

NBER WORKING PAPER SERIES

MONETARY POLICY ANALYSIS WHEN PLANNING HORIZONS ARE FINITE

Michael Woodford

Working Paper 24692

<http://www.nber.org/papers/w24692>

NATIONAL BUREAU OF ECONOMIC RESEARCH

1050 Massachusetts Avenue

Cambridge, MA 02138

June 2018

Prepared for the NBER Macroeconomics Annual conference, April 12-13, 2018. I would like to thank Roger Farmer, Xavier Gabaix, Mariana Garcia-Schmidt, Marc Giannoni, Jennifer La'O, Guido Lorenzoni, Bentley MacLeod, Benoit Mojon, and Argia Sbordone for helpful discussions. The views expressed herein are those of the author and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2018 by Michael Woodford. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Monetary Policy Analysis when Planning Horizons are Finite
Michael Woodford
NBER Working Paper No. 24692
June 2018
JEL No. E03,E52

ABSTRACT

It is common to analyze the effects of alternative monetary policy commitments under the assumption of fully model-consistent expectations. This implicitly assumes unrealistic cognitive abilities on the part of economic decision makers. The relevant question, however, is not whether the assumption can be literally correct, but how much it would matter to model decision making in a more realistic way. A model is proposed, based on the architecture of artificial intelligence programs for problems such as chess or go, in which decision makers look ahead only a finite distance into the future, and use a value function learned from experience to evaluate situations that may be reached after a finite sequence of actions by themselves and others. Conditions are discussed under which the predictions of a model with finite-horizon forward planning are similar to those of a rational expectations equilibrium, and under which they are instead quite different. The model is used to re-examine the consequences that should be expected from a central-bank commitment to maintain a fixed nominal interest rate for a substantial period of time. “Neo-Fisherian” predictions are shown to depend on using rational expectations equilibrium analysis under circumstances in which it should be expected to be unreliable.

Michael Woodford
Department of Economics
Columbia University
420 W. 118th Street
New York, NY 10027
and NBER
mw2230@columbia.edu

It has become commonplace — certainly in the scholarly literature, but also increasingly in central banks and other policy institutions — to analyze the predicted effects of possible monetary policies using dynamic stochastic general equilibrium models, in which both households and firms are assumed to make optimal decisions under rational expectations. Since the methodological revolution in macroeconomics initiated by Kydland and Prescott (1982), this has come to mean assuming that economic agents formulate complete state-contingent intertemporal plans over an infinite future. Yet such a postulate is plainly heroic, as the implicit assumptions made about the knowability of all possible future situations, the capacity of people to formulate detailed plans before acting, and the ability of individuals to solve complex optimization problems in real time are well beyond the capabilities even of economists, let alone members of society in general.

Most if not all macroeconomists who use models of this kind probably do so on the assumption that such models represent a useful idealization — that while not literally correct, their predictions are approximately correct, while their logical simplicity makes them convenient to use in thinking through a variety of thought experiments of practical interest. Yet their use in this way requires that one have some basis for judgment about the degree to which, and the circumstances under which, one should expect the predictions of an admittedly idealized model to nonetheless be approximately correct. The issue of the conditions under which an idealized model can approximate a more complex reality deserves analysis, rather than simply being a matter of faith (or badge of professional identity), as it too often is.

Here we propose an approach to macroeconomic analysis that makes less extreme cognitive demands than conventional rational expectations equilibrium analysis, and thus allows us to pose the question of the degree to which the conclusions of the conventional analysis should be at least approximately valid even in a world in which people are only boundedly rational. It allows us to identify circumstances under which the predictions of the conventional analysis can be correct, or at least approximately correct, without people having to have quite such extraordinary cognitive capacities as the rational expectations analysis would seem on its face to require.

It can also address a conceptual problem with rational expectations analysis, which is providing a ground for selection of a particular solution as the relevant prediction of one's model, under circumstances where an infinite-horizon model admits a large number of potential rational expectations equilibria. The boundedly rational solution concept proposed here is necessarily unique, and so, in cases in which it

coincides with a rational expectations equilibrium (or approaches one as the limit on computational complexity is relaxed), it provides a reason for using that particular rational expectations equilibrium as the predicted effect of the policy in question.

At the same time, the proposed approach will not always result in predictions similar to those of any rational expectations equilibria; in such cases, it provides a reason to doubt the practical relevance of conclusions from rational expectations analysis. In particular, we will argue that conclusions about the effects of a central-bank “forward guidance” based on rational expectations analysis are sometimes quite misleading, as they depend on assuming the validity of rational expectations analysis under circumstances under which a more realistic (though still highly sophisticated) model of human decision making would lead to quite different conclusions.

Our proposed approach proceeds from the observation that in the case of complex intertemporal decision problems, people — even experts — are not able to “solve” such problems using the sort of backward induction or dynamic programming approaches that are taught in economics classes. Rather than beginning by considering all possible final situations, valuing them, and then working back from such judgments about the desirability of endpoint to reach a conclusion about the best first action to take in one’s current situation, it posits that people actually start from the specific situation that they are in, and work forward from it to some finite extent, considering alternative situations that can be reached through some finite sequence of possible actions; but that they necessarily truncate this process of “forward planning” before all of the consequences of their earlier actions have been realized.¹

And rather than supposing that people should be able to deductively compute a correct value function for possible interim situations that they might be able to reach — through some algorithm such as value function iteration, which requires that a decision maker begin by specifying the set of possible states for which values must be computed — it recognizes that while people have some ability to learn the values of particular situations by observing their average consequences over some body of actual or simulated experience, this necessarily requires a coarse classification of possible situations in order to make such averaging feasible. It is because of the coarseness of the state space for which a value function can be learned, relative to the more fine-grained information about one’s current situation that can be made

¹Earlier proposed models of boundedly rational economic decision making in this spirit include Jéhiel (1995), MacLeod (2002), and Gabaix *et al.* (2006).

use of in a forward planning exercise, that forward planning is useful, even when only feasible to some finite distance into the future. Our proposed approach makes use of both (finite-depth) forward planning and (coarse) value-function learning, in order to take advantage of the strengths of each while mitigating the most important weaknesses of each.

The paper proceeds as follows. Section 1 introduces the basic approach to modeling boundedly rational intertemporal decision making that we propose. Section 2 then shows how this approach can be applied to monetary policy analysis, in the context of a simple but relatively standard micro-founded New Keynesian model. In the analysis developed in this section, the coarse value function that decision makers use to value potential situations at the horizon at which their forward planning is truncated is taken as given, though motivated as one that would be optimal in a certain kind of relatively simple environment. Section 3 applies the framework developed in section 2 to the specific problem of analyzing the effects of an announcement that a new approach to monetary policy will be followed for a period of time, as in recent experiments with “forward guidance,” and compares the conclusions from our boundedly rational analysis with conventional rational expectations analyses. Section 4 then extends the analytical framework to also consider how value functions are learned from experience, allowing them to eventually change in response to a sufficiently persistent change in either policy or fundamentals. This allows us to consider the validity of the proposition that the Fisher equation should hold in the long run, regardless of how inflationary or deflationary monetary policy may be, and of the “neo-Fisherian” conclusions that are sometimes drawn from this proposition. Section 5 concludes.

1 How Are Complex Intertemporal Decisions Made?

In practice, even in highly structured environments such as the games of chess or go — where clear rules mean that the set of possible actions in any situation can be completely enumerated, and the set of situations that can ever possibly be encountered is also finite, so that in principle all possible strategies can be exhaustively studied — it is not possible even for the most expert players (whether human or AIs) to discern the optimal strategy and simply execute it. (Indeed, tournament play would not be interesting, and the challenge of designing better programs would not continue to en-

gauge computer scientists, were this the case.) This fact reveals something about the limitations of the kinds of computational strategies that economists use to compute optimal decision rules in classroom exercises.

But it is also worth considering how the best players approach these problems in practice — in particular, the approaches used by state-of-the-art artificial intelligence programs, since these are now the best players in the world, and (more to the point) we know how they work. If we wish to assume in economic models that the people who make up the economy should be highly rational and do a good job of choosing strategies that serve their interests — but not that they have magical powers — then it would seem reasonable to assume that they make decisions using methods similar to those used by the most effective artificial intelligence programs.²

Programs such as Deep Mind for chess (Campbell *et al.*, 2002) or AlphaGo for the game of go (Silver *et al.*, 2016) have the following basic structure. Whenever it is the computer’s turn to move, it begins by observing a precise description of the current state of the board. Starting from this state, it considers the states that it is possible to move to through a legal move; the possible situations that can arise as a result of any legal responses by the opponent in any such state; the possible states that can be moved to through a legal move from the situation created by the opponent’s move, and so on — creating a tree structure with the current state of the board as its root.

Once the tree is created, values are assigned to reaching the different possible terminal nodes (that is, the nodes at which the process of “tree search” is truncated). Different hypothetical sequences of moves, extending forward until a terminal node is reached, can then be valued according to the value of the terminal node that they would allow one to reach. This allows the selection of a preferred sequence of moves (a finite-horizon plan — though not a plan for the entire rest of the game). The move that is taken is then the first move in the preferred sequence. However, the finite-horizon plan chosen at one stage in the game need not be continued; instead, the forward planning exercise is repeated each time another move must be selected (looking farther into the future as the game progresses, and hence possibly choosing a new plan that does not begin by continuing the one selected at the time of the previous move).

²The kind of decision-making algorithm proposed here — called “planning-to-habit” by Keramati *et al.* (2016) — has also been used to describe the behavior of non-expert human decision makers in settings where less extensive training has been possible.

Such a tree search procedure would be fully rational if the complete game tree (terminating only at nodes at which the game has ended) were considered. But except in special circumstances (possibly near the end of a game), this is not feasible. Hence a tree of only a finite “depth” must be considered before choosing a current action. The best programs use sophisticated rules to decide when to search further down particular branches of the game tree and when to truncate the search earlier, in order to deploy finite computational resources more efficiently. In the model proposed below, however, we simply assume a uniform depth of search k ; that is, a decision maker is assumed to consider all of the possible states that can be reached through a feasible sequence of actions over the next k periods. Our focus here is on comparing a model with finite-horizon forward planning to one in which the complete (unbounded) future is considered, and on considering how the length of the finite horizon matters.

Another crucial aspect of such a program is the specification of the function that is used to evaluate the different terminal nodes. It is important to note that the answer cannot be that the value assigned to a terminal node should be determined by looking at the states farther down the game tree that can be reached from it; the whole point of having a value function with which to evaluate terminal nodes is to allow the program to avoid having to look farther into the future (and thus have to consider an even larger number of possible outcomes). The value function must be learned in advance (before a particular game is played), on the basis of an extensive database of actual or simulated play, and represents essentially an empirical average of the values observed to follow from reaching particular states.

If sufficient prior experience were available to allow a correct value function (taking into account a precise description of the situation that has been reached) to be learned, then truncation of the forward planning at a finite depth would not result in sub-optimal decisions. Indeed, there would be no need for multi-stage forward planning at all; one could simply consider the positions to which it is possible to move from one’s current position, evaluate them, and choose the best move on this basis. The only reason that forward planning (to the depth that is feasible) is useful is that in practice, a completely accurate value function cannot be learned, even from a large database of experience; there are too many possible states that might in principle need to be evaluated for it to be possible to observe all of the outcomes that might result from each one of them and tabulate the average values of each. Thus in practice, the value function used by such a program must evaluate a situation based on a certain

set of *features*, which provide a coarse description of the situation but do not uniquely identify it.

The degree to which forward planning should be used, before resorting to the use of a value function learned from prior experience to evaluate the situations that may be reached under alternative finite-horizon action plans, reflects a trade-off between the respective strengths and weaknesses of the two approaches. Evaluation of possible situations using the value function is quick and inexpensive (once the value function has once been learned); but it has the disadvantage that (in order for it to be feasible to ever learn the value function) the value function can take into account only a coarse description of each of the possible situations.

Forward planning via tree search can instead take into account very fine-grained information about the particular situation in which one currently finds oneself, because it is only implemented for a particular situation once one is in it; but it has the disadvantage that the process of considering all possible branches of the decision tree going forward rapidly becomes computationally burdensome as the depth of search increases. Finite-horizon forward planning to an appropriate depth makes use of fine-grained information when it is especially relevant and not overly costly (that is, when thinking about the relatively near future), but switches to the use of a more coarse-grained empirical value function to evaluate possible situations when thinking about the farther future.

A final feature of such algorithms deserves mention. If the intertemporal decision problem is not an individual decision problem, but instead one where outcomes for the decision maker depend on the actions of others as well — an opponent, in the case of chess or go, or the other households and firms whose actions determine market conditions, in a macroeconomic model — then the algorithm must include a model of others' behavior, in order to deduce the consequences of choosing a particular sequence of actions. It makes sense to assume that those others will also behave rationally; but it will not be possible to compute their predicted behavior using an algorithm that is as complex as the forward planning algorithm that one uses to make one's own decision.

In particular, if the algorithm used to choose one's own plan of action looks forward to possible situations after k successive moves, it cannot also model the opponent's choice after one's first move by assuming that the opponent will look forward to possible situations after k successive moves, and consider what one should

do oneself after its first reply by simulating the result of looking forward to possible situations after k more moves, and so on. For continuation of such a chain of reasoning would amount to reasoning about possible situations that can be reached after many more than just k successive moves. In order for the complexity of the decision tree that must be considered to be bounded by looking out only to some finite depth, it becomes necessary to assume an even shorter horizon in the forward planning that one simulates on the part of other people whose behavior at a later stage of the tree must be predicted. This idea is made concrete below in the context of a general equilibrium analysis.

Our proposed approach has certain similarities to models of boundedly rational decision making discussed by Branch *et al.* (2012). Branch *et al.* assume that decision makers use econometric models to forecast the future evolution of variables the future values of which matter to their intertemporal decision problem, and compare a variety of assumptions about how those forecasts may be used to make decisions; in particular, they discuss models in which decision makers solve only a finite-horizon problem, and hence only need to forecast over a finite horizon. A crucial difference between their models and the one we propose here is that in those models, the *same* econometric model is used both to forecast conditions during the (near-term) period for which a finite-horizon plan is chosen and to estimate the value of reaching different possible terminal nodes. Instead, we emphasize that the types of reasoning involved in finite-horizon forward planning on the one hand, and in the evaluation of terminal nodes on the other, are quite different, and that the sources of information that are taken into account for the two purposes are accordingly quite different. This has important consequences for our analysis of the effects of central-bank forward guidance — which we assume is taken into account in forward planning (based on the decision maker’s complete information about the current situation), but not in the value function (that necessarily classifies situations using only a limited set of features, with which there must have been extensive prior experience).

2 A New Keynesian DSGE Model with Finite-Horizon Planning

We now illustrate how the proposed approach can be applied to monetary policy analysis, by deriving the equations of a New Keynesian DSGE model similar to the basic model developed in Woodford (2003), but replacing the standard assumption of infinite-horizon optimal planning by a more realistic assumption of finite-horizon planning. We begin by deriving boundedly rational analogs of the two key structural equations of the textbook New Keynesian model — the “New Keynesian IS relation” and the “New Keynesian Phillips curve” — and then discuss the implications of the modified equations for an analysis of the effects of forward guidance regarding monetary policy.

2.1 Household Expenditure with a Finite Planning Horizon

We assume an economy made up of a large number of identical households, each of which represents a “dynasty” of individuals that share a single intertemporal budget constraint, and earn income and spend over an infinite horizon. At any point in time t , household i wishes to maximize its expected utility from then on,

$$\hat{E}_t^i \sum_{\tau=t}^{\infty} \beta^{\tau-t} [u(C_\tau^i; \xi_\tau) - w(H_\tau^i; \xi_\tau)],$$

where C_τ^i is the expenditure of i in period τ on a composite good, H_τ^i is hours worked in period τ , and ξ_τ is a vector of exogenous disturbances, that can include disturbances to the urgency of current expenditure or the disutility of working. As usual, we suppose that for each value of the disturbance vector, $u(\cdot; \xi)$ is an increasing, strictly concave function, and that $w(\cdot; \xi)$ is an increasing, convex function; and that the discount factor satisfies $0 < \beta < 1$. The composite good is a Dixit-Stiglitz aggregate

$$C_\tau^i \equiv \left[\int_0^1 (C_\tau^i(f))^{\frac{\theta-1}{\theta}} df \right]^{\frac{\theta}{\theta-1}} \quad (2.1)$$

of the household’s expenditure $C_\tau^i(f)$ on each of a continuum of differentiated goods indexed by f , where $\theta > 1$. The operator $\hat{E}_t^i[\cdot]$ indicates the expected value under the subjective expectations of household i at time t , which we have yet to specify.

In the present subsection, we shall be concerned purely with the household's planning of its state-contingent expenditure C_τ^i on the composite good, given the expected evolution of the price P_τ of the composite good (a Dixit-Stiglitz index of the prices of the individual goods) and the household's income from working and from its share in the profits of firms. Both the question of how the household allocates its spending across the different individual goods and how its hours of work are determined are left for later. Here we note simply that we assume an organization of the labor market under which each household is required to supply its share of the aggregate labor H_τ demanded by firms; hence the expected evolution of $H_\tau^i = H_\tau$ is independent of household i 's intentions with regard to spending and wealth accumulation. Moreover, each household's total income other than from its financial position (saving or borrowing) will simply equal its share of the total value Y_τ of production of the composite good. The evolution of this income variable is outside the control of an individual household i .

We further simplify the household's problem by supposing that there is a single kind of traded financial claim, a one-period riskless nominal debt contract, promising a nominal interest rate i_τ (that is, one dollar saved in period τ buys a claim to $1 + i_\tau$ dollars in period $\tau + 1$) that is controlled by the central bank. We denote the financial wealth carried into period t by household i by the variable B_t^i , defined as the nominal value of claims maturing in period t deflated by the price index P_{t-1} ; this definition makes B_t^i a real variable that is purely predetermined (dependent only on decisions made at date $t - 1$). The household's financial position evolves in accordance with a flow budget constraint

$$B_{\tau+1}^i = (1 + i_\tau)[B_\tau^i/\Pi_\tau + Y_\tau - C_\tau] \quad (2.2)$$

for each period $\tau \geq t$, where $\Pi_\tau \equiv P_\tau/P_{\tau-1}$ is the gross inflation rate between period $\tau - 1$ and period τ .

The problem considered in this subsection is the household's choice of an intended path of expenditure $\{C_\tau^i\}$, where spending in any period τ may depend on the aggregate state s_τ at that time, together with the associated path for its financial position $\{B_{\tau+1}^i\}$ implied by (2.2). (Here we use the notation s_τ for the complete state vector, including both the exogenous disturbances ξ_τ and any policy decisions that have been announced as of date τ .) These are chosen to maximize the expected discounted sum of utility from expenditure (since the disutility of working can be taken as given, for

present purposes), given an existing financial position B_t^i and the expected evolution of the variables $\{i_\tau, \Pi_\tau, Y_\tau\}$ in periods $\tau \geq t$.

2.1.1 An Optimal Finite-Horizon Plan

Rather than assuming that the household chooses a complete infinite-horizon state-contingent expenditure plan — and that it must accordingly consider the possible paths of the variables $\{i_\tau, \Pi_\tau, Y_\tau\}$ over an infinite future — we shall suppose that the household engages in explicit forward planning for only k periods into the future. This means that in period t (and given the state s_t at that time), the household chooses state-contingent values $\{C_\tau^i(s_\tau)\}$ only for the possible states s_τ that may be reached at dates $t \leq \tau \leq t+k$.

This plan is chosen to maximize the finite-horizon objective

$$\mathbb{E}_t^k \left[\sum_{\tau=t}^{t+k} \beta^{\tau-t} u(C_\tau^i; \xi_\tau) + \beta^{k+1} v(B_{t+k+1}^i; s_{t+k}) \right], \quad (2.3)$$

where the evolution of $B_{\tau+1}^i$ for $t \leq \tau \leq t+k$ under the finite-horizon plan is given by (2.2), and $v(B_{t+k+1}^i; s_{t+k})$ is the *value function* that the household uses to estimate the continuation value of its problem at each of the possible states s_{t+k} at which it truncates the forward planning exercise. Here the operator $\mathbb{E}_t^k[\cdot]$ indicates the expectations at time t of a decision maker that plans k periods ahead regarding the probabilistic evolution of the variables outside her control. The information set of such a decision maker is assumed to include the current exogenous state s_t , and the equilibrium realizations of all endogenous variables at dates $t-1$ or earlier; the conditional probabilities of future exogenous states are assumed to be correctly known, while the values of endogenous variables at date t or later (conditional on the exogenous state that is reached) are computed using the model's structural equations (as discussed further below). These expectations differ from model-consistent or “rational” expectations only because of the truncation of the household's planning horizon, and so are the same for all households with planning horizon k .

The household's finite-horizon plan will satisfy a set of first-order conditions,

$$u_c(C_\tau^i; \xi_\tau) = \beta \mathbb{E}_t^k [((1+i_\tau)/\Pi_{\tau+1}) u_c(C_{\tau+1}^i; \xi_{\tau+1}) | s_\tau] \quad (2.4)$$

for each possible state s_τ (given the state s_t at the time of the planning) at each date

$t \leq \tau \leq t + k - 1$, and

$$u_c(C_{t+k}^i; \xi_{t+k}) = \beta(1 + i_{t+k}) v_B(B_{t+k+1}^i; s_{t+k}) \quad (2.5)$$

for each possible state s_{t+k} at which the forward planning is truncated. (The expectation operator $E_t^k[\cdot | s_\tau]$ refers to the expectations that the decision maker at date t expects to have, conditional on reaching state s_τ .) Conditions (2.4)–(2.5) together with the budget constraints (2.2) determine a state-contingent plan $\{C_\tau^i, B_{\tau+1}^i\}$ for periods $t \leq \tau \leq t + k$.

If we were to assume that the household's expectations $E_t^k[\cdot]$ are fully model-consistent expectations, and that the value function $v(B; s_{t+k})$ corresponds to the true (model-consistent) value of the household's continuation problem for any level of net saving B chosen in state s_{t+k} , then this system of equations would characterize the household's optimal infinite-horizon expenditure plan. We assume instead that the household's plan is only boundedly rational, in two respects.

First, we assume that while in period t , the household chooses planned expenditure for periods t through $t + k$, it does not then implement this plan (beyond the level of spending C_t^i chosen for the current period); and when it reoptimizes in the following period, it will not generally choose to continue with the plan chosen in period t (since in period $t + 1$, it looks forward to period $t + k + 1$, rather than truncating the planning at period $t + k$). This is however neglected by the household when it chooses a spending plan in period t . Indeed, we cannot assume that the household in period t has model-consistent expectations about its spending in the different states that may be reached in period $t + 1$, for this would require that the household use the model structural equations to calculate what should be expected to happen in period $t + k + 1$, rather than truncating the deductive forward planning in period $t + k$. Instead, we suppose that the household calculates as if in period $t + 1$ it will plan forward only $k - 1$ periods into the future, in period $t + 2$ it will plan forward only $k - 2$ periods into the future, and so on.

And while we assume that the household correctly understands the equations of the structural model (including the policy rule announced by the central bank), and uses them in the planning exercise to deduce the values of Π_τ, Y_τ , and i_τ that should be expected in each possible state s_τ for $t \leq \tau \leq t + k$, this does not suffice to imply model-consistent expectations of those variables. According to our model, aggregate expenditure in period $t + j$ (for some $1 \leq j \leq k$) is determined by the

planning decisions of households in period $t + j$ that (if they have $k - period$ planning horizons) look forward to the anticipated model-based determination of variables as far in the future as period $t + j + k$. But in order for the household's planning not to require it to consider what the model equations imply about states farther in the future than period $t + k$, we assume that the household assumes that in period $t + j$, spending and pricing decisions will be made by households and firms with only $(k - j)$ -period planning horizons.

Thus if we let Π_t^k, Y_t^k, i_t^k , be the model-consistent solutions for the endogenous variables in period t (given the state s_t reached at that time) in a model where all decision makers are assumed to have planning horizons of length k , then the expectations used in period t by a household with a k -period planning horizon are assumed to satisfy

$$E_t^k[Z_\tau] = E_t Z_\tau^{t+k-\tau}$$

for any endogenous variable Z_τ in period τ ($t \leq \tau \leq t + k$), where $E_t[\cdot]$ now refers to the model-consistent expectation conditional on being in state s_t . Similarly, we assume that the expectations about the following period that the household expects (in its period- t planning exercise) to hold in any future state s_τ ($t + 1 \leq \tau \leq t + k$) are given by

$$E_t^k[Z_{\tau+1}|s_\tau] = E_\tau Z_{\tau+1}^{t+k-\tau}.$$

Hence the household's Euler equation (2.4) can alternatively be written

$$u_c(C_\tau^j; \xi_\tau) = \beta E_\tau[(1 + i_\tau^j)/\Pi_{\tau+1}^{j-1}] u_c(C_{\tau+1}^{j-1}; \xi_{\tau+1}), \quad (2.6)$$

for any planning horizon $j \geq 1$. We can also use (2.5) to obtain a corresponding FOC for a household with a zero-period planning horizon:

$$u_c(C_\tau^0; \xi_\tau) = \beta(1 + i_\tau) v_B(B_{\tau+1}^0; s_\tau), \quad (2.7)$$

where $B_{\tau+1}^0$ is the wealth carried into period $\tau + 1$ by a household with a zero-period planning horizon in period τ . We now have a system of equations (consisting of (2.6) for each of the periods $t \leq \tau \leq t + k$, and (2.7) for period $\tau = t + k$), now involving only model-consistent conditional expectations, to determine the state-contingent plan chosen in period t by a household with a k -period planning horizon.

The second respect in which we depart from fully model-consistent expectations is that we do not assume that the value function $v(B; s_{t+k})$ used to evaluate possible

situations at the point where forward planning is truncated necessarily corresponds to the model-implied continuation value of the household's discounted expected utility conditional upon reaching state s_{t+k} . As discussed above, we suppose that the value function is learned by averaging past experience, rather than by using the model structural equations to deduce what should happen farther in the future; and it will not be practical to learn the value of wealth conditioning on all details of the complete state vector s_{t+k} .

To simplify the current presentation, we suppose that the value function is not state-contingent at all, though households are assumed to correctly learn the average continuation value $v(B)$ associated with a given level of real wealth B .³ In particular, we assume in our treatment of forward guidance below that the value function $v(B)$ does not take any account of the consequences of any announcement by the central bank of a change in the monetary policy that will be implemented at dates beyond the planning horizon (though this would be part of the complete state vector s_{t+k}). We do assume that the value function reflects one simple kind of state-dependence: households are assumed to recognize that it is their real financial position, rather than their nominal position, that should determine the value of their continuation problem (so that the price level anticipated for period $t + k$ is taken into account).

We defer until section 4 a discussion of how the value function is assumed to be learned from experience. We first examine the equilibrium consequences of finite-horizon planning under the assumption of a given value function $v(B)$, setting aside the question of how the function should eventually shift over time in response to further experience.⁴ For the sake of concreteness, we further suppose that the economy has for a long time been in a stationary equilibrium in which there have been no real disturbances and monetary policy has maintained a constant inflation rate $\bar{\Pi}$; that as a result, output and the nominal interest rate have been constant as well, with values

³It is not an essential feature of our method that the value function be so simple. The key assumption for our analysis below of forward guidance is that we suppose that an announcement of a new monetary policy does not change the value functions that households or firms use in their finite-horizon planning. It is particularly plausible to suppose that the value functions should ignore this aspect of the state, as in the case of a policy that has never been used previously, there will have been no opportunity to learn the consequences of this kind of policy change from experience.

⁴The issue of the endogeneity of the value function can be abstracted from when we are concerned only with the effects of relatively transitory disturbances, including short-lived changes in monetary policy.

\bar{Y} and \bar{i} respectively; and that households and firms have eventually learned value functions that are correct for this stationary environment. (We assume that $\bar{\Pi}^* > \beta$, so that the implied stationary nominal interest rate $\bar{i} = \beta^{-1}\bar{\Pi} - 1$ is positive.)

In such an environment, a household's correct continuation value function is the function $v(B)$ that solves the Bellman equation

$$v(B) = \max_C \{u(C) + \beta v(B')\} \quad \text{s.t. } B' = \beta^{-1}[B + (\bar{Y} - C)\bar{\Pi}],$$

where \bar{Y} is the stationary equilibrium level of output when households and firms optimize using model-consistent expectations (perfect foresight). The solution to this problem is easily seen to be

$$v(B) = (1 - \beta)^{-1} u(\bar{Y} + (1 - \beta)B/\bar{\Pi}). \quad (2.8)$$

This is the value function that we shall assume that households use, until our consideration of learning dynamics in section 4.

2.1.2 Log-Linear Approximation of the Optimal Plan

Suppose now that there are no real disturbances and that the central bank uses its policy instrument to maintain the target inflation rate $\bar{\Pi}$ at all times. If households and firms act on the basis of finite-horizon plans, but use correct value functions, then actions are the same as if they chose their actions on the basis of infinite-horizon optimization, and hence all aspects of the equilibrium will be the same. If all households start with identical financial positions ($B_t^i = 0$, since financial claims are in zero net supply⁵), then in this equilibrium $C_t^i = \bar{C} = \bar{Y}$ for each household at all times.

If instead we allow for real disturbances and/or time variation in monetary policy, equilibrium dynamics with finite-horizon planning will not generally coincide with the predictions of rational expectations equilibrium analysis. However, in the case of small enough departures from the assumptions of the perfect foresight steady state, we can approximately characterize these dynamics through a perturbation of the solution just computed in the case of zero disturbances and the constant policy. A first-order

⁵In the present analysis, we abstract from both government spending and government debt issuance; extension of the model to consider the effects of fiscal policy as well is left for a future study.

perturbation solution for the representative household's finite-horizon spending plan is obtained by linearizing the structural equations (2.2), (2.5) and (2.6) around the stationary solution.

We write the linearized equations in terms of percentage deviations from the stationary values of the variables, using the notation

$$c_t \equiv \log(C_t/\bar{C}), \quad y_t \equiv \log(Y_t/\bar{Y}), \quad b_t^i \equiv B_t^i/(\bar{\Pi}\bar{Y}),$$

$$\pi_t \equiv \log(\Pi_t/\bar{\Pi}), \quad \hat{i}_t \equiv \log(1 + i_t/1 + \bar{i}).$$

Then if we let $\sigma \equiv -u_c(\bar{C})/(u_{cc}(\bar{C})\bar{C}) > 0$ be the intertemporal elasticity of substitution of household expenditure, and parameterize disturbances to the urgency of spending by the quantity g_t such that to first order,

$$\log(u_c(C_t^i; \xi_t)/u_c(\bar{C}; \bar{\xi})) = -\sigma^{-1}(c_t^i - g_t),$$

(2.6) can be linearized to yield

$$c_\tau^j - g_\tau = E_\tau[c_{\tau+1}^{j-1} - g_{\tau+1}] - \sigma[\hat{i}_\tau^j - E_\tau\pi_{\tau+1}^{j-1}]. \quad (2.9)$$

Here the value of c_τ^j for any horizon j and period τ is understood to depend not only on the aggregate state s_τ at that time, but also on the financial position that households with a planning horizon of j periods are assumed to start the period with.

Similarly, (2.8) implies that to a log-linear approximation, the marginal value of wealth used in evaluating potential financial positions when forward planning is terminated is equal to

$$\log(v'(B_{\tau+1}^0)/v'(\bar{B})) = -(1 - \beta)\sigma^{-1}b_{\tau+1}^0.$$

Hence (2.7) can be linearized to yield

$$c_\tau^0 - g_\tau = -\sigma\hat{i}_\tau^0 + (1 - \beta)b_{\tau+1}^0. \quad (2.10)$$

Here the solutions for c_τ^0 , \hat{i}_τ^0 and $b_{\tau+1}^0$ all depend on both the aggregate state s_τ and the financial position of the household at the beginning of period τ .

Finally, the flow budget constraint (2.2) can be linearized to yield

$$b_{\tau+1}^j = \beta^{-1}[b_\tau^{j+1} + y_\tau^j - c_\tau^j] \quad (2.11)$$

for any horizon $j \geq 0$. Conditions (2.9) and (2.11), setting $\tau = t + k - j$ for each $0 \leq j \leq k$, and condition (2.10) for $\tau = t + k$, together with model-consistent solutions for the variables $\{i_{t+k-j}^j, y_{t+k-j}^j, \pi_{t+k-j+1}^{j-1}\}$ for $0 \leq j \leq k$, then provide a system of simultaneous linear equations to solve for a linear approximation to the optimal plan of a household with a k -period planning horizon in period t . The solution for optimal expenditure in period t (the period in which the planning is undertaken) is given by

$$c_t^k = g_t + (1 - \beta)b_t + (1 - \beta) \sum_{j=0}^k \beta^j \mathbb{E}_t[y_{t+j}^{k-j} - g_{t+j}] - \sigma \sum_{j=0}^{k-1} \beta^{j+1} \mathbb{E}_t[\hat{l}_{t+j}^{k-j} - \pi_{t+j+1}^{k-j-1}] - \sigma \beta^{k+1} \mathbb{E}_t \hat{l}_{t+k}^0, \quad (2.12)$$

where b_t is the household's financial position at the beginning of period t . This is the aspect of the plan chosen in period t that is actually implemented.

In order to determine aggregate expenditure in period t , it is also necessary to determine the state-contingent values $\{i_{t+k-j}^j, y_{t+k-j}^j, \pi_{t+k-j+1}^{j-1}\}$ of aggregate state variables used in the planning exercise. Let us assume for now that all households have identical planning horizons,⁶ and also start with identical financial positions. Then they make identical decisions at all times, and thus necessarily begin each period with a financial position $b_t = 0$. It is assumed that each household has the structural knowledge required to deduce this, and correctly assumes that $b_t^i = 0$ for each of the other households as well. Aggregate demand (and aggregate real income) is then given by $y_t^k = c_t^k$. It follows from (2.11) that $b_{t+1}^k = 0$ for each household, and each household is assumed to have the structural knowledge required to deduce that this will be true for all other households. Thus in predicting the spending decisions of other households in period $t + 1$ (which are assumed then to have planning horizons of only $k - 1$ periods), these households are each assumed to enter period $t + 1$ with financial positions $b_{t+1}^i = 0$. The level of spending that each is expected to choose in period $t + 1$ is then given by equation (2.12), except with t replaced by $t + 1$ and k replaced by $k - 1$, but again assuming initial financial wealth of zero.

The same argument applies in each successive period of the household's planning exercise. It follows that the value y_{t+j}^{k-j} assumed in the planning exercise (for any $1 \leq j \leq k$ and any possible state s_{t+j}) is equal to the value c_{t+j}^{k-j} implied by (2.12)

⁶The analysis is generalized to allow for heterogeneous planning horizons in section 2.3 below.

when t replaced by $t + j$, k is replaced by $k - j$, and b_{t+j} is set equal to zero. The value of π_{t+j}^{k-j} assumed is the one implied by the model of firm planning described in the next section, when firms assume this model of aggregate demand determination.

And finally, interest-rate expectations are assumed to be based on knowledge of the central bank's reaction function, which we assume also to be a log-linear relation of the form

$$\hat{i}_t = i_t^* + \phi_{\pi,t}\pi_t + \phi_{y,t}y_t, \quad (2.13)$$

where the coefficients may be time-dependent (though independent of the realizations of the endogenous variables), in order to allow for "forward guidance" experiments, as discussed in section 3. The interest-rate expectations \hat{i}_{t+j}^{k-j} used in a forward planning exercise at time t are thus assumed to satisfy

$$\hat{i}_\tau^j = i_\tau^* + \phi_{\pi,\tau}\pi_\tau^j + \phi_{y,\tau}y_\tau^j, \quad (2.14)$$

for any planning horizon $j \geq 0$ and any date $t \leq \tau \leq t + j$; in the latter expression, the time-dependent coefficients are assumed to be the ones implied by policy announcements as of period t .

It follows that for any $j \geq 1$, the aggregate expenditure of households with a planning horizon of j periods must satisfy the recursive relationship

$$y_t^j - g_t = E_t[y_{t+1}^{j-1} - g_{t+1}] - \sigma(\hat{i}_t^j - E_t\pi_{t+1}^{j-1}), \quad (2.15)$$

while the aggregate expenditure of households with a zero-period planning horizon must satisfy

$$y_t^0 - g_t = -\sigma\hat{i}_t^0. \quad (2.16)$$

These equations can be solved recursively to determine predicted aggregate demand. Equation (2.16) can be solved for y_{t+k}^0 , given a solution for π_{t+k}^0 and the central bank's reaction function (2.14). Using this solution for y_{t+k}^0 , equation (2.15) can then be solved for y_{t+k-1}^1 , given solutions for π_{t+k-1}^1 and π_{t+k}^0 and the central bank's reaction function. Using this solution for y_{t+k-1}^1 , equation (2.15) can then be solved for y_{t+k-2}^2 , and so on. The model's prediction for aggregate expenditure y_t^k in period t is then given by the solution for y_t^k .

2.2 Price-Setting by Firms with a Finite Planning Horizon

We turn now to inflation determination through the price-setting decisions of firms. Each differentiated good f is assumed to be sold by a monopolistically competitive

producer (also indexed by f), that sets the price P_τ^f of good f each period. The objective of each firm is assumed to be maximization of the average value to shareholders of the stream of earnings generated by the firm's pricing policy.

We assume for simplicity that shares in the firms are not traded, and that each household i receives an equal share of the earnings of all firms. Firms are assumed to value incremental earnings in different aggregate states in proportion to the average marginal utility of additional real income to their shareholders in those different states, averaging the marginal utilities of the different households;⁷ that is, in proportion to the quantity

$$\lambda_t \equiv \int u_c(C_t^i; \xi_t) di. \quad (2.17)$$

Hence a firm's objective is assumed to be the maximization of

$$\hat{E}_t^f \sum_{\tau=t}^{\infty} \beta^{\tau-t} \lambda_\tau H(P_\tau^f / P_\tau; Z_\tau), \quad (2.18)$$

where $H(r_\tau^f; Z_\tau)$ represents the real profits of firm f in period τ if its relative price is $r_\tau^f \equiv P_\tau^f / P_\tau$; Z_τ is a vector of real state variables at date τ (specified below) that are outside the control of the individual firm (under the assumption of monopolistic competition) but that matter for the firm's real profits; and the operator $\hat{E}_t^f[\cdot]$ indicates the expectations used by firm f in a planning exercise at date t .

2.2.1 Staggered Price Adjustment

As in the models of Calvo (1983) and Yun (1996), we assume that each period only a fraction $1 - \alpha$ of all goods prices are reconsidered in any period, with the particular prices that are reconsidered a random selection from the set of goods. In a period t in which the price of good f is not reconsidered, we assume that $P_t^f = P_{t-1}^f \cdot \bar{\Pi}$. That is, we assume (as in Yun, 1996) that prices are automatically increased at the target inflation rate when the optimality of this default pricing rule is not considered. This assumption implies that in the absence of aggregate disturbances or changes

⁷In the exposition in the previous section, we have assumed that all households solve an identical problem, with a k -period planning horizon; if so, in equilibrium all households value additional income in the same way, and we can simply refer to the marginal utility of income of the representative household. However, the exposition in this section allows for possible heterogeneity in households' planning horizons, in preparation for the discussion in section 2.3.

in monetary policy, the equilibrium prices of all goods will be the same as in a perfect foresight equilibrium with flexible prices, despite the fact that not all prices are reconsidered each period. Staggered pricing matters for equilibrium dynamics only to the extent that there are disturbances and/or temporary shifts in the central bank's reaction function.

In a period t in which the price of good f is reconsidered, we assume that the firm chooses a new price P_t^f so as to maximize its subjective assessment of the objective (2.18). But as in the previous section, we assume that the firm plans ahead for only k periods, and evaluates possible situations in period $t+k$ using a value function to estimate the value of its continuation problem. Furthermore, when choosing a new price in period t , the firm need only consider the consequences for future states in which it has not yet reconsidered its price more recently than period t . Hence P_t^f is chosen to maximize

$$\hat{E}_t^f \left[\sum_{\tau=t}^{t+k} (\alpha\beta)^{\tau-t} \lambda_\tau H(P_t^f \bar{\Pi}^{\tau-t} / P_\tau; Z_\tau) + (\alpha\beta)^{k+1} \tilde{v}(P_t^f \bar{\Pi}^k / P_{t+k}; s_{t+k}) \right], \quad (2.19)$$

where $\tilde{v}(r_{t+k}^f; s_{t+k})$ represents the firm's estimate of the value of discounted real profits from period $t+k+1$ onward, conditional on reaching state s_{t+k} in period $t+k$, not reconsidering its price in any of the periods between $t+1$ and $t+k+1$, and having a relative price of r_{t+k}^f in period $t+k$.

In a stationary perfect foresight equilibrium in which the central bank maintains a constant inflation rate $\bar{\Pi}$, the real variables Z_t have constant values \bar{Z} that satisfy

$$H_r(1; \bar{Z}) = 0, \quad (2.20)$$

and the price of each good satisfies $P_t^j = P_t$ at all times. (The allocation of resources in this stationary equilibrium is the same as in a stationary equilibrium with perfectly flexible prices.) As in the previous section, we assume that the value function used by firms is optimal for this stationary equilibrium (assumed to have prevailed for some time prior to the period in which we seek to analyze the effects of a change in monetary policy). But again we simplify by assuming that the value function is not state-dependent; thus we assume that

$$\tilde{v}(r) = (1 - \alpha\beta)^{-1} \bar{\lambda} H(r; \bar{Z}), \quad (2.21)$$

where $\bar{\lambda} \equiv u_C(\bar{C}; \bar{\xi})$ is the constant value of λ_t in the perfect-foresight steady state.

It follows that the first-order condition for maximization of (2.19) is given by

$$\hat{E}_t^f \left[\sum_{\tau=t}^{t+k} (\alpha\beta)^{\tau-t} \lambda_\tau H_r(P_t^f \bar{\Pi}^{\tau-t}/P_\tau; Z_\tau) \frac{P_t \bar{\Pi}^{\tau-t}}{P_\tau} + \frac{(\alpha\beta)^{k+1}}{1-\alpha\beta} \bar{\lambda} H_r(P_t^f \bar{\Pi}^k/P_{t+k}; \bar{Z}) \frac{P_t \bar{\Pi}^k}{P_{t+k}} \right] = 0.$$

Log-linearizing this condition around the values that hold in the perfect-foresight steady state yields

$$\hat{E}_t^f \left\{ \sum_{\tau=t}^{t+k} (\alpha\beta)^{\tau-t} [p_t^f - \sum_{s=t}^{\tau} \pi_s - m_\tau] + \frac{(\alpha\beta)^{k+1}}{1-\alpha\beta} [p_t^f - \sum_{s=t}^{t+k} \pi_s] \right\} = 0, \quad (2.22)$$

using the notation

$$p_t^f \equiv \log \left(\frac{P_t^f}{P_{t-1} \bar{\Pi}} \right), \quad m_t \equiv -H_{rr}(1; \bar{Z})^{-1} H_r(1; Z_t).$$

(Note that m_t is a function of real variables outside the control of firm f , and is of only first order in the deviations of these variables from their steady-state values. To a linear approximation, it measures the percentage deviation of the average real marginal cost of supplying output from its steady-state level.)

As in the previous section, we assume that the firm's expectations $\hat{E}_t^f[\cdot]$ are deduced from the model structural equations, but assuming that endogenous variables determined in period τ are determined by the decisions of households and firms with planning horizons of $t+k-\tau$ periods, for any $t \leq \tau \leq t+k$. (The endogenous variables referred to here now include the variables Z_t , along with i_t , Π_t , and Y_t .) It then follows from (2.22) that any firm f that reconsiders its price in period t , and plans ahead for k periods, will choose a new relative price $p_t^f = p_t^{*k}$, where the optimized relative price is given by

$$p_t^{*k} = E_t \sum_{\tau=t}^{t+k} (\alpha\beta)^{\tau-t} [\pi_\tau^{t+k-\tau} + (1-\alpha\beta) m_\tau^{t+k-\tau}]. \quad (2.23)$$

As usual, the Dixit-Stiglitz price index is defined as

$$P_t \equiv \left[\int (P_t^f)^{1-\theta} df \right]^{\frac{1}{1-\theta}}.$$

This implies that when we log-linearize around the stationary equilibrium with constant inflation rate $\bar{\Pi}$, we obtain (to a first-order approximation)

$$\pi_t^k = (1-\alpha) p_t^{*k}.$$

Equation (2.23) then implies that

$$\pi_t^k = (1 - \alpha) \mathbb{E}_t \sum_{\tau=t}^{t+k} (\alpha\beta)^{\tau-t} [\pi_\tau^{t+k-\tau} + (1 - \alpha\beta)m_\tau^{t+k-\tau}].$$

A similar equation holds if we replace k by any horizon $j \geq 0$. We then see that the $\{\pi_t^j\}$ for different finite horizons satisfy a simple recursion of the form

$$\pi_t^j = \frac{(1 - \alpha)(1 - \alpha\beta)}{\alpha} m_t^j + \beta \mathbb{E}_t \pi_{t+1}^{j-1} \quad (2.24)$$

for any $j \geq 1$, with the special form

$$\pi_t^0 = \frac{(1 - \alpha)(1 - \alpha\beta)}{\alpha} m_t^0 \quad (2.25)$$

when the planning horizon is of length zero. Equations (2.24)–(2.25) can be solved forward to yield

$$\pi_t^k = \frac{(1 - \alpha)(1 - \alpha\beta)}{\alpha} \sum_{\tau=t}^{t+k} \beta^{\tau-t} \mathbb{E}_t m_\tau^{t+k-\tau}$$

as the solution for equilibrium inflation given expectations regarding the determination of m_τ . If all firms have a k -period planning horizon, then the actual inflation rate will be given by $\pi_t = \pi_t^k$.

2.2.2 Determinants of Real Marginal Cost

It remains to discuss how the variable m_t is jointly determined along with the other variables in our model. This requires that we consider further the form of the firm's profit function. Real profits each period are equal to real sales revenues minus the firm's real wage bill. As usual, Dixit-Stiglitz preferences (2.1) imply that the demand for good f is equal to $Y_t(r_t^f)^{-\theta}$, where Y_t is aggregate demand for the composite good; hence the real revenues of firm f are equal to $Y_t(r_t^f)^{1-\theta}$.

We assume that each firm produces its differentiated good using labor as the only variable input, with a production function $Y_t(f) = A_t H_t(f)^{1/\phi}$, where $H_t(f)$ is the labor hired by firm f , A_t is an exogenous common productivity factor, and $\phi \geq 1$ indicates the degree of diminishing returns to scale. Real labor costs are therefore equal to $W_t(Y_t(f)/A_t)^\phi$, where W_t is the real wage in period t .

We assume that wages are determined in the following way. The household suppliers of labor are represented in wage negotiations by representatives that each bundle

the labor of a representative sample of the different household types in the economy, and offer to supply a certain number of total hours by members of their group at a given wage; when a given number of hours are agreed upon for a given wage, each household in the group must supply that number of hours, and receives the same wage. There are a large number of such representatives (each bargaining on behalf of an identical group of households), so that none has any market power. The representative chooses a number of hours H_t that the group will offer to work so as to maximize the average utility of the households in the group; this results in a first-order condition for optimal labor supply of the form

$$v_H(H_t; \xi_t) = \lambda_t W_t,$$

where λ_t is again defined by (2.17). Note that this is the relationship between wages and hours that would hold in a representative-household model of the kind assumed in Woodford (2003); here, as in Woodford (2013), we assume a labor market organization that implies a similar relationship even when the subjective marginal utilities of income of households may differ because of their boundedly-rational expectations.

Real profits are then given by

$$H(r_t^j; Z_t) \equiv Y_t (r_t^f)^{1-\theta} - \lambda_t^{-1} w((Y_t/A_t)^\phi (r_t^f)^{-\theta\phi}; \xi_t),$$

where

$$w(H_t; \xi_t) \equiv v_H(H_t; \xi_t) \cdot H_t.$$

(The vector Z_t is now seen to have as its elements the endogenous variables Y_t and λ_t , and the exogenous variables A_t and the elements of ξ_t that affect the disutility of labor.)

Differentiating this profit function to obtain H_r , we can use (2.20) together with the requirement that $\bar{\lambda} = u_C(\bar{Y}; \bar{\xi})$ to determine the steady-state values \bar{Y} and $\bar{\lambda}$. We can also differentiate the function H_r , obtaining

$$m_t = [1 + (\phi - 1)\theta]^{-1} \cdot \left[(1 + \nu)\phi(y_t - a_t) - y_t - \hat{\lambda}_t + \xi_t^m \right],$$

where

$$\begin{aligned} \nu &\equiv \bar{H} w_{HH} / w_H > 0, & a_t &\equiv \log(A_t / \bar{A}), & \hat{\lambda}_t &\equiv \log(\lambda_t / \bar{\lambda}), \\ \xi_t^m &\equiv \log(w_H(\bar{H}; \xi_t) / w_H(\bar{H}; \bar{\xi})). \end{aligned}$$

If all households have a planning horizon j , we have $\lambda_t^j = -\sigma^{-1}(y_t^j - g_t)$, from which it follows that

$$m_t^j = \xi (y_t^j - y_t^*),$$

where y_t^* is a linear combination of the exogenous terms g_t, a_t and ξ_t^m identifying (to a linear approximation) the percentage change in the flexible-price equilibrium level of output,⁸ and

$$\xi \equiv \frac{(\phi - 1) + \nu\phi}{1 + (\phi - 1)\theta} > 0.$$

It follows that the recursive system (2.24)–(2.25) can alternatively be written as

$$\pi_t^j = \kappa (y_t^j - y_t^*) + \beta \mathbf{E}_t \pi_{t+1}^{j-1} \quad (2.26)$$

for any $j \geq 1$, and

$$\pi_t^0 = \kappa (y_t^0 - y_t^*), \quad (2.27)$$

where

$$\kappa \equiv \frac{(1 - \alpha)(1 - \alpha\beta)}{\alpha} \cdot \xi > 0.$$

The system of equations consisting of (2.14), (2.15)–(2.16), and (2.26)–(2.27) can then be solved recursively to obtain solutions for the evolution of the variables $\{\hat{v}_t^j, \pi_t^j, y_t^j\}$ for any planning horizon $j \geq 0$. Finally, if the actual planning horizon of firms is always k periods, then the model's prediction for equilibrium inflation will be $\pi_t = \pi_t^k$.

2.3 Heterogeneous Planning Horizons

Thus far, we have assumed that the planning horizon of all households and firms is of some finite length, k periods. However, there is no reason to suppose that this must be uniform across decision makers, or even that it must be the same each time that a decision is made, for a given decision maker. We can easily extend the model to allow for heterogeneity in the length of planning horizons (while still treating this as an exogenous parameter for the decision maker, rather than another decision).

We have discussed above how to compute y_t^j , the overall spending that would be undertaken by households with a planning horizon of j periods, and π_t^j , the average rate of price increase in excess of the target inflation rate on the part of firms with

⁸Note that in a flexible-price equilibrium, one would have $H_r(1; Z_t) = 0$ at all times, meaning that (in a linear approximation) $m_t = 0$ at all times, which requires that $y_t = y_t^*$ to first order.

a planning horizon of j periods, for arbitrary values of j . These calculations did not depend on whether any households or firms actually have planning horizons of length j . (In the discussion above, we have used the notation k for the length of actual planning horizons; but in order to determine the behavior of households and firms with planning horizons of length k , it has been necessary to consider the counterfactual behavior of households and firms with planning horizons $0 \leq j \leq k$.) In the calculation of y_t^j or π_t^j , the fact that actual planning horizons are of length k was never used. (Households with planning horizons of length j are assumed to expect that their current-period income will result from the spending decisions of other households that also have planning horizons of length j , that their income in the following period will result from the decisions of other households whose planning decisions are all of length $j - 1$, and so on — even if actual planning horizons are all of length k . This was necessary in order to allow decisions to be made without having to think about what anyone should think about conditions more than k periods in the future.)

Hence the equations stated above for the determination of $\{y_t^j, \pi_t^j, \hat{v}_t^j\}$ for different horizons j continue to apply, even if we assume that a variety of different lengths of planning horizons are actually used. Suppose that each period, a fraction ω_j of households have planning horizons of length j (for $j \geq 0$), and similarly that a fraction $\tilde{\omega}_j$ of firms have planning horizons of length j , where the sequences $\{\omega_j, \tilde{\omega}_j\}$ satisfy $\sum_j \omega_j = \sum_j \tilde{\omega}_j = 1$. The particular households with planning horizons of a given length may or may not remain the same from period to period; what matters is that we assume stable population fractions. Then aggregate real expenditure (and hence real income) in period t will be given by

$$y_t = \sum_j \omega_j y_t^j, \quad (2.28)$$

and the overall rate of inflation will be given by

$$\pi_t = \sum_j \pi_t^j, \quad (2.29)$$

to a log-linear approximation in each case.

The assumption of heterogeneous planning horizons does introduce a complication worth mentioning, relative to the discussion above. Households with planning horizon k will assume that they will receive income in the current period equal to y_t^k (because they assume the planning horizons of others to be the same as their own),

but they will actually receive income y_t given by (2.28); and so, while if they start with a financial position $b_t^i = 0$, they choose a level of expenditure that they expect to imply $b_{t+1}^i = 0$ as well, this will generally not be the case, as their income will be different than they had expected. However, it continues to be the case that the *aggregate* financial position of households will be zero, as planned. And because the approximate expenditure rules derived above are linear in b_t^i , the heterogeneous evolution of financial positions across the population of households is of no consequence for the evolution of aggregate demand, aggregate income, or wage setting (and hence the evolution of real marginal costs).

It should also be noted that the existence of equilibrium requires that the sums in (2.28)–(2.29) converge. If there exists a finite upper bound on the planning horizons of both households and firms, then this is not an issue, since each of the individual quantities y_t^j and π_t^j is necessarily well-defined and finite. In the case that there is no upper bound on planning horizons, convergence of the sums depends on how y_t^j and π_t^j behave for large j . If these quantities converge as j is made large, there is again no problem; but it is possible to define policies (such as the case of a permanent interest-rate peg, discussed below) under which y_t^j and π_t^j grow explosively as j is increased. In such a case, the existence of an equilibrium depends on ω_j and $\tilde{\omega}_j$ going to zero rapidly enough for large j . Since forward planning over extremely long horizons is unlikely to be within anyone’s cognitive capacities, we regard this as a plausible assumption, even in the case of a policy commitment of such an extreme kind.

Allowing for potentially unbounded length of planning horizons is convenient for at least one reason, which is that it allows us to analyze the consequences of finite planning horizons in a case in which the state space of our model is no more complex than under the rational expectations analysis. In the special case in which $\omega_j = \tilde{\omega}_j = (1 - \rho)\rho^j$ for some $0 < \rho < 1$, we can average equations (2.15)–(2.16) to obtain

$$y_t - g_t = \rho \mathbf{E}_t[y_{t+1} - g_{t+1}] - \sigma (\hat{i}_t - \rho \mathbf{E}_t \pi_{t+1}). \quad (2.30)$$

Similarly, we can average equations (2.26)–(2.27) to obtain

$$\pi_t = \kappa (y_t - y_t^*) + \beta \rho \mathbf{E}_t \pi_{t+1}. \quad (2.31)$$

These two conditions, together with a specification of the central-bank reaction function (2.13), provide a complete system of three equations per period to solve for the evolution of the three endogenous variables $\{y_t, \pi_t, \hat{i}_t\}$.

Of course, while the conditions (2.30)–(2.31) are *necessary* conditions for given paths to be consistent with finite-horizon optimization by households and firms, they are not *sufficient* conditions. It is still necessary to validate any candidate solution by computing the paths $\{y_t^j, \pi_t^j, \hat{v}_t^j\}$ for each planning horizon j , and verifying that the sums (2.28)–(2.29) converge. (This will often be the case, but need not be, as discussed in the next section.)

The conditions (2.30)–(2.31) are quite similar to the equilibrium conditions of the standard New Keynesian model with rational expectations (discussed further in section 3), differing only in the appearance here of the factor $\rho < 1$ that decreases the influence of the expectational terms in both equations. Exactly this kind of modification of the standard model is also proposed by Gabaix (2017), albeit on somewhat different grounds.⁹ Our proposal is not equivalent to that of Gabaix, however, because of the requirement that the sums (2.28)–(2.29) must converge; this leads to importantly different conclusions about the consequences of an interest-rate peg, for example.¹⁰

3 Comparison with Rational Expectations Equilibrium Analysis

Under the assumption of infinite-horizon optimization with rational (i.e., model-consistent) expectations, the variables $\{y_t, \pi_t, \hat{v}_t\}$ must instead evolve in accordance with the equations

$$y_t - g_t = E_t[y_{t+1} - g_{t+1}] - \sigma(\hat{v}_t - E_t\pi_{t+1}), \quad (3.1)$$

$$\pi_t = \kappa(y_t - y_t^*) + \beta E_t\pi_{t+1} \quad (3.2)$$

for all t , along with the central-bank reaction function (2.13). Solutions to this system of equations under different assumptions about policy have been extensively discussed (see, e.g., Woodford, 2003).

⁹The interpretation proposed by Gabaix is one in which people solve infinite-horizon decision problems, but under distorted beliefs about the laws of motion of variables that need to be forecasted, that are biased so as not to differ too much from a simpler “default” model.

¹⁰The discussion by Gabaix of how the “default” model should endogenously respond to experience is also different from the model of learning proposed in section 4 below, and this also matters for our conclusions about the long-run effects of a permanent policy change.

If the limiting values

$$y_t^\infty \equiv \lim_{k \rightarrow \infty} y_t^k, \quad \pi_t^\infty \equiv \lim_{k \rightarrow \infty} \pi_t^k, \quad \hat{i}_t^\infty \equiv \lim_{k \rightarrow \infty} \hat{i}_t^k \quad (3.3)$$

are well-defined, then it follows from (2.15) and (2.26) that these limiting processes must satisfy both of equations (3.1)–(3.2) each period; that is, they must describe a rational expectations equilibrium. It might seem from this that standard analyses of the rational expectations equilibria consistent with a given policy commitment are therefore equivalent to the predictions of our model with finite-horizon planning, in the case that planning horizons are assumed to be long.

This is not quite true, however, for two reasons. First of all, in general the equation system consisting of (2.13) together with (3.1)–(3.2) admits of a large multiplicity of solutions, but not all of these solutions can correspond to the large- k limit of the predictions of the finite-horizon model, for the predictions of the finite horizon model are always *uniquely determined* for any horizon k , and if the sequence converges the limit must be unique as well. Thus in cases where the limiting values are well-defined, these limiting values provide an interpretation of how the behavior described by a rational expectations equilibrium can arise; but they also provide a *selection criterion* that identifies a single solution among the large set of possible rational-expectations solutions as the one that one should expect to see. This clarifies an important issue for monetary policy analysis.

And second, it is not always the case that the sequence of finite-horizon decisions converges as k is made large. In such a case, *none* of the rational-expectations solutions are similar to the behavior predicted by our model, regardless of what one might think is a realistic range of values for k .

To see why, it is useful to write our equilibrium relations in vector form. If we use (2.13) to substitute for \hat{i}_t in (3.1), we can write the equation system defining a rational expectations equilibrium as a two-dimensional system,

$$x_t = M_t \cdot E_t x_{t+1} + N_t \cdot u_t, \quad (3.4)$$

using the vector notation¹¹

$$x_t \equiv \begin{bmatrix} y_t - g_t \\ \pi_t \end{bmatrix}, \quad u_t \equiv \begin{bmatrix} i_t^* + \phi_y g_t \\ g_t - y_t^* \end{bmatrix},$$

and where M_t and N_t are 2×2 matrices of coefficients that depend on the coefficients of the monetary policy reaction function (but are time-invariant if the reaction function is time-invariant).

Using the same notation, the conditions that define an equilibrium with finite-horizon planning can be written in the form

$$x_t^j = M_t \cdot E_t x_{t+1}^{j-1} + N_t \cdot u_t, \quad (3.5)$$

for any $j \geq 1$. For the case $j = 0$, we instead have simply

$$x_t^0 = N_t \cdot u_t. \quad (3.6)$$

3.1 Announcement of a Change in Monetary Policy

We can illustrate the use of this apparatus by considering our model's predictions about a "forward guidance" experiment of the following sort. Suppose that it is announced at date $t = 0$ that from period zero until some horizon $t = T$, monetary policy will follow a rule of the form (2.13) with some constant coefficients (i^*, ϕ_π, ϕ_y) ; we further suppose that $i^* \neq 0$, so that the "new" policy rule is not consistent with the steady state with inflation rate $\bar{\Pi}$ (that we suppose has prevailed prior to the policy experiment). However, it is also understood that from $t = T$ onward, policy will revert to the central bank's "normal" rule, which involves $i_\tau^* = 0$ for all $\tau \geq T$. We wish to consider the effects of announcing a new policy that will be maintained for a specified period of time.¹² Note that we will have constant matrices $M_t = M, N_t = N$ for all $t < T$. We assume for simplicity that there are no real disturbances.

For any $\tau \geq T$, we observe that since $u_\tau = 0$, (3.6) implies that $x_\tau^0 = 0$. Then, proceeding recursively, we can use (3.5) to show that $x_\tau^j = 0$ for any $j \geq 0$. Hence

¹¹We define the first element of x_t to be $y_t - g_t$, rather than simply y_t or the "output gap" $y_t - y_t^*$, for convenience in writing equation (4.30) below. Note that if we can solve the system of equilibrium conditions for the path of $y_t - g_t$, then we also have the solution for the path of y_t or $y_t - y_t^*$, simply by adding the appropriate exogenous variable.

¹²See García-Schmidt and Woodford (2015) for discussion of the connection between such a thought experiment and recent debates about the effects of forward guidance.

regardless of the forecast horizon, for all periods $\tau \geq T$ the outcome will be the steady state with inflation equal to $\bar{\Pi}$; note that this is *one* of the possible rational expectations equilibria for the period after time T .

In fact, there is a unique rational expectations equilibrium with $x_\tau = 0$ for $\tau \geq T$, namely the solution in which

$$x_t^{RE} = (I + M + M^2 + \dots + M^{T-t-1})Nu^* = (I - M^{T-t})(I - M)^{-1}Nu^* \quad (3.7)$$

for all $t < T$. Here the final expression is valid only in the case that $I - M$ is non-singular, which is true as long as

$$\phi_\pi + \frac{1 - \beta}{\kappa}\phi_y \neq 1;$$

but the first equality holds more generally. In these expressions, u^* is the constant value of u_t for all $t < T$, and M, N are the constant values of the matrices M_t, N_t for all $t < T$.

We can then solve the system (3.5)–(3.6) recursively for periods $t < T$, to show that $x_t^j = x_t^{RE}$ for any $j \geq T - t - 1$, while instead $x_t^j = x_{T-j-1}^{RE}$ for any $j < T - t - 1$. (The former result is also true in the case of exogenous disturbances, as long as they are sufficiently transitory; it requires only that $E_t u_\tau = 0$ for all $T \geq t$. The latter result is more special, as it relies on our assumption that $u_\tau = u^*$, a constant vector, for all $\tau < T$.) It follows that in the case of any temporary policy change (i.e., case in which T is finite), the limits (3.3) exist, and are given by $x_t^\infty = x_t^{RE}$ for all t .

Thus in the case of a policy experiment of this kind, finite-horizon planning leads to the same predictions as a rational expectations analysis, *if* people’s planning horizons are long enough ($j \geq T - t - 1$ for everyone), *and* one uses the right selection criterion to choose from among the large set of possible rational expectations solutions. If the policy change is relatively transitory (or at least, people are given no reason to expect anything different from the “normal” reaction function, except over a relatively near future), the length of people’s planning horizons need not even be very long. (And of course, it is only necessary that *most* people’s planning horizons be long enough, in order for the predicted outcome to be *approximately* the same as the rational expectations prediction.)

Matters are more complex in the case of a policy change that is expected to last for a long time. Consider the case in which (contrary to the assumption above) the

change in policy is *permanent*. If the new policy satisfies the “Taylor principle,” i.e., the coefficients of the reaction function satisfy

$$\phi_\pi + \frac{1 - \beta}{\kappa} \phi_y \neq 1; \tag{3.8}$$

then both eigenvalues of M are inside the unit circle.¹³ In this case, (3.7) implies that

$$\lim_{T \rightarrow \infty} x_t^{RE} = x_{ss}^{RE} \equiv (I - M)^{-1} Nu^*.$$

This identifies one of the possible rational expectations equilibria consistent with such a policy commitment: one in which the economy moves immediately to the new stationary equilibrium consistent with the new policy.

It then follows as well that

$$\lim_{k \rightarrow \infty} x_t^k = x_{ss}^{RE}$$

as well, for all $t \geq 0$ (all dates after the announcement of the permanent change in policy). Thus we again find that the limits (3.3) are well-defined; hence we again justify selection of a particular rational expectations equilibrium as the one that approximates what will happen if people have only finite planning horizons, both those horizons are sufficiently long. And it is again the case that, as long as we are confident that most people’s planning horizons are not too short, we should expect an outcome that is approximately the same as in a (suitably chosen) rational expectations equilibrium. (How long horizons must be in order to be “not too short” depends on the largest eigenvalue of M , which depends on the strength of the policy feedback coefficients.)

If instead the inequality in (3.8) is reversed (as will be the case if neither response coefficient is very large), the results are quite different. Because in this case M has an eigenvalue greater than 1, the rational expectations solution x_t^{RE} does not have a well-defined limit as T is made large (it grows explosively). This is a troubling feature of the rational expectations analysis, even when applied to the case of a policy change of long but finite duration; it implies that the predicted outcome is very different depending on the exact value of T , which means that changes in expected policy that change only what is expected about policy *very far in the future* can have a substantial

¹³See Woodford (2003) or García-Schmidt and Woodford (2015) for demonstrations of this, and further discussion.

effect on immediate outcomes — an intuitively unappealing conclusion. Moreover, the large eigenvalue also implies that x_t^k does not have a well-defined limit as k is made large. Hence the outcome with finite-horizon planning need not be similar to the predictions of (any) rational expectations equilibrium, even if one supposes that the planning horizons of most households are quite long.

3.2 The Case of an Interest-Rate Peg

As a case of particular interest in which (3.8) is not satisfied, suppose that $\phi_\pi = \phi_y = 0$. This is the case in which the central bank promises to fix the short-term nominal interest rate at some level i^* up until date T , as in the case of a central bank that announces that its policy rate will remain at its effective lower bound for a stated period of time. Suppose furthermore that $i^* < 0$, meaning not that the nominal interest rate is negative, but that it is lower than its level in the steady state with constant inflation at the target rate $\bar{\Pi}$.

In this case, $Nu^* \ll 0$ and $M \gg 0$, so that (3.7) implies that both elements of x_t^{RE} are positive and monotonically increasing as T is increased. Moreover, because M has an eigenvalue that is greater than 1, both elements of x_t^{RE} are predicted to grow without bound for large enough T . This implies that a commitment to keep the interest rate at a low level should be a stimulative policy, increasing both output and inflation. Moreover, even if a real disturbance would (in the absence of a countervailing change in monetary policy) significantly lower output and inflation, and even if the shock is so severe that a contemporaneous response of monetary policy alone cannot offset it because of the constraint imposed by the effective lower bound on nominal interest rates, it should be possible to fully offset the contractionary effects of the shock by committing to keep the nominal interest rate at the lower bound for a sufficiently long time. This is because, under the rational expectations analysis, the effects of a commitment to keep the nominal interest rate low can be unboundedly large as long as T is long enough.

This result, however — that forward guidance should not only be effective but should have effects that can be unboundedly large (and that grow explosively with the length of the commitment) — has met with some skepticism, so that the results predicted by standard models under rational expectations have been termed a “forward guidance puzzle” (Del Negro *et al.*, 2015). An analysis under the assumption

of finite-horizon planning also predicts that forward guidance should be effective, up to a point; as long as $T \leq t + k + 1$, the model predicts that $x_t^k = x_t^{RE}$, so that both elements of x_t^k will be increased by increasing T . But once $T = t + k + 1$, further increases in T are predicted to have no further effect on x_t^k , which will be given by

$$x_t^k = (I - M^{k+1})(I - M)^{-1} Nu^*, \quad (3.9)$$

for all $T \geq t + k + 1$. The predicted effects of forward guidance are bounded, no matter how long the commitment might be.¹⁴ This is a more empirically plausible result; it also avoids the uncomfortable prediction of the rational expectations analysis, that changes in policy commitments far in the future (leaving expected policy over the next several decades unchanged) should have any material effect on macroeconomic outcomes now.

The conclusions from a rational expectations analysis are even more paradoxical in the case of a thought experiment in which the central bank commits to peg the nominal interest rate forever. Cochrane (2017) suggests that standard New Keynesian models imply that such a policy should have perfectly well-behaved effects, on the ground that there are well-behaved (non-explosive, stationary) rational expectations equilibria consistent with an expectation that such a policy rule will be followed forever. However, all of these rational expectations equilibria have the property that at least eventually (for large enough t) the average inflation rate should be *lower*, the lower the nominal interest rate that the central bank commits to maintain. And many of them involve lower inflation immediately, and not merely in the long run; indeed, if one uses a “minimum-state-variable criterion” to select the rational expectations equilibrium that is expected to occur, a permanent interest-rate peg with $i^* < 0$ (and no current or expected future real disturbances) should lead to an immediate jump to the new stationary equilibrium with a constant inflation rate and the interest rate i^* ¹⁵ — which, because of the Fisher equation, will be an inflation rate that is lower the lower the pegged interest rate.

¹⁴The discussion here assumes a finite planning horizon k for all decision makers, but it suffices that there be some finite upper bound on the length of planning horizons. If instead there is no finite upper bound, but the distribution of planning horizons $\{\omega_j\}$ is the same for both households and firms, then the predicted effects of forward guidance will still be bounded as long as $\sum_j \omega_j M^j$ is still a convergent sum; i.e., as long as ω_j approaches zero sufficiently rapidly for large j .

¹⁵See García-Schmidt and Woodford (2015) for further discussion of this argument.

Such observations suggest that a commitment to a permanently low nominal interest rate should be expected to be a deflationary policy, rather than an inflationary one; and one might then wonder (on the principle that changes in expected policy far in the future should make little difference in the present) if a commitment to a low nominal interest rate for a long though finite time should not be deflationary as well. The conclusions from an analysis that assumes that planning horizons are finite are, however, quite different.

As indicated above, if people’s planning horizons are of some finite length k , then a commitment to peg the nominal interest rate at a lower level than the one consistent with the steady state with inflation rate $\bar{\Pi}$ necessarily results in higher output and inflation. These stimulative effects are predicted to be larger, the longer the length T of the commitment; but once $T \geq t + k + 1$, there are no further effects in period t of lengthening the commitment. The predicted effects are also given by (3.9) in the case of a commitment to a *permanent* interest-rate peg. These effects are positive, though bounded; they are quite unlike the effects that should be observed under *any* of the rational expectations equilibria consistent with such a commitment, at least as regards the levels of output and inflation predicted as t increases.

Thus an analysis based on finite planning horizons provides no support for the “neo-Fisherian” proposal that the way for an economy experiencing chronic low inflation to raise its inflation rate is by committing to peg the nominal interest rate at a *higher* level. However, the prediction just derived for the case of a commitment to a permanent interest-rate peg is not entirely satisfactory. It implies that the inflation rate and nominal interest rate should fail to conform to the Fisher equation, even in the long run, and even though (in the thought experiment just presented) both are forever *constant* at levels incompatible with the Fisher equation.

How can this be? In a situation where output and inflation are given by (3.9) each period (or by an average of this quantity for different values of k , in accordance with the distribution of planning horizons in the population), and the interest rate is pegged at i^* , households are modeled as choosing a constant level of expenditure each period despite facing a constant real rate of interest that differs from their rate of time preference. The reason that this is possible, despite their intertemporal planning, is that they are modeled as using a value function $v(B)$ — to evaluate potential levels of financial wealth at times where they truncate their forward planning — that would be sensible in a stationary equilibrium with an inflation rate of $\bar{\Pi}$, but that does *not*

represent a correct evaluation of the value to them of a given level of financial wealth, even on average, in the environment in which they find themselves after the change in policy.

And while we have justified our assumption of a particular value function by assuming that prior to the policy experiment, households have had considerable experience with a regime in which inflation was kept close to $\bar{\pi}$ and the economy was stable enough for them to come to have correct expectations (including a correct estimate of the value of financial wealth) in that environment, it is not plausible to suppose that they would continue to use this value function if the new policy regime is maintained forever — especially if macroeconomic conditions under the new regime are simple and predictable, as implied by the above thought experiment. A more satisfactory analysis of the effects of long-lasting policy changes (or for that matter, of long-lasting changes in other fundamentals, such as the effects of a permanent change in productivity) requires that we consider how the estimated value functions of households and firms should be shaped by further experience.

4 Learning the Value Functions

In the analysis thus far, we have treated the value functions $v(B)$ and $\tilde{v}(r)$ as fixed, and equal to the correct continuation value functions in a stationary environment with no real disturbances and a constant inflation rate $\bar{\pi}$. If we are interested in analyzing the effects of relatively transitory departures from such an environment — the effects of transitory real disturbances and/or transitory changes in monetary policy — then it may suffice to assume that these value functions continue to be used in the face of such disturbances. But if we wish to analyze the effects of more persistent changes — as in the discussion above of the consequences of a permanent interest-rate peg — then the assumption that the value functions should remain forever equal to these ones is unappealing.

The value functions are intended to represent values that decision makers have learned by averaging their past experience with different states, and a sufficient amount of experience with an environment that persistently differs from the stationary equilibrium with inflation rate $\bar{\pi}$ should eventually cause the value functions to change. In particular, if a new stationary equilibrium is eventually established, it makes sense to suppose that (at least in the long run) the value functions should be

optimal for that new stationary equilibrium, and not for some previous stationary equilibrium far in the past.

4.1 Updating Beliefs

We now illustrate how adaptive learning of the value functions can be incorporated into our analysis. Again we consider first the problem of a household. We cannot suppose that $v(B)$ is a simple average of the household's utility levels on previous occasions when real wealth was equal to B . First, there is the problem that many possible values of B (that need to be considered in the planning exercise) will never have been previously experienced. And more importantly, there is the problem that the value function is intended to assess the value of the household's (infinite-horizon) continuation problem — the expected discounted utility flow over an unbounded sequence of subsequent periods — and the actual value of the household's discounted utility over an unbounded period of time is never observed. But both of these problems can be solved by assuming that what the household averages is not its actual discounted utility following a period in which it has a particular financial position, but rather an estimate of its discounted utility that it computes as part of the finite-horizon forward planning exercise.

Suppose that household i enters period t with a financial position B_t^i , learns the current state s_t , and engages in forward planning using its current estimate $v_t(B)$ of its value function. (The time subscript indicates that we no longer assume that the same value function is used at all times.) Through this exercise (described in section 2.1), it chooses a state-contingent expenditure plan for periods t through $t+k$ to maximize the estimated value of the objective (2.3). As part of this calculation, it obtains an estimated value for its continuation utility from period t onward, the maximized value of (2.3). Furthermore, the household can perform this same calculation for any hypothetical value of B . In this way, the household obtains an estimated value function $v_t^{est}(B)$ for any value of B . (Note that this calculation is performed only for the particular state s_t in which the household finds itself, and not for all of the possible states that it might ever be in — so that the calculation remains much less expensive than a computation of the true value function.)

We may then suppose that the household revises its estimate of the value function (for use in future periods' forward planning exercises) based on a comparison of

its new estimate $v_t^{est}(B)$ with the assumption $v_t(B)$ used in its forward planning.¹⁶ Specifically, let us suppose that

$$v_{t+1}(B) = \gamma v_t^{est}(B) + (1 - \gamma)v_t(B), \quad (4.1)$$

where the “gain parameter” $0 < \gamma < 1$ indicates the rate at which discrepancies between the assumed value function and the new estimate are corrected by adjusting the assumed value function.

Similarly, suppose that a firm f that reconsiders its price in period t engages in forward planning using its current estimate $\tilde{v}_t(r)$ of its value function. Through this exercise (described in section 2.2), it chooses a new price to maximize the estimated value of the objective (2.19). As part of this exercise, it must compute an estimate of what the value of (2.19) would be for any choice of P_t^f ; let this estimate be denoted $\tilde{v}_t^{est}(P_t^f/(P_{t-1}\bar{\Pi}))$. This then implies an estimate for the continuation value function used in forward planning, for any value of r . We may then suppose that the firm revises its estimate of its value function using an error-correction rule of the form

$$\tilde{v}_{t+1}(r) = \tilde{\gamma}\tilde{v}_t^{est}(r) + (1 - \tilde{\gamma})\tilde{v}_t(r), \quad (4.2)$$

where the gain parameter $\tilde{\gamma}$ of firms need not be the same as that of households.

In the case that there are no real disturbances and monetary policy maintains a constant inflation rate $\bar{\Pi}$, assumption of a value function $v^*(B)$ defined by (2.8) on the part of households, and of a value function $\tilde{v}^*(r)$ defined by (2.21) on the part of firms, will result in estimated value functions $v^{est}(B)$ and $\tilde{v}^{est}(r)$ that are also equal to $v^*(B)$ and $\tilde{v}^*(r)$ respectively. Hence the value functions (v^*, \tilde{v}^*) constitute a fixed point of the dynamics defined by (4.1)–(4.2), in such a situation. We wish now to consider a local approximation to the dynamics implied by (4.1)–(4.2), through a perturbation of this solution.

¹⁶This model of learning is related to the model of “value function learning” proposed by Evans and McGough (2015). Like us, they assume that an estimated value function v_t^{est} is computed each period as part of a finite-horizon forward planning calculation using the currently assumed value function. The value function v_t used in the forward planning is then estimated econometrically, using the sequence of past calculated values $\{v_{t-1}^{est}, v_{t-2}^{est}, \dots\}$ as data. Our specification (4.1) can be viewed as a constant-gain recursive estimation procedure for such a problem.

4.2 Log-Linearization of the Learning Dynamics

We first consider a local approximation to (4.1). The structural equations defining the household's optimal finite-horizon plan involve the derivative $v'(B)$ of the value function; hence a log-linear approximation to those equations, of the kind used above to approximate the optimal plan, will involve a log-linear approximation to $v'(B)$. We parameterize this as

$$\log(v'_t(B)/v^{*'}(0)) = -\sigma^{-1} [\nu_t + \chi_t \cdot b].$$

Using this approximation, we can as in section 2.1.2 compute a log-linear approximation to the household's optimal finite-horizon plan in period t , as a function of the assumed coefficients (ν_t, χ_t) . This solution gives approximate values for variables such as c_t^i that are linear functions of b_t^i .

If we let $C_t^i(B)$ be the optimal expenditure plan of the household under the counter-factual assumption $B_t^i = B$, then the derivative of the estimated value function will equal

$$v_t^{est'}(B) = \hat{E}_t^i[u_C(C_t^i(B); \xi_t)/\Pi_t].$$

Hence to a log-linear approximation,

$$\log(v_t^{est'}(B)/v^{*'}(0)) = -\sigma^{-1}(c_t^k(b) - g_t) - \pi_t^k,$$

where k is the length of the planning horizon of the household (and the planning horizon assumed for the firms that revise their prices in period t). Our log-linear approximation to the optimal household plan, $c_t^k(b) = c_t^k(0) + (c_t^k)' \cdot b$, allows us to express the right-hand side of this equation as a linear function of b . Approximating the left-hand side as $-\sigma^{-1} [\nu_t^{est} + \chi_t^{est} \cdot b]$, and equating coefficients, we obtain

$$\nu_t^{est} = y_t^k - g_t + \sigma \pi_t^k, \tag{4.3}$$

$$\chi_t^{est} = (c_t^k)'. \tag{4.4}$$

A log-linear approximation to (4.1) can be written as

$$[\nu_{t+1} + \chi_{t+1} b] = \gamma [\nu_t^{est} + \chi_t^{est} b] + (1 - \gamma) [\nu_t + \chi_t b]. \tag{4.5}$$

Equating coefficients on the two sides of (4.5), we obtain separate updating equations for ν_t and χ_t .

The implied learning dynamics for χ_t turn out to be independent of the pricing behavior of firms. If the household's planning horizon is k periods, the right-hand side of (4.4) is equal to $g_k(\chi_t)$, where

$$g_k(\chi) \equiv \frac{\chi}{\beta^{k+1} + \left(\frac{1-\beta^{k+1}}{1-\beta}\right)\chi},$$

and (4.5) then implies that

$$\chi_{t+1} = \gamma g_k(\chi_t) + (1 - \gamma)\chi_t.$$

This is an autonomous nonlinear difference equation for the evolution of χ_t . One observes furthermore that for any $\chi > 0$, $g_k(\chi)$ is greater than, less than, or equal to χ if and only if χ is less than, greater than, or equal to $1 - \beta$. Hence the difference equation implies monotonic convergence of χ_t to the fixed point $1 - \beta$, from any initial condition $\chi_0 > 0$; this is true for any value of the gain parameter, and is unaffected by exogenous shocks or shifts in monetary policy.

Since there is necessarily eventual convergence of this parameter, we assume in our analysis that convergence has already occurred, and let $\chi_t = 1 - \beta$ at all times. With this simplification, our analysis of learning dynamics reduces to an analysis of the adjustment of the coefficient ν_t . Equation (4.5) implies that

$$\nu_{t+1} = \gamma \nu_t^{est} + (1 - \gamma)\nu_t, \quad (4.6)$$

where ν_t^{est} is given by (4.3).

We can compute a similar local approximation to (4.2). The FOC characterizing the firm's optimal price adjustment depends on the derivative $\tilde{v}'(r)$ of the firm's value function, and log-linearization of this condition requires a log-linear approximation to $\tilde{v}'(r)$. Suppose that we parameterize this as

$$\tilde{v}'_t(r) = -\frac{\bar{\lambda}}{1 - \alpha\beta} H_{rr}(1; \bar{Z}) \cdot [\tilde{\nu}_t - \tilde{\chi}_t \cdot \log r].$$

The log-linearized FOC can then be solved for a linear approximation to the solution for p_t^f , as a function of the coefficients $\tilde{\nu}_t$ and $\tilde{\chi}_t$ used in the approximate value function.

The firm's estimated value function $\tilde{v}_t^{est}(P_t^f / (P_{t-1}\bar{\Pi}))$ is simply the estimated value of the objective (2.19). The derivative $\tilde{v}_t^{est'}(r)$ is obtained by differentiating this

expression. Linearizing this, as in the derivation of (2.22), we obtain

$$\tilde{v}_t^{est}(r) = -\bar{\lambda}H_{rr} \cdot \left(\frac{1 + (\alpha\beta)^{k+1}(\tilde{\chi}_t - 1)}{1 - \alpha\beta} \right) \cdot [p_t^{*k} - \log r],$$

where p_t^{*k} is the optimal log relative price (now dependent on $\tilde{\nu}_t$ and $\tilde{\chi}_t$), the value of $\log r$ that maximizes $\tilde{v}_t^{est}(r)$.

Then if we write $\tilde{v}_t^{est}(r)$ in log-linear form,

$$\tilde{v}_t^{est}(r) = -\frac{\bar{\lambda}}{1 - \alpha\beta} H_{rr}(1; \bar{Z}) \cdot [\tilde{\nu}_t^{est} - \tilde{\chi}_t^{est} \cdot \log r],$$

and equate coefficients, we obtain

$$\tilde{\chi}_t^{est} = \tilde{g}(\chi_t) \equiv [1 + (\alpha\beta)^{k+1}(\tilde{\chi}_t - 1)], \quad (4.7)$$

$$\tilde{\nu}_t^{est} = g(\chi_t) \cdot p_t^{*k}. \quad (4.8)$$

Equating coefficients in a log-linear approximation to (4.2), we obtain updating equations for the coefficients of the form

$$\tilde{\nu}_{t+1} = \tilde{\gamma} \tilde{\nu}_t^{est} + (1 - \tilde{\gamma}) \tilde{\nu}_t, \quad (4.9)$$

$$\tilde{\chi}_{t+1} = \tilde{\gamma} \tilde{\chi}_t^{est} + (1 - \tilde{\gamma}) \tilde{\chi}_t. \quad (4.10)$$

Substituting (4.7) into (4.10), we see that the evolution of $\tilde{\chi}_t$ is determined by an autonomous linear difference equation,

$$\tilde{\chi}_{t+1} = \tilde{\gamma} \tilde{g}(\tilde{\chi}_t) + (1 - \tilde{\gamma}) \tilde{\chi}_t.$$

Moreover, we observe that $\tilde{g}(\tilde{\chi})$ is greater than, less than, or equal to $\tilde{\chi}$ if and only if $\tilde{\chi}$ is less than, greater than, or equal to 1. Hence the updating equation implies monotonic convergence of $\tilde{\chi}_t$ to the fixed point of 1, starting from any initial estimate $\tilde{\chi}_0$, and regardless of the paths of exogenous disturbances or of monetary policy.

We shall accordingly assume in our analysis that convergence has already occurred, and that $\tilde{\chi}_t = 1$ at all times. In this case, (4.8) reduces to

$$\tilde{\nu}_t^{est} = p_t^{*k} = (1 - \alpha)^{-1} \pi_t^k, \quad (4.11)$$

and the dynamic evolution of $\tilde{\nu}_t$ is then given by (4.9) with this substitution. The log-linearized learning dynamics are then described by the system of equations consisting of (4.6) and (4.9) for the evolution of ν_t and $\tilde{\nu}_t$ respectively, where the right-hand sides of both equations can be expressed as linear functions of the current values of the coefficients $(\nu_t, \tilde{\nu}_t)$.

4.3 Equilibrium Dynamics with Learning

In order to describe the complete dynamics of both actions and beliefs, we must consider how the endogenous variables y_t^k and π_t^k (that appear in the expressions (4.3) and (4.11) for the coefficients of the estimated value functions) are affected by variations in the coefficients $(\nu_t, \tilde{\nu}_t)$. This requires us to review the derivations of our log-linear approximations to the optimal decision rules, now allowing a more general specification of the value functions.

Let the predicted equilibrium evolution of each of the endogenous variables $(y_t^j, \pi_t^j, \hat{v}_t^j)$ be expressed as a sum of two components,

$$y_t^j = \tilde{y}_t^j + \bar{y}_t^j, \quad \pi_t^j = \tilde{\pi}_t^j + \bar{\pi}_t^j, \quad \hat{v}_t^j = \tilde{v}_t^j + \bar{v}_t^j,$$

where in each case the tilde component means the predicted value for the variable under the assumption that $\nu_t = \tilde{\nu}_t = 0$ in all periods, but taking account of exogenous shocks and policy changes, while the bar component represents the discrepancy from this prediction as a result of variation in ν_t and $\tilde{\nu}_t$. The evolution of the variables $\{\tilde{y}_t^j, \tilde{\pi}_t^j, \tilde{v}_t^j\}$ for all horizons $j \geq 0$ then continues to be described by the equations derived in section 2. It remains only to compute the values of the bar terms, i.e., the effects on the endogenous variables of perturbation of the value functions.

The log-linear approximations to the FOCs for the household's problem remain as stated in section 2.1.2, except that log-linearization of (2.7) now yields

$$c_\tau^0 - g_\tau = -\sigma \hat{v}_\tau^0 + (1 - \beta)b_{\tau+1}^0 + \nu_t \tag{4.12}$$

as a generalization of (2.10). Here the variables all refer to state-contingent values at date τ that are contemplated by the household in its planning exercise at date t ; in that exercise, the household assumes that all households will use the value function parameterized by ν_t in evaluating terminal states, even in decisions that it imagines them making in periods $\tau > t$.

Because the household's calculations incorporate the requirement that $y_\tau^j = c_\tau^j$, the log-linearized household Euler equations (2.9) imply that

$$\bar{y}_t^j = \bar{y}_t^{j-1} - \sigma[\bar{v}_t^j - \bar{\pi}_t^{j-1}] \tag{4.13}$$

for each $j \geq 1$. Note that the effects of ν_t and $\tilde{\nu}_t$ on the household's period- t calculation of y_τ^{j-1} or π_τ^{j-1} for dates $\tau > t$ are identical to the effects of those belief shifts

on the value of y_t^{j-1} and π_t^{j-1} ; this allows us to make reference purely to period- t variables in (4.13).

The household's log-linearized flow budget constraint also continues to be given by (2.11). Because the household understands that $y_\tau^j = c_\tau^j$, the flow budget constraint implies that in the household's optimal plan, $b_{\tau+1}^j = 0$ each period. Hence (4.12) requires that

$$\bar{y}_t^0 = -\sigma \bar{v}_t^0 + \nu_t. \quad (4.14)$$

The system consisting of (4.13)–(4.14) can be solved recursively to obtain \bar{y}_t^j for any $j \geq 0$.

We turn next to optimal price-setting by firms. The FOC for the pricing decision of a firm that reconsiders its price in period t , and has a planning horizon of k periods, can be log-linearized as above to yield the solution

$$p_t^{*k} = E_t^k \sum_{\tau=t}^{t+k} (\alpha\beta)^{\tau-t} [\pi_\tau + (1 - \alpha\beta)m_\tau] + (\alpha\beta)^{k+1} \tilde{v}_t, \quad (4.15)$$

generalizing (2.23). Now that we now use the operator $E_t^k[\cdot]$ rather than $E_t[\cdot]$, because the predictions about decisions made in periods $\tau > t$ used in the firm's forward planning in period t are no longer model-consistent; this is because the firm assumes in period t that value functions parameterized by ν_t and \tilde{v}_t will also be used in periods $\tau > t$, while genuinely model-consistent expectations would take account of the predictable evolution of beliefs.

The solution (4.15) implies as before that for any $j \geq 1$,

$$p_t^{*j} = E_t^k [\pi_t + (1 - \alpha\beta)m_t + \alpha\beta p_{t+1}^{*j-1}].$$

From this it follows that

$$\bar{\pi}_t^j = \kappa \bar{y}_t^j + \beta \bar{\pi}_t^{j-1} \quad (4.16)$$

for all $j \geq 1$. (Again we use the fact that the effects of ν_t and \tilde{v}_t on the household's period- t calculation of y_τ^{j-1} or π_τ^{j-1} for dates $\tau > t$ are identical to the effects of those belief shifts on the value of y_t^{j-1} and π_t^{j-1} .) The solution (4.15) also implies that

$$\bar{\pi}_t^0 = \kappa \bar{y}_t^0 + (1 - \alpha)\beta \tilde{v}_t. \quad (4.17)$$

The system consisting of (4.16)–(4.17) can be solved recursively to obtain $\bar{\pi}_t^j$ for any $j \geq 0$.

Finally, the central-bank reaction function (2.13) implies that

$$\tilde{\nu}_t^j = \phi_{\pi,t} \bar{\pi}_t^j + \phi_{y,t} \bar{y}_t^j \quad (4.18)$$

for any $j \geq 0$. Equations (4.13)–(4.14), (4.16)–(4.17), and (4.18) form a complete system that can be solved for the values of $\{\bar{y}_t^j, \bar{\pi}_t^j, \tilde{\nu}_t^j\}$ for all $j \geq 0$, as linear functions of ν_t and $\tilde{\nu}_t$. Note that these are all static relationships, as they relate purely to the way that given perturbations of the value functions influence the calculations of households and firms in the forward planning that takes place in a single period t .

Once we have solved for both the tilde variables and the bar variables, we have obtained complete solutions for the variables $\{y_t^j, \pi_t^j, \hat{y}_t^j\}$ at any point in time t , as linear functions of the belief coefficients ν_t and $\tilde{\nu}_t$. Updating of the belief coefficients then requires that we calculate the implied values of the estimates ν_t^{est} and $\tilde{\nu}_t^{est}$.

We recall that ν_t^{est} is given by (4.3), if all households have planning horizons of length k . We can, however, allow for heterogeneity in the length of planning horizons; in this case, there will also be heterogeneity in the updating of value functions. However, in the linear equations (4.14) and (4.17), it is only the population averages of the belief coefficients ν_t and $\tilde{\nu}_t$ that matter for the determination of aggregate variables such as y_t and π_t , and we shall assume from here on that the variables ν_t and $\tilde{\nu}_t$ refer to these averages. The linear updating equations (4.6) and (4.9) continue to hold when we interpret ν_t and $\tilde{\nu}_t$ as population averages, as long as ν_t^{est} and $\tilde{\nu}_t^{est}$ are also now understood to refer to population averages.

In this case, (4.3) must be replaced by the more general form

$$\nu_t^{est} = \sum_j \omega_j [y_t^j - g_t + \sigma \pi_t^j], \quad (4.19)$$

where ω_j is the fraction of households each period with planning horizons of length j . Similarly, (4.11) must be replaced by the more general form

$$\tilde{\nu}_t^{est} = (1 - \alpha)^{-1} \sum_j \tilde{\omega}_j \pi_t^j, \quad (4.20)$$

where $\tilde{\omega}_j$ is the fraction of firms each period with planning horizons of length j .

The complete system of equations to describe the evolution of output, inflation and interest rates, taking into account learning dynamics, then consists of the following sets of equations: (i) equations (2.14), (2.15)–(2.16), and (2.26)–(2.27) comprise a

forward-looking system of equations that can however be solved recursively to obtain solutions for the evolution of the variables $\{\tilde{y}_t^j, \tilde{\pi}_t^j, \tilde{v}_t^j\}$ as functions of the exogenous disturbances and changes in monetary policy; (ii) equations (4.13)–(4.14), (4.16)–(4.17), and (4.18) comprise a static system of linear equations to solve for the values $\{\bar{y}_t^j, \bar{\pi}_t^j, \bar{v}_t^j\}$ as linear functions of ν_t and $\tilde{\nu}_t$; (iii) equations (4.19)–(4.20) allow the estimates ν_t^{est} and $\tilde{\nu}_t^{est}$ to be computed from the solution for the evolution of the variables $\{y_t^j, \pi_t^j\}$ along with the exogenous disturbances; and (iv) equations (4.6) and (4.9) describe the evolution of the belief variables ν_t and $\tilde{\nu}_t$ given these estimates.

4.4 A Useful Special Case

This is in general a relatively complex and high-dimensional system of linear equations, though the causal structure is relatively simple, and (at least if there is a finite upper bound on the planning horizons of both households and firms) there is necessarily a unique solution that can be computed without inverting any large matrices. Further insight into the kind of dynamics implied by this system can be obtained by considering again the special case in which $\omega_j = \tilde{\omega}_j = (1 - \rho)\rho^j$ for some $0 < \rho < 1$. In this case, we can decompose each of our aggregate variables into two components,

$$y_t = \tilde{y}_t + \bar{y}_t, \quad \pi_t = \tilde{\pi}_t + \bar{\pi}_t, \quad \hat{v}_t = \tilde{v}_t + \bar{v}_t,$$

where we define $\tilde{y}_t \equiv (1 - \rho) \sum_j \rho^j \tilde{y}_t^j$, and similarly for the other variables.

The paths of the variables $\{\tilde{y}_t, \tilde{\pi}_t, \tilde{v}_t\}$ then must satisfy equations (2.13) and (2.30)–(2.31). This means that the actual dynamics $\{y_t, \pi_t, v_t\}$ must satisfy

$$y_t - g_t - \bar{y}_t = \rho \mathbf{E}_t[y_{t+1} - g_{t+1} - \bar{y}_{t+1}] - \sigma [(\hat{v}_t - \bar{v}_t) - \rho \mathbf{E}_t(\pi_{t+1} - \bar{\pi}_{t+1})], \quad (4.21)$$

$$\pi_t - \bar{\pi}_t = \kappa (y_t - y_t^* - \bar{y}_t) + \beta \rho \mathbf{E}_t[\pi_{t+1} - \bar{\pi}_{t+1}], \quad (4.22)$$

$$\hat{v}_t - \bar{v}_t = \hat{v}_t^* + \phi_{\pi,t}(\pi_t - \bar{\pi}_t) + \phi_{y,t}(y_t - \bar{y}_t). \quad (4.23)$$

This provides a purely forward-looking system of equations to solve for the deviations of the variables (y_t, π_t, \hat{v}_t) from their “trend” components $(\bar{y}_t, \bar{\pi}_t, \bar{v}_t)$. In the case that the coefficients (ϕ_π, ϕ_y) do not vary over time, it is furthermore a linear system with constant coefficients. Using (4.23) to substitute for $\hat{v}_t - \bar{v}_t$ in (4.21), we can write this as a two-dimensional system,

$$[x_t - \bar{x}_t] = \rho M \cdot \mathbf{E}_t[x_{t+1} - \bar{x}_{t+1}] + N \cdot u_t, \quad (4.24)$$

using the same matrix-vector notation as in (3.4), but now also defining the vector

$$\bar{x}_t \equiv \begin{bmatrix} \bar{y}_t \\ \bar{\pi}_t \end{bmatrix}.$$

Similarly, we can average equations (4.13)–(4.14) over the different horizons j to obtain

$$\bar{y}_t = \rho \bar{y}_t - \sigma [\bar{\nu}_t - \rho \bar{\pi}_t] + (1 - \rho) \nu_t,$$

which can be written more simply as

$$\bar{y}_t = -\frac{\sigma}{1 - \rho} [\bar{\nu}_t - \rho \bar{\pi}_t] + \nu_t. \quad (4.25)$$

And we can average equations (4.16)–(4.17) over the different horizons j to obtain

$$\bar{\pi}_t = \kappa \bar{y}_t + \beta \rho \bar{\pi}_t + (1 - \rho)(1 - \alpha) \beta \tilde{\nu}_t,$$

which can be written more simply as

$$\bar{\pi}_t = \frac{\kappa}{1 - \beta \rho} \bar{y}_t + \frac{(1 - \rho)(1 - \alpha) \beta}{1 - \beta \rho} \tilde{\nu}_t. \quad (4.26)$$

And finally, (4.18) can be averaged to obtain

$$\bar{\nu}_t = \phi_{\pi,t} \bar{\pi}_t + \phi_{y,t} \bar{y}_t. \quad (4.27)$$

This is a system of three simultaneous equations that can be solved for $(\bar{y}_t, \bar{\pi}_t, \bar{\nu}_t)$ as linear functions of ν_t and $\tilde{\nu}_t$.

In particular, we obtain a solution of the form

$$\bar{x}_t = \Xi \begin{bmatrix} \nu_t \\ \tilde{\nu}_t \end{bmatrix} \quad (4.28)$$

for the elements of \bar{x}_t , where (as long as $\phi_\pi, \phi_y \geq 0$) Ξ is an invertible 2×2 matrix. We can then write an evolution equation for the trend components \bar{x}_t through a linear transformation of the laws of motion (4.6) and (4.9) for ν_t and $\tilde{\nu}_t$.

We note that the system (4.19)–(4.20) can be written in the form

$$\begin{bmatrix} \nu_t^{est} \\ \tilde{\nu}_t^{est} \end{bmatrix} = \Phi x_t, \quad (4.29)$$

where Φ is a 2×2 matrix of coefficients. Then the evolution equation for the trend components can be written as

$$\bar{x}_{t+1} = \Lambda \bar{x}_t + Q x_t, \quad (4.30)$$

where

$$\Lambda \equiv \Xi[I - \Gamma]\Xi^{-1}, \quad Q \equiv \Xi\Gamma\Phi,$$

and Γ is the 2×2 diagonal matrix with diagonal elements $(\gamma, \tilde{\gamma})$. Note that the eigenvalues of Λ are $1 - \gamma$ and $1 - \tilde{\gamma}$; thus both eigenvalues are inside the unit circle, and if the variables x_t remain constant over time, the trend variables \bar{x}_t necessarily converge to constant values as well, though this convergence may be slow (if γ and $\tilde{\gamma}$ are small). The dynamics of x_t and \bar{x}_t are then completely determined by (4.24) and (4.30), given a specification of the exogenous disturbance processes, monetary policy, and an initial condition for \bar{x}_0 .

We thus obtain a “hybrid” New Keynesian model, in which deviations of output and inflation from their “trend” values are determined in a purely forward-looking way (though the system (4.24) is somewhat less forward-looking than in the standard model, if ρ is significantly less than 1), but in which there are persistent fluctuations in the “trend” values (quite persistent, if the revision of estimated value functions is slow), determined in a purely backward-looking way. The model thus produces persistent dynamics of both output and inflation, without any need for hypotheses of habit-persistence in preferences, costs of adjusting the rate of investment spending, or automatic indexation of prices to past inflation, of the kind often assumed in econometric New Keynesian models. Like the models of Milani (2007) and Slobodyan and Wouters (2012), the model proposed here generates persistence as a result of learning from past experience. However, unlike those models, the model proposed here does not make expectations purely backward-looking, so that forward guidance (and other special, circumstantial sources of information) is not implied to be irrelevant.

4.5 Long-Run Equilibrium, the Fisher Equation, and the Neo-Fisherian Fallacy

We return now to consideration of the validity of the proposition that the Fisher equation should hold in a long-run equilibrium. In our model that has been augmented to allow adaptive learning of the value functions, this proposition is correct.

Consider a situation in which the central bank's reaction function is constant over time, but not necessarily consistent with the inflation rate $\bar{\Pi}$ assumed in the stationary equilibrium around which we have log-linearized our model equations; and suppose that all exogenous states are also constant over time ($\xi_t = \bar{\xi}$ for all t). Given this, let us consider whether it is possible to have a stationary equilibrium in which the endogenous variables (y_t, π_t, \hat{i}_t) are all constant over time.

If so, the values of ν_t^{est} and $\tilde{\nu}_t^{est}$ will also be constant over time, from which it follows that ν_t and $\tilde{\nu}_t$ will necessarily converge and eventually be constant as well. This in turn means that the “trend” components $(\bar{y}_t, \bar{\pi}_t, \bar{i}_t)$ must eventually take constant values as well. Hence we consider the possibility of stationary solutions in which each of these variables takes a value that is independent of time. (For simplicity, we treat here only the case of exponentially distributed planning horizons just discussed, though a version of the Fisher equation holds in more general cases as well.)

In such a solution, (4.21) requires that

$$(1 - \rho)(y - \bar{y}) = -\sigma[(\hat{i} - \bar{i}) - \rho(\pi - \bar{\pi})]. \quad (4.31)$$

In addition, (4.25) requires that

$$\bar{y} = -\frac{\sigma}{1 - \rho} [\bar{i} - \rho\bar{\pi}] + \nu.$$

Moreover, (4.6) requires that $\nu = \nu^{est}$, which using (4.19) can be seen to imply that

$$\nu = y + \sigma\pi.$$

Using these latter two equations to substitute for \bar{y} and ν , we find that (4.31) requires that

$$\hat{i} = \pi.$$

That is, deviations of the constant inflation rate from the value $\bar{\Pi}$ of the steady state around which we have log-linearized must be associated with deviations of the nominal interest rate of exactly the same size. The Fisher equation must hold in any long-run stationary equilibrium, once we take account of the endogenous adjustment of the beliefs that are reflected in the value function of households.

Does this mean, then, that a commitment to maintain a constant nominal interest rate forever must, at least eventually, bring about a level of inflation consistent with the Fisher equation — so that commitment to maintaining a low nominal interest rate

must eventually be disinflationary or even deflationary, while commitment to keep the nominal interest rate must eventually result in correspondingly high inflation, in accordance with “neo-Fisherian” reasoning? No; for while we have shown that *if* such a policy were to lead, at least in the long run, to a stationary equilibrium, it would have to be one consistent with the Fisher equation, there is no reason to expect that such a policy — however credible it may be that it will be maintained in perpetuity — should lead to a stationary equilibrium, even in the long run.

Let us return to the thought experiment considered in section 3.2 above, in which the central bank pegs the nominal interest permanently at a level lower than the constant level associated with the stationary equilibrium with inflation rate $\bar{\Pi}$; but let us now consider the consequences of adaptation of the estimated value functions. The outcome calculated in section 3.2 assumes that $\nu_t = \tilde{\nu}_t = 0$. If this were to remain true forever, then one would have a constant values for output and inflation given by (3.9), which implies that both inflation and output are *higher* than their values in the stationary equilibrium with the target inflation rate. (In fact, π_t^k and y_t^k are higher for arbitrary k ; so this conclusion is true regardless of the assumed distribution of forecast horizons.) If this were to remain true permanently, the inflation rate and nominal interest rate would fail to conform to the Fisher equation, even in the long run.

But in such a situation, the estimated value functions should not remain the ones that were appropriate to the previous steady state. As discussed in section 3.2, in this thought experiment $Nu^* \ll 0$. Moreover, in the case in which ρ is small enough for the infinite sums $\sum_j \omega_j \tilde{x}_t^j$ to converge, we have $(I - \rho M)^{-1} \gg 0$, and we can solve (4.24) forward to obtain

$$[x_t - \bar{x}_t] = \tilde{x} \equiv [I - \rho M]^{-1} Nu^* \gg 0$$

at each date t . Substitution of this system of equations together with (4.28) into (4.29) yields

$$\begin{bmatrix} \nu_t^{est} \\ \tilde{\nu}_t^{est} \end{bmatrix} = \Phi \tilde{x} + \Phi \Xi \cdot \begin{bmatrix} \nu_t \\ \tilde{\nu}_t \end{bmatrix} \quad (4.32)$$

for each date.

We further observe that in the case of an interest-rate peg, $\Xi \gg 0$ and hence $\Phi \Xi \gg 0$. It then follows from (4.32) that in the case of any beliefs satisfying $\nu_t, \tilde{\nu}_t \geq 0$, we must have $\nu_t^{est}, \tilde{\nu}_t^{est} > 0$. Then since ν_{t+1} is specified to be a weighted average

of ν_t and ν_t^{est} , we must have $\nu_{t+1} > 0$; and similarly for $\tilde{\nu}_{t+1}$. We can then show recursively that starting from initial values $\nu_0 = \tilde{\nu}_0 = 0$ (beliefs consistent with the previous steady state), we must have $\nu_t, \tilde{\nu}_t > 0$ for all $t > 0$. It then follows from (4.28) that $\bar{x}_t \gg 0$ for all $t > 0$, and hence that

$$x_t = \tilde{x} + \bar{x}_t \gg 0$$

for all dates t after the policy change. But this means that the levels of output and inflation can never converge to the long-run steady state consistent with the Fisher equation, since in this steady state both elements of x_t must be *negative*.

Since the dynamics implied by our system of equations are linear, the fact that there is no convergence to the unique steady state means that the dynamics must diverge explosively. Thus once learning dynamics are taken into account, the model predicts an explosive inflationary spiral that should continue until the interest-rate peg is abandoned, as in the famous analysis of an interest-rate peg by Friedman (1968). Similarly, the model implies that commitment to maintaining a fixed high interest rate should never succeed in bringing about a correspondingly high rate of inflation.

Similar conclusions about the instability of learning dynamics under an interest-rate peg have been obtained in the context of New Keynesian models based on intertemporal optimization by authors including Bullard and Mitra (2002), Preston (2005), and Evans and McGough (2017).¹⁷ The present model illustrates, however, that such conclusions can be obtained without modeling expectations as purely backward-looking, as these authors do. The present analysis does allow central-bank announcements about intended future policy to influence behavior immediately, even before any change in actual central-bank behavior occurs, because it is assumed that households and firms should both take into account such information in their forward planning exercises. But this does not imply dynamics that converge to a stationary

¹⁷Gabaix (2017) obtains a different conclusion about the long-run effects of an announced permanent peg of the nominal interest rate at a different level than has prevailed in the past. However, this is because he assumes that the interest-rate peg is accompanied by central-bank “guidance” that directly affects people’s expectations (because their “default model” is influenced by central-bank guidance), and that the way that this guidance, to the extent that it is accepted, implies that one should expect *the inflation rate that makes the new interest-rate peg consistent with the Fisher equation*. That is, Gabaix assumes that people’s beliefs should be at least partially neo-Fisherian, even when their experience points in the opposite direction.

equilibrium consistent with the Fisher equation.

5 Conclusions

The analysis above shows that care must be used in drawing conclusions about contemplated monetary policies using rational expectations equilibrium analysis. We do not mean to suggest that such analysis is never useful. In some cases, the rational expectations equilibrium outcome (with a suitable equilibrium selection) *should* provide a reasonable approximation to what a more realistic model with finite-horizon forward planning would imply, at least if many people are somewhat forward-looking.

For example, this should be the case if one is interested in computing predicted responses to economic disturbances that are (i) relatively transitory, and (ii) recurrent enough for people to have learned their serial correlation from experience, when (iii) the central bank’s reaction function satisfies the “Taylor Principle” (3.8). In such a case, the decisions that result from forward planning are not very sensitive to the length of the horizon over which people plan; and as a consequence, the limit as the horizon length k is made unboundedly large is well-defined, and corresponds to a particular selection from among the rational expectations equilibrium solutions.

But for some questions, there is no selection criterion under which rational expectations equilibrium analysis provides reliable predictions. The question of predicting the effects of a central-bank commitment to maintain its nominal interest-rate target at a low level for a considerable period of time is such a case. We have seen that in this case, finite-horizon forward planning does not result in outcomes similar to any of the rational expectations equilibria consistent with such a policy, no matter how long people’s planning horizons may be, or how rapidly they may adjust their estimated value functions to reflect more recent experience.¹⁸

Avoiding misleading conclusions is only possible by considering the implications of explicit models of boundedly rational cognition, and examining the extent to which they lead to results similar to those of the more familiar rational expectations analyses. The example provided here is intended to show how such analyses can be tractable, in a setting that is no less general than those often used in rational ex-

¹⁸The fragility of rational expectations equilibrium results in this case is also shown by the related work of García-Schmidt and Woodford (2015), who relax the assumption of optimization under rational expectations in a different way.

pectations analyses of alternative policies. It is offered in the hope that analyses in this style will become more common in the monetary policy literature, and that more robust conclusions about policy can be reached as a result.

References

- [1] Branch, William, George W. Evans, and Bruce McGough, “Finite Horizon Learning,” in T.J. Sargent and J. Vilmunen, eds., *Macroeconomics at the Service of Public Policy*, Oxford: Oxford University Press, 2012.
- [2] Bullard, James, and Kaushik Mitra, “Learning About Monetary Policy Rules,” *Journal of Monetary Economics* 49: 1105-1129 (2002).
- [3] Calvo, Guillermo A., “Staggered Prices in a Utility-Maximizing Framework,” *Journal of Monetary Economics* 12: 383-398 (1983).
- [4] Campbell, Murray, A. Joseph Hoane, Jr., and Feng-Hsiung Hsu, “Deep Blue,” *Artificial Intelligence* 134: 57-83 (2002).
- [5] Cochrane, John H., “Michelson-Morley, Fisher, and Occam: The Radical Implications of Stable, Quiet Inflation at the Zero Bound,” unpublished working paper, Stanford University, June 2017.
- [6] 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.
- [7] Evans, George W., and Bruce McGough, “Learning to Optimize,” unpublished working paper, University of Oregon, July 2015.
- [8] Evans, George W., and Bruce McGough, “Interest Rate Pegs in New Keynesian Models,” unpublished working paper, University of Oregon, March 2017.
- [9] Friedman, Milton, “The Role of Monetary Policy,” *American Economic Review* 58: 1-17 (1968).
- [10] Gabaix, Xavier, “A Behavioral New Keynesian Model,” unpublished working paper, Harvard University, February 2017.
- [11] Gabaix, Xavier, David Laibson, Guillermo Moloche, and Steven Weinberg, “Costly Information Acquisition: Experimental Analysis of a Boundedly Rational Model,” *American Economic Review* 96: 1043-1068 (2006).

- [12] García-Schmidt, Mariana, and Michael Woodford, “Are Low Interest Rates Deflationary? A Paradox of Perfect-Foresight Analysis,” NBER Working Paper no. 21614, October 2015.
- [13] Jéhiel, Philippe, “Limited Horizon Forecast in Repeated Alternate Games,” *Journal of Economic Theory* 67: 497-519 (1995).
- [14] Keramati, Mehdi, Peter Smitenaar, 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 USA* 113: 12868-12873 (2016).
- [15] Kydland, Finn E., and Edward C. Prescott, “Time to Build and Aggregate Fluctuations,” *Econometrica* 50: 1345-1370 (1982).
- [16] MacLeod, W. Bentley, “Complexity, Bounded Rationality, and Heuristic Search,” *Contributions to Economic Analysis and Policy* 1(1), article 8 (2002).
- [17] Milani, Fabio, “Expectations, Learning and Macroeconomic Persistence,” *Journal of Monetary Economics* 54: 2065-2082 (2007).
- [18] Preston, Bruce, “Learning About Monetary Policy Rules when Long-Horizon Expectations Matter,” *International Journal of Central Banking* 1: 81-126 (2005).
- [19] Silver, David, *et al.*, “Mastering the Game of Go with Deep Neural Networks and Tree Search,” *Nature* 529: 484-489 (2016).
- [20] Slobodyan, Sergey, and Raf Wouters, “Learning in a Medium-Scale DSGE Model with Expectations Based on Small Forecasting Models,” *American Economic Journal: Macroeconomics* 4: 65-101 (2012).
- [21] Woodford, Michael, *Interest and Prices: Foundations of a Theory of Monetary Policy*, Princeton: Princeton University Press, 2003.
- [22] Woodford, Michael, “Macroeconomic Analysis Without the Rational Expectations Hypothesis,” *Annual Review of Economics* 5: 303-346 (2013).
- [23] Yun, Tack, “Nominal Price Rigidity, Money Supply Endogeneity, and Business Cycles,” *Journal of Monetary Economics* 37: 345-370 (1996).