Breaking the New Keynesian Dichotomy: Asset Market Segmentation and the Monetary Transmission Mechanism

Robert G. King: Boston University and NBER
Julia K. Thomas: Federal Reserve Bank of Philadelphia and NBER[1]

September 23, 2007

[1]email: rking@bu.edu, mail@juliathomas.net. We thank Aubhik Khan, Frank Schorfheide, and session participants at the 2007 Midwest Macro and SED meetings for useful comments and suggestions. The views expressed here are those of the authors and do not represent the views of Federal Reserve Bank of Philadelphia or the Federal Reserve System.

Abstract

We develop a general framework to examine how the presence of a monetary transmission mechanism shapes aggregate responses to shocks and the effects of monetary policy. Our framework nests two leading monetary models: a text-book New Keynesian setting and a setting where small transactions costs associated with adjustments in households' money balances lead to an evolving distribution of money across households. In the text-book New Keynesian model, the effective isolation of a single condition determining aggregate money demand imposes a dichotomy that eliminates any role played by the demand for money in the determination of aggregate demand whenever the monetary authority uses an interest rate rule. As such, it rationalizes a narrow attention to direct links between interest rate setting and objectives such as desired paths for inflation and real activity in a wide range of current discussions involving monetary policy. In this paper, we argue that the simplicity of the Keynesian dichotomy is not an inevitable or desirable feature of a tractable monetary model.

The basic mechanism implying non-neutralities in our second nested model does not permit the dichotomy raised above. Rather, households' decisions regarding consumption and labor supply in this model are intimately affected by both their individual money holdings and, through wages and interest rates, the entire distribution of money balances. This implies a rich monetary transmission mechanism, in that the level of aggregate demand depends crucially on monetary factors. Examining a composite setting where the pricing frictions of New Keynesian monetary models are allowed to interact with the rich money demand mechanism implied by households' inventory-theoretic portfolio management, we find that the resulting model is not only tractable, but also has very desirable properties from an empirical standpoint. When solved under a money stock rule, it implies a path for the nominal interest rate that initially declines in the face of a monetary expansion, in keeping with the liquidity effect documented across a broad range of empirical studies. Moreover, when solved under a standard interest rate rule, our model implies greater persistence in the dynamic responses following shocks to monetary policy, as well as nonmonotone responses to real shocks. These desirable implications emerge precisely because it is not possible to describe aggregate demand without reference to money demand in our model. The distribution of transactions balances across individuals is an essential part of the transmission mechanism from monetary policy actions to real economic activity.

1 Introduction

Most small modern macroeconomic models used for conceptual and quantitative monetary policy analysis have the property that the demand for money is irrelevant to the determination of aggregate demand, when the monetary authority is using an interest rate rule.[1] This property – which we label the Keynesian dichotomy – has a long history in macroeconomic analysis. Indeed, the earliest generation of quantitative monetary policy models, built by various research teams in the 1950s and 1960s, did not even include a demand for money. However, its role has been strongly reinforced by the currently dominant set of macroeconomic models, in which it is nearly always a key ingredient. Further, in a wide range of current discussions of monetary policy – theoretical, applied, and practical – this property is used to rationalize a focus entirely on the links between settings of the short-term interest rate and objectives such as desired paths for inflation and real activity. In this paper, we argue that the Keynesian dichotomy is not an inevitable or desirable feature of macroeconomic models, developing a fully articulated macroeconomic model in which portfolio adjustment costs destroy the dichotomy and provide the basis for a nonstandard interplay between interest rates and real activity.

1.1 An illustration of the dichotomy

As a reference point for the discussion below, consider the following textbook modern macroeconomic model. There is an "IS curve" which takes the form

$$y_t = E_t y_{t+1} + s[r_t - r],$$

where y_t represents aggregate demand/output, $E_t y_{t+1}$ is expected future output, and r_t is the real interest rate. Next, there is a Fisher equation,

$$i_t = r_t + E_t \pi_{t+1},$$

which links the nominal interest rate, i, to the the real interest rate and expected inflation. Finally, there is a description of "inflation dynamics" or a "Phillips curve" that relates current inflation to expected future inflation and a deviation of output from a natural rate level.

$$\pi_t = \beta E_t \pi_{t+1} + h[y_t - y_t^*]$$

In the specific three-equation model above, monetary variables do not enter either in the IS curve or elsewhere. It is this property that we label the Keynesian dichotomy. It is a characteristic

[1] Analytical examples of these models may be found in Clarida, Gali and Gertler (1999) , McCallum and Nelson (1999), and Woodford (2007). To stress the nature of monetary policy in this setting, Kerr and King (1996) describe the framework as an "IS model". Larger models used for policy evaluation – along the lines of Smets and Wouters (2003) for the Euro area and Christiano, Eichenbaum and Evans (2005) for the US – also display something close to the dichotomy, although these researchers add other elements affecting the dynamics of aggregate demand. The Smets and Wouters (2003) model can be discussed without reference to monetary aggregates. Christiano et. al. have a monetary constraint of a one period form that plays some role in the determination of aggregate demand.

shared by a large class of models used by central banks around the world, which elaborate the aggregate demand equation into a block of equations and the inflation equation into a block of equations. Thus, monetary variables are seen to enter the determination of aggregate activity only if the central bank is operating using a monetary quantity instrument or if the central bank's rule for setting an interest rate instrument rule places weight on monetary variables.

1.2 Developing an alternative model

We provide a framework that features a role for monetary variables in the determination of aggregate demand and real activity, but also nests a textbook fully articulated New Keynesian macroeconomic model as a limiting special case. That special case generates the simple textbook model above as a linear approximation result around zero inflation. In the textbook setting, there is monetary non-neutrality in the short run solely because firms implicitly face small menu costs of price adjustment. In this model and many variants, the dynamic adjustment process to real and nominal shocks is heavily influenced by the fact that some firms make price adjustments quickly, while others do not, so that there is an evolving distribution of nominal prices. This distribution implies that adjustments in the aggregate price level following monetary disturbances are gradualized, so that nominal shocks have real consequences over the short-run.

In its most common form, the Keynesian dichotomy is imposed by assuming a role for money that is self-contained, effectively quarantined from other variables in the model. We use a simple and popular version in which real balances appear in the economy only as an additively separable source of household utility. Figure 1 illustrates how the perfect dichotomy therein leads to a complete irrelevance of money demand when monetary policy takes on an active stabilization role implemented through interest rate targeting, the Taylor rule so frequently analyzed throughout modern monetary economics. There, we see that the initially limited role of money in the textbook economy is effectively eliminated to accommodate a policy rule that maps quite directly from interest rate setting to realized objectives for the paths of inflation and output. As a result, there is effectively no monetary transmission mechanism between instrument and goals to complicate, or enrichen, our analysis of the model economy's dynamics.

The framework we develop in section 2 also houses, as a second special case, a flexible-price model in which households face small transactions costs of making adjustments in the monetary balances that they use to finance expenditure. In this class of models, the dynamic adjustment process to real and nominal shocks is heavily influenced by the fact that some households make portfolio adjustments quickly, while others do not. Because this setting yields an evolving non-trivial distribution of money, the Keynesian dichotomy is forcefully broken. There, the level of aggregate demand depends crucially on monetary factors.

In this second special case model, monetary non-neutrality arises from time-varying and heterogeneous money spending rates on the part of households, rather than from heterogeneous nominal prices on the part of firms. A transactions-based role for money is imposed through the assumption that all goods and labor market transactions must be conducted with money, but the environment differs from that in a traditional cash-in-advance model in two important respects. First, households are able to adjust their money holdings after the resolution of all uncertainty within a period, so they cannot be forced to hold an undesirable quantity of money within the period purely as a result of an unforeseen action on the part of the monetary authority. Second, however, this ability to adjust money balances is not without frictions, so households do not completely undo the effects of a money injection with proportional rises in aggregate expenditure yielding immediate one-for-one adjustment in the aggregate price level, as they would in a perfect-foresight cash-in-advance setting. Instead, households are assumed to face fixed costs of transferring wealth between interest-bearing assets and money. Given these transactions costs, there is *endogenous asset market segmentation* in that households choose to access their interest income infrequently.

To implement their infrequent asset market participation, households carry inventories of money in excess of current spending to finance their spending over future dates. As a result, most households do not exhaust their available money within any given period, so that aggregate velocity deviates from 1. Moreover, this aggregate spending rate varies over time, because households are able to change both the timing of their participation in asset markets in response to real and nominal shocks, as well as the individual spending rates they adopt given current money holdings.

It is well known that market segmentation implies that open market operations can have real effects, because they directly involve only a subset of households.[2] The advantage of the endogenous version of the framework we examine is that it allows changes in the fractions of households participating in the asset markets – i.e., changes in the extent of market segmentation – over time in response to aggregate disturbances. These changes can produce long-lasting disruptions in the distribution of household money holdings that lend added persistence to movements in real and nominal variables.

Following the model descriptions in sections 2-3, and a brief summary of functional forms and parameter values in section 4, we consider the implications of our alternative model containing both mechanisms discussed above, the evolving distribution of nominal prices, as well as the rich distribution of relative money holdings across households. In section 5, we show that this model has some very desirable properties. When solved under a money stock rule, for example, it predicts that the nominal interest rate initially declines in the face of a monetary expansion, in contrast to the inevitable rise in the pure New Keynesian model. When it is solved under a standard interest rate rule, we observe much more protracted quantity responses to monetary policy shocks, entirely due to the richer dynamics of aggregate demand.

Most economists would view these features as desirable implications of a macroeconomic model, bringing it closer to conventional viewpoints about the implications of actual policy interventions. However, these implications come precisely because it is not possible to describe aggregate demand

[2] See, for example, Alvarez, Atkeson and Edmond (2003), Alvarez, Lucas and Weber (2001), and the references therein. Khan and Thomas (2007) build on the work in this literature, and that of Alvarez, Atkeson and Kehoe (2002), to develop a flexible-price environment wherein asset market segmentation evolves endogenously.

without reference to money demand. The distribution of transactions balances across individuals is a key part of the transmission mechanism from monetary policy actions to real economic activity. In particular, rich dynamics emerge in our model from changes in the distribution of money, which themselves can deliver powerful propagation of shocks, adding persistence to the movements in output, employment, inflation and real interest rates following an open market operation, as well as nonmonotone responses in these variables following a persistent shock to productivity in the presence of an interest rate rule. While interest rate targets continue to shape economic activity in our model, the targets are themselves influenced by both current and future changes that they induce in money demand. Thus, the convenient mapping from the policy instrument to the economy's resulting aggregate dynamics is destroyed.

2 Model

This section presents a model containing both the core elements of a New Keynesian model and those elements that distinguish models with segmented asset markets. The New Keynesian aspects of our model lie in the production side of the model, where imperfect competition among firms combines with staggered nominal price adjustments to imply non-neutralities and, in particular, a short-run Phillips curve relationship. The market segmentation aspects, by contrast, are in the household side of our model. There, small financial market frictions give rise to a nontrivially evolving distribution of money holdings across households that, through changes in velocity, also delivers short-term non-neutralities.

Our full model embeds a typical New Keynesian production side of the economy, where the timing of monopolistically competitive intermediate goods producers' nominal price-setting is governed by a (flat) Calvo hazard, together with an endogenously segmented asset markets side of the economy, where households maintain inventories of money to finance their spending across multiple periods, due to fixed transactions costs incurred when they transform their less-liquid, higher return assets into money (and vice versa).

In the following subsections, we first describe the production side of the economy, then the household side, and follow next with the equilibrium conditions that connect the two. As we lay out the details below, we will take care to note the modifications required to obtain the textbook New Keynesian and the flexible-price segmented asset markets environments. The mechanics of these two special case models may be useful for reference as we examine the dynamics of our full model in section 5.

2.1 Firms: New Keynesian production environment

A perfectly competitive representative producer supplies the economy's final consumption good using a continuum of intermediate inputs, y_i, $i \in [0,1]$. We assume a constant elasticity of substitution across intermediate inputs in the production of final goods; specifically, the producer's total output is $Y = \left(\int_0^1 \left[y_i^{\frac{\varepsilon-1}{\varepsilon}}\right] di\right)^{\frac{\varepsilon}{\varepsilon-1}}$, where $\varepsilon > 1$. Each intermediate input is produced by a single monopolistic competitor using labor, n_i, as the sole factor of production. All intermediate producers (henceforth, *firms*) have the common technology $y_i = zn_i^\nu$, where z is a persistent aggregate productivity level, and $\nu \in (0,1]$. The aggregate state of the economy is $s = (\boldsymbol{\kappa}, z, \mu)$, where $\boldsymbol{\kappa}$ represents the endogenous aggregate vector, and μ is the current growth rate of the aggregate money supply. Firms and households take the aggregate state as given, as well as its evolution over time according to a law of motion $s' = F(s)$, which we must determine in equilibrium.

2.1.1 Final good producer

In any date, the final good producer solves the following profit maximization problem, taking as given the nominal price level associated with final goods, P, and the nominal prices of intermediate inputs, P_i.

$$\max_{y_i, i\in[0,1]} P\left[\int_0^1 y_i^{\frac{\varepsilon-1}{\varepsilon}} di\right]^{\frac{\varepsilon}{\varepsilon-1}} - \int_0^1 P_i y_i di$$

The resulting first order conditions with respect to each individual input i are easily rearranged to yield demand functions of the form $y_i = \left(\frac{P_i}{P(\mathbf{s})}\right)^{-\varepsilon} Y(s)$. Thus, for each intermediate input, the price elasticity of demand is $-\varepsilon$. Next, we define the *relative price* of the i^{th} intermediate as $p_i \equiv \frac{P_i}{P(\mathbf{s})}$. Using this notation, we may write the demand functions facing firms as:

$$d(p_i, \mathbf{s}) = p_i^{-\varepsilon} Y(\mathbf{s}). \tag{1}$$

Finally, before leaving this subsection, we use the demand functions above to calculate the nominal price of a unit of final output consistent with the final good producer's zero profit condition. This gives us the following expression for the aggregate price level, P.

$$P \equiv \left[\int_0^1 P_i^{1-\varepsilon} di\right]^{\frac{1}{1-\varepsilon}} \tag{2}$$

2.1.2 Intermediate input firms

The firms that supply intermediate inputs are monopolistic competitors, each setting the nominal price of its good (occasionally). Given current relative price p_i, the aggregate scale of production, $Y(s)$, and the demand functions from (1) above, the flow profit of the i^{th} firm is:

$$\pi(p_i; \mathbf{s}) = p_i^{1-\varepsilon} Y(\mathbf{s}) - e\left(y_i(p_i, \mathbf{s}); w(\mathbf{s}), z\right) \tag{3}$$

where $e\left(y; w, z\right)$ represents its cost of producing y units of output. Given the production function $y = zn^\nu$, this total cost function is:

$$e\left(y_i(p_i, \mathbf{s}); w(\mathbf{s}), z\right) = w(\mathbf{s}) \left(\frac{y_i}{z}\right)^{\frac{1}{\nu}}, \tag{4}$$

and the corresponding marginal cost is $\frac{w(\mathbf{s})}{\nu z} \left(\frac{y_i}{z}\right)^{\frac{1-\nu}{\nu}}$.

If our intermediate goods firms faced no frictions in setting their nominal prices, each would simply maximize its static profits in each period, setting its relative price as a familiar markup over its marginal cost of production, with the markup given by $\frac{\varepsilon}{\varepsilon-1}$.[3] However, in the New Keynesian production environment we consider here, firms must be more forward-looking in their decisions, because they are able to change the nominal price of their output only infrequently. In particular, we assume that each firm faces a constant (and common) probability, $\boldsymbol{\alpha} \in (0,1)$, of having a price-adjustment opportunity. Of course, there are alternatives to this assumption. For instance, we could adopt the approach developed by Dotsey, King and Wolman (1999) to explicitly derive staggered price-setting among firms by assuming they encounter fixed costs each time they adjust their prices. Instead, we choose to directly impose firm-level price rigidity using the common Calvo (1983) price-setting assumption to allow greater comparability with the bulk of the New Keynesian literature.

We denote by p_0 the relative price selected by a firm that is currently able to reset its price. The problem of such a price-setting firm is given by (5) - (6) below. First, we have:

$$\begin{aligned} V^0(\mathbf{s}) &= \max_{p_0} \Omega(\mathbf{s})\,\pi(p_0;\mathbf{s}) + \beta\Big[\boldsymbol{\alpha}\int V^0(\mathbf{s}')\,dF(\mathbf{s},d\mathbf{s}') \\ &\qquad + (1-\boldsymbol{\alpha})\int V^1\left(p_0\frac{P}{P'};\mathbf{s}'\right)dF(\mathbf{s},d\mathbf{s}')\Big], \end{aligned} \tag{5}$$

where $V^0(\mathbf{s}')$ represents the value of the firm next period if it is able to again select its price, and $V^1\left(p_0\frac{P}{P'};\mathbf{s}'\right)$ represents its next period value otherwise. The term $\Omega(\mathbf{s})$ in the functional equation above represents the current real value placed on a unit of final output. Its presence as a weight on current profit flows places the firm's value in units of marginal utility, so that the firm is seen to discount future profits by the constant (household subjective) discount factor β, rather than the stochastic discount factor $\frac{\beta\Omega(\mathbf{s}')}{\Omega(\mathbf{s})}$. In the pure New Keynesian special case model, Ω is, in equilibrium, simply the representative household's marginal utility of consumption. In our full model with limited participation, by contrast, the relevant marginal utility determining Ω is that of a household currently active in the asset markets, as it is only such households that are able to convert assets into money and hence into goods. While we will occasionally suppress dependencies for expositional convenience below, this marginal valuation is, of course, a function of the aggregate state, s, as is the case for Y, total production, alongside the real wage and the price level, w and P.

With probability $\boldsymbol{\alpha}$, the firm now setting its price will be able to do so again in the next period, and there will be no future implications of this period's price. However, with probability $(1-\boldsymbol{\alpha})$, the firm will be unable to adjust its price again in the next period. In that case, given the current relative price, p_0, and corresponding nominal price P_0, its relative price will be $p_0\frac{P}{P'} = \frac{P_0}{P'}$. The

[3] In the special case of linear production most commonly considered, the marginal cost itself would be constant with respect to production, implying $p(i) = \frac{\varepsilon}{\varepsilon-1}\frac{w}{z}$.

value of a firm with relative price p that is currently unable to adjust its price is:

$$
\begin{aligned}
V^{1}(p;\mathbf{s}) \quad &= \quad \Omega(\mathbf{s})\,\pi(p;\mathbf{s}) + \beta\Big[\boldsymbol{\alpha}\int V^{0}(\mathbf{s}')\,dF(\mathbf{s},d\mathbf{s}') \\
&\qquad + (1-\boldsymbol{\alpha})\int V^{1}\left(p\frac{P}{P'};\mathbf{s}'\right)dF(\mathbf{s},d\mathbf{s}')\Big].
\end{aligned} \tag{6}
$$

We characterize the solution to the firm's problem (5 - 6) in Appendix A, deriving an expression that determines the optimal price as a function of expected future interest rates alongside current and expected future demand and marginal cost conditions. Following the result obtained there (in equation 31), we replace the relative price with its nominal counterpart, $P_{0,t} = p_{0,t}P_t$, and simplify the resulting expression to arrive at:

$$
\begin{aligned}
&\left(\frac{P_{0t}}{P_{t-1}}\right)^{\frac{\nu(1-\varepsilon)+\varepsilon}{\nu}}\left(\mathbf{E}_t\sum_{j=0}^{\infty}\left[\beta(1-\boldsymbol{\alpha})\right]^{j}\Omega_{t+j}Y_{t+j}\left(\frac{P_{t+j}}{P_{t-1}}\right)^{\varepsilon-1}\right) \\
&\qquad = \frac{\varepsilon}{\varepsilon-1}\mathbf{E}_t\sum_{j=0}^{\infty}\left[\beta(1-\boldsymbol{\alpha})\right]^{j}\Omega_{t+j}Y_{t+j}^{\frac{1}{\nu}}\frac{w_{t+j}}{\nu}z_{t+j}^{\frac{-1}{\nu}}\left(\frac{P_{t+j}}{P_{t-1}}\right)^{\frac{\varepsilon}{\nu}}.
\end{aligned} \tag{7}
$$

We can make this price setting rule recursive by defining two forward-looking summary variables Δ_t^0 and Δ_t^1 as follows.

$$
\begin{aligned}
\Delta_t^0 \quad &\equiv \quad \mathbf{E}_t\sum_{j=0}^{\infty}\left[\beta(1-\boldsymbol{\alpha})\right]^{j}\Omega_{t+j}Y_{t+j}\left(\frac{P_{t+j}}{P_{t-1}}\right)^{\varepsilon-1} \\
\Delta_t^1 \quad &\equiv \quad \frac{\varepsilon}{\varepsilon-1}\mathbf{E}_t\sum_{j=0}^{\infty}\left[\beta(1-\boldsymbol{\alpha})\right]^{j}\Omega_{t+j}Y_{t+j}^{\frac{1}{\nu}}\frac{w_{t+j}}{\nu}z_{t+j}^{\frac{-1}{\nu}}\left(\frac{P_{t+j}}{P_{t-1}}\right)^{\frac{\varepsilon}{\nu}}
\end{aligned}
$$

Given these definitions, we can re-write (7) in first-order form using the following three equations.

$$
\left(\frac{P_{0t}}{P_{t-1}}\right)^{\frac{\nu(1-\varepsilon)+\varepsilon}{\nu}}\Delta_t^0 = \Delta_t^1, \tag{8}
$$

where the definitions of Δ^0 and Δ^1 imply

$$
\Delta_t^0 \quad = \quad \Omega_t Y_t\left(\frac{P_t}{P_{t-1}}\right)^{\varepsilon-1} + \beta(1-\boldsymbol{\alpha})\left(\frac{P_t}{P_{t-1}}\right)^{\varepsilon-1}\mathbf{E}_t\Delta_{t+1}^0 \tag{9}
$$

$$
\Delta_t^1 \quad = \quad \frac{\varepsilon}{\varepsilon-1}\Omega_t Y_t^{\frac{1}{\nu}}\frac{w_t}{\nu}z_t^{\frac{-1}{\nu}}\left(\frac{P_t}{P_{t-1}}\right)^{\frac{\varepsilon}{\nu}} + \beta(1-\boldsymbol{\alpha})\left(\frac{P_t}{P_{t-1}}\right)^{\frac{\varepsilon}{\nu}}\mathbf{E}_t\Delta_{t+1}^1. \tag{10}
$$

These equations, (8), (9) and (10), completely summarize optimal price-setting in our model. We will next turn to consider the growth rate of the aggregate price level in this setting, and thereafter derive expressions for aggregate employment demand and profits.

2.1.3 Inflation

From equation (2) above, we know that the aggregate price level is an integral of firms' nominal prices that implies $P_t^{1-\varepsilon} = \int_0^1 P_t^{1-\varepsilon}(i)$. In the current period, $\boldsymbol{\alpha}$ fraction of firms set their nominal price to $p_{0t}P_t = P_{0,t}$. Of the remaining $1-\boldsymbol{\alpha}$ fraction of firms, $\boldsymbol{\alpha}$ fraction are firms that set their nominal price last period (to $P_{0,t-1}$), and so forth. Based on these observations, the fraction of firms producing this period with a nominal price last selected j periods ago is $\boldsymbol{\alpha}(1-\boldsymbol{\alpha})^j$, which brings us to the following equation.

$$P_t^{1-\varepsilon} = \boldsymbol{\alpha} P_{0,t}^{1-\varepsilon} + (1-\boldsymbol{\alpha})\,\boldsymbol{\alpha} P_{0,t-1}^{1-\varepsilon} + (1-\boldsymbol{\alpha})^2\,\boldsymbol{\alpha} P_{0,t-2}^{1-\varepsilon} + \cdots + (1-\boldsymbol{\alpha})^j\,\boldsymbol{\alpha} P_{0,t-j}^{1-\varepsilon} + \cdots$$

Substituting into this equation the lagged version of itself, we arrive at a law of motion for the aggregate price level,

$$P_t = \left(\boldsymbol{\alpha} P_{0t}^{1-\varepsilon} + (1-\boldsymbol{\alpha})\,P_{t-1}^{1-\varepsilon}\right)^{\frac{1}{1-\varepsilon}},$$

from which we obtain an expression for (gross) inflation rates.

$$\pi_t \equiv \frac{P_t}{P_{t-1}} = \left(\boldsymbol{\alpha}\left(\frac{P_{0t}}{P_{t-1}}\right)^{1-\varepsilon} + (1-\boldsymbol{\alpha})\right)^{\frac{1}{1-\varepsilon}} \tag{11}$$

Notice that this expression links the inflation rate to the nominal price selected by firms that are able to change their price this period. Equations (8) - (10) and (11) jointly determine $\frac{P_{0t}}{P_{t-1}}$, Δ_t^0, Δ_t^1 and $\frac{P_t}{P_{t-1}}$ given $(Y_t, w_t, \Omega_t, z_t)$. In general, given non-zero inflation, these variables will be functions of the distribution of employment across firms, which in turn depends upon the degree of nominal price dispersion.

2.1.4 Aggregate employment demand and profits

Let n_{jt} denote the labor demanded by a firm that last set its price j periods ago. Recalling the total cost function $e(y;w,z) = w\left(\frac{y(i)}{z}\right)^{\frac{1}{\nu}}$, and noting that the current relative price of any such firm is $\frac{P_{0,t-j}}{P_t}$, we may write its labor demand as a function of aggregate production; $n_{j,t} = \left[\frac{Y_t}{z_t}\left(\frac{P_{0,t-j}}{P_t}\right)^{-\varepsilon}\right]^{\frac{1}{\nu}}$. Noting again the constant fraction of firms able to set their prices in each date, $\boldsymbol{\alpha}$, aggregate employment demand may be written as $N_t^D = \sum_{j=0}^{\infty} n_{jt}\boldsymbol{\alpha}(1-\boldsymbol{\alpha})^j$, or:

$$N_t^D = \left(\frac{Y_t}{z_t}\right)^{\frac{1}{\nu}} \left(\frac{P_{t-1}}{P_t}\right)^{-\frac{\varepsilon}{\nu}} \sum_{j=0}^{\infty} \boldsymbol{\alpha}(1-\boldsymbol{\alpha})^j \left(\frac{P_{0,t-j}}{P_{t-1}}\right)^{-\frac{\varepsilon}{\nu}}.$$

To make the expression above recursive, we define the (time $t+1$) state variable $\Delta_t^N \equiv \sum_{j=0}^{\infty} \boldsymbol{\alpha}(1-\boldsymbol{\alpha})^j \left(\frac{P_{0,t-j}}{P_{t-1}}\right)^{-\frac{\varepsilon}{\nu}}$. Using this definition, we have:

$$N_t^D = \left(\frac{Y_t}{z_t}\right)^{\frac{1}{\nu}} \left(\frac{P_{t-1}}{P_t}\right)^{-\frac{\varepsilon}{\nu}} \Delta_t^N, \text{ where} \tag{12}$$

$$\Delta_t^N = \boldsymbol{\alpha}\left(\frac{P_{0,t}}{P_{t-1}}\right)^{-\frac{\varepsilon}{\nu}} + (1-\boldsymbol{\alpha})\left(\frac{P_{t-2}}{P_{t-1}}\right)^{-\frac{\varepsilon}{\nu}} \Delta_{t-1}^N. \tag{13}$$

Equations (12) and (13) determine total employment demand, given $\left(\frac{P_{0,t}}{P_{t-1}}\right)$, the lagged inflation rate $\left(\frac{P_{t-1}}{P_{t-2}}\right)$, Δ^N_{t-1}, aggregate output, Y_t and the technology shock z_t. Finally, it is straightforward to show that here, just as in a flexible-price setting, aggregate profits are simply the difference between aggregate output and real wage payments,

$$\Pi_t = Y_t - w_t N_t^D. \tag{14}$$

2.2 Households: Endogenous asset market segmentation environment

As we proceed to describe the environment in this side of our economy, we introduce three sets of agents: a unit measure of ex-ante identical households, a perfectly competitive financial intermediary, and a monetary authority. Of these, only households require discussion in any depth. In our full model, we abandon the reduced-form representation of money demand typically adopted by New Keynesian monetary models in favor of a more explicit transactions-based demand for money. Here, we adopt the household environment developed by Khan and Thomas (2007) to examine the implications of endogenous asset market segmentation.[4] Thus, we replace the assumption of a representative household with a nontrivial distribution of households. These households differ in their current money holdings, but they are able to ensure their bond holdings by pooling idiosyncratic risk period-by-period within an extended family of which they all are members. This family of households will be described further below, and is based on that derived from individual households' lifetime utility maximization problems in Khan and Thomas (2007).

Any given household in our economy is distinguished by its individual history of realized financial market transactions costs, ξ, which are independently drawn at the start of each period from a time-invariant distribution. Money is required for all transactions in the goods market. However, inside any period, a household may only exchange assets in the bond market for money in its bank account upon payment of its fixed transactions cost. Thus, the household undertakes such a trade, becoming *active* in the asset markets, only if its current transactions cost is sufficiently low. Households differ increasingly in their money holdings over time as those encountering relatively high transactions costs, and thus avoiding trades, for many periods see their real balances further and further eroded relative to those of recently active households.

Each infinitely-lived household values consumption and leisure in each period according to a period utility function defined over consumption and labor, $u(c, n)$, where $n = 1 - L$, and each discounts future utility by the constant subjective discount factor $\beta \in (0, 1)$. Households have two means of saving. First, they have access to a complete set of state-contingent nominal bonds, which they maintain in interest-bearing brokerage accounts. Next, they hold money in non-interest-bearing bank accounts in order to purchase consumption goods.[5] The assumption

[4] The Khan and Thomas model is a special case of our environment in which prices are fully flexible and markets are perfectly competitive. We describe this special case model further in section 3.2.

[5] The distinction between bonds and money here is sharp in that money earns a zero nominal rate of return. More generally, we view the variable termed bonds as relatively illiquid, high-yield assets, and that termed money as more liquid assets that are substantially dominated by bonds in their average rate of return of return.

that all trades in the goods market require money (given current nominal wage and profit income is delivered only at the end of the period) ensures that all households carry money within a period. However, as mentioned above, we also assume that households must pay fixed transactions costs each time they transfer assets between their brokerage and bank accounts. This additional friction ensures that households deliberately hold money across periods. In particular, they choose to carry inventories of money, managing them according to generalized (S,s) rules, in order to limit the frequency of their transfers.

The transactions costs that give rise to asset market segmentation, ξ, are fixed in that they are independent of the size of the current account transfer. However, they vary over time and across households. Here, for convenience, we subsume the idiosyncratic features distinguishing households directly in their transactions costs by assuming that each household draws its own current cost from the time invariant distribution $H(\xi)$ upon entering each period. While households are ex-ante identical, a household's current ξ influences the decision of whether to undertake an account transfer within the period. Thus, it affects the household's current consumption and labor supply and, in turn, the money with which it exits the period. It is for this reason that, at any date t, households are distinguished by their histories of these draws, $\xi^t = (\xi_1, \xi_2, \cdots, \xi_t)$.

Let the economy's aggregate history be denoted by $\mathbf{s}^t$. Given date-event history, $(\mathbf{s}^t, \xi^t)$, a household has the following available assets as it enters a period. In its brokerage account, it has nominal bonds $B(\mathbf{s}^t, \xi^t)$, which it purchased in the previous period at price $q(\mathbf{s}^{t-1}, \mathbf{s}_t, \xi_t)$ from the perfectly competitive financial intermediary. It is essential to note that these bonds are contingent on both aggregate *and individual* state variables. Households' access to these fully state-contingent bonds imply that they may perfectly insure themselves in their brokerage accounts. In addition to these assets, each household also has available in its brokerage account a fraction $(1-\lambda)$ of its nominal wage and lump-sum profit income from the previous period, $P(\mathbf{s}^{t-1})[w\left(\mathbf{s}^{t-1}\right) n(\mathbf{s}^{t-1}, \xi^{t-1}) + \Pi(\mathbf{s}^{t-1})]$. The remaining λ fraction of that income is available in the household's bank account, and supplements the money it retained there from the previous period, $A(\mathbf{s}^{t-1}, \xi^{t-1})$. We represent the total start-of-period nominal bank balance by $M(\mathbf{s}^{t-1}, \xi^{t-1})$, where

$$M(\mathbf{s}^{t-1}, \xi^{t-1}) = A(\mathbf{s}^{t-1}, \xi^{t-1}) + \lambda P(\mathbf{s}^{t-1})[w\left(\mathbf{s}^{t-1}\right) n(\mathbf{s}^{t-1}, \xi^{t-1}) + \Pi(\mathbf{s}^{t-1})].$$

With knowledge of its current portfolio and transactions cost, as well as the current aggregate state, a household chooses whether it will shift some assets across its accounts before the second half of the period, when production and shopping take place. The table below summarizes this financial market trading decision in nominal terms.

	withdrawal from brokerage	*post-transfer bank balance*
active	$P(\mathbf{s}^t)\xi_t + x\left(\mathbf{s}^t, \xi^t\right)$	$M\left(\mathbf{s}^{t-1}, \xi^{t-1}\right) + x\left(\mathbf{s}^t, \xi^t\right)$
inactive	0	$M\left(\mathbf{s}^{t-1}, \xi^{t-1}\right)$

Row 1 describes what happens if the household chooses to become *active* in managing its assets. In this case, it pays its current nominal transactions cost from its brokerage account, and it selects a nominal transfer, $x\left(\mathbf{s}^t, \xi^t\right)$, to be made to its bank account. The only restriction on the transfer is that it result in non-negative balances in both accounts;

$$x\left(\mathbf{s}^t, \xi^t\right) \in [-M\left(\mathbf{s}^{t-1}, \xi^{t-1}\right), B(\mathbf{s}^t, \xi^t) + (1-\lambda)P(\mathbf{s}^{t-1})[w\left(\mathbf{s}^{t-1}\right)n(\mathbf{s}^{t-1}, \xi^{t-1}) + \Pi(\mathbf{s}^{t-1})].$$

Given its transfer, the household then enters the current production/shopping sub-period with $M\left(\mathbf{s}^{t-1}, \xi^{t-1}\right) + x\left(\mathbf{s}^t, \xi^t\right)$ available in its bank balance. Alternatively, the household may choose to remain *inactive* (row 2), undertaking no account transfer and thus entering shopping with its start-of-period nominal balances. Regardless of which option it selects, the household's post-transfer bank balance must finance all current consumption expenditure, $P(\mathbf{s}^t)c\left(\mathbf{s}^t, \xi^t\right)$, because the money it retains for next period, $A(\mathbf{s}^t, \xi^t)$, is required to be non-negative.

2.2.1 Some simplifying results

An appealing aspect of our endogenous limited participation environment is that the heterogeneity among households is computationally manageable. We do not re-derive the results from Khan and Thomas (2007) here, however we do summarize the relevant economic points. Each of the main theoretical results that simplify the model's solution may be traced to one of three essential assumptions in the model. First, we have said that all households have access to a complete set of state-contingent bonds in their brokerage accounts. Next, we further assume that they are able to purchase these bonds in an initial period 0 when they are identical. In that period, the government has some claims against it, which are distributed evenly across households. Households simply use their initial wealth to purchase bonds for period 1 when transactions costs will start to distinguish them, and they will begin working and consuming. Finally, the third assumption is the time-invariant distribution from which transactions costs are drawn each period, that is, the fact that these costs are not serially correlated.

Given the assumptions above, it is immediate that all households have both a common lifetime budget constraint in their brokerage accounts and common expectations as they select their state-contingent lifetime plans for consumption, labor supply, transfers, and saving. Thus, they select the same lifetime plans, and are able to eliminate the effects of their past individual shocks whenever they choose to access their brokerage accounts. This makes the household-side of our model equivalent to one where households pool their brokerage account risk period-by-period and together hold the aggregate portfolio of government bonds in a family brokerage account to which they all have equal claim. As a result, *money is the single household state variable.* Because households' bond holdings are perfectly insured against idiosyncratic risk, a household's current transactions cost ξ will affect its decision of whether to undertake a transfer, but will not affect its brokerage balance. While the transfer decision in turn influences the household's consumption and the money it saves, the fact that transactions costs are serially independent means that ξ_t gives no information about ξ_{t+1}. Thus, it has no influence upon future variables beyond the money the household takes into the next period. As a result, we need not track households' shock histories over time, so long as we know the money balances with which they enter the period.

The second important result is that even start-of-period money balances cease to matter for a household once it decides to become active. *All households currently active make common decisions*, because they have access to their perfectly insured brokerage accounts, and thus can eliminate the effects of their individual histories. Thus, they choose the same current consumption and labor supply. Further, combining this fact with their common expectations over future shocks, we know that they also choose the same money to retain for next period. This means that they will enter into the next period effectively identical. Moreover, of this group of currently active households, those that are inactive next period will share common money holdings and thus again make common decisions. As such, they will remain identical going into the following period.

From the observations above, it becomes clear that all households last actively trading in the asset markets at some common date, say j periods in the past, enter into the current period effectively identical in that (a) the relevant differences across households are limited to their money balances, given perfect insurance in the bond market, and (b) once the decision of whether or not to pay the current transactions cost to access the brokerage account has been made for the current period, all households that last traded bonds for money or vice-versa at the same time have the same current balances, and thus make the same current decisions. Thus, we can *track the distribution of households by grouping them together according to their time-since-(last)-active.* All we need know are the fractions of households in each group, along with the group-specific money balances with which every member of any one such group enters the period. Moreover, because we assume a finite upper support on the distribution of transactions cost, this distribution is summarized by the population fractions and money holdings of a finite number of time-since-active groups, since all households eventually trade when they find their money holdings sufficiently far from their desired, or target, real balances. Finally, it is straightforward to show that *households adopt threshold rules in determining whether to become active.* That is, given the current aggregate state and its current money holdings, a household will choose to pay its transactions cost only if it lies at or below some maximum cost that it is willing to pay.

2.2.2 Family state vector and constraints

We will be brief in summarizing our shift to a period-by-period risk-sharing extended-family representation of the household side of our economy, and refer the reader to Khan and Thomas (2007) for further details. Nonetheless, we take some care in describing the timing and disbursement of households' wage and profit income into their individual bank accounts and the family's joint brokerage account. As noted above, all such incomes are paid nominally at the end of a period, so they cannot be used until the subsequent period. We are also, for now, agnostic about the fractions of these incomes paid into the households' individual bank accounts (λ_N and λ_Π) versus those paid into the family brokerage account, ($1-\lambda_N$ and $1-\lambda_\Pi$), allowing for differences in the disbursement of wage income versus profit income.

Let w_{t-1} and Π_{t-1} represent the real wage and real aggregate profits from date $t-1$. Consider a household entering date t as a member of time-since-active group j, having worked $n_{j,t-1}$

hours last period. At the start of the current period, this household receives a real payment of $[\lambda_N(w_{t-1}n_{j,t-1}) + \lambda_\Pi \Pi_{t-1}]\frac{P_{t-1}}{P_t}$ into its bank account, and has $[(1-\lambda_N)(w_{t-1}n_{j,t-1}) + (1-\lambda_\Pi)\Pi_{t-1}]\frac{P_{t-1}}{P_t}$ paid to the family brokerage account. We will represent the current-type-j household's wage earnings with which it ended the previous period as $e_{jt} \equiv w_{t-1}n_{j,t-1}$ in the problem that follows. The other wealth specific to a household is the real value of the money it chose to save in its bank account from the previous period. For a household currently of type j, let m_{jt} represent its real money savings as of the end of the previous period. These savings imply real balances of $m_{jt}\frac{P_{t-1}}{P_t}$ in the household's bank account at the start of this period. Thus, the extended family of all households has the following predetermined state variables as it begins date t: $[\{\theta_{jt}, m_{jt}, e_{jt}\}_{j=1}^{J}, \Pi_{t-1}, \chi_t]$. In fact, the final variable is superfluous. For convenience only, we use χ_t to summarize total real income deposited into the family brokerage account at the end of date $t-1$, which has real value $\chi_t \frac{P_{t-1}}{P_t}$ in the current period.

It is a convenient fiction to envision the family determining which households actively trade in the bond markets in any period and which do not. We may use this fiction here, because the family actually implements the allocation that arises when households individually implement their state-contingent lifetime utility maximization plans, given their access to complete insurance in the brokerage accounts. Given this alternative view of household decision making, within any particular time-since-active group of households, j, we can isolate the maximum transactions cost that the risk-sharing family will be willing to pay from the family bond account on behalf of a type j household to allow it to return to the family brokerage account and replenish or shed money balances. Given the common distribution from which transactions costs are drawn, and denoting the threshold cost associated with households of type j by ξ^T_{jt}, the fraction of group j households becoming active is $\alpha_{jt} = H(\xi^T_{jt})$. Alternatively, the family can directly choose the fractions of each group that will become active, α_{jt}, with knowledge of the associated threshold cost. Beyond these decisions, the family also selects the real balances with which all currently active households will leave the family account, m_{0t}. Finally, towards a convenient summary of how the distribution of household money balances evolves over time, we denote the fraction of all households entering the current period as members of time-since-active group j as θ_{jt}, and let J represent the maximum number of periods before which any currently active household will again be active. This implies that the distribution of households over time-since-last active will be tracked using the vector $[\theta_{1t}, ..., \theta_{Jt}]$, where, for $j = 2, ..., J$, $\theta_{jt} = \theta_{j-1,t-1}(1-\alpha_{j-1,t-1})$ and $\theta_{1t} = \sum_{j=1}^{J} \theta_{j-1,t-1}\alpha_{j-1,t-1}$. Note that members of group 1 this period were active last period, and thus made their consumption, labor supply and money savings decisions within that period as members of time-since-last active group 0.

In each period, the family's brokerage income, plus the money savings and bank account income returning to the family brokerage account with currently active households, together with any new balances injected by the monetary authority, must fully finance the total balances exiting the account with active households, as well as the associated total adjustment costs. We list this family budget constraint below, and will attach to it the multiplier Ω_t.

$$\frac{P_{t-1}}{P_t}\left[\chi_t + \sum_{j=1}^{J}\theta_{jt}\alpha_{jt}[m_{jt} + \lambda_N e_{jt} + \lambda_\Pi \Pi_{t-1}] + \mu_t \frac{M_{t-1}}{P_{t-1}}\right] \geq m_{0t}\sum_{j=1}^{J}\theta_{jt}\alpha_{jt} + \sum_{j=1}^{J}\theta_{jt}\varphi(\alpha_{jt}), \quad (15)$$

where $\varphi(\alpha_{jt}) \equiv \int_0^{H^{-1}(\alpha_{jt})} xg(x)dx$, and $\varphi'(\alpha_{jt}) = \xi^T(\alpha_{jt})$. The constraint determining the family's end-of-date t brokerage income follows, and will carry the multiplier H_t.

$$(1-\lambda_\Pi)\Pi_t + (1-\lambda_N)w_t\left[\sum_{j=1}^{J}\theta_{jt}\alpha_{jt}n_{0t} + \sum_{j=1}^{J-1}\theta_{jt}(1-\alpha_{jt})n_{jt}\right] \geq \chi_{t+1} \quad (16)$$

Next, the constraints below represent the evolution of the time-since-active distribution of households, which enter the family's problem with multipliers q_{0t} and $\{q_{jt}\}_{j=1}^{J-1}$.

$$\sum_{j=1}^{J}\theta_{jt}\alpha_{jt} \geq \theta_{1,t+1} \quad (17)$$

$$\theta_{jt}(1-\alpha_{jt}) \geq \theta_{j+1,t+1} \text{ for } j = 1, ..., J-1 \quad (18)$$

Because the family acts to maximize the weighted sum of households' utilities (with each household's period utility flow being $u(c, n)$), the bank account constraints and bank balance evolution facing individual households are also relevant to the family. Consumption for active and inactive households, respectively, will be constrained by the following.

$$m_{0t} - m_{1,t+1} \geq c_{0t} \quad (19)$$

$$\frac{P_{t-1}}{P_t}[m_{jt} + \lambda_N e_{jt} + \lambda_\Pi \Pi_{t-1}] - m_{j+1,t+1} \geq c_{jt} \text{ for } j = 1, ..., J-1 \quad (20)$$

Because these constraints on consumption bind in every period, we substitute them directly into the family's objective.

Two further points should be made regarding (19) and (20). First, m_{0t} is unique relative to every other m_{jt} variable, in that (a) it is a current choice variable rather than a predetermined state and (b) it represents current-date real balances. Second, for all j, the choices of $m_{j+1,t+1}$ are subject to non-negativity constraints. The final set of constraints links each household's current labor supply to the labor income it is entitled to at the start of the next period:

$$w_t n_{jt} \geq e_{j+1,t+1} \text{ for } j = 0, ..., J-1. \quad (21)$$

These constraints enter the family's problem with multipliers $\theta_{j+1,t+1}r_{jt}$, for $j = 0, ..., J-1$.

2.2.3 Family problem

The family takes as given, in each date, $\frac{M_{t-1}}{P_{t-1}}$, the total supply of real balances existing at the end of $t-1$, alongside current aggregate total factor productivity, z_t, and the currently growth rate of the aggregate money supply, μ_t, as well as the resulting prices and profits, w_t, P_t, Π_t, at all dates. Its choice variables are summarized by Ψ_t, where

$$\Psi_t \equiv \left[\{\alpha_{jt}\}_{j=1}^{J-1}, \{\theta_{j+1,t+1}\}_{j=0}^{J-1}, \{n_{jt}\}_{j=0}^{J-1}, \{e_{j+1,t+1}\}_{j=0}^{J-1}, m_{0t}, \{m_{j+1,t+1}\}_{j=0}^{J-1}, \chi_{t+1}\right].$$

Noting that the supply of aggregate real balances evolves as $\frac{M_t}{P_t} = \frac{M_{t-1}}{P_{t-1}}(1+\mu_t)\frac{P_{t-1}}{P_t}$, and given the constraints described above, we may express the family's optimization problem as follows.

$$\begin{aligned}
V\left(\{\theta_{jt}, m_{jt}, e_{jt}\}_{j=1}^{J}, \chi_t, \Pi_{t-1}; \frac{M_{t-1}}{P_{t-1}}, z_t, \mu_t\right) = \max_{\Psi_t}\Bigg[& u(m_{0t} - m_{1,t+1}, 1 - n_{0t}) \sum_{j=1}^{J} \theta_{jt}\alpha_{jt} \\
& + \sum_{j=1}^{J} \theta_{jt}(1-\alpha_{jt}) u\left(\frac{P_{t-1}}{P_t}\left[m_{jt} + \lambda_N e_{jt} + \lambda_\Pi \Pi_{t-1}\right] - m_{j+1,t+1}, 1 - n_{jt}\right) \\
& + \beta \int_{zx\mu} V\left(\{\theta_{j,t+1}, m_{j,t+1}, e_{j,t+1}\}_{j=1}^{J}, \chi_{t+1}, \Pi_t; \frac{M_t}{P_t}, z_{t+1}, \mu_{t+1}\right) F\left([z_t, \mu_t] d[z_{t+1}, \mu_{t+1}]\right) \\
& + \Omega_t \left[\frac{P_{t-1}}{P_t}\left(\chi_t + \mu_t \overline{\mathbf{m}}_t + \sum_{j=1}^{J} \theta_{jt}\alpha_{jt}[m_{jt} + \lambda_N e_{jt} + \lambda_\Pi \Pi_{t-1}]\right) - m_{0t}\sum_{j=1}^{J}\theta_{jt}\alpha_{jt} - \sum_{j=1}^{J}\theta_{jt}\varphi(\alpha_{jt})\right] \\
& + H_t \left[(1-\lambda_\Pi)\Pi_t + (1-\lambda_N) w_t \left(\sum_{j=1}^{J} \theta_{jt}\alpha_{jt} n_{0t} + \sum_{j=1}^{J-1} \theta_{jt}(1-\alpha_{jt}) n_{jt}\right) - \chi_{t+1}\right] \\
& + q_{0t}\left[\sum_{j=1}^{J} \theta_{jt}\alpha_{jt} - \theta_{1,t+1}\right] + \sum_{j=1}^{J-1} q_{jt}\Big[\theta_{jt}(1-\alpha_{jt}) - \theta_{j+1,t+1}\Big] \\
& + r_{0t}\sum_{j=1}^{J}\theta_{jt}\alpha_{jt}\Big[w_t n_{0t} - e_{1,t+1}\Big] + \sum_{j=1}^{J-1}\theta_{jt}(1-\alpha_{jt}) r_{jt}\Big[w_t n_{jt} - e_{j+1,t+1}\Big]\Bigg]
\end{aligned}$$

Appendix B lists the efficiency conditions characterizing the solution to this problem.

Finally, before specifying the conditions that connect households and firms in the equilibrium of our model, it is useful to define some household-side aggregates. Aggregate labor supply, consumption and output demand, and demand for real balances, respectively, are as listed below.

$$\begin{aligned}
N_t^S &= n_{0t}\sum_{j=1}^{J}\theta_{jt}\alpha_{jt} + \sum_{j=1}^{J-1}\theta_{jt}(1-\alpha_{jt}) n_{jt} \\
C_t &= c_{0t}\sum_{j=1}^{J}\theta_{jt}\alpha_{jt} + \sum_{j=1}^{J-1}\theta_{jt}(1-\alpha_{jt}) c_{jt} \\
Y_t^D &= C_t + \sum_{j=1}^{J}\theta_{jt}\varphi(\alpha_{jt}) \\
\left(\frac{M_t}{P_t}\right)^D &= \sum_{j=1}^{J}\theta_{jt}\alpha_{jt}[c_{0t} + m_{1,t+1}] + \sum_{j=1}^{J-1}\theta_{jt}(1-\alpha_{jt})[c_{jt} + m_{j+1,t+1}] + \sum_{j=1}^{J}\theta_{jt}\varphi(\alpha_{jt})
\end{aligned}$$

2.3 Market clearing

The equilibrium sequence of wages, w_t, aggregate price levels, P_t, and nominal interest rates, $i_t = \frac{\Omega_t}{\beta \boldsymbol{E}_t \Omega_{t+1} \frac{P_{t+1}}{P_t}} - 1$, ensure that the optimizing choices made by firms and households clear the markets for real balances (bonds), final output, and labor in each period. These sequences of market clearing conditions are as follow.

$$
\begin{aligned}
\frac{M_{t-1}}{P_{t-1}}(1+\mu_t)\frac{P_{t-1}}{P_t} &= \left(\frac{M_t}{P_t}\right)^D \\
Y_t^D &= Y_t \\
N_t^S &= N_t^D
\end{aligned}
$$

Finally, to conclude this section, we provide a complete list of the endogenous state variables in our economy. These are: $\{m_{jt}\}_{j=1}^{J-1}$, $\{\theta_{jt}, e_{jt}\}_{j=1}^{J}$, χ_t, Π_{t-1}, $\frac{M_{t-1}}{P_{t-1}}$, $\frac{P_{t-1}}{P_{t-2}}$, and Δ_{t-1}^N.

3 Two special cases of our model

Our two special case models may be described quite simply. First, when examining the textbook New Keynesian model, we drop the description of households from section 2.2, replacing the distribution of households there with instead a representative household that directly values real balances as a source of utility. Second, in the case of the flexible-price segmented asset markets model, we instead eliminate the description of production from 2.1, replacing the Calvo price-setting intermediate input firms there instead with a single perfectly competitive representative firm.

3.1 Pure New Keynesian case

To examine a pure New Keynesian environment delivering the perfect dichotomy discussed in section 1, we need only add a representative household to the description of the production side of the economy from section 2.1 above. To avoid studying the influence of monetary policy in a cashless economy, we assume real balances as a direct source of utility. While not explicit, this common device for sustaining money in the model may be viewed as a proxy reflecting a transactions-based demand for real balances.

In this special-case model, we determine the real wage, w_t, and the real price of output in each period, Ω_t, by appending the production-side of the economy with the optimization problem of an infinitely-lived representative household that derives utility from consumption, C, and from real balances $\frac{M}{P}$, and derives disutility from hours worked, N. Given the bond holdings and money with which it enters its initial period of life, (B_0, M_{-1}), and given its nominal profit income in each period, $P_t \Pi_t$, the household solves the following lifetime utility maximization problem:

$$
\max \ E_0 \sum_{t=0}^{\infty} \beta^t \left(U\left(C_t, N_t\right) + V\left(\frac{M_t}{P_t}\right)\right),
$$

subject to

$$P_tC_t + \frac{B_{t+1}}{1+i_t} + M_t \leq P_tw_tN_t + B_t + M_{t-1} + P_t\Pi_t.$$

Denoting the LaGrange multiplier for the household's problem by $\frac{\Omega_t}{P_t}$, we arrive at the following first-order conditions with respect to C_t, N_t, B_{t+1} and M_t, respectively.

$$\begin{aligned}
D_1U\left(C_t, N_t\right) &= \Omega_t \\
w_tD_1U\left(C_t, 1-N_t\right) &= -D_2U\left(C_t, N_t\right) \\
\frac{\Omega_t}{P_t}\frac{1}{1+i_t} &= \beta\mathbf{E}_t\frac{\Omega_{t+1}}{P_{t+1}} \\
\frac{\Omega_t}{P_t} &= DV\left(\frac{M_t}{P_t}\right)\frac{1}{P_t} + \beta\mathbf{E}_t\frac{\Omega_{t+1}}{P_{t+1}}
\end{aligned}$$

Notice that the money demand condition above may be re-written as: $DV\left(\frac{M_t}{P_t}\right) = \Omega_t - \beta\mathbf{E}_t\frac{P_t\Omega_{t+1}}{P_{t+1}} = \Omega_t\left(1 - \frac{1}{1+i_t}\right)$. Further, the absence of any fixed transactions costs implies that the household's only use for output is consumption; thus goods market clearing requires that $Y_t^D = C_t$ in equilibrium. This leaves us with the following four equations replacing the series of conditions from section 2.2.

$$D_1U\left(Y_t, N_t\right) = \Omega_t \tag{22}$$

$$w_t\Omega_t = -D_2U\left(Y_t, N_t\right) \tag{23}$$

$$\Omega_t\frac{1}{1+i_t} = \beta\mathbf{E}_t\frac{P_t\Omega_{t+1}}{P_{t+1}} \tag{24}$$

$$DV\left(\frac{M_t}{P_{t-1}}\frac{P_{t-1}}{P_t}\right) = D_1U\left(Y_t, N_t\right)\frac{i_t}{1+i_t} \tag{25}$$

The first equation determines the household's current valuation of output, and thus determines the stochastic discount factor, $\frac{\beta\Omega_{t+1}}{\Omega_t}$. The second equation determines household labor supply. The third equation is the household Euler equation, while the fourth determines the nominal interest rate, given the supply of real balances.

We may fully characterize the equilibrium of this model using 11 equations in the variables $\left(\frac{P_{0,t}}{P_{t-1}}, \Delta_t^0, \Delta_t^1, \frac{P_{t-1}}{P_{t-2}}, N_t, \Delta_{t-1}^N, \Omega_t, Y_t, w_t, i_t, \frac{M_{t-1}}{P_{t-1}}\right)$, where $\frac{P_{t-1}}{P_{t-2}}$, Δ_{t-1}^N, and $\frac{M_{t-1}}{P_{t-1}}$ are predetermined state variables. First, we have the four household efficiency conditions above, (22) - (25). Next, from the production-side of the model, we have the three equations determining optimal price setting, (8), (9) and (10), the two equations determining aggregate employment demand, (12) and (13), and the inflation equation, (11). Finally, the dynamic system is completed by appending these household and firm equations with a money supply rule of the form $\overline{M}_t = (1+\mu_t)\overline{M}_{t-1}$.

3.2 Pure segmented asset markets case

To examine the special case model with endogenous asset market segmentation, but no price-setting frictions, we maintain the description of households from section 2.2, and replace the description of production from section 2.1 with a representative final goods firm. This perfectly competitive firm hires labor from the households of section 2.2, and produces final output directly via the Cobb-Douglas production function, $Y = zN^{\nu}$. Maximizing its profit flows, period-by-period, the firm solves $\max_{N_t^D} z_t \left(N_t^D\right)^{\nu} - w_t N_t^D$, arriving at the following first order condition.

$$w_t = \nu z_t \left(N_t^D\right)^{\nu-1}. \tag{26}$$

Equation 26 determines the real wage in each period, given $N_t^D = N_t^S$, as the firm's marginal product of labor. Finally, as in our composite model, aggregate profits in each period are $\Pi_t = Y_t - w_t N_t^D$.

4 Functional forms and parameter values

Throughout the examples that follow, we set the length of a period to one quarter, and, wherever comparable, our parameter values are typical of those found in the literature (see, for example, Dotsey and King (2005)). First, we choose the household discount factor, β, at 0.9925 to imply a steady-state annual real interest rate of 3 percent, and also select the steady-state rate of money growth, μ^*, at 0.0075 to imply an average inflation rate of 3 percent per year. We assume that households' preferences over consumption, labor supply take the form $u(c,n) = \frac{c^{1-\sigma}}{1-\sigma} - A_L \frac{n^{1+\eta}}{1+\eta}$, where $\sigma = 2$. We select η at 2/3 to imply a labor supply elasticity of 1.5, and choose A_L to imply that households work on average one-third of their available time. When we examine the pure New Keynesian version of our model, we will also include an additively-separable preference for real balances in period utility, $v(m) = A_m \log(m)$, in order to sustain money in the model economy. This particular specification implies a unit elasticity of money demand with respect to its rental price, $\frac{i}{1+i}$, and we select A_m in each case to imply the same steady-state velocity as results in the full model economy to which it is being compared.

Intermediate inputs firms have a constant returns to scale production function ($\nu = 1$), and we set $\varepsilon = 11$. This relatively high elasticity across the inputs that firms supply implies that a firm able to adjust its price perfectly flexible in every period would change a 10 percent markup of price over marginal cost. Next, the degree of imposed firm-level price stickiness in these examples is a relatively conventional choice, with $\boldsymbol{\alpha} = \frac{1}{5}$.

In terms of the aspects of the model specific to the inclusion of market segmentation, we make the following choices in our baseline parameter set. We assume that 60 percent of all income is received by households in their bank accounts, setting $\lambda_N = \lambda_\Pi = 0.6$. Next we choose the form of the transactions cost distribution to imply that households have relatively little uncertainty about the timing of their account transfers on average, by selecting a sharply right-skewed beta distribution, $\beta(3, 1/3, 0.33)$. The upper support of this distribution, at 0.33, implies an annual velocity of 1.98, which corresponds well to the average velocity of a broad money aggregate in the U.S. ($M2$ less money market mutual funds) over the postwar period. See Khan and Thomas (2007) for a discussion of the monetary aggregate and the parameters identifying segmentation.

5 Results

5.1 Household portfolio management in steady state

Before we begin to examine the dynamics of our model, it may be useful to briefly consider the steady-state implications of the new aspects introduced with the inclusion of asset market segmentation and an explicit transactions role for money. Non-neutralities no longer arise solely from the price stickiness imposed on firms in this model. Now there is an additional source of non-neutrality arising from the fact that households generally do not exhaust their available money balances in current consumption expenditure, but instead retain inventories of money across time to avoid repeated transactions costs. Thus, total nominal expenditure can differ from the money supply, and changes in aggregate velocity will be allowed to take a role in determining how the economy responds to disturbances.

The top panel of Figure 2 shows households' spending rates and shopping-time distribution in the steady state of our baseline model. In this case, households' mean duration of inactivity is 7.6 quarters, and they anticipate a maximum duration of 8 quarters. Households are organized into the time-since-last active groupings described in section 2.2 above, and they are seen here as they enter into the shopping and production sub-period within each of these groups. Households in group 0 are currently active, having exchanged bonds for money at the beginning of the current period, households in group 1 are those that were last active one period ago, and so on. The dashed line shows the fraction of households shopping in each group. This curve is negatively sloped, because the fractions of households choosing to become active rise in the time-since-last-active, because willingness to pay transactions costs rises as real balances are eroded by the effects of inflation and previous expenditures. (The slope is fairly weak here, however, because the sharply right-skewed cost distribution implies that few households actually undertake portfolio adjustments before the mean duration of 7.6 quarters.)

The dash-dot line in the top panel represents the individual spending rate, or individual velocity, among households in each group. Because households that are currently active expect to wait the longest until their next money transfer, they choose to spend the smallest fraction of their current bank balance, roughly 30 percent, retaining the remaining 70 percent for coming periods over which they may remain isolated from the asset markets. Thereafter, as we look out across groups, we see these spending rates rise in time since last active. Note that aggregate velocity in this model is a weighted sum of the individual velocities within each time-since-active group (with weights determined by the population sizes and relative money holdings of these members of each of these groups), alongside a small term associated with total transactions costs. When a greater share of real balances is held by members of lowered numbered groups, aggregate velocity is pulled downward, and conversely. Thus, the shape of the individual velocity profile has the potential for generating important implications for the dynamics of our model. In the steady state, its implication is to keep households' consumption profiles essentially flat, as shown in the lower panel of the figure.

5.2 Dynamics following a persistent money growth shock

As we begin to consider our model's dynamics, we will first examine the responses to an unanticipated, persistent rise in the growth rate of the money supply. We begin by briefly reviewing the responses in the pure New Keynesian model. Thereafter, we turn to examine how our full model differs relative to this reference and briefly consider why these differences arise, making use of the observations in the section above.

5.2.1 Pure New Keynesian case

Figure 3 shows the dynamics of aggregate output, employment, wages, interest rates and inflation following a persistent rise in the money growth rate in the pure New Keynesian model, absent any financial frictions.[6] In contrast to Milton Friedman's observations on responses in an actual economy, we see here that nominal interest rate in the lower left panel actually rises (very slightly) at the impact of this shock. This absence of a liquidity effect in the model arises directly from the Fisher relation. Specifically, the real interest rate does not fall sufficiently to offset the rise in expected inflation, so that the nominal rate fails to decline with an expansionary open market operation.

The rise in expected inflation seen in the upper left panel of the figure is intimately linked to the short-run Phillips curve that underlies the appeal of this simple model. At the impact of a surprise money injection, four-fifths of firms are unable to respond with a rise in their nominal prices. Instead, these firms see their relative prices eroded, as firms that do have a price-setting opportunity raise their nominal prices in response to the increased money supply, thus driving up the aggregate price level. Given that the large group of firms with fixed nominal prices are forced to satisfy demand at those prices, we see a rise in employment and output in the upper middle panel. However, in future periods, as a growing fraction of firms have the ability to reset their prices, the effects of the monetary shock will shift away from quantities and into the aggregate price level.

5.2.2 Full model

In Figure 4, we examine the responses to the same money growth shock in the baseline version of our full model, where the New Keynesian production aspects described above are combined with an explicit transactions role for money and a segmented asset markets environment yielding household inventories of money. Viewed as a whole, this figure shows broadly similar responses relative to those in Figure 3. Nonetheless, there are certainly differences to be noted. The most striking one, perhaps, is the fact that the nominal interest rate in the lower left panel of this figure falls in response to the shock, and indeed remains below average for 8 quarters. From the Fisher equation, we know that this response must be explained by the responses in inflation and the real interest rate. The former shows relatively minor differences in comparison with the pure New

[6] For comparability, we set the persistence of this shock at 0.57 as in Chari, Kehoe and McGrattan (2000).

Keynesian model. Looking closely, we see slightly less inflation at the impact of the shock, and that the inflation response dies off a bit more slowly over time. In the real interest rate response, by contrast, the differences are pronounced. Beyond an obvious nonmonotonicity absent in Figure 3 (to which we will return below), notice that the real rate falls here by roughly 200 annualized basis points at the shock's impact (while it fell 140 points in Figure 3). Thus, our model has a liquidity effect both because asset market segmentation reinforces price stickiness somewhat, and because it sharply exacerbates the decline in the real interest rate.

Turning to the figure's upper middle panel, notice that the impact-date responses in output, employment and the real wage are all smaller in our full model relative to the previous figure. For example, output rises by only about 0.5 percent here, while it rose roughly 1 percent in the pure New Keynesian model. Moreover, all of these series now decay more slowly than they did in Figure 3. Indeed, we actually see a hump shape in the employment response, and there is a slight secondary upturn in output in period 8. Thus, in addition to the liquidity effect, this model generally dampens initial responses and increases persistence in output, employment, wages and inflation. A final distinction is that, here, aggregate velocity actually falls in response to the shock, in keeping with Friedman's observations.

To explain these differences arising with the introduction of segmented asset markets in our full model, we begin with the most essential observation regarding any model with asset market segmentation. *The only direct participants in any open market operation are households that are currently active - trading bonds for money.* This implies that any money injection does not fall over households evenly, but instead goes solely to those households currently active in the bond markets. Thus, households active at the date of the shock carry unusually high real balances, and they experience a rise in their permanent income that implies unusually high current consumption. This may be seen in the second panel of Figure 5, where the consumption of active households rises roughly 1.25 percent above its steady-state level at the impact of the shock.

Next, we remind the reader that the real interest rate in any segmented markets model is determined by the marginal utilities of *active* households in adjacent periods, as these are the only households able to exchange bonds for consumption, and thus the only ones relevant for pricing bonds. Given their unusually high consumption, we see in Figure 5 that active households in date 1 have a below-average marginal utility of consumption. Indeed, the greater rise in these households' consumption relative to that of the representative household in the Pure New Keynesian model causes their marginal utility to fall roughly 1 percent more relative to steady-state than it did in Figure 3. This fully explains the more substantial drop in our model's real interest rate, and thus (given marginally less initial inflation) how it succeeds in delivering a liquidity effect.

Given the preceding discussion, we can easily understand why aggregate velocity drops in our model in response to the shock (in contrast to the rise seen in the pure New Keynesian model). This too is a direct implication of how real balances are redistributed across households at the impact of the money shock. We noted above in section 5.1 that households currently active in asset markets have the lowest individual spending rates among all households, so they lower aggregate velocity.

At the shock's impact, these households carry an unusually high fraction of the economy's real balances, and hence pull aggregate velocity below its mean. This explains why our aggregate price level in the lower right panel of Figure 4 is initially slightly more sluggish relative to that in Figure 3. Moreover, velocity is dragged further downward in several subsequent dates, as these recently active households continue to have relatively low velocities, while new groups of households coming active also leave the bond accounts with above-average real balances (as we will explain below).

Changes in the distribution of money across households also explain the lesser initial responses in output and employment (and thus wages) in our model, since we are observing a demand-driven expansion. These differences again may be traced to increased inventories of real balances among active households. Because those becoming active at the date of the shock retain much of the extra money to finance their spending in later dates, the rise in aggregate spending is gradual relative to the rise in the money supply. This behavior directly implies that there is a lesser rise in the real balances that are circulating in the goods market, and thus a lesser rise in total production. However, the altered shape of the subsequent dates' employment and output responses still remain to be explained; for this, we must look more carefully at the series plotted in Figure 5.

We noted above that households active in periods after the impact of the shock also carry above-average real balances. This is partly explained by the persistence of the shock, of course. However, it is also partly attributable to small changes in activity rates in response to the shock. When the shock hits, households have an incentive to return to bond market early in order to obtain raised real balances (at the expense of inactive households). This implies a rise in the overall participation rate (the total fraction of all households becoming active) at the date of the shock. However, that very rise implies that, in date 2, an usually large fraction of households have just adjusted their balances and thus are relatively unwilling to suffer transactions costs to return to the bond markets again. As a result, the overall participation rate begins to fall despite the incentives arising from the fact that money growth is still unusually high, and indeed falls below average in dates 3 through 8.

The persistence of the shock, together with falling participation rates, imply that households active after date 1 also achieve unusually high consumption, as is seen in the middle panel of Figure 5. However, the wealth effect that goes to these households is not as large as that for the initially active households, and this has an important influence on the labor they supply. Initially, despite the rise in the real wage, the wealth effect is dominant in determining active households' labor supply, so that n_0 falls slightly relative to steady state. This restrains the early response in aggregate employment. However, as the rise in lifetime wealth grows smaller for households active in later dates, while the real wage remains high, we see n_0 rising. By date 3, labor supply among active households rises above average, and it continues to rise in the subsequent 5 periods. Combining these observations with the fact that each household's hours worked over dates of inactivity are typically quite close to that selected at its last trading date, we have an explanation for the humped shape in our model's aggregate employment response.

By contrast to the discussion above, Figure 4 shows that output monotonically declines over the periods while employment is rising. This would be impossible in a single-factor production setting with a representative firm. However, given the New Keynesian production environment, it matters here how production is spread across firms. Because production is largely concentrated among the 4/5 of firms unable to alter their nominal prices at the date of the shock, a rise in employment is quite effective in raising aggregate output. However, in subsequent dates, as more firms are able to adjust their prices, the low-relative-price firms that undertake most of the economy's production become a smaller subset of firms. This implies a greater dispersion in production, which reduces the aggregate marginal product of labor, so that output begins to fall despite rising labor input.

There are various echo effects in our model's interest rate and quantity responses around date 8, when the bulk of households that were active at the impact of the shock are returning again to the brokerage accounts. We have noted that, at the shock's impact, active households saw a rise in their real balances and corresponding a rise in their lifetime wealth. At that initial date, given their preference for smooth consumption profiles, these households insist on finding a way to spread the benefits of their windfall over their entire lifetime - beyond simply achieving high consumption in their first shopping cycle. Indeed, if there was no way to do this, they would simply have put much of the extra money into bonds (in effect refusing the injection). Instead, however, these households are able to again carry above-average real balances after they revisit the brokerage accounts, so that they experience raised consumption throughout their next shopping cycle, and so forth. This is achieved with the second fall in the nominal interest rate occurring in date 8, and immediately reversed in date 9. While somewhat subtle in Figure 5, this allows consumption of active households in date 8 that is slightly higher than the average of the consumptions of households active in the two adjacent periods, which causes the temporary drop again in the real rate in date 8. Finally, the secondary upturn of the aggregate output series in Figure 4 is explained by the endogenous return to the brokerage accounts of many of the unusually large number of households active at the date of the shock, and their combined influence in raising aggregate demand. As a whole, it is clear from Figures 4 and 5 that the distribution of money balances in our economy plays a substantial role in shaping the dynamics of not only interest rates, but also aggregate quantities.

5.3 Responses to a productivity shock under an interest rate rule

We next turn to consider a supply-driven expansion to see whether the distribution of money balances will continue to play a role in shaping aggregate quantities and relative prices when the monetary authority adopts an interest rate rule. Again beginning with the pure New Keynesian model as a reference, we examine our economy's dynamics following a persistent positive shock to productivity in the presence of an active Taylor rule placing 1.5 weight on inflation deviations, and 0.5 weight on output deviations, relative to steady state.

5.3.1 Pure New Keynesian case

Given our discussion in section 1 regarding the perfect dichotomy that arises in the pure New Keynesian model, we know that the interest rate rule itself completely dictates the response to the productivity shock. We need know nothing about average money demand in this economy to explain the dynamics following the shock. For reference, the aggregate responses of this model are reviewed in Figure 6.

In response to the productivity shock, firms produce more output, and pay a higher real wage. This, in turn, generates raised demand for real balances to accompany the representative household's raised consumption of both goods and leisure. As a result, inflation is pushed below trend. Given the Taylor rule, the monetary authority's response to these events is to reduce the nominal rate in sufficient degree as to lower the real rate, thereby dampening the fall in inflation. Finally, given the relatively large group of firms that cannot lower their prices to sell more output in response to the shock, we see that aggregate output does not rise by the full one-percent otherwise implied by the shock. Correspondingly, with this large subset of firms using less labor to produce at an unchanged level, total employment actually falls in response to the shock, and recovers only as a sufficient fraction of firms are able to adjust their prices.

Given this summary of the dynamics in our special case model with a perfect Keynesian dichotomy, we now turn to two final subsections of results to understand how well this convenient model proxies for the dynamics arising in our full model. In particular, we seek to answer the following question: *Are the predictions in this figure robust to richer settings where money is actually allowed the possibility of some role in the economy's dynamics?*

5.3.2 Full model: Baseline case

Figure 7 presents the responses following the same productivity shock as above, under the same interest rate rule, in the baseline case of our full model. As in our comparison of the two models' responses to the persistent money growth shock in section 5.2.2, the responses to the productivity shock here are broadly similar in shape and magnitude to those in the pure New Keynesian model. However, despite the presence of the interest rate rule, there are again clear differences in the details as we compare the series in this figure to those in Figure 6.

Here again, the monetary authority's response to the shock, given the Taylor rule, is to lower the nominal interest rate. Given the reduced return on bonds, the fraction of households actively trading bonds for money rises. However, just as we saw in Figure 5, this participation rate falls over subsequent dates, given relatively low activity rates among an unusually high fraction of households recently active. Thereafter, aggregate participation rises once again when the bulk of the initial surge in active households returns to the bond markets.

The initial disturbance to the distribution of money holdings across households and the echoes it generates through the distributions in subsequent periods have real consequences distinguishing our model's responses in aggregate quantities and prices. First, the adjustment implied for the nominal interest rate with the fall in inflation at the shock's impact is roughly 40 annualized basis points greater in our model. Correspondingly, the initial decline in the real rate is also considerably greater; however, there is very little persistence in this series relative to its pure New Keynesian counterpart. While the inflation responses are fairly close across the two models, ours falls a few basis points further initially, despite the larger interest rate movements designed to cushion it.

Another notable feature in the lower panels of this figure is that the monetary authority in our economy achieves a smooth inflation response following the shock only by deviating from the smooth returns in interest rates that would be implied in the economy without asset market segmentation. In those dates where the bulk of households active in the initial period of the shock are returning to the bond markets, we see echo effects in both interest rate series. These echoes are generated by permanent income considerations, and are primarily associated with the households active at the date of the productivity shock, just as we discussed above in section 5.2.2. However, they are more pronounced in this case. Here, the policy intervention exacerbates the 8-period cycles in the total participation rate discussed above. The rule itself, by dictating a money injection at the date of the shock, causes an initial surge of active households. This surge disperses far more slowly relative to what was shown in Figure 5, because of the endogeneity of the interest rate rule. Each time the bulk of initially active households returns to the bond markets, there is further interest rate intervention, again stimulating participation rates, and, in turn, output and wages. This explains the small repeated peaks in these series.[7]

The limited participation in our model delivers several other notable changes beyond the echo effects in output, employment and the real wage. Note that output and the wage are initially dampened in their responses relative to those in Figure 6, while employment falls by more. This is because the real balances in circulation in our economy actually affect what is produced and sold. While the fraction of households returning to the bond markets does rise at the impact of the shock, there is nonetheless a sizeable group of households constrained by their existing money balances. Thus (as we saw above in the responses to the money growth shock), aggregate expenditures are initially restrained, and hence so is output. Given these additional frictions restraining output (and real wages), it is not surprising that aggregate employment falls by more in our model than was necessary in the pure New Keynesian case.

Finally, the most striking qualitative difference between the responses in the full model and those in the standard New Keynesian model must certainly be the humped shape in the output (and consequently real wage) series. Its explanation lies again in the transactions role for money and limited participation in our model, and follows immediately from the explanation of why output is initially more restrained here. Over time, as more households become active in the bond markets, they return to the goods market with the additional liquidity needed to purchase the extra output made possible by the productivity shock. Thus, the initial restraint on total production begins to give way, and the series continues to rise for 6 or 7 quarters after the productivity shock hits, in contrast to the monotone response we saw in Figure 6.

Given the changes we have seen here, it appears that the convenience of the dichotomy implied

[7] It is worth noting that the participation rate response is far smaller, initially negative, and smoother following the productivity shock under a constant money growth rate rule. In that case, the echoes seen throughout Figure 7 disappear, giving way to smooth responses in output and employment, as well as the real wage. However, this still does not remove the hump in output. With or without the Taylor rule, this hump is a natural consequence of the transactions role for money and limited participation in our model, as will be explained below.

by the pure New Keynesian model may come at some substantive cost. In particular, it is arguable that the predictions of the pure model are not robust to at least one setting where money has an explicit role in the economy, and thus can no longer be quarantined from real variables. One natural question one might ask, however, is: *Does the particular form of the limited participation (or other aspects influencing the steady state distribution of money and its dynamics) matter in the response under an interest rate rule - or do we have simply a modified dichotomy in the presence of any segmentation?*

5.3.3 Full model: Alternative high velocity case

To answer the question above, we briefly consider an alternative case of our full model. Here, we weaken the financial frictions in our economy by both lowering the upper support on the distribution and imposing a uniform CDF (which implies both greater uncertainty among households regarding their transactions costs and a greater fraction of households drawing costs from the lower end of the support in each period). While steady-state output, employment, hours and real wage are essentially unaffected (e.g., hours falls from 0.33 to 0.31), the transactions costs paid relative to GDP fall from 10 percent to only 3 percent in this alternative setting. Moreover, participation is less limited. Households return to the bond markets far more frequently; their mean duration of inactivity is now only 3.1 quarters (versus 7.6 above), so that they now carry less money on average. As a result, the average annual velocity of money falls from roughly 2.5 to 2.

Figure 8 presents the responses in this alternative case of our model, following the same productivity shock (under the same interest rate rule) as in Figure 7. In this case, we see that the echoes so apparent in the Figure 7 responses have largely disappeared. Now, with greater randomness in transactions costs, it is no longer true that the bulk of households becoming active at the date of the shock together remain inactive for the next 7 periods. Instead, this bulge in the time-since-last-active distribution dissipates fairly quickly over subsequent periods, which removes the abrupt echo effect.

Another important change in this figure is that the initial rise in output is more brisk. Given lesser financial market frictions, more households return to the bond market early in response to the interest rate intervention in this case of the model than did in the baseline case. Thus, more real balances go into circulation faster here, allowing aggregate expenditure and thus output to rise more quickly. This, in turn, implies that aggregate employment falls by less. Finally, it also implies a less pronounced hump shape in the overall responses of output and wages.

Given the quantitative and qualitative differences we have seen across these final two figures, it is clear that, irrespective of an interest rate rule, our model provides a setting where money demand does matter for the economy's aggregate response to shocks, and, moreover, the details of the distribution of money demand across households matter as well. By combining an explicit transactions role for money and endogenous limited participation, we have broken the Keynesian dichotomy.

6 Two suggestive policy experiments

In the textbook New Keynesian model, at least near zero inflation, it is possible to view the optimal monetary policy response to productivity shocks as ensuring that *both* the average markup of price over marginal cost and the price level remain constant (see, for example, the discussion in Goodfriend and King (2001). While a complete derivation of optimal policy using the recursive methods of Khan, King and Wolman (2003) would be feasible in our economy, it is beyond the scope of the current paper. In such an analysis, the money stock or interest rate rule becomes a means of supporting a "second best" optimal allocation, rather than being specified *ex ante* as above.

In this section, we use a similar approach to study two rules requiring that monetary policy respond to productivity shocks so as to accomplish a specific objective for an endogenous variable.[8] We first require that monetary policy fully stabilize the path of the price level. Next, we instead require that monetary policy stabilize the average markup at its steady-state level. These experiments may be informative as to the implications of our environment for desirable monetary policy.

Both experiments are conducted around a zero-inflation steady-state, where the simple linear model from section 1 is an analytical log-linear approximation to our pure New Keynesian model. Continuing to consider the model without asset market segmentation, it is straightforward to similarly produce results for the behavior of output and labor under flexible prices.[9] These are:

$$
\begin{aligned}
\widehat{n} &= \frac{(1-\sigma)}{[\eta+\sigma\nu+(1-\nu)]}\widehat{z} \\
\widehat{c} &= \frac{(1+\eta)}{[\eta+\sigma\nu+(1-\nu)]}\widehat{z}
\end{aligned}
$$

Recalling our standard parameter choices ($\sigma = 2$, $\eta = 0.67$, and $\nu = 1$), employment in the flexible-price model declines in response to a productivity increase with an elasticity of -0.375 (as income effects dominate substitution effects), while consumption rises with an elasticity of 0.625. Further, the real interest rate is linked to the growth rate of consumption by $r_t - r = \sigma[E_t\widehat{c}_{t+1} - \widehat{c}_t] = \sigma(\rho-1)\widehat{c}_t$, with $\rho < 1$, so that it should decline with a positive but ultimately temporary productivity shock. In the pure New Keynesian model, both of the policies that we consider exactly stabilize inflation to produce these neutral responses. Thus, the nominal interest

[8] These two examples are easily implemented, as may be described in the context of the simple model from section 1. There, the (log) price level is added as $P_t = P_{t-1} + \pi_t$, and the model can be solved under an interest rate or money stock rule by adding a policy equation. Alternatively, the model can be solved by adding the equation $P_t = P_{t-1}$, which specifies a path for the price level. The resulting money or interest rate paths are then interpretable as the paths for the policy variables required to implement the stable price level.

[9] Recall that the period utility function over consumption and labor takes the form $u(c,n) = \frac{c^{1-\sigma}}{1-\sigma} - A_L\frac{n^{1+\eta}}{1+\eta}$. Under a constant markup and without financial frictions, the household's efficiency conditions imply that $wc^{-\sigma} = A_L(1+\eta)n^{\eta}$, where w is the real wage, and the firm's cost minimization implies that $\mu = w/[z\nu n^{\nu-1}]$, where μ is the level of the markup. Imposing goods-market clearing, $c = zn^{\nu}$, and log-linearizing, we arrive at: $\widehat{z}+(\nu-1)\widehat{n}-\sigma\widehat{c} = \eta\widehat{n}$ and $\widehat{c} = \widehat{z}+\nu\widehat{n}$. The solution to these expressions yields the results in the text.

rate should exactly mirror the real rate. With this as background, we now turn to examine the behavior following a productivity shock in our complete model around zero inflation under three policies: (a) the Taylor rule; (b) inflation targeting; and (c) markup stabilization.

6.1 Taylor rule

Figure 9 shows the behavior of economic activity in our model under the Taylor rule. In response to the persistent positive productivity shock, consumption/output rises, labor declines, and the real interest rate falls. Thus, all three responses are consistent with the responses described above. However, the Taylor rule also introduces some negative inflation in response to the productivity shock and some related deviations of output and employment from the flexible price paths. That is, consumption/output should mirror the smooth response of productivity with a coefficient of 0.625, but it responds somewhat less positively to the shock. Symmetrically, labor should mirror the smooth response of productivity with a coefficient of -0.375, but it responds more negatively to the shock. Overall, the Taylor rule smooths output relative to the neutral solution, but destabilizes labor.

6.2 Price level stabilization and monetary dynamics

Figure 10 shows the dynamic response of our model economy under the inflation-targeting rule. In the current context of a zero inflation steady state, this corresponds to full price level stabilization. In this case, we see that the solutions for consumption and labor input correspond more closely to their reference neutral values, with consumption rising on impact by about 0.6 percent and labor dropping by about 0.4 percent. While the monetary dynamics do introduce some small ripples around the targeted paths for real activity, the overall shapes are more closely in accord with the neutral solution.

Interestingly and importantly, there appears to be a greater departure with respect to the path of the nominal interest rate (identical to the real rate because inflation is always zero). Rather than following the neutral path, which implies

$$i_t - i = (\rho - 1)\frac{\sigma(1+\eta)}{[\eta + \sigma\nu + (1-\nu)]}\widehat{z}_t,$$

the interest rate dynamics required to support the inflation target are more complicated in our model economy. The initial interval of low rates is much shorter than would be the case in the standard model. This suggests that, even if the introduction of portfolio adjustment frictions does not imply substantial changes for the nature of desirable monetary policy, it may nonetheless imply important changes to the way in which this policy is implemented - here, in the interest rate settings necessary to bring about the desired real outcomes.

6.3 Markup stabilization and monetary dynamics

We need not examine any further figures to study the consequences of stabilizing the markup in our complete model, for a simple reason. Around zero inflation, the dynamics of Calvo price-setting can be described as

$$\pi_t = \beta E_t \pi_{t+1} + \gamma \widehat{\psi}_t$$

where ψ_t is real marginal cost. In our economy, as in many others, real marginal cost is the reciprocal of the average markup. Hence, the implications of a markup stabilization policy and that of stabilizing the inflation rate are identical in our model.

7 Concluding remarks

In the sections above, we have constructed a model economy where two frictions frequently suggested as individual rationalizations for monetary non-neutrality are simultaneously introduced, namely frictions leading to infrequent price and portfolio adjustments. We have stressed that the latter of these frictions leads the distribution of money balances to become part of the economy's state vector, and that this distribution is linked to both the history of monetary injections by the central bank and the nature of optimal portfolio adjustments among households. We have also stressed how this, in turn, leads to a breakdown of the Keynesian dichotomy, so that it is no longer possible to determine aggregate demand without reference to monetary variables. Viewed another way, in our economy, there is no longer a separate IS curve of the old or new Keynesian form. Examining the nature of dynamic responses arising in our economy, we have found that, to our eyes, they are more interesting and plausible than those arising in the nested special cases that include only one of our two frictions.

There are two critical directions of work as we continue our work on this project. First, it will be important to discipline the introduction of the portfolio adjustment friction by reference to its macroeconomic and microeconomic implications. That is, the nature of portfolio adjustment costs must be structured in a manner that leads our model to deliver adjustment dynamics consistent with those in actual economies. On the microeconomic side, this calls for an understanding of the frequency and likelihood of individual portfolio adjustments within the model and in actual economies. While there is a regrettably small literature on the microeconomics of monetary adjustment, it is growing (for example, with the important study of households' frequency of trade in high-yield, risk-bearing assets undertaken by Vissing-Jørgensen (2002)). On the macroeconomic side, we will need to explore our model's consistency with the large literature on money demand regressions. (The recent studies of Ball (2001,2007) on long-run and short-run money demand seem appropriate descriptions of the state of this research.) Some researchers, among them Goodfriend (1985), are skeptical that transactions costs are large enough to produce an important difference between short-run and long-run money demand, as our theory implies. As such, an essential test of our model will be whether it is successful in simultaneously reconciling the micro and macro data. Second, we have highlighted that our model economy delivers a breakdown of the Keynesian dichotomy, so that it is no longer true that the IS curve does not depend on monetary variables. It will be important to determine the magnitude of the departures that our model economy predicts.

References

[1] Alvarez, Fernando, Andrew Atkeson, and Chris Edmond (2003) "On the Sluggish Response of Prices to Money in an Inventory-Theoretic Model of Money Demand" NBER Working Paper no. 10016.

[2] Alvarez, Fernando, Andrew Atkeson, and Patrick Kehoe (2002) "Money, Interest Rates and Exchange Rates with Endogenously Segmented Markets", *Journal of Political Economy* 110(1), pages 73 - 112.

[3] Alvarez, Fernando, Robert E. Lucas, Jr., and Warren Weber (2001) 'Interest Rates and Inflation' *American Economic Review* 91(2), pages 219 - 225.

[4] Ball, Laurence (2001), "Another look at long-run money demand," *Journal of Monetary Economics*, Volume 47, Issue 1, pages 31-44.

[5] Ball, Laurence (2007), "Short-run money demand," Johns Hopkins University Working paper under revision for *Journal of Monetary Economics*.

[6] Calvo, Guillermo A. (1983), "Staggered prices in a utility maximizing framework", *Journal of Monetary Economics* 12, pages 383-98.

[7] Chari, V.V., Patrick J. Kehoe and Ellen R. McGrattan (2000), "Sticky Price Models of the Business Cycle: Can the Contract Multiplier Solve the Persistence Problem?" *Econometrica* 68(5), pages 1151-79.

[8] Christiano, Lawrence J., Martin Eichenbaum and Charles L. Evans (2005), "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy", Journal of Political Economy, *Journal of Political Economy*, vol. 113, no. 1, pages 1-45.

[9] Clarida, Richard, Jordi Galí and Mark Gertler (1999), "The Science of Monetary Policy: A New Keynesian Perspective," *Journal of Economic Literature*, Vol. 37, No. 4, pages 1661-1707.

[10] Dotsey, Michael, Robert G. King, and Alexander L. Wolman (1999) "State-Dependent Pricing and the General Equilibrium Dynamics of Money and Output", *Quarterly Journal of Economics* 114, pages 655-90.

[11] Dupor, Bill, Jing Han and Yi-Chan Tsai (2007), "What do technology shocks tell us about the New Keynesian Paradigm," presented at the NBER Economic Fluctuations and Growth Conference, Cambridge, July 2007.

[12] Gali, Jordi (1992), "How Well Does the IS-LM Model Fit Postwar U.S. Data?" *The Quarterly Journal of Economics*, Vol. 107, No. 2, pages 709-38.

[13] Goodfriend, Marvin S. (1985), "Reinterpreting Money Demand Regressions." Carnegie-Rochester Conference Series on Public Policy 22, pages 207-42.

[14] Goodfriend, Marvin and King, Robert G. (2001), "The Case for Price Stability", NBER Working Paper No. W8423.

[15] Goodfriend, Marvin S. and Bennett T. McCallum, ""

[16] Keen, Benjamin D. (2004), "In Search of the Liquidity Effect in a Modern Monetary Model," *Journal of Monetary Economic*s, 51(7), pages 1467-94.

[17] Kerr, William and Robert G. King (1996), "Limits on Interest Rules in the IS Model," *Federal Reserve Bank of Richmond Quarterly Review*, vol. 82, no. 2, pages 47-75.

[18] Khan, Aubhik, Robert G. King and Alexander L. Wolman (2003), "Optimal Monetary Policy," *Review of Economic Studies*, 70 (4), pages 825-60.

[19] Khan, Aubhik and Julia K. Thomas (2007) "Inflation and Interest Rates with Endogenous Market Segmentation", Working Paper, March 2007.

[20] King, Robert G. and Julia K. Thomas (2006), "Partial Adjustment without Apology," *International Economic Review* 47 (3), pages 779-809.

[21] McCallum, Bennett T. and Edward M. Nelson (1999), "An Optimizing IS-LM Specification for Monetary Policy and Business Cycle Analysis," *Journal of Money, Credit, and Banking* Vol. 31(3), pages 296–316.

[22] Smets, Frank and Raf Wouters (2003), "An estimated dynamic stochastic general equilibrium model of the Euro area," *Journal of the European Economic Association*, 1(5), pages 1123-75.

[23] Vissing-Jørgensen, Annette (2002), "Towards an Explanation of Household Portfolio Choice Heterogeneity: Nonfinancial Income and Participation Cost Structures" NBER working paper 8884.

[24] Woodford, Michael (2007), "How important is money in the conduct of monetary policy?", NBER working paper 13335.

FIGURE 1: Responses to a Productivity Shock under a Taylor Rule in Two New Keynesian Models

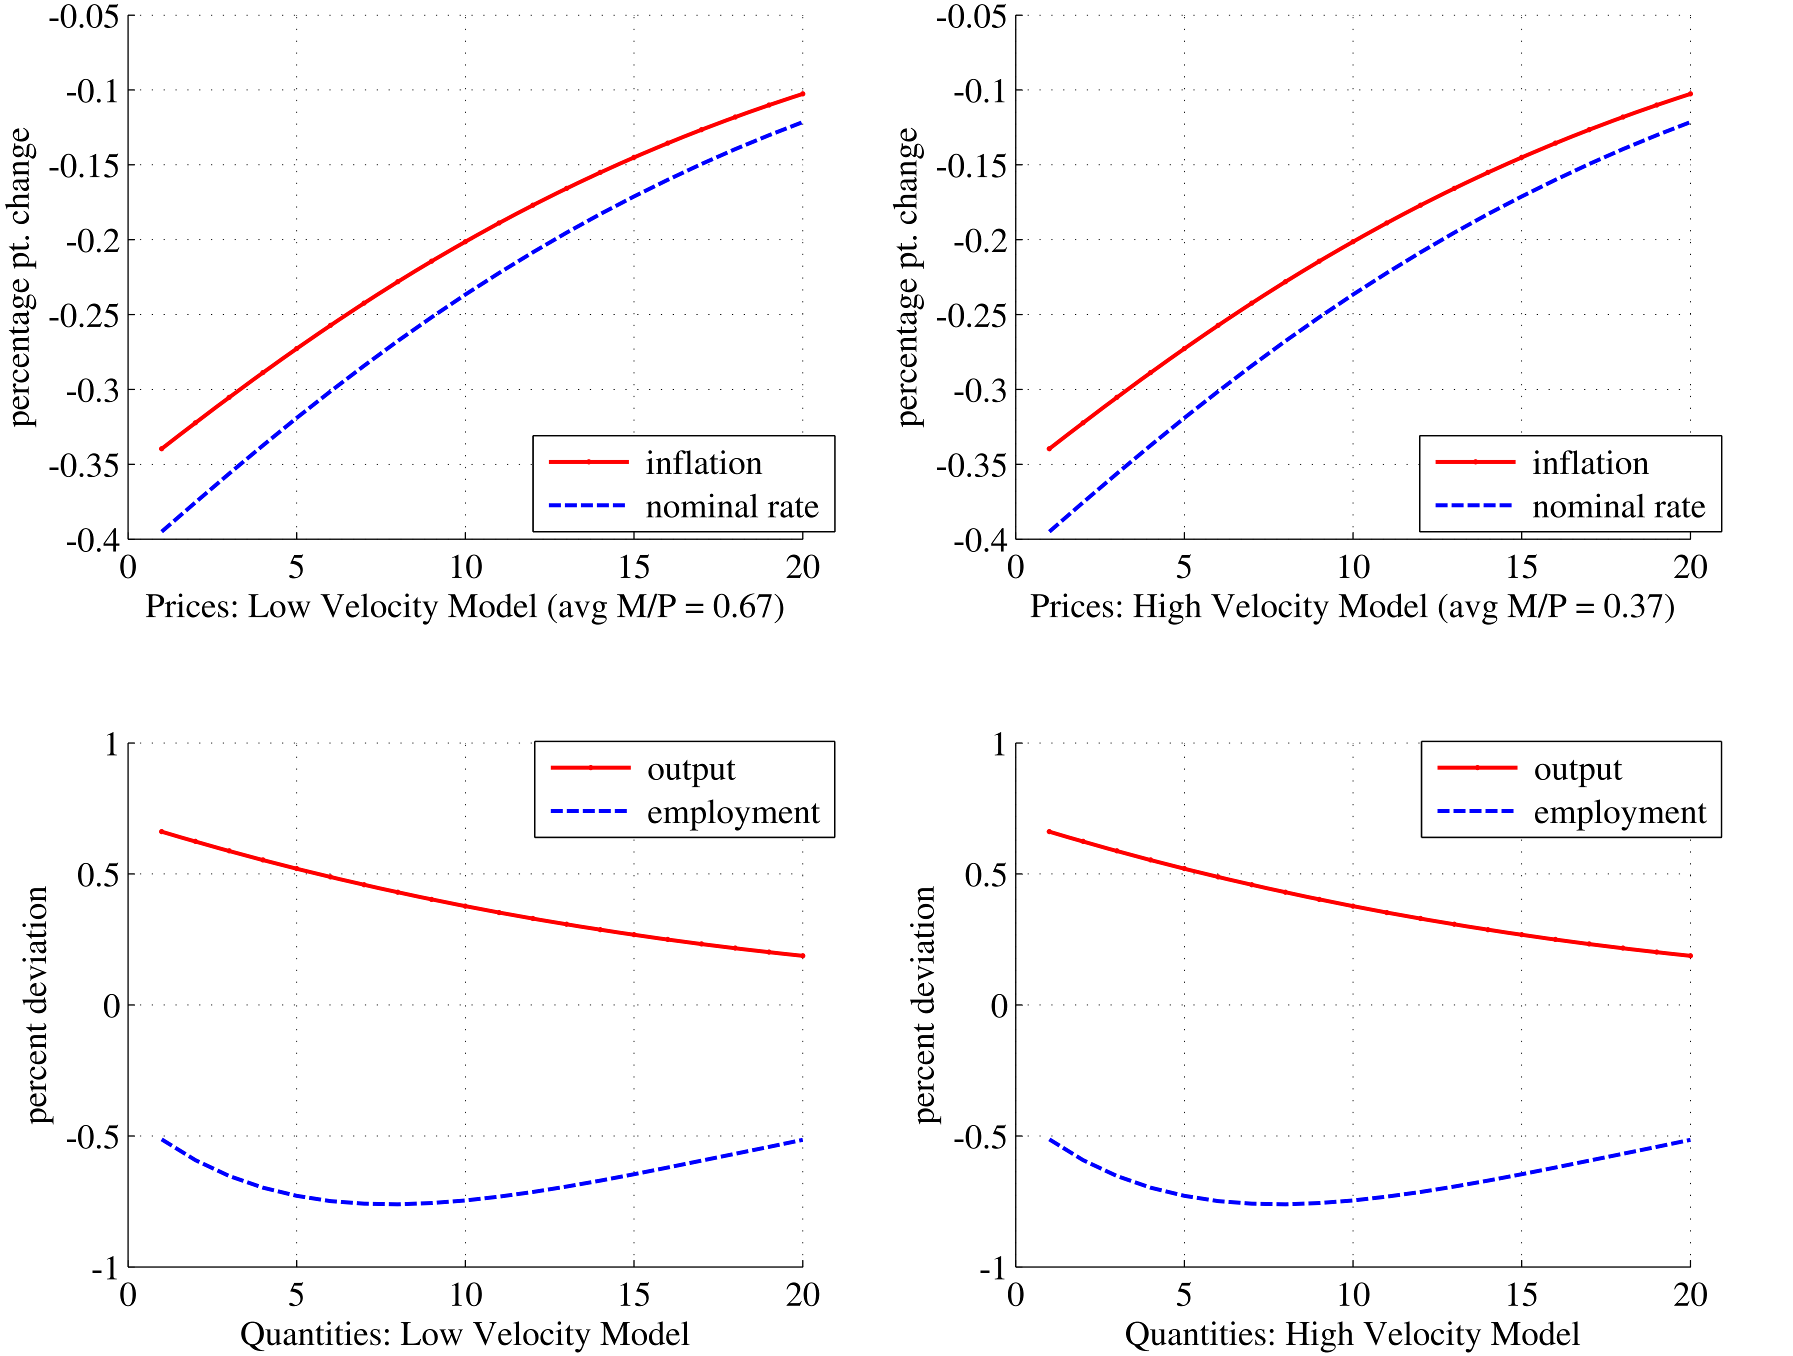

FIGURE 2: Baseline case of Full Model - Households in the steady state

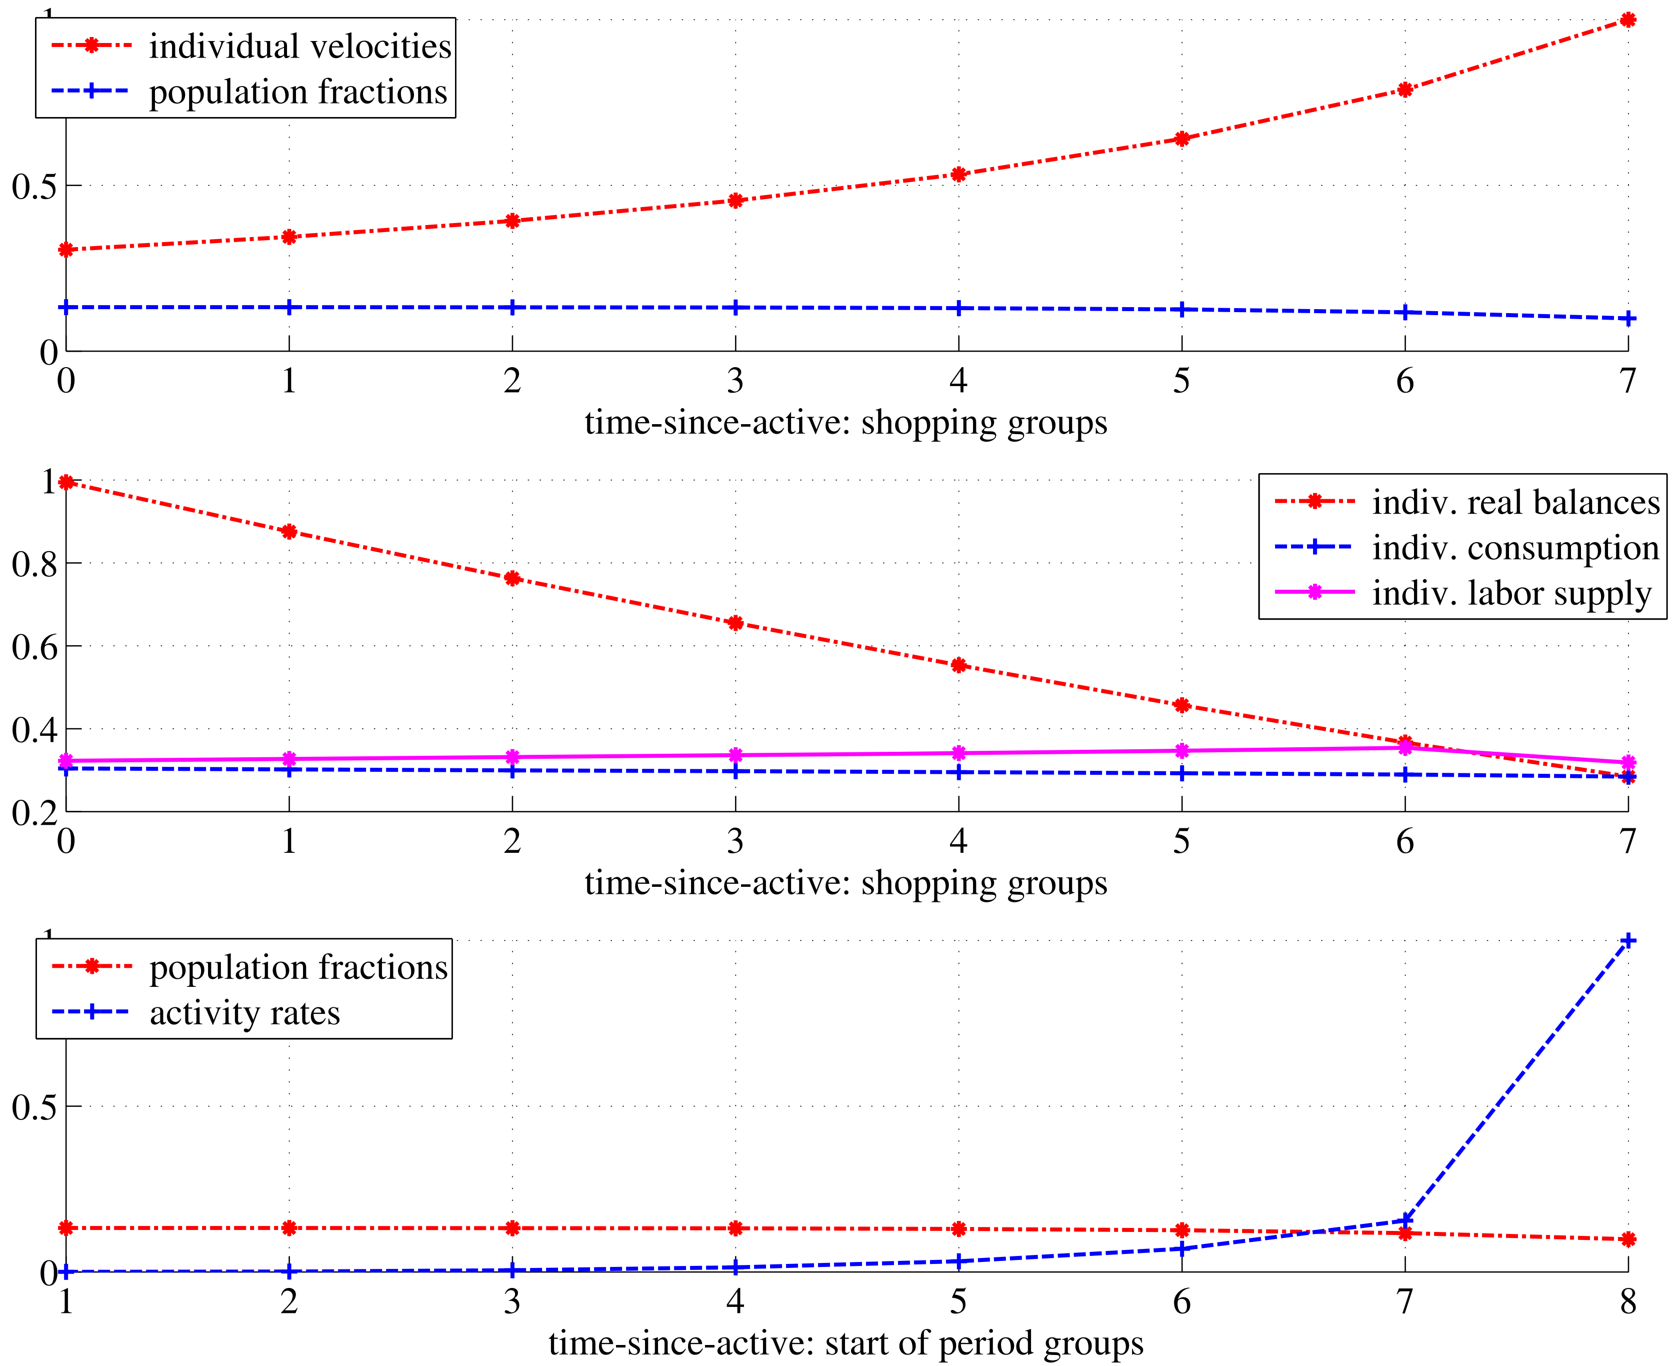

FIGURE 3: Money Growth Shock in the Pure New Keynesian Model

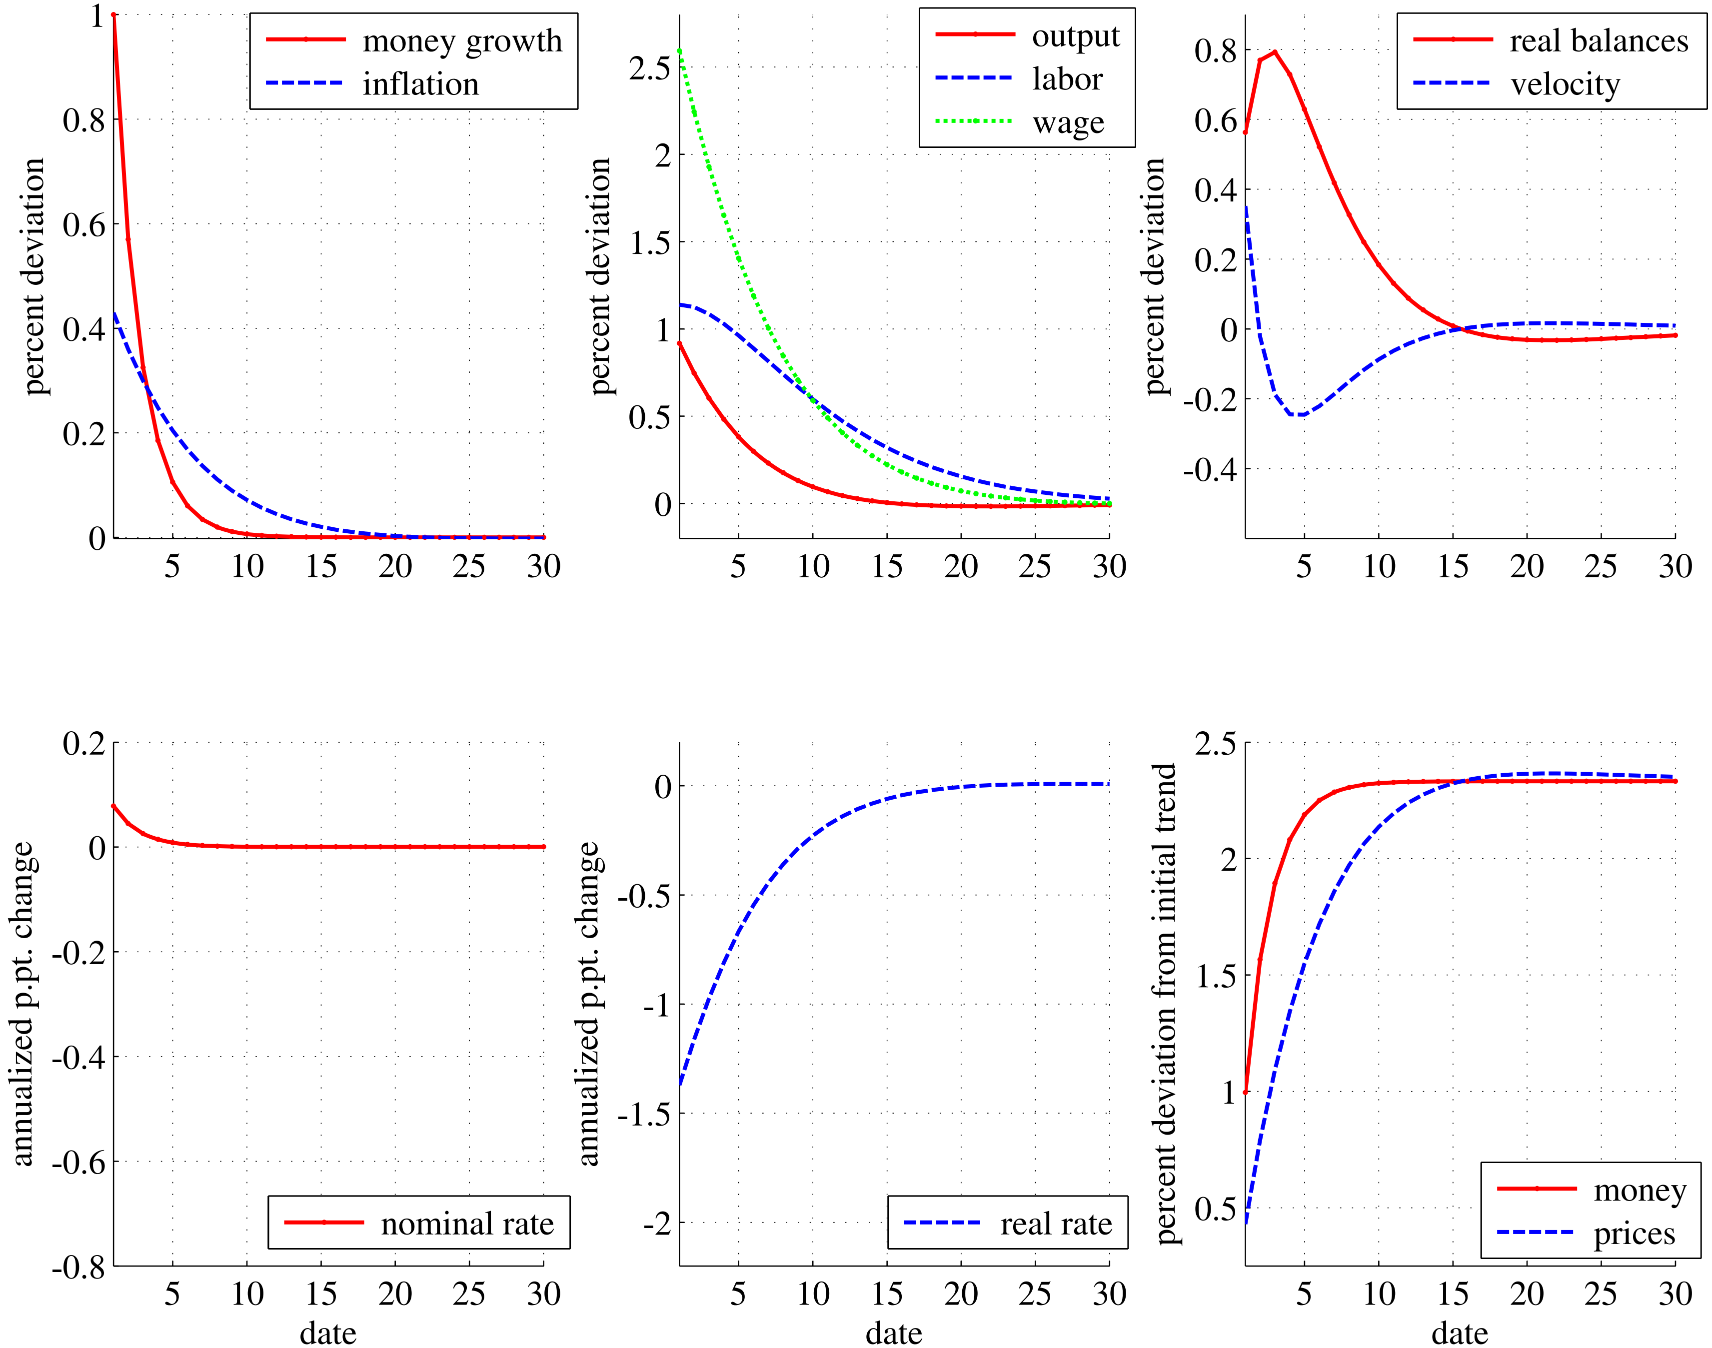

FIGURE 4: Money Growth Shock in Full Model with limited participation

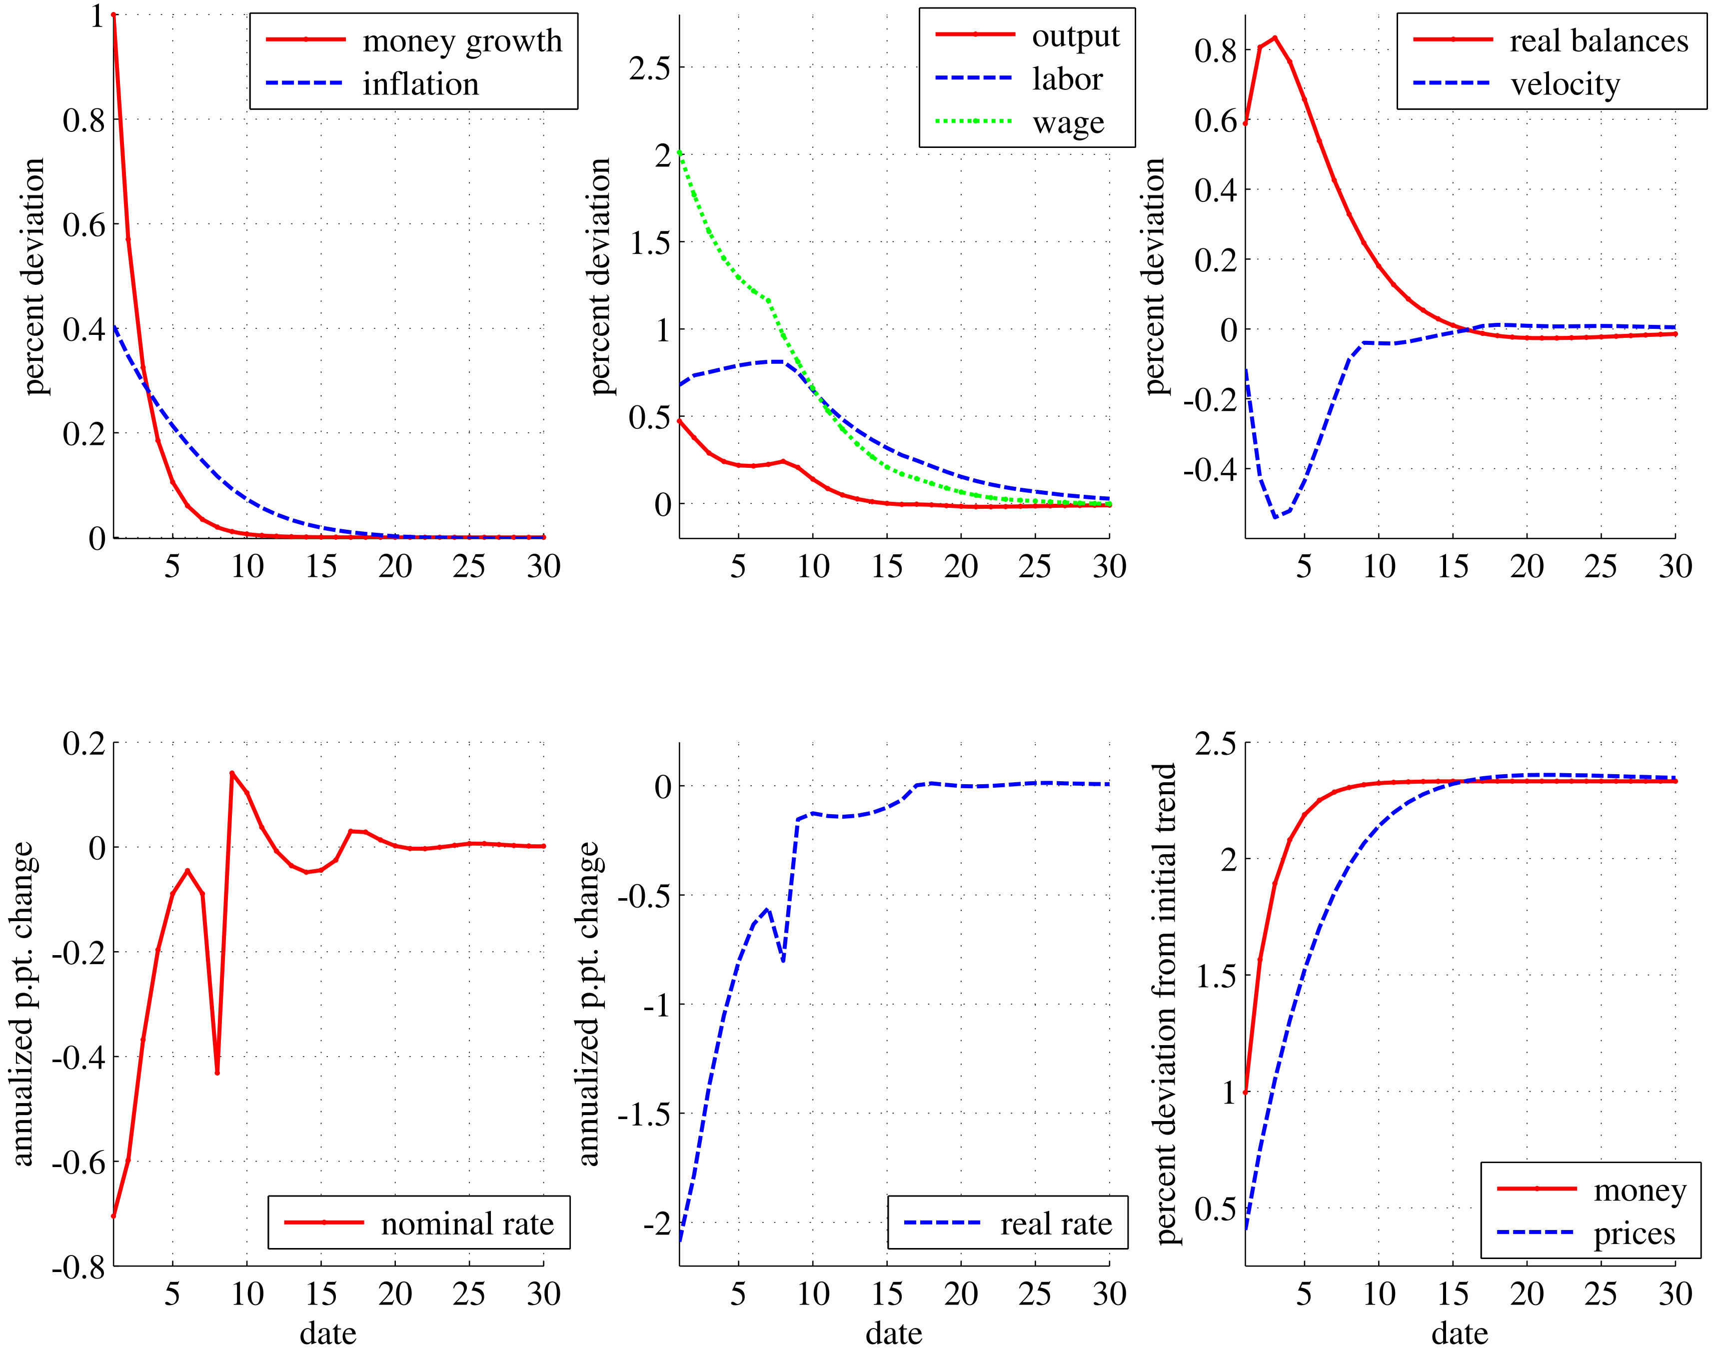

FIGURE 5: Money Growth Shock in Full Model - Decisions Among Active Households

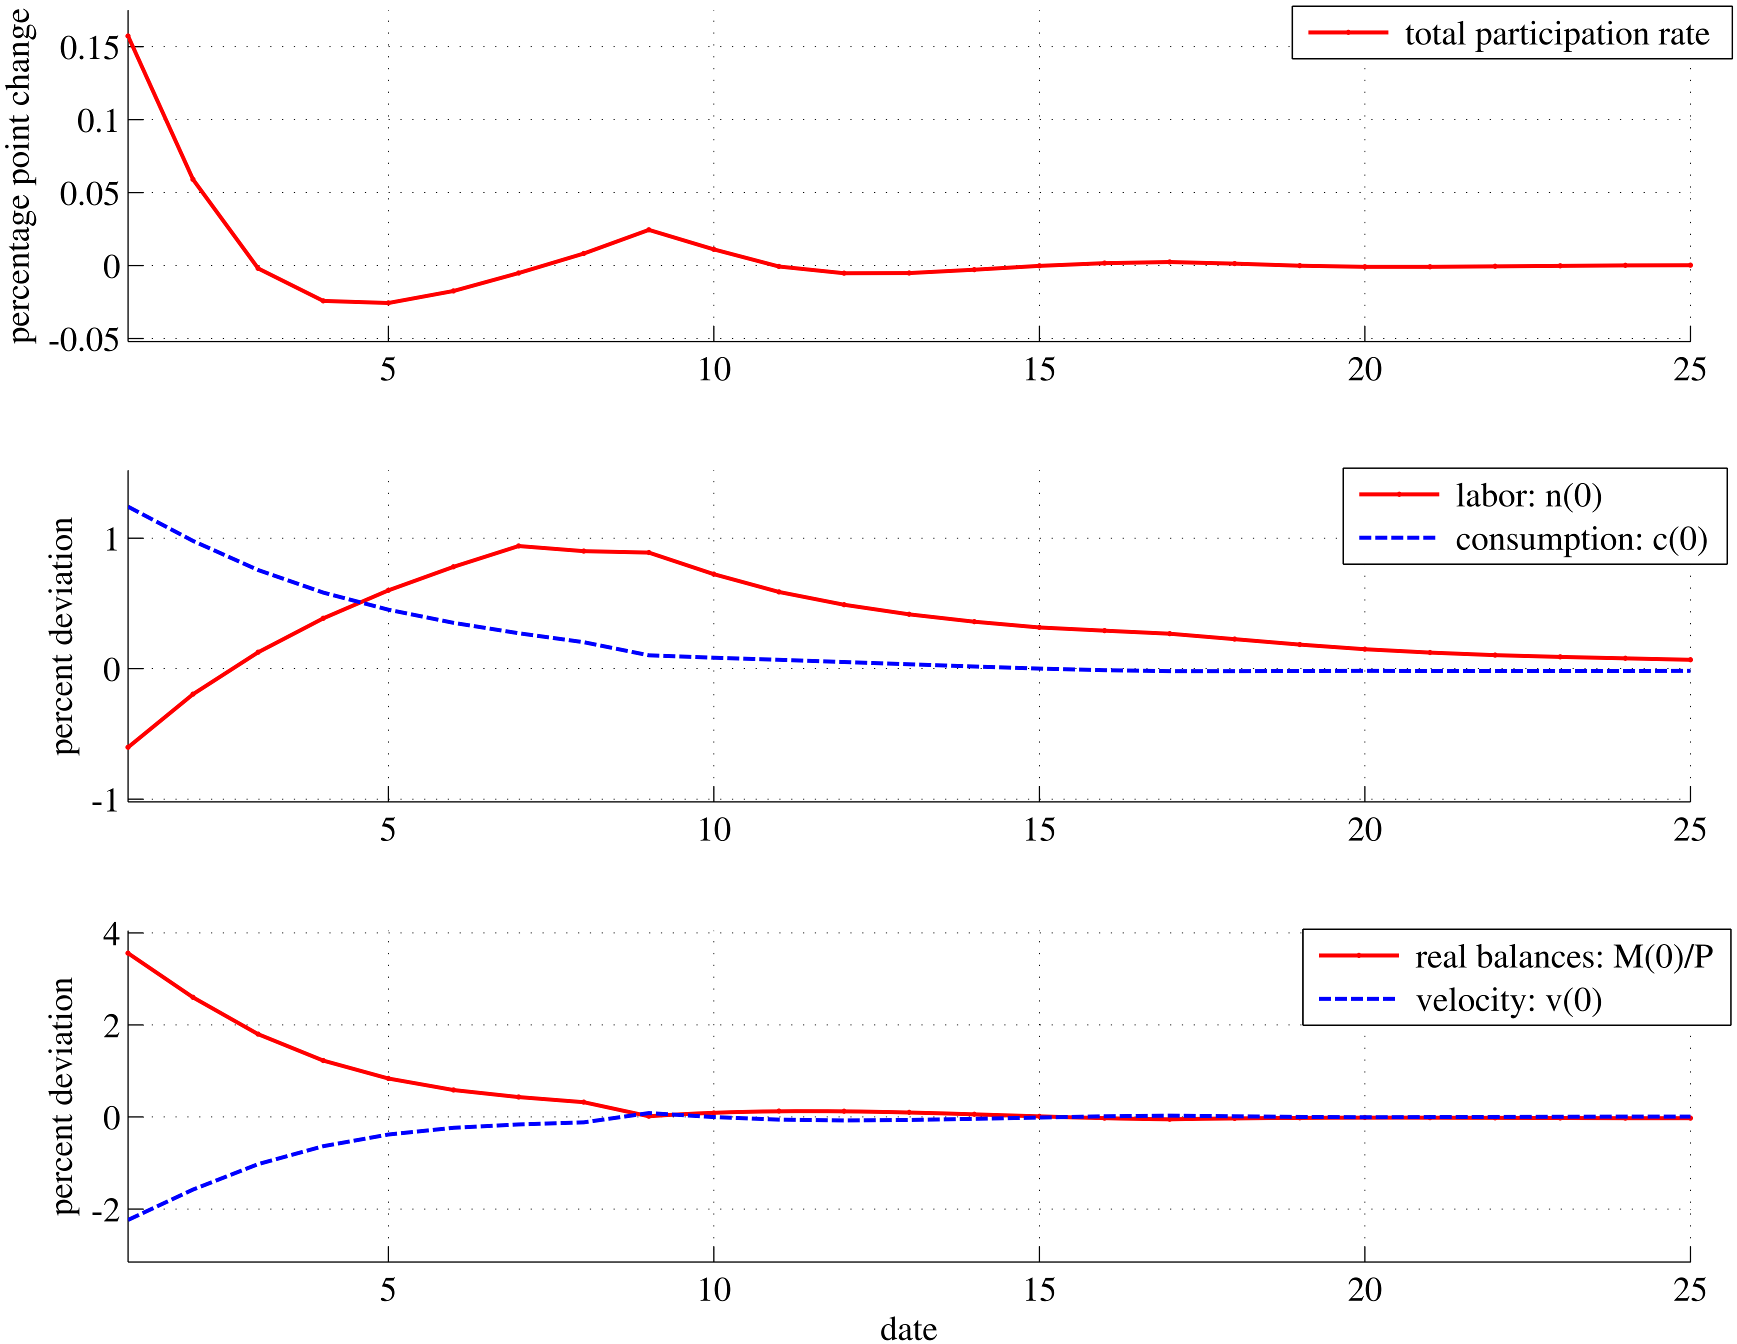

FIGURE 6: Productivity Shock under Taylor Rule in the Pure New Keynesian Model

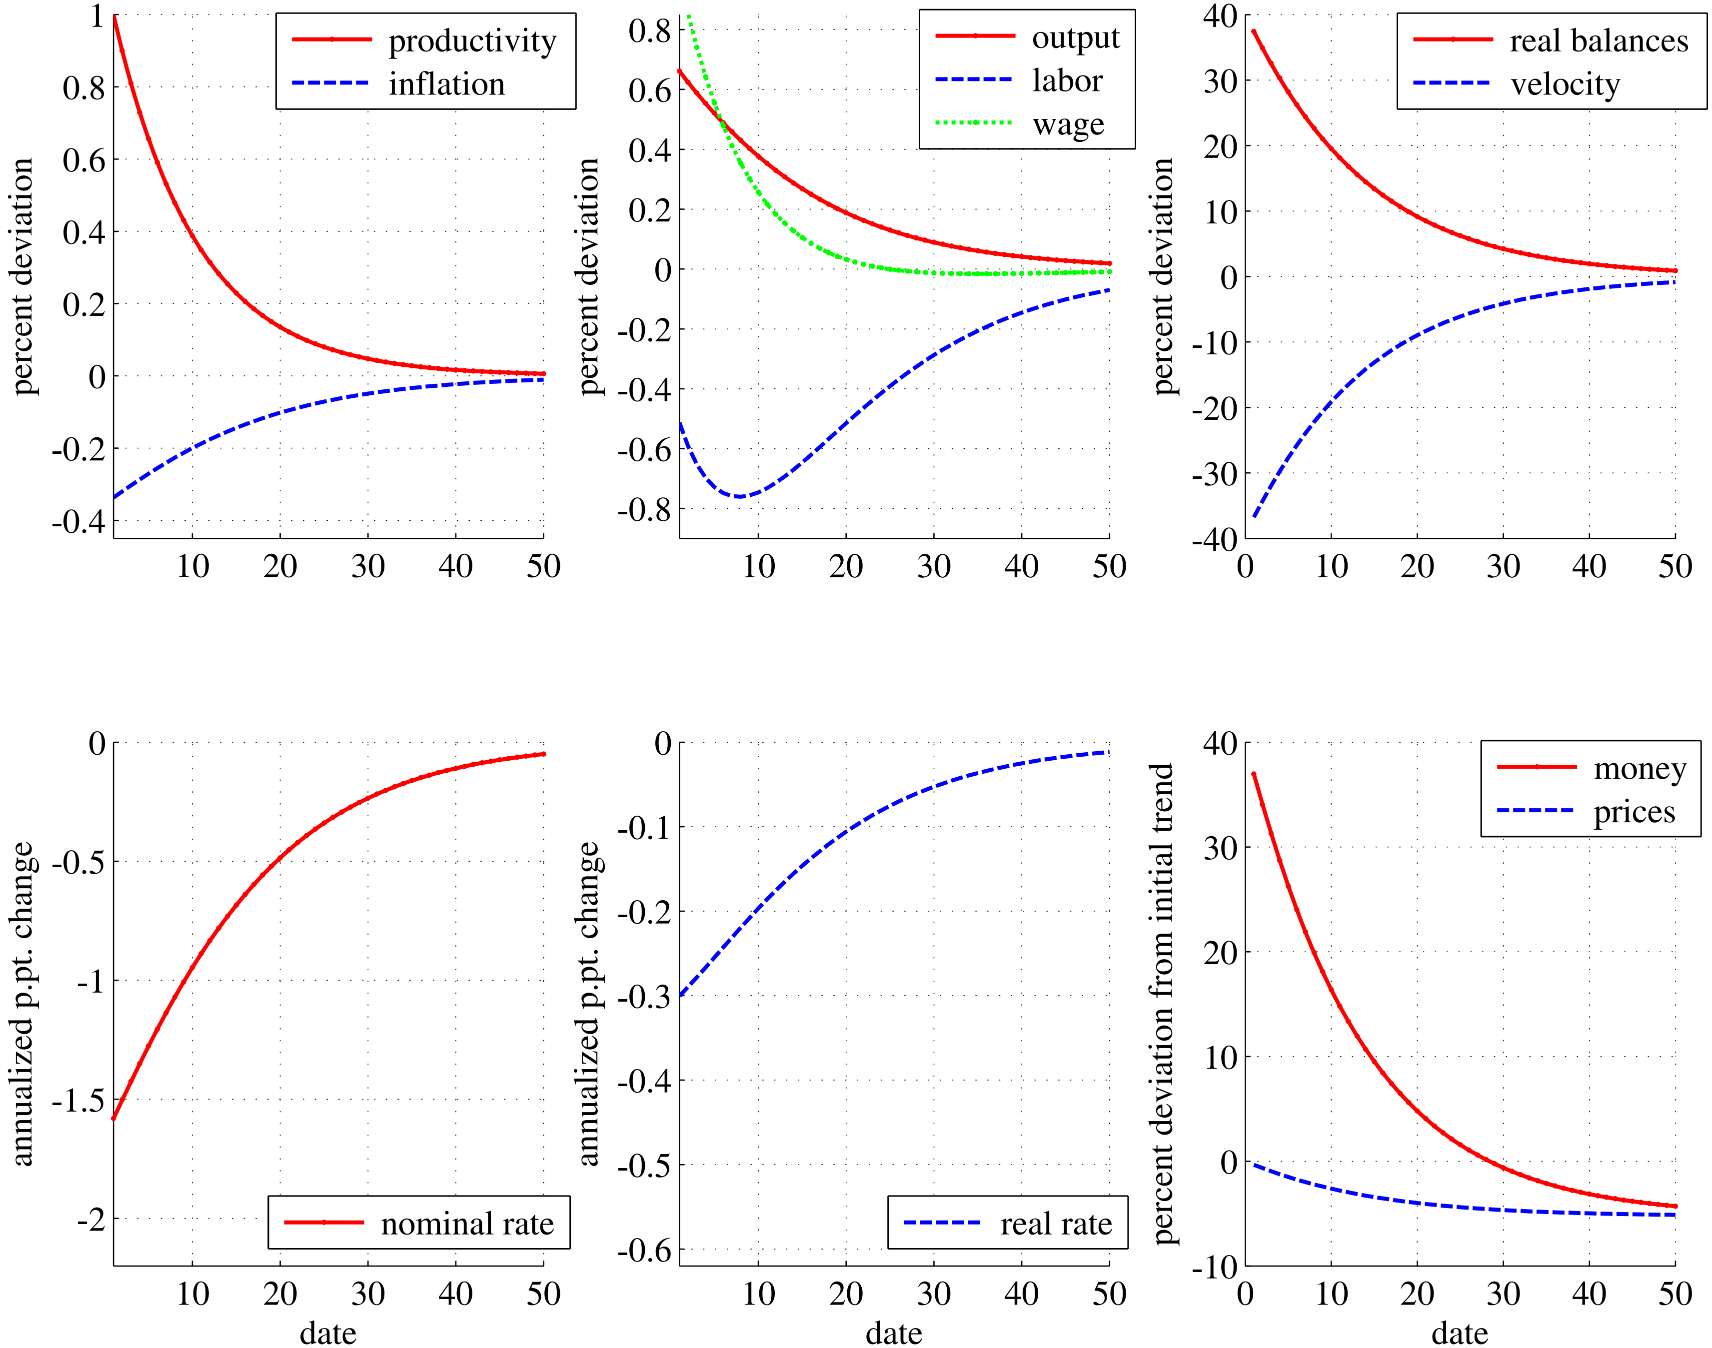

FIGURE 7: Productivity Shock under Taylor Rule in Full Model - Baseline Case

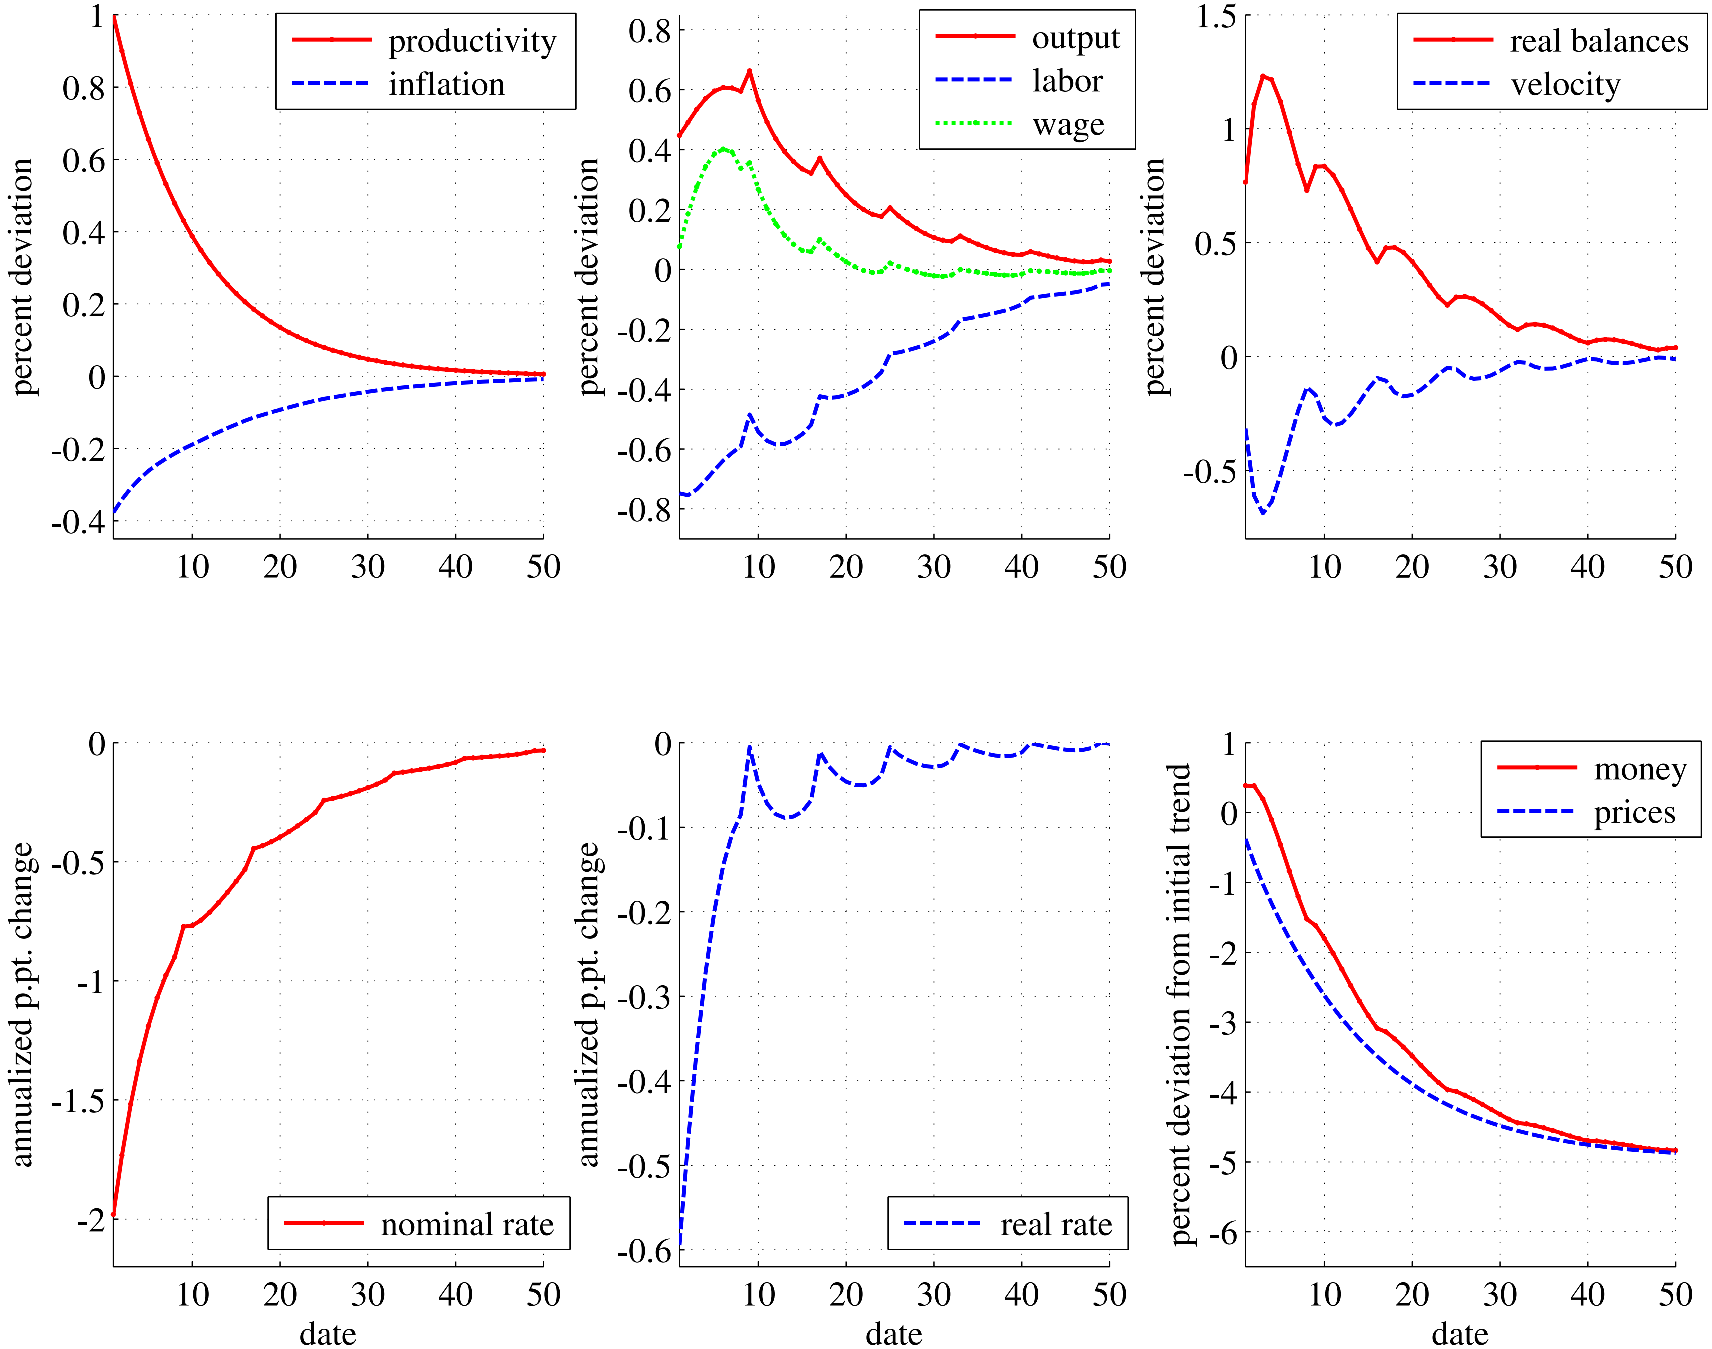

FIGURE 8: Productivity Shock under Taylor Rule in Full Model - Alternate Case

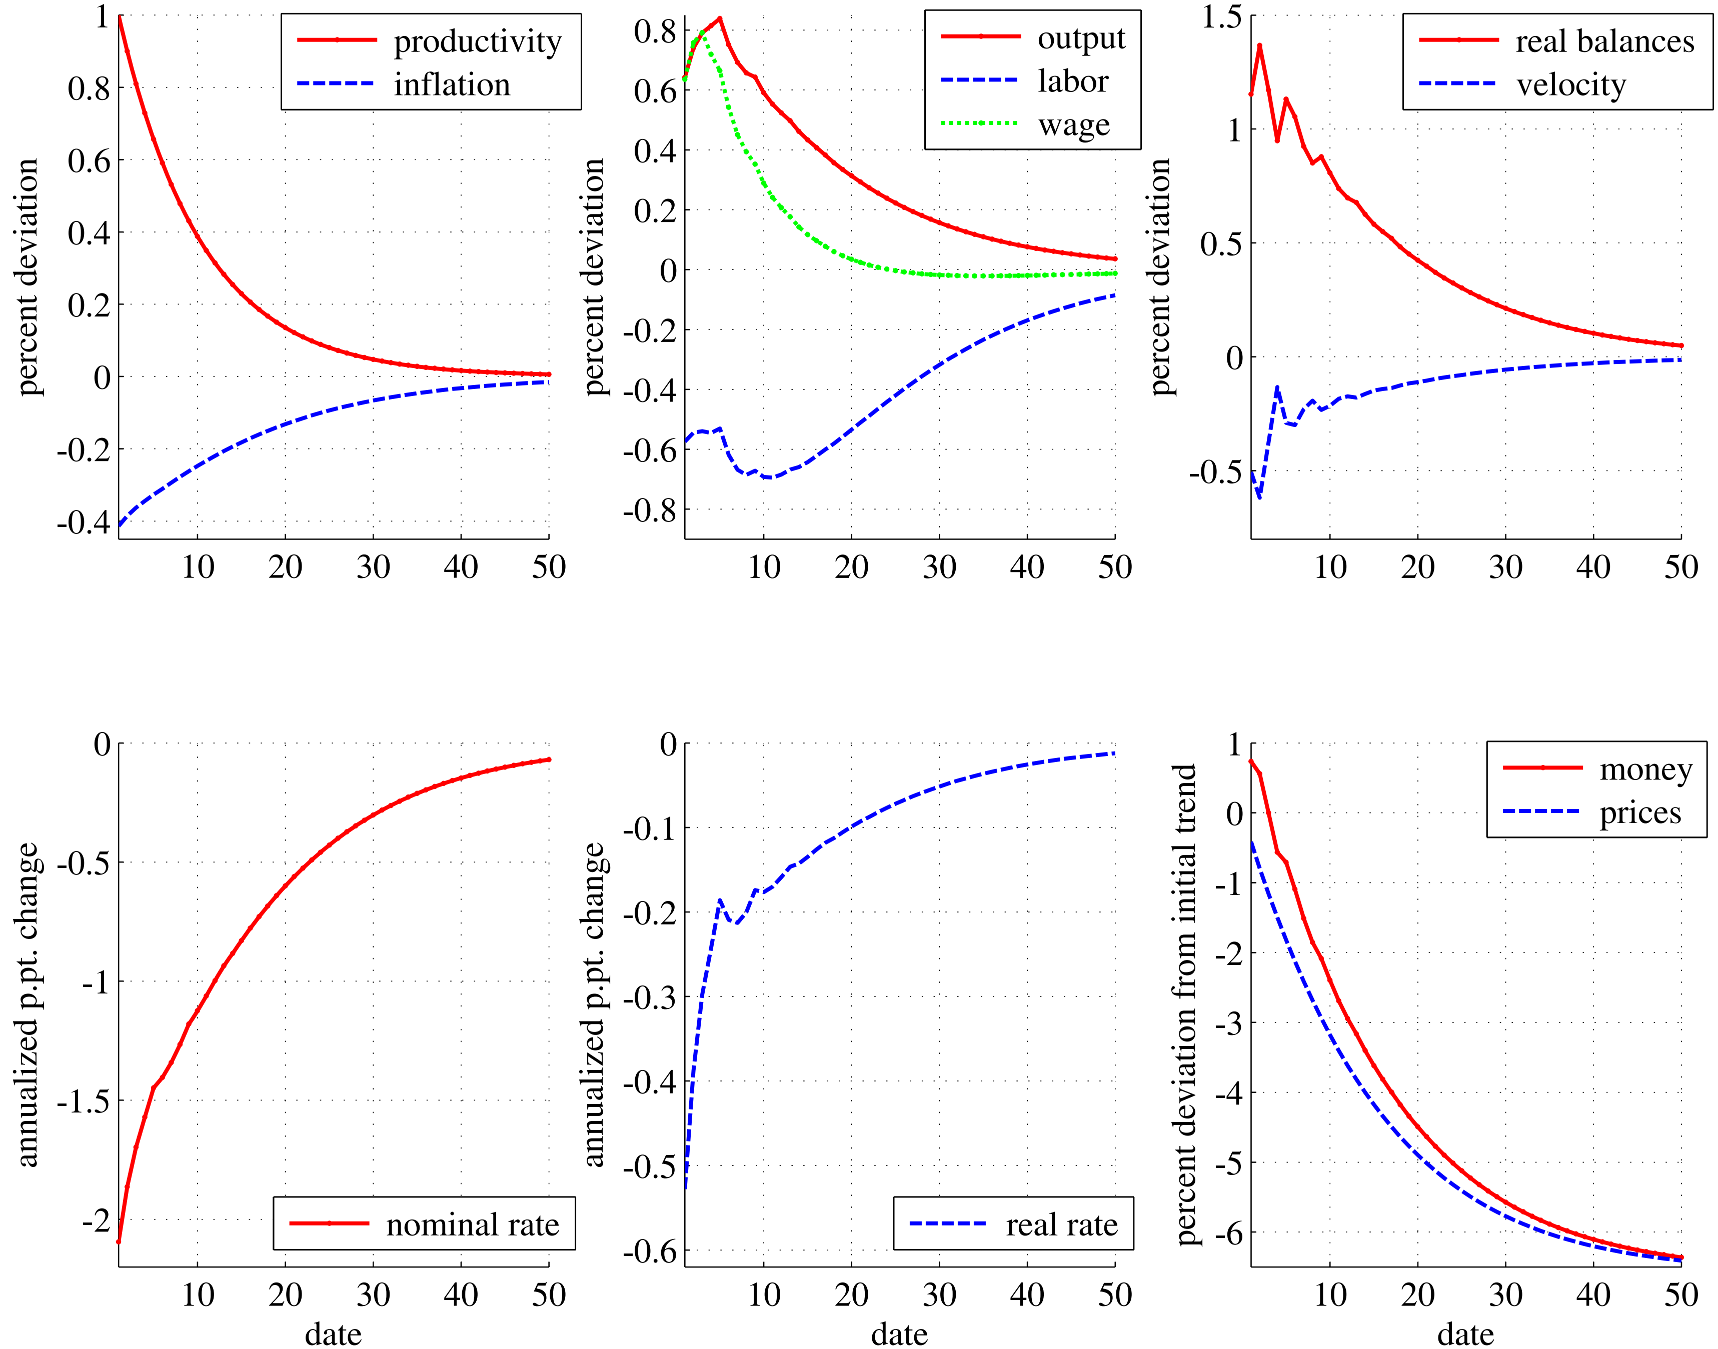

FIGURE 9: Full Model under a Taylor rule (zero steady-state inflation case)

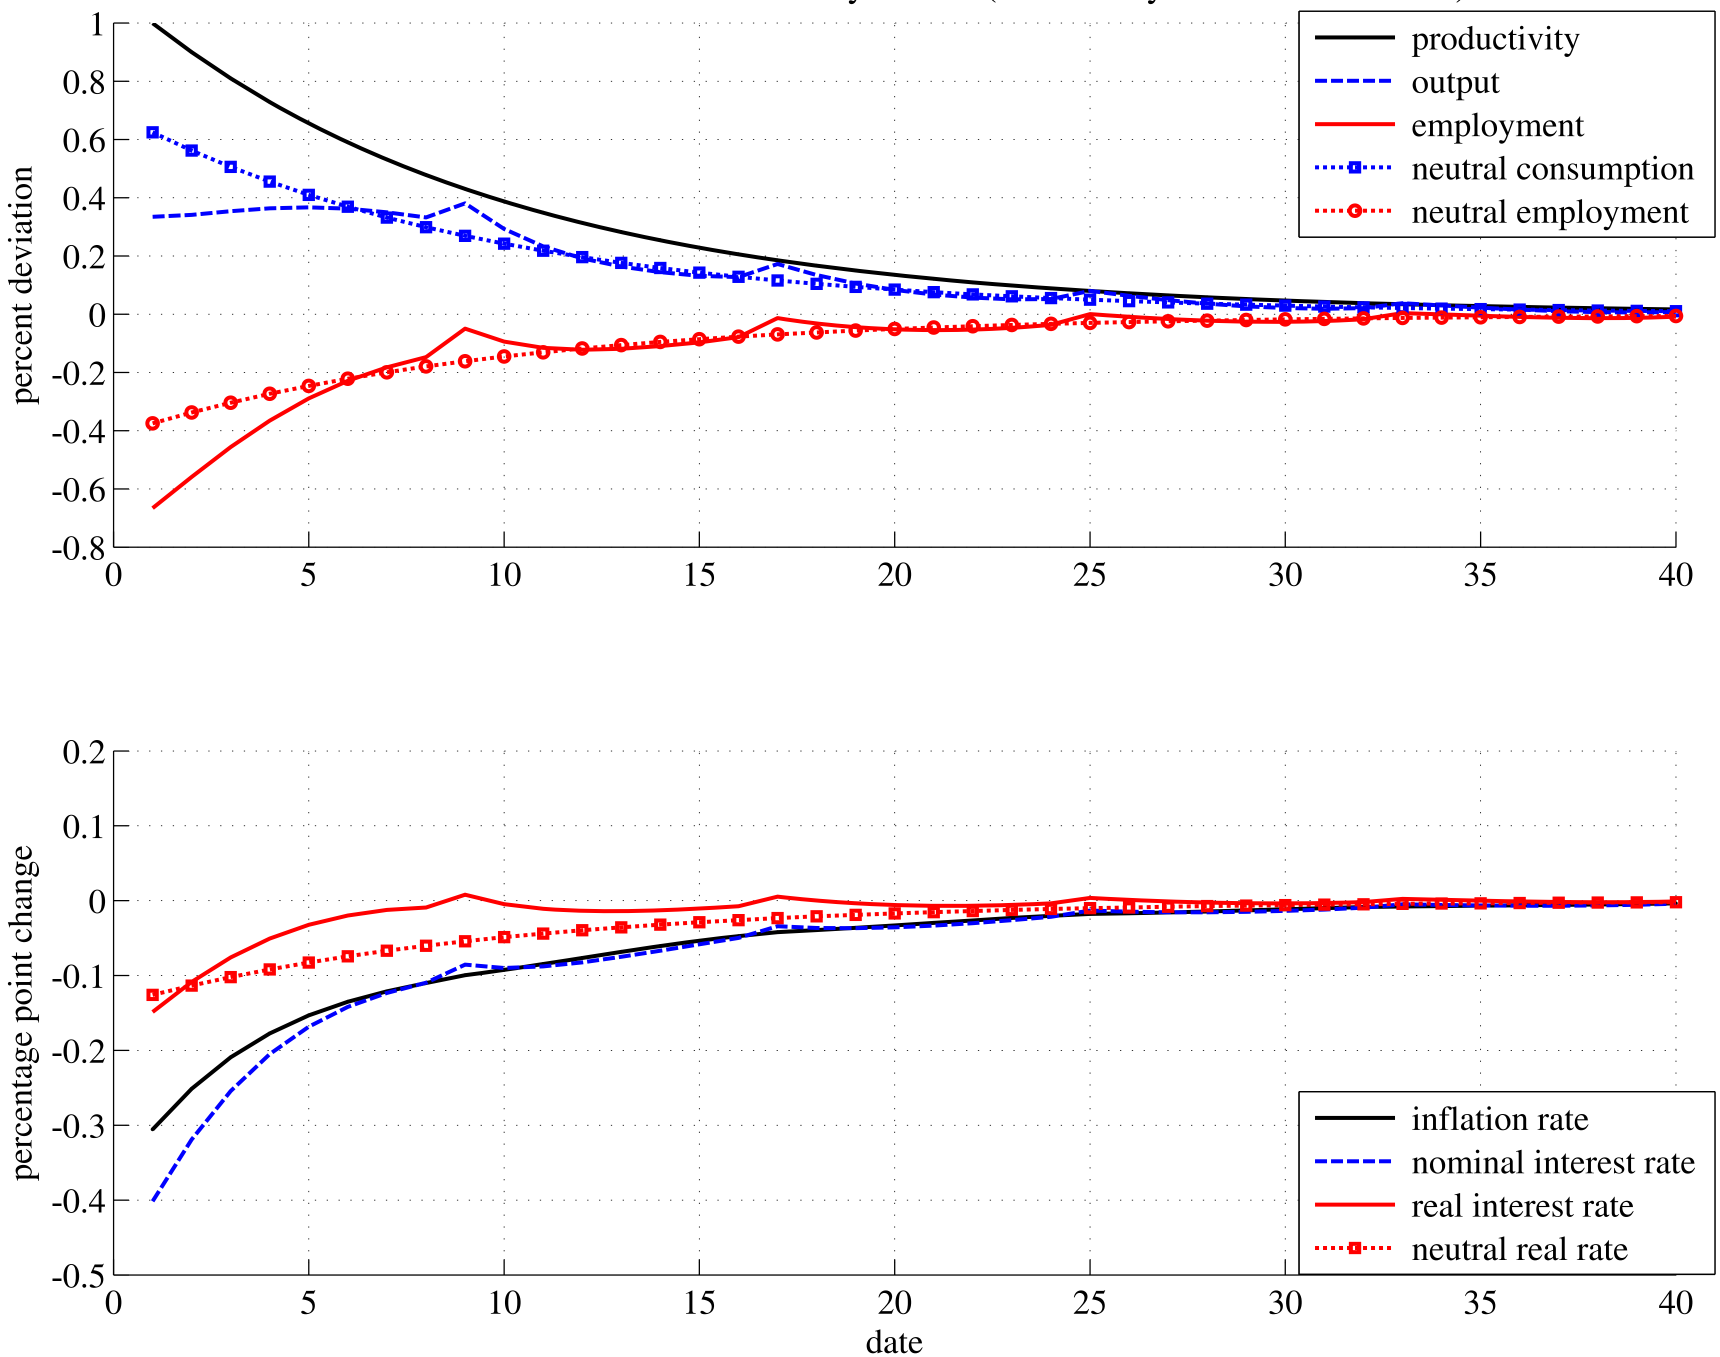

FIGURE 10: Full Model under Inflation Targeting (zero steady-state inflation case)

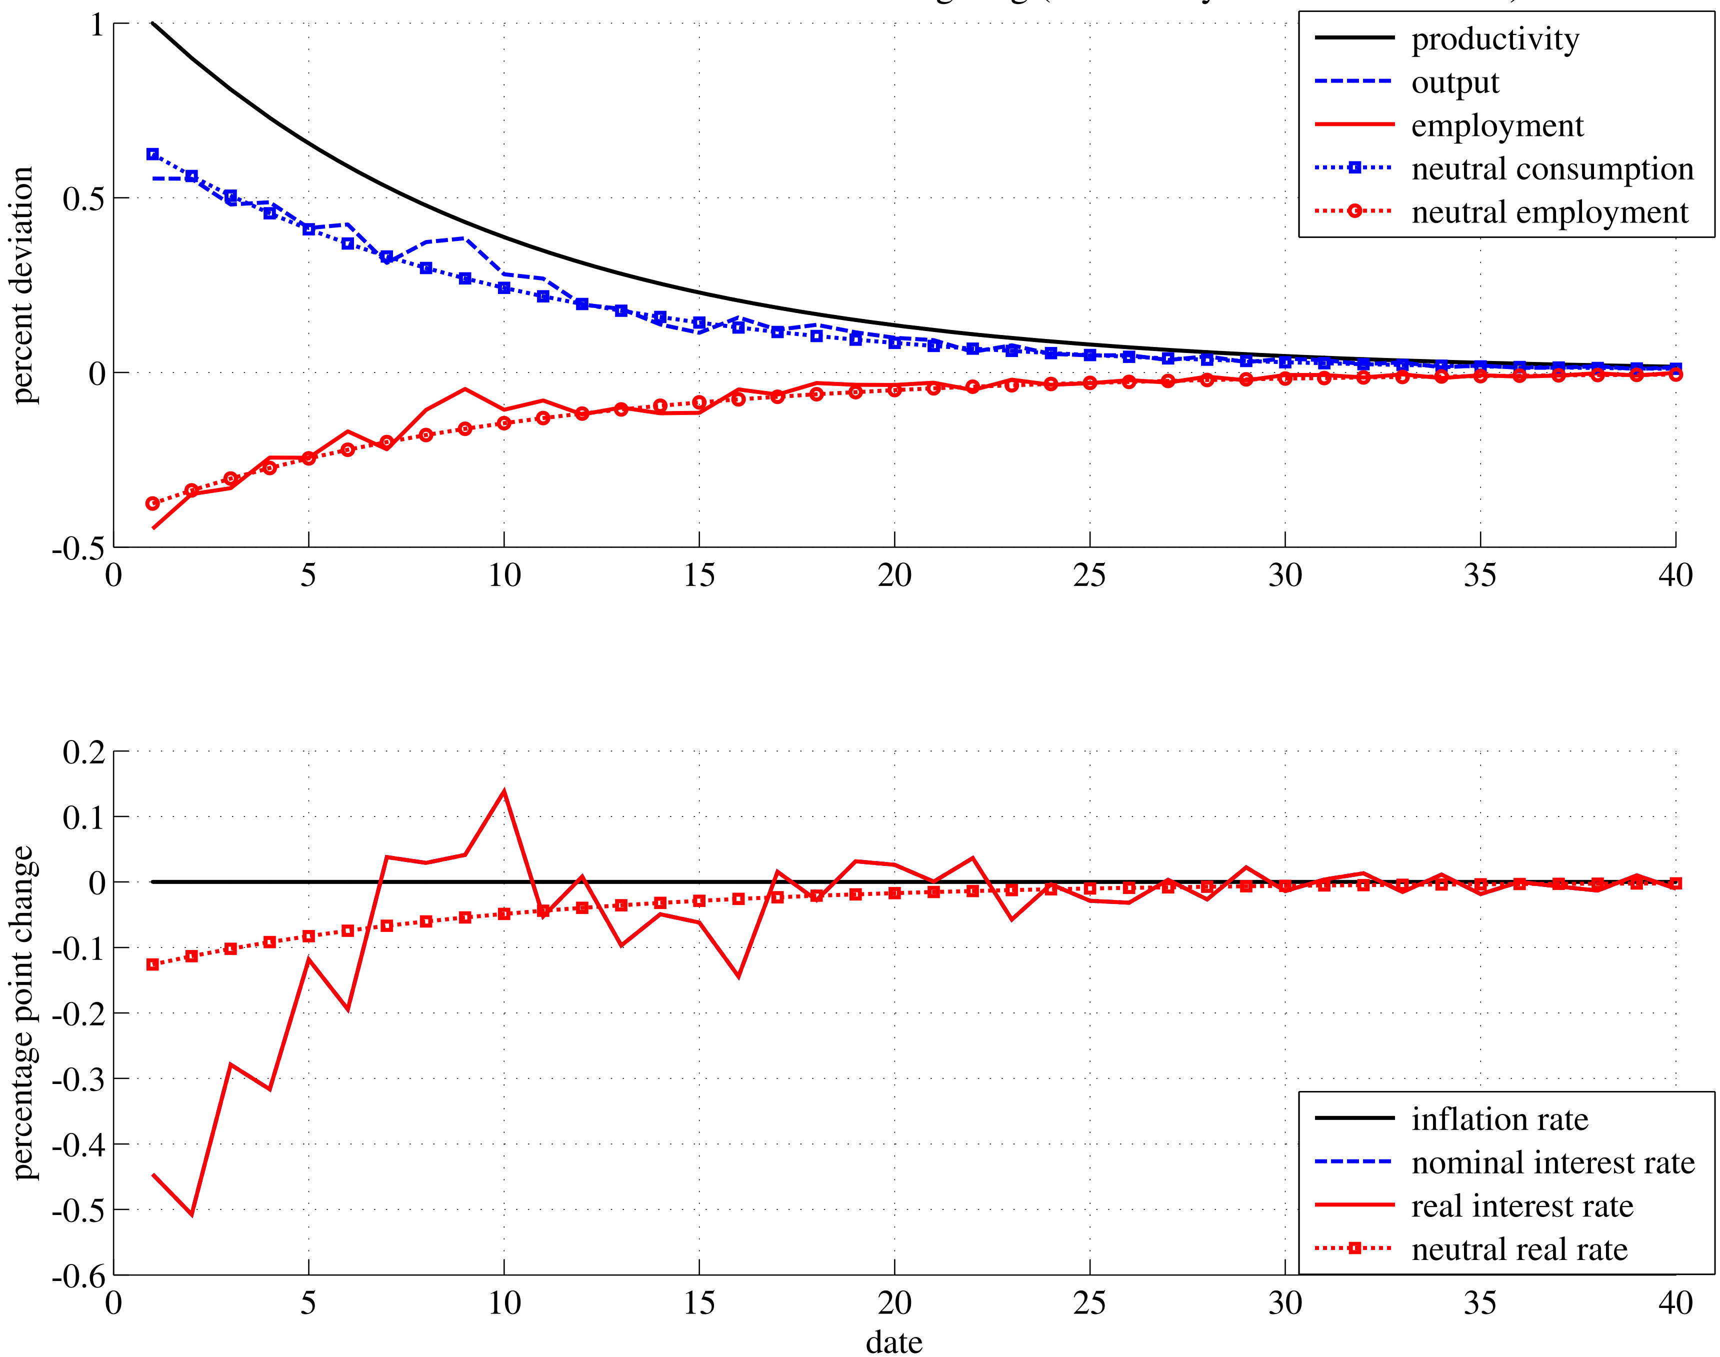

Appendix A Solution to firm's price-setting problem

As we characterize the solution to the price-setting firm's problem in this appendix, it will be convenient to begin with the nominal counterparts to equations 5 and 6. These are listed below.

$$
\begin{aligned}
W^0(\mathbf{s}) &= \max_{P_0} \Omega(\mathbf{s})\, \pi(\frac{P_0}{P};\mathbf{s}) + \beta\Big[\boldsymbol{\alpha}\int W^0(\mathbf{s}')\, dF(\mathbf{s}, d\mathbf{s}') \\
&\qquad + (1-\boldsymbol{\alpha})\int W^1(P_0;\mathbf{s}')\, dF(\mathbf{s}, d\mathbf{s}')\Big]
\end{aligned}
$$

$$
\begin{aligned}
W^1\left(\widetilde{P};\mathbf{s}\right) &= \Omega(\mathbf{s})\, \pi(\frac{\widetilde{P}}{P};\mathbf{s}) + \beta\Big[\boldsymbol{\alpha}\int W^0(\mathbf{s}')\, dF(\mathbf{s}, d\mathbf{s}') \\
&\qquad + (1-\boldsymbol{\alpha})\int W^1\left(\widetilde{P};\mathbf{s}'\right) dF(\mathbf{s}, d\mathbf{s}')\Big]
\end{aligned}
$$

The optimal choice for the firm currently selecting its nominal price will satisfy the first-order condition below, with the Benveniste-Scheinkman condition delivering the expression for $D_1W^1\left(\widetilde{P};\mathbf{s}'\right)$.

$$
0 = \frac{\Omega}{P} D_1\pi(\frac{P_0}{P};\mathbf{s}) + \beta(1-\boldsymbol{\alpha})\int D_1W^1(P_0;\mathbf{s}')\, dF(\mathbf{s}, d\mathbf{s}'), \tag{27}
$$

$$
D_1W^1\left(\widetilde{P};\mathbf{s}\right) = \frac{\Omega}{P} D_1\pi(\frac{\widetilde{P}}{P};\mathbf{s}) + \beta(1-\boldsymbol{\alpha})\int D_1W^1\left(\widetilde{P};\mathbf{s}'\right) dF(\mathbf{s}, d\mathbf{s}'). \tag{28}
$$

With repeated substitutions of (28) into (27), we arrive at a forward-looking condition characterizing the optimal current nominal price selected by the firm, P_{0t}. Using the notation $\mathbf{E}_t$ as a shorthand to represent the firm's expectations over future aggregate states, given current state s_t, we have:

$$
0 = \frac{\Omega_t}{P_t} D_1\pi(\frac{P_{0t}}{P_t};\mathbf{s}_t) + \beta(1-\boldsymbol{\alpha})\,\mathbf{E}_t\frac{\Omega_{t+1}}{P_{t+1}} D_1\pi(\frac{P_{0t}}{P_{t+1}};\mathbf{s}_{t+1}) + \beta^2(1-\boldsymbol{\alpha})^2\,\mathbf{E}_t\frac{\Omega_{t+2}}{P_{t+2}} D_1\pi(\frac{P_{0t}}{P_{t+2}};\mathbf{s}_{t+2}) + \cdots,
$$

which may be written more concisely as:

$$
E_t\sum_{j=0}^{\infty}(1-\boldsymbol{\alpha})^j\left[\frac{\beta^j\Omega(\mathbf{s}_{t+j})}{\Omega(\mathbf{s}_t)}\frac{P_t}{P_{t+j}} D_1\pi(\frac{P_{0t}}{P_{t+j}};\mathbf{s}_{t+j})\right] = 0.
$$

Given the constant-elasticity demand function in (1), and the cost function in (4), the firm's marginal profit function is $D_1\pi(p_t;s) = (1-\varepsilon)d(p_t,s) - [\frac{\partial e(p_t,\mathbf{s},\mathbf{z})}{\partial d(p_t,\mathbf{s})}][\frac{-\varepsilon d(p_t,\mathbf{s})}{p_t}]$. Abbreviating demand by $d_t \equiv d(p_t,s)$ and marginal cost by $e_t' \equiv \frac{\partial e(p_i,\mathbf{s},\mathbf{z})}{\partial d(p_i,\mathbf{s})}$, this expression becomes $D_1\pi(p_t;s) = [1-\varepsilon+\varepsilon e_t'/p_t]d_t$. Using this information in the efficiency condition above, we see that the firm sets its current price as a function of expected future interest rates (determined by the path of Ω) and current and future demand and marginal cost conditions (summarized in the paths of P, d, e') to satisfy:

$$
\mathbf{E}_t\sum_{j=0}^{\infty}(1-\boldsymbol{\alpha})^j\left[\frac{\beta^j\Omega(\mathbf{s}_{t+j})}{\Omega(\mathbf{s}_t)}\frac{P_t}{P_{t+j}}\left(1-\varepsilon+\varepsilon\frac{P_{t+j}}{P_{0t}}e_{t+j}'\right)d_{t+j}\right] = 0. \tag{29}
$$

Recall that the firm's demand in date $t+j$ is given by $d_{t+j} = Y_{t+j} p_{t+j}^{-\varepsilon}$. If it has not yet adjusted its nominal price away from that selected in date t, the firm's demand will be $d_{t+j} = Y_{t+j} \left[\frac{p_{0t}}{\boldsymbol{\pi}_{t,t+j}}\right]^{-\varepsilon}$, where $\boldsymbol{\pi}_{t,t+j} = \frac{P_{t+j}}{P_t}$ denotes the ratio of the aggregate price level in period $t+j$ relative to that in period t. Given these substitutions, we can move from equation (29) to the following expression for the firm's optimal relative price.

$$p_{0t} = \frac{\varepsilon}{\varepsilon - 1} \frac{\mathbf{E}_t \sum_{j=0}^{\infty} \left[\beta\left(1-\boldsymbol{\alpha}\right)\right]^j \frac{\Omega_{t+j}}{\boldsymbol{\pi}_{t,t+j}} Y_{t+j} e'_{t+j} \left(Y_{t+j}\left[\frac{p_{0t}}{\boldsymbol{\pi}_{t,t+j}}\right]^{-\varepsilon}\right) \boldsymbol{\pi}_{t,t+j}^{\varepsilon+1}}{\mathbf{E}_t \sum_{j=0}^{\infty} \left[\beta\left(1-\boldsymbol{\alpha}\right)\right]^j \frac{\Omega_{t+j}}{\boldsymbol{\pi}_{t,t+j}} Y_{t+j} \boldsymbol{\pi}_{t,t+j}^{\varepsilon}} \tag{30}$$

Notice that the expression above provides an explicit solution only if the firm's marginal cost, e'_{t+j}, is constant in its level of production. More generally, when $\nu \neq 1$, e'_{t+j} will depend upon the firm's level of production, $y(\frac{p_{0t}}{\boldsymbol{\pi}_{t,t+j}}) = Y_{t+j}\left[\frac{p_{0t}}{\boldsymbol{\pi}_{t,t+j}}\right]^{-\varepsilon}$. In this case, we have $e'_{t+j} = \frac{w_{t+j}}{\nu} z_{t+j}^{\frac{-1}{\nu}} Y_{t+j}^{\frac{1-\nu}{\nu}} \boldsymbol{\pi}_{t,t+j}^{\frac{\varepsilon(1-\nu)}{\nu}} p_{0t}^{-\frac{\varepsilon(1-\nu)}{\nu}}$, which may be substituted into (30) to obtain

$$p_{0t} = \frac{\varepsilon}{\varepsilon - 1} \frac{\mathbf{E}_t \sum_{j=0}^{\infty} \left[\beta\left(1-\boldsymbol{\alpha}\right)\right]^j \Omega_{t+j} Y_{t+j}^{\frac{1}{\nu}} \frac{w_{t+j}}{\nu} z_{t+j}^{\frac{-1}{\nu}} \boldsymbol{\pi}_{t,t+j}^{\frac{\varepsilon}{\nu}}}{\mathbf{E}_t \sum_{j=0}^{\infty} \left[\beta\left(1-\boldsymbol{\alpha}\right)\right]^j \Omega_{t+j} Y_{t+j} \boldsymbol{\pi}_{t,t+j}^{\varepsilon-1}} p_{0t}^{-\frac{\varepsilon(1-\nu)}{\nu}}, \tag{31}$$

making the solution for the firm's relative price direct.

Appendix B Family efficiency conditions

This appendix lists the complete set of efficiency conditions arising from the optimization problem facing the extended family of all households, following some minor intermediate algebra. First, we have the m_{0t} first order condition:

$$D_1 u(c_{0t}, 1 - n_{0t}) = \Omega_t. \tag{32}$$

Next, we have the first order conditions with regard to the $m_{j+1,t+1}$ choices:

$$D_1 u(c_{jt}, 1 - n_{jt}) = \beta \boldsymbol{E}_t \frac{P_t}{P_{t+1}} \Big[\alpha_{j+1,t+1}\Omega_{t+1} + (1 - \alpha_{j+1,t+1}) D_1 u(c_{j+1,t+1}, 1 - n_{j+1,t+1})\Big], \quad \text{for } j = 0, ..., J-2, \tag{33}$$

though the analogous first order condition for $m_{J,t+1}$is replaced by $m_{J,t} = 0$, because non-negativity binds on this choice in the presence of positive nominal interest rates.

$$D_1 u(c_{J-1,t}, 1 - n_{J-1,t}) = \beta \boldsymbol{E}_t \frac{P_t}{P_{t+1}} \Omega_{t+1} \text{ is instead: } m_{J,t} = 0. \tag{34}$$

Next, there are the n_{jt} first order conditions, followed by the $e_{j+1,t+1}$ first order conditions.

$$D_2 u(c_{jt}, 1 - n_{jt}) = (1 - \lambda_N) w_t H_t + w_t r_{jt}, \text{ for } j = 0, ..., J-1 \tag{35}$$

$$r_{jt} = \beta \boldsymbol{E}_t \frac{P_t}{P_{t+1}} \lambda_N \Big[\alpha_{j+1,t+1}\Omega_{t+1} + (1-\alpha_{j+1,t+1})D_1 u(c_{j+1,t+1}, 1-n_{j+1,t+1})\Big], \text{ for } j=0,...,J-2, \text{ and,}$$

$$r_{J-1,t} = \beta \boldsymbol{E}_t \frac{P_t}{P_{t+1}} \lambda_N \Omega_{t+1}$$

However, notice that the right-hand-side in each equation above matches its counterpart from equations 33 and 34 perfectly, but for the λ_N. So, for convenience, we write these as

$$r_{jt} = \lambda_N D_1 u(c_{jt}, 1-n_{jt}), \text{ for } j=0,...,J-2, \text{ and} \tag{36}$$

$$r_{J-1,t} = \beta \boldsymbol{E}_t \frac{P_t}{P_{t+1}} \lambda_N \Omega_{t+1} \quad (\text{provided } n_{J-1,t} \geq 0). \tag{37}$$

The α_{jt} first order conditions are:

$$\begin{aligned}
&[u(c_{0t}, 1-n_{0t}) - u(c_{jt}, 1-n_{jt})] \\
&\qquad + \Omega_t \Big[\frac{P_{t-1}}{P_t}\Big(m_{jt} + \lambda_N e_{jt} + \lambda_\Pi \Pi_{t-1}\Big) - m_{0t} - \varphi'(\alpha_{jt})\Big] \\
&\qquad + H_t(1-\lambda_N) w_t [n_{0t} - n_{jt}] + [q_{0t} - q_{jt}] = 0, \text{ for } j=1,...,J-1.
\end{aligned} \tag{38}$$

Next, we have the $\theta_{j+1,t+1}$ conditions, followed by the first order condition with respect to χ_{t+1}.

$$\begin{aligned}
q_{jt} = \beta \boldsymbol{E}_t \Bigg[&\alpha_{j+1,t+1} u(c_{0,t+1}, 1-n_{0,t+1}) + (1-\alpha_{j+1,t+1}) u(c_{j+1,t+1}, 1-n_{j+1,t+1}) \\
&+ \Omega_{t+1} \Big[\frac{P_t}{P_{t+1}} \alpha_{j+1,t+1}\Big(m_{j+1,t+1} + \lambda_N e_{j+1,t+1} + \lambda_\Pi \Pi_t\Big) - \alpha_{j+1,t+1} m_{0,t+1} - \varphi(\alpha_{j+1,t+1})\Big] \\
&+ H_{t+1}(1-\lambda_N) w_{t+1}\Big(\alpha_{j+1,t+1} n_{0,t+1} + (1-\alpha_{j+1,t+1}) n_{j+1,t+1}\Big) \\
&+ \alpha_{j+1,t+1} q_{0,t+1} + (1-\alpha_{j+1,t+1}) q_{j+1,t+1} \Bigg], \text{ for } j=0,...,J-2, \text{ and}
\end{aligned} \tag{39}$$

$$\begin{aligned}
q_{J-1,t} = \beta \boldsymbol{E}_t \Bigg[&u(c_{0,t+1}, 1-n_{0,t+1}) + \Omega_{t+1} \Big[\frac{P_t}{P_{t+1}}\Big(m_{J,t+1} + \lambda_N e_{J,t+1} + \lambda_\Pi \Pi_t\Big) \\
&- m_{0,t+1} - \varphi(1)\Big] + H_{t+1}(1-\lambda_N) w_{t+1} n_{0,t+1} + q_{0,t+1} \Bigg]
\end{aligned} \tag{40}$$

$$H_t = \beta \boldsymbol{E}_t \frac{P_t}{P_{t+1}} \Omega_{t+1} \tag{41}$$

A set of additional essential conditions governing the behavior on the household side of the economy are the constraints themselves; these are reiterated here for completeness.

$$\frac{P_{t-1}}{P_t}\left(\chi_t + \mu_t \overline{\mathbf{m}}_t + \sum_{j=1}^{J} \theta_{jt}\alpha_{jt}[m_{jt} + \lambda_N e_{jt} + \lambda_\Pi \Pi_{t-1}]\right) = m_{0t} \sum_{j=1}^{J} \theta_{jt}\alpha_{jt} + \sum_{j=1}^{J} \theta_{jt}\varphi(\alpha_{jt}) \tag{42}$$

$$\chi_{t+1} = (1-\lambda_\Pi)\Pi_t + (1-\lambda_N) w_t \left(\sum_{j=1}^{J} \theta_{jt}\alpha_{jt} n_{0t} + \sum_{j=1}^{J-1} \theta_{jt}(1-\alpha_{jt}) n_{jt} \right) \tag{43}$$

$$\theta_{1,t+1} = \sum_{j=1}^{J} \theta_{jt}\alpha_{jt} \tag{44}$$

$$\theta_{j+1,t+1} = \theta_{jt}(1-\alpha_{jt}), \quad \text{for } j = 1, ..., J-1 \tag{45}$$

$$e_{j+1,t+1} = w_t n_{jt}, \quad \text{for } j = 0, ..., J-1 \tag{46}$$

$$c_{0t} = m_{0t} - m_{1,t+1} \tag{47}$$

$$c_{jt} = \frac{P_{t-1}}{P_t}\left[m_{jt} + \lambda_N e_{jt} + \lambda_\Pi \Pi_{t-1}\right] - m_{j+1,t+1}, \text{ for } j = 1, ..., J-1 \tag{48}$$