DSGE Models for Monetary Policy*

Lawrence J. Christiano[†], Mathias Trabandt[‡], Karl Walentin[§]

Preliminary and not for distribution.
Prepared for Handbook on Monetary Economics, edited by Friedman and Woodford.

October 28, 2009

Abstract

We review the consensus New Keynesian model, and how its various features are designed to replicate VAR-based facts about the response of aggegate variables to economic shocks. We discuss extensions of the model to better capture labor market variables such as unemployment. We estimate and evaluate our various models using the method of matching impulse response functions. We show how that procedure can be cast as a Bayesian procedure.
Keywords: DSGE, labor market frictions, endogenous separations, unemployment, Bayesian estimation

JEL codes: E2, E3, E5, J6

*The views expressed in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Executive Board of the European Central Bank or of Sveriges Riksbank.

[†]Northwestern University, Department of Economics, 2001 Sheridan Road, Evanston, Illinois 60208, USA. Phone: +1-847-491-8231. E-mail: l-christiano@northwestern.edu.

[‡]European Central Bank and Sveriges Riksbank. Contact Addresses: European Central Bank, Kaiserstrasse 29, 60311 Frankfurt am Main, Germany. Phone: +49 69 1344 6321. E-mail: mathias.trabandt@ecb.int. Sveriges Riksbank, Research Division, 103 37 Stockholm, Sweden. Phone: +46-8-787 0438. E-mail: mathias.trabandt@riksbank.se

[§]Sveriges Riksbank, Research Division, 103 37 Stockholm, Sweden. Phone: +46-8-787 0491. E-mail: karl.walentin@riksbank.se

Contents

1. Introduction

There has been enormous progress in recent years in the development of dynamic, stochastic general equilibrium (DSGE) models for the purpose of monetary policy analysis. These models have been shown to fit aggregate data well by conventional econometric measures. For example, they have been shown to do as well or better than simple atheoretical statistical models at forecasting outside the sample of data on which they were estimated. In part because of these successes, a consensus has formed around a particular model structure, the New Keynesian model.

In this paper, we begin by reviewing the New Keynesian model. We describe and motivate its various features. These features are primarily motivated by researchers' beliefs about how the economy responds to monetary policy shocks.[1] In many cases, though not all, researchers' beliefs are motivated by results based on identified vector autoregressions (VARs). We explain why the current consensus model adopts habit persistence in preferences, adjustment costs in terms of the change in investment and frictions in the setting of prices and wages.

Using US macroeconomic data, we show how the parameters of the consensus DSGE model are estimated by matching model and VAR-based impulse response functions. The advantage of this econometric approach is transparency and focus. The transparency reflects that the estimation strategy has a simple graphical representation, involving objects - impulse response functions - about which economists have strong intuition. The advantage of focus comes from the possibility of studying the empirical properties of a model without having to specify a full set of shocks. An important methodological development of recent years is the adoption of Bayesian methods of econometric inference. We show how to implement the impulse response matching strategy using Bayesian methods. As a result, all the machinery of priors and posteriors, as well as the marginal likelihood as a measure of model fit, is available to researchers doing inference about DSGE models based on matching model and VAR-based impulse response functions.

We show how well the New Keynesian model is able to replicate VAR-based evidence on the response of the economy to three shocks: a monetary policy shock, a neutral technology shock and a capital embodied technology shock. For example, a version of the model with plausible parameter values is able to replicate evidence that aggregate inflation responds slowly to a monetary policy shock while output, employment and other aggregate quantities respond strongly and in a hump-shape way. As an example of what we mean here by

[1] See, for example, Green's ( ) discussion of Woodford's ( ) book, which is the basic text of the New Keynesian model. Green asks the rhetorical question, 'what made the profession evolve from the earlier frictionless models of aggregate fluctuations, to the New Keynsian model with all its special features?' The answer, Green suggests, lies in the desire to construct models that generate plausible dynamic responses to shocks.

'plausible', is that the model does not need a larger amount of price and wage stickiness than seems warranted by micro evidence on prices.

Interestingly, our analysis does bring to light a price tension in the model. The New Keynesian model has difficulty simultaneously accounting for the rapid estimated response of inflation to a technology shock and the slow response of inflation to a monetary policy shock. Although the problem is perhaps not large in a statistical sense, it is potentially large in an economic sense. The price tension may be evidence of fundamental misspecification in the New Keynesian model, warranting the exploration of alternative model structures. Apart from this price tension, however, the confrontation of the New Keynesian model with impulse response functions is mostly a happy experience for the model.

But, there remain important challenges for DSGE models. The financial crisis of the past two years has alerted researchers to the absence of a serious financial sector in the consensus model. Progress on rectifying this problem is now well underway. Another challenge is that the consensus model has nothing to say about labor market variables like unemployment, the labor force, vacancies, etc. These variables are also of substantial interest to policy makers. We explore recent efforts to improve the labor market implications of the consensus DSGE model.

We describe two basic extensions of the consensus New Keynesian model which introduce an interesting model of the labor market. In each case, we apply the same impulse response matching methodology that was applied to the consensus New Keynesian model. This permits a close comparison of the different models.

The two basic extensions on the labor market explored here offer two different answers to the question, 'why there is unemployment?' Given the structure of the standard New Keynesian model, the simplest and most natural extension is one recently proposed by Gali (2009). In this case, the theory of unemployment is that wages are too high because of the presence of monopoly power in wage setting. This is a natural extension of the New Keynesian model because, as Gali shows, it primarily involves a reinterpretation of the variables in that model. Since wages are set by an entity with monopoly power in the New Keynesian model, wages are on average higher than the marginal cost of labor.[2] As a result, there will always be workers who respond 'yes' to the question, 'given the prevailing structure of wages, would you take a job if one were offered?'.

The second labor market extension that we explore takes a different position on the reason for unemployment. Its position is that people are constantly separating from their jobs for various reasons, while there are frictions in finding new work. The model also integrates the frictions that are so important in the fit of the standard New Keynesian model, for example, price and wage frictions. The labor market extension explored in this part of the paper

[2] The wage part of that model was proposed by Erceg, Henderson and Levin (2000).

builds on the line of work associated with Mortensen and Pissarides (1994), and the more recent work of Hall (2005a,b,c) and Shimer (2005a,b). Important recent contributions are those of Gertler and Trigari (2008), and Gertler, Sala and Trigari (2008) (henceforth GST). The latter work lays out a strategy for integrating the search and matching framework into an monetary model of aggregate fluctuations. Our presentation here builds specifically on the work of Christiano, Ilut, Motto, and Rostagno (2007) (CIMR), which follows GST.[3]

For the most part, the existing search and matching literature focuses on models in which employment fluctuations reflect variations in the job finding rate. The job separation rate is assumed to be constant over the business cycle. However, as emphasized by den Haan, Ramey and Watson (2000) and Fujita and Ramey (2008), this assumption is counterfactual.[4] We endogenize the job finding rate and find that how exactly this is done matters for how well the model reproduces aggregate fluctuations.

All the models are compared on the basis of their ability to replicate VAR-based facts about the response of standard macroeconomic variables to our three economic shocks. In addition, the models that are extended to account for the labor market are also evaluated on the basis of their ability to account for VAR-based facts about the response of labor market variables to shocks.

The paper is organized as follows. Section 2 describes our baseline model, the standard New Keynesian model. Section 3 discusses Gali's proposal to add unemployment to the model. Second 4 discusses the incorporation of labor adjustment costs into our baseline model. The search-based model of unemployment is developed in section 5. Section 6 describes our estimation strategy. Section 7 presents the results, which are not included in this draft. Section 8 will contain concluding remarks.

2. Baseline Model

Our baseline model is a version of the model in Christiano, Eichenbaum and Evans (2005) (CEE). We describe the objectives and constraints of the agents in the model, and leave the derivation of the equilibrium conditions to the appendix. We also discuss the motivation for key features of the model. In practice, those features are motivated by the VAR-based evidence on the dynamic response of macroeconomic variables to a monetary policy shock.

[3] Other work that integrates search-based labor market frictions into a DSGE model includes Thomas (2009).

[4] The job separation rate is defined as the number of people that move from employment to unemployment divided by the number of employed people.

2.1. Goods Production

An important feature of the baseline model is the price-setting frictions, which are designed to address the evidence of inertia in aggregate inflation. Obviously, the presence of price-setting frictions requires that firms have the power to set prices, and this in turn requires the presence of monopoly power. A challenge is to have an environment in which there is monopoly power, without running into conflict that there is a very large number of firms in actual economies. The Dixit-Stiglitz framework of production handles this challenge very nicely, because it has a very large number of price-setting monopolist firms.

In the Dixit-Stiglitz setup a homogeneous good, Y_t, is produced using

$$Y_t = \left[\int_0^1 Y_{i,t}^{\frac{1}{\lambda_f}} di \right]^{\lambda_f}, \; 1 \leq \lambda_f < \infty. \tag{2.1}$$

The good is produced by a competitive, representative firm which takes the price of output, P_t, and the price of inputs, $P_{i,t}$, as given.[5] The first order necessary condition associated with optimization is:

$$\left(\frac{P_t}{P_{i,t}} \right)^{\frac{\lambda_f}{\lambda_f - 1}} Y_t = Y_{i,t}. \tag{2.2}$$

A useful result is obtained by substituting out for Y_{it} in (2.1) from (2.2):

$$P_t = \left[\int_0^1 (P_{i,t})^{\frac{-1}{\lambda_f - 1}} di \right]^{-(\lambda_f - 1)}. \tag{2.3}$$

The firms with the monopoly power are the producers of the intermediate goods, $Y_{i,t}$, $i \in (0,1)$. There is one monopolist corresponding to each $i \in (0,1)$. Equation (2.2) represents the demand for the i^{th} good. The two variables chosen by the i^{th} monopolist are $P_{i,t}$ and $Y_{i,t}$, respectively. The intermediate good producer has no impact on the aggregate quantities, P_t and Y_t. These variables are the integral over all prices and quantities, respectively, and the i^{th} intermediate producer is correct to treat these as independent of the values of $P_{i,t}$ and $Y_{i,t}$.

The value of λ_f determines how much monopoly power the i^{th} firm has. If λ_f is large, then each intermediate good is a poor substitute for the others, and monopoly supplier of good i has a lot of market power. Consistent with this, note that if λ_f is large, then the demand for $Y_{i,t}$ is relatively price inelastic (see (2.2)). If λ_f is close to unity, so that each

[5] The idea of a representative firm is used often in macroeconomics. It is invoked when the technology is linear homogeneous and firms are competitive. In this case, it is easy to verify that economy-wide average output (Y_t in our case) is a function only of the aggregate quantity of each inputs, Y_{it}, $i \in (0,1)$. In particular, the size distribution of individual firms is indeterminate. This indeterminacy is accommodated by what is in effect a normalizing assumption: the assumption that all firms are identical. This firm is referred to as the 'representative firm'.

$Y_{i,t}$ is virtually a perfect substitute for the other, then i^{th} firm faces a demand curve that is virtually infinitely elastic. In this case, the firm has virtually no market power.

The i^{th} intermediate good producer has access to the following production technology:

$$Y_{i,t} = (z_t H_{i,t})^{1-\alpha} K_{i,t}^{\alpha} - z_t^+ \phi, \tag{2.4}$$

where $K_{i,t}$ denotes capital services used for production by the i^{th} intermediate good producer. Also, $\log(z_t)$ is a technology shock whose first difference has a positive mean and ϕ denotes a fixed production cost. The economy has two sources of growth: the positive drift in $\log(z_t)$ and a positive drift in $\log(\Psi_t)$, where Ψ_t is the state of an investment-specific technology shock discussed below. The object, z_t^+, in (2.4) is defined as follows:

$$z_t^+ = \Psi_t^{\frac{\alpha}{1-\alpha}} z_t.$$

Along a non-stochastic steady state growth path, Y_t/z_t^+ and $Y_{i,t}/z_t^+$ converge to constants. That is, z_t^+ represents the trend in the data, according to the model.

The shocks, z_t and Ψ_t, are specified to be unit root processes in order to be consistent with the assumptions we use in our VAR analysis to identify the dynamic response of the economy to neutral and capital-embodied technology shocks.

We assume that there is no entry or exit by intermediate good producers. The no entry assumption would be implausible if firms enjoyed large and persistent profits. The fixed cost in (2.4) is introduced to minimize the incentive to enter. We set ϕ so that intermediate good producer profits are zero in steady state. This requires that the fixed cost grows at the same rate as the growth rate of economic output, and this is why ϕ is multiplied by z_t^+ in (2.4). A potential empirical advantage of including fixed costs of production is that, by introducing some increasing returns to scale, the model can in principle account for evidence that labor productivity rises in the wake of a positive monetary policy shock.

In (2.4), $H_{i,t}$ denotes homogeneous labor services hired by the i^{th} intermediate good producer. Firms must borrow a fraction of the wage bill, so that one unit of labor costs is given by

$$W_t R_t^f,$$

where

$$R_t^f = \nu^f R_t + 1 - \nu^f. \tag{2.5}$$

Here, W_t denotes the aggregate wage rate, R_t denotes the gross nominal interest rate on working capital loans, and ν^f denotes the fraction of the wage bill that must be financed in advance. The assumption that firms require working capital was introduced by CEE as a way to help dampen the rise in inflation after an expansionary shock to monetary policy. An expansionary shock to monetary policy drives R_t down and - other things the same -

this reduces firm marginal cost. Inflation is dampened because marginal cost is the key input into firms' price-setting decision. Indirect evidence consistent with the working capital assumption includes the frequently-found VAR-based results, suggesting that inflation drops for a little while after a positive monetary policy shock.[6] It is hard to think of an alternative to the working capital assumption to explain this evidence, apart from the possibility that the estimated response reflects some kind of econometric specification error.[7]

Another motivation for treating interest rates as part of the cost of production has to do with the 'dis-inflationary boom' critique made by Ball (1994) of models that do not include interest rates in costs. Ball's critique focuses on the famous Phillips curve reduced form associated with New Keynesian models like the one studied here:

$$\pi_t = \beta E_t \pi_{t+1} + \gamma s_t,$$

where π_t denotes current aggregate inflation in P_t and s_t represents marginal cost, while $\gamma > 0$ is a reduced form parameter and β is nearly unity. According to the above equation, if the monetay authority announces it will fight inflation by strategies which (plausibly) bring down future inflation more than present inflation, then s_t must jump. In simple models s_t is directly related to the volume of output. High output requires more intense utilization of scare resources, their price goes up, driving up marginal cost, s_t. Ball criticized theories that do not include the interest rate in marginal cost on the grounds that we do not observe booms during disinflations. Including the interest rate in marginal cost potentially avoids the Ball critique because the high s_t may simply reflect the high interest rate that corresponds to the disinflationary policy, and not higher output.

If there were no price frictions, the i^{th} firm would simply set output in each period so that marginal cost equals marginal revenue, and price would be what is implied by the demand curve. However, we assume that firms cannot set price in each period. The mechanism determining when firms set their price follows the one proposed in Calvo ( ). In particular, with probability ξ_p the intermediate good firm cannot reoptimize its price, in which case it is assumed to set its price according to the following rule:

$$P_{i,t} = \tilde{\pi}_{f,t} P_{i,t-1}, \ \tilde{\pi}_{f,t} \equiv (\pi_{t-1})^{\kappa_f} (\bar{\pi})^{1-\kappa_f}, \tag{2.6}$$

where $\kappa_f \in (0,1)$ is a parameter, π_{t-1} is lagged (gross) inflation rate and $\bar{\pi}$ is the steady state inflation rate. Note that in steady state, firms that do not optimize price raise prices at the general rate of inflation. Firms that do optimize price in a steady state growth path

[6] Reference that evidence here. Reference evidence in the survey of manufacturers which suggests that short term borrowing is substantial, at least among manufacturing firms. See also Barth and Ramey and Christiano, Eichenbaum and Evans ( ).

[7] This possibility was suggested by Sims (1992) and explored further in Christiano, Eichenbaum and Evans (1999).

choose to also raise their price at the steady state rate of inflation. This is a key reason why all firms' prices are the same in the steady state of the model.

CEE introduced inflation indexation in (2.6) as a device to help moderate the immediate price response to a monetary policy shock. A firm that is able to reoptimize its price in the period that there is a positive monetary policy shock has less of an incentive to raise price immediately in the wake of such a shock when there is price indexing. Such a firm knows that in case it cannot reoptimize its price in later periods, its price will rise automatically with the slow rise in inflation that is expected to occur after a monetary policy shock.[8]

With probability $1-\xi_p$ the intermediate good firm can reoptimize its price. Apart from the fixed cost, the i^{th} intermediate good producer's profits are:

$$E_t \sum_{j=0}^{\infty} \beta^j \upsilon_{t+j} \{P_{i,t+j} Y_{i,t+j} - s_{t+j} P_{t+j} Y_{i,t+j}\},$$

where s_t denotes the marginal cost of production, denominated in units of the homogeneous good. The object, s_t, is a function only of the costs of capital and labor, and is described in Appendix A.4. Marginal cost is independent of the level of $Y_{i,t}$ because of the linear homogeneity of the first expression on the right of (2.4). In the firm's discounted profits, $\beta^j \upsilon_{t+j}$ is the multiplier on the household's nominal period $t+j$ budget constraint. Because they are the owners of the intermediate good firms, households are the recipients of firm profits. In this way, it is natural that the firm should weigh profits in different dates and states of nature using $\beta^j \upsilon_{t+j}$. In states of nature when the firm can reoptimize price, it does so to maximize its discounted profits, subject to the price setting frictions and to the requirement that it satisfy demand, (2.2). The first order necessary conditions associated with this optimization problem are reported in Appendix A.6.

Goods market clearing dictates that the homogeneous output good is allocated among alternative uses as follows:

$$Y_t = G_t + C_t + \tilde{I}_t. \tag{2.7}$$

Here, C_t denotes household consumption, G_t denotes exogenous government consumption and $\tilde{I}_t$ is a homogenous investment good which is defined as follows:

$$\tilde{I}_t = \frac{1}{\Psi_t} \left(I_t + a(u_t) \bar{K}_t \right). \tag{2.8}$$

The investment goods, I_t, are used to by households to add to the physical stock of capital, $\bar{K}_t$. The remaining investment goods are used to cover maintenance costs, $a(u_t)\bar{K}_t$, arising from capital utilization, u_t. The cost function, $a(\cdot)$, is increasing and convex, and has the

[8] CEE actually introduced 'full indexation', $\kappa_f = 1$. SW introduced the possibility, $0 < \kappa_f < 1$.

property that in steady state, $u_t = 1$ and $a(1) = 0$. The relationship between the utilization of capital, u_t, capital services, K_t, and the physical stock of capital, $\bar{K}_t$, is as follows:

$$K_t = u_t \bar{K}_t.$$

The investment and capital utilization decisions are discussed in section 2.2. See Appendix A.2 for the functional form of the capital utilization cost cost function. Finally, Ψ_t in (2.8) denotes a unit root investment specific technology shock with positive drift.

2.2. Households

2.2.1. Preferences

There is a continuum of households indexed by $j \in (0,1)$. The j^{th} household has the following preferences:

$$E_0 \sum_{t=0}^{\infty} \beta^t \left[\ln\left(C_t - bC_{t-1}\right) - A_L \frac{\left(h_{j,t+i}\right)^{1+\sigma_L}}{1+\sigma_L} \right], \ A_L, \sigma_L > 0, \ \beta \in (0,1). \tag{2.9}$$

Here, $h_{j,t}$ denotes the quantity of the j^{th} type of labor service supplied. The j^{th} household is the sole supplier of this type of labor service, which is imperfectly substitutable with the i^{th} household's labor service, for $i \neq j$. We explain the household's participation in the labor market in the next subsection. In (2.9), C_t and C_{t-1} denote the j^{th} household's consumption at dates t and $t-1$, respectively. As explained below, it is the presence of the appropriate insurance markets which guarantees that individual household consumption is the same for each $j \in (0,1)$.

The presence of $b > 0$ in (2.9) is motivated by VAR-based evidence like that displayed below, which suggests that a positive monetary policy triggers (i) a hump-shape response in consumption and (ii) a persistent reduction in the real rate of interest. With $b = 0$ and a utility function separable in labor and consumption like the one above, (i) and (ii) are difficult to reconcile. A positive monetary policy shock that triggers an increase in expected future consumption would be associated with rise in the real rate of interest, not a fall. Alternatively, a fall in the real interest rate would cause people to rearrange consumption intertemporally, so that consumption is relatively high right after the monetary shock and low later. Intuitively, one way to reconcile (i) and (ii) is to suppose the marginal utility of consumption is inversely proportional not to the level of consumption, but to its derivative. To see this, it is useful to recall the intertemporal Euler equation implied by household optimization:

$$E_t \beta \frac{u_{c,t+1}}{u_{c,t}} \frac{R_t}{\pi_{t+1}} = 1.$$

Here, $u_{c,t}$ denotes the marginal utility of consumption at time t. From this expression, we see that a low R_t/π_{t+1} tends to produce a high $u_{c,t+1}/u_{c,t}$, i.e., a rising trajectory for the marginal utility of consumption. This illustrates the problematic implication of the model when $u_{c,t}$ is inversely proportional to C_t as in (2.9) with $b = 0$. To fix this implication we need a model change which has the property that a rising $u_{c,t}$ path implies hump-shape consumption. A hump-shaped consumption path corresponds to a scenario in which the slope of the consumption path is falling, suggesting that (i) and (ii) can be reconciled if $u_{c,t}$ is proportional to the slope of consumption. The notion that marginal utility is inversely proportional to the slope of consumption corresponds loosely to $b > 0$.[9] The fact that (i) and (ii) can be reconciled with the assumption of habit persistence is of special interest, because there is evidence from other places that also favors the assumption of habit persistence, for example in asset pricing (see, for example, Constantinedes ( ) and Boldrin, Christiano and Fisher ( )) and growth (see Carroll ( )). In addition, there may be a solid foundation in psychology for this specification of preferences.[10]

The logic associated with the intertemporal Euler equation above suggests that there are other ways to reconcile (i) and (ii). For example, Guerron-Quintana (2008) shows that non-separability between consumption and labor in (2.9) can help reconcile (i) and (ii). He points out that if the marginal utility of consumption is an increasing function of labor and the model predicts that employment rises with a hump shape after a positive monetary injection, then it is possible that consumption itself rises with a hump-shape.

2.2.2. Wage Setting by Households

The model incorporates Calvo-style wage setting frictions along the lines spelled out in EHL. Because wages are an important component of costs, wage setting frictions help slow the response of inflation to a monetary policy shock. As in the case of prices, wage setting frictions require that there be market power. To make sure there is not too much market

[9] In particular, suppose first that lagged consumption in (2.9) represents aggregate, economy wide consumption and $b > 0$. This corresponds to the so-called 'external habit' case, where it is the lagged consumption of others that enters utility. In that case, the marginal utility of houeshold C_t is $1/\left(C_t - bC_{t-1}\right)$, which corresponds to the inverse of the slope of the consumption path, at least if b is large enough. In our model we think of C_{t-1} as corresponding to the household's own lagged consumption (that's why we use the same notation for current and lagged consumption), the so-called 'internal habit' case. In this case, the marginal utility of C_t also involves future terms, in addition to the inverse of the of the slope of consumption from $t = 1$ to t. The intuition described in the text, which implicitly assumed external habit, also applies roughly to the external habit case that we consider.

[10] Anyone who has gone swimming has experienced the psychological aspect of habit persistence. It is usually very hard at first to jump into the water because it seems so cold. The swimmer who jumps (or is pushed!) into the water after much procrastenation initially experiences a tremendous shock with the sudden drop in temperature. However, after only a few minutes the new, lower temperature is perfectly comfortable. In this way, the lagged temperature seems to influence one's experience of current temperature, as in habit persistence.

power, we follow EHL in adopting the Dixit-Stiglitz type framework used in the context of price-setting. The many households with specialized labor inputs, $h_{j,t}$ in (2.9) correspond to the many intermediate good firms producing specialized inputs.

We suppose that, with probability $1-\xi_w$, the j^{th} worker is able to reoptimize its wage and with probability ξ_w that worker sets $W_{j,t}$ according to the following rule:

$$W_{j,t+1} = \tilde{\pi}_{w,t+1} W_{j,t} \tag{2.10}$$
$$\tilde{\pi}_{w,t+1} = (\pi_t)^{\kappa_w} (\bar{\pi})^{(1-\kappa_w)} \mu_{z^+}, \tag{2.11}$$

where $\kappa_w \in (0,1)$. Note that in steady state, non-optimizing workers raise their real wage at the rate of growth of the economy. Because optimizing workers also do this in steady state, it follows that in the steady state, the wage of each type of worker is the same.

To understand the problem of the $1-\xi_w$ households which have the opportunity to reoptimize their wage, it is useful to understand the source of labor demand. We suppose that the labor power hired by intermediate good firms is homogeneous labor that is 'produced' by competitive labor contractors. Labor contractors produce homogeneous labor by aggregating the different types of specialized labor, $j \in (0,1)$, as follows:

$$H_t = \left[\int_0^1 (h_{t,j})^{\frac{1}{\lambda_w}} dj \right]^{\lambda_w}, \ 1 \leq \lambda_w < \infty. \tag{2.12}$$

Labor contractors take the wage rate of H_t and $h_{t,j}$ as given and equal to W_t and $W_{t,j}$, respectively. Profit maximization by labor contractors leads to the following first order necessary condition:

$$W_{j,t} = W_t \left(\frac{H_t}{h_{t,j}} \right)^{\frac{\lambda_w - 1}{\lambda_w}}. \tag{2.13}$$

Equation (2.13) is the demand curve for the j^{th} household's type of labor. We assume that this demand curve must be satisfied at each point in time, whether or not the household has the opportunity to reoptimize its wage. In considering (2.13), the j^{th} household correctly treats H_t and W_t as given and beyond its control.[11]

In principle, the idiosyncratic experiences of individual households will, over time, cause them to have different wealth holdings and therefore also different levels of consumption. Under these circumstances, aggregate economic outcomes may be dependent on the distribution of wealth across households. If so, then the distribution of wealth and the law of motion of that distribution must be solved for as part of the solution of the model. In practice, it is probably the case that solving models of the size considered in this paper is infeasible when there is non-trivial heterogeneity among households. For this reason, we follow EHL

[11] Substituting out for $h_{t,j}$ in (2.12) using (2.13) we obtain an expression relating W_t to $W_{j,t}$ for $j \in (0,1)$ that is analogous to (2.3).

in adopting the assumption that there are insurance markets on the realization of the Calvo uncertainty determining whether the household can or cannot adjust its wage.

It is possible, however, that the heterogeneity induced by idiosyncratic uncertainty in the setting of wages may only have a negligible impact on aggregate outcomes. This possibility is suggested by the recent literature on versions of the baseline model in which intermediate good firms have a state variable, such as the capital stock or a stock of employment. In these models, idiosyncratic uncertainty about the timing of price reoptimization gives rise to a distribution across firms of their state variable. In principle, this matters for aggregate outcomes in the same way that idiosyncratic uncertainty about the timing of wage reoptimization on the part of households might matter. In the context of intermediate good firms, Woodford (2004) has shown that as long as (i) the variables in a stochastic equilibrium are not too far from their value in non-stochastic steady state and (ii) agents in the non-stochastic steady state are identical, then standard linearization methods can be applied and the details of the distribution of firms by their state variable do not matter for economic aggregates.[12] The intuition is simple. Condition (i) guarantees that individual firm decision rules are well approximated by linearizing about the steady state and condition (ii) guarantees that all those decision rules have the same intercept and slope coefficients. To understand how this guarantees that the details of the microeconomic distribution of state variables does not matter, consider the old-fashioned Keynesian consumption function:

$$C = \alpha + \beta Y,$$

where C denotes consumption and Y denotes household income. In reality, different households have different levels of income and in principle poor households have different β's than rich ones. If this were the case, then a simple relationship like the one above relating aggregate consumption to aggregate income would not exist: to predict C one would have to know not only aggregate Y, but also how it is distributed among rich and poor people. However, if everyone - poor and rich alike - all had consumption functions with the same slope and intercept terms, then aggregate consumption would be determined from aggregate income as in the old-fashioned Keynesian consumption function. This insight has been applied with success to models in which firms have idiosyncratic state variables, and it may also work in models like the present one in which households have different wealth levels because of the effects of idiosyncratic realizations of the ability to set wages. A challenge for the approach would be to ensure condition (i). This requires that the model incorporate forces that pull

[12] A write-up of the method in a simplified model appears in Christiano ( ). An application in a medium-scaled model with size approximating the size of our model appears in ACEL. See also Sveen and Weinke ( ) for an alternative strategy for solving a model with firm-specific factors. Other studies of models with firm-specific factors include....

the distribution of wealth across households back together after a disturbance has pulled it apart.

2.2.3. Accumulation of Capital

The household owns the economy's physical stock of capital, sets the utilization rate of capital and rents the services of capital in a competitive market. The household accumulates capital using the following technology:

$$\bar{K}_{t+1} = (1-\delta)\,\bar{K}_t + F\left(I_t, I_{t-1}\right) + \Delta_t, \tag{2.14}$$

where Δ_t denotes physical capital purchased in a market with other households. Since all households are the same in terms of capital accumulation decisions, $\Delta_t = 0$ in equilibrium. We nevertheless include Δ_t so that we can assign a price to installed capital. In (2.14), $\delta \in (0,1)$ and we use the specification suggested in CEE:

$$F\left(I_t, I_{t-1}\right) = \left(1 - \tilde{S}\left(\frac{I_t}{I_{t-1}}\right)\right) I_t,\ \ \tilde{S}\left(x\right) \equiv \frac{S''}{2}\left(x - \mu_I\right)^2. \tag{2.15}$$

Here, $S'' > 0$ and μ_I denotes the nonstochastic steady state growth rate of I_t. The technology, (2.14) and (2.15), represent the technology for producing capital from investment goods, and is conventionally referred to as adjustment costs.

Let $P_t P_{k',t}$ denote the nominal market price of Δ_t. For each unit of $\bar{K}_{t+1}$ acquired in period t, the household receives X_{t+1}^k in net cash payments in period $t+1:$

$$X_{t+1}^k = u_{t+1} P_{t+1} r_{t+1}^k - \frac{P_{t+1}}{\Psi_{t+1}} a(u_{t+1}). \tag{2.16}$$

The first term is the gross nominal period $t+1$ rental income from a unit of $\bar{K}_{t+1}$. The second term represents the cost of capital utilization, $a(u_{t+1})P_{t+1}/\Psi_{t+1}$. Here, P_{t+1}/Ψ_{t+1} is the nominal price of the investment goods absorbed by capital utilization. That P_{t+1}/Ψ_{t+1} is the equilibrium market price of investment goods follows from the technology specified in (2.7) and (2.8), and the assumption that investment goods are produced from homogeneous output goods by competitive firms.

The introduction of variable capital utilization is motivated by a desire to explain the slow response of inflation to a monetary policy shock. In the baseline model, prices are heavily influenced by costs. These in turn are influenced by the elasticity of the factors of production. If factors can be rapidly expanded with a small rise in cost, then inflation will not rise much after a monetary policy shock. Allowing for variable capital utilization is a way to make the services of capital elastic. If there is very little curvature in the a function, then households are able to expand capital services without much increase in cost.

The the form of the investment adjustment costs in (2.14) is motivated by a desire to reproduce VAR-based evidence that investment has a hump-shaped response to a monetary policy shock. Alternative specifications include $F \equiv I_t$ and

$$F = I_t - \frac{S''}{2}\left(\frac{I_t}{K_t} - \delta\right)^2 K_t. \tag{2.17}$$

Specification (2.17) has a long history in macroeconomics, and has been in use since at least Lucas and Prescott ( ). To understand why DSGE models generally use the adjustment cost specification in (2.15) rather than (2.17), it is useful to define the rate of return on invesment:

$$R_{t+1}^k = \frac{x_{t+1}^k + \left[1 - \delta + S''\left(\frac{I_{t+1}}{K_{t+1}} - \delta\right)\frac{I_{t+1}}{K_{t+1}} - \frac{S''}{2}\left(\frac{I_{t+1}}{K_{t+1}} - \delta\right)^2\right]P_{k',t+1}}{P_{k',t}}. \tag{2.18}$$

The numerator is the one-period payoff from an exra unit of $\bar{K}_{t+1}$, and the denominator is the corresponding cost, both in consumption units. In (2.18), $x_{t+1}^k \equiv X_{t+1}^k/P_{t+1}$ denotes the earnings net of costs. The term in square brackets is the quantity of additional $\bar{K}_{t+2}$ made possible by the additional unit of $\bar{K}_{t+1}$. This is composed of the undepreciated part of $\bar{K}_{t+1}$ left over after production in period $t+1$, plus the impact of $\bar{K}_{t+1}$ on $\bar{K}_{t+2}$ via the adjustment costs. The object in square brackets is converted to consumption units using $P_{k',t+1}$, which is the market price of $\bar{K}_{t+2}$ denominated in consumption goods. Finally, the denominator is the price of the extra unit of $\bar{K}_{t+1}$.

The price of extra capital, in competitive markets corresponds to the marginal cost of production. Thus,

$$\begin{aligned} P_{k',t} &= -\frac{dC_t}{d\bar{K}_{t+1}} = -\frac{dC_t}{dI_t} \times \frac{dI_t}{d\bar{K}_{t+1}} \\ &= \frac{1}{\Psi_t}\left[\frac{1}{\frac{d\bar{K}_{t+1}}{dI_t}}\right] = \frac{1}{\Psi_t}\begin{cases} 1 & F = I \\ \frac{1}{1 - S'' \times \left(\frac{I_t}{K_t} - \delta\right)} & F \text{ in } (2.17) \end{cases}. \end{aligned} \tag{2.19}$$

The derivatives in the first line correspond to marginal rates of technical transformation. The marginal rate of technical transformation between consumption and investment is implicit in (2.7) and (2.8). The marginal rate of technical transformation between I_t and $\bar{K}_{t+1}$ is given by the capital accumulation equation. The relation in the second line of (2.19) is referred to as 'Tobin's q' relation, where Tobin's q here corresponds to $\Psi_t P_{k',t}$. This is the market value of capital divided by the price of investment goods. Here, q can differ from unity due to the investment adjustment costs.

We are now in a position to convey the intuition about why DSGE models have generally abandoned the specification in (2.17) in favor of (2.14). The key reason has to do with VAR-based evidence that suggests the real interest rate falls persistently after a positive monetary policy shock, while investment responds in a hump-shaped pattern. Any model that is capable of producing this type of response will have the property that the real return on capital, (2.18) - for arbitrage reasons - also falls after an expansionary monetary policy shock. Suppose, to begin, that $S'' = 0$, so that there are no adjustment costs at all and $P_{k',t} = 1$. In this case, the only component in R_t^k that can fall is x_{t+1}^k, which is dominated by the marginal product of capital. That is, approximately, the rate of return on capital is:

$$K_{t+1}^{\alpha-1} H_{t+1}^{1-\alpha} + 1 - \delta.$$

In steady state this object is $1/\beta$ (ignoring growth), which is roughly 1.03 in annual terms. At the same time, the object, $1 - \delta$, is roughly 0.9 in annual terms, so that the endogenous part of the rate of return of capital is a very small part of that rate of return. As a result, any given drop in the return on capital requires a very large percentage drop in the endogenous part, $K_{t+1}^{\alpha-1} H_{t+1}^{1-\alpha}$. An expansion in investment can accomplish this, but it has to be a very substantial surge. To see this, note that the endogenous part of the rate of return is not only small, but the capital stock receives a weight substantially less than unity in that expression. Moreover, a model that successfuly reproduces the VAR-based evidence that employment rises after a positive monetary policy implies that hours worked rises. This pushes the endogenous component up, increasing the burden on the capital stock to bring down the rate of return on investment. For these reasons, models without adjustment costs generally imply a counterfactually strong surge in invesment in the wake of a positive shock to monetary policy.

With $S'' > 0$ the endogenous component of the rate of return on capital is much larger. However, in practice models that adopt the adjustment cost specification, (2.17), generally imply that the biggest investment response occurs in the period of the shock, and not later. To gain intuition into why this is so, suppose the contrary: that investment does exhibit a hump-shape response in investment. Equation (2.19) implies a similar hump-shape pattern in the price of capital, $P_{k',t}$.[13] This is because that $P_{k',t}$ is primarily determined by the contemporaneous flow of investment. So, under our supposition about the investment response, a positive the monetary policy shock generates a rise in $P_{k',t+1}/P_{k',t}$ over at least several periods in the future. According to (2.18), creates the expectation of future capital gains, $P_{k',t+1}/P_{k',t} > 1$ and increases the immediate response of the rate of return on capital. Thus, households would be induced to substitute away from a hump shape response, towards one in which the immediate response is much stronger. In practice, this means that in equilibrium,

[13] Note from (2.19) that the price of capital increases as investment rises above its level in steady state, which is the level required to just meet the depreciation in the capital stock. Our assertion that the price of capital follows the same hump shape pattern as investment after a positive monetary policy shock reflects our implicit assumption tht the shock occurs when the economy is in a steady state. This will be true on average, but not at each date.

the biggest response of investment occurs in the period of the shock, with later responses converging to zero.

The adjustment costs in (2.15) do have the implication that investment responds in a hump-shape. The reason is (2.15)'s implication that a quick rise in investment from previous levels is expensive.

There are other reasons to take the specification in (2.15) seriously. Luca ( ) and Matsuyama (1982) have described interesting theoretical foundations which produce (2.15) as a reduced form. For example, in Matsuyama, shifting production between goods and capital involves a learning by doing process, which makes quick movements in either direction expensive. Also, Matsuyama explains how the abundance of empirical evidence that appears to reject (2.17) may be consistent with (2.15). Consistent with (2.15), Rosen ( ) argues that data on housing construction cannot be understood without using a cost function that involves the change in the flow of housing constuction.

2.2.4. Household Optimization Problem

The j^{th} household's period t budget constraint is as follows:

$$P_t \left(C_t + \frac{1}{\Psi_t} I_t \right) + B_{t+1} \leq W_{t,j} h_{t,j} dj + X_t^k \bar{K}_t + R_{t-1} B_t + a_{jt} \tag{2.20}$$

where $W_{t,j}$ represents the wage earned by the j^{th} household, B_{t+1} denotes the quantity of risk-free bonds purchased by the household, R_t denotes the gross nominal interest rate on bonds purchased in period $t-1$ which pay off in period t, and a_{jt} denotes the payments and receipts associated with the insurance on the timing of wage reoptimization. The household's problem is to select sequences, $\left\{ C_t, I_t, W_{t,j}, B_{t+1}, \bar{K}_{t+1} \right\}$, to maximize (2.9) subject to (2.13), (2.10), (2.11), (2.14), (2.16), (2.20) and the mechanism determining when wages can be reoptimized.

2.3. Fiscal and Monetary Authorities, and Equilibrium

We suppose that monetary policy follows a Taylor rule of the following form:

$$\log \left(\frac{R_t}{R} \right) = \rho_R \log \left(\frac{R_{t-1}}{R} \right) + (1 - \rho_R) \left[r_\pi \log \left(\frac{\pi_{t+1}}{\bar{\pi}} \right) + r_y \log \left(\frac{gdp_t}{gdp} \right) \right] + \varepsilon_{R,t}. \tag{2.21}$$

where gdp_t denotes scaled real GDP defined as:

$$gdp_t = \frac{G_t + C_t + I_t}{z_t^+}, \tag{2.22}$$

and gdp denotes the nonstochastic steady state value of gdp_t, and G_t denotes government consumption. We adopt the model of government spending suggested in Christiano and Eichenbaum (1992), in which

$$G_t = gz_t^+.$$

In principle, g could be a random variable, though our focus in this paper is just on monetary policy and technology shocks. So, we set g to a constant. Lump-sum transfers are assumed to balance the government budget.

An equilibrium is a stochastic process for the prices and quantities which has the property that the household and firm problems are satisfied, and goods and labor markets clear.

2.4. Adjusment Cost Functions and Shock Processes

We adopt the following functional forms. The capacity utilization cost function is:

$$a(u) = 0.5b\sigma_a u^2 + b(1-\sigma_a)u + b((\sigma_a/2) - 1),$$

where b is selected so that $a(1) = a'(1) = 0$ in steady state and σ_a is a parameter that controls the curvature of the cost function. The closer σ_a is to zero, the less curvature there is and the easier it is to change utilization. The investment adjustment cost function takes the following form:

$$\begin{aligned} \tilde{S}(x_t) &= \frac{1}{2}\left\{\exp\left[\sqrt{\tilde{S}''}(x_t - \mu_{z^+}\mu_\Psi)\right] + \exp\left[-\sqrt{\tilde{S}''}(x_t - \mu_{z^+}\mu_\Psi)\right] - 2\right\}, \\ &= 0,\ x = \mu_{z^+}\mu_\Psi. \end{aligned}$$

where $x_t = I_t/I_{t-1}$ and $\mu_{z^+}\mu_\Psi$ is the growth rate of investment in steady state. With this adjustment cost function, $\tilde{S}(\mu_{z^+}\mu_\Psi) = \tilde{S}'(\mu_{z^+}\mu_\Psi) = 0$. Also, $\tilde{S}'' > 0$ is a parameter having the property that it is the second derivative of $\tilde{S}(x_t)$ evaluated at $x_t = \mu_{z^+}\mu_\Psi$. Because of the nature of the above adjustment cost functions, the curvature parameters have no impact on the model's steady state.

We assume that the neutral technology shock evolves as follows:

$$\begin{aligned} \Delta\log z_t &= \mu_z + \rho_z\Delta\log z_{t-1} + \varepsilon_t^z,\ E(\varepsilon_t^z)^2 = (\sigma_z)^2 \\ \Delta\log\Psi_t &= \mu_\Psi + \rho_\Psi\Delta\log\Psi_{t-1} + \varepsilon_t^\Psi,\ E\left(\varepsilon_t^\Psi\right)^2 = (\sigma_\Psi)^2. \end{aligned}$$

Note that with these specifications, shocks to ε_t^z and ε_t^Ψ have permanent effects on the levels of $\log z_t$ and $\log\Psi_t$, respectively. This means that the shocks also affect the level of labor productivity in the long run, making the model consistent with the the identifying assumptions used in the VAR analysis to deduce the dynamic effects in the data in response to disturbances in ε_t^z and ε_t^Ψ.[14]

[14] Interestingly, our assumption implies that there are 'two unit roots' in the data, not just one.

3. Extending the Baseline Model to Include Unemployment

A shortcoming of the baseline model is that it is silent about labor market variables like unemployment and the labor force, even though these are objects of central interest in monetary policy analysis. In a recent paper, Gali (2009) shows how the labor market variable in the baseline model can be interpreted as a quantity of people working.[15] With his adjustment, Gali (2009) shows that the baseline model very naturally yields a theory of the labor force and of unemployment. Unemployment is high and fluctuating because monopoly power keeps the wage too high on average and because of frictions in price adjustment.

Interestingly, Gali's adjustment to the baseline model only involves a reinterpretation of the model variables. The equations that characterize the equilibrium in the baseline model are left unaffected. Gali's suggestion allows one to solve the baseline model exactly as before, and provides additional equilibrium conditions that allow one to recursively deduce the model implications for the labor force and the unemployment rate. His suggestions does not increase the number of model parameters. From an econometric standpoint, it increases the set of data that can be used to pin down the values of the baseline model parameters.

Gali assumes that there is a large number of people capable of offering each different type of labor service, $j \in (0, 1)$. Consider a fixed $j \in (0, 1)$ that corresponds, say, to plumbers. Suppose the prevailing wage rate for plumbers is $W_{j,t}$. Given $W_{j,t}$, the demand for plumbers, (2.13), determines how many plumbers, $h_{j,t}$, are working. There is a number of plumbers, $U_{j,t}$, that are not working, though they would respond 'yes' if asked wether they would like to work at the wage, $W_{j,t}$. The presence of employed and unemployed plumbers requires some sort of worker selection mechanism. Gali assumes the selection is done on the basis of a particular type of worker heterogeneity.

Every worker of every specific type j has an index, $l \geq 0$, written on their forehead which communicates two different pieces of information. First, it indicates how much disutility the worker suffers by working:

$$\zeta_t A_L l^{\sigma_L},\ \sigma_L, A_L > 0.$$

Here, ζ_t is a labor preference shifter, which we discuss below. The index, l, on a worker's forehead also indicates how many workers have a lower disutility of work. Given the quantity of type j worker demand, $h_{j,t}$, workers with index, $l = 0$ up to $l = h_{j,t}$ go to work. Because there is market power in setting the wage rate, some workers, $l > h_{j,t}$, would also like to go to work given the wage rate, $W_{j,t}$. These are the unemployed workers, $U_{j,t}$. The total disutility of work experienced by type j workers if $h_{j,t}$ are employed is given by:

$$\zeta_t A_L \int_0^{h_{j,t}} l^{\sigma_L} dl = \zeta_t A_L \frac{h_{j,t}^{\sigma_L + 1}}{\sigma_L + 1}.$$

[15] In his specification, there is no variation in the number of hours per worker.

Note the similarity of the term on the right of the equality to the labor term in (2.9).

A worker of type j is defined as unemployed if his index is $l > h_{j,t}$ (i.e., he is not working) and he answers 'yes' to the question, "if you were offered a job at the prevailing wage for your type of labor, would you take it?" He would answer yes to this question if the wage rate, $W_{j,t}$, exceeded his marginal cost of working, that is, if

$$W_{j,t} > \frac{\zeta_t A_L l^{\sigma_L}}{\upsilon_t},$$

where υ_t is the multiplier on the household budget constraint. That is, υ_t converts the marginal disutility of work (the numerator in the above expression) into currency units for comparison with the nominal wage, $W_{j,t}$. Let $l_{t,j}^*$ denote the value of l that enforces the above equation hold as a strict equality. The object, $l_{t,j}^*$, is the disutility index of the marginal worker who is just indifferent between working and not working. This index is also the quantity of workers with disutility index equal to $l_{t,j}^*$ or less. As a result, $l_{t,j}^*$ is the workforce of type j workers. Type j workers with disutility indices in the interval, $\left(h_{j,t}, l_{j,t}^*\right)$ are unemployed. The economy's overall unemployment rate is defined as:

$$u_{n,t} = \frac{\int_0^1 \left[l_{j,t}^* - h_{j,t}\right] dj}{\int_0^1 l_{j,t}^* dj}.$$

The quantity of people unemployed, $l_{j,t}^* - h_{j,t}$, is expected to always be positive, because of the way the wage rate is set in this economy. We turn to this now.

It remains to discuss how wages are determined. To fill in the required details, we have to take one of two different directions. We could suppose - as we did in the baseline model - that each household has only type j workers, or that each household has every type of worker. The equilibrium conditions associated with these alternative views are identical and so the two directions are observationally equivalent from the point of view of the aggregate data. Still, different researchers may feel more comfortable with one or the other interpretation based on other evidence.

The above alternative interpretations are also available in the baseline model. However, in the baseline model the assumption that each household has only a type j worker seems obviously more compelling. This is because of the realism of the notion that there is a small number (unity, in this case) of workers in each household. However, with the Gali interpretation, the small household assumption is not an option because now there are many workers offering each type, j, of labor services. As a result, we have a mild preference for Gali (2009)'s own interpretation, in which each household contains every type, $j \in (0, 1)$, of worker. Thus, each household is identical since it includes workers with every type of labor service, j, and every type of disutility of work, l.

Returning to the j that corresponds to plumbers, we assume that all plumbers in each household gather in a single plumbers union. That union takes the demand for plumbers, (2.13), as given and chooses the wage rate to maximize the discounted utility of its members, subject to Calvo wage setting frictions. With probability $1-\xi_w$, the wage rate is set according to (2.10) and (2.11). With the complementary probability, the union has the opportunity to reset its wage. When it does so, it need only concern itself with future histories in which it does not have the opportunity to reoptimize. That is, it the wage is set to optimize:

$$E_t \sum_{i=0}^{\infty} (\beta\xi_w)^i \left[-\zeta_t A_L \frac{(h_{j,t+i})^{1+\sigma_L}}{1+\sigma_L} + \upsilon_{t+i} W_{j,t+i} h_{j,t+i} \right],$$

subject to (2.13), and $W_{j,t+i}$, $i > 0$, satisfying (2.10) and (2.11). Here, υ_{t+i} is the multiplier on the household budget constraint, which measures the marginal value to a worker of a unit of currency. The mechanism that determines the wage rate in this model is identical to the mechanism by which it is set in the baseline model. This is why the equilibrium allocations in the two models coincide, as long as $\zeta_t \equiv 1$. The difference between this model and the baseline model is that with the reinterpretation of variables, it is now possible to deduce implications for unemployment and the labor force.

To gain intuition about the determinants of unemployment in the model, it is useful to see the connection between unemployment and the markup. Let $\lambda_{w,j,t}$ denote the markup of the wage, $W_{j,t}$, over the marginal type j worker's cost of providing labor. In steady state, $\lambda_{w,j,t} = \lambda_w$, the labor substitutability parameter in (2.12). The markup, $\lambda_{w,j,t}$, fluctuates over t and j in response to shocks because of wage-setting frictions. By the definition of the markup,

$$W_{j,t} = \lambda_{w,j,t}\zeta_t A_L \frac{h_{j,t}^{\sigma_L}}{\upsilon_t}.$$

Recalling the definition of the labor force of type j workers, $l_{j,t}^*$:

$$W_{j,t} = A_L \zeta_t \frac{\left(l_{j,t}^*\right)^{\sigma_L}}{\upsilon_t}. \tag{3.1}$$

Using the second equation to substitute out for $W_{j,t}$ in the first equation and rearranging, we obtain:

$$\frac{l_{j,t}^* - h_{j,t}}{l_{j,t}^*} = 1 - \frac{1}{\lambda_{w,j,t}^{\frac{1}{\sigma_L}}}.$$

The expression on the left of the equality is the unemployment rate of type j workers in steady state. Note that it is positive as long as $\lambda_{w,j,t} > 1$. The latter condition holds in steady state and also in response to shocks in case λ_w is big enough and the shocks are small enough. We conclude that in the steady state,

$$u_n = 1 - \lambda_w^{-\frac{1}{\sigma_L}}. \tag{3.2}$$

Thus, the model provides a link between the (Frisch) labor supply elasticity, $1/\sigma_L$, the markup, λ_w, and the unemployment rate, u_n. As labor supply is less elastic, a higher markup is required to support a given level of unemployment.

Herein lies a potential problem for the model. Empirically, the labor force exhibits relatively little variation. However, according to (3.1), the labor force equation is static and implies an immediate response to flucttutations in wages and the value of work, υ_t. This dependence could in principle be removed simply by setting σ_L to a large number. However, (3.2) indicates that this would require a large value of λ_w if the model is to be consistent with an empirically plausible average rate of unemployment.

An alternative strategy works with ζ_t. Suppose the labor supply shifter has the following form:

$$\zeta_t = \left(\frac{l_t^*}{l^*}\right)^{\psi}, \ \psi \geq 0,$$

where l_t^* is the economy-wide average labor force and l^* is its steady state value. Note that the monopoly union does not internalize its impact on l_t^*, since the j^{th} monopoly union is small relative to the economy as a whole. A large value of ψ has an impact on (3.1) that resembles the impact of simply adding ψ to σ_L in that equation. In particular, with ψ positive, the 'labor force supply elasticity' in the labor force equation becomes zero, and the labor force is roughly constant. Of course, $\psi > 0$ also reduces the labor supply elasticity for labor in the model, and this can be expected to have effects elsewhere in the model. However, with ξ_w sufficiently high, we may suppose that labor supply has very little impact on the equilibrium. For example, consider the extreme case where the wage rate is literally fixed. In this case, because labor is demand determined, the labor supply equation is completely irrelevant.

In practice, we also considered other specifications of ζ_t. In particular, we considered the case in which ζ_t is a function of the level and growth rate of aggregate employment, h_t.

4. Baseline Model With Employment Adjustment Costs

We modify the baseline model so that the intermediate good firm's marginal cost is increasing in its own output. ACEL found that without such a change, an estimation exercise like ours is likely to result in an implausibly high estimated value for ξ_p, the degree of stickiness in goods prices. Because of our focus on labor markets in this paper, a natural way to make marginal cost positively sloped is to suppose that it is difficult to quickly expand the amount of labor services used in production.

We assume that the i^{th} firm's current employment, $H_{i,t}$, is a state variable. In period t

the only thing it can do about employment is to adjust $H_{i,t+1}$ by purchasing

$$z_t^+ V\left(\frac{H_{i,t+1}}{H_{i,t}}\right),\ V(x) = \frac{a}{2}(x-1)^2,\ a \geq 0,$$

final goods.[16] We assume that these costs of adjusting labor are internal to the firm. For example, it might be that the firm needs to build and furnish cubicles for people to work in if they want to expand jobs and they need to dismantle the cubicles when they reduce employment. We continue to maintain the assumption that firms hire labor in competitive markets at the exogenously specified wage rate, W_t. To accommodate this change, we must also change the clearing condition in the market for final goods:

$$Y_t = G_t + C_t + \tilde{I}_t + \int_0^1 z_t^+ V\left(\frac{H_{i,t+1}}{H_{i,t}}\right) di.$$

The fact that an intermediate good firm's employment is a state variable at time t means that the only way it can accommodate an unexpected rise in demand is to hire more capital services. Diminishing returns in capital services imply higher marginal production costs as output expands. It is well known that a positively sloped marginal cost curve implies that a firm's price responds less to an outside shock to supply (say, a rise in the rental rate of capital or a rise in the wage rate).[17] In estimated Phillips curves, the slope coefficient on marginal cost is typically estimated to be very small. Absent a real rigidity like the one we employ, the small slope coefficient must be explained by a stickiness in prices. With a real rigidity, the need for stickiness in prices is reduced. Another way to see the potential importance of the type of real rigidity we introduce here, is that it helps account for the small rise in inflation that seems to occur in the wake of a positive monetary policy shock.[18]

In our empirical analysis we use firm vacancies as indicators that firms are undertaking efforts to increase employment. From the perspective of the model used here, the vacancies are only indicators that efforts are being undertaking inside the firm to raise employment. They are not treated as indicators of the importance of search costs in changing the level of employment. These costs are assumed to be non-existent in the model of this section and the previous section. We measure the time t vacancies of the i^{th} firm as $H_{i,t+1} - H_{i,t}$ if this

[16] For a similar approach, see Thomas (2009). In our paper, the real rigidity comes about because of internal adjustment costs in employment. In Thomas, it comes about because of adjustment costs arising from search frictions.

[17] The intuition is simple. Suppose something causes the firm to contemplate an increase in its price. Given the slope of its demand curve, this translates into a reduction in output. With less output, marginal cost is lower. Other things the same, the reduced marginal cost cuts into the original incentive to raise price.

[18] For another paper that explores the potential for firm-specific factors to account for the slow response of inflation to shocks, see de Walque, Smets and Wouters (2006). For other real rigidities can have similar effects see Ball and Romer (1990), Eichenbaum and Fisher (2007), and Kimball (1995),

term is positive. That is economy-wide vacancies are given by:

$$\int_0^1 I\left(H_{i,t+1} - H_{i,t}\right) di,$$

where $I(x) \equiv x$ if $x \geq 0$ and $I(x) \equiv 0$ otherwise. In the steady state of the model there are no vacancies, since $H_{i,t+1} = H_{i,t}$. This is an artifact of the fact that all our model variables are expressed in per capita terms. We expect that the discrepancy between the level of vacancies in our model's steady state and the level of vacancies in the data does not influence the results of our analysis. That is because we compare our model's predictions for the change in vacancies in response to a shock, with a corresponding empirical estimate.

An advantage of evaluating the model in this section using vacancy data is that it helps put the model on a comparable basis with the models that emphasize search frictions in the labor market in the next section.

Because price-setting intermediate good firms are heterogeneous in terms of their histories of Calvo price frictions, in a stochastic equilibrium each firm's state variable is different. This introduces a potentially formidable technical problem for solving the model. The standard solution to models with Calvo-style price frictions is massively simplified by the assumption that firms have no state variable. What makes solving our model potentially very difficult is that the distribution across firms of their state variable potentially matters for the response of macroeconomic aggregates to a shock. For example, suppose that the number of people working in the economy is a given magnitude. The inflation effect of a monetary policy disturbance may be one thing when those people are distributed roughly evenly among firms, and a very different thing if a disproportionate number of people are concentrated in a small number of firms. This would be the case if the response of a firm with a lot of employment to a shock is very different from the response of a firm with a small amount of employment.

Woodford (2004) has shown that as long as (i) the variables in a stochastic equilibrium are not too far from their value in non-stochastic steady state and (ii) firms in the non-stochastic steady state are identical, then standard linearization methods can be applied and the details of the distribution of firms by their state variable do not matter for economic aggregates.[19] Condition (i) guarantees that individual firm decision rules are well approximated by linearizing about the steady state and condition (ii) guarantees that all those decision rules have the same intercept and slope coefficients. To understand how this guarantees that the details of the microeconomic distribution of state variables does not matter, consider the

[19] A write-up of the method in a simplified model appears in Christiano ( ). An application in a medium-scaled model with size approximating the size of our model appears in ACEL. See also Sveen and Weinke ( ) for an alternative strategy for solving a model with firm-specific factors. Other studies of models with firm-specific factors include....

old-fashioned Keynesian consumption function:

$$C = \alpha + \beta Y,$$

where C denotes consumption and Y denotes household income. In reality, different households have different levels of income and in principle poor households have different β's than rich ones. If this were the case, then a simple relationship like the one above relating aggregate consumption to aggregate income would not exist: to predict C one would have to know not only aggregate Y, but also how it is distributed among rich and poor people. However, if everyone - poor and rich alike - all had consumption functions with the same slope and intercept terms, then aggregate consumption would be determined from aggregate income as in the old-fashioned Keynesian consumption function. The technical details of how the model described in this section is solved appear in the appendix.

5. Search-based Representation of the Labor Market

An active line of research on the labor market is the one associated with the work of Mortensen and Pissarides (1994) and, more recently, Hall (2005a,b,c) and Shimer (2005a,b). This work pursues the idea that fluctuations in unemployment reflect frictions associated with finding jobs. The details of the approach taken here follows the version of the Gertler, Sala and Trigari (2006) (henceforth GST) strategy implemented in Christiano, Ilut, Motto, and Rostagno (2007) (CIMR). This approach integrates wage-setting frictions into the search and bargaining framework by building on an insight in Hall (2005a).

GST adopt Calvo-style wage frictions, following the approach suggested by EHL and implemented in the baseline model. We follow the CIMR approach of instead using so-called Taylor contracts. These are labor contracts that have a fixed duration and are which are negotiated at fixed points in time. We adopt this approach here because of its relative simplicity. In addition, the rigid structure of Taylor wage setting may have some appeal based on realism.

The CIMR and GST approaches both assume that job separations are exogenous. Motivated in part by the empirical work of Fujita and Ramey (2006), who show that separations have an important cyclical component, we endogenize job separations. This requires introducing worker heterogeneity as well as a criterion for determining who leaves and who stays. We adopt two alternative selection criteria. The first, the *Total Surplus Criterion*, specifies that workers for whom the total surplus is non-negative remain while the others separate. The second, the *Firm Surplus Criterion*, specifies that the separation decision is based on the firm's surplus only.

Wages are determined by Nash bargaining. In our formulation, bargaining occurs between an individual worker and a representative of the firm, taking all other worker bargaining outcomes within the firm as given.

5.1. Sketch of the Model

Figures 1 and 2 provide graphical illustrations of the structure and the timing in our search-based representation of the labor market. We adopt the Dixit-Stiglitz specification of production, as in the baseline model. A representative, competitive retail firm aggregates differentiated intermediate goods into a homogeneous good using (2.1). Intermediate goods are supplied by monopolists, who hire labor and capital services in competitive factor markets. The intermediate good firms are assumed to be subject to the same Calvo price setting frictions in the baseline model.

In the baseline model, homogeneous labor services are produced from differentiated labor using the aggregator, (2.12). Here, we assume that labor arrives in the labor market in homogeneous form, sent by competitive employment agencies (see Figure 1). Key labor market activities - vacancy postings, layoffs, labor bargaining, setting the intensity of labor effort - are all carried out inside the employment agencies. An alternative, perhaps more realistic, specification would assume these labor market functions are performed inside the intermediate input firms. We adopt our specification because, like in the standard problem, the intermediate good firm has no state variable. We explained in section 4, how adding a state variable to the intermediate good firm substantially complicates the equilibrium conditions associated with price optimization.[20]

Each household is composed of many workers, each of which is in the labor force. A worker begins the period either unemployed or employed with a particular employment agency. Unemployed workers do undirected search. They find a job with a particular agency with a probability that is proportional to the efforts made by the agency to attract workers. Workers are separated from employment agencies either exogenously, or because they are actively cut. Workers pass back and forth between unemployment and employment with an agency. There are no agency to agency transitions.

The events during the period in an employment agency are displayed in Figure 2. Each employment agency begins a period with a stock of workers. That stock is immediately reduced by exogenous separations and it is increased by new arrivals that reflect the agency's recruiting efforts in the previous period. Then, the economy's technology shocks are realized.

At this point, each agency's wage is set. The agencies are allocated permanently into N equal-sized cohorts and each period $1/N$ agencies establish a new wage by Nash bargaining. When a new wage is set, it evolves over the subsequent $N-1$ periods according to (2.10) and

[20] For work that confronts these complications head on, see Thomas (2009).

(2.11). The wage negotiated in a given period covers all workers employed at an agency for each of the subsequent $N-1$ periods, even those that will not arrive until later. There are two types of bargaining arrangements. In the monopoly union arrangement, all workers are represented by a single negotiator. The union bargains on behalf of the current membership. As a result, the union discounts the future heavily, to take into account that some of the current membership will leave during the period of the contract. In the atomistic bargaining arrangement, each worker bargains separately with a representative of the employment agency. Under this arrangement workers who arrive in later periods receive the agency-wide wage if they arrive in a period when no bargaining occurs.

Next, each worker draws an idiosyncratic productivity shock. A cutoff level of productivity is determined, and workers with lower productivity are layed off. We consider two mechanism by which the cutoff is determined. One is based on the total surplus of a given worker and the other is based purely on the employment agency's interest. Finally, the intensity of each worker's labor effort is determined by an efficiency criterion.

After the endogenous layoff decision, the employment agency posts vacancies and the intensity of work effort is chosen. At this point, the current period monetary policy shock is realized and the employment agency supplies labor to the labor market.

We now describe these various labor market activities in greater detail. We begin with the decisions at the end of the period and work backwards to the bargaining problem. This is a convenient way to develop the model because the bargaining problem internalizes everything that comes after. The actual equilibrium conditions are displayed in the appendix.

5.2. Labor Hours

Labor intensity is chosen to equate the value of labor services to the employment agency with the cost of providing it by the household. To explain the latter, we display the utility function of the household, which is a modified version of (2.9):

$$E_t \sum_{l=0}^{\infty} \beta^{l-t} \{ \log(C_{t+l} - bC_{t+l-1}) - \zeta_{t+l} A_L \sum_{i=0}^{N-1} \frac{(\varsigma_{i,t+l})^{1+\sigma_L}}{1+\sigma_L} \left[1 - \mathcal{F}\left(\bar{a}^i_{t+l}\right) \right] l^i_{t+l} \}. \tag{5.1}$$

Here, $i \in \{0, ..., N-1\}$ indexes the cohort to which the employment agency belongs. The index, $i = 0$ corresponds to the cohort whose employment agency renegotiates the wage in the current period, $i = 1$ corresponds to the cohort that renegotiated in the previous period, and so on. The object, l^i_t denotes the number of workers in cohort i, after exogenous separations and new arrivals from unemployment have occured. Let a^i_t denote the idiosyncratic productivity shock drawn by a worker in cohort i. Then, $\bar{a}^i_t$, denotes the endogenously-determined cutoff such that all workers with $a^i_t < \bar{a}^i_t$ are layed off from the firm. Also, let

$$\mathcal{F}\left(\bar{a}^i_t\right) = P\left[a^i_t < \bar{a}^i_t\right]$$

denote the cumulative distribution function of the idiosyncratic productivity shock. (In practice, we assume $\mathcal{F}$ is lognormal with $Ea = 1$ and standard deviation of $log\,(a)$ equal to σ_a.) Then,

$$\left[1 - \mathcal{F}\left(\bar{a}_t^i\right)\right] l_t^i \tag{5.2}$$

denotes the number of workers with an employment agency in the i^{th} cohort who survive the endogenous layoffs.

Let $\varsigma_{i,t}$ denote the number of hours supplied by a worker in the i^{th} cohort. The absence of the index, a, on $\varsigma_{i,t}$ reflects our assumption that each worker who survives endogenous layoffs in cohort i works the same number of hours, regardless of the realization of their idiosyncratic level of productivity. The disutility experienced by a worker that works $\varsigma_{i,t}$ hours is:

$$\zeta_t A_L \frac{\left(\varsigma_{i,t}\right)^{1+\sigma_L}}{1+\sigma_L}.$$

The utility function in (5.1) sums the disutility experienced by the workers in each cohort.

Although the individual worker's labor market experience - whether employed or unemployed - is determined by idiosyncratic shocks, each household has sufficiently many workers that the total fraction of workers employed,

$$L_t = \sum_{i=0}^{N-1} \left[1 - \mathcal{F}\left(\bar{a}_t^i\right)\right] l_t^i,$$

as well as the fractions allocated among the different cohorts, $\left[1 - \mathcal{F}\left(\bar{a}_t^i\right)\right] l_t^i$, $i = 0, ..., N-1$, are the same for each household. We suppose that all the household's workers are supplied inelastically to the labor market (i.e., labor force participation is constant).

The household's currency receipts arising from the labor market are:

$$\left(1 - L_t\right) P_t b^u z_t^+ + \sum_{i=0}^{N-1} W_t^i \left[1 - \mathcal{F}\left(\bar{a}_t^i\right)\right] l_t^i \varsigma_{i,t} \tag{5.3}$$

where W_t^i is the nominal wage rate earned by workers in cohort $i = 0, ..., N-1$. The presence of the term involving b^u indicates the assumption that unemployed workers, $1 - L_t$, receive a payment of $b^u z_t^+$ final consumption goods. These unemployment benefits are financed by lump sum taxes.

Let W_t denote the price received by employment agencies for supplying one unit of labor service. It represents the marginal gain to the employment agency that occurs when an individual worker increases time spent working by one unit. Because the employment agency is competitive in the supply of labor services, it takes W_t as given. We treat W_t as an unobserved variable in the data. In practice, it is the shadow value of an extra worker supplied by the human resources department to a firm.

Following GST, we assume that labor hours are chosen to equate the worker's marginal cost of working with the agency's marginal benefit:

$$W_t \frac{\mathcal{E}_t^i}{1-\mathcal{F}_t^i} = \zeta_t A_L \varsigma_{i,t}^{\sigma_L} \frac{1}{\upsilon_t} \tag{5.4}$$

for $i = 0, ..., N-1$. Here,

$$\mathcal{E}_t^i \equiv \mathcal{E}\left(\bar{a}_t^j\right) \equiv \int_{\bar{a}_t^j}^{\infty} a d\mathcal{F}(a) \tag{5.5}$$

$$\mathcal{F}_t^i = \mathcal{F}\left(\bar{a}_t^j\right) = \int_0^{\bar{a}_t^j} d\mathcal{F}(a). \tag{5.6}$$

To understand the expression on the right of (5.4), note that the marginal cost, in utility terms, to an individual worker who increases labor intensity by one unit is $\zeta_t^h A_L \varsigma_{i,t}^{\sigma_L}$. This is converted to currency units by dividing by the multiplier, υ_t, on the household's nominal budget constraint. The left side of (5.4) represents the increase in revenues to the employment agency from increasing hours worked by one unit (recall, all workers who survive endogenous layoffs work the same number of hours.) Division by $1-\mathcal{F}_t^i$ is required in (5.4) so that the expectation is relative to the distribution of a conditional on $a \geq \bar{a}_t^j$.

Labor intensity is potentially different across cohorts because $\mathcal{E}_t^i / \left(1-\mathcal{F}_t^i\right)$ in (5.4) is indexed by cohort. When the wage rate is determined by Nash bargaining, it is taken into account that labor intensity is determined according to (5.4) and that some workers will endogenously separate.

5.3. Vacancies and the Employment Agency Problem

The employment agency in the i^{th} cohort determines how many employees it will have in period $t+1$ by choosing vacancies, v_t^i. The vacancy posting costs associated with v_t^i are:

$$\frac{\kappa z_t^+}{\varphi} \left(\frac{Q_t^\iota v_t^i}{\left[1-\mathcal{F}\left(\bar{a}_t^i\right)\right] l_t^i} \right)^\varphi \left[1-\mathcal{F}\left(\bar{a}_t^i\right)\right] l_t^i,$$

units of the domestic homogeneous good. The parameter φ determines the curvature of the cost function and in practice we set $\varphi = 2$. Also, $\kappa z_t^+/\varphi$ is a cost parameter which is assumed to grow at the same rate as the overall economic growth rate and, as noted above, $\left[1-\mathcal{F}\left(\bar{a}_t^i\right)\right] l_t^i$ denotes the number of employees in the i^{th} cohort after endogenous separations have occurred. Also, Q_t is the probability that a posted vacancy is filled, a quantity that is exogenous to an individual employment agency. The functional form of our cost function reduces to the function used in GT and GST when $\iota = 1$. With this parameterization, costs are a function of the number of people hired, not the number of vacancies per se. We interpret this as reflecting that the GT and GST specifications emphasize internal costs (such as training and other) of adjusting the work force, and not search costs. In models used in the search literature (see, e.g., Shimer (2005a)), vacancy posting costs are independent of Q_t, i.e., they set $\iota = 0$. To understand the implications for our type of empirical analysis, consider a shock that triggers an economic expansion and also produces a fall in the probability of filling a vacancy, Q_t. We expect the expansion to be smaller in a version of the model that emphasizes search costs (i.e., $\iota = 0$) than in a version that emphasizes internal costs (i.e., $\iota = 1$).

To further describe the vacancy decisions of the employment agencies, we require their objective function. We begin by considering $F\left(l_t^0, \omega_t\right)$, the value function of the representative employment agency in the cohort, $i = 0$, that negotiates its wage in the current period. The arguments of F are the agency's workforce after beginning-of-period exogenous separations and new arrivals, l_t^0, and an arbitrary value for the nominal wage rate, ω_t. That is, we consider the value of the firm's problem after the wage rate has been set.

We suppose that the firm chooses a particular monotone transform of vacancy postings, which we denote by $\tilde{v}_t^i$:

$$\tilde{v}_t^i \equiv \frac{Q_t^{\iota} v_t^i}{\left(1 - \mathcal{F}_t^j\right) l_t^i},$$

where $1 - \mathcal{F}_t^j$ denotes the fraction of the beginning-of-period t workforce in cohort j which survives endogenous separations. The agency's hiring rate, χ_t^i, is related to $\tilde{v}_t^i$ by:

$$\chi_t^i = Q_t^{1-\iota} \tilde{v}_t^i. \tag{5.7}$$

To construct $F\left(l_t^0, \omega_t\right)$, we must derive the law of motion of the firm's work force, during the period of the wage contract. If l_t^i is the period t work force just after exogenous separations and new arrivals, then (5.2) is the size of the workforce after endogenous separations. The time $t+1$ workforce of the representative agency in the i^{th} cohort at time t is denoted l_{t+1}^{i+1}. That workforce reflects the endogenous separations in period t as well as the exogenous separations and new arrivals at the start of period $t+1$. Let ρ denote the probability that an individual worker attached to an employment agency at the start of a period survives the exogenous separation. Then, given the hiring rate, χ_t^i, we have

$$l_{t+1}^{j+1} = \left(\chi_t^j + \rho\right)\left(1 - \mathcal{F}_t^j\right) l_t^j, \tag{5.8}$$

for $j = 0, 1, ..., N-1$, with the understanding here and throughout that $j = N$ is to be interpreted as $j = 0$. Expression (5.8) is deterministic, reflecting the assumption that the representative employment agency in cohort j employs a large number of workers.

The value function of the firm is:

$$\begin{aligned} F\left(l_t^0,\omega_t\right) &= \sum_{j=0}^{N-1}\beta^j E_t\frac{\upsilon_{t+j}}{\upsilon_t}\max_{\left(\tilde{v}_{t+j}^j,\bar{a}_{t+j}^j\right)}[\int_{\bar{a}_{t+j}^j}^{\infty}\left(W_{t+j}a-\Gamma_{t,j}\omega_t\right)\varsigma_{j,t+j}d\mathcal{F}\left(a\right) \\ &\quad -P_{t+j}\frac{\kappa z_{t+j}^{+}}{\varphi}\left(\tilde{v}_t^j\right)^{\varphi}\left(1-\mathcal{F}_{t+j}^j\right)]l_{t+j}^j \\ &\quad +\beta^N E_t\frac{\upsilon_{t+N}}{\upsilon_t}F\left(l_{t+N}^0,\tilde{W}_{t+N}\right), \end{aligned} \tag{5.9}$$

where l_t^j evolves according to (5.8), $\varsigma_{j,t}$ satisfies (5.4) and

$$\Gamma_{t,j}=\begin{cases}\tilde{\pi}_{w,t+j}\cdots\tilde{\pi}_{w,t+1}, & j>0\\ 1 & j=0\end{cases}. \tag{5.10}$$

Here, $\tilde{\pi}_{w,t}$ is defined in (2.11). The term, $\Gamma_{t,j}\omega_t$, represents the wage rate in period $t+j$, given the wage rate was ω_t at time t and there have been no wage negotiations in periods $t+1$, $t+2$, up to and including period $t+j$. In (5.9), $\tilde{W}_{t+N}$ denotes the Nash bargaining wage that is negotiated in period $t+N$, which is when the next round of bargaining occurs. At time t, the agency takes the state $t+N-$contingent function, $\tilde{W}_{t+N}$, as given. The vacancy decision of employment agencies solve the maximization poblem in (5.9).

It is easily verified using (5.9) that $F\left(l_t^0,\omega_t\right)$ is linear in l_t^0:

$$F\left(l_t^0,\omega_t\right)=J\left(\omega_t\right)l_t^0, \tag{5.11}$$

where $J\left(\omega_t\right)$ is not a function of l_t^0. The function, $J\left(\omega_t\right)$, is the surplus that a firm bargaining in the current period enjoys from a match with an individual worker, when the current wage is ω_t. Although later in the period workers become heterogeneous when they draw an idiosyncratic shock to productivity, the fact that that draw is iid over time means that workers are all identical when (5.11) is evaluated.

5.4. Worker Value Functions

In order to discuss the endogenous separation decisions, as well as the bargaining problem, we must have the value functions of the individual worker. For the bargaining problem, we require the worker's value function before he knows what his idiosyncratic productivity draw is. For the endogenous separation problem, we need to know the worker's value function after he knows he has survived the endogenous separation. For both the bargaining and separation problem, we need to know the value of unemployment to the worker.

Let V_t^i denote the period t value of being a worker in an agency in cohort i, after that worker has survived that period's endogenous separation:

$$\begin{aligned} V_t^i &= \Gamma_{t-i,i}\tilde{W}_{t-i}\varsigma_{i,t}-\zeta_t A_L\frac{\varsigma_{i,t}^{1+\sigma_L}}{\left(1+\sigma_L\right)\upsilon_t} \\ &\quad +\beta E_t\frac{\upsilon_{t+1}}{\upsilon_t}\left[\rho\left(1-\mathcal{F}_{t+1}^{i+1}\right)V_{t+1}^{i+1}+\left(1-\rho+\rho\mathcal{F}_{t+1}^{i+1}\right)U_{t+1}\right], \end{aligned} \tag{5.12}$$

for $i=0,1,...,N-1$. In (5.12), $\tilde{W}_{t-i}$ denotes the wage negotiated i periods in the past, and $\Gamma_{t-i,i}\tilde{W}_{t-i}$ represents the wage received in period t by workers in cohort i. The two terms after the equality in (5.12) represent a worker's period t flow utility, converted into units of currency.[21] The terms in square brackets in (5.12) correspond to utility in the two possible period $t+1$ states of the world. With probability $\rho\left(1-\mathcal{F}_{t+1}^{i+1}\right)$ the worker survives the exogenous and endogenous separations in period $t+1$, in which case its value function in $t+1$ is V_{t+1}^{i+1}. With the complementary probability, $1-\rho+\rho\mathcal{F}_{t+1}^{i+1}$, the worker separates into unemployment in period $t+1$, and enjoys utility, U_{t+1}.

The (currency) value of being unemployed in period t is:

$$U_t = P_t z_t^+ b^u + \beta E_t \frac{\upsilon_{t+1}}{\upsilon_t}[f_t V_{t+1}^x + (1-f_t)\, U_{t+1}]. \tag{5.13}$$

Here, f_t is the probability that an unemployed worker matches with an employment agency at the start of period $t+1$. Also, V_{t+1}^x is the period $t+1$ value function of a worker who knows that he has matched with an employment agency at the start of $t+1$, but does not know which one. In particular,

$$V_{t+1}^x = \sum_{i=0}^{N-1} \frac{\chi_t^i\,(1-\mathcal{F}_t^i)\, l_t^i}{m_t} \tilde{V}_{t+1}^{i+1}. \tag{5.14}$$

Here, total new matches at the start of period $t+1$, m_t, is given by:

$$m_t = \sum_{j=0}^{N-1} \chi_t^j \left(1-\mathcal{F}_t^j\right) l_t^j. \tag{5.15}$$

In (5.14),

$$\frac{\chi_t^i\,(1-\mathcal{F}_t^i)\, l_t^i}{m_t}$$

is the probability of finding a job in $t+1$ in an agency belonging to cohort i in period t. Note that this is a proper probability distribution because it is positive for each i and it sums to unity by (5.15).

In (5.14), $\tilde{V}_{t+1}^{i+1}$ is the analog of V_{t+1}^{i+1}, except that the former is defined before the worker knows if he survives the endogenous productivity cut, while the latter is defined after survival. The superscript $i+1$ appears on $\tilde{V}_{t+1}^{i+1}$ because the probabilities in (5.14) refer to activities in a particular agency cohort in period t, while in period $t+1$ the index of that cohort is incremented by unity.

We complete the definition of U_t in (5.13) by giving the formal definition of $\tilde{V}_t^j$:

$$\tilde{V}_t^j = \mathcal{F}_t^j U_t + \left(1-\mathcal{F}_t^j\right) V_t^j. \tag{5.16}$$

[21] Note the division of the disutility of work in (5.12) by υ_t, the multiplier on the budget constraint of the household optimization problem.

That is, at the start of the period, the worker has probability $\mathcal{F}_t^j$ of returning to unemployment, and the complementary probability of surviving in the firm to work and receive a wage in period t.

5.5. Separation Decision

This section describes how the separation decision is made. We consider the problem of the $j = 0$ cohort of agencies which renegotiate the wage in the current period. The other cohorts are symmetric. Just prior to the realization of idiosyncratic worker uncertainty, the number of workers attached to the firm is l_t^0. Each of these workers draws from the cumulative distribution function, $\mathcal{F}$, and those workers who draw $a \leq \bar{a}_t^0$ will be separated from the firm. This section discusses the determination of $\bar{a}_t^0$. We select $\bar{a}_t^0$ by optimizing alternative measures of the surplus associated with the l_t^0 workers.

The sum of the worker surplus across the l_t^0 workers is

$$\left(V_t^0 - U_t\right)\left(1 - \mathcal{F}_t^0\right) l_t^0.$$

Each worker in the fraction, $\left(1 - \mathcal{F}_t^0\right)$, of workers with $a \geq \bar{a}_t^0$ experiences the same surplus, $V_t^0 - U_t$. This is the worker's surplus because the worker's outside option is U_t. The fraction, $\mathcal{F}_t^0$, of workers enjoys zero surplus. Note that $\mathcal{F}_t^0$ is a direct function of $\bar{a}_t^0$, while $\bar{a}_t^0$ affects V_t^0 indirectly via its impact on $\varsigma_{0,t}$ (see (5.4) and (5.12)). The derivative of worker surplus with respect to $\bar{a}_t^0$ is straightforward to evaluate.

The surplus enjoyed by the employment agency before idiosyncratic worker uncertainty is realized and when the workforce is l_t^0, is given by (5.9). According to (5.11) firm surplus per worker in l_t^0 is given by $J(\omega_t)$ which is defined under the assumption that $\bar{a}_t^0$ is set to its optimized value. To express $J(\omega_t)$ as a function of an arbitrary value of $\bar{a}_t^0$, express $J(\omega_t)$ as follows:

$$J(\omega_t) = \max_{\bar{a}_t^0} \tilde{J}\left(\omega_t; \bar{a}_t^0\right)\left(1 - \mathcal{F}_t^0\right),$$

where

$$\tilde{J}\left(\omega_t; \bar{a}_t^0\right) = \max_{\tilde{v}_t^0} \left\{ \left(W_t \frac{\mathcal{E}_t^0}{1 - \mathcal{F}_t^0} - \omega_t \right) \varsigma_{0,t} - P_t z_t^+ \frac{\kappa}{\varphi}\left(\tilde{v}_t^0\right)^{\varphi} + \beta \frac{\upsilon_{t+1}}{\upsilon_t}\left(\tilde{v}_t^0 Q_t^{1-\iota} + \rho\right) J_{t+1}^1(\omega_t) \right\}.$$

Here J_{t+1}^1, agency profits as a fraction of its workforce at the start of $t+1$, is not a function of $\bar{a}_t^0$.[22] Thus, firm surplus associated with a workforce, l_t^0, as a function of an arbitrary value of $\bar{a}_t^0$ is:

$$\tilde{J}\left(\omega_t;\bar{a}_t^0\right)\left(1-\mathcal{F}_t^0\right)l_t^0.$$

Consider the following measure of surplus, which integrates both the firm and the worker:

$$\left[s_w\left(V_t^0-U_t\right)+s_e\tilde{J}\left(\omega_t;\bar{a}_t^0\right)\right]\left(1-\mathcal{F}_t^0\right)l_t^0. \tag{5.17}$$

The parameters $s_w,\ s_e\in\{0;1\}$ allow for a variety of interesting surplus measures. If $s_w=0$ and $s_e=1$ we have employer surplus. If $s_w=1$ and $s_e=1$ we have total surplus. We suppose that $\bar{a}_t^0$ is selected to maximize (5.17). The employer surplus model is the one based on $s_w=0,\ s_e=1$ and the total surplus model is the one based on $s_w=s_e=1$.

5.6. Bargaining Problem

We suppose that bargaining occurs among a continuum of worker-agency representative pairs. Each bargaining session takes the outcomes of all other bargaining sessions as given. Because each bargaining session is atomistic, each session ignores its impact on the wage earned by workers arriving in the future during the contract. We assume that those future workers are simply paid the average of the outcome of all bargaining sessions. Since each bargaining problem is identical, the wage that solves each problem is the same and so the average wage coincides with the wage that solves the individual bargaining problem. Because each bargaining session is atomistic, it also ignores the impact of the wage bargain on decisions like vacancies and separations, taken by the firm.

The Nash bargaining problem that determines the wage rate is a combination of the worker surplus and firm surplus

$$\max_{\omega_t}\left(\tilde{V}_t^0-U_t\right)^{\eta}J\left(\omega_t\right)^{(1-\eta)}.$$

Here, the firm surplus, $J\left(\omega_t\right)$, reflects that the outside option of the firm in the bargaining problem is zero. The above problem has an interesting structure. Note first (ignoring the

[22] Here,

$$J_{t+1}^1\left(\omega_t\right)=\max_{\left\{\bar{a}_{t+j}^j,\tilde{v}_{t+j}^j\right\}_{j=1}^{N-1}}\left\{\left[\left(W_{t+1}\frac{\mathcal{E}_{t+1}^1}{1-\mathcal{F}_{t+1}^1}-\Gamma_{t,1}\omega_t\right)\varsigma_{1,t+1}-P_{t+1}z_{t+1}^+\frac{\kappa}{\varphi}\left(\tilde{v}_{t+1}^1\right)^{\varphi}\right]\left(1-\mathcal{F}_{t+1}^1\right)\right.$$
$$+\beta\frac{\upsilon_{t+2}}{\upsilon_{t+1}}\left[\left(W_{t+2}\mathcal{E}_{t+2}^2-\Gamma_{t,2}\omega_t\left(1-\mathcal{F}_{t+2}^2\right)\right)\varsigma_{2,t+2}-P_{t+2}z_{t+2}^+\frac{\kappa}{\varphi}\left(\tilde{v}_{t+2}^2\right)^{\varphi}\left(1-\mathcal{F}_{t+2}^2\right)\right]\times$$
$$\left(\tilde{v}_{t+1}^1Q_{t+1}^{1-\iota}+\rho\right)\left(1-\mathcal{F}_{t+1}^1\right)$$
$$+...+$$
$$\left.+\beta^{N-1}\frac{\upsilon_{t+N}}{\upsilon_{t+1}}J\left(\tilde{W}_{t+N}\right)\left(\tilde{v}_{t+1}^1Q_t^{1-\iota}+\rho\right)\cdots\left(\tilde{v}_{t+N-1}^{N-1}Q_{t+N-1}^{1-\iota}+\rho\right)\left(1-\mathcal{F}_{t+N-1}^{N-1}\right)\cdots\left(1-\mathcal{F}_{t+1}^1\right)\right\}.$$

impact of ω_t on the vacancy decision):

$$\begin{aligned} J_{w,t} &= -\left(1-\mathcal{F}_t^0\right)\varsigma_{0,t} \\ &+\beta\frac{\upsilon_{t+1}}{\upsilon_t}\left[-\Gamma_{t,1}\varsigma_{1,t+1}\rho\left(1-\mathcal{F}_{t+1}^1\right)\left(1-\mathcal{F}_t^0\right)\right] \\ &+\beta^2\frac{\upsilon_{t+2}}{\upsilon_t}\left[-\Gamma_{t,2}\varsigma_{2,t+2}\right]\rho^2\left(1-\mathcal{F}_{t+2}^2\right)\left(1-\mathcal{F}_{t+1}^1\right)\left(1-\mathcal{F}_t^0\right) \\ &+\ldots+ \\ &+\beta^{N-1}\frac{\upsilon_{t+N-1}}{\upsilon_t}\left[-\Gamma_{t,N-1}\varsigma_{N-1,t+N-1}\right]\rho^{N-1}\left(1-\mathcal{F}_{t+N-1}^{N-1}\right)\cdots\left(1-\mathcal{F}_t^0\right), \end{aligned}$$

where $J_{w,t}$ denotes the derivative of the surplus with respect to the wage rate. Note that a rise in the wage reduces J_t only in future states of the world in which the worker survives both exogenous and endogenous separation. It is easy to verify that $J_{w,t} = -\tilde{V}_{w,t}$. That is, a contemplated increase in the wage simply reallocates resources between the firm and the worker.[23] If we had formulated the bargaining problem as one between a single union representing the workers and a single firm representative, then $J_{w,t}$ and $\tilde{V}_{w,t}$ would have been different. This is because the worker and the firm discount the future differently in the union case. The union, which is only concerned about the welfare of the current membership, discounts the future relatively heavily like in the atomistic case because some of the membership will separate. The firm, on the other hand, understands that the wage negotiated now will also be paid to newly arriving future workers. In the atomistic case, the firm representative correctly sees no such connection. So the firm, in the union bargaining case, discounts the future less heavily. Under these circumstances, if there were no restrictions on the intertemporal pattern of wage payments in the union model, our Nash bargaining criterion would imply that it is desirable to front load wage payments to the present. This would in effect shift surplus away from future workers and towards the workers involved in the current bargaining. In the atomistic bargaining case that we consider both the worker and the firm discount the future in the same way and these incentives to intertemporally reallocate wages are not present.

While the union bargaining case might be more in line with European labor markets, the atomistic individual bargaining case appears to be more reasonable for the USA. Since in this paper we will estimate the models using US data we proceed with the individual bargaining case from now on. In addition, by decoupling the impact of the wage bargain on the firm separation decision, the atomistic bargaining approach is relatively simple to implement.[24]

[23] This result is not robust to the treatment of taxes, which we ignore here.

[24] In principle, the cutoff, $\bar{a}_t^0$, also affects the vacancy posting decision in the union bargaining problem. However, that effect can be ignored by the envelope condition, because the (optimized) vacancy decision only appears in the firm surplus and not the worker surplus. In effect, the vacancy decision optimizes the same criterion that is optimized in the Nash bargaining problem. In the union problem, there is no envelope condition that allows us to ignore the impact of the wage on the endogenous separation decision. This is because the criterion optimized by the optimization decision is different from the Nash bargaining problem.

Until now we have implicitly assumed that the negotiated wage paid by an employment agency which has renegotiated most recently i periods in the past is always inside the bargaining set, $[\underline{w}_t^i, \bar{w}_t^i]$, $i = 0, 1, ..., N-1$. In other words, the wage paid is not lower than the workers reservation wage and not higher than the wage an employment agency is willing to pay. In appendix A.12 we describe how we check that the paid wage is within the bargaining set.

5.7. Resource Constraint in the Search-Based Unemployment Model

We assume that the posting of vacancies uses the homogeneous domestic good. We leave the production technology equation, (A.27), unchanged, and we alter the resource constraint:

$$y_t = g_t + c_t + \left(i_t + a(u_t) \frac{\bar{k}_t}{\mu_{\psi,t}\mu_{z^+,t}} \right) + \frac{\kappa}{\varphi} \sum_{j=0}^{N-1} \left(\tilde{v}_t^j\right)^{\varphi} \left[1 - \mathcal{F}_t^j\right] l_t^j. \quad (5.18)$$

Total job matches must also satisfy the following matching function:

$$m_t = \sigma_m (1 - L_t)^{\sigma} v_t^{1-\sigma}, \quad (5.19)$$

where

$$L_t = \sum_{j=0}^{N-1} \left(1 - \mathcal{F}_t^j\right) l_t^j. \quad (5.20)$$

and σ_m is the productivity of the matching technology.

In our environment, there is a distinction between effective hours and measured hours. Effective hours is the hours of each person, adjusted by their productivity, a. The average productivity of a worker in working in cohort j (i.e., who has survived the endogenous productivity cut) is $\mathcal{E}_t^j / \left(1 - \mathcal{F}_t^j\right)$. The number of workers who survive the productivity cut in cohort j is $\left(1 - \mathcal{F}_t^j\right) l_t^j$, so that our measure of total effective hours is:

$$H_t = \sum_{j=0}^{N-1} \varsigma_{j,t} \mathcal{E}_t^j l_t^j, \quad (5.21)$$

$$\mathcal{E}\left(\bar{a}_t^j; \sigma_{a,t}\right) = \int_{\bar{a}_t^j}^{\infty} a d\mathcal{F}(a; \sigma_{a,t}) = 1 - prob\left[v < \frac{\log\left(\bar{a}_t^j\right) + \frac{1}{2}\sigma_{a,t}^2}{\sigma_{a,t}} - \sigma_{a,t}\right], \quad (5.22)$$

where $prob$ refers to the standard normal distribution. We also need:

$$\mathcal{F}\left(\bar{a}^j; \sigma_a\right) = \int_0^{\bar{a}^j} d\mathcal{F}(a; \sigma_a) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\frac{\log\left(\bar{a}^j\right) + \frac{1}{2}\sigma_a^2}{\sigma}} \exp^{\frac{-v^2}{2}} dv \quad (5.23)$$
$$= prob\left[v < \frac{\log\left(\bar{a}^j\right) + \frac{1}{2}\sigma_a^2}{\sigma_a}\right].$$

Total measured hours is:

$$H_t^{meas} = \sum_{j=0}^{N-1} \varsigma_{j,t} \left(1 - \mathcal{F}_t^j\right) l_t^j.$$

The job finding rate is:

$$f_t = \frac{m_t}{1 - L_t}. \tag{5.24}$$

The probability of filling a vacancy is:

$$Q_t = \frac{m_t}{v_t}. \tag{5.25}$$

Total vacancies v_t are related to vacancies posted by the individual cohorts as follows:

$$v_t = \frac{1}{Q_t^{\iota}} \sum_{j=0}^{N-1} \tilde{v}_t^j \left(1 - \mathcal{F}_t^j\right) l_t^j.$$

Note however, that this equation does not add a constraint to the model equilibrium. In fact, it can be derived from the equilibrium equations (5.25), (5.15) and (5.7). This completes the discussion of the model with a search-based approach to the labor market.

6. Estimation Strategy

Our estimation strategy is a Bayesian version of the two-step impulse response matching approach applied by Rotemberg and Woodford (1997) and CEE. We begin with a discussion of the two steps. After that, we discuss the computation of a particular weighting matrix used in the analysis. Finally, we describe the use of the Laplace approximation to compute posterior distributions and an overall measure of model fit.

6.1. VAR Step

We estimate the dynamic responses of a set of aggregate variables to three shocks, using standard vector autoregression methods. The three shocks are the monetary policy shock, the innovation to the permanent technology shock, z_t, and the innovation to the investment-specific technology shock, Ψ_t. The contemporaneous and 20 lagged responses to each of $N = 11$ macroeconomic variables to the three shocks are stacked in the vector, $\hat{\psi}$. The Y_t

vector of variables in the VAR is:

$$\underbrace{Y_t}_{12\times 1} = \begin{pmatrix} \Delta\ln(\text{relative price of investment}_t) \\ \Delta\ln(GDP_t/\text{Hours}_t) \\ \Delta\ln(GDP\text{ deflator}_t) \\ \text{unemployment rate}_t \\ \text{capacity utilization}_t \\ \ln(\text{Hours}_t) \\ \ln(GDP_t/\text{Hours}_t) - \ln(W_t/P_t) \\ \ln(C_t/GDP_t) \\ \ln(I_t/GDP_t) \\ \text{vacancies}_t \\ \log(\text{Hours}_t/\text{Labor force}_t) \\ \text{Federal Funds Rate}_t \end{pmatrix} \tag{6.1}$$

An extensive general discussion of identification in VAR's appears in Christiano, Eichenbaum and Evans ( ). The specific technical details of how we simultaneously identify the responses to all three structural shocks in our model appear in ACEL. Briefly, the data in our analysis are quarterly and seasonally adjusted and cover the period 1951Q1 to 2008Q4. In the estimation, we lose the first observation to first differencing and the next four observations to the four lags included in the VAR. Thus, the estimation period of the analysis is 1952Q2 to 2008Q4. Our identification assumptions are as follows. The only variable that the monetary policy shock affects contemporaneously is the Federal Funds Rate. We make two assumptions to identify the dynamic response to the technology shocks: (i) the only shocks that affect labor productivity in the long run are the two technology shocks and (ii) the only shock that affects the price of investment relative to consumption is the innovation to the investment specific shock. All these identification assumptions are satisfied in our model. Details of our strategy for computing impulse response functions imposing the shock identification are reported in ACEL.[25]

Our data set extends over a long range, while we estimate a single set of impulse response functions and model parameters. In effect, we suppose that there has been no parameter break over this long period. This is an issue that has been debated. For example, it has been argued that the parameters of the monetary policy rule have not been constant over this period. We do not review this debate here. Implicitly, our analysis sides with the conclusions of those that argue that the evidence of parameter breaks is not strong. For example, Sims

[25] The identification assumption for the monetary policy shock by itself imposes no restriction on the VAR parameters. Similarly, Fisher ( ) showed that the identification assumptions for the technology shocks when applied without simultaneously applying the monetary shock identification, also imposes no restriction on the VAR parameters. However, ACEL showed that when all the identification assumptions are imposed at the same time, then there are restrictions on the VAR parameters. We found that the test of the overidentifying restrictions on the VAR rejects the null hypothesis that the restrictions are satisfied at 1 percent critical level. We decided to continue the analysis anyway.

and Zha ( ) argue that the evidence is consistent with the idea that monetay policy rule parameters have been unchanged over the sample. Christiano, Eichenbaum and Evans ( ) argue that the evidence is consistent with the proposition that the dynamic effects of a monetary policy shock have not changed during this sample. Standard lag-length selection criteria led us to work with a VAR with 2 lags.

The number of elements in $\hat{\psi}$ corresponds to the number of impulses estimated. Since we consider $n = 20$ lags in the impulses, there are in principle 3 (i.e., the number of shocks) times 11 (number of variables) times 20 (number of lags) = 660 elements in $\hat{\psi}$. However, we do not include in $\hat{\psi}$ the 10 contemporaneous responses to the monetary policy shock that are required to be zero by our monetary policy identifying assumption. Taking the latter into account, the vector $\hat{\psi}$ has 650 elements.

According to standard classical asymptotic sampling theory, when the number of observations, T, is large, we have

$$\sqrt{T}\left(\hat{\psi} - \psi\left(\theta_0\right)\right) \overset{a}{\sim} N\left(0, W\left(\theta_0, \zeta_0\right)\right).$$

We find it convenient to express this result in the following form:

$$\hat{\psi} \overset{a}{\sim} N\left(\psi\left(\theta_0\right), V\left(\theta_0, \zeta_0, T\right)\right), \tag{6.2}$$

where

$$V\left(\theta_0, \zeta_0, T\right) \equiv \frac{W\left(\theta_0, \zeta_0\right)}{T}.$$

Note that V is a function of the parameters that appear in our analysis, and also of the parameters, ζ, that we do not consider.

6.2. Impulse Response Matching Step

In the second step of our analysis, we treat $\hat{\psi}$ as 'data' and we choose a value of θ to make $\psi\left(\theta\right)$ as close as possible, in a specific metric, to $\hat{\psi}$. We give our strategy an approximate Bayesian interpretation or, in the case that we do not use priors, an approximate maximum likelihood interpretation.[26] This interpretation uses (6.2) to define an approximate likelihood of the data, $\hat{\psi}$, as a function of θ :

$$f\left(\hat{\psi}|\theta, V\left(\theta_0, \zeta_0, T\right)\right) = \left(\frac{1}{2\pi}\right)^{\frac{N}{2}} \left|V\left(\theta_0, \zeta_0, T\right)\right|^{-\frac{1}{2}} \exp\left[-\frac{1}{2}\left(\hat{\psi} - \psi\left(\theta\right)\right)' V\left(\theta_0, \zeta_0, T\right)^{-1}\left(\hat{\psi} - \psi\left(\theta\right)\right)\right].$$

As we explain below, we treat the true value of $V\left(\theta_0, \zeta_0, T\right)$ as a known object. Under these circumstances, the value of θ that maximizes the above function represents an approximate

[26] Our approach follows in the spirit of Chernozhukov and Hong (2003). For a careful classical analysis of the impulse response methodology, see

maximum likelihood estimator of θ. It is approximate for two reasons: (i) the central limit theorem underlying (6.2) only holds exactly as $T \to \infty$ and (ii) the value of $V(\theta_0, \zeta_0, T)$ that we use is guaranteed to be correct only for $T \to \infty$. Interestingly, our approximation does not require (as in standard Bayesian analysis, which works with a Normal likelihood) that the data underlying the VAR, Y_t, be Normal. This is an advantage of the method, because the Normality assumption is not a good one for macroeconomic variables (see Christiano ( )).

Treating the function, f, as the likelihood of $\hat{\psi}$, it follows that the Bayesian posterior of θ conditional on $\hat{\psi}$ and $V(\theta_0, \zeta_0, T)$ is:

$$f\left(\theta|\hat{\psi}, V(\theta_0, \zeta_0, T)\right) = \frac{f\left(\hat{\psi}|\theta, V(\theta_0, \zeta_0, T)\right) p(\theta)}{f\left(\hat{\psi}|V(\theta_0, \zeta_0, T)\right)}, \tag{6.3}$$

where $p(\theta)$ denotes the priors on θ and $f\left(\hat{\psi}|V(\theta_0, \zeta_0, T)\right)$ denotes the marginal density of $\hat{\psi}$:

$$f\left(\hat{\psi}|V(\theta_0, \zeta_0, T)\right) = \int f\left(\hat{\psi}|\theta, V(\theta_0, \zeta_0, T)\right) p(\theta)\, d\theta.$$

As usual, the mode of the posterior distribution of θ can be computed by simply maximizing the value of the numerator in (6.3), since the denominator is not a function of θ. The marginal density of $\hat{\psi}$ is required when we want an overall measure of the fit of our model and when we want to report the shape of the posterior marginal distribution of individual elements in θ. To compute the marginal likelihood, we can use an MCMC algorithm or the Laplace Approximation. We briefly review the latter in the last subsection below.

6.3. Computation of $V(\theta_0, \zeta_0, T)$

A crucial ingredient in our empirical methodology is the matrix, $V(\theta_0, \zeta_0, T)$. The logic of our approach requires that we have an at least approximately consistent estimator of $V(\theta_0, \zeta_0, T)$. A variety of approaches are possible here. We use a bootstrap approach. Using our estimated VAR and its fitted disturbances, we generate a set of M bootstrap realizations for the impulse responses. We denote these by ψ_i, $i = 1, ..., M$, where ψ_i denotes the i^{th} realization of the 650×1 vector of impulse responses.[27] Consider

$$\bar{V} = \frac{1}{M} \sum_{i=1}^{M} \left(\psi_i - \bar{\psi}\right) \left(\psi_i - \bar{\psi}\right)',$$

[27] To compute a given bootstrap realization, ψ_i, we first simulate an artificial data set, $Y_1, ..., Y_T$. We do this by simulating the response of our estimated VAR to an iid sequence of 11×1 shock vectors that are drawn randomly with replacement from the set of fitted shocks. We then fit a 4-lag VAR to the artificial data set using the same procedure used on the actual data. The resulting estimated VAR is then used to compute the impulse responses, which we stack into the 650×1 vector, ψ_i.

where $\bar{\psi}$ is the mean of ψ_i, $i = 1, ..., M$ and M is large. The object, $\bar{V}$, is a 650 by 650 matrix, and we assume that the small sample (in the sense of T) properties of this way (or any other way) of estimating $V(\theta_0, \zeta_0, T)$ are poor. To improve small sample efficiency, we proceed in a way that is analogous to the strategy taken in the estimation of frequency-zero spectral densities (see Newey and West (1987)). In particular, rather than working with the raw variance-covariance matrix, $\bar{V}$, we instead work with $\hat{\bar{V}}$:

$$\hat{\bar{V}} = f\left(\bar{V}, T\right).$$

The transformation, f, has the property that it converges to the identity transform, as $T \to \infty$. In particular, $\hat{\bar{V}}$ damps some elements in $\bar{V}$, and the damping factor is removed as the sample grows large. The matrix, $\hat{\bar{V}}$, has on its diagonal, the diagonal elements of $\bar{V}$. The entries in $\hat{\bar{V}}$ that correspond to the correlation between the l^{th} lagged response and the j^{th} lagged response in a given variable to a given shock equals the corresponding entry in $\bar{V}$, multiplied by

$$\left[1 - \frac{|l-j|}{n}\right]^{\theta_{1,T}}, \; l, j = 1, ..., n.$$

Now consider the components of $\bar{V}$ that correspond to the correlations between components of different impulse response functions, either because a different variable is involved or because a different shock is involved, or both. We damp these entries in a way that damps more the greater is τ, the separation in time of the two impulses. In particular, we adopt the following damping factors for these entries:

$$\beta_T \left[1 - \frac{|\tau|}{n}\right]^{\theta_{2,T}}, \; \tau = 0, 1, ..., n.$$

We suppose that

$$\beta_T \to 1, \; \theta_{i,T} \to 0, \; T \to \infty, \; i = 1, 2,$$

where the rate of convergence is whatever is required to ensure consistency of $\hat{\bar{V}}$. These conditions leave completely open what values of β_T, $\theta_{1,T}$, $\theta_{2,T}$ we use in our sample. At one extreme, we have

$$\beta_T = 0, \; \theta_{1,T} = \infty,$$

and $\theta_{2,T}$ unrestricted. This corresponds to the approach in CEE and ACEL, in which $\hat{\bar{V}}$ is simply a diagonal matrix composed of the diagonal components of $\bar{V}$. At the other extreme, we could set β_T, $\theta_{1,T}$, $\theta_{2,T}$ at their $T \to \infty$ values, in which $\hat{\bar{V}} = \bar{V}$. In this draft, we work with the approach taken in CEE and ACEL. This has the important advantage of making our estimator particularly transparent. It corresponds to selecting θ so that the model implied impulse responses lie inside a confidence tunnel around the estimated impulses. When non-diagonal terms in $\bar{V}$ are also used, then the estimator worries not just about getting inside a confidence tunnel about the point estimates, but it is also concerned about the pattern of 'misses' across different impulse responses. Precisely how the off-diagonal components of $\bar{V}$ give rise to concerns about cross-impulse response patterns of misses is virtually impossible to understand. This is both because $\bar{V}$ is an enormous matrix and because it is not $\bar{V}$ itself that enters our criterion but its inverse.

6.4. Laplace Approximation

Because the likelihood we work with is only approximate, it is perhaps appropriate that we also work with an approximation to the posterior distribution. This is not essential, however, since Monte Carlo algorithms apply perfectly well in our setting, for computing marginal posteriors or θ and the marginal likelihood of $\hat{\psi}$.

To derive the Laplace approximation to $f\left(\theta|\hat{\psi}, V\left(\theta_0, \zeta_0, T\right)\right)$, define

$$g\left(\theta\right) \equiv \log f\left(\hat{\psi}|\theta, V\left(\theta_0, \zeta_0, T\right)\right) + \log p\left(\theta\right).$$

Let θ^* denote the mode of the posterior distribution and define the following Hessian matrix:

$$g_{\theta\theta} = -\frac{\partial^2 g\left(\theta\right)}{\partial\theta\partial\theta'}|_{\theta=\theta^*}.$$

Note that the matrix, $g_{\theta\theta}$, is an automatic by-product of standard gradient methods for computing the mode, θ^*. The second order Taylor series expansion of g about $\theta = \theta^*$ is:

$$g\left(\theta\right) = g\left(\theta^*\right) - \frac{1}{2}\left(\theta - \theta^*\right)' g_{\theta\theta}\left(\theta - \theta^*\right),$$

where the slope term is zero if θ^* is an interior optimum, which we assume. Then,

$$f\left(\hat{\psi}|\theta, V\left(\theta_0, \zeta_0, T\right)\right) p\left(\theta\right) \approx f\left(\hat{\psi}|\theta^*, V\left(\theta_0, \zeta_0, T\right)\right) p\left(\theta^*\right) \exp\left[-\frac{1}{2}\left(\theta - \theta^*\right)' g_{\theta\theta}\left(\theta - \theta^*\right)\right].$$

Note:

$$\frac{1}{\left(2\pi\right)^{\frac{m}{2}}}\left|g_{\theta\theta}\right|^{\frac{1}{2}} \exp\left[-\frac{1}{2}\left(\theta - \theta^*\right)' g_{\theta\theta}\left(\theta - \theta^*\right)\right]$$

is the $m-$variable Normal distribution for the m random variables, θ, with mean θ^* and variance-covariance matrix, $g_{\theta\theta}^{-1}$. By the standard property of a density function,

$$\int \frac{1}{\left(2\pi\right)^{\frac{n}{2}}}\left|g_{\theta\theta}\right|^{\frac{1}{2}} \exp\left[-\frac{1}{2}\left(\theta - \theta^*\right)' g_{\theta\theta}\left(\theta - \theta^*\right)\right] d\theta = 1. \tag{6.4}$$

Bringing together the previous results, we obtain:

$$\begin{aligned} f\left(\hat{\psi}|V\left(\theta_0, \zeta_0, T\right)\right) &= \int f\left(\hat{\psi}|\theta, V\left(\theta_0, \zeta_0, T\right)\right) p\left(\theta\right) d\theta \\ &\approx \int f\left(\hat{\psi}|\theta^*, V\left(\theta_0, \zeta_0, T\right)\right) p\left(\theta^*\right) \exp\left[-\frac{1}{2}\left(\theta - \theta^*\right)' g_{\theta\theta}\left(\theta - \theta^*\right)\right] d\theta \\ &= \left(2\pi\right)^{\frac{n}{2}}\left|g_{\theta\theta}\right|^{-\frac{1}{2}} f\left(\hat{\psi}|\theta^*, V\left(\theta_0, \zeta_0, T\right)\right) p\left(\theta^*\right), \end{aligned}$$

by (6.4). We now have the marginal distribution for $\hat{\psi}$. We can use this to compare the fit of different models for $\hat{\psi}$. In addition, we have an approximation to the marginal posterior distribution for an arbitrary element of θ, say θ_i :

$$\theta_i \tilde{} N\left(\theta_i^*, \left[g_{\theta\theta}^{-1}\right]_{ii}\right),$$

where $[A]_{ii}$ denotes the i^{th} diagonal element of the matrix, A.

7. Results

This section presents the results of estimating and evaluating the various models described in earlier sections. Table 1 displays the results for model parameters. It reports the name and symbol for the parameters as well as the parameter priors, the prior means and standard deviations, and the posterior modes and corresponding standard deviations. When posterior and prior standard deviations take on a similar value, this indicates there is little information in the data set about our parameter. Figures 3 - xx display the estimated dynamic responses of data to our three shocks: shocks to monetary policy, to the state of neutral technology and to the state of embodied technology.

7.1. VAR Results

We now briefly describe the impulse response functions implied by the VAR. The solid line in the figures indicates the point estimates of the impulse response function, while the gray area displays the corresponding two standard error confidence bounds.[28]

7.1.1. Monetary Policy Shocks

We first make sixth observations about the estimated dynamic effects of monetary policy shocks, displayed in Figure 3. The first observation about monetary policy shocks concerns the response of inflation. Here, there are two important things to note: the the price puzzle and the delayed and gradual response of inflation.[29] Note that in the very short run the point estimates indicate that inflation moves in a seemingly perverse direction in response to the expansionary monetary policy shock. This transitory drop in inflation in the immediate aftermath of a monetary policy shock has been widely commented on, and has been dubbed the 'price puzzle'. Christiano, Eichenbaum and Evans ( ) review the argument that the

[28]The standard errors were computed by the bootstrap procedure described in ACEL.

[29]Here, we have borrowed the adjects used by Mankiw (2000) to characterize the nature of the response of inflation to a monetary policy shock. Though he writes 10 years ago and he cites a wide range of evidence, Mankiw's conclusion about how inflation responds to a monetary policy shock resembles our VAR evidence very closely. He argues that the response of inflation to a monetary policy shock is gradual in the sense that it does not peak for 9 quarters.

puzzle may be the outcome of the sort of econometric specification error suggested by Sims ( ), and find evidence that is consistent with that view. Here, we follow ACEL and CEE in taking the position that there is no econometric specification error. It is important to note that the 'price puzzle' is not statistically significant. It nevertheless deserves comment because it has potentially very great economic significance. For example, the presence of a price puzzle in the data complicates the political problem associated with using high interest rates as a strategy to fight inflation. High interest rates and the consequent slowdown in economic growth is politically painful and if the public sees it producing higher inflation in the short run, support for the policy may evaporate unless the price puzzle has been explained.[30] Regarding the slow response of inflation, note how inflation reaches a peak after two years. Of course, the exact timing of the peak is not very well pinned down (note the wide confidence intervals). However, the evidence does suggest a sluggish response of inflation. This is consistent with the views of others, arrived at by other methods, about the slow response of inflation to a monetary policy shock.[31] It has been argued that this is a major puzzle for macroeconomics. For example, Mankiw (2000) argues that with price frictions of the type used here, the only way to explain the delayed and gradual response of inflation to a monetary policy shock is to introduce a degree of stickiness in prices that exceeds by far what can be justified based on the micro evidence. For this reason, when we study the ability of our models to match the estimated impulse response functions, we must be wary of the possibility that this is done only by making prices and wages very sticky. In addition, we must be wary of the possibility that the econometrics leans too hard on other features (such as variable capital utilization) to explain the gradual and delayed response of inflation to a monetary policy shock.

The third observation is that output, consumption, investment and hours worked all display a slow, hump-shape response to a monetary policy shock, peaking a little over one year after the shock. As emphasized in section 2, these hump-shape observations are the reason that researchers introduce habit persistence and costs of adjustment in the flow of investment into the baseline model. In addition, note that the effect of the monetary shock

[30] There is an important historical example of this political problem. In the early 1970s, at the start of the Great Inflation in the US, Arthur Burns was chairman of the US Federal Reserve and Wrigth-Patman was chairman of the House Banking Committee. Wright-Patman had the opinion that, by raising costs of production, high interest rates increase inflation. That is, he believed in the price puzzle, though he did not think it was only a transitory phenomenon, as VAR evidence suggests. Wright-Patman's belief had enormous significance because he was influential in writing the wage and price control legislation at the time. He threatened Burns that if Burns tried to raise interest rates to fight inflation, Wrigt-Patman would see to it that interest rates were brought under the control of the wage-price control board (cite Newsweek article about this).

[31] For example, Mankiw (2000) cites Hume's 1752 essay 'Of Money', in which Hume says that an increase in the money supply '..must first quicken the diligence of every individual, before it increases the price of labour.'

on the interest rate is roughly gone after one year, yet the economy continues to respond well after that. This suggests that to understand the dynamic effects of a monetary policy shock, one must have a model that displays considerable sources of internal propagation.

A fourth observation concerns the response of capacity utilization. Recall from the discussion of section 2 that the magnitude of the empirical response of this variable represents an important discipline on the analysis. In effect, those data constrain how heavily we can lean on variable capital utilization to explain the slow response of inflation to a monetary policy shock. The evidence in Figure 3 suggests that capacity utilization responds very sharply to a positive monetary policy shock. For example, it rises three times as much as employment. In interpreting this finding, we must bear in mind that the capital utilization numbers we have are for the manufacturing sector. To the extent that the data are influenced by the durable part of manufacturing, they may overstate the volatility of capacity utilization generally in the economy.

Our fifth set of observations have to do with the response of labor market variables such as unemployment, vacancies and the labor force, to a monetary policy shock. The labor force and vacancies both rise, while the rate of unemployment falls. The response of the labor force is very small, peaking at about 0.05 percent compared with the 0.2 percent peak in hours worked. With regard to vacancies and unemployment, we have in effect computed a 'monetary policy Beveridge curve'. The Beveridge curve is typically displayed as a graph with the percent deviation from the mean in vacancies on the vertical axis and the deviation in the unemployment rate (in percentage point terms) on the horizontal axis. When this is done with the raw data (see, e.g., Rob Shimer's website) one gets a very nearly straight line with slope (minus) 1/3: a three percentage point rise in unemployment is associated with a one percent change in vacancies. Our estimates indicate that monetary policy shocks alone also generate a negatively sloped Beveridge curve, but its slope is considerably greater. Figure 3 indicates that that slope is over to 20.[32]

Our sixth observation concerns the price of investment. In our model, this price is unaffected by shocks other than those to the technology for converting homogeneous output into investment goods. Figure 3 indicates that the price of investment rises in response to a positive shock to investment, contrary to our model. This suggests that it would be worth exploring modifications to the technology for producing investment goods so that the trade-off between consumption and investment is nonlinear.[33] Under these conditions, the rise in

[32] By quarter 6, the unemployment rate is down roughly 0.12 percentage points while vacancies are up roghly 2.5 percent.

[33] For example, instead of specifying a resource constraint in which $C_t + I_t$ appears, we could adopt one in which C_t and I_t appear in a CES function, i.e.,

$$\left[a_1 C_t^{1/\rho} + a_2 I_t^{1/\rho}\right]^{\rho}.$$

the investment to consumption ratio that appears to occur in response to an expansionary monetary policy shock would be associated with a rise in the price of invesment.

7.1.2. Technology Shocks

Figures 4 and 5 display the responses to neutral and embodied technology shocks. Overall, the confidence intervals are wide. A minute's reflection on the nature of the question being asked here suggests that this is to be expected. The VAR is povided with a remarkably subtle piece of information: that there are shocks in the data which have a long run effect on labor productivity. Imagine for a moment staring at a data plot and thinking about how one might detect such shocks. The VAR is not just asked to find these shocks, but it is also asked to determine its short run dynamic effects on a long list of variables. It's no wonder that in many cases, the VAR comes back with the answer, 'I don't know how this variable responds to a shock'. This is what the wide confidence intervals are about. For example, nothing much can be said about the response of capacity utilization to a neutral technology shock.

Though confidence intervals are wide, there does appear to be some information. For example, there is a significan rise in consumption, output, and hours worked in response to a neutral shock, according to these results. Also, while in many cases the confidence intervals include zero, in these cases they often rule out a significant fall or rise. For example, the results suggest that there is not a strong rise in the Federal Funds rate after a neutral technology shock.

One striking observation emerges from Figure 4. The results indicate that there is an immediate drop in inflation in the wake of positive shock to neutral technology. One might wonder if there is a tension here between Figure 3 and Figure 4. The former indicates that inflation is slow to respond to a monetary policy shock, and this is one important motivation for the assumption of frictions in price setting. But, price-setting frictions seem to suggest that prices should be sluggish in the face of all shocks. Alternatives to the wage/price frictions approach to macroeconomics are being developed and if there is a tension between Figures 3 and 4 on this dimension, this may be the tipoff that these other approaches are on the right track, while the focus on price frictions may be wrong-headed.[34] This is an example of how impulse responses based on VARs can provide potentially valuable input for researchers attempting to select between different models of the economy. These observations illustrate the argument Sims ( ) originally made for the use of impulse responses from VARs:

The standard linear specification is a special case of this one, with $a_1 = a_2 = \rho = 1$.

[34] Alternatives that could in principle account for a very different response of inflation to different shocks include those based on information confusion. See, for example, B and Wiederholt ( ), Lorenzoni ( ), and Mendes ( ).

they represent a useful way to organize data for economists wishing to distinguish between different economic theories.

7.2. Model Results

7.2.1. The Baseline Model

Consider first the results in Figures 3-5. Consider inflation first. Note that the baseline model does very well on that. It even captures the price puzzle. This reflects the effects of the working capital channel in the model. Also, inflation displays the delayed and gradual response that we see in the data. To see whether this is due to leaning on variable capital utilization too much, consider the response of capital utilization in the model. Note that the model misses badly on this variable. From an economic perspective, however, the key thing is that it misses because it understates the movement in utilization. This is a far less disturbing result than if the model had overshot, because that would have placed a dark cloud over the model's apparent success at accounting for the inertial response of inflation to a monetary shock. Note that in terms of the other variables, the model does quite well. For example, it captures the hump shape in the response of output, consumption, investment and hours worked. The model predicts no particular change in the real wage, reflecting the presence of frictions both in wages and prices. This is consistent with the data, although in part this is due to the very wide confidence intervals on this variable. As noted before, the model misses on the response of the price of investment to a monetary policy shock.

Turning to the technology shocks, note that the model does fine on virtually all the variables. Of course, to some extent this reflects the relative lack of precision in our estimates of the response of variables to shocks. An interesting exception is the response of inflation to a neutral technology shock. Note that the model understates the drop in inflation in the first period. In the periods afterward, the model lies in the gray area. In a statistical sense, the miss between model and data is perhaps small. It involves only one impulse. Still, as discussed above, the miss might be very important from an economic point of view if it persuades researchers to look for an entirely different model of aggregate fluctuations.

To investigate the significance of the miss, we increased the weight on the inflation impact effect of a neutral technology shock in our estimation criterion. We did this by reducing the variance of that impact effect by a factor of 10, just enough so that the impact effect on inflation lies on the boundary of the gray area when the model is re-estimated. The resulting impulse response functions of the resulting model are the ones that correspond to the 'baseline model' in Figure 6-8. Note that the model impulse response functions now lie inside all the gray areas. The forced success on the inflation impact of a neutral technology shock was not obtained by a serious miss on other dimensions. Provisionally, we infer that the baseline model's failure to reproduce the impact effect of inflation in Figure 4 is not necessarily of great economic significance.

Whether we take seriously the apparent good fit of the baseline model in Figures 3-8 depends on what parameter values were required to achieve that fit. To see this, we turn to Table 1. That table reports the priors and posteriors on each parameter, including the associated standard deviations. Note that the mode of the posterior distribution on ξ_p implies that there are on average 1.5 quarters between price reoptimizations. This is a tiny amount of price stickiness. At the same time, the degree of price stickiness on wages is unrealistically large. The mode of ξ_w implies that there are on average about 8.5 quarters between wage reoptimizations. We can also see what happened to the parameters when we adjusted the criterion to force the baseline model to hit the contemporaneous response of inflation to a technology shock. The primary effect was to reduce the degree of price indexation on inflation, κ_f, to basically zero.

7.2.2. Baseline Model Extended to Include Unemployment

Figures 3-5 display the impulse responses of the version of the baseline model proposed by Gali (2009) (see 'base. + unemp'). This model was fit to the same impulse responses as the baseline model, plus the impulse responses of the labor force and the unemployment rate. Interestingly, this model does better than the baseline model. First, does much better on the inflation response to a neutral shock. Second, the fact that the model also addresses labor market variables is very important. Moreover, Figures 3-5 indicate that the model does reasonably well one those variables

The model also does better in terms of some of the parameters (see Table 1). For example, the model does reasonably well on ξ_p and ξ_w. On other dimensions the model parameter values are similar to those of the baseline model.

However, these results do not imply a clean success for the model. This is because the success required including the labor externality described in section 3. We specified the externality, ζ_t, as follows:

$$\zeta_t = \left(\frac{L_t}{L}\right)^{-\gamma_1} \left(\frac{L_t}{L_{t-1}}\right)^{-\gamma_2},$$

where L_t represents aggregate hours of work (see (5.20)). Note that each of γ_1 and γ_2 has a posterior mode that is about 10 times greater than its posterior standard deviation. This reflects our discussion in section 3, where we noted that when γ_1=$\gamma_2 = 0$, the model implies a large fall in the labor force in the wake of a positive technology shock and an extremely large drop in unemployment. This is because the value of working falls as consumption rises after a positive monetary policy shock. When γ_1=$\gamma_2 = 0$ the estimation procedure drives $\sigma_L \to \infty$ so that the labor force ceases to respond to the monetary policy shock. Interestingly, the effect of driving the labor supply elasticity to zero has little effect elsewhere because the presence of sticky wages effectly removes labor supply from the system. However, driving $\sigma_L \rightarrow \infty$ does have the consequence of either (i) driving steady state unemployment to zero or, (ii) if λ_w is selected to always enforce an empirical estimate of the unemployment rate, then $\lambda_w \rightarrow \infty$ as $\sigma_L \rightarrow \infty$. The problem is technically avoided with the externality. However, absent an interesting economic interpretation of the externality, the externality must be interpreted as a measure of the lack of fit of the model.

7.2.3. Baseline Model with Search and Matching

Figures 9-11 display the impulse responses of the baseline model with search and matching, evaluated at the mode of its parameter values. The corresponding parameter values are reported in Table 1. Because we had difficulty estimating all the parameters in the model, in this draft of the paper we report results based on estimating a subset of the parameters. The four monetary policy rule parameters and the four parameters controlling the stochastic process driving the exogenous shocks were simply fixed at their modes in the baseline model. This at least gives us an ability to compare results with that model.

Our assumption of four quarter Taylor contracts on wages in this model limits the degree of wage stickiness. That a lot of wage stickiness is required to explain the data is suggested by the results for the baseline model. Evidence that the search and matching model is short on wage frictions can be seen in the fact that that model substantially overshoots the real wage response to a monetary policy shock. The search and matching model appears to have responded to the enforced lack of wage frictions by selecting a very high degree of stickiness in prices (see Table 1). Despite this, the search and matching model's account of the inertial response of inflation to a monetary policy shock is not as good as the baseline model's.

Given its predictions about labor costs, it is perhaps not surprising that the search and matching model undershoots the rise in total hours worked in the wake of a monetary policy shock. Interestingly, the model nevertheless overshoots output in the immediate aftermath of a monetary shock. This is accomplished by a counterfactually large jump in capacity utilization.

Turning to the labor market variables, the search and matching model mimics the 'monetary policy Beveridge curve' fairly well. The initial percent change in vacancies is about 20 times the percentage point change in unemployment. However, the model undershoots the response of both unemployment and vacancies.

Turning to the neutral technology shock, the model does very poorly on the response of inflation to a neutral technology shock. This is perhaps not surprising in view of the very high degree of price stickyness in the model. Generally, the model predicts far too little response in variables to a neutral technology shock. Because confidence intervals are generally high in the case of the neutral technology shock, this leaves the model impulse responses in the gray area for many of the variables. Three places where the model lies outside the gray area are consumption, hours worked and inflation. In terms of the 'neutral technology shock Beveridge curve', the model does reasonably well. However, the sign and magnitude of the vacancy and unemployment responses to a neutral technology shock are both wrong. The model expects a fall in unemployment and a rise in vacancies in the wake of a positive technology shock, but the data seem to indicate the opposite. The rise in unemployment in the data appears to reflect a rise in the labor force, something that is held constant in the model. It is perhaps puzzling that the data show a fall in vacancies. A shortcoming of the VAR analysis is that one has to always wonder whether a given impulse response function is correctly interpreted (in this case, as reflecting that vacancies actually do fall in response to a neutral shock), or whether there is specification error in the identifying assumptions.

Finally, turning to the investment specific shock the model once again predicts little response in the variables. Here, the gray areas are all wide enough that this does not set off any alarms about model failure.

Consider model parameters in Table 1. We have already mentioned the price stickyness parameter. A parameter that receives a great deal of attention in the search and matching literature is the 'replacement ratio': the ratio of the payment to unemployed workers, b^u, to the flow utility of unemployed workers, in steady state. In our analysis we treat this ratio as a parameter. Conditional on the other model parameters, we use the value of the replacement ratio to back out b^u. Thus, we drop b^u from the parameter space and replace it by the replacement ratio. Our priors constrain the value of the replacement ratio to lie very close to 0.70. Presumably, if we allowed this parameter to be bigger, the model would predict a stronger response of vacancies to a shock.

8. Conclusion

The baseline model performs fairly well in the econometric inference exercise done here. We extended the model to include unemployment, and it is clear that there remain challenges here.

References

[1] Altig, David, Christiano, Lawrence J, Martin Eichenbaum and Jesper Linde (2004), "Firm-Specific Capital, Nominal Rigidities and the Business Cycle", manuscript, Northwestern University.

[2] Ball, Laurence, 1994, 'Credible Disinflation with Staggered Price Setting', American Economic Review, vol. 84, March, pp. 282-289.

[3] Barro, Robert (1977), "Long-Term Contracting, Sticky Prices and Monetary Policy," Journal of Monetary Economics, Vol 3(3), 305-316.

[4] Chernozhukov, Victor and Han Hong, 2003, An MCMC approach to classical estimation, Journal of Econometrics 115, pp 293 – 346

[5] Christiano, Lawrence J., 2004, "Firm-Specific Capital and Aggregate Inflation Dynamics in Woodford's Model," manuscript.

[6] Christiano, Lawrence J., 2007, discussion of del Negro, Schorfheide, Smets and Wouters, Journal of Business and Economic Statistics, April, Vol. 25, No. 2.

[7] Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans, 1999, "Monetary Policy Shocks: What Have We Learned and to What End?", Handbook of Macroeconomics, vol. 1A, eds., Michael Woodford and John Taylor, Amsterdam, New York and Oxford: Elsevier Science, North-Holland.

[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, 113(1): 1–45.

[9] Christiano, Lawrence J., Cosmin Ilut, Roberto Motto and Massimo Rostagno (2007), "Monetary Policy and Stock Market Boom-Bust Cycles", manuscript, Northwestern University.

[10] den Haan, Wouter, Garey Ramey and Joel Watson (2000), "Job Destruction and Propagation of Shocks", American Economic Review, Vol 90(3), 482-498.

[11] del Negro, Marco, Frank Schorfheide, Frank Smets and Rafael Wouters, "On the Fit of New Keynesian Models", Journal of Business & Economic Statistics, April 2007, Vol. 25, No. 2.

[12] de Walque, Gregory, Frank Smets and Raf Wouters, 2006, "Firm-specific Production Factors in a DSGE Model with Taylor Price Setting," International Journal of Central Banking, 2 (3) 107-154.

[13] Eichenbaum, Martin and Jonas Fisher, 2007, "Estimating the Frequency of Price Reoptimization in Calvo-style Models," Journal of Monetary Economics, 54 (7) 2032-2047.

[14] Erceg, Christopher J., Henderson, Dale, W., and Andrew T. Levin (2000), "Optimal Monetary Policy with Staggered Wage and Price Contracts," Journal of Monetary Economics, Vol 46, 281-313.

[15] Fisher, Jonas (1998), "Credit market imperfections and the heterogeneous response of firms to monetary shocks", Federal Reserve Bank of Chicago, Working Paper Series, 96-23.

[16] Fujita, Shigeru and Garey Ramey (2006) "The Cyclicality of Job Loss and Hiring," manuscript.

[17] Gertler, Mark and Antonella Trigari (2008), "Unemployment Fluctuations with Staggered Nash Bargaining," Journal of Political Economy, forthcoming.

[18] Gertler, Mark, Luca Sala and Antonella Trigari (2008), "An Estimated Monetary DSGE Model with Unemployment and Staggered Nominal Wage Bargaining," Journal of Money, Credit and Banking, December.

[19] Guerron-Quintana, Pablo, 2008, "Refinements on Macroeconomic Modeling: The Role of Non-Separability and Heterogeneous Labor Supply," Journal of Economic Dynamics and Control 32, pp: 3613 – 3630.

[20] Hall, Alastair, Atsushi Inoue, James M. Nason and Barbara Rossi, 2009, "Information Criteria for Impulse Response Function Matching Estimation of DSGE Models," manuscript, Duke University.

[21] Hall, Robert (2005a), "Employment Fluctuations with Equilibrium Wage Stickiness," American Economic Review, Vol 95(1), 50-65.

[22] Hall, Robert (2005b), "Employment Efficiency and Sticky Wages: Evidence from Flows in the Labor Market," Review of Economics and Statistics, Vol 87(3), 397-407.

[23] Hall, Robert (2005c), "Job Loss, Job Finding, and Unemployment in the U.S. Economy over the Past Fifty Years," NBER Macroeconomics Annual, M. Gertler and K. Rogoff editors, MIT Press, 101-137.

[24] Kimball, Miles, 1995, "The Quantitative Analytics of the Basic Neomonetarist Model,' Journal of Money Credit and Banking, 27 (4), part 2, 1241-1277.

[25] Mankiw, N. Gregory, 2000, 'The Inexorable and Mysterious Tradeoff Between Inflation and Unemployment,' National Bureau of Economic Research Working Paper 7884.

[26] Mortensen, Dale and Christopher Pissarides (1994), "Job Creation and Job Destruction in the Theory of Unemployment," Review of Economic Studies, Vol 61, 397-415.

[27] Newey, Whitney K. and Kenneth D. West (1987) "A Simple, Positive Semi-definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix", Econometrica, vol. 55(3), 703–708.

[28] Rotemberg, Julio and Michael Woodford (1997): "An Optimization-Based Econometric Framework for the Evaluation of Monetary Policy", in Bernanke, B. and J. Rotemberg, eds., NBER, Macroeconomics Annual, Cambridge, MA: MIT Press.

[29] Shimer, Robert (2005a), "The Cyclical Behavior of Equilibrium Unemployment Vacancies, and Wages: Evidence and Theory," American Economic Review, Vol 95(1), 25-49.

[30] Shimer, Robert (2005b), "Reassessing the Ins and Outs of Unemployment," manuscript, University of Chicago.

[31] Sims, Christopher A. (1992), 'Interpreting the Macroeconomic Time Series Facts: The Effects of Monetary Policy', European Economic Review, Vol. 36, No. 5, pages 975-1000.

[32] 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): 1123–75.

[33] Smets, Frank, and Raf Wouters (2007), "Shocks and Frictions in US Business Cycles", American Economic Review, Vol. 97 No. 3: 586-606.

[34] Thomas, Carlos, 2009, "Search frictions, real rigidities and inflation dynamics," manuscript, Bank of Spain.

[35] Woodford, Michael, 2004, "Inflation and Output Dynamics with Firm-Specific Investment," manuscript.

[36] Yun, Tack (1996), "Nominal Price Rigidity, Money Supply Endogeneity, and Business Cycles," Journal of Monetary Economics, Vol 37(2), 345-370.

Table 1: Economic Parameters – Priors and Posteriors

Parameter		Prior Distr	Prior mean (Std Dev)	Posterior Mode (Std Dev)				
				Baseline	**Baseline + Galí**	**Baseline + high weight**	**Basline + high σ^L**	**Baseline + Search & Matching**
Price and wage setting parameters								
Price Stickiness	ξ_f	Beta	0.75	0.343	0.759	0.615	0.361	0.894
			(0.15)	(0.094)	(0.048)	(0.035)	(0.097)	(0.005)
Wage Stickiness	ξ_w	Beta	0.75	0.882	0.781	0.871	0.871	-.-
			(0.15)	(0.012)	(0.02)	(0.014)	(0.014)	-.-
Price Markup	λ_f	Normal	1.2	1.183	1.342	1.146	1.19	1.592
			(0.1)	(0.057)	(0.059)	(0.056)	(0.057)	(0.067)
Wage Markup	λ_w	Normal	1.2	1.199	-.-	1.199	1.199	-.-
			(0.05)	(0.05)	-.-	(0.05)	(0.05)	-.-
Price Indexation	κ_f	Beta	0.5	0.445	0.16	0.031	0.416	0.963
			(0.2)	(0.265)	(0.09)	(0.017)	(0.259)	(0.019)
Monetary Authority parameters								
Taylor Rule: Interest Smoothing	ρ_R	Beta	0.8	0.868	0.864	0.868	0.865	0.868
			(0.1)	(0.011)	(0.013)	(0.011)	(0.011)	(0.0)
Taylor Rule: Inflation Coef.	r_π	Gamma	1.6	1.507	1.423	1.403	1.522	1.51
			(0.15)	(0.118)	(0.103)	(0.099)	(0.121)	(0.0)
Taylor Rule: GDP Coef.	r_y	Gamma	0.2	0.065	0.066	0.075	0.062	0.065
			(0.15)	(0.017)	(0.017)	(0.021)	(0.017)	(0.0)
Household Parameters								
Consumption Habit	b	Beta	0.75	0.772	0.76	0.769	0.768	0.757
			(0.15)	(0.013)	(0.015)	(0.013)	(0.014)	(0.012)
Labor Disutil. Curv.	σ_L	Gamma	1.0	-.-	2.754	-.-	-.-	-.-
			(0.5)	-.-	(0.601)	-.-	-.-	-.-
Capacity Adj. Costs Curv.	σ_a	Gamma	1.0	0.265	0.133	0.219	0.223	0.189
			(0.75)	(0.049)	(0.061)	(0.047)	(0.047)	(0.028)
Inv. Adj. Costs Curv.	S''	Gamma	8.0	15.367	17.281	14.884	15.089	11.44
			(5.0)	(2.67)	(2.653)	(2.681)	(2.698)	(1.763)
Working Capital Fraction	ν^f	Beta	0.5	0.225	0.452	0.434	0.167	0.225
			(0.2)	(0.065)	(0.147)	(0.146)	(0.06)	(0.0)
Labor Externality - Level	γ_1	Normal	0.0	-.-	2.415	-.-	-.-	-.-
			(5.0)	-.-	(0.302)	-.-	-.-	-.-
Labor Externality - Growth Rate	γ_2	Normal	0.0	-.-	7.237	-.-	-.-	-.-
			(5.0)	-.-	(0.647)	-.-	-.-	-.-

-.- means parameter was fixed at its prior mean.

Table 1, Continued

Parameter		Prior Distr	Prior mean (Std Dev)	Posterior Mode (Std Dev)				
				Baseline	**Baseline + Galí**	**Baseline + high weight**	**Basline + high σ^L**	**Baseline + Search & Matching**
Search and Matching								
SS Endog. Separation Rate (in %)	$F(\bar{a})$	Gamma	0.5	-/-	-/-	-/-	-/-	0.108
			(0.1)	-/-	-/-	-/-	-/-	(0.023)
Replacement Ratio	$\frac{b_u}{\text{flow utility}}$	Beta	0.7	-/-	-/-	-/-	-/-	0.73
			(0.01)	-/-	-/-	-/-	-/-	(0.01)
SS Recr't costs/Output(in %)	$\frac{\text{recr't cost}}{\text{output}}$	Gamma	0.5	-/-	-/-	-/-	-/-	0.697
			(0.1)	-/-	-/-	-/-	-/-	(0.099)
Share in Matching Function	σ_m	Beta	0.5	-/-	-/-	-/-	-/-	0.386
			(0.1)	-/-	-/-	-/-	-/-	(0.045)
Shocks								
Autocorr. Neutral Tech.	ρ_z	Beta	0.75	0.858	0.851	0.754	0.865	0.853
			(0.15)	(0.011)	(0.025)	(0.045)	(0.011)	(0.006)
Autocorr. Invest. Tech.	ρ_ψ	Beta	0.75	0.786	0.742	0.764	0.797	0.776
			(0.15)	(0.028)	(0.048)	(0.035)	(0.027)	(0.011)
Std. Neutral Tech. Shock	σ_z	Inv. Gamma	0.05	0.048	0.051	0.074	0.046	0.048
			(4.0)	(0.003)	(0.007)	(0.01)	(0.002)	(0.0)
Std. Invest. Tech. Shock	σ_ψ	Inv. Gamma	0.05	0.093	0.105	0.101	0.089	0.093
			(4.0)	(0.009)	(0.015)	(0.011)	(0.009)	(0.0)
Std. Monetary Shock	σ_R	Inv. Gamma	0.15	0.128	0.128	0.126	0.129	0.128
			(4.0)	(0.006)	(0.007)	(0.006)	(0.006)	(0.0)

-/- means parameter is not defined in indicated variant of the model.

Standard deviation of 0 means parameter was fixed at the value indicated.

Table 2: Diagnosing the Fit of the Baseline Model

Parameter		Prior Distr	Prior mean (Std Dev)	Posterior Mode (Std Dev)			
				Baseline	**Baseline MP Shock only**	**Baseline – Techn. Shock only**	**Basline – Inv. Techn. Shock only**
Price and wage setting parameters							
Price Stickiness	ξ_f	Beta	0.75	0.343	0.581	0.823	0.798
			(0.15)	(0.094)	(0.101)	(0.068)	(0.087)
Wage Stickiness	ξ_w	Beta	0.75	0.882	0.798	0.903	-.-
			(0.15)	(0.012)	(0.024)	(0.014)	-.-
Price Markup	λ_f	Normal	1.2	1.183	1.1	1.2	1.239
			(0.1)	(0.057)	(0.086)	(0.094)	(0.098)
Wage Markup	λ_w	Normal	1.2	1.199	1.201	1.195	-.-
			(0.05)	(0.05)	(0.05)	(0.05)	-.-
Price Indexation	κ_f	Beta	0.5	0.445	0.807	0.226	0.595
			(0.2)	(0.265)	(0.152)	(0.141)	(0.374)
Monetary Authority parameters							
Taylor Rule: Interest Smoothing	ρ_R	Beta	0.8	0.868	0.875	0.909	0.808
			(0.1)	(0.011)	(0.015)	(0.027)	(0.073)
Taylor Rule: Inflation Coef.	r_π	Gamma	1.6	1.507	1.452	1.535	1.62
			(0.15)	(0.118)	(0.118)	(0.135)	(0.145)
Taylor Rule: GammaDP Coef.	r_y	Gamma	0.2	0.065	0.097	0.117	0.171
			(0.15)	(0.017)	(0.07)	(0.034)	(0.072)
Household parameters							
Consumption Habit	b	Beta	0.75	0.772	0.814	0.551	0.469
			(0.15)	(0.013)	(0.019)	(0.137)	(0.129)
Capacity Adj. Costs Curv.	σ_a	Gamma	1.0	0.265	0.101	2.069	0.088
			(0.75)	(0.049)	(0.075)	(0.758)	(0.048)
Inv. Adj. Costs Curv	S''	Gamma	8.0	15.367	10.066	9.59	9.967
			(5.0)	(2.67)	(2.031)	(4.502)	(4.555)
Working Capital Fraction	ν^f	Beta	0.5	0.225	0.639	0.488	0.606
			(0.2)	(0.065)	(0.153)	(0.271)	(0.26)
Labor Disutil. Curv.	σ_L	Gamma	1.0	-/-	-/-	-/-	-/-
			(0.5)	-/-	-/-	-/-	-/-
Labor Externality – Level	γ_1	Normal	0.0	-/-	-/-	-/-	-/-
			(5.0)	-/-	-/-	-/-	-/-
Labor Externality – Growth Rate	γ_2	Normal	0.0	-/-	-/-	-/-	-/-
			(5.0)	-/-	-/-	-/-	-/-

-.- means parameter was fixed at its prior mean.

-/- means parameter is not defined in indicated variant of the model.

Table 2, continued

Parameter		Prior Distr	Prior mean (Std Dev)	Posterior Mode (Std Dev)			
				Baseline	**Baseline MP Shock only**	**Baseline – Techn. Shock only**	**Basline – Inv. Techn. Shock only**
Search and Matching							
SS Endog Separation Rate (in %)	$F(\bar{a})$	Gamma	0.5	-/-	-/-	-/-	-/-
			(0.1)	-/-	-/-	-/-	-/-
Replacement Ratio	$\frac{b_u}{\text{flow utility}}$	Beta	0.7	-/-	-/-	-/-	-/-
			(0.01)	-/-	-/-	-/-	-/-
SS Recr't costs/Output(in %)	$\frac{\text{recr't cost}}{\text{output}}$	Gamma	0.5	-/-	-/-	-/-	-/-
			(0.1)	-/-	-/-	-/-	-/-
Share in Matching Fun.	σ_m	Beta	0.5	-/-	-/-	-/-	-/-
			(0.1)	-/-	-/-	-/-	-/-
Shocks							
Autocorr. Neutral Tech.	ρ_z	Beta	0.75	0.858	-/-	0.766	-/-
			(0.15)	(0.011)	-/-	(0.062)	-/-
Autocorr. Invest. Tech.	ρ_ψ	Beta	0.75	0.786	-/-	-/-	0.715
			(0.15)	(0.028)	-/-	-/-	(0.057)
Std. Neutral Tech. Shock	σ_z	Inv. Gamma	0.05	0.048	-/-	0.086	-/-
			(4.0)	(0.003)	-/-	(0.019)	-/-
Std. Invest. Tech. Shock	σ_ψ	Inv. Gamma	0.05	0.093	-/-	-/-	0.115
			(4.0)	(0.009)	-/-	-/-	(0.018)
Std. Monetary Shock	σ_R	Inv. Gamma	0.15	0.128	0.127	-/-	-/-
			(4.0)	(0.006)	(0.008)	-/-	-/-

-/- means parameter is not defined in indicated variant of the model.

Figure 3: Response to Monetary Policy Shock

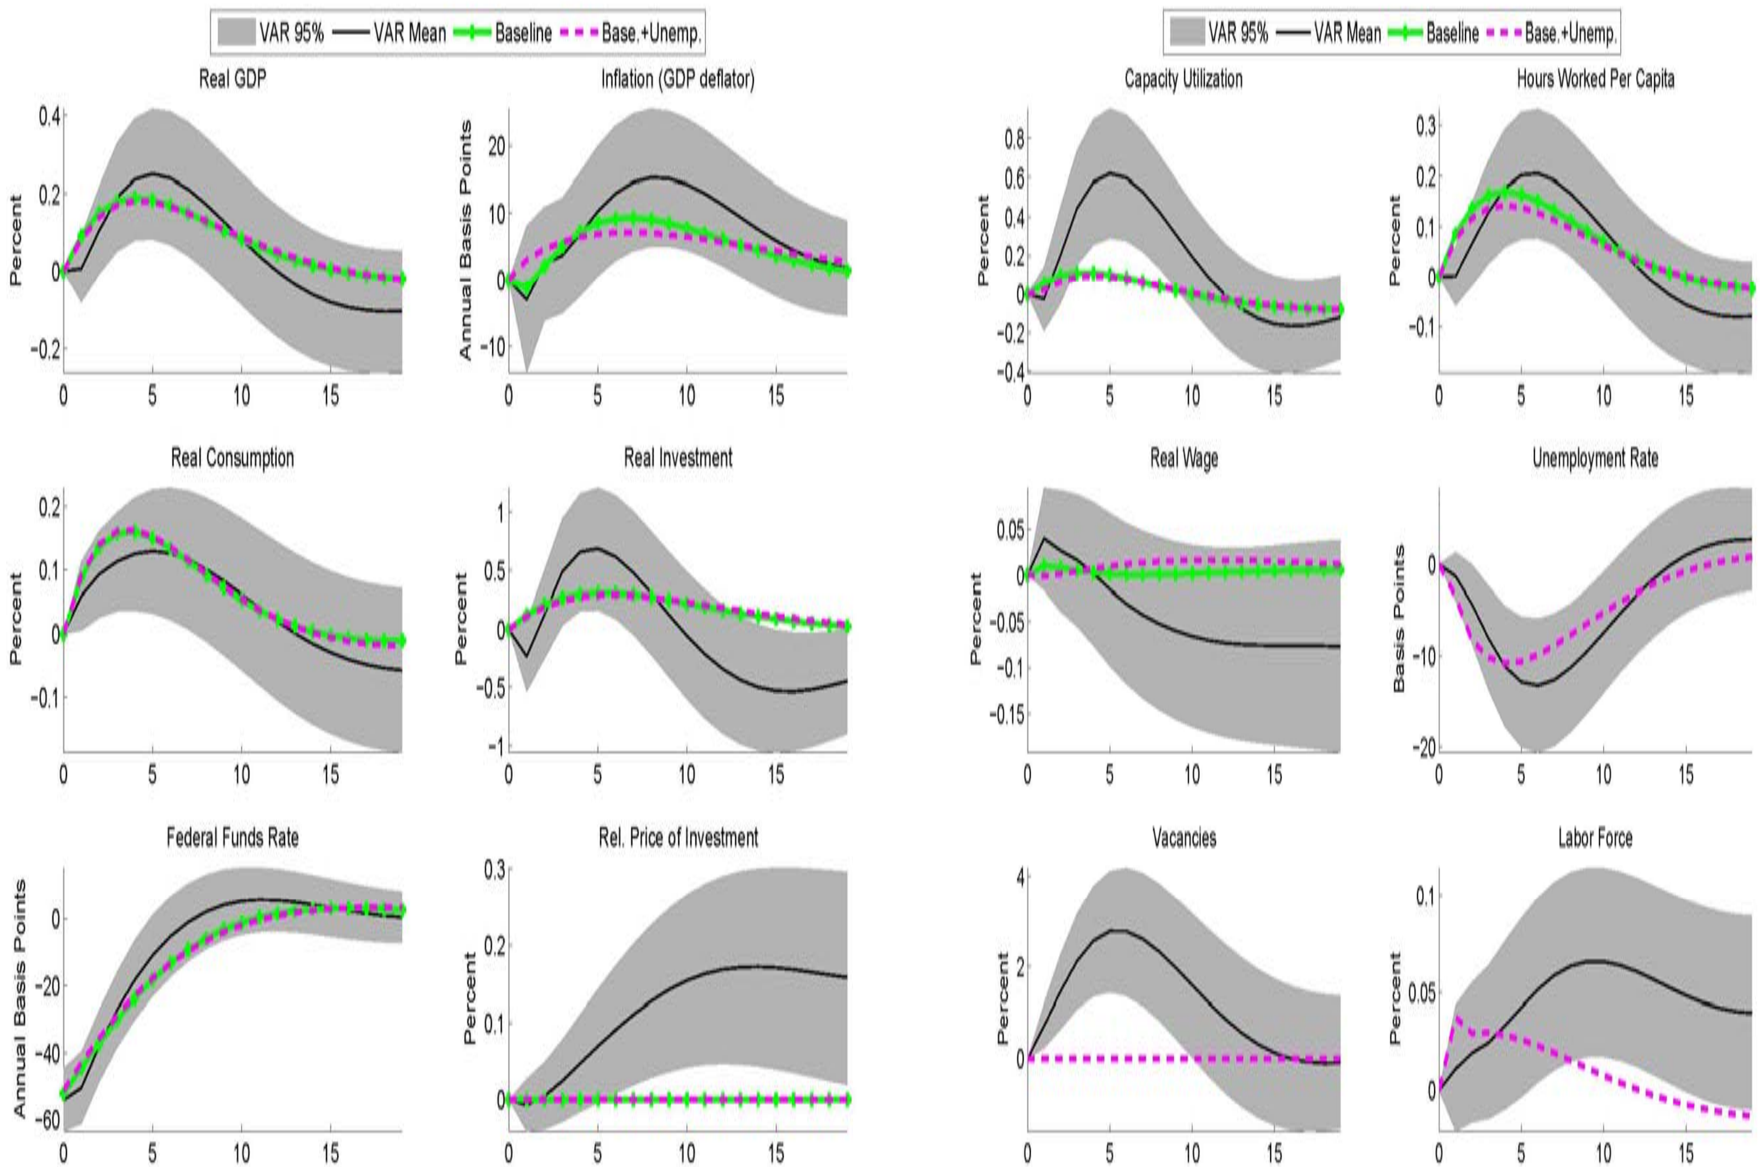

Figure 4: Response to Neutral Technology Shock

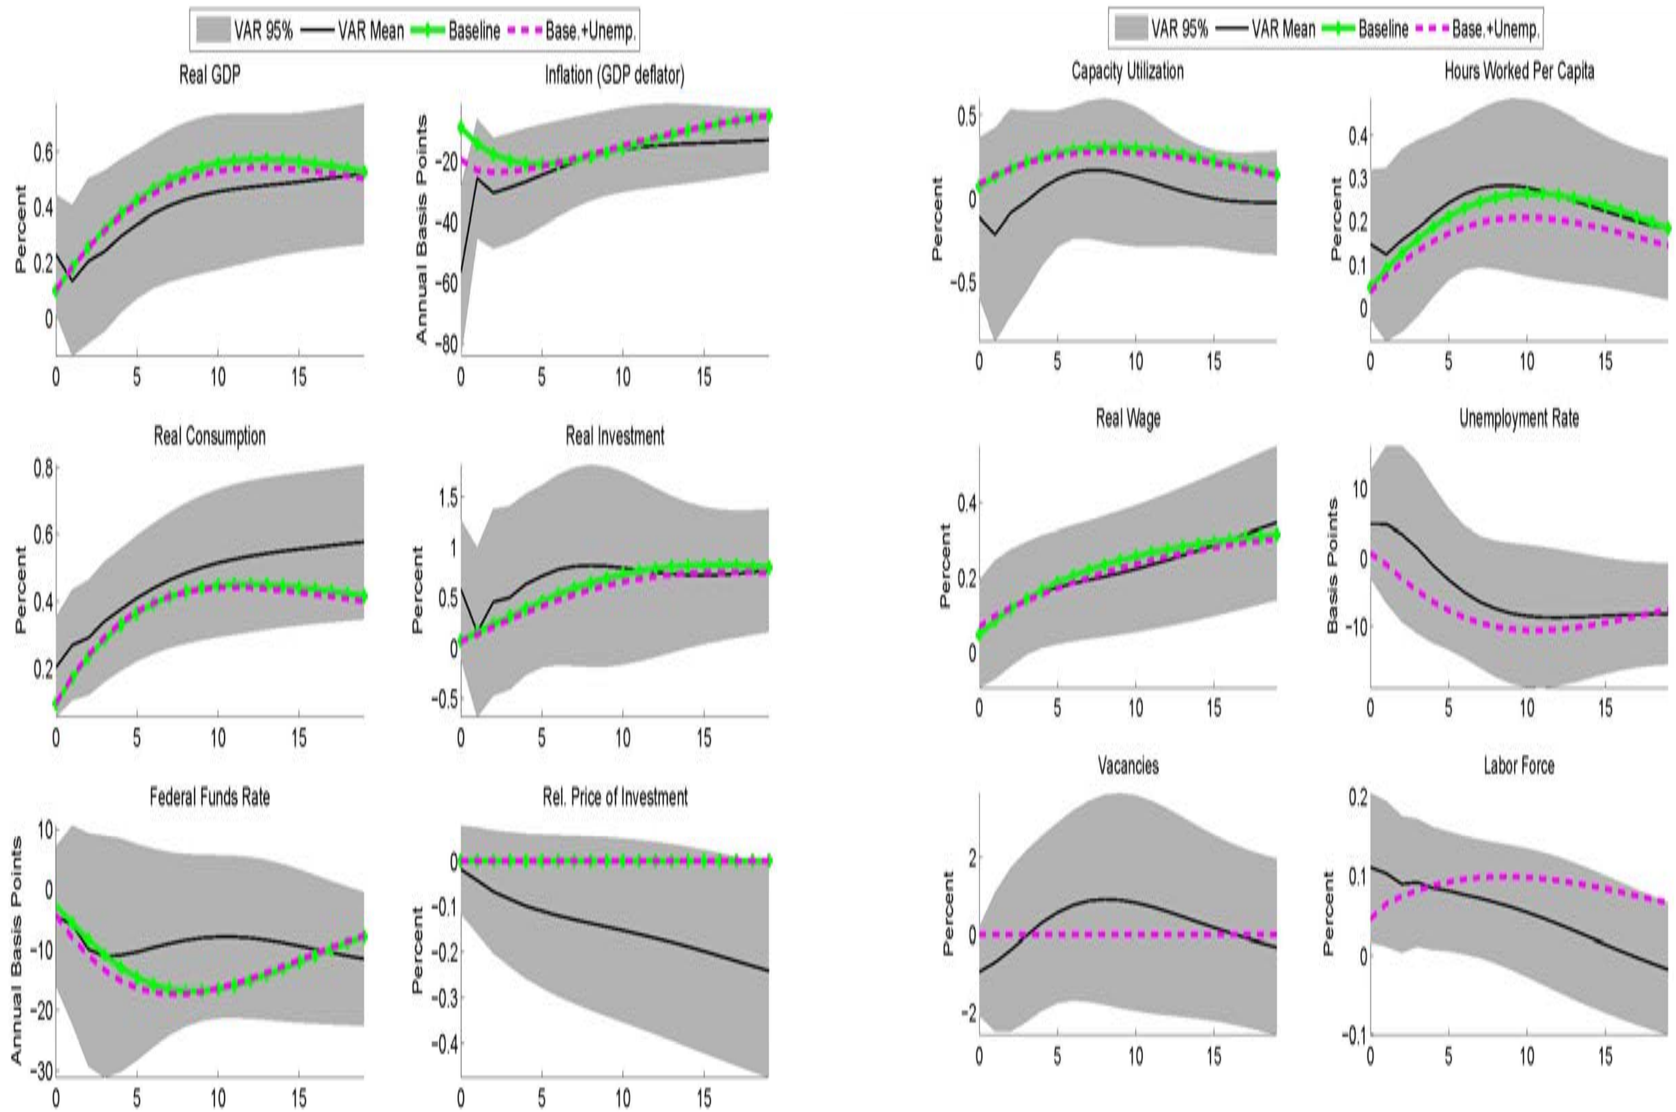

Figure 5: Response to Capital Embodied Shock

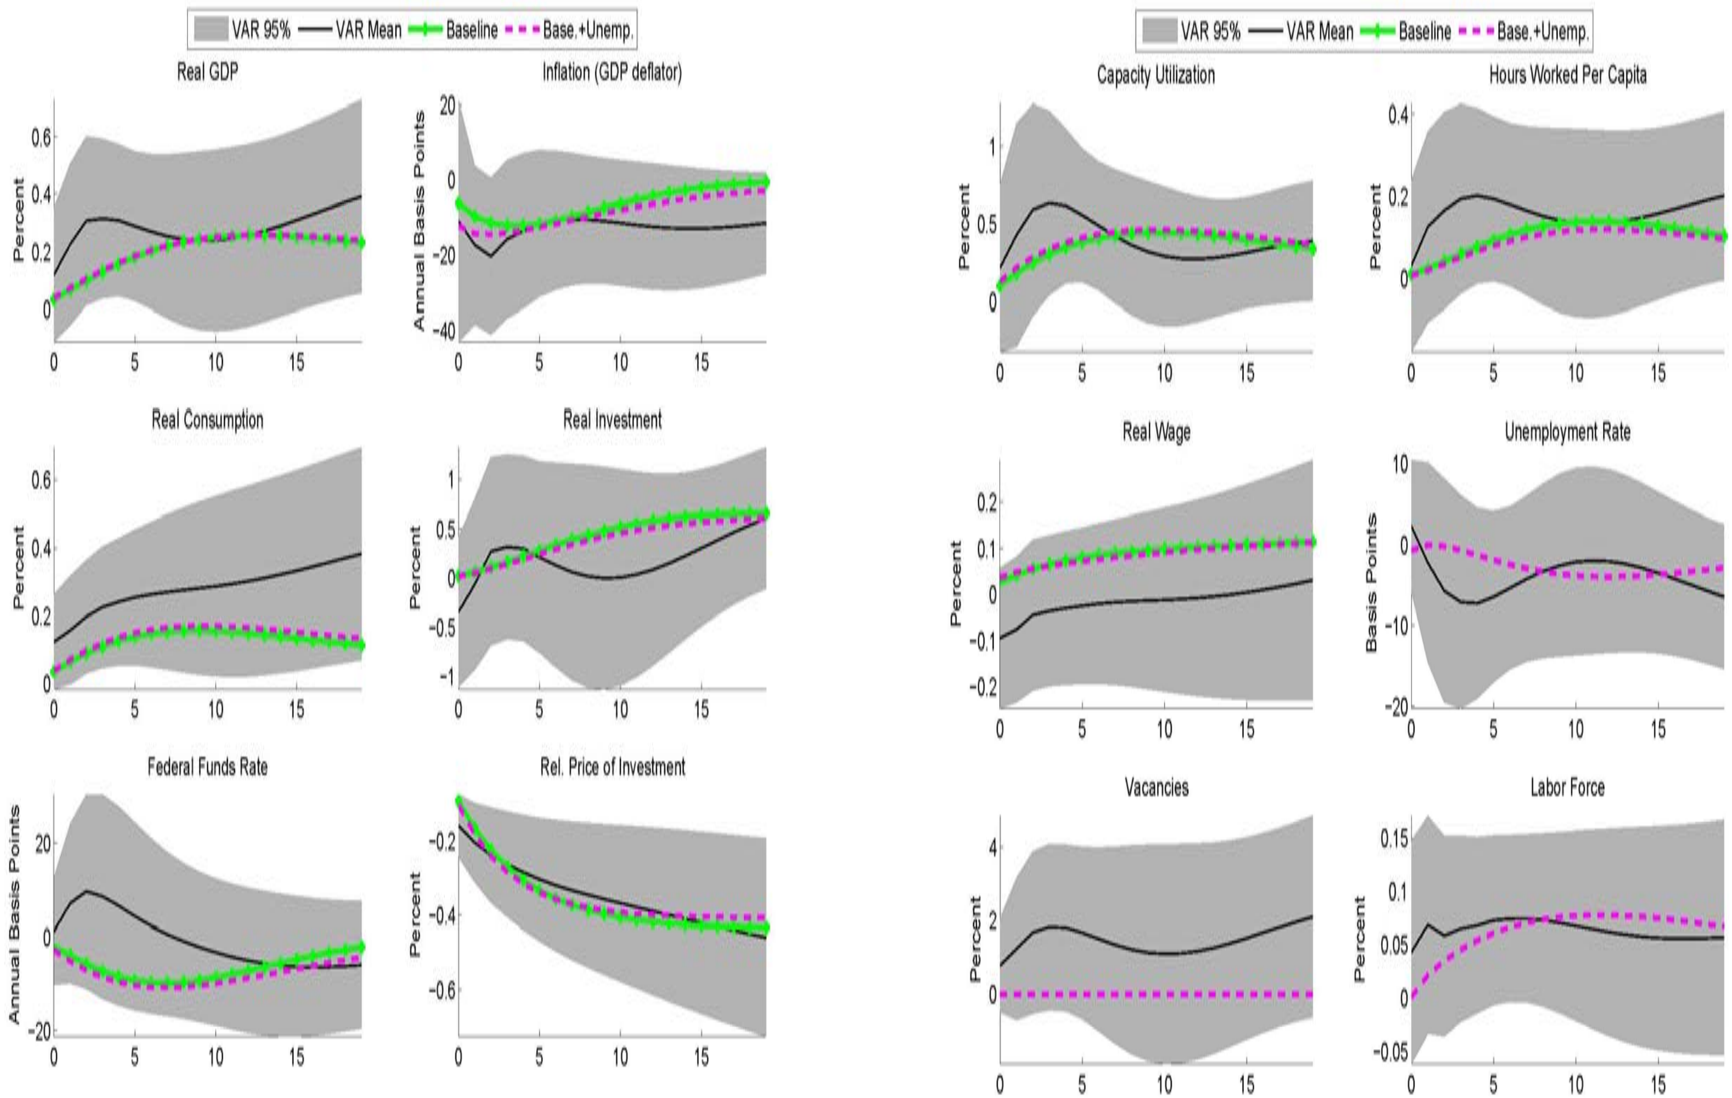

Figure 6: Response to Money Shock, Baseline W/High Inflation Weight

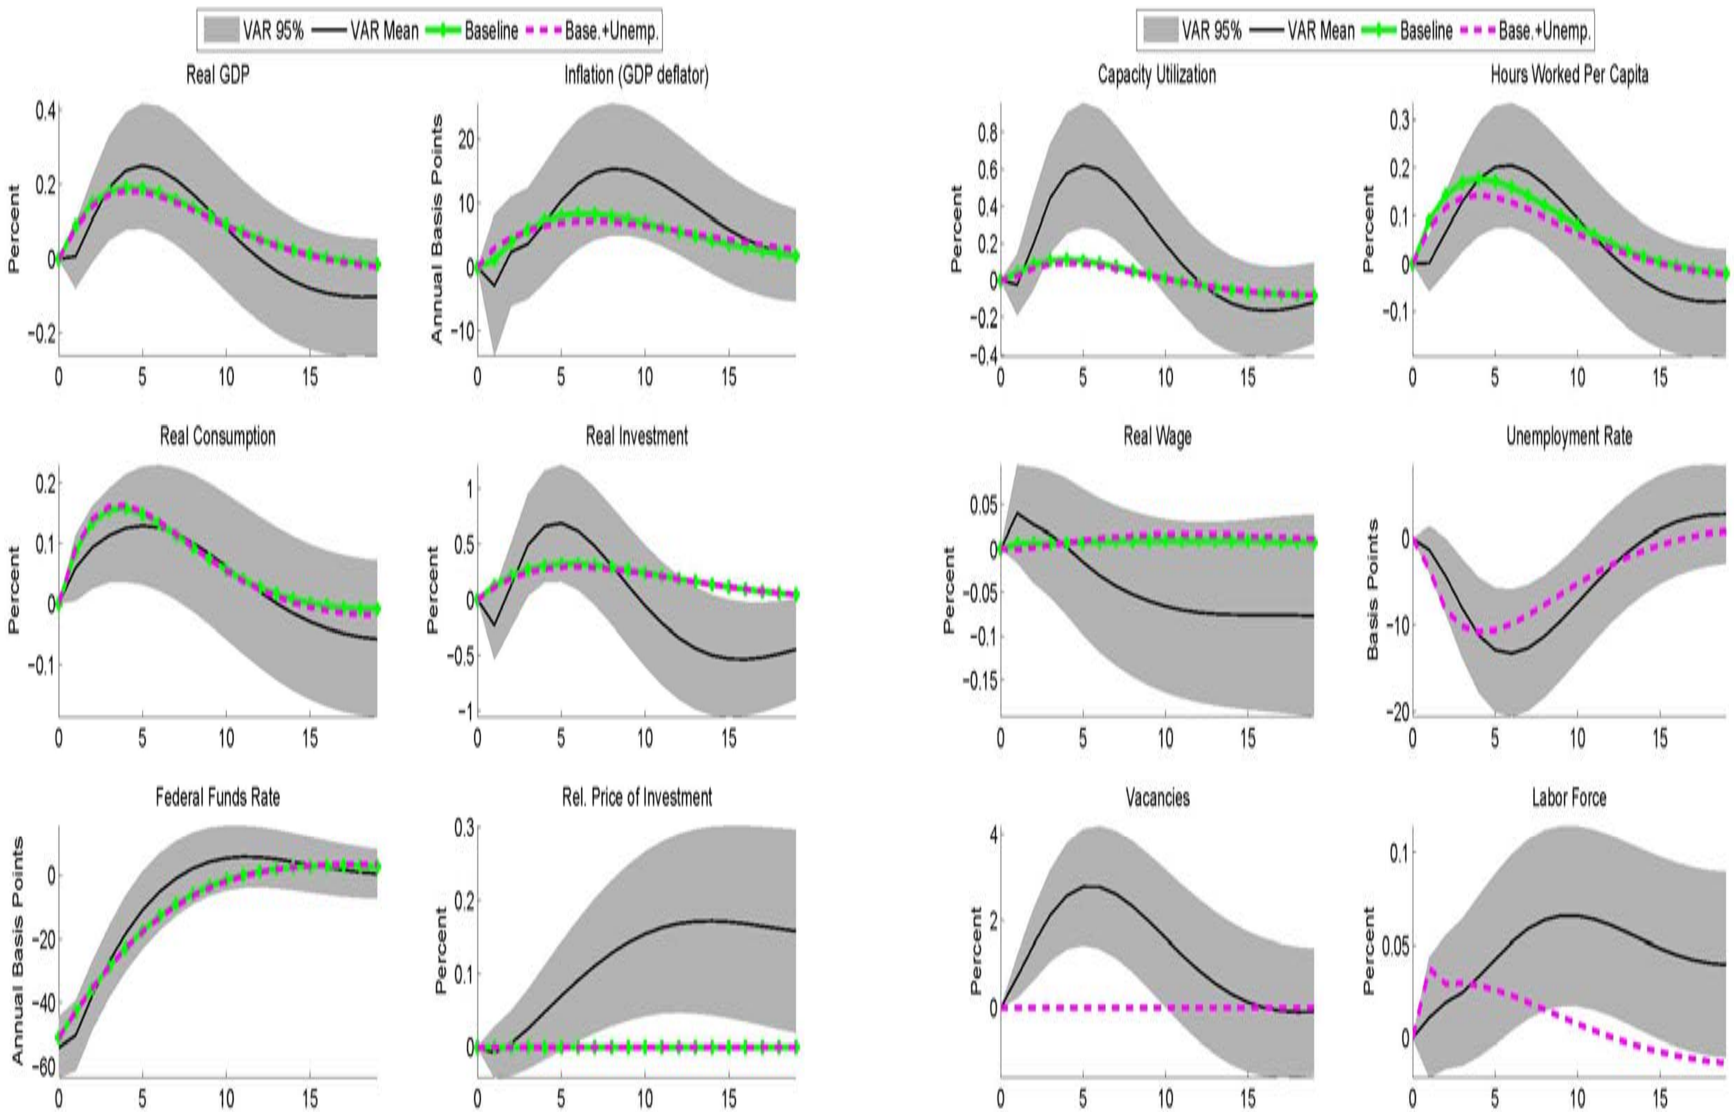

Figure 7: Response to Neutral Shock, Baseline W/High Inflation Weight

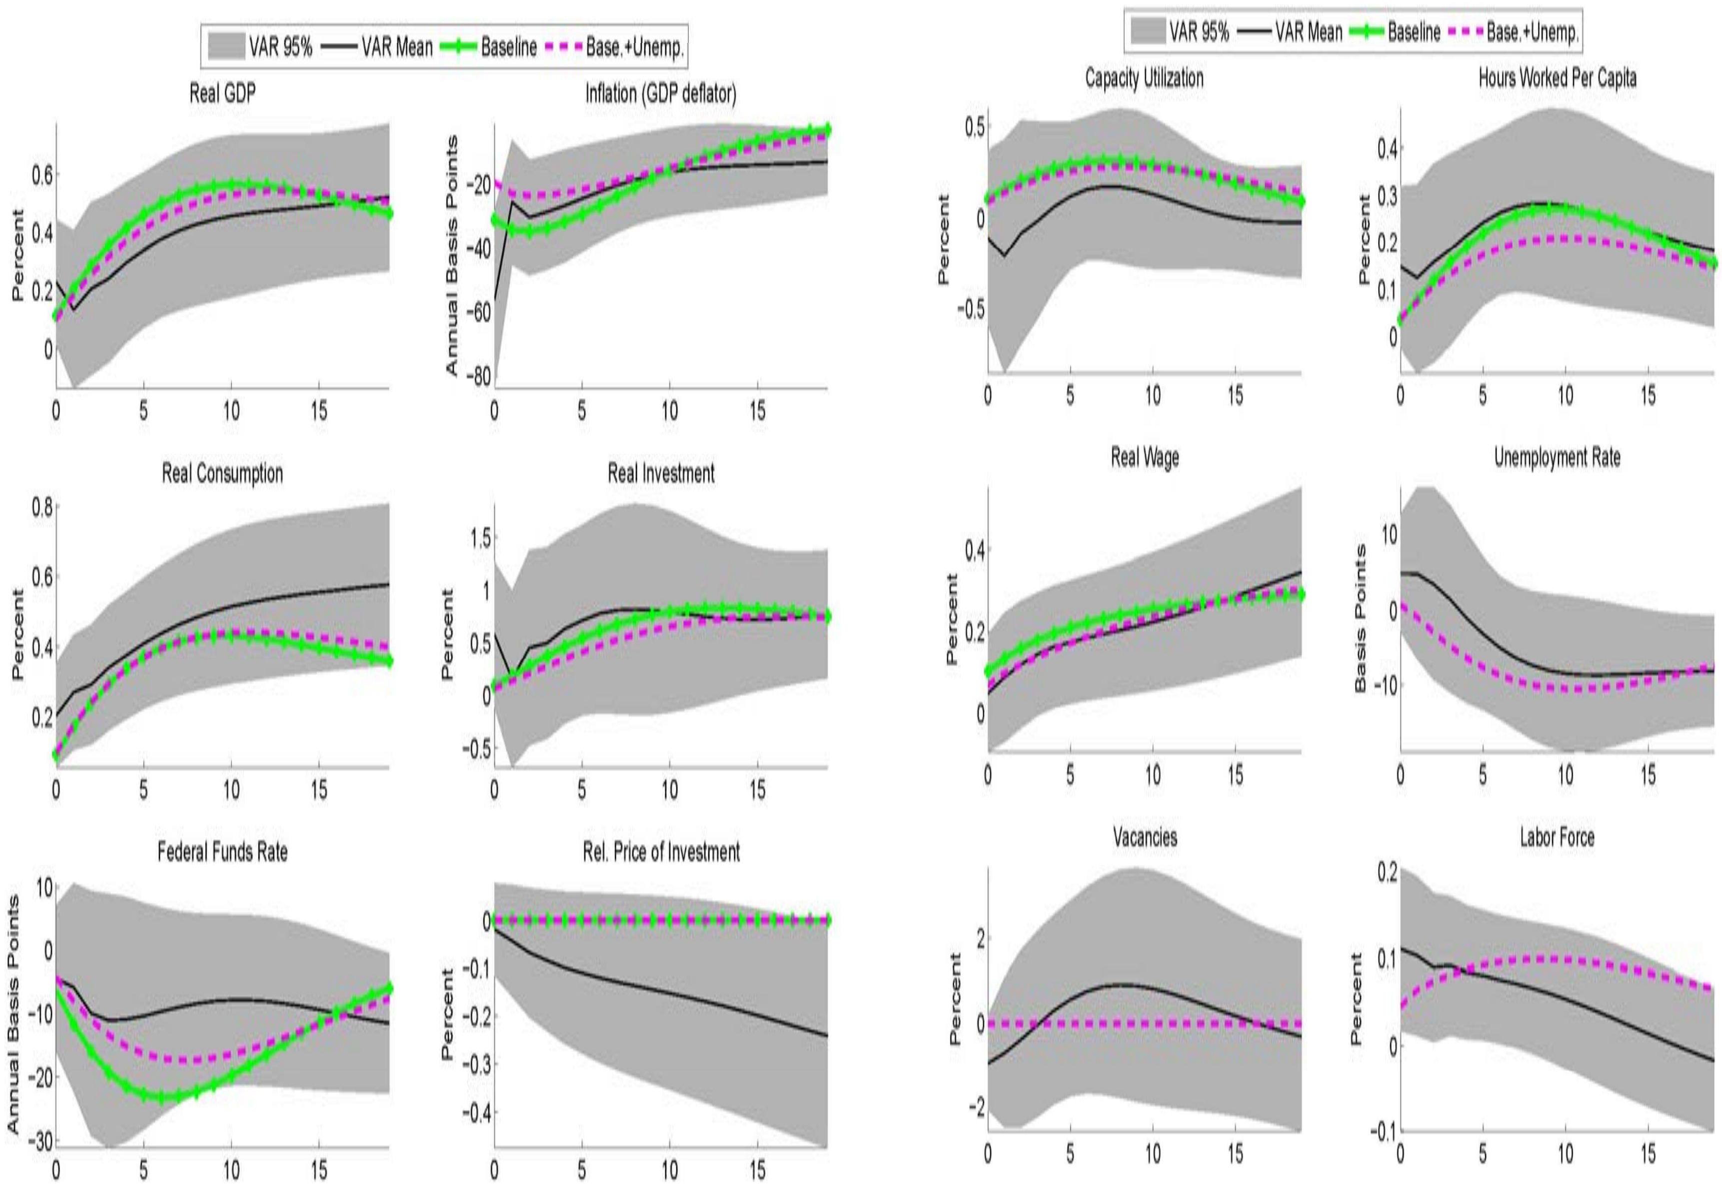

Figure 8: Response to Investment Shock, Baseline W/High Inflation Weight

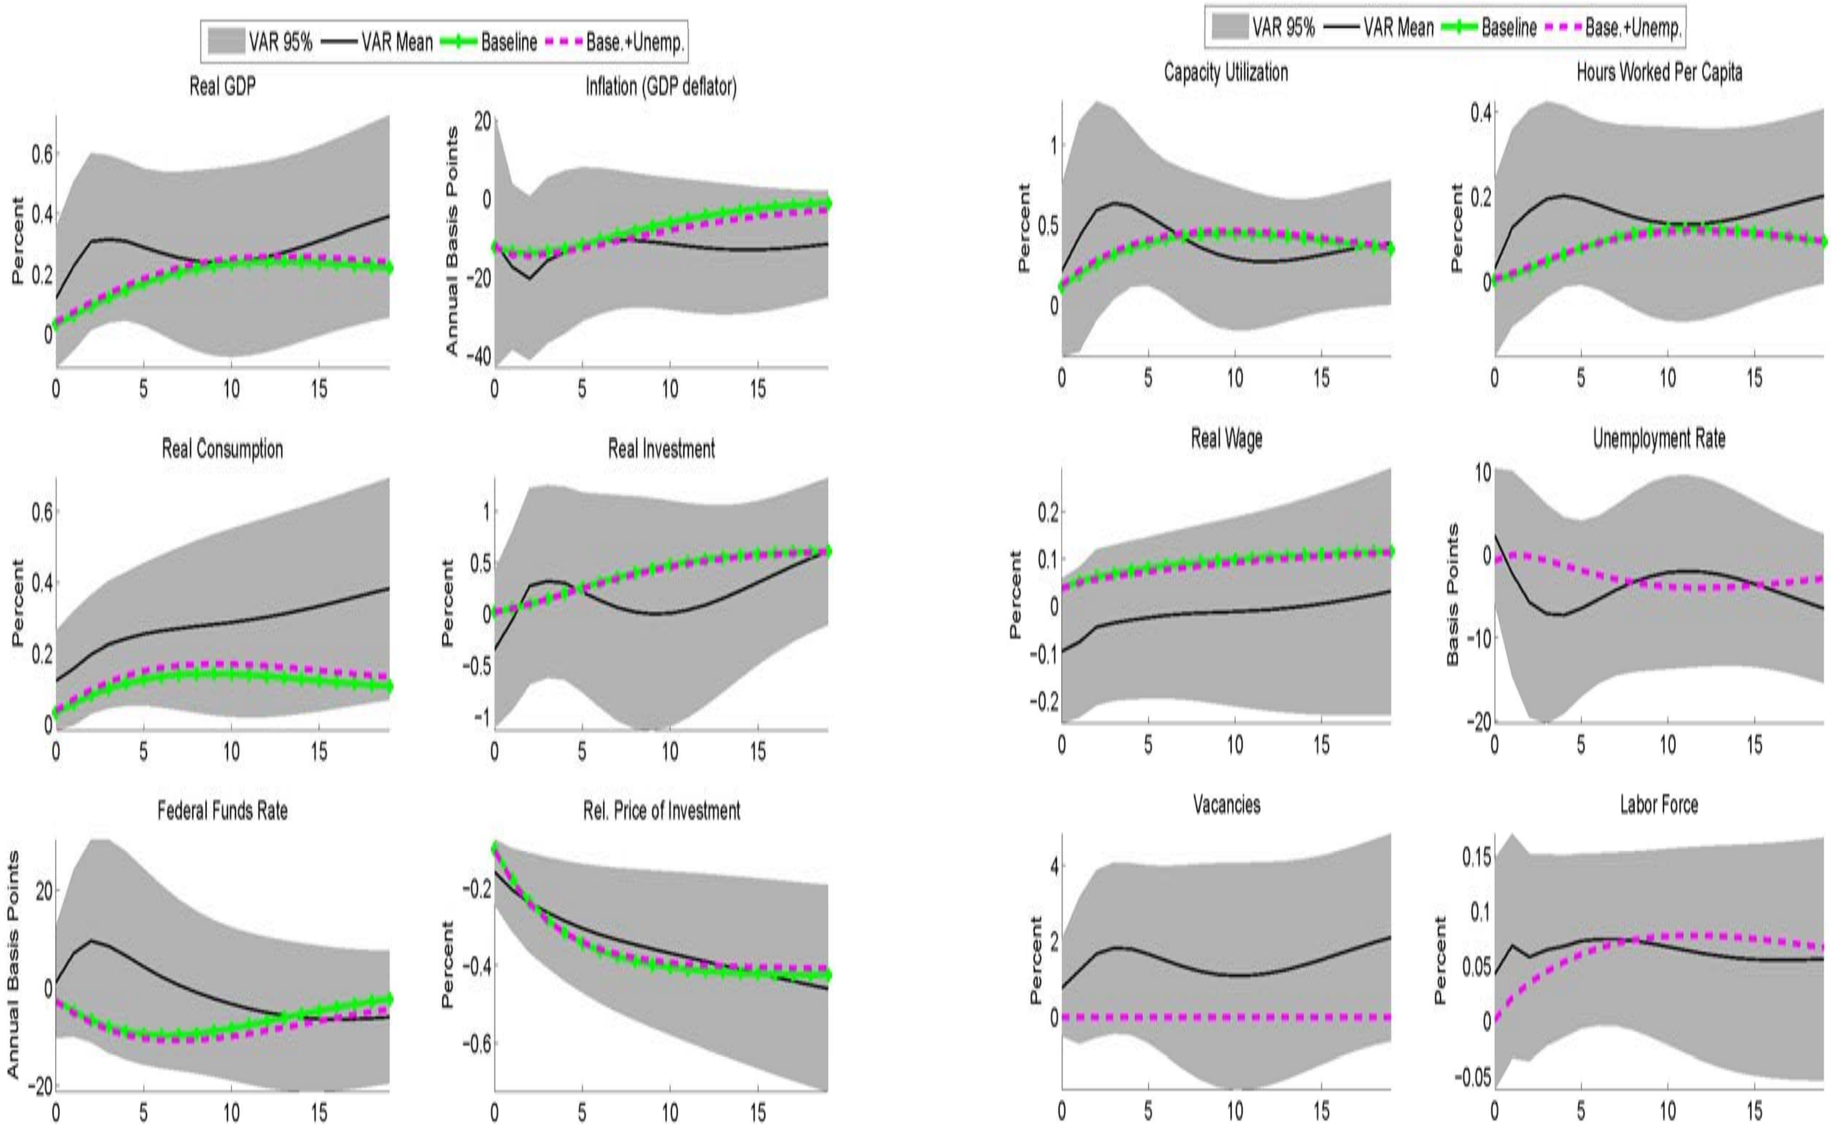

Figure 9: response to monetary shock, Search and Matching

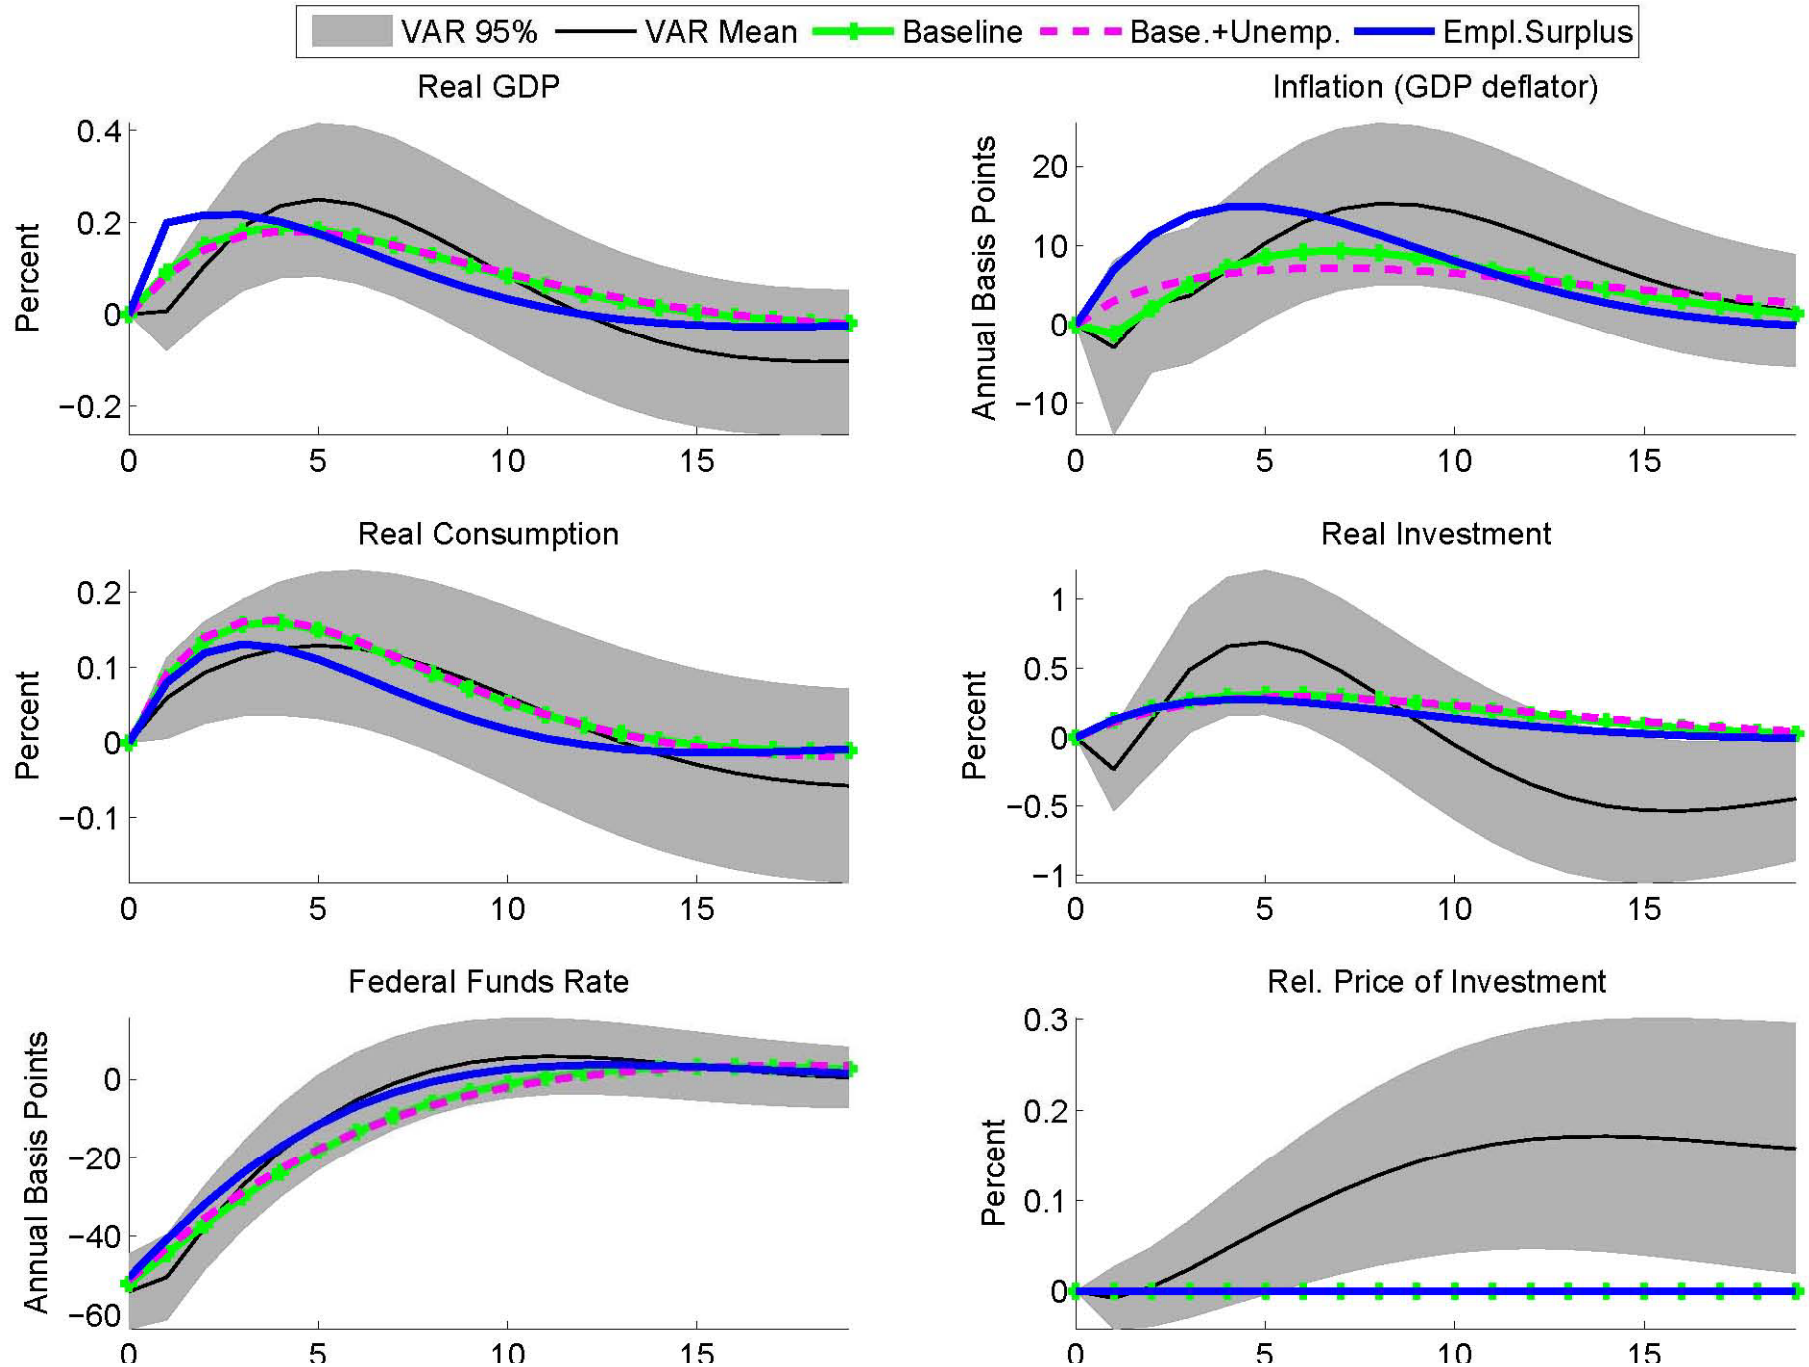

Figure 9, continued

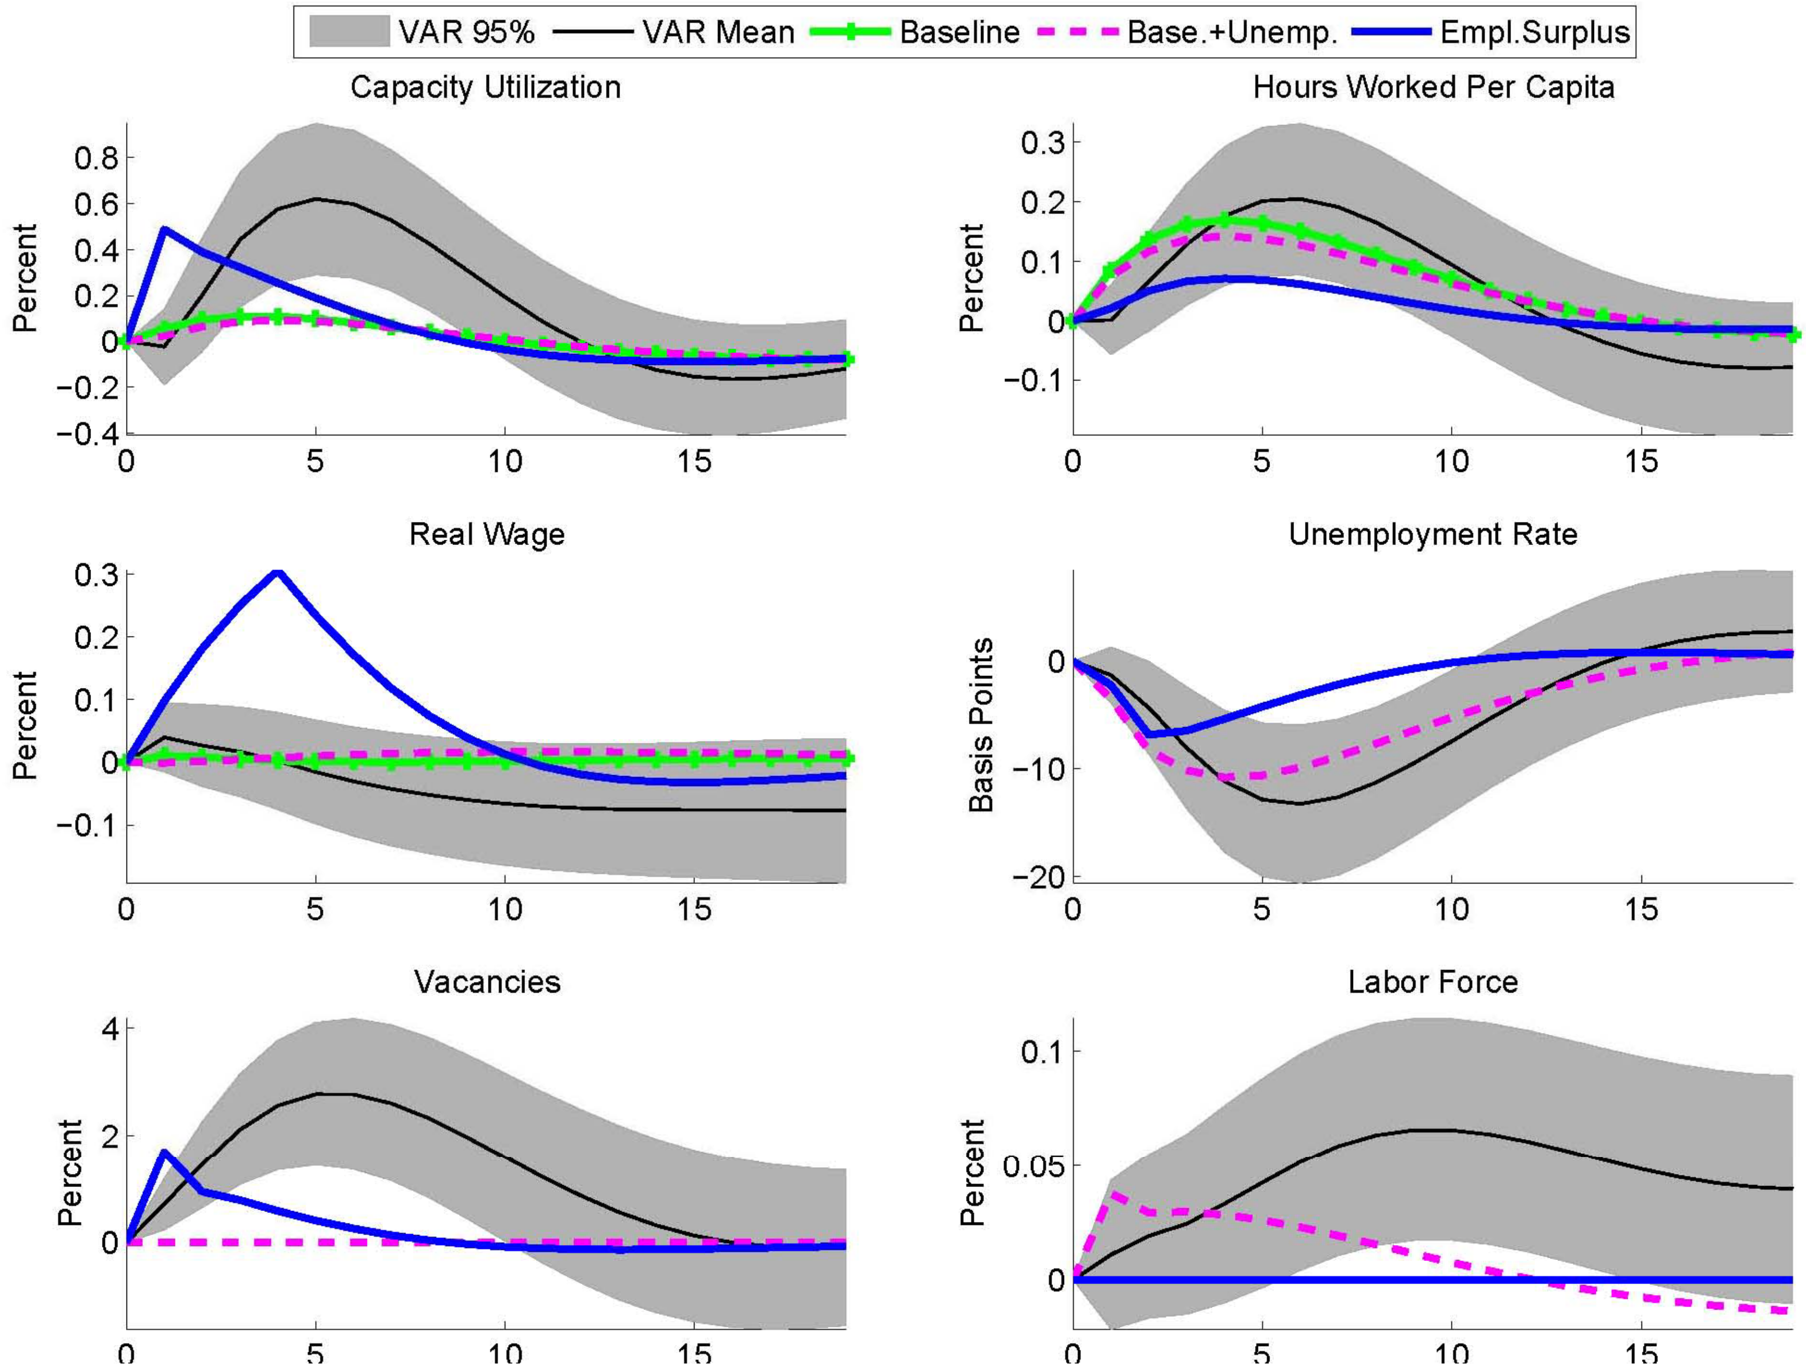

Figure 10: Response to Neutral Shock, Search and Matching Model

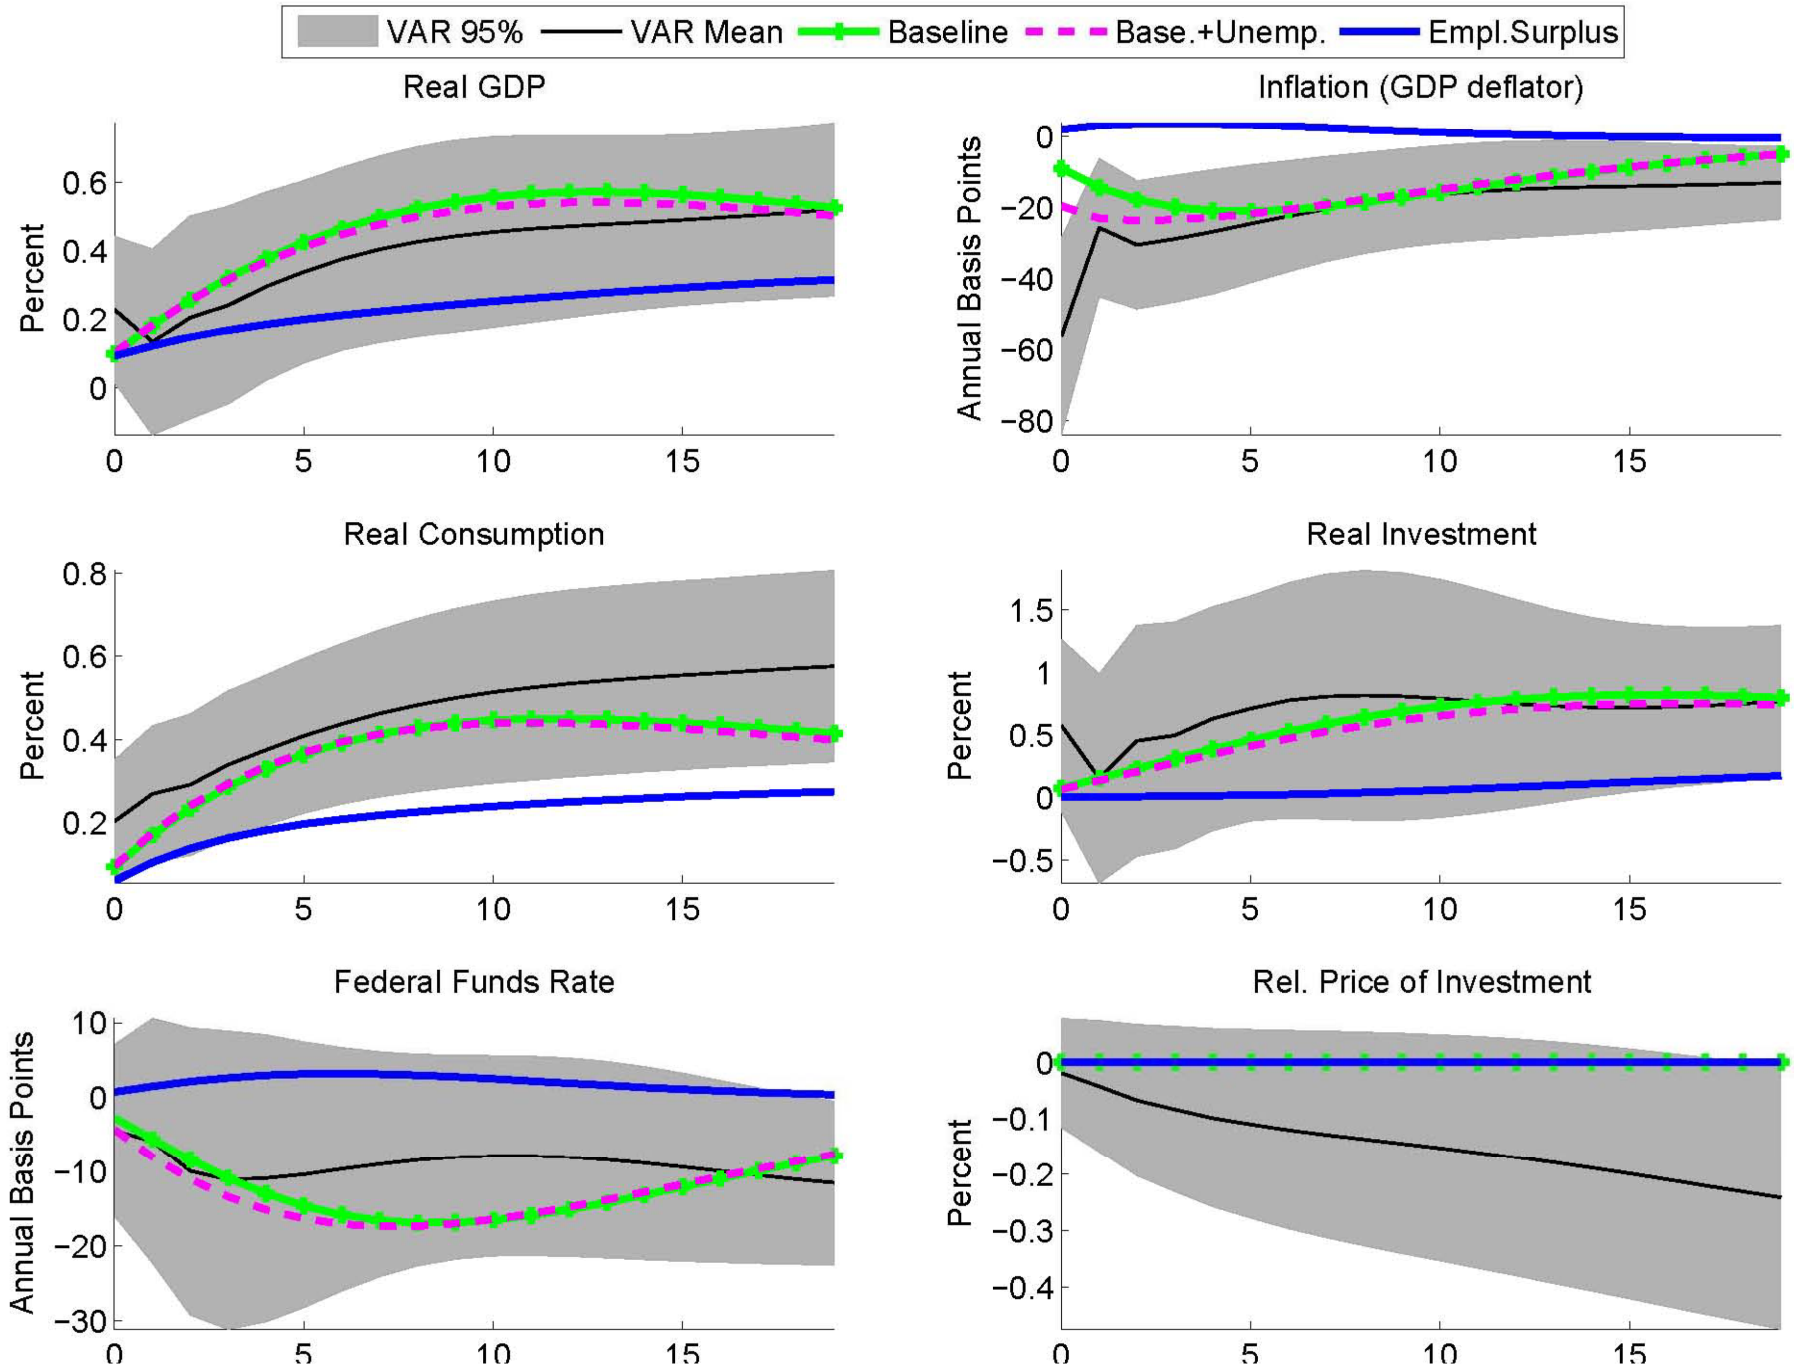

Figure 10, continued

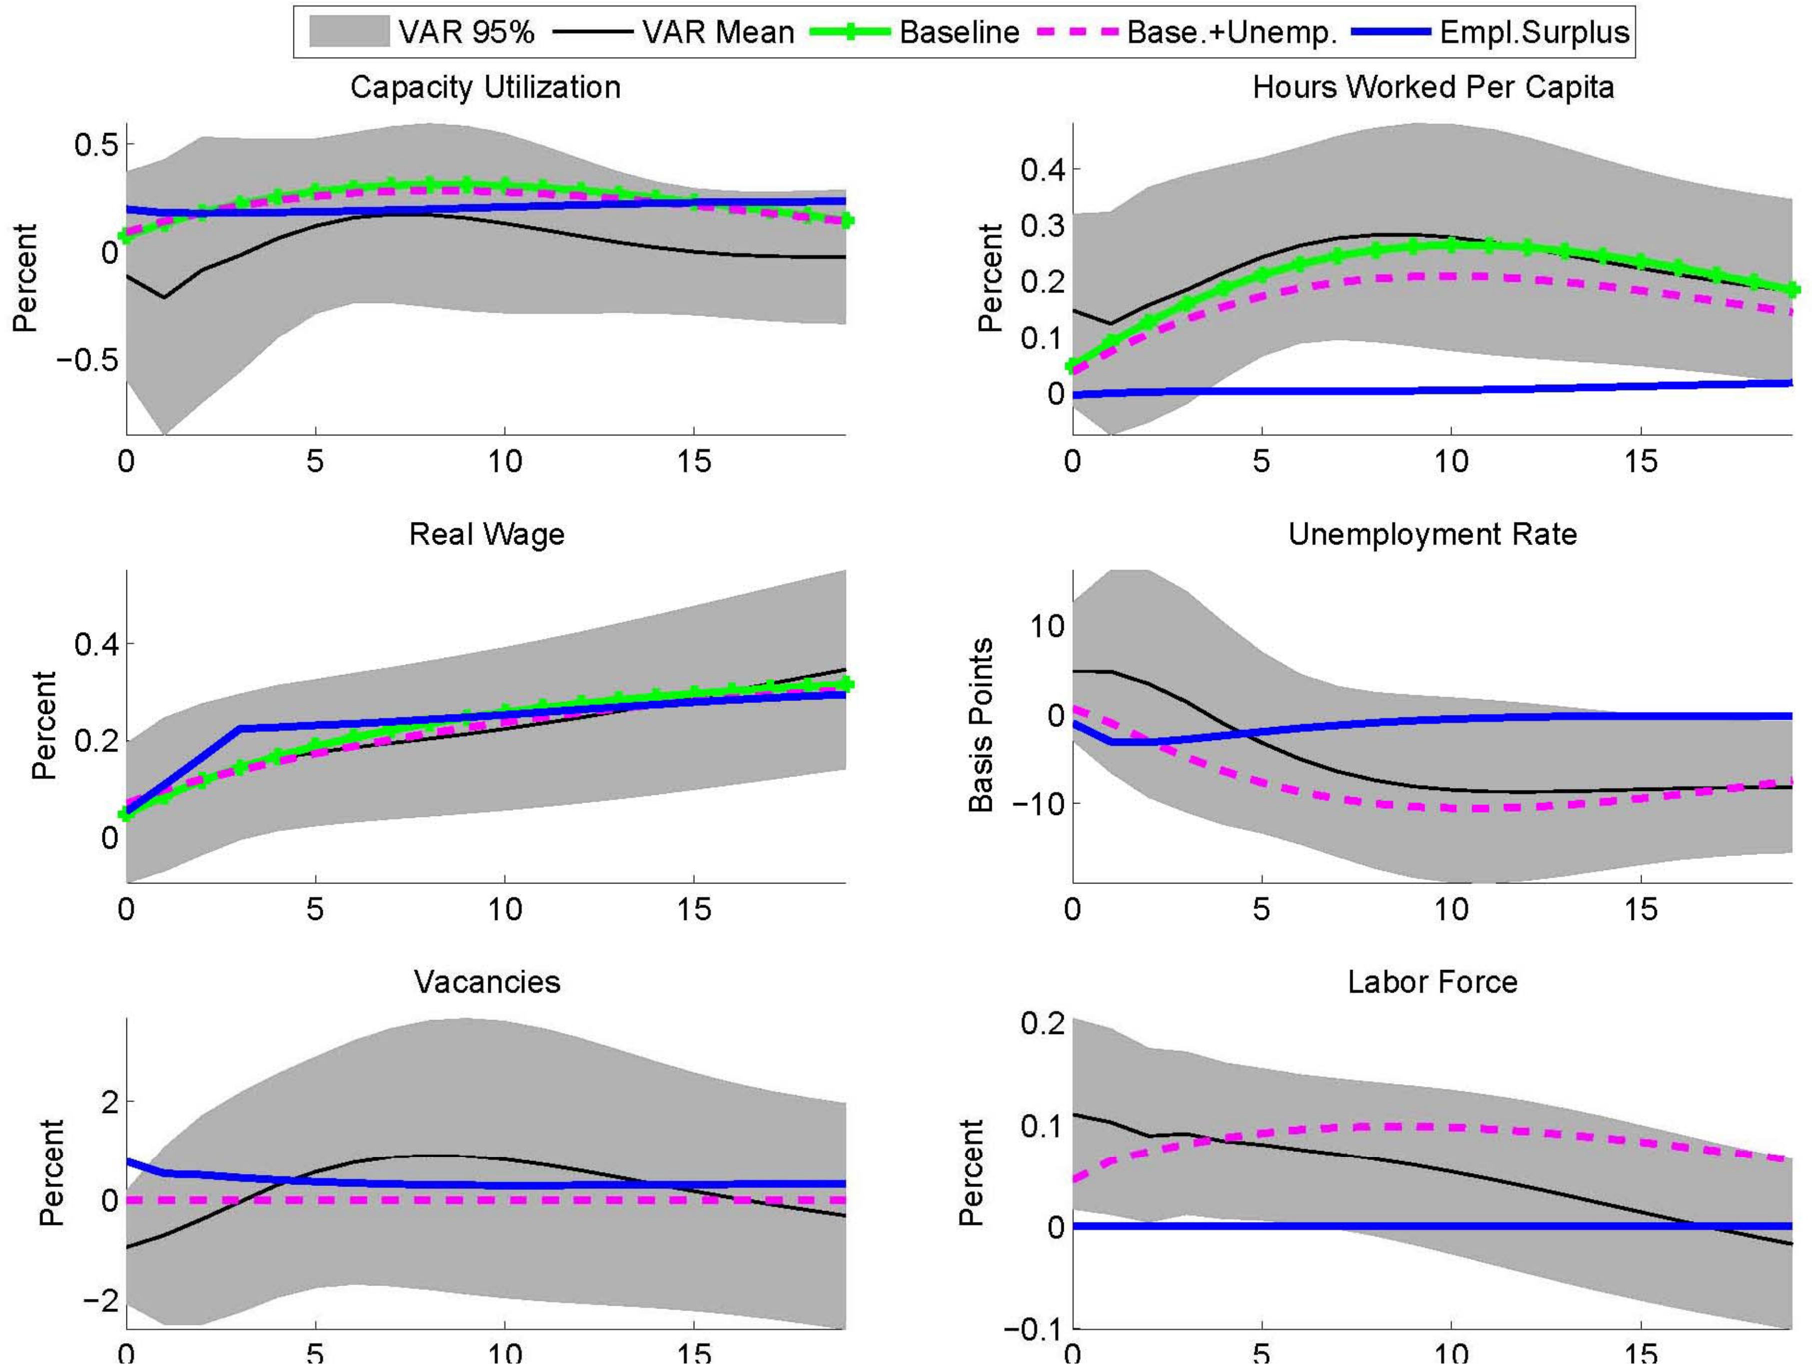

Figure 11, Response to Investment Shock, Search and Matching

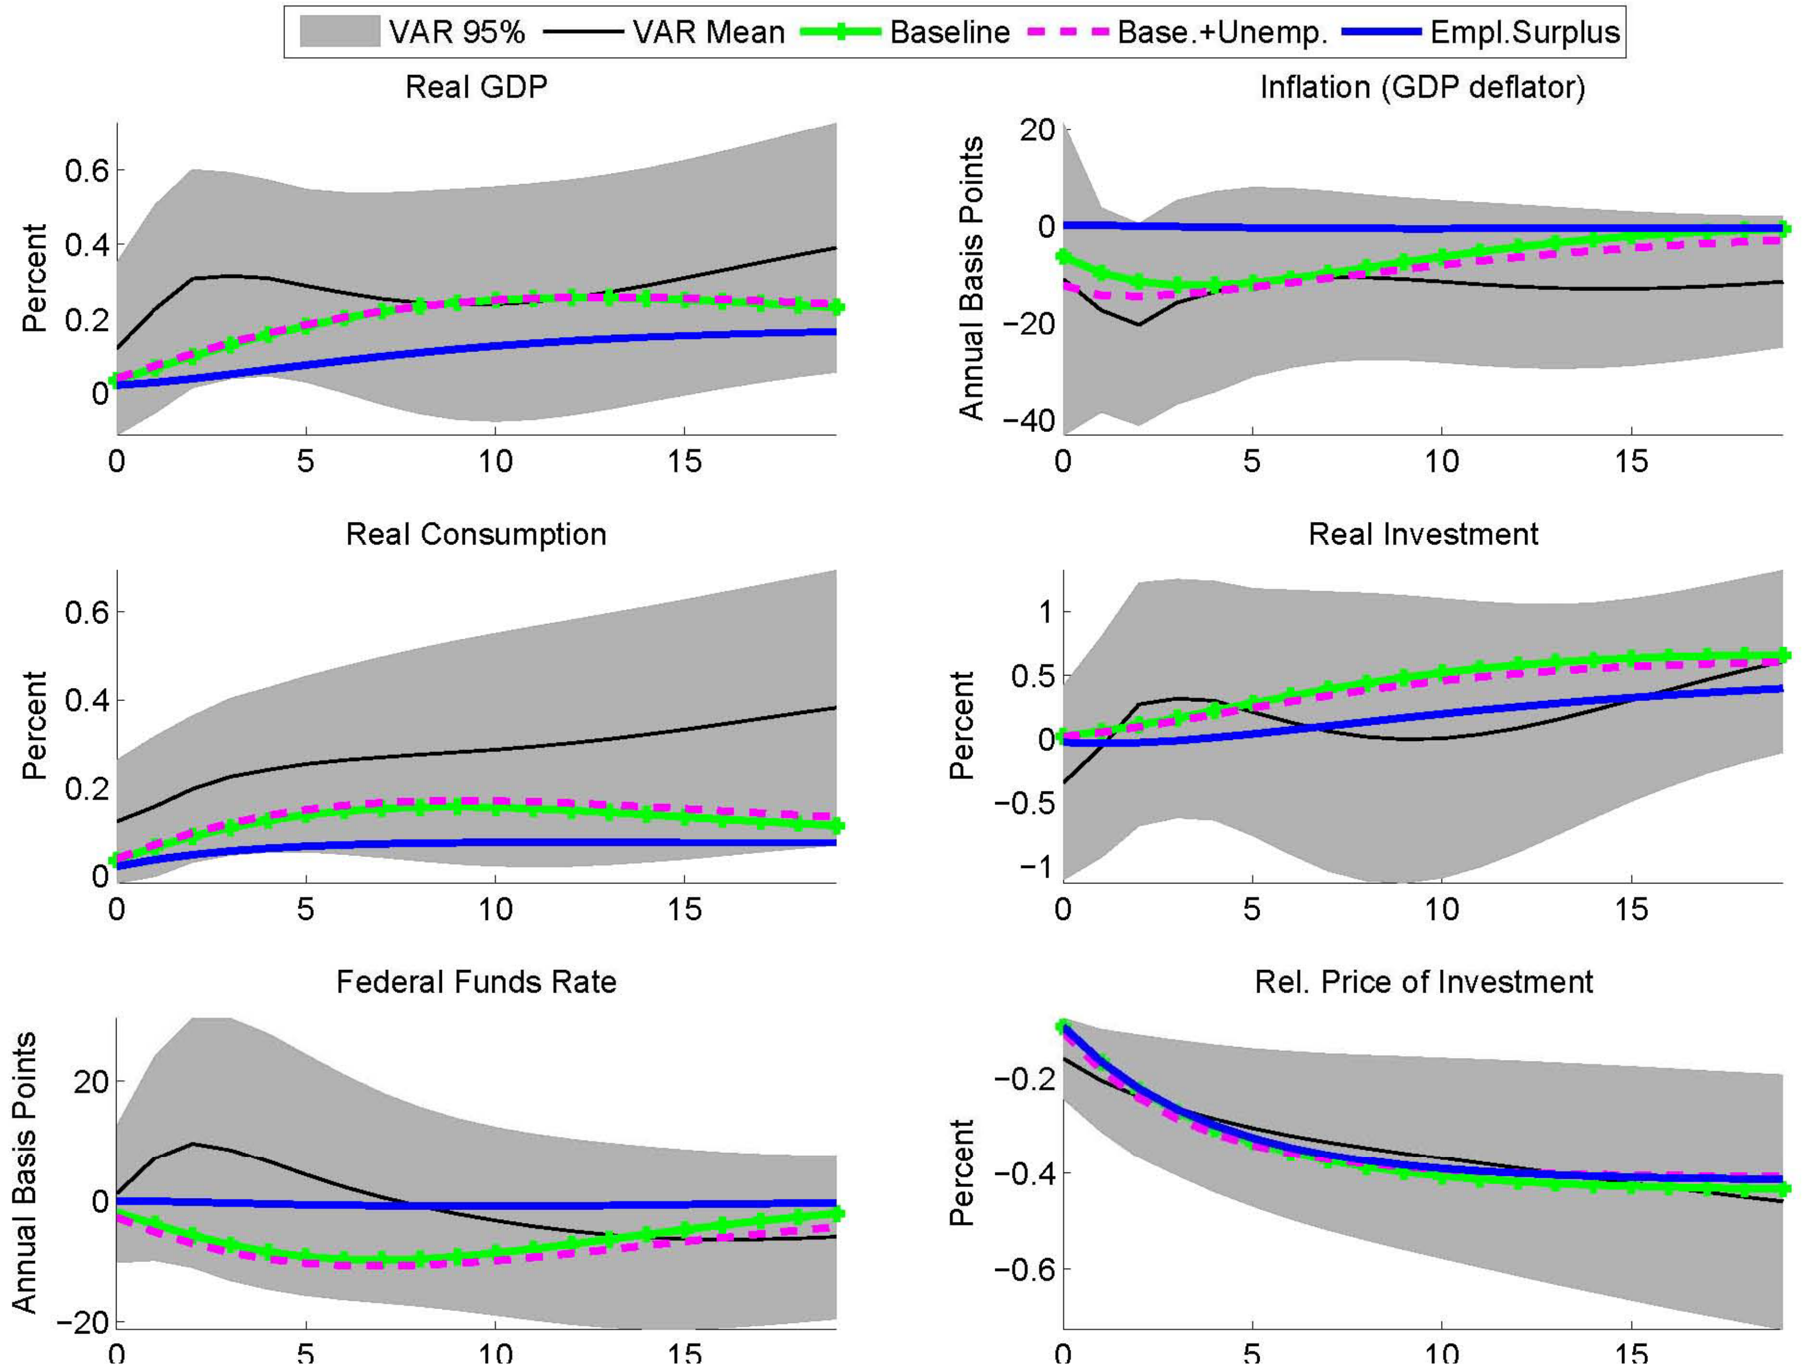

Figure 11, continued.