” includes all the terms that are independent of policy.

Note that $Y_t = \Sigma_h \omega_h Y_t^h$, and then to a second-order approximation, we have

$$\hat{Y}_t = \Sigma_h \omega_h \hat{Y}_t^h + \frac{1}{2} [\Sigma_h \omega_h (\hat{Y}_t^h)^2 - (\Sigma_h \omega_h \hat{Y}_t^h)^2],$$

which implies

$$\begin{aligned}
\Sigma_h \omega_h [\hat{Y}_t^h + \frac{1}{2}(1 - \sigma^{-1})(\hat{Y}_t^h)^2] &= \hat{Y}_t - \frac{1}{2} [\Sigma_h \omega_h (\hat{Y}_t^h)^2 - (\Sigma_h \omega_h \hat{Y}_t^h)^2] \\
&\quad + \frac{1}{2}(1 - \sigma^{-1}) [\Sigma_h \omega_h (\hat{Y}_t^h)^2].
\end{aligned}$$

⁶⁰We have assumed a tax policy that exactly offsets the distortion due to market power so that the steady-state output level is efficient.

⁶¹Mathematically, we define $\sigma \equiv -u_c/\bar{Y}u_{cc}$, $\omega \equiv \tilde{v}_y\bar{Y}/\tilde{v}_{yy}$, $g_t \equiv -u_{c\xi}\xi_t/\bar{Y}u_{cc}$, and $q_t \equiv -\tilde{v}_{y\xi}\xi_t/\bar{Y}\tilde{v}_{yy}$.

By plugging the above equation into the expression (D.9), we obtain

$$\begin{aligned}
U_t &= \bar{Y}u_c \left\{ -\frac{1}{2}\sigma^{-1}\Sigma_h\omega_h(\hat{Y}_t^h)^2 + \frac{1}{2}(\Sigma_h\omega_h\hat{Y}_t^h)^2 + \sigma^{-1}g_t(\Sigma_h\omega_h\hat{Y}_t^h) \right. \\
&\quad \left. - \frac{1}{2}(1+\omega)(\Sigma_h\omega_h\hat{Y}_t^h)^2 + \omega q_t(\Sigma_h\omega_h\hat{Y}_t^h) - \frac{1}{2}(\theta^{-1}+\omega)\text{var}_i(\hat{y}_t(i)) \right\} + \text{t.i.p.} \\
&= -\frac{\bar{Y}u_c}{2} \left\{ \sigma^{-1}\Sigma_h\omega_h(\hat{Y}_t^h)^2 + \omega(\Sigma_h\omega_h\hat{Y}_t^h)^2 + (\sigma^{-1}+\omega)\hat{Y}_t^n(\Sigma_h\omega_h\hat{Y}_t^h) \right. \\
&\quad \left. + (\theta^{-1}+\omega)\text{var}_i(\hat{y}_t(i)) \right\} + \text{t.i.p.}, \tag{D.10}
\end{aligned}$$

where θ is the elasticity of substitution between differentiated goods (under the assumption of CES preferences) and $\hat{Y}_t^n \equiv (\sigma^{-1}g_t + \omega q_t)/(\sigma^{-1} + \omega)$ denotes the log of the natural rate of output (i.e., the equilibrium level of output under flexible price).

Note that the assumption of CES preferences implies that

$$\text{var}_i \log y_t(i) = \theta^2 \text{var}_i \log p_t(i),$$

and then by substituting this term into the expression (D.10), we have

$$\begin{aligned}
U_t &= -\frac{\bar{Y}u_c}{2} \left\{ \sigma^{-1}\Sigma_h\omega_h(\hat{Y}_t^h)^2 + \omega(\Sigma_h\omega_h\hat{Y}_t^h)^2 + (\sigma^{-1}+\omega)\hat{Y}_t^n(\Sigma_h\omega_h\hat{Y}_t^h) \right. \\
&\quad \left. + \theta(1+\omega\theta)\text{var}_i(\log p_t(i)) \right\} + \text{t.i.p.} \tag{D.11}
\end{aligned}$$

Letting

$$\tilde{P}_t \equiv E_i \log p_t(i), \quad V_t \equiv \text{var}_i \log p_t(i),$$

and if one assumes a fraction $0 < \alpha < 1$ of all prices that remain unchanged each period as in the Calvo (1983) pricing model, it follows that

$$\begin{aligned}
V_t &= \text{var}_i[\log p_t(i) - \tilde{P}_{t-1}] \\
&= E_i\{[\log p_t(i) - \tilde{P}_{t-1}]^2\} - (E_i \log p_t(i) - \tilde{P}_{t-1})^2 \\
&= \alpha E_i\{[\log p_t(i) - \tilde{P}]^2\} + (1-\alpha)\Sigma_h\omega_h(\log p_t^{*h} - \tilde{P}_{t-1})^2 - (\tilde{P}_t - \tilde{P}_{t-1})^2 \\
&= \alpha V_{t-1} + \frac{1}{1-\alpha}\Sigma_h\omega_h(\pi_t^h)^2 - (\Sigma_h\omega_h\pi_t^h)^2,
\end{aligned}$$

where p_t^{*h} is the optimal price chosen at date t by firms (with planning horizon h) who can choose a new price at that date. Here we use the definition of $\pi_t^h = \frac{1}{1-\alpha}(\log p_t^{*h} - \tilde{P}_{t-1})$, and observe that it follows that $\tilde{P}_t - \tilde{P}_{t-1} = \Sigma_h\omega_h\pi_t^h$.

Thus if we take the discounted value of V_t over all periods $t \geq 0$, it follows that

$$\begin{aligned}\Sigma_{t=0}^{\infty} \beta^t V_t &= \frac{1}{1 - \alpha\beta} \Sigma_{t=0}^{\infty} \beta^t \left[\frac{1}{1 - \alpha} \Sigma_h \omega_h (\pi_t^h)^2 - (\Sigma_h \omega_h \pi_t^h)^2 \right] \\ &= \frac{\alpha}{(1 - \alpha)(1 - \alpha\beta)} \Sigma_{t=0}^{\infty} \beta^t \left[\frac{1}{\alpha} \Sigma_h \omega_h (\pi_t^h)^2 - \frac{1 - \alpha}{\alpha} (\Sigma_h \omega_h \pi_t^h)^2 \right],\end{aligned}$$

and hence, together with equation (D.11), we obtain

$$\Sigma_{t=0}^{\infty} \beta^t U_t = -\Omega \Sigma_{t=0}^{\infty} \beta^t L_t + \text{t.i.p.},$$

where Ω is a positive constant and the period quadratic welfare loss function L_t is given by

$$\begin{aligned}L_t &= \frac{1}{\alpha} \Sigma_h \omega_h (\pi_t^h)^2 - \frac{1 - \alpha}{\alpha} (\Sigma_h \omega_h \pi_t^h)^2 \\ &\quad + [\lambda_{disp} \Sigma_h \omega_h (y_t^h)^2 + (\lambda_{agg} - \lambda_{disp}) (\Sigma_h \omega_h y_t^h)^2],\end{aligned}$$

with the relative weight $\lambda_{disp} = \lambda_{agg} \sigma^{-1} / (\sigma^{-1} + \omega)$ and $\lambda_{agg} = \kappa / \theta$. Here we use the notation $y_t^h = \hat{Y}_t^h - \hat{Y}_t^n$.

Since the aggregate inflation rate and output are given by $\pi_t = \Sigma_h \omega_h \pi_t^h$ and $y_t = \Sigma_h \omega_h y_t^h$, we can rewrite the expression of L_t as

$$L_t = \pi_t^2 + \frac{1}{\alpha} \text{var}(\pi_t^h) + \lambda_{agg} y_t^2 + \lambda_{disp} \text{var}(y_t^h), \quad (\text{D.12})$$

where the dispersion of the values of π_t^h and y_t^h are respectively given by

$$\text{var}(\pi_t^h) = \Sigma_h \omega_h (\pi_t^h)^2 - (\Sigma_h \omega_h \pi_t^h)^2, \quad \text{var}(y_t^h) = \Sigma_h \omega_h (y_t^h)^2 - (\Sigma_h \omega_h y_t^h)^2.$$

Thus both the composition of aggregate output y_t and the composition of aggregate inflation rate π_t matter for welfare — it is not enough to only stabilize y_t and π_t . Note that in the case that all households and firms have a common planning horizon h , $\text{var}(\pi_t^h) = \text{var}(y_t^h) = 0$ at all times, regardless of the policy chosen, and in this case the welfare loss function of the form (D.12) reduces to the simpler loss function (4.4) in section 4.2.

E Optimal Fiscal-Monetary Coordination with an Exponential Distribution of Planning Horizons: Numerical Methods

In this section, we propose a numerical method to compute the solutions for optimal exogenous state-contingent fiscal transfer policy and interest-rate policy, which minimize the accumulated present value of average-utility welfare loss. As in section 4.2 and 4.3, we assume that the shock of financial wedge lasts for T periods and its path is perfectly predictable, i.e., $\hat{\Delta}_t = -\hat{i} + \underline{\Delta}$ for $0 \leq t < T - 1$, where $\underline{\Delta} > 0$ is the excess financial wedge that cannot be offset by a reduction in nominal interest rate, and $\hat{\Delta} = 0$ for all $t \geq T$.

Because all uncertainty about fundamentals is resolved in period $t = 0$, we optimize over policies specified by deterministic paths $\{\hat{i}_\tau, b_{\tau+1}\}$ consistent with the ZLB constraint (2.12) for all $\tau \geq 0$, announced at time $t = 0$ and similarly understood by all decision makers that plan far enough ahead (at any date $t \geq 0$) to model economic conditions at date τ . In other words, the optimization problem for the policymakers (with commitment) is to choose the set of possible sequences $\{\hat{i}_\tau, b_{\tau+1}\}_{\tau=0}^\infty$ so as to minimize

$$E_0 \sum_{t=0}^\infty \beta^t L_t$$

subject to the ZLB constraint (2.12) for all $\tau \geq 0$. Here the average-utility welfare loss function L_t is given by equation (D.12) and the paths $\{y_t^h, \pi_t^h\}$ for all $t \geq 0$ implied by the policy $\{\hat{i}_\tau, b_{\tau+1}\}$ are given by (C.5). In the case of an exponential distribution of planning horizons, the share of agents with planning horizon h is given by $\omega_h = (1 - \rho)\rho^h$ for all $h \geq 0$.

To numerically solve this optimization problem, we assume that after long-enough time $T_{max} > T$, all equilibrium variables will be staying at the steady state under the optimal policy.⁶² Furthermore, to approximate the exponential distribution of planning horizons, we assume that for those agents looking forward beyond h_{max} periods, their planning horizon is h_{max} .⁶³ Thus the optimization problem of the

⁶²In the numerical example we presented in section 4.3, given that the shock lasts for $T = 10$ periods in a quarter model, we choose $T_{max} = 200$ quarters. The numerical results are robust if we take a larger value of T_{max} .

⁶³With an average planning horizon being 8 quarters, we choose $h_{max} = 60$ quarters. In this case,

policymakers reduces to choosing $\{\hat{i}_\tau, b_{\tau+1}\}_{\tau=0}^{T_{max}+h_{max}}$ so as to minimize $\sum_{t=0}^{T_{max}} \beta^t L_t$, subject to the ZLB constraint (2.12) for all $\tau \geq 0$. Note that in this case, the loss function L_t is only determined by $\{y_t^h, \pi_t^h\}$ with $0 \leq h \leq h_{max}$.

Since this is a deterministic system with finite choice variables in the optimization problem, we shall be able to numerically find the optimal policy, i.e., the sequences of $\{\hat{i}_\tau, b_{\tau+1}\}_{\tau=0}^{T_{max}+h_{max}}$ (with $\hat{i}_\tau = b_{\tau+1} = 0$ for all $\tau > T_{max} + h_{max}$). We utilize the optimization package “fmincon” provided by Matlab to solve such a minimization problem with constraints. For robustness check, we have confirmed that with this algorithm, the same results can be obtained as in Appendix C if we assume a common planning horizon instead of an exponential distribution.

the share of agents that look beyond 60 quarters is less than 0.1% of the whole population. The numerical results are robust if we take a larger value of h_{max} .