

Understanding the Great Recession*

Lawrence J. Christiano[†] Martin S. Eichenbaum[‡] Mathias Trabandt[§]

April 2, 2014

Abstract

We argue that the vast bulk of movements in aggregate real economic activity during the Great Recession were due to financial frictions interacting with the zero lower bound. We reach this conclusion looking through the lens of a New Keynesian model in which firms face moderate degrees of price rigidities and no nominal rigidities in the wage setting process. Our model does a good job of accounting for the joint behavior of labor and goods markets, as well as inflation, during the Great Recession. According to the model the observed fall in total factor productivity and the rise in the cost of working capital played critical roles in accounting for the small size of the drop in inflation that occurred during the Great Recession.

*The views expressed in this paper are those of the authors and do not necessarily reflect those of the Board of Governors of the Federal Reserve System or of any other person associated with the Federal Reserve System. We are grateful for discussions with Gadi Barlevy.

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

[‡]Northwestern University, Department of Economics, 2001 Sheridan Road, Evanston, Illinois 60208, USA. Phone: +1-847-491-8232. E-mail: eich@northwestern.edu.

[§]Board of Governors of the Federal Reserve System, Division of International Finance, Trade and Financial Studies Section, 20th Street and Constitution Avenue N.W., Washington, D.C. 20551, USA, E-mail: mathias.trabandt@gmail.com.

1. Introduction

The Great Recession has been marked by extraordinary contractions in output, investment and consumption. Mirroring these developments, per capita employment and the labor force participation rate have dropped substantially and show little sign of improving. The unemployment rate has declined from its Great Recession peak. But, this decline primarily reflects a sharp drop in the labor force participation rate, not an improvement in the labor market. Indeed, while vacancies have risen to their pre-recession levels, this rise has not translated into an improvement in employment. Despite all this economic weakness, the decline in inflation has been relatively modest.

We seek to understand the key forces driving the US economy in the Great Recession. To do so, we require a model that provides an empirically plausible account of key macroeconomic aggregates, including labor market outcomes like employment, vacancies, the labor force participation rate and the unemployment rate. To this end, we extend the medium-sized dynamic, stochastic general equilibrium (DSGE) model in Christiano, Eichenbaum and Trabandt (2013) (CET) to endogenize the labor force participation rate. To establish the empirical credibility of our model, we estimate its parameters using pre-2008 data. We argue that the model does a good job of accounting for the dynamics of twelve key macroeconomic variables over this period.

We show that four shocks can account for the key features of the Great Recession. Two of these shocks capture in a reduced form way frictions which are widely viewed as having played an important role in the Great Recession. The first of these is motivated by the literature stressing a reduction in consumption as a trigger for a zero lower bound (ZLB) episode (see Eggertsson and Woodford (2003), Eggertsson and Krugman (2012) and Guerrieri and Lorenzoni (2012)). For convenience, we capture this idea as in Smets and Wouters (2007) and Fisher (2014), by introducing a perturbation to agents' intertemporal Euler equation governing the accumulation of the risk-free asset. We refer to this perturbation as the *consumption wedge*. The second friction shock is motivated by the sharp increase in credit spreads observed in the post-2008 period. To capture this phenomenon, we introduce a wedge into households' first order condition for optimal capital accumulation. Simple financial friction models based on asymmetric information with costly monitoring imply that credit market frictions can be captured in a reduced form way as a wedge in the household's first order condition for capital (see Christiano and Davis 2006). We refer to this wedge as the *financial wedge*. Also, motivated by models like e.g. Bigio (2013), we allow the financial wedge to impact on the cost of working capital.

The third shock in our analysis is a neutral technology shock that captures the observed decline, relative to trend, in total factor productivity (TFP). The final shock in our analy-

sis corresponds to the changes in government consumption that occurred during the Great Recession.

Our main findings can be summarized as follows. First, our model can account, quantitatively, for the key features of the Great Recession, including the ongoing decline in the labor force participation rate. Second, according to our model the vast bulk of the decline in economic activity is due to the financial wedge and, to a somewhat smaller extent, the consumption wedge.¹ The rise in government consumption associated with the American Recovery and Reinvestment Act of 2009 did have a peak multiplier effect in excess of 2. But, the rise in government spending was too small to have a substantial effect. In addition, for reasons discussed in the text, we cannot attribute the long duration of the Great Recession to the substantial decline in government consumption that began around the start of 2011. Third, consistent with the basic findings in CET, we are able to account for the observed behavior of real wages during the Great Recession, even though we do not allow for sticky wages. Fourth, our model can account for the relatively small decline in inflation with only a moderate amount of price stickiness.

Our last finding is perhaps surprising in light of arguments by Hall (2011) and others that New Keynesian (NK) models imply inflation should have been much lower than it was during the Great Recession.² Del Negro et al. (2014) argue that Hall's conclusions do not hold if the Phillips curve is sufficiently flat.³ In contrast, our model accounts for the behavior of inflation after 2008 by incorporating two key features of the data into our analysis: (i) the prolonged slowdown in TFP growth during the Great Recession and (ii) the rise in the cost of firms' working capital as measured by the spread between the corporate-borrowing rate and the risk-free interest rate. In our model, these forces drive up firms' marginal costs, exerting countervailing pressures on the deflationary forces operative during the post 2008 period.

Our paper may be of independent interest from a methodological perspective for three reasons. First, our analysis of the Great Recession requires that we do stochastic simulations of a model that is highly non-linear in several respects: (i) we work with the actual nonlinear equilibrium conditions; (ii) we confront the fact that the ZLB on the nominal interest rate is binding in parts of the sample and not in others; and (iii) our characterization of monetary policy allows for forward guidance, a policy rule that is characterized by regime switches in response to the values taken on by endogenous variables. The one approximation that we use in our solution method is certainty equivalence. Second, as we explain below, our analysis of

¹The findings with respect to the financial wedge is consistent Del Negro, Giannoni and Schorfheide (2014), who reach their conclusion using a different methodology than ours.

²In a related criticism Dupor and Li (2013) argue that the behavior of actual and expected inflation during the period of the American Recovery and Reinvestment Act is inconsistent with the predictions of NK style models.

³Christiano, Eichenbaum and Rebelo (2011) reach a similar conclusion based on data up to the end of 2010.

the Great Recession requires that we adopt an unobserved components representation for the growth rate of neutral technology. This leads to a series of challenges in solving the model and deriving its implications for the data. Third, we note that traditional analyses of vacancies and unemployment based on the Beveridge curve would infer that there was a deterioration in the efficiency of labor markets during the Great Recession. We argue that this conclusion is based on a technical assumption which is highly misleading when applied to data from the Great Recession.

The remainder of this paper is organized as follows. The next section describes our model. The following two sections describe the data, methodology and results for estimating our model on pre-2008 data. In the next two sections, we use our model to study the Great Recession. We close with a brief conclusion. Many technical details of our analysis are relegated to a separate technical appendix that is available on request.

2. The Model

In this section, we describe a medium-sized DSGE model whose structure is, with one important exception, the same as the one in CET. The exception is that we modify the framework to endogenize labor force participation rates.

2.1. Households and Labor Force Dynamics

The economy is populated by a large number of identical households. Each household has a unit measure of members. Members of the household can be engaged in three types of activities: (i) $(1 - L_t)$ members specialize in home production in which case we say they are not in the labor force and that they are in the non-participation state; (ii) l_t members of the household are in the labor force and are employed in the production of a market good, and (iii) $(L_t - l_t)$ members of the household are unemployed, i.e. they are in the labor force but do not have a job.

We now describe aggregate flows in the labor market. We derive an expression for the total number of people searching for a job at the end of a period. This allows us to define the job finding rate, f_t , and the rate, e_t , at which workers transit from non-participation into labor force participation.

At the end of each period a fraction $1 - \rho$ of randomly selected employed workers is separated from the firm with which they had been matched. Thus, at the end of period $t - 1$ a total of $(1 - \rho)l_{t-1}$ workers separate from firms and ρl_{t-1} workers remain attached to their firm. Let u_{t-1} denote the unemployment rate at time $t - 1$, so that the number of unemployed

workers at time $t - 1$ is $u_{t-1}L_{t-1}$. The sum of separated and unemployed workers is given by:

$$\begin{aligned} (1 - \rho)l_{t-1} + u_{t-1}L_{t-1} &= (1 - \rho)l_{t-1} + \frac{L_{t-1} - l_{t-1}}{L_{t-1}}L_{t-1} \\ &= L_{t-1} - \rho l_{t-1}. \end{aligned}$$

We assume that a separated worker and an unemployed worker have an equal probability, $1 - s$, of exiting the labor force. It follows that s times the number of separated and unemployed workers, $s(L_{t-1} - \rho l_{t-1})$, remain in the labor force and search for work. We refer to s as the ‘staying rate’.

The household chooses r_t , the number of workers that it transfers from non-participation into the labor force. Thus, the labor force in period t is:

$$L_t = s(L_{t-1} - \rho l_{t-1}) + \rho l_{t-1} + r_t. \quad (2.1)$$

By its choice of r_t the household in effect chooses L_t . The total number of workers searching for a job at the start of t is:

$$s(L_{t-1} - \rho l_{t-1}) + r_t = L_t - \rho l_{t-1}. \quad (2.2)$$

Here we have used (2.1) to substitute out for r_t on the left hand side of (2.2).

It is of interest to calculate the probability, e_t , that a non-participating worker is selected to be in the labor force. We assume that the $(1 - s)(L_{t-1} - \rho l_{t-1})$ workers who separate exogenously into the non-participation state do not return home in time to be included in the pool of workers relevant to the household’s choice of r_t . As a result, the universe of workers from which the household selects r_t is $1 - L_{t-1}$. It follows that e_t is given by:⁴

$$e_t = \frac{r_t}{1 - L_{t-1}} = \frac{L_t - s(L_{t-1} - \rho l_{t-1}) - \rho l_{t-1}}{1 - L_{t-1}}. \quad (2.3)$$

The law of motion for employment is:

$$l_t = (\rho + x_t)l_{t-1} = \rho l_{t-1} + x_t l_{t-1}. \quad (2.4)$$

⁴We include the staying rate, s , in our analysis for a substantive as well as a technical reason. The substantive reason is that, in the data, workers move in both directions between unemployment, non-participation and employment. The gross flows are much bigger than the net flows. Setting $s < 1$ helps the model account for these patterns. The technical reason for allowing $s < 1$ can be seen by setting $s = 1$ in (2.3). In that case, if the household wishes to make $L_t - L_{t-1} < 0$, it must set $e_t < 0$. That would require withdrawing from the labor force some workers who were unemployed in $t - 1$ and stayed in the labor force as well as some workers who were separated from their firm and stayed in the labor force. But, if some of these workers are withdrawn from the labor force then their actual staying rate would be lower than the fixed number, s . So, the actual staying rate would be a non-linear function of $L_t - L_{t-1}$ with the staying rate below s for $L_t - L_{t-1} < 0$ and equal to s for $L_t - L_{t-1} \geq 0$. This kink point is a non-linearity that would be hard to avoid because it occurs precisely at the model’s steady state. Even with $s < 1$ there is a kink point, but it is far from steady state and so it can be ignored when we solve the model.

The job finding rate is the ratio of the number of new hires divided by the number of people searching for work, given by (2.2):

$$f_t = \frac{x_t l_{t-1}}{L_t - \rho l_{t-1}}. \quad (2.5)$$

2.2. Household Maximization

Members of the household derive utility from a market consumption good and a good produced at home.⁵ The home good is produced using the labor of individuals that are not in the labor force, $1 - L_t$, and the labor of the unemployed, $L_t - l_t$:

$$C_t^H = \eta_t^H (1 - L_t)^{1-\alpha_c} (L_t - l_t)^{\alpha_c} - \mathcal{F}(L_t, L_{t-1}; \eta_t^L). \quad (2.6)$$

The term $\mathcal{F}(L_t, L_{t-1}; \eta_t^L)$ captures the idea that it is costly to change the number of people who specialize in home production,

$$\mathcal{F}(L_t, L_{t-1}; \eta_t^L) = 0.5 \eta_t^L \phi_L (L_t/L_{t-1} - 1)^2 L_t. \quad (2.7)$$

We assume $\alpha_c < 1 - \alpha_c$, so that in steady state the unemployed contribute less to home production than do people who are out of the labor force. Finally, η_t^H and η_t^L are processes that ensure balanced growth. We discuss these processes in detail below. We included the adjustment costs in L_t so that the model can account for the gradual and hump-shaped response of the labor force to a monetary policy shock (see subsection 4.3).

Workers experience no disutility from working and supply their labor inelastically. An employed worker brings home the wages that he earns. Unemployed workers receive government-provided unemployment compensation which they give to the household. Unemployment benefits are financed by lump-sum taxes paid by the household. The details of how workers find employment and receive wages are explained below. All household members have the same concave preferences over consumption, so each is allocated the same level of consumption.

The representative household maximizes the objective function:

$$E_0 \sum_{t=0}^{\infty} \beta^t \ln(\tilde{C}_t), \quad (2.8)$$

where

$$\tilde{C}_t = [(1 - \omega) (C_t - b\bar{C}_{t-1})^\chi + \omega (C_t^H - b\bar{C}_{t-1}^H)^\chi]^\frac{1}{\chi}.$$

⁵Erceg and Levin (2013) also exploit this type of tradeoff in their model of labor force participation. However, their households find themselves in a very different labor market than ours do. In our analysis the labor market is a version of the Diamond-Mortensen-Pissarides model, while in their analysis, the labor market is a competitive spot market.

Here, C_t and C_t^H denote market consumption and consumption of the good produced at home. The elasticity of substitution between C_t and C_t^H is $1/(1 - \chi)$ in steady state. The parameter b controls the degree of habit formation in household preferences. We assume $0 \leq b < 1$. A bar over a variable indicates its economy-wide average value.

The flow budget constraint of the household is as follows:

$$\begin{aligned} & P_t C_t + P_{I,t} I_t + B_{t+1} \\ \leq & (R_{K,t} u_t^K - a(u_t^K) P_{I,t}) K_t + (L_t - l_t) P_t \eta_t^D D_t + l_t W_t + R_{t-1} B_t - T_t. \end{aligned} \quad (2.9)$$

The variable T_t denotes lump-sum taxes net of transfers and firm profits, B_{t+1} denotes beginning-of-period t purchases of a nominal bond which pays rate of return R_t at the start of period $t + 1$, and $R_{K,t}$ denotes the nominal rental rate of capital services. The variable u_t^K denotes the utilization rate of capital. As in Christiano, Eichenbaum and Evans (2005) (CEE), we assume that the household sells capital services in a perfectly competitive market, so that $R_{K,t} u_t^K K_t$ represents the household's earnings from supplying capital services. The increasing convex function $a(u_t^K)$ denotes the cost, in units of investment goods, of setting the utilization rate to u_t^K . The variable $P_{I,t}$ denotes the nominal price of an investment good and I_t denotes household purchases of investment goods. In addition, the nominal wage rate earned by an employed worker is W_t and $\eta_t^D D_t$ denotes exogenous unemployment benefits received by unemployed workers from the government. The term η_t^D is a process that ensures balanced growth and will be discussed below.

When the household chooses L_t it takes the aggregate job finding rate, f_t , and the law of motion linking L_t and l_t as given:

$$l_t = \rho l_{t-1} + f_t (L_t - \rho l_{t-1}). \quad (2.10)$$

Relation (2.10) is consistent with the actual law of motion of employment because of the definition of f_t (see (2.5)).

The household owns the stock of capital which evolves according to,

$$K_{t+1} = (1 - \delta_K) K_t + [1 - S(I_t/I_{t-1})] I_t. \quad (2.11)$$

The function $S(\cdot)$ is an increasing and convex function capturing adjustment costs in investment. We assume that $S(\cdot)$ and its first derivative are both zero along a steady state growth path.

The household chooses state-contingent sequences, $\{C_t^H, L_t, l_t, C_t, B_{t+1}, I_t, u_t^K, K_{t+1}\}_{t=0}^\infty$, to maximize utility, (2.8), subject to, (2.6), (2.7), (2.9), (2.10) and (2.11). The household takes $\{K_0, B_0, l_{-1}\}$ and the state and date-contingent sequences, $\{R_t, W_t, P_t, R_{K,t}, P_{I,t}\}_{t=0}^\infty$, as given. As in CEE, we assume that the $C_t^H, L_t, l_t, C_t, I_t, u_t^K, K_{t+1}$ decisions are made

before the realization of the current period monetary policy shock and after the realization of the other shocks. This assumption captures the notion that monetary policy shocks occur at a higher frequency of time than the other shocks discussed below.

2.3. Final Good Producers

A final homogeneous market good, Y_t , is produced by competitive and identical firms using the following technology:

$$Y_t = \left[\int_0^1 (Y_{j,t})^{\frac{1}{\lambda}} dj \right]^\lambda, \quad (2.12)$$

where $\lambda > 1$. The representative firm chooses specialized inputs, $Y_{j,t}$, to maximize profits:

$$P_t Y_t - \int_0^1 P_{j,t} Y_{j,t} dj,$$

subject to the production function (2.12). The firm's first order condition for the j^{th} input is:

$$Y_{j,t} = \left(\frac{P_t}{P_{j,t}} \right)^{\frac{\lambda}{\lambda-1}} Y_t. \quad (2.13)$$

2.4. Retailers

As in Ravenna and Walsh (2008), the j^{th} input good is produced by a monopolist *retailer*, with production function:

$$Y_{j,t} = k_{j,t}^\alpha (z_t h_{j,t})^{1-\alpha} - \eta_t^\phi \phi. \quad (2.14)$$

The retailer is a monopolist in the product market and is competitive in the factor markets. Here $k_{j,t}$ denotes the total amount of capital services purchased by firm j . Also, $\eta_t^\phi \phi$ represents an exogenous fixed cost of production, where ϕ is a positive scalar and η_t^ϕ is a process, discussed below, that ensures balanced growth. We calibrate the fixed cost so that retailer profits are zero along the balanced growth path. In (2.14), z_t is a technology shock whose properties are discussed below. Finally, $h_{j,t}$ is the quantity of an intermediate good purchased by the j^{th} retailer. This good is purchased in competitive markets at the price P_t^h from a *wholesaler*. Analogous to CEE, we assume that to produce in period t , the retailer must borrow a share \varkappa of $P_t^h h_{j,t}$ at the interest rate, R_t , that he expects to prevail in the current period. In this way, the marginal cost of a unit of $h_{j,t}$ is

$$P_t^h (\varkappa R_t + (1 - \varkappa)), \quad (2.15)$$

where \varkappa is the fraction of the intermediate input that must be financed. The retailer repays the loan at the end of period t after receiving sales revenues. The j^{th} retailer sets its price, $P_{j,t}$,

subject to the demand curve, (2.13), and the Calvo sticky price friction (2.16). In particular,

$$P_{j,t} = \begin{cases} P_{j,t-1} & \text{with probability } \xi \\ \tilde{P}_t & \text{with probability } 1 - \xi \end{cases} . \quad (2.16)$$

Here, \tilde{P}_t denotes the price set by the fraction $1 - \xi$ of producers who can re-optimize. We assume these producers make their price decision before observing the current period realization of the monetary policy shock, but after the other time t shocks. Note that, unlike CEE, we do not allow the non-optimizing firms to index their prices to some measure of inflation. In this way, the model is consistent with the observation that many prices remain unchanged for extended periods of time (see Eichenbaum, Jaimovich and Rebelo, 2011, and Klenow and Malin, 2011).

2.5. Wholesalers and the Labor Market

A perfectly competitive representative wholesaler firm produces the intermediate good using labor only. Let l_{t-1} denote employment of the wholesaler at the end of period $t - 1$. Consistent with our discussion above, a fraction $1 - \rho$ of workers separates exogenously from the wholesaler at the end of period. A total of ρl_{t-1} workers are attached to the wholesaler at the start of period t . To meet a worker at the beginning the period, the wholesaler must pay a fixed cost, $\eta_t^\kappa \kappa$, and post a suitable number of vacancies. Here, κ is a positive scalar and η_t^κ is a process, discussed below, that ensures balanced growth. To hire $x_t l_{t-1}$ workers, the wholesaler must post $x_t l_{t-1} / Q_t$ vacancies where Q_t denotes the aggregate vacancy filling rate which the representative firm takes as given. Posting vacancies is costless. We assume that the representative firm is large, so that if it posts $x_t l_{t-1} / Q_t$ vacancies, then it meets exactly $x_t l_{t-1}$ workers.

Because of the linearity of the firm's problem, in equilibrium it must make zero profits. That is, the cost of a worker must equal the value, J_t , of a worker:

$$\eta_t^\kappa \kappa = J_t, \quad (2.17)$$

where the objects in (2.17) are expressed in units of the final good.

At the beginning of the period, the representative wholesaler is in contact with a total of l_t workers (see equation (2.4)). This pool of workers includes workers with whom the firm was matched in the previous period, plus the new workers that the firm has just met. Each worker in l_t engages in bilateral bargaining with a representative of the wholesaler, taking the outcome of all other negotiations as given. The equilibrium real wage rate,

$$w_t \equiv W_t / P_t,$$

is the outcome of the bargaining process described below. In equilibrium all bargaining sessions conclude successfully, so the representative wholesaler employs l_t workers. Production begins immediately after wage negotiations are concluded and the wholesaler sells the intermediate good at the real price, $\vartheta_t \equiv P_t^h/P_t$.

Consistent with Hall and Milgrom (2008) and CET, we assume that wages are determined according to the alternating offer bargaining protocol proposed in Rubinstein (1982) and Binmore, Rubinstein and Wolinsky (1986). Let w_t^p denote the expected present discounted value of the wage payments by a firm to a worker that it is matched with:

$$w_t^p = w_t + \rho E_t m_{t+1} w_{t+1}^p.$$

Here m_t is the time t household discount factor which firms and workers view as an exogenous stochastic process beyond their control.

The value of a worker to the firm, J_t , can be expressed as follows:

$$J_t = \vartheta_t^p - w_t^p.$$

Here ϑ_t^p denotes the expected present discounted value of the marginal revenue product associated with a worker to the firm:

$$\vartheta_t^p = \vartheta_t + \rho E_t m_{t+1} \vartheta_{t+1}^p. \quad (2.18)$$

Let V_t denote the value to a worker of being matched with a firm that pays w_t in period t :

$$\begin{aligned} V_t = & w_t + E_t m_{t+1} [\rho V_{t+1} + (1 - \rho) s (f_{t+1} \bar{V}_{t+1} + (1 - f_{t+1}) U_{t+1}) \\ & + (1 - \rho) (1 - s) N_{t+1}]. \end{aligned} \quad (2.19)$$

Here, \bar{V}_{t+1} denotes the value of working for another firm in period $t + 1$. In equilibrium, $\bar{V}_{t+1} = V_{t+1}$. Also, U_{t+1} in (2.19) is the value of being an unemployed worker in period $t + 1$ and N_{t+1} is the value of being out-of-the labor force in period $t + 1$. The objects, s , ρ and f_{t+1} were discussed in the previous section. Relation (2.19) reflects our assumption that an employed worker remains in the same job with probability ρ , transits to another job without passing through unemployment with probability $(1 - \rho) s f_{t+1}$, to unemployment with probability $(1 - \rho) s (1 - f_{t+1})$ and to non-participation with probability $(1 - \rho) (1 - s)$.

It is convenient to rewrite (2.19) as follows:

$$V_t = w_t^p + A_t, \quad (2.20)$$

where

$$\begin{aligned} A_t = & (1 - \rho) E_t m_{t+1} [s f_{t+1} \bar{V}_{t+1} + s (1 - f_{t+1}) U_{t+1} + (1 - s) N_{t+1}] \\ & + \rho E_t m_{t+1} A_{t+1}. \end{aligned} \quad (2.21)$$

According to (2.20), V_t consists of two components. The first is the expected present value of wages received by the worker from the firm with which he is currently matched. The second corresponds to the expected present value of the payments that a worker receives in all dates and states when he is separated from that firm.

The value of unemployment, U_t , is given by,

$$U_t = \eta_t^D D_t + \tilde{U}_t. \quad (2.22)$$

Recall that $\eta_t^D D_t$ represents unemployment compensation at time t . The variable, \tilde{U}_t , denotes the continuation value of unemployment:

$$\tilde{U}_t \equiv E_t m_{t+1} [s f_{t+1} V_{t+1} + s(1 - f_{t+1}) U_{t+1} + (1 - s) N_{t+1}]. \quad (2.23)$$

Expression (2.23) reflects our assumption that an unemployed worker finds a job in the next period with probability $s f_{t+1}$, remains unemployed with probability $s(1 - f_{t+1})$ and exits the labor force with probability $1 - s$.

The value of non-participation is:

$$N_t = E_t m_{t+1} [e_{t+1} (f_{t+1} V_{t+1} + (1 - f_{t+1}) U_{t+1}) + (1 - e_{t+1}) N_{t+1}]. \quad (2.24)$$

Expression (2.24) reflects our assumption that a non-participating worker is selected to join the labor force with probability e_t , defined in (2.3).

The structure of alternating offer bargaining is the same as it is in CET.⁶ Each matched worker-firm pair (both those who just matched for the first time and those who were matched in the past) bargain over the current wage rate, w_t . Each time period (a quarter) is subdivided into M periods of equal length, where M is even. The firm makes a wage offer at the start of the first subperiod. It also makes an offer at the start of every subsequent odd subperiod in the event that all previous offers have been rejected. Similarly, workers make a wage offer at the start of all even subperiods in case all previous offers have been rejected. Because M is even, the last offer is made, on a take-it-or-leave-it basis, by the worker. When the firm rejects an offer it pays a cost, $\eta_t^\gamma \gamma$, of making a counteroffer. Here γ is a positive scalar and η_t^γ is a process that ensures balanced growth.

In subperiod $j = 1, \dots, M - 1$, the recipient of an offer can either accept or reject it. If the offer is rejected the recipient may declare an end to the negotiations or he may plan to make a counteroffer at the start of the next subperiod. In the latter case there is a probability, δ , that bargaining breaks down and the wholesaler and worker revert to their outside option. For the firm, the value of the outside option is zero and for the worker the outside option

⁶When bargaining breaks down, we assume that workers are sent to unemployment, not out-of-the labor force.

is unemployment.⁷ Given our assumptions, workers and firms never choose to terminate bargaining and go to their outside option.

It is always optimal for the firm to offer the lowest wage rate subject to the condition that the worker does not reject it. To know what that wage rate is, the wholesaler must know what the worker would counteroffer in the event that the firm's offer was rejected. But, the worker's counteroffer depends on the firm's counteroffer in case the worker's counteroffer is rejected. We solve for the firm's initial offer beginning with the worker's final offer and working backwards. Since workers and firms know everything about each other, the firm's opening wage offer is always accepted.

Our environment is sufficiently simple that the solution to the bargaining problem has the following straightforward characterization:

$$\alpha_1 J_t = \alpha_2 (V_t - U_t) - \alpha_3 \eta_t^\gamma \gamma + \alpha_4 (\vartheta_t - \eta_t^D D_t) \quad (2.25)$$

where $\beta_i = \alpha_{i+1}/\alpha_1$, for $i = 1, 2, 3$ and,

$$\begin{aligned} \alpha_1 &= 1 - \delta + (1 - \delta)^M \\ \alpha_2 &= 1 - (1 - \delta)^M \\ \alpha_3 &= \alpha_2 \frac{1 - \delta}{\delta} - \alpha_1 \\ \alpha_4 &= \frac{1 - \delta}{2 - \delta} \frac{\alpha_2}{M} + 1 - \alpha_2. \end{aligned}$$

The technical appendix contains a detailed derivation of (2.25) and describes the procedure that we use for solving the bargaining problem.

To summarize, in period t the problem of wholesalers is to choose the hiring rate, x_t , and to bargain with the workers that they meet. These activities occur before the monetary policy shock is realized and after the other shocks are realized.

2.6. Innovations to Technology

In this section we describe the laws of motion of technology. Turning to the investment-specific shock, we assume that $\ln \mu_{\Psi,t} \equiv \ln(\Psi_t/\Psi_{t-1})$ follows an AR(1) process

$$\ln \mu_{\Psi,t} = (1 - \rho_{\Psi}) \ln \mu_{\Psi} + \rho_{\Psi} \ln \mu_{\Psi,t-1} + \sigma_{\Psi} \varepsilon_{\Psi,t}.$$

Here, $\varepsilon_{\Psi,t}$ is the innovation in $\ln \mu_{\Psi,t}$, i.e., the error in the one-step-ahead forecast of $\ln \mu_{\Psi,t}$ based on the history of past observations of $\ln \mu_{\Psi,t}$.

⁷We could allow for the possibility that when negotiations break down the worker has a chance of leaving the labor force. To keep our analysis relatively simple, we do not allow for that possibility here.

For reasons explained later, it is convenient for our post-2008 analysis to adopt a components representation for neutral technology.⁸ In particular, we assume that the growth rate of neutral technology is the sum of a permanent ($\mu_{P,t}$) and a transitory ($\mu_{T,t}$) component:

$$\ln(\mu_{z,t}) = \ln(z_t/z_{t-1}) = \ln(\mu_z) + \mu_{P,t} + \mu_{T,t}, \quad (2.26)$$

where

$$\mu_{P,t} = \rho_P \mu_{P,t-1} + \sigma_P \varepsilon_{P,t}, \quad |\rho_P| < 1, \quad (2.27)$$

and

$$\mu_{T,t} = \rho_T \mu_{T,t-1} + \sigma_T (\varepsilon_{T,t} - \varepsilon_{T,t-1}), \quad |\rho_T| < 1. \quad (2.28)$$

In (2.27) and (2.28), $\varepsilon_{P,t}$ and $\varepsilon_{T,t}$ are mean zero, unit variance, iid shocks. To see why (2.28) is the transitory component of $\ln(z_t)$, suppose $\mu_{P,t} \equiv 0$ so that $\mu_{T,t}$ is the only component of technology and (ignoring the constant term) $\ln(\mu_{z,t}) = \mu_{T,t}$, or

$$\ln(\mu_{z,t}) = \ln(z_t) - \ln(z_{t-1}) = \rho_T (\ln(z_{t-1}) - \ln(z_{t-2})) + \sigma_T (\varepsilon_{T,t} - \varepsilon_{T,t-1}).$$

Diving by $1 - L$, where L denotes the lag operator, we have:

$$\ln(z_t) = \rho_T \ln(z_{t-1}) + \sigma_T \varepsilon_{T,t}.$$

Thus, a shock to $\varepsilon_{T,t}$ has only a transient effect on the forecast of $\ln(z_t)$. By contrast a shock, say $\Delta\varepsilon_{P,t}$, to $\varepsilon_{P,t}$ shifts $E_t \ln(z_{t+j})$, $j \rightarrow \infty$ by the amount, $\Delta\varepsilon_{P,t}/(1 - \rho_P)$.

We assume that when there is a shock to $\ln(z_t)$, agents do not know whether it reflects the permanent or the temporary component. As a result, they must solve a signal extraction problem when they adjust their forecast of future values of $\ln(z_t)$ in response to an unanticipated move in $\ln(z_t)$. Suppose, for example, there is a shock to $\varepsilon_{P,t}$, but that agents believe most fluctuations in $\ln(z_t)$ reflect shocks to $\varepsilon_{T,t}$. In this case they will adjust their near term forecast of $\ln(z_t)$, leaving their longer-term forecast of $\ln(z_t)$ unaffected. As time goes by and agents see that the change in $\ln(z_t)$ is too persistent to be due to the transitory component, the long-run component of their forecast of $\ln(z_t)$ begins to adjust. Thus, a disturbance in $\varepsilon_{P,t}$ triggers a sequence of forecast errors for agents who cannot observe whether a shock to $\ln(z_t)$ originates in the temporary or permanent component of $\ln(\mu_{z,t})$.

Because agents do not observe the components of technology directly, they do not use the components representation to forecast technology growth. For forecasting, they use the univariate Wold representation that is implied by the components representation. The shocks to the permanent and transitory components of technology enter the system by perturbing

⁸Unobserved components representations have played an important role in macroeconomic analysis. See, for example, Erceg and Levin (2003) and Edge, Laubach and Williams (2007).

the error in the Wold representation. To clarify these observations we first construct the Wold representation.

Multiply $\ln(\mu_{z,t})$ in (2.26) by $(1 - \rho_P L)(1 - \rho_T L)$, where L denotes the lag operator:

$$(1 - \rho_P L)(1 - \rho_T L) \ln(\mu_{z,t}) = (1 - \rho_T L) \sigma_P \varepsilon_{P,t} + (1 - \rho_P L) (\sigma_T \varepsilon_{T,t} - \sigma_T \varepsilon_{T,t-1}). \quad (2.29)$$

Let the stochastic process on the right of the equality be denoted by \mathcal{W}_t . Evidently, \mathcal{W}_t has a second order moving average representation, which we express in the following form:

$$\mathcal{W}_t = (1 - \theta_1 L - \theta_2 L^2) \sigma_\eta \eta_t, \quad E\eta_t = 1. \quad (2.30)$$

We obtain a mapping from $\rho_P, \rho_T, \sigma_P, \sigma_T$ to $\theta_1, \theta_2, \sigma_\eta$ by first computing the variance and two lagged covariances of the object to the right of the first equality in (2.29). We then find the values of θ_1, θ_2 , and σ_η for which the variance and two lagged covariances of \mathcal{W}_t and the object on the right of the equality in (2.29) are the same. In addition, we require that the eigenvalues in the moving average representation of \mathcal{W}_t , (2.30), lie inside the unit circle. The latter condition is what guarantees that the shock in the Wold representation is the innovation in technology. In sum, the Wold representation for $\ln(\mu_{z,t})$ is:

$$(1 - \rho_P L)(1 - \rho_T L) \ln(\mu_{z,t}) = (1 - \theta_1 L - \theta_2 L^2) \sigma_\eta \eta_t. \quad (2.31)$$

The mapping from the structural shocks, $\varepsilon_{P,t}$ and $\varepsilon_{T,t}$, to η_t is obtained by equating the objects on the right of the equalities in (2.29) and (2.30):

$$\eta_t = \theta_1 \eta_{t-1} + \theta_2 \eta_{t-2} + \frac{\sigma_P}{\sigma_\eta} (\varepsilon_{P,t} - \rho_T \varepsilon_{P,t-1}) + (1 - \rho_P L) \frac{\sigma_T}{\sigma_\eta} (\varepsilon_{T,t} - \varepsilon_{T,t-1}). \quad (2.32)$$

According to this expression, if there is a positive disturbance to $\varepsilon_{P,t}$, this triggers a sequence of one-step-ahead forecast errors for agents, consistent with the intuition described above.

When we estimate our model, we treat the innovation in technology, η_t , as a primitive and are not concerned with the decomposition of η_t into the $\varepsilon_{P,t}$'s and $\varepsilon_{T,t}$'s. In effect, we replace the unobserved components representation of the technology shock with its representation in (2.31). That representation is an autoregressive, moving average representation with two autoregressive parameters, two moving average parameters and a standard deviation parameter. Thus, in principle it has five free parameters. But, since the Wold representation is derived from the unobserved components model, it has only four free parameters. Specifically, we estimate the following parameters: ρ_P, ρ_T, σ_P and the ratio $\frac{\sigma_T}{\sigma_P}$.

Although we do not make use of the decomposition of the innovation, η_t , into the structural shocks when we estimate our model, it turns out that the decomposition is very useful for interpreting the post-2008 data.

2.7. Market Clearing, Monetary Policy and Functional Forms

The total supply of the intermediate good is given by l_t which equals the total quantity of labor used by the wholesalers. So, clearing in the market for intermediate goods requires

$$h_t = l_t, \quad (2.33)$$

where

$$h_t \equiv \int_0^1 h_{j,t} dj.$$

The capital services market clearing condition is:

$$u_t^K K_t = \int_0^1 k_{j,t} dj.$$

Market clearing for final goods requires:

$$C_t + (I_t + a(u_t^K)K_t)/\Psi_t + \eta_t^\kappa \kappa x_t l_{t-1} + G_t = Y_t. \quad (2.34)$$

The right hand side of the previous expression denotes the quantity of final goods. The left hand side represents the various ways that final goods are used. Homogeneous output, Y_t , can be converted one-for-one into either consumption goods, goods used to hire workers, or government purchases, G_t . In addition, some of Y_t is absorbed by capital utilization costs. Homogeneous output, Y_t can also be used to produce investment goods using a linear technology in which one unit of the final good is transformed into Ψ_t units of I_t . Perfect competition in the production of investment goods implies,

$$P_{I,t} = \frac{P_t}{\Psi_t}.$$

Finally, clearing in the loan market requires that the demand for loans by wholesalers, $\varkappa h_t P_t^h$, equals the supply, B_{t+1} :

$$\varkappa h_t P_t^h = B_{t+1}.$$

We adopt the following specification of monetary policy:

$$\begin{aligned} \ln(R_t/R) &= \rho_R \ln(R_{t-1}/R) \\ &+ (1 - \rho_R) \left[0.25r_\pi \ln\left(\frac{\pi_t^A}{\pi^A}\right) + r_y \ln\left(\frac{\mathcal{Y}_t}{\mathcal{Y}_t^*}\right) + 0.25r_{\Delta y} \ln\left(\frac{\mathcal{Y}_t}{\mathcal{Y}_{t-4}\mu_y^A}\right) \right] + \sigma_R \varepsilon_{R,t}, \end{aligned} \quad (2.35)$$

where $\pi_t^A \equiv P_t/P_{t-4}$ and π^A is the monetary authority's inflation target. The object, π^A is also the value of π_t^A in nonstochastic steady state. The shock, $\varepsilon_{R,t}$, is a unit variance, zero mean and serially uncorrelated disturbance to monetary policy. The variable, \mathcal{Y}_t , denotes Gross Domestic Product (GDP):

$$\mathcal{Y}_t = C_t + I_t/\Psi_t + G_t,$$

where G_t denotes government consumption, which is assumed to have the following representation:

$$G_t = \eta_t^g g_t. \quad (2.36)$$

Here, η_t^g is a process that guarantees balanced growth and g_t is an exogenous stochastic process. The variable, \mathcal{Y}_t^* , is defined as follows:

$$\mathcal{Y}_t^* = \eta_t^y \iota, \quad (2.37)$$

where η_t^y is a process that guarantees balanced growth and ι is a constant chosen to guarantee that $\ln(\mathcal{Y}_t/\mathcal{Y}_t^*)$ converges to zero in nonstochastic steady state. The constant, $\mu_{\mathcal{Y}}^A$, is the value of $\mathcal{Y}_t/\mathcal{Y}_{t-4}$ in nonstochastic steady state. Also, R denotes the steady state value of R_t .

The sources of long-term growth in our model are the neutral and investment-specific technological progress shocks discussed in the previous subsection. The growth rate in steady state for the model variables is a composite, Φ_t , of these two technology shocks:

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

The variables Y_t/Φ_t , C_t/Φ_t , w_t/Φ_t and $I_t/(\Psi_t\Phi_t)$ converge to constants in nonstochastic steady state.

If objects like the fixed cost of production, the cost of hiring, etc., were constant, they would become irrelevant over time. To avoid this implication, it is standard in the literature to suppose that such objects are proportional to the underlying source of growth, which is Φ_t in our setting. However, this assumption has the unfortunate implication that technology shocks of both types have an immediate effect on the vector of objects

$$\Omega_t = \left[\eta_t^y, \eta_t^g, \eta_t^D, \eta_t^\gamma, \eta_t^\kappa, \eta_t^\phi, \eta_t^L, \eta_t^H \right]'. \quad (2.38)$$

Such a specification seems implausible and so we instead proceed as in Christiano, Trabandt and Walentin (2012) and Schmitt-Grohé and Uribe (2012). In particular, we suppose that the objects in Ω_t are proportional to a long moving average of composite technology, Φ_t :

$$\Omega_{i,t} = \Phi_{t-1}^\theta (\Omega_{i,t-1})^{1-\theta}, \quad (2.39)$$

where $\Omega_{i,t}$ denotes the i^{th} element of Ω_t , $i = 1, \dots, 8$. Also, $0 < \theta \leq 1$ is a parameter to be estimated. Note that $\Omega_{i,t}$ has the same growth rate in steady state as GDP. When θ is very close to zero, $\Omega_{i,t}$ is virtually unresponsive in the short-run to an innovation in either of the two technology shocks, a feature that we find very attractive on a priori grounds.

We adopt the investment adjustment cost specification proposed in CEE. In particular, we assume that the cost of adjusting investment takes the form:

$$\begin{aligned} S(I_t/I_{t-1}) &= 0.5 \exp \left[\sqrt{S''} (I_t/I_{t-1} - \mu_\Phi \times \mu_\Psi) \right] \\ &\quad + 0.5 \exp \left[-\sqrt{S''} (I_t/I_{t-1} - \mu_\Phi \times \mu_\Psi) \right] - 1. \end{aligned}$$

Here, μ_Φ and μ_Ψ denote the steady state growth rates of Φ_t and Ψ_t . The value of I_t/I_{t-1} in nonstochastic steady state is $(\mu_\Phi \times \mu_\Psi)$. In addition, S'' represents a model parameter that coincides with the second derivative of $S(\cdot)$, evaluated in steady state. It is straightforward to verify that $S(\mu_\Phi \times \mu_\Psi) = S'(\mu_\Phi \times \mu_\Psi) = 0$. Our specification of the adjustment costs has the convenient feature that the steady state of the model is independent of the value of S'' .

We assume that the cost associated with setting capacity utilization is given by,

$$a(u_t^K) = 0.5\sigma_a\sigma_b(u_t^K)^2 + \sigma_b(1 - \sigma_a)u_t^K + \sigma_b(\sigma_a/2 - 1)$$

where σ_a and σ_b are positive scalars. We normalize the steady state value of u_t^K to unity, so that the adjustment costs are zero in steady state, and σ_b is equated to the steady state of the appropriately scaled rental rate on capital. Our specification of the cost of capacity utilization and our normalization of u_t^K in steady state has the convenient implication that the model steady state is independent of σ_a .

Finally, we discuss the determination of the equilibrium vacancy filling rate, Q_t . We posit a standard matching function:

$$x_t l_{t-1} = \sigma_m (L_t - \rho l_{t-1})^\sigma (l_{t-1} v_t)^{1-\sigma}, \quad (2.40)$$

where $l_{t-1} v_t$ denotes the economy-wide average number of vacancies and v_t denotes the aggregate vacancy rate. Then,

$$Q_t = \frac{x_t}{v_t}. \quad (2.41)$$

3. Data and Econometric Methodology for Pre-2008 Sample

We estimate our model using a Bayesian variant of the strategy in CEE that minimizes the distance between the dynamic response to three shocks in the model and the analog objects in the data. The latter are obtained using an identified VAR for post-war quarterly U.S. times series that include key labor market variables. The particular Bayesian strategy that we use is the one developed in Christiano, Trabandt and Walentin (2011), henceforth CTW.

CTW estimate a 14 variable VAR using quarterly data that are seasonally adjusted and cover the period 1951Q1 to 2008Q4. To facilitate comparisons, our analysis is based on the same VAR that CTW use. As in CTW, we identify the dynamic responses to a monetary policy shock by assuming that the monetary authority observes the current and lagged values of all the variables in the VAR, and that a monetary policy shock affects only the Federal Funds Rate contemporaneously. As in Altig, Christiano, Eichenbaum and Linde (2011), Fisher (2006) and CTW, we make two assumptions to identify the dynamic responses to the technology shocks: (i) the only shocks that affect labor productivity in the long-run are the innovations to the neutral technology shock, η_t , and the innovations to the investment-specific

technology shock $\varepsilon_{\Psi,t}$, and (ii) the only shocks that affects the price of investment relative to consumption in the long-run are the innovations to the investment-specific technology shock $\varepsilon_{\Psi,t}$. These assumptions are satisfied in our model. Standard lag-length selection criteria lead CTW to work with a VAR with 2 lags.⁹ The assumptions used to identify the effects of monetary policy and technology shocks are satisfied in our model.

We include the following variables in the VAR:

$$\left(\begin{array}{c} \Delta \ln(\text{relative price of investment}_t) \\ \Delta \ln(\text{real } GDP_t/\text{hours}_t) \\ \Delta \ln(GDP \text{ deflator}_t) \\ \text{unemployment rate}_t \\ \ln(\text{capacity utilization}_t) \\ \ln(\text{hours}_t) \\ \ln(\text{real } GDP_t/\text{hours}_t) - \ln(\text{real wage}_t) \\ \ln(\text{nominal } C_t/\text{nominal } GDP_t) \\ \ln(\text{nominal } I_t/\text{nominal } GDP_t) \\ \ln(\text{job vacancies}_t) \\ \text{job separation rate}_t \\ \text{job finding rate}_t \\ \ln(\text{hours}_t/\text{labor force}_t) \\ \text{federal funds rate}_t \end{array} \right). \quad (3.1)$$

See section A of the technical appendix in CTW for details about the data. Here, we briefly discuss the job vacancy data. Our time series on vacancies splices together a help-wanted index produced by the Conference Board with a job openings measure produced by the Bureau of Labor Statistics in their Job Openings and Labor Turnover Survey (JOLTS). According to JOLTS, a ‘job opening’ is a position that the firm would fill in the event that a suitable candidate appears. A job vacancy in our model corresponds to this definition of a ‘job opening’. To see this, recall that in our model the representative firm is large. We can think of our firm as consisting of a large number of plants. Suppose that the firm wants to hire z people per plant when the vacancy filling rate is Q . The firm instructs each plant to post z/Q vacancies with the understanding that each vacancy which generates a job application will be turned into a match.¹⁰ This is the sense in which vacancies in our model meet the JOLTS definition of a job opening. Of course, it is possible that the people responding to the JOLTS survey report job opening numbers that correspond more closely to z . To the extent that this is true, the JOLTS data should be thought of as a noisy indicator of vacancies in our model. This measurement issue is not unique to our model. It arises in the standard search and matching model (see, for example, Shimer (2005)).

⁹See CTW for a sensitivity analysis with respect to the lag length of the VAR.

¹⁰Some plants will hire more than z people and others will hire fewer. By the law of large numbers, there is no uncertainty at the firm level about how many people will be hired.

Given an estimate of the VAR we compute the implied impulse response functions to the three structural shocks. We stack the contemporaneous and 14 lagged values of each of these impulse response functions for 13 of the variables listed above in a vector, $\hat{\psi}$. We do not include the job separation rate because that variable is constant in our model. We include the job separation rate in the VAR to ensure the VAR results are not driven by an omitted variable bias.

The logic underlying our model estimation procedure is as follows. Suppose that our structural model is true. Denote the true values of the model parameters by θ_0 . Let $\psi(\theta)$ denote the model-implied mapping from a set of values for the model parameters to the analog impulse responses in $\hat{\psi}$. Thus, $\psi(\theta_0)$ denotes the true value of the impulse responses whose estimates appear in $\hat{\psi}$. According to standard classical asymptotic sampling theory, when the number of observations, T , is large, we have

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

Here, ζ_0 denotes the true values of the parameters of the shocks in the model that we do not formally include in the analysis. Because we solve the model using a log-linearization procedure, $\psi(\theta_0)$ is not a function of ζ_0 . However, the sampling distribution of $\hat{\psi}$ is a function of ζ_0 . We find it convenient to express the asymptotic distribution of $\hat{\psi}$ in the following form:

$$\hat{\psi} \overset{a}{\sim} N(\psi(\theta_0), V), \tag{3.2}$$

where

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

For simplicity our notation does not make the dependence of V on θ_0 , ζ_0 and T explicit. We use a consistent estimator of V . Motivated by small sample considerations, that estimator has only diagonal elements (see CTW). The elements in $\hat{\psi}$ are graphed in Figures 1 – 3 (see the solid lines). The gray areas are centered, 95 percent probability intervals computed using our estimate of V .

In our analysis, we treat $\hat{\psi}$ as the observed data. We specify priors for θ and then compute the posterior distribution for θ given $\hat{\psi}$ using Bayes' rule. This computation requires the likelihood of $\hat{\psi}$ given θ . Our asymptotically valid approximation of this likelihood is motivated by (3.2):

$$f(\hat{\psi}|\theta, V) = \left(\frac{1}{2\pi} \right)^{\frac{N}{2}} |V|^{-\frac{1}{2}} \exp \left[-\frac{1}{2} \left(\hat{\psi} - \psi(\theta) \right)' V^{-1} \left(\hat{\psi} - \psi(\theta) \right) \right]. \tag{3.3}$$

The value of θ that maximizes the above function represents an approximate maximum likelihood estimator of θ . It is approximate for three reasons: (i) the central limit theorem

underlying (3.2) only holds exactly as $T \rightarrow \infty$, (ii) our proxy for V is guaranteed to be correct only for $T \rightarrow \infty$, and (iii) $\psi(\theta)$ is calculated using a linear approximation.

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

$$f(\theta|\hat{\psi}, V) = \frac{f(\hat{\psi}|\theta, V) p(\theta)}{f(\hat{\psi}|V)}. \quad (3.4)$$

Here, $p(\theta)$ denotes the priors on θ and $f(\hat{\psi}|V)$ denotes the marginal density of $\hat{\psi}$:

$$f(\hat{\psi}|V) = \int f(\hat{\psi}|\theta, V) p(\theta) d\theta.$$

The mode of the posterior distribution of θ is computed by maximizing the value of the numerator in (3.4), since the denominator is not a function of θ .

4. Empirical Results, Pre-2008 Sample

This section presents results for the estimated model. First, we discuss the priors and posteriors of structural parameters. Second, we discuss the ability of the model to account for the dynamic response of the economy to a monetary policy shock, a neutral technology shock and an investment-specific technology shock.

4.1. Calibration and Parameter Values set a Priori

We set the values for a subset of the model parameters a priori. These values are reported in Panel A of Table 1. We also set the steady state values of five model variables, listed in Panel B of Table 1. We specify β so that the steady state annual real rate of interest is three percent. The depreciation rate on capital, δ_K , is set to imply an annual depreciation rate of 10 percent. The growth rate of composite technology, μ_Φ , is equated to the sample average of real per capita GDP growth. The growth rate of investment-specific technology, μ_Ψ , is set so that $\mu_\Phi\mu_\Psi$ is equal to the sample average of real, per capita investment growth. We assume the monetary authority's inflation target is 2 percent per year and that the profits of intermediate good producers are zero in steady state. We set the steady state value of the vacancy filling rate, Q , to 0.7, as in den Haan, Ramey and Watson (2000) and Ravenna and Walsh (2008). The steady state unemployment rate, u , is set to the average unemployment rate in our sample, 0.05. We assume the parameter M to be equal to 60 which roughly corresponds to the number of business days in a quarter. We set $\rho = 0.9$, which implies a match survival rate that is consistent with both HM and Shimer (2012). Finally, we assume that the steady state value of the ratio of government consumption to gross output is 0.20.

Two additional parameters pertain to the household sector. We set the elasticity of substitution in household utility between home and market produced goods, $1/(1 - \chi)$, to 3. This magnitude is similar to what is reported in Aguiar, Hurst and Karabarbounis (2013).¹¹ We set the steady state labor force to population ratio, L , to 0.67.

To make the model consistent with the 5 calibrated values for L , Q , G/Y , u , and profits, we select values for 5 parameters: the weight of market consumption in the utility function, ω ; the constant in front of the matching function, σ_m ; the fixed cost of production, ϕ ; the cost for the firm to make a counteroffer, γ ; and the scale parameter, g , in government consumption. The values for these parameters, evaluated at the posterior mode of the set of parameters that we estimate, are reported in Table 3.

4.2. Parameter Estimation

The priors and posteriors for the model parameters about which we do Bayesian inference are summarized in Table 2. A number of features of the posterior mode of the estimated parameters of our model are worth noting.

First, the posterior mode of ξ implies a moderate degree of price stickiness, with prices changing on average once every 4 quarters. This value lies within the range reported in the literature.

Second, the posterior mode of δ implies that there is a roughly 0.05 percent chance of an exogenous break-up in negotiations when a wage offer is rejected.

Third, the posterior modes of our model parameters, along with the assumption that the steady state unemployment rate equals 5.5 percent, implies that it costs firms 0.81 days of marginal revenue to prepare a counteroffer during wage negotiations (see Table 3).

Fourth, the posterior mode of steady state hiring costs as a percent of gross output is equal to 0.5 percent. This result implies that steady state hiring costs as a percent of total wages of newly-hired workers is equal to 7 percent. Silva and Toledo (2009) report that, depending on the exact costs included, the value of this statistic is between 4 and 14 percent, a range that encompasses the corresponding statistic in our model

Fifth, the posterior mode of the replacement ratio is 0.19. To put this number in perspective, consider the following narrow measure of the fraction of unemployment benefits to wages in the data. The numerator of the fraction is total payments of the government for

¹¹We take our elasticity of substitution parameter from the literature to maintain comparability. However, there is a caveat. To understand this, recall the definition of the elasticity of substitution. It is the percent change in C/C^H in response to a one percent change in the corresponding relative price, say λ . From an empirical standpoint, it is difficult to obtain a direct measure of this elasticity because we do not have data on C^H or λ . As a result, structural relations must be assumed, which map from observables to C^H and λ . Since estimates of the elasticity are presumably dependent on the details of the structural assumptions, it is not clear how to compare values of this parameter across different studies, which make different structural assumptions.

unemployment insurance divided by the total number of unemployed people. The denominator of the fraction is total compensation of labor divided by the number of employees, i.e. the average wage per worker. The average of the numerator divided by denominator in our sample period is 0.14. This fraction represents the lower bound on the average replacement rate because it leaves out some other government contributions that unemployed people are eligible for. HM summarize the literature and report a range of estimates from 0.12 to 0.36 for the replacement ratio. It is well known that Diamond (1982), Mortensen (1982) and Pissarides (1985) (DMP) style models require a replacement ratio in excess of 0.9 to account for fluctuations in labor markets (see e.g. CET for an extended discussion). For the reasons stressed in CET, alternating offer bargaining between workers and firms mutes the sensitivity of real wages to aggregate shocks. This property underlies our model's ability to account for the estimated response of the economy to monetary policy shocks and shocks to neutral and investment-specific technology with a low replacement ratio.

Sixth, the posterior mode of s implies that a separated or unemployed worker leaves the labor force with probability $1 - s = 0.2$.

Seventh, the posterior mode for α_c is 0.02, implying that people out-of-the labor force account for virtually all of home production.

Eighth, the posterior mode of θ which governs the responsiveness of the elements of Ω_t to technology shocks, is small (0.16). So, variables like government purchases and unemployment benefits are very unresponsive in the short-run to technology shocks.

Ninth, the posterior modes of the parameters governing monetary policy are similar to those reported in the literature (see for example Justiniano, Primiceri, and Tambalotti, 2010).

Tenth, we turn to the parameters of the unobserved components representation of the neutral technology shock. According to the posterior mode, the standard deviation of the shock to the transient component is roughly 5 times the standard deviation of the permanent component. So, according to the posterior mode, most of the fluctuations (at least, at a short horizon) are due to the transitory component of neutral technology. This result is driven primarily by our prior, the rationale for which is discussed in section 5.2. The permanent component of neutral technology has an autocorrelation of roughly 0.8, so that a one percent shock to the permanent component eventually drives the level of technology up by about 5 percent. The temporary component is also fairly highly autocorrelated.

Many authors conclude that the growth rate of neutral technology follows roughly a random walk (see, for example, Prescott, 1986). Our model is consistent with this view. We find that the first order autocorrelation of $\ln(z_t/z_{t-1})$ in our model is 0.06, which is very close to zero. For discussions of how a components representation, in which the components are both highly autocorrelated, can nevertheless generate a process that looks like a random walk, see Christiano and Eichenbaum (1990) and Quah (1990).

Table 4 reports the frequency with which workers transit between the three states that they can be in. The table reports the steady state frequencies implied by the model and the analog statistics calculated from data from the Current Population Survey. Note that we did not use these data statistics when we estimated or calibrated the model.¹² Nevertheless, with two minor exceptions, the model does very well at accounting for those statistics of the data. The exceptions are that the model somewhat understates the frequency of transition from unemployment into unemployment and slightly overstates the frequency of transition from unemployment to out-of-the labor force. Finally, we note that in the data over half of newly employed people are hired from other jobs (see Diamond 2010, page 316). Our model is consistent with this fact: in the steady state of the model, 51 percent of newly employed workers in a given quarter come from other jobs.¹³ Overall, we view these findings as additional evidence in support of the notion that our model of the labor market is empirically plausible.

4.3. Impulse Response Functions

The solid black lines in Figures 1-3 present the impulse response functions to a monetary policy shock, a neutral technology shock and an investment-specific technology shock implied by the estimated VAR. The grey areas represent 95 percent probability intervals. The solid blue lines correspond to the impulse response functions of our model evaluated at the posterior mode of the structural parameters. Figure 1 shows that the model does very well at reproducing the estimated effects of an expansionary monetary policy shock, including the hump-shaped rises in real GDP and hours worked, the rise in the labor force participation rate and the muted response of inflation. Notice that real wages respond by much less than hours worked to a monetary policy shock. Even though the maximal rise in hours worked is roughly 0.14 percent, the maximal rise in real wages is only 0.06 percent. Significantly, the model accounts for the hump-shaped fall in the unemployment rate as well as the rise in the job finding rate and vacancies that occur after an expansionary monetary policy shock. The model does understate the rise in the capacity utilization rate. The sharp rise of capacity utilization in the estimated VAR may reflect that our data on the capacity utilization rate pertains to the manufacturing sector, which probably overstates the average response across all sectors in the economy.

¹²Our data does include the job finding rate. However, our impulse response matching procedure only uses the dynamics of that variable and not its level.

¹³We reached this conclusion as follows. Workers starting a new job at the start of period t come from three states: employment, unemployment and not-in-the labor force. The quantities of these people are $(1 - \rho)l_{t-1}sf_t$, $f_tsu_{t-1}L_{t-1}$ and $f_t e_t (1 - L_t)$, respectively. We computed these three objects in steady state using the information in Tables 1, 2 and 3. The fraction reported in the text is the ratio of the first number to the sum of all three.

From Figure 2 we see that the model does a good job of accounting for the estimated effects of a negative innovation, η_t , to neutral technology (see (2.31)). Note that the model is able to account for the initial fall and subsequent persistent rise in the unemployment rate. The model also accounts for the initial rise and subsequent fall in vacancies and the job finding rate after a negative shock to neutral technology. The model is consistent with the relatively small response of the labor force participation rate to a technology shock.

Turning to the response of inflation after a negative neutral technology shock, note that our VAR implies that the maximal response occurs in the period of the shock.¹⁴ Our model has no problem reproducing this observation. See CTW for intuition.

Figure 3 reports the VAR-based estimates of the responses to an investment-specific technology shock. The figure also displays the responses to $\varepsilon_{\Psi,t}$ implied by our model evaluated at the posterior mode of the parameters. Note that in all cases the model impulses lie in the 95 percent probability interval of the VAR-based impulse responses.

Viewed as a whole, the results of this section provide evidence that our model does well at accounting for the cyclical properties of key labor market and other macro variables in the pre-2008 period.

5. Modeling The Great Recession

In this section we provide a quantitative characterization of the Great Recession. We suppose that the economy was buffeted by a sequence of shocks that began in 2008Q3. Using simple projection methods, we estimate how the economy would have evolved in the absence of those shocks. The difference between how the economy would have evolved and how it did evolve is what we define as the Great Recession. We then extend our modeling framework to incorporate four candidate shocks that in principle could have caused the Great Recession. In addition, we provide an interpretation of monetary policy during the Great Recession, allowing for a binding ZLB and forward guidance. Finally, we discuss our strategy for stochastically simulating our model.

5.1. Characterizing the Great Recession

The solid line in Figure 4 displays the behavior of key macroeconomic variables since 2001. To assess how the economy would have evolved absent the large shocks associated with the Great Recession, we adopt a simple and transparent procedure. With five exceptions, we fit a linear trend from 2001Q1 to 2008Q2, represented by the dashed red line. To characterize what the data would have looked like absent the shocks that caused the financial crisis and Great

¹⁴This finding is consistent with results in e.g. Altig, Christiano, Eichenbaum and Linde (2011) and Paciello (2011).

Recession, we extrapolate the trend line (see the thin dashed line) for each variable. According to our model, all the nonstationary variables in the analysis are difference stationary. Our linear extrapolation procedure implicitly assumes that the shocks in the period 2001-2008 were small relative to the drift terms in the time series. Given this assumption, our extrapolation procedure approximately identifies how the data would have evolved, absent shocks after 2008Q2.

Four of the exceptions to our extrapolation method are inflation, the federal funds rate, the unemployment rate and the job finding rate. For these variables, we assume a no-change trajectory after 2008Q2. In the case of these four variables, our linear projection procedure led to implausible results. For example, the federal funds rate would be projected to be almost 6 percent in 2013Q2.

The fifth exception pertains to a measure of the spread between the corporate-borrowing rate and the risk-free interest rate. The measure that we use corresponds to the one provided in Gilchrist and Zakrajšek (2012) (GZ). These data are displayed in the (3,4) element of Figure 4. Our projection for the period after 2008Q2 is that the GZ spread falls linearly to 1 percent, its value during the relatively tranquil period, 1990-1997.

There are of course many alternative procedures for projecting the behavior of the economy. For example, we could use separate ARMA time series models for each of the variables or we could use multivariate methods including the VAR estimated with pre-2008Q2 data. A challenge for a multivariate approach is the nonlinearity associated with the ZLB. Still, it would be interesting to pursue alternative projection approaches in the future.

The projections for the labor force and employment after 2008Q2 are controversial because of ongoing demographic changes in the U.S. population. Our procedure attributes roughly 1.5 percentage points of the 2.5 percentage points decline in the labor force participation rate since 2008 to cyclical factors. Projections for the labor force to population ratio published by the Bureau of Labor Statistics in November 2007 suggest that the cyclical component in the decline in this ratio was roughly 2 percentage points.¹⁵ Reifschneider, Wascher and Wilcox (2013) and Sullivan (2013) estimate that the cyclical component of the decline in the labor force to population ratio is equal to 1 percentage point and 0.75 percentage points, respectively. So, we are roughly at the mid-point of these estimates.

According to Figure 4, the employment to population ratio fell by about 5 percentage points from 2008 to 2013. According to our procedure only a small part, about 0.5 percentage points, of this drop is accounted for by non-cyclical factors. Krugman (2014) argues that 1.6 percentage points of the 5 percentage points are due to demographic factors. So, like us, he ascribes a small portion of the decline in the employment to population ratio to non-cyclical factors. In contrast, Tracy and Kapon (2014) argue that the cyclical component of the decline

¹⁵See Erceg and Levin (2013), Figure 1.

was smaller. To the extent that this is true, it would be easier for our model to account for the data. In this sense we are adopting a conservative stance.

The distances between the solid lines and the thin-dashed lines in Figure 4 represents our estimates of the economic effects of the shocks that hit the economy in 2008Q3 and later. In effect, these distances are the Great Recession targets that we seek to explain.

Some features of the targets are worth emphasizing. First, there is large drop in log GDP. While some growth began in late 2009, per capita GDP has still not returned to its pre-crisis level as of the end of our sample. Second, there was a very substantial decline in consumption and investment. While the latter showed strong growth since late 2009 it has not yet surpassed its pre-crisis peak in per capita terms. Strikingly, although per capita consumption initially grew starting in late 2009, it stopped growing around the middle of 2012. The stop of consumption growth is mirrored by a slowdown in the growth rate of GDP and investment at around the same time.

An obvious candidate for a macro shock during this time period were the events surrounding the debt ceiling crisis and the sequester. It is obviously difficult to pick one particular date for when agents took seriously the possibility that the U.S. government was going to fall off the fiscal cliff. Still, it is interesting to note that in Spring 2012, Chairman Bernanke warned lawmakers of a ‘massive fiscal cliff’ involving year-end tax increases and spending cuts.¹⁶

Third, vacancies dropped sharply before 2009 and then rebounded almost to their pre-recession levels. At the same time, unemployment rose sharply, but then only fell modestly. Kocherlakota (2010) interprets these observations as implying that firms had positions to fill, but the unemployed workers were simply not suitable. This explanation is often referred to as the mismatch hypothesis. Davis, Faberman, and Haltiwanger (2012) provide a different interpretation of these observations. In their view, what matters for filling jobs is the intensity of firms’ recruiting efforts, not vacancies per se. They argue that the intensity with which firms recruited workers after 2009 did not rebound in the same way that vacancies did. Perhaps surprisingly, our model can account for the joint behavior of unemployment and vacancies, even though the forces stressed by Kocherlakota and Davis, et al. are absent from our framework.

Finally, we note that despite the steep drop in GDP, inflation dropped by only about 1/2 - 1 percentage points. Authors like Hall (2011) argue that this joint observation is particularly challenging for NK models.

¹⁶According to the Huffington Post (http://www.huffingtonpost.com/2012/12/27/fiscal-cliff-2013_n_2372034.html) in Autumn of 2012, many economists warned that if left unaddressed, concerns about the ‘fiscal cliff’, could trigger a recession.

5.2. The Shocks Driving the Great Recession

We suppose that the Great Recession was triggered by four shocks. Two of these shocks are wedges which capture in a reduced form way frictions which are widely viewed as having been important during the Great Recession. The other sources of shocks that we allow for are government consumption and technology.

Financial Shocks

We begin by discussing the two financial shocks. The first is a shock to households preferences for safe and/or liquid assets. We capture this shock by introducing a perturbation, Δ_t^b , to agents' intertemporal Euler equation associated with saving via risk-free bonds. The object, Δ_t^b , is the *consumption wedge* we discussed in the introduction. The Euler equation associated with the nominally risk-free bond is given by:

$$1 = (1 + \Delta_t^b)E_t m_{t+1} R_t / \pi_{t+1}.$$

See Fisher (2014) for a discussion of how a positive realization of Δ_t^b can, to a first-order approximation, be interpreted as reflecting an increase in the demand for risk-free bonds.¹⁷

We do not have data on Δ_t^b . We suppose that in 2008Q3, agents think that Δ_t^b goes from zero to a constant value, 0.33 percent per quarter, for 20 quarters, i.e. until 2013Q2. They expected Δ_t^b to return to zero after that date (see the dashed line in the (2,2) element in Figure 7). We then assume that in 2012Q3, agents revised their expectations and thought that Δ_t^b would remain at 0.33 percent until 2014Q3. We interpret this revision to expectations as a response to the events associated with the fiscal cliff and the sequester. We chose the particular value of Δ_t^b to help the model achieve our targets.

To assess the magnitude of the shock to Δ^b , it is useful to think of Δ^b as a shock to agents' discount rate. Recall from Table 1 that $\beta = 0.9968$, implying an annual discount rate of about 1.3 percent.¹⁸ With $\Delta^b > 0$, the discount factor is in effect $(1 + \Delta^b)\beta$, which implies an annual discount rate of roughly zero percent.¹⁹ So, our Δ^b shock implies that the annual discount rate drops 1.3 percentage points. This drop is substantially smaller than the 6 percentage point drop assumed by Eggertsson and Woodford (2003).

The second financial shock represents a wedge to agents' intertemporal Euler equation for capital accumulation:

$$1 = (1 - \Delta_t^k)E_t m_{t+1} R_{t+1}^k / \pi_{t+1}.$$

¹⁷The shock is also similar to the 'flight-to-quality' shock found to play a substantial role in the start of the Great Depression in Christiano, Motto and Rostagno (2003).

¹⁸In particular, $100 [1/(\beta) - 1] = 1.3$, after rounding.

¹⁹In particular, $100 [1/(\beta(1 + \Delta^b)) - 1] = 0$, after rounding.

A simple financial friction model based on asymmetric information with costly monitoring implies that credit market frictions can be captured as a tax, Δ_t^k , on the gross return on capital (see Christiano and Davis (2006)). The object, Δ_t^k , is the *financial wedge* that we discussed in the introduction.

Recall that firms finance a fraction, \varkappa , of the intermediate input in advance of revenues (see (2.15)). In contrast to the existing DSGE literature, we allow for a *risky working capital channel* in the sense that the financial wedge also applies to working capital loans. Specifically, we replace (2.15) with

$$P_t^h [\varkappa R_t (1 + \Delta_t^k) + (1 - \varkappa)], \quad (5.1)$$

where $\varkappa = 1/2$, as before. The risky working capital channel captures in a reduced form way the frictions modeled in e.g. Bigio (2013).

We measure the financial wedge using the GZ interest rate spread. The latter is based on the average credit spread on senior unsecured bonds issued by non-financial firms covered in Compustat and by the Center for Research in Security Prices. The average and median duration of the bonds in GZ's data is 6.47 and 6.00 years, respectively. We interpret the GZ spread as an indicator of Δ_t^k . We suppose that the Δ_t^k 's are related to the GZ spread as follows:

$$\Gamma_t = E \left[\frac{\Delta_t^k + \Delta_{t+1}^k + \dots + \Delta_{t+23}^k}{6} \middle| \Omega_t \right], \quad (5.2)$$

where Γ_t denotes the GZ spread minus the projection of that spread as of 2008Q2. Also, Ω_t denotes the information available to agents at time t . In (5.2) we sum over Δ_{t+j}^k for $j = 0, \dots, 23$ because Δ_t^k is a tax on the one quarter return to capital while Γ_t applies to $t + j$, $j = 0, 1, \dots, 23$ (i.e., 6 years). Also, we divide the sum in (5.2) by 6 to take into account that Δ_t^k is measured in quarterly decimal terms while our empirical measure of Γ_t is measured in annual decimal terms.

We feed the difference between the projected and actual Γ_t 's, for $t \geq 2008Q3$, to the model. The projected and actual Γ_t 's are displayed in the (3,4) element of Figure 4. The difference is displayed in the (1,1) element of Figure 7. We assume that at each date agents must forecast the future values of the Γ_t 's. They do so using a mean zero, first order autoregressive representation (AR(1)), with autoregressive coefficient, $\rho_\Gamma = 0.5$. This low value of ρ_Γ captures the idea that agents thought the sharp increase in the financial wedge was transitory in nature. This belief is consistent with the actual behavior of the GZ spread. To solve their problem, agents actually work with the Δ_t^k 's. But, for any sequence, $\Gamma_t, E_t \Gamma_{t+j}, j = 1, 2, \dots$, they can compute a sequence, $\Delta_t^k, E_t \Delta_{t+j}^k, j = 1, 2, 3, \dots$ that satisfies (5.2).²⁰

Total Factor Productivity (TFP) Shocks

²⁰In performing this computation, we impose that $E_t \Gamma_{t+j} \rightarrow \Gamma_{ss}$ as $j \rightarrow \infty$ and $E_t \Delta_{t+j}^k \rightarrow \Delta_{ss}^k$, where the subscript *ss* signifies the nonstochastic steady state.

We now turn to a discussion of TFP. Various measures produced by the Bureau of Labor Statistics (BLS) are reported in the (1,1) panel of Figure 5. Each measure is the log of value-added minus the log of capital and labor services weighted by their shares in the income generated in producing the measure of value-added.²¹ In each case, we report a linear trend line fitted to the data from 2001Q1 through 2008Q2. We then project the numbers forward after 2008Q2. We do the same for three additional measures of TFP in the (1,2) panel of Figure 5. Two are taken from Fernald (2012) and the third is taken from the Penn World Tables. The bottom panel of Figure 5 displays the post-2008Q2 projection for log TFP minus the log of its actual value. Note that, with one exception, (i) TFP is below its pre-2008 trend during the Great Recession, and (ii) it remains well below its pre-2008 trend all the way up to the end of our data set. The exception is Fernald’s (2012) utilization adjusted TFP measure, which briefly rises above trend in 2009. Features (i) and (ii) of TFP play an important role in our empirical results.

To assess the robustness of (i) and (ii), we redid our calculations using an alternative way of computing the trend lines. Figure 6 reproduces the basic calculations for three of our TFP measures using a linear trend that is constructed using data starting in 1982Q2. While there are some interesting differences across the figures, they have all share the two key features, (i) and (ii), discussed above. Specifically, it appears that TFP was persistently low during the Great Recession.

We now explain why we adopt an unobserved components time series representation of $\ln(z_t)$. If we assume that agents knew in 2008Q3 that the fall in TFP would turn out to be so persistent, then our model generates a counterfactual surge in inflation. We infer that agents only gradually became aware of the persistence in the decline of TFP. The notion that it took agents time to realize that the drop in TFP was highly persistent is consistent with other evidence. For example, Figure 4 in Swanson and Williams (2013) shows that professional forecasts consistently underestimated how long it would take the economy to emerge from the ZLB.

The previous considerations are the reason that we work with the unobserved components representation for $\ln(z_t)$ in (2.26). In addition, these considerations underlie our prior that the standard deviation of the transitory shock is substantially larger than the standard deviation of the permanent shock. We imposed this prior in estimating the model on pre-2008 data.

At this point, it is worth to repeat the observation made in section 4.2 that we have not assumed anything particularly exotic about technology growth. As noted above, our model implies that the growth rate of technology is roughly a random walk, in accordance with a long tradition in business cycle theory. What our analysis in effect exploits is that a process

²¹The BLS measure is only available at an annual frequency. We interpolate the annual data to a quarterly frequency using a standard interpolation routine described in Boot, Feibes, and Lisman (1967).

that is as simple as a random walk can have components that are very different from a random walk.

Our analysis involves simulating the response of the model to shocks. So, we must compute a sequence of realized values of $\ln(z_t)$. Unlike government spending and interest rate spreads, we do not directly observe $\ln(z_t)$. In our model log TFP does not coincide with $\ln(z_t)$. The principle reason for this is the presence of the fixed cost in production in our model. But, the behavior of model-implied TFP is sensitive to $\ln(z_t)$.

To our initial surprise, the behavior is also very sensitive to $\ln(z_t)$. So, from this perspective both inflation and TFP contain substantial information about $\ln(z_t)$. These observations led us to choose a sequence of realized values for $\ln(z_t)$ that, conditional on the other shocks, allows the model to account reasonably well for inflation and log TFP.

The bottom panel of Figure 5 reports the measure of TFP for our model, computed using a close variant of the Bureau of Labor Statistics' procedure.²² The black line with dots displays the model's simulated value of TFP relative to trend (how we detrend and solve the model is discussed below). Note that model TFP lies within the range of empirical measures reported in Figure 5. The bottom panel of Figure 6 shows that we obtain the same result when we detrend our three empirical measures of TFP when using a trend that begins in 1982.

Nonlinear versions of the standard Kalman smoothing methods could be used to combine our model, the values for its parameters, and our data to estimate the sequence of $\ln(z_t)$ (and, Δ_t^b) in the post 2008Q2 data. In practice, this approach is computationally challenging and we defer it to future work.²³ For convenience, we assume there was a one-time shock to $\ln(z_t)$ in 2008Q3. For the reasons given above, we assume that the shock was to the permanent component of $\ln(z_t)$, i.e., ε_t^P . We selected a value of -0.25 percent for that shock so that, in conjunction with our other assumptions, the model does a reasonable job of accounting for post 2008Q2 inflation and log TFP. This one-time shock leads to a persistent move in $\ln(z_t)$ which eventually puts z_t roughly 1.2 percent below the level it would have been in the absence of the shock. The shock to ε_t^P also leads to a sequence of one-step-ahead forecast errors for agents, via (2.32). Our specification of $\ln(z_t)$ captures features (i) and (ii) of the TFP data that were discussed above.

Government Consumption Shocks

Next we consider the shock to government consumption. The variable η_t^g defined in (2.36) is computed using the simulated path of neutral technology, $\ln(z_t)$ (see (2.38) and (2.39)).²⁴

²²Our measure of TFP is the ratio of GDP (i.e., $C + I + G$) to capital and labor services, each raised to a power that corresponds to their steady state share of total income.

²³See e.g. Gust, Lopez-Salido and Smith (2013) who estimate a nonlinear DSGE model subject to an occasionally binding ZLB constraint.

²⁴For simplicity, in our calculations we assume that the investment-specific technology shock remains on its steady state growth path after 2008.

Then, g_t is measured by dividing actual government consumption in Figure 4, by η_t^g . Agents forecast the period t value of η_t^g using current and past realizations of the technology shocks. We assume that agents forecast g_t by using the following AR(2) process:

$$\ln(g_t/g) = 1.6 \ln(g_{t-1}/g) - .64 \ln(g_{t-2}/g) + \varepsilon_t^G,$$

where ε_t^G is a mean zero, unit variance iid shock. We chose the roots for the AR(2) process such that the first and second order autocorrelations of $\Delta \ln G_t$ in our estimated model are close to the data for the sample 1951Q1 to 2008Q2.

5.3. Monetary Policy

We make two key assumptions about monetary policy during the post 2008Q3 period. We assume that the Fed initially followed a version of the Taylor rule that respects the ZLB on the nominal interest rate. We assume that there was an unanticipated regime change in 2011Q3, when the Fed switched to a policy of forward guidance.

5.3.1. Taylor Rule

We now define our version of the Taylor rule that takes the non-negativity constraint on the nominal interest rate into account. Let Z_t denote a gross ‘shadow’ nominal rate of interest, which satisfies the following Taylor-style monetary policy rule:

$$\ln(Z_t) = \ln(R) + r_\pi \ln(\pi_t^A/\pi^A) + 0.25r_y \ln(\mathcal{Y}_t/\mathcal{Y}_t^*) + 0.25r_{\Delta y} \ln(\mathcal{Y}_t/(\mathcal{Y}_{t-4}\mu_{\mathcal{Y}}^A)). \quad (5.3)$$

The actual policy rate, R_t , is determined as follows:

$$\ln(R_t) = \max\{\ln(R/a), \rho_R \ln(Z_{t-1}) + (1 - \rho_R) \ln(Z_t)\}. \quad (5.4)$$

In 2008Q2, the federal funds rate was close to two percent (see Figure 4). Consequently, because of the ZLB, the federal funds rate could only fall by at most two percentage points. To capture this in our model, we set the scalar a to 1.004825.

Absent the ZLB constraint, the policy rule given by (5.3)-(5.4) coincides with (2.35), the policy rule that we estimated using pre-2008 data.

5.3.2. Forward Guidance

We interpret forward guidance as a monetary policy that commits to keeping the nominal interest rate zero until there is substantial improvement in the state of the economy. Initially, in 2011Q3 the Fed did not quantify what they meant by ‘substantial improvement’. Instead, they reported how long they thought it would take until economic conditions would have improved substantially. In December 2012 the Fed became more explicit about what the

state of the economy would have to be for them to consider leaving the ZLB. In particular, the Fed said that it would keep the interest rate at zero as long as inflation remains below 2.5 percent and unemployment remains above 6.5 percent. They did not commit to any particular action in case one or both the thresholds are breached.

In modeling forward guidance we begin with the period, 2011Q3-2012Q4. We do not know what the Fed's thresholds were during this period. But, we do know that in 2011Q3, the Fed announced that it expected the interest rate to remain at zero until mid-2013 (see Campbell, et al., 2012). According to Swanson and Williams (2013), when the Fed made its announcement the number of quarters that professional forecasters expected the interest rate to remain at zero jumped from 4 quarters to 7 or more quarters. We assume that forecasters believed the Fed's announcement and thought that the nominal interest rate would be zero for about 8 quarters. Interestingly, Swanson and Williams (2013) also report that forecasters continued to expect the interest rate to remain at zero for 7 or more quarters in each month through January 2013. Clearly, forecasters were repeatedly revising upwards their expectation of how long the ZLB episode would last. To capture this scenario in a parsimonious way we assume that in each quarter, beginning in 2011Q3 and ending in 2012Q4, agents believed the ZLB would remain in force for another 8 quarters. Thereafter, we suppose that they expected the Fed to revert back to the Taylor rule, (5.3) and (5.4).²⁵

Beginning in 2013Q1, we suppose that agents believed the Fed switched to an explicit threshold rule. Specifically, we assume that agents thought the Fed would keep the Federal Funds rate close to zero until either the unemployment rate fell below 6.5 percent or inflation rose above 2.5 percent. We assume that as soon as these thresholds are met, the Fed switches back to our estimated Taylor rule, (5.3) and (5.4). The latter feature of our rule is an approximation because, as noted above, the Fed did not announce what it would do when the thresholds were met.

5.4. Solving the Model

Our specification of monetary policy includes a non-negativity constraint on the nominal interest rate, as well as regime switching. A subset of the latter depend on realizations of endogenous variables. We search for a solution to our model in the space of sequences.²⁶ The solution satisfies the equilibrium conditions which take the form of a set of stochastic difference equations that are restricted by initial and end-point conditions. Our solution

²⁵Our model of monetary policy is clearly an approximation. For example, it is possible that in our stochastic simulations the Fed's actual thresholds are breached before 8 quarters. Since we do not know what those thresholds were, we do not see a way to substantially improve our approach. Later, in December 2013, the Fed did announce thresholds, but there is no reason to believe that those were their thresholds in the earlier period.

²⁶Our procedure is related to the one proposed in Fair and Taylor (1983).

strategy makes one approximation: certainty equivalence. That is, wherever an expression like $E_t f(x_{t+j})$ is encountered, we replace it by $f(E_t x_{t+j})$, for $j > 0$.

Let y_t denote the vector of shocks operating in the post-2008Q2 period:

$$y_t = \left(\Delta_t^b \quad \Delta_t^k \quad z_t \quad g_t \right)'$$

The law of motion and agents' information sets about y_t have been discussed above.

Let ϱ_t denote the $N \times 1$ vector of period t endogenous variables, appropriately scaled to account for steady growth. We express the equilibrium conditions of the model as follows:

$$E \left[f \left(\varrho_{t+1}, \varrho_t, \varrho_{t-1}, y_t, y_{t+1} \right) \mid \Omega_t \right] = 0. \quad (5.5)$$

Here, the information set is given by

$$\Omega_t = \left\{ \varrho_{t-1-j}, y_{t-j}, j \geq 0 \right\}.$$

Our solution strategy proceeds as follows. As discussed above, we fix a sequence of values for y_t for the periods after 2008Q2. We suppose that at date t agents observe y_{t-s} , $s \geq 0$ for each t after 2008Q2. At each such date t , they compute forecasts, $y_{t+1}^t, y_{t+2}^t, y_{t+2}^t, \dots$, of the future values of y_t . It is convenient to use the notation $y_t^t \equiv y_t$.

We adopt an analogous notation for ϱ_t . In particular, denote the expected value of ϱ_{t+j} formed at time t by ϱ_{t+j}^t , where $\varrho_t^t \equiv \varrho_t$. The equilibrium value of ϱ_t is the first element in the sequence, ϱ_{t+j}^t , $j \geq 0$. To compute this sequence we require y_{t+j}^t , $j \geq 0$, and ϱ_{t-1} . For t greater than 2008Q3 we set $\varrho_{t-1} = \varrho_{t-1}^{t-1}$. For t corresponding to 2008Q3, we set ϱ_{t-1} to its non-stochastic steady state value.

We now discuss how we computed ϱ_{t+j}^t , $j \geq 0$. We do so by solving the equilibrium conditions and imposing certainty equivalence. In particular, ϱ_t^t must satisfy:

$$\begin{aligned} & E \left[f \left(\varrho_{t+1}, \varrho_t, \varrho_{t-1}, y_t, y_{t+1} \right) \mid \Omega_t \right] \\ & \simeq f \left(\varrho_{t+1}^t, \varrho_t^t, \varrho_{t-1}, y_t^t, y_{t+1}^t \right) = 0. \end{aligned}$$

Evidently, to solve for ϱ_t^t requires ϱ_{t+1}^t . Relation (5.5) implies:

$$\begin{aligned} & E \left[f \left(\varrho_{t+2}, \varrho_{t+1}, \varrho_t, y_{t+1}, y_{t+2} \right) \mid \Omega_t \right] \\ & \simeq f \left(\varrho_{t+2}^t, \varrho_{t+1}^t, \varrho_t^t, y_{t+1}^t, y_{t+2}^t \right) = 0. \end{aligned}$$

Proceeding in this way, we obtain a sequence of equilibrium conditions involving ϱ_{t+j}^t , $j \geq 0$. Solving for this sequence requires a terminal condition. We obtain this condition by imposing that ϱ_{t+j}^t converges to the non-stochastic steady state value of ϱ_t . With this procedure it is straightforward to implement our assumptions about monetary policy.

6. The Great Recession: Empirical Results

In this section we analyze the behavior of the economy from 2008Q3 to the end of our sample, 2013Q2. First, we investigate how well our model accounts for the data. Second, we use our model to assess which shocks account for the Great Recession. In addition, we also investigate the role of the risky working capital channel and forward guidance.

6.1. The Model's Implications for the Great Recession

Figure 8 displays our empirical characterization of the Great Recession, i.e., the difference between how the economy would have evolved absent the post 2008Q2 shocks and how it did evolve. In addition, we display the relevant model analogs. For this, we assume that the economy would have been on its steady state growth path in the absence of the post-2008Q2 shocks. This is an approximation that simplifies the analysis and is arguably justified by the fact that the volatility of the economy is much greater after 2008 than it was before. The model analog to our empirical characterization of the Great Recession is the log difference between the variables on the steady state growth path and their response to the post-2008Q2 shocks.

Figure 8 indicates that the model does quite well at accounting for the behavior of our 11 endogenous variables during the Great Recession. Notice in particular that the model is able to account for the modest decline in real wages despite the absence of nominal rigidities in wage setting. Also, notice that the model accounts very well for the average level of inflation despite the fact that our model incorporates only a moderate degree of price stickiness: firms change prices on average once a year. In addition, the model also accounts well for the key labor market variables: labor force participation, employment, unemployment, vacancies and the job finding rate.

Figure 9 provides another way to assess the model's implications for vacancies and unemployment. There, we report a scatter plot with vacancies on the vertical axis and unemployment on the horizontal. The variables in Figure 9 are taken from the 2,1 and 4,1 panels of Figure 8. Although the variables are expressed in deviations from trend, the resulting Beveridge curve has the same key features as those in the raw data (see, for example, Diamond 2010, Figure 4). In particular, notice how actual vacancies fall and unemployment rises from late 2008 to late 2009. This downward relationship is referred to as the Beveridge curve. After 2009, vacancies rise but unemployment falls by less than one would have predicted based on the Beveridge curve that existed before 2009. That is, it appears that after 2009 there was a shift up in the Beveridge curve. This shift is often interpreted as reflecting a deterioration in match efficiency, captured in a simple environment like ours by a fall in the parameter governing productivity in the matching function (see σ_m in (2.40)). This interpretation reflects

a view that models like ours imply a stable downward relationship between vacancies and unemployment, which can only be perturbed by a change in match efficiency. However, this downward relationship is in practice derived as a steady state property of models, and is in fact not appropriate for interpreting quarterly data. To explain this, we consider a simple example.²⁷

Suppose that the matching function is given by:

$$h_t = \sigma_{m,t} V_t^\alpha U_t^{1-\alpha}, \quad 0 < \alpha < 1,$$

where h_t , V_t and U_t denote hires, vacancies and unemployment, respectively. Also, $\sigma_{m,t}$ denotes a productivity parameter that can potentially capture variations in match efficiency. Dividing the matching function by the number of unemployed, we obtain the job finding rate $f_t \equiv h_t/U_t$, so that:

$$f_t = \sigma_{m,t} (V_t/U_t)^\alpha.$$

The simplest search and matching model assumes that the labor force is constant so that:

$$1 = l_t + U_t,$$

where l_t denotes employment and the labor force is assumed to be of size unity. The change in the number of people unemployed is given by:

$$U_{t+1} - U_t = (1 - \rho) l_t - f_t U_t,$$

where $(1 - \rho) l_t$ denotes the employed workers that separate into unemployment in period t and $f_t U_t$ is the number of unemployed workers who find jobs. In steady state, $U_{t+1} = U_t$, so that:

$$U_t = (1 - \rho) / (f_t + 1 - \rho).$$

Combining this expression with the definition of the finding rate and solving for V_t , we obtain:

$$V_t = \left[\frac{(1 - \rho)(1 - U_t)}{\sigma_{m,t} U_t^{1-\alpha}} \right]^{\frac{1}{\alpha}} \quad (6.1)$$

This equation clearly implies (i) a negative relationship between U_t and V_t and (ii) the only way that relationship can shift is with a change in the value of $\sigma_{m,t}$ or in the value of the other matching function parameter, α .²⁸ Results (i) and (ii) are apparently very robust, as they do not require taking a stand on many key relations in the overall economy. In the technical

²⁷We include this example for completeness. It can be found in other places, for example, Yashiv (2008).

²⁸In principle, a change in the separation rate, $1 - \rho$, could also have shifted the Beveridge curve during the Great Recession. This explanation does not work because the separation rate fell from an average level of 3.7 percent before the Great Recession to an average of 3.1 percent after 2009. These numbers were calculated using JOLTS data available at the BLS website.

appendix, we derive a similar result for our model, which also does not depend on most of our model details, such as the costs for arranging meetings between workers and firms, the determination of the value of a job, etc.

While the steady state Beveridge curve described in the previous paragraph may be useful for many purposes, it is misleading for interpreting data from the Great Recession, when the steady state condition, $U_{t+1} - U_t$, is far from being satisfied. To see this, note in Figure 9 that our model is able to account for the so-called shift in the Beveridge curve, even though the productivity parameter in our matching function is constant. The only difference between the analysis in Figure 9 and our model's steady state Beveridge curve is that we do not impose the $U_{t+1} - U_t = 0$ condition. Thus, according to our analysis the data on vacancies and unemployment present no reason to suppose that there has been a deterioration in match efficiency. No doubt such a deterioration has occurred to some extent, but it does not seem to be a first order feature of the Great Recession.

We conclude this section by noting that our model does not fully account for the Great Recession targets in the case of two variables. First, it does not capture the full magnitude of the collapse in investment in 2009. In the data, the maximal drop is a little less than 40 percent while in the model the drop is a little over 25 percent. After 2010 the model and data line up reasonably well with respect to investment. Second, the model does not account for the sharp drop in consumption relative to trend that began in 2012.

6.2. The Causes of the Great Recession

Figures 10 through 14 decompose the impact of the different shocks and the risky working capital channel on the economy in the post 2008Q3 period. We determine the role of a shock by setting that shock to its steady state value and redoing the simulations underlying Figure 8. The resulting decomposition is not additive because of the nonlinearities in the model.

Figure 10 displays the effect of the neutral technology shock on post-2008 simulations. For convenience, the solid line reproduces the corresponding solid line in Figure 8. The dashed line displays the behavior of the economy when neutral technology shock is shut down (i.e., $\varepsilon_t^p = 0$ in 2008Q3). Comparing the solid and dashed lines, we see that the neutral technology slowdown had a significant impact on inflation. Had it not been for the *decline* in neutral technology, there would have been substantial deflation, as predicted by very simple NK models that do not allow for a drop in technology during the ZLB period. The negative technology shock also pushes up output, investment, consumption, employment and the labor force. Abstracting from wealth effects on labor force participation, a fall in neutral technology raises marginal cost and, hence, inflation. In presence of the ZLB, the latter effect lowers the real interest rate, driving up aggregate spending and, hence, output and employment. In fact, the wealth effect of a negative technology shock does lead to an increase in the labor

force participation rate. Other things the same, this effect exerts downward pressure on the wage and, hence, on marginal cost. Evidently, this effect is outweighed by the direct effect of the negative technology shock, so that marginal costs rise.²⁹

Medium-sized DSGE models typically abstract from the working capital channel. A natural question is: how important is that channel in allowing our model to account for the moderate degree of inflation during the Great Recession? To answer that question, we redo the simulation underlying Figure 8, replacing (5.1) with (2.15). The results are displayed in Figure 11. We find that the risky working capital channel plays a very important role in allowing the model to account for the moderate decline in inflation that occurred during the Great Recession. In the presence of a risky working capital requirement, a higher interest rate due to a positive financial wedge shock directly raises firms' marginal cost. Other things equal, this rise leads to inflation. Gilchrist, Schoenle, Sim and Zakrajšek (2013) provide firm-level evidence consistent with the importance of our risky working capital channel. They find that firms with bad balance sheets raise prices relative to firms with good balance sheets. From our perspective, firms with bad balance sheets face a very high cost of working capital and therefore, high marginal costs.

Taken together, the negative technology shocks and the risky working capital channel explain the relatively modest disinflation that occurred during the Great Recession. Essentially they exerted countervailing pressure on the disinflationary forces that were operative during the Great Recession. The output effects of the risky working capital channel are much weaker than those of the neutral technology shocks. In part this reflects the fact that the working capital risk channel works via the financial wedge shocks and these are much less persistent than the technology shocks.

Figures 12 and 13 report the effects of the financial and consumption wedges, respectively. The latter plays an important role in driving the economy into the ZLB and has substantial effects on real quantities and inflation. The fact that the nominal interest rate remains at zero after 2011 when there is no consumption wedge reflects our specification of monetary policy. The financial wedge has a relatively small impact on inflation and on the interest rate, but it has an enormous impact on real quantities. For example, the financial wedge is overwhelmingly the most important shock for investment. Notice that the model attributes the substantial drop in the labor force participation rate almost entirely to the consumption and financial wedges. This reflects that these wedges lead to a sharp deterioration in labor market conditions: drops in the job vacancy and finding rates and in the real wage. We do not think these wedge shocks were important in the pre-2008 period. In this way, the model is consistent with the fact that labor force participation rates are not very cyclical during

²⁹There are other forces at work in the ZLB that can cause a persistent decrease in technology to generate more inflation than a transitory decrease.

normal recessions, while being very cyclical during the Great Recession.³⁰

We now turn to Figure 14, which analyzes the role of government consumption in the Great Recession. Government consumption passes through two phases (see Figure 7). The first phase corresponds to the expansion associated with the American Recovery and Reinvestment Act of 2009. The second phase involves a contraction that began at the start of 2011. The first phase involves a maximum rise of 3 percent in government consumption (i.e., 0.6 percent relative to steady state GDP) and a maximum rise of 1.4 percent in GDP. This implies a maximum government consumption multiplier of $1.4/0.6$ or 2.17. In the second phase the decline in government spending is much more substantial, falling a maximum of nearly 10 percent, or 2 percent relative to steady state GDP. At the same time, the resulting drop in GDP is about 1.5 percent (see Figure 14). So, in the second phase, the government spending multiplier is only $1.5/2$ or 0.75. In light of this result, it is difficult to attribute the long duration of the Great Recession to the recent decline in government consumption.

The second phase findings may at first seem inconsistent with existing analyses, which suggest that the government consumption multiplier may be very large in the ZLB. Indeed, Christiano, Eichenbaum and Rebelo (2011) show that a rise in government consumption that is expected to not extend beyond the ZLB has a large multiplier effect. But, they also show that a rise in government consumption that is expected to extend beyond the ZLB has a relatively small multiplier effect. The intuition for this is straightforward. An increase in spending after the ZLB ceases to bind has no direct impact on spending in the ZLB. But, it has a negative impact on household consumption in the ZLB because of the negative wealth effects associated with the (lump-sum) taxes required to finance the increase in government spending. A feature of our simulations is that the increase in government consumption in the first phase is never expected by agents to persist beyond the ZLB. In the second phase the decrease in government consumption is expected to persist beyond the end of the ZLB.

Figure 15 displays the impact of the switch to forward guidance in 2011. The dashed line represents the model simulation with all shocks, when the Taylor rule is in place throughout the period. The figure indicates that without forward guidance the Fed would have started raising the interest rate in 2012. By keeping the interest rate at zero, the monetary authority caused output to be 2 percent higher and the unemployment rate to be one percentage point lower. Interestingly, this relationship is consistent with Okun's law.

We also examined the role, in our simulations, of the unexpected extension in the duration of the consumption wedge, Δ_t^b , in 2012Q3. To save space, we simply report the key results. The extension has two important effects. First, it helps the model account for the slowdown in consumption that occurred around end-2011 (see panel 2,4 of Figure 4). Second, it helps

³⁰See Erceg and Levin (2013), for an analysis which reaches a qualitatively similar conclusion using a small scale, calibrated model.

the model account for fact that inflation remains low for so long.

7. Conclusion

This paper argues that the bulk of movements in aggregate real economic activity during the Great Recession were due to financial frictions interacting with the ZLB. We reach this conclusion looking at the data through the lens of a New Keynesian model in which firms face moderate degrees of price rigidities and no nominal rigidities in the wage setting process. Our model does a good job of accounting for the joint behavior of labor and goods markets, as well as inflation, during the Great Recession. According to the model the observed fall in TFP relative to trend and the rise in the cost of working capital played key roles in accounting for the small size of the drop in inflation that occurred during the Great Recession.

References

- [1] Aguiar, Mark, Erik Hurst and Loukas Karabarbounis, 2012, “Time Use During the Great Recession,” *American Economic Review*, forthcoming.
- [2] Altig, David, Lawrence Christiano, Martin Eichenbaum and Jesper Linde, 2011, “Firm-Specific Capital, Nominal Rigidities and the Business Cycle,” *Review of Economic Dynamics*, Elsevier for the Society for Economic Dynamics, vol. 14(2), pages 225-247, April.
- [3] Binmore, Ken, Ariel Rubinstein, and Asher Wolinsky, 1986, “The Nash Bargaining Solution in Economic Modelling,” *RAND Journal of Economics*, 17(2), pp. 176-88.
- [4] Bigio, Saki, 2013, “Endogenous Liquidity and the Business Cycle,” manuscript, Columbia Business School.
- [5] Boot, J. and W. Feibes and J. Lisman, 1967, “Further Methods of Derivation of Quarterly Figures from Annual Data,” *Applied Statistics* 16, pp. 65 – 75.
- [6] Campbell, Jeffrey R., Charles L. Evans, Jonas D. M. Fisher, and Alejandro Justiniano, 2012, “Macroeconomic Effects of Federal Reserve Forward Guidance,” *Brookings Papers on Economic Activity*, Economic Studies Program, The Brookings Institution, vol. 44, Spring, pages 1-80.
- [7] Christiano, Lawrence J. and Joshua M. Davis, 2006, “Two Flaws In Business Cycle Accounting,” NBER Working Paper No. 12647, National Bureau of Economic Research, Inc.
- [8] Christiano, Lawrence J., Roberto Motto and Massimo Rostagno, 2003, “The Great Depression and the Friedman-Schwartz Hypothesis,” *Journal of Money, Credit and Banking*, Vol. 35, No. 6, Part 2: Recent Developments in Monetary Economics, December, pp. 1119-1197.
- [9] Christiano, Lawrence J., and Martin S. Eichenbaum, 1990, “Unit Roots in GNP: Do We Know and Do We Care?”, *Carnegie-Rochester Conference Series on Public Policy* 32, pp. 7-62.

- [10] Christiano, Lawrence J., Martin S. Eichenbaum and Charles L. Evans, 2005, “Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy,” *Journal of Political Economy*, 113(1), pp. 1-45.
- [11] Christiano, Lawrence J., Martin S. Eichenbaum and Sergio Rebelo, 2011, “When Is the Government Spending Multiplier Large?,” *Journal of Political Economy*, University of Chicago Press, vol. 119(1), pages 78 - 121.
- [12] Christiano, Lawrence J., Martin S. Eichenbaum and Mathias Trabandt, 2013, “Unemployment and Business Cycles,” NBER Working Paper No. 19265, National Bureau of Economic Research, Inc.
- [13] Christiano, Lawrence J., Martin S. Eichenbaum and Robert Vigfusson, 2007. “Assessing Structural VARs,” NBER Chapters, in: *NBER Macroeconomics Annual 2006*, Volume 21, pages 1-106 National Bureau of Economic Research, Inc.
- [14] Christiano, Lawrence J., Mathias Trabandt and Karl Walentin, 2011, “DSGE Models for Monetary Policy Analysis,” in Benjamin M. Friedman, and Michael Woodford, editors: *Handbook of Monetary Economics*, Vol. 3A, The Netherlands: North-Holland.
- [15] Christiano, Lawrence J., Mathias Trabandt and Karl Walentin, 2012, “Involuntary Unemployment and the Business Cycle,” manuscript, Northwestern University, 2012.
- [16] Davis, Steven, J., R. Jason Faberman, and John C. Haltiwanger, 2012, “Recruiting Intensity During and After the Great Recession: National and Industry Evidence,” *American Economic Review: Papers and Proceedings*, vol. 102, no. 3, May.
- [17] den Haan, Wouter, Garey Ramey, and Joel Watson, 2000, “Job Destruction and Propagation of Shocks,” *American Economic Review*, 90(3), pp. 482-98.
- [18] Del Negro, Marco, Marc P. Giannoni and Frank Schorfheide, 2014, “Inflation in the Great Recession and New Keynesian Models,” Federal Reserve Bank of New York Staff Report no. 618, January.
- [19] Dupor, Bill and Rong Li, 2012, “The 2009 Recovery Act and the Expected Inflation Channel of Government Spending,” Federal Reserve Bank of St. Louis Working Paper No. 2013-026A.
- [20] Diamond, Peter A., 1982, “Aggregate Demand Management in Search Equilibrium,” *Journal of Political Economy*, 90(5), pp. 881-894.
- [21] Diamond, Peter A., 2010, “Unemployment, Vacancies and Wages,” Nobel Prize lecture, December 8.
- [22] Edge, Rochelle M., Thomas Laubach, John C. Williams, 2007, “Imperfect credibility and inflation persistence,” *Journal of Monetary Economics* 54, 2421–2438.
- [23] Eichenbaum, Martin, Nir Jaimovich and Sergio Rebelo, 2011, “Reference Prices and Nominal Rigidities,” *American Economic Review*, 101(1), pp. 242-272.
- [24] Eggertsson, Gauti B. and Paul Krugman, 2012, “Debt, Deleveraging, and the Liquidity Trap: A Fisher-Minsky-Koo Approach,” *The Quarterly Journal of Economics*, Oxford University Press, vol. 127(3), pages 1469-1513.
- [25] Eggertsson, Gauti and Michael Woodford, 2003, “The zero interest-rate bound and optimal monetary policy,” *Brookings Papers on Economic Activity*, Economic Studies Program, The Brookings Institution, vol. 34(1), pages 139-235.

- [26] Erceg, Christopher J., and Andrew T. Levin, 2003, “Imperfect credibility and inflation persistence,” *Journal of Monetary Economics*, 50, 915–944.
- [27] Erceg, Christopher J., and Andrew T. Levin, 2013, “Labor Force Participation and Monetary Policy in the Wake of the Great Recession”, CEPR Discussion Paper No. DP9668, September.
- [28] Fair, Ray C., and John Taylor, 1983, “Solution and Maximum Likelihood Estimation of Dynamic Nonlinear Rational Expectations Models,” *Econometrica*, vol. 51, no. 4, July, pp. 1169-1185.
- [29] Fernald, John, 2012, “A Quarterly, Utilization-Adjusted Series on Total Factor Productivity,” Federal Reserve Bank of San Francisco Working Paper 2012-19.
- [30] Fisher, Jonas, 2006, “The Dynamic Effects of Neutral and Investment-Specific Technology Shocks,” *Journal of Political Economy*, 114(3), pp. 413-451.
- [31] Fisher, Jonas, 2014, “On the Structural Interpretation of the Smets-Wouters ‘Risk Premium’ Shock,” manuscript, Federal Reserve Bank of Chicago.
- [32] Gilchrist, Simon and Egon Zakrajšek, 2012, Credit Spreads and Business Cycle Fluctuations, *American Economic Review*, 102(4): 1692–1720.
- [33] Gilchrist, Simon, Raphael, Schoenle, Jae Sim and Egon Zakrajšek (2013), “Inflation Dynamics during the Financial Crisis,” manuscript, December.
- [34] Gust, Christopher J., J. David Lopez-Salido, and Matthew E. Smith (2013), “The Empirical Implications of the Interest-Rate Lower Bound,” Finance and Economics Discussion Series 2012-83. Board of Governors of the Federal Reserve System.
- [35] Hall, Robert E., 2005, “Employment Fluctuations with Equilibrium Wage Stickiness,” *American Economic Review*, 95(1), pp. 50-65.
- [36] Hall, Robert E. and Paul R. Milgrom, 2008, “The Limited Influence of Unemployment on the Wage Bargain,” *The American Economic Review*, 98(4), pp. 1653-1674.
- [37] Hall, Robert E., 2011, “The Long Slump,” *American Economic Review*, 101, 431-469.
- [38] Justiniano, Alejandro, Giorgio E. Primiceri, and Andrea Tambalotti, 2010, “Investment Shocks and Business Cycles,” *Journal of Monetary Economics* 57(2), pp. 132-145.
- [39] Kapon, Samuel and Joseph Tracy, 2014, “A Mis-Leading Labor Market Indicator,” Federal Reserve Bank of New York, <http://libertystreeteconomics.newyorkfed.org/2014/02/a-mis-leading-labor-market-indicator.html>, February 3, 2014.
- [40] Klenow, Pete and Benjamin Malin, 2011, “Microeconomic Evidence on Price-Setting,” in the *Handbook of Monetary Economics* 3A, editors: B. Friedman and M. Woodford, Elsevier, pp. 231-284.
- [41] Kocherlakota, Narayana, 2010, “Inside the FOMC,” Federal Reserve Bank of Minneapolis, http://www.minneapolisfed.org/news_events/pres/speech_display.cfm?id=4525.
- [42] Krugman Paul, 2014, “Demography and Employment,” <http://krugman.blogs.nytimes.com/2014/02/03/demography-and-employment-wonkish>. February 3, 2014.

- [43] Lorenzoni, Guido and Veronica Guerrieri, (2012), “Credit Crises, Precautionary Savings and the Liquidity Trap,” *Quarterly Journal of Economics*.
- [44] Mortensen, Dale T., 1982, “Property Rights and Efficiency in Mating, Racing, and Related Games,” *American Economic Review*, 72(5), pp. 968–79.
- [45] Paciello, Luigi, (2011), “Does Inflation Adjust Faster to Aggregate Technology Shocks than to Monetary Policy Shocks,” *Journal of Money, Credit and Banking*, 43(8).
- [46] Pissarides, Christopher A., 1985, “Short-Run Equilibrium Dynamics of Unemployment, Vacancies, and Real Wages,” *American Economic Review*, 75(4), pp. 676-90.
- [47] Prescott, Edward, C., 1986, “Theory Ahead of Business Cycle Measurement,” Federal Reserve Bank of Minneapolis *Quarterly Review*, Fall, vol. 10, no. 4.
- [48] Quah, Danny, 1990, “Permanent and Transitory Movements in Labor Income: An Explanation for ‘Excess Smoothness’ in Consumption,” *Journal of Political Economy*, Vol. 98, No. 3, June, pp. 449-475.
- [49] Ravenna, Federico and Carl Walsh, 2008, “Vacancies, Unemployment, and the Phillips Curve,” *European Economic Review*, 52, pp. 1494-1521.
- [50] Reifschneider, David, William Wascher and David Wilcox, 2013, “Aggregate Supply in the United States: Recent Developments and Implications for the Conduct of Monetary Policy,” Federal Reserve Board *Finance and Economics Discussion Series* No. 2013-77, December.
- [51] Rubinstein, Ariel, 1982, “Perfect Equilibrium in a Bargaining Model,” *Econometrica*, 50(1), pp. 97-109.
- [52] Schmitt-Grohé, Stephanie, and Martin Uribe, 2012, “What’s News in Business Cycles?,” *Econometrica*, 80, pp. 2733-2764.
- [53] Silva, J. and M. Toledo, 2009, “Labor Turnover Costs and the Cyclical Behavior of Vacancies and Unemployment,” *Macroeconomic Dynamics*, 13, Supplement 1.
- [54] Shimer, Robert, 2005, “The Cyclical Behavior of Equilibrium Unemployment and Vacancies,” *The American Economic Review*, 95(1), pp. 25-49.
- [55] Shimer, Robert, 2012, “Reassessing the Ins and Outs of Unemployment,” *Review of Economic Dynamics*, 15(2), pp. 127-48.
- [56] Smets, Frank and Rafael Wouters, 2007, “Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach,” *American Economic Review*, 97(3), pp. 586-606.
- [57] Sullivan, Daniel, 2013, “Trends In Labor Force Participation, Presentation, Federal Reserve Bank of Chicago,” https://www.chicagofed.org/digital_assets/others/people/research_resources/sullivan_daniel/sullivan_cbo_labor_force.pdf.
- [58] Swanson, Eric T. and John C. Williams, 2013, ‘Measuring the Effect of the Zero Lower Bound On Medium- and Longer-Term Interest Rates,’ forthcoming, *American Economic Review*.
- [59] Yashiv, Eran, 2008, “The Beveridge Curve,” *The New Palgrave Dictionary of Economics*, Second Edition, Edited by Steven N. Durlauf and Lawrence E. Blume.

Table 1: Non-Estimated Model Parameters and Calibrated Variables

Parameter	Value	Description
<i>Panel A: Parameters</i>		
δ_K	0.025	Depreciation rate of physical capital
β	0.9968	Discount factor
ρ	0.9	Job survival probability
M	60	Maximum bargaining rounds per quarter
$(1 - \chi)^{-1}$	3	Elasticity of substitution market and home consumption
$100 (\pi^A - 1)$	2	Annual net inflation rate target
$400 \ln(\mu_\Phi)$	1.7	Annual output per capita growth rate
$400 \ln(\mu_\Phi \times \mu_\Psi)$	2.9	Annual investment per capita growth rate
<i>Panel B: Steady State Values</i>		
<i>profits</i>	0	Intermediate goods producers profits
Q	0.7	Vacancy filling rate
u	0.055	Unemployment rate
L	0.67	Labor force to population ratio
G/Y	0.2	Government consumption to gross output ratio

Table 2: Priors and Posteriors of Model Parameters

	Prior			Posterior	
	Distribution	Mean,Std.	Mode	Std.	
<i>Price Setting Parameters</i>					
Price Stickiness	ξ	Beta	0.66,0.15	0.737	0.022
Price Markup Parameter	λ	Gamma	1.20,0.05	1.322	0.042
<i>Monetary Authority Parameters</i>					
Taylor Rule: Interest Rate Smoothing	ρ_R	Beta	0.75,0.15	0.792	0.015
Taylor Rule: Inflation Coefficient	r_π	Gamma	1.70,0.10	1.672	0.093
Taylor Rule: GDP Gap Coefficient	r_y	Gamma	0.01,0.01	0.012	0.007
Taylor Rule: GDP Growth Coefficient	$r_{\Delta y}$	Gamma	0.20,0.05	0.184	0.048
<i>Preferences and Technology</i>					
Market and Home Consumption Habit	b	Beta	0.50,0.15	0.889	0.013
Capacity Utilization Adjustment Cost	σ_a	Gamma	0.50,0.30	0.036	0.028
Investment Adjustment Cost	S''	Gamma	8.00,2.00	12.07	1.672
Capital Share	α	Beta	0.33,0.03	0.247	0.018
Technology Diffusion	θ	Beta	0.50,0.20	0.115	0.024
<i>Labor Market Parameters</i>					
Probability of Bargaining Breakup	100δ	Gamma	0.50,0.20	0.051	0.015
Replacement Ratio	D/w	Beta	0.40,0.10	0.194	0.058
Hiring Cost to Output Ratio	s_l	Gamma	1.00,0.30	0.474	0.146
Labor Force Adjustment Cost	ϕ_L	Gamma	100,50.0	134.7	28.34
Unemployed Share in Home Production	α_{cH}	Beta	0.03,0.01	0.015	0.005
Probability of Staying in Labor Force	s	Beta	0.85,0.05	0.816	0.060
Matching Function Parameter	σ	Beta	0.50,0.10	0.506	0.039
<i>Shocks</i>					
Standard Deviation Monetary Policy Shock	$400\sigma_R$	Gamma	0.65,0.05	0.650	0.035
AR(1) Persistent Component of Neutral Techn.	ρ_P	Gamma	0.50,0.07	0.792	0.041
Stdev. Persistent Component of Neutral Techn.	$100\sigma_P$	Gamma	0.15,0.04	0.037	0.004
AR(1) Transitory Component of Neutral Techn.	ρ_T	Beta	0.75,0.07	0.927	0.033
Stdev. Ratio Transitory and Perm. Neutral Techn.	σ_T/σ_P	Gamma	6.00,0.45	4.916	0.403
AR(1) Investment Technology	ρ_Ψ	Beta	0.75,0.10	0.714	0.056
Standard Deviation Investment Technology Shk.	$100\sigma_\Psi$	Gamma	0.10,0.05	0.114	0.017

Notes: s_l denotes the steady state hiring to gross output ratio (in percent).

Table 3: Model Steady States and Implied Parameters

Variable	At Estimated Posterior Mode	Description
K/Y	7.01	Capital to gross output ratio (quarterly)
C/Y	0.57	Market consumption to gross output ratio
I/Y	0.22	Investment to gross output ratio
l	0.63	Employment to population ratio
R	1.0125	Gross nominal interest rate (quarterly)
R^{real}	1.0075	Gross real interest rate (quarterly)
mc	0.76	Marginal cost (inverse markup)
σ_b	0.036	Capacity utilization cost parameter
Y	0.83	Gross output
ϕ/Y	0.32	Fixed cost to gross output ratio
σ_m	0.66	Level parameter in matching function
f	0.63	Job finding rate
ϑ	0.98	Marginal revenue of wholesaler
x	0.1	Hiring rate
J	0.06	Value of firm
V	197.1	Value of work
U	193.3	Value of unemployment
N	185.1	Value of not being in the labor force
v	0.18	Vacancy rate
e	0.06	Probability of leaving non-participation
ω	0.47	Home consumption weight in utility
C^H	0.31	Home consumption
w	0.97	Real wage
$\gamma/(\vartheta/M)$	0.81	Counteroffer costs as share of daily revenue

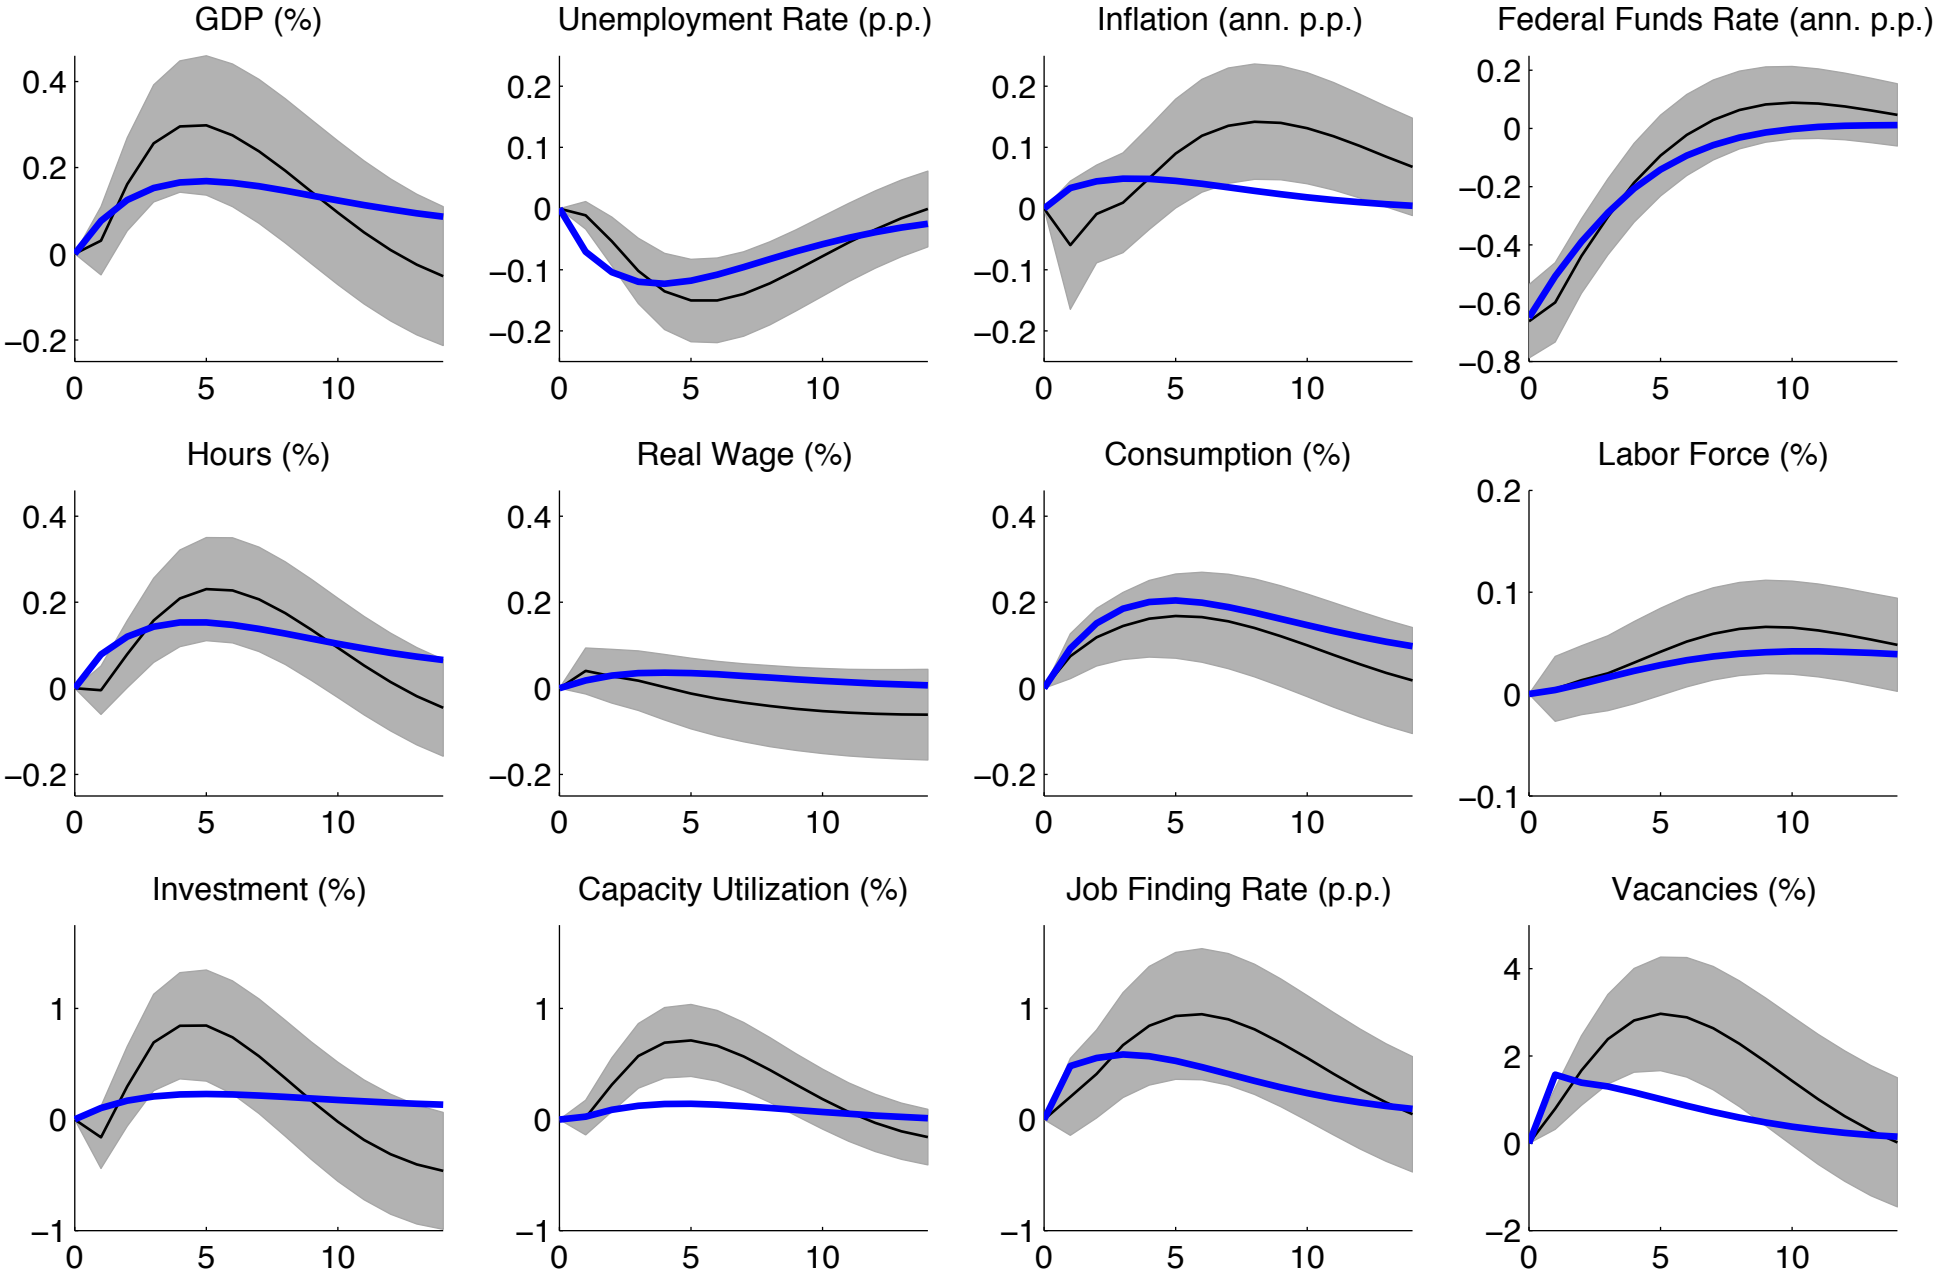
Table 4: Labor Market Status Transition Probabilities

		To					
		E		U		N	
$From$	E	Data	Model	Data	Model	Data	Model
	U	0.89	0.95	0.03	0.03	0.08	0.02
	N	0.46	0.52	0.17	0.30	0.37	0.18
	N	0.14	0.04	0.05	0.02	0.81	0.94

Notes: Transition probabilities between employment (E), unemployment (U) and non-participation (N). Model refers to transition probabilities in steady state at estimated parameter values. Data are based on Current Population Survey. We take the average of monthly transition probabilities from January 1990 to December 2013. To convert from monthly to quarterly frequency we take the average monthly transition probability matrix to the power of three.

Figure 1: Impulse Responses to an Expansionary Monetary Policy Shock


■ VAR 95% — VAR Mean — Model



Notes: x-axis in quarters.

Figure 2: Impulse Responses to Negative Innovation in Neutral Technology


■ VAR 95% — VAR Mean — Model



Notes: x-axis in quarters.

Figure 3: Impulse Responses to Negative Innovation in Investment-Specific Technology

VAR 95% VAR Mean Model



Notes: x-axis in quarters.

Figure 4: The Great Recession in the U.S.

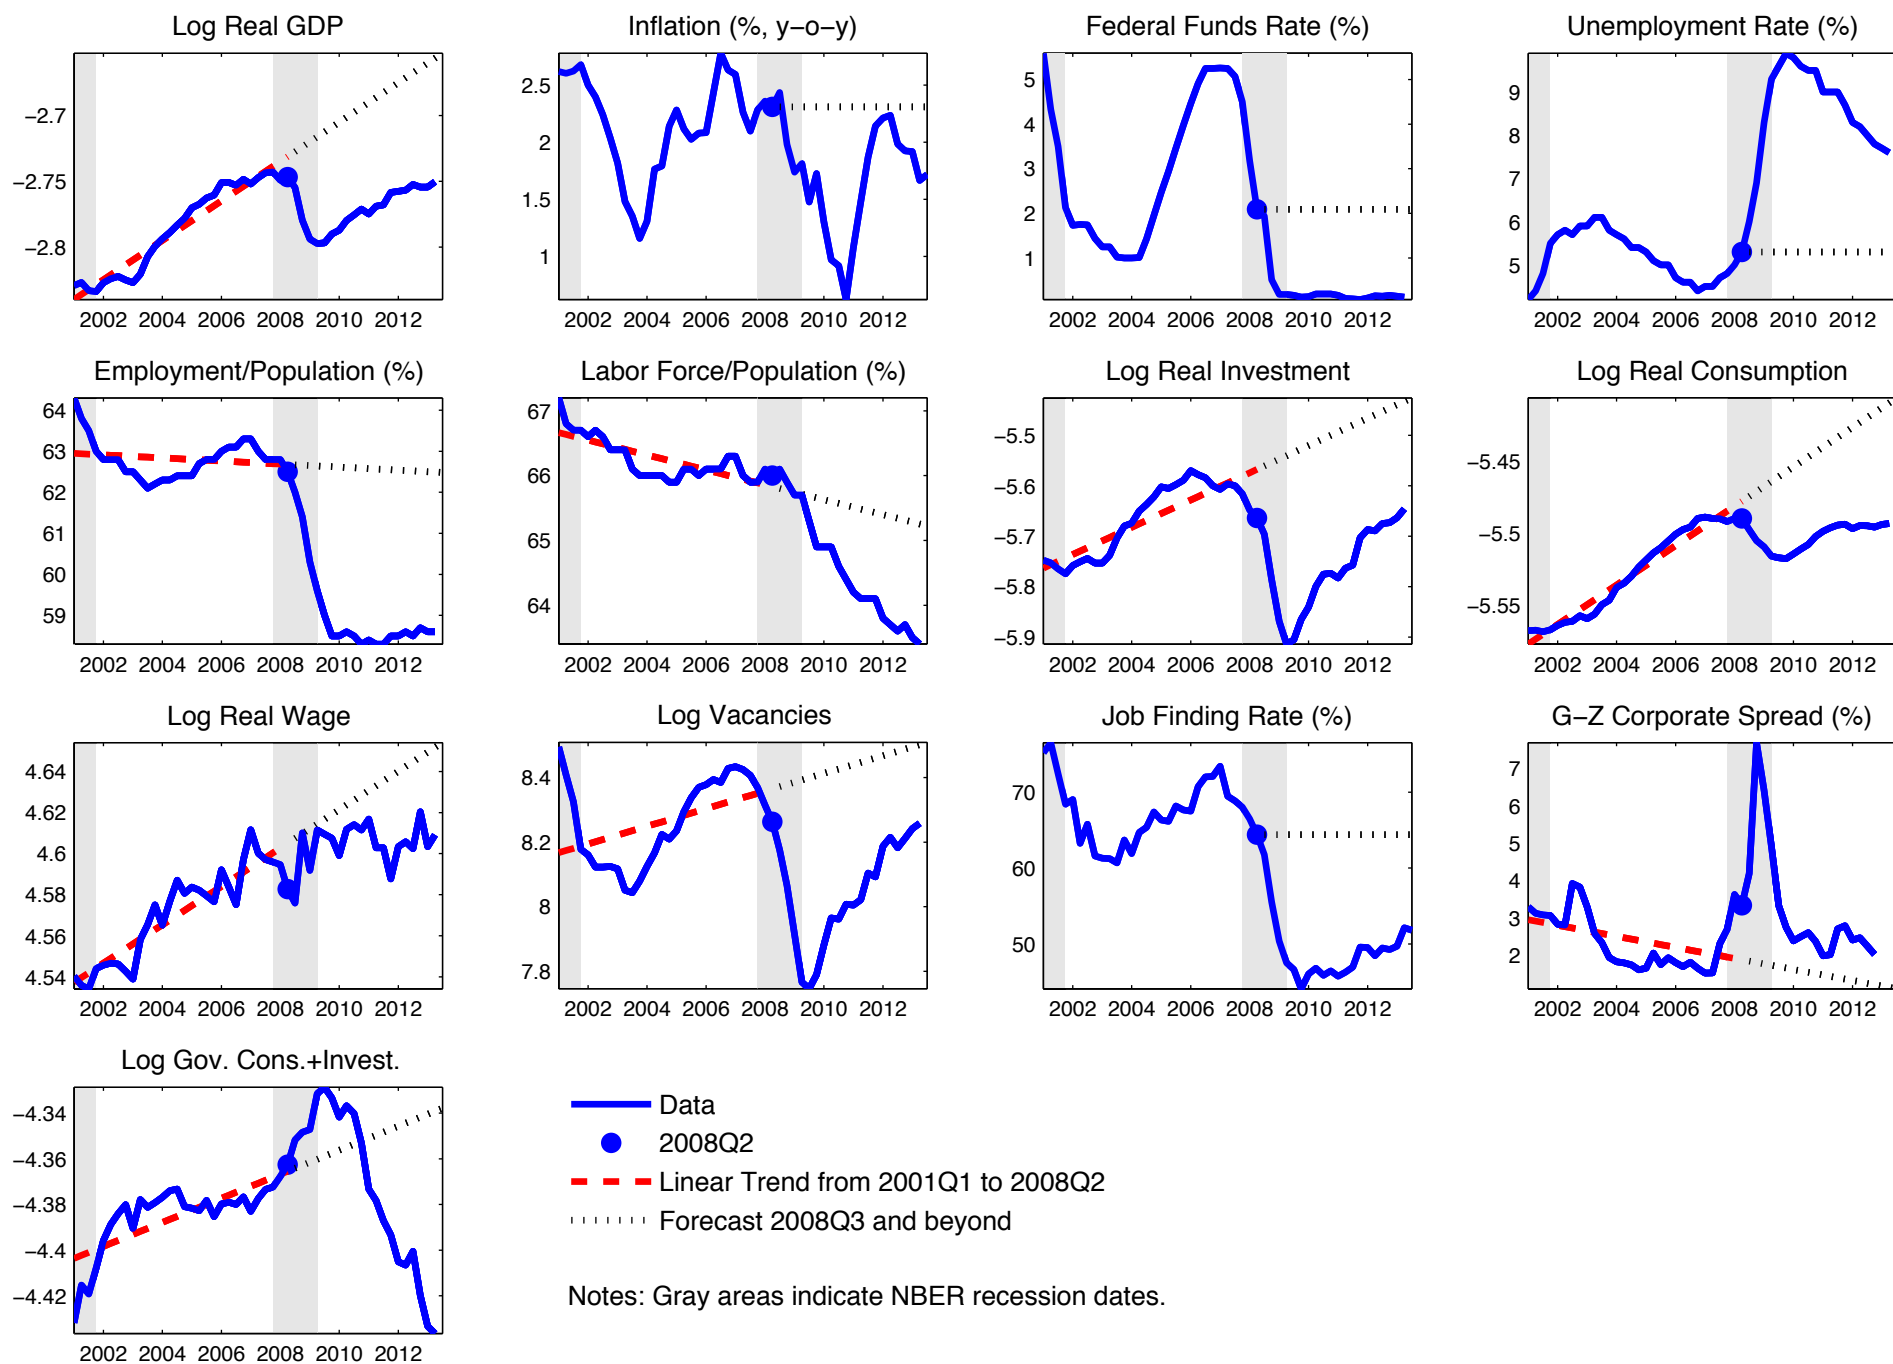




Figure 5: Measures of Total Factor Productivity (TFP): 2001 to 2013

Log TFP



Log TFP



Notes: Linear trend from 2001Q1–2008Q2 (dashed–dotted). Forecast 2008Q3 and beyond based on linear trend (dotted).

TFP (% Deviation from Trend)





Figure 6: Measures of Total Factor Productivity: 1982– 2013

Log Total Factor Productivity (TFP): 1982–2013



Notes: Linear trend from 1982Q1–2008Q2 (dashed–dotted). Forecast from 2008Q3 based on linear trend (dotted).

TFP (% Deviation from Trend)









Figure 7: The U.S. Great Recession: Exogenous Variables



Notes: Data are the differences between raw data and forecasts, see Figure 4.

Figure 8: The U.S. Great Recession: Data vs. Model



Notes: Data are the differences between raw data and forecasts, see Figure 4.

Figure 9: Beveridge Curve: Data vs. Model




Figure 10: The U.S. Great Recession: Effects of Neutral Technology




Figure 11: The U.S. Great Recession: Effects of Spread on Working Capital

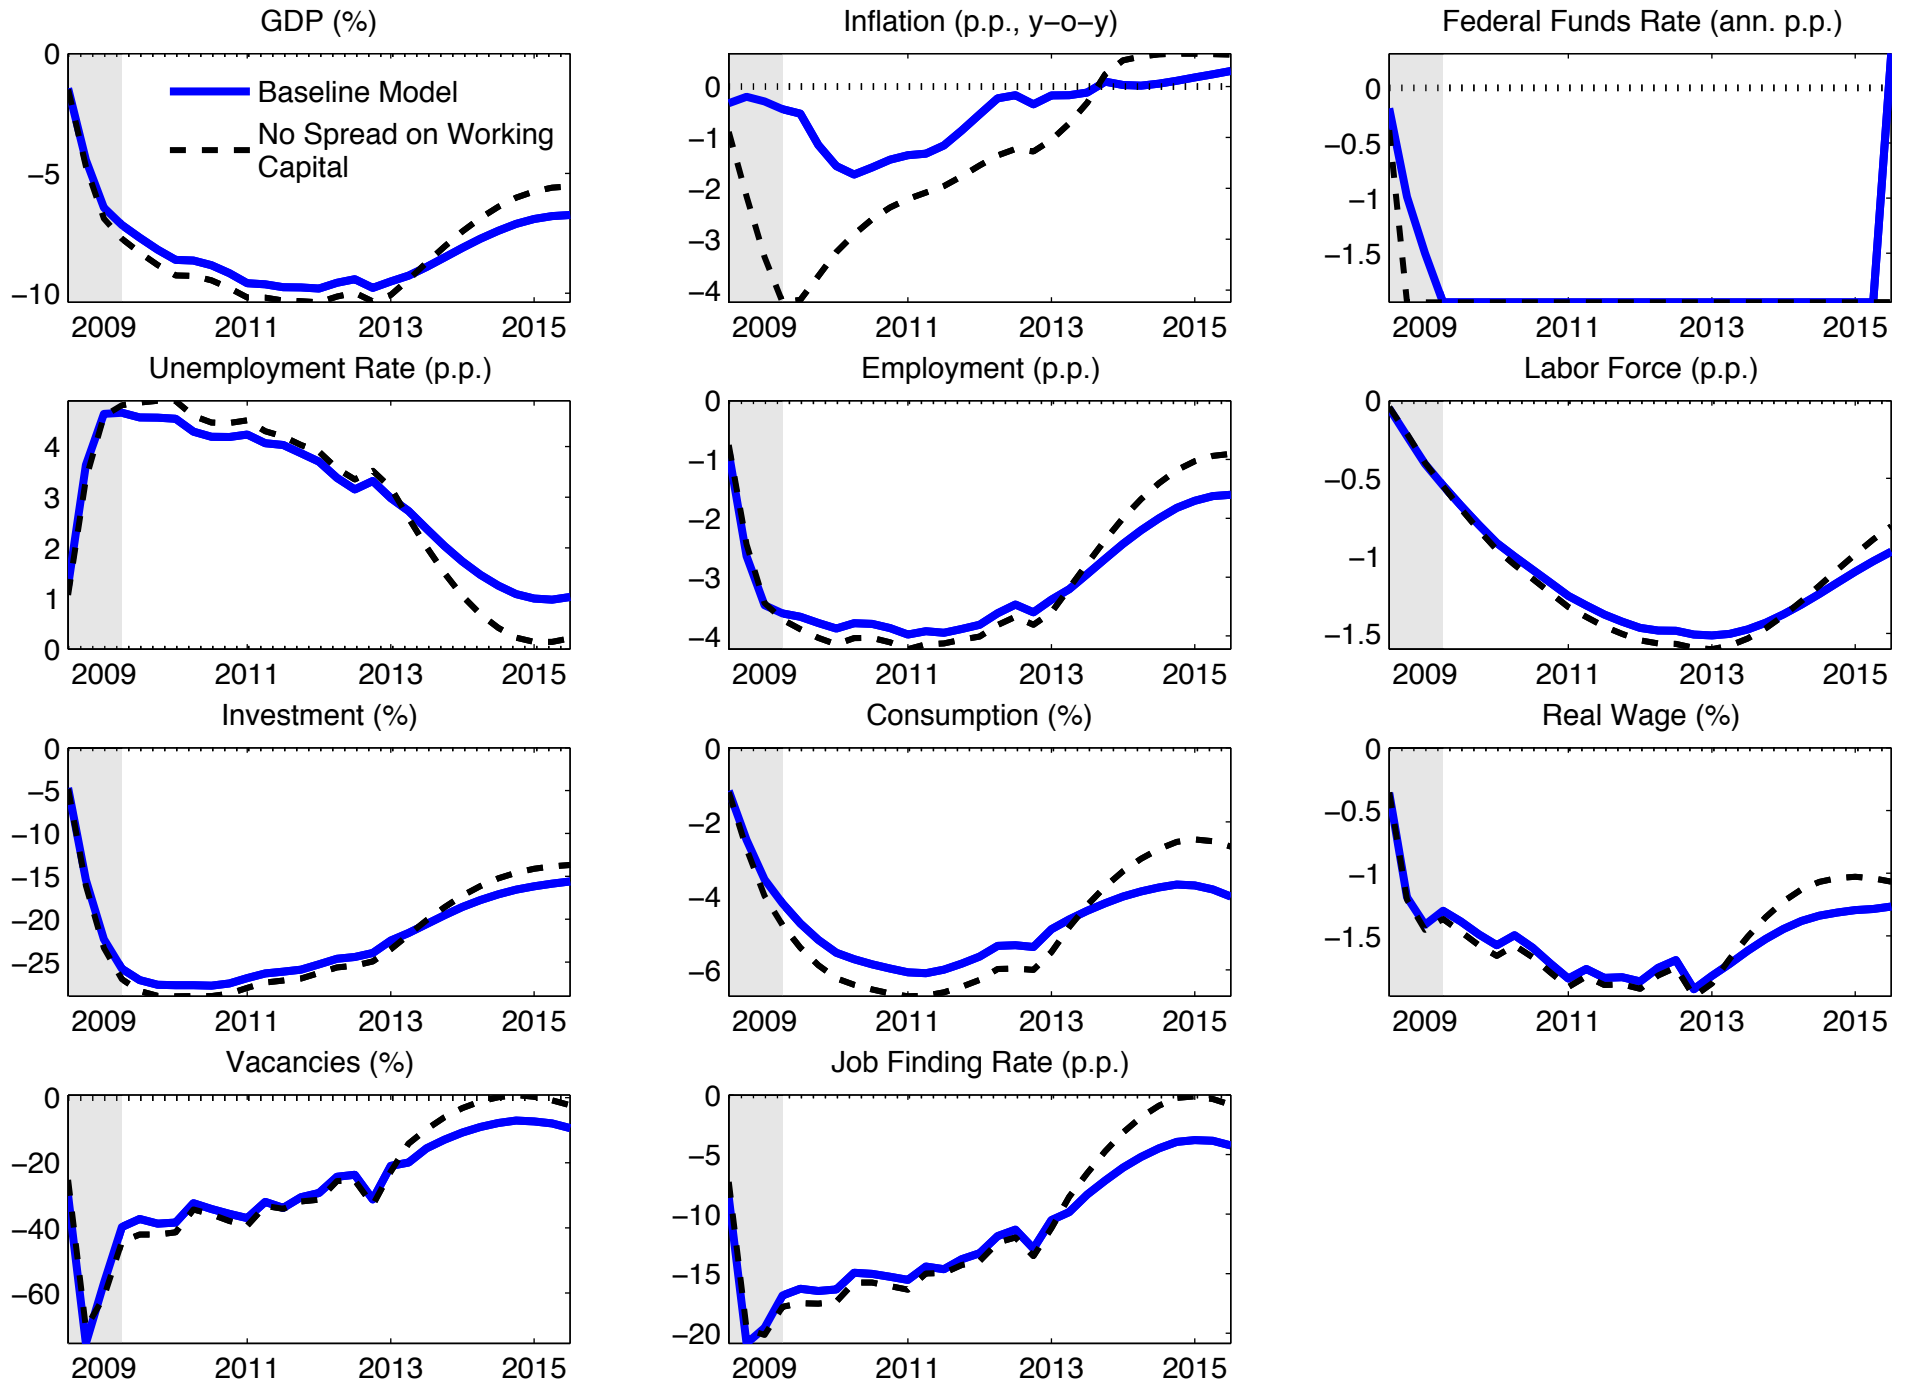


Figure 12: The U.S. Great Recession: Effects of Financial Wedge

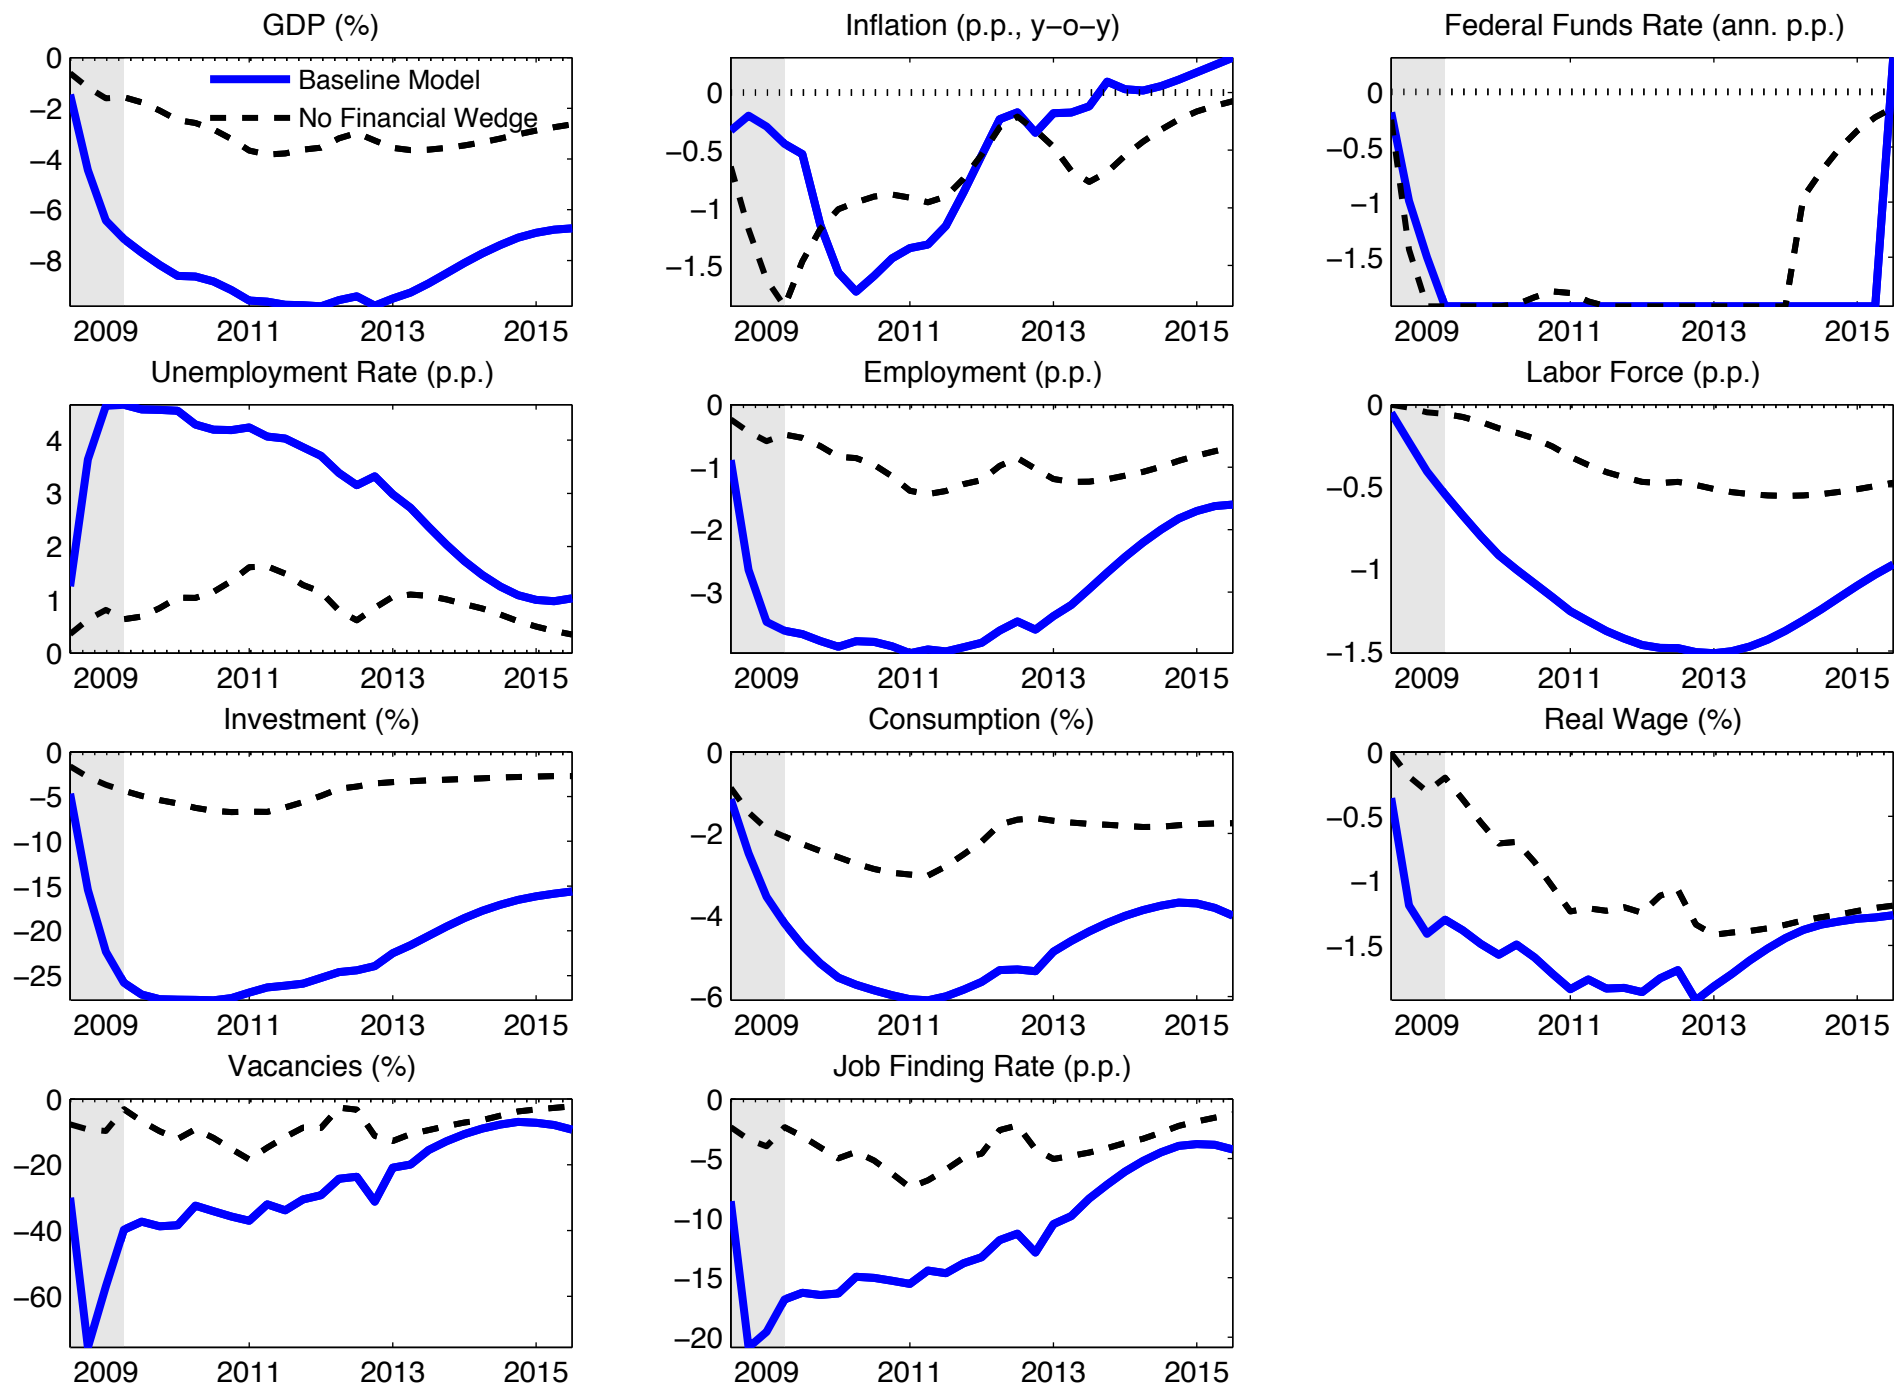


Figure 13: The U.S. Great Recession: Effects of Consumption Wedge

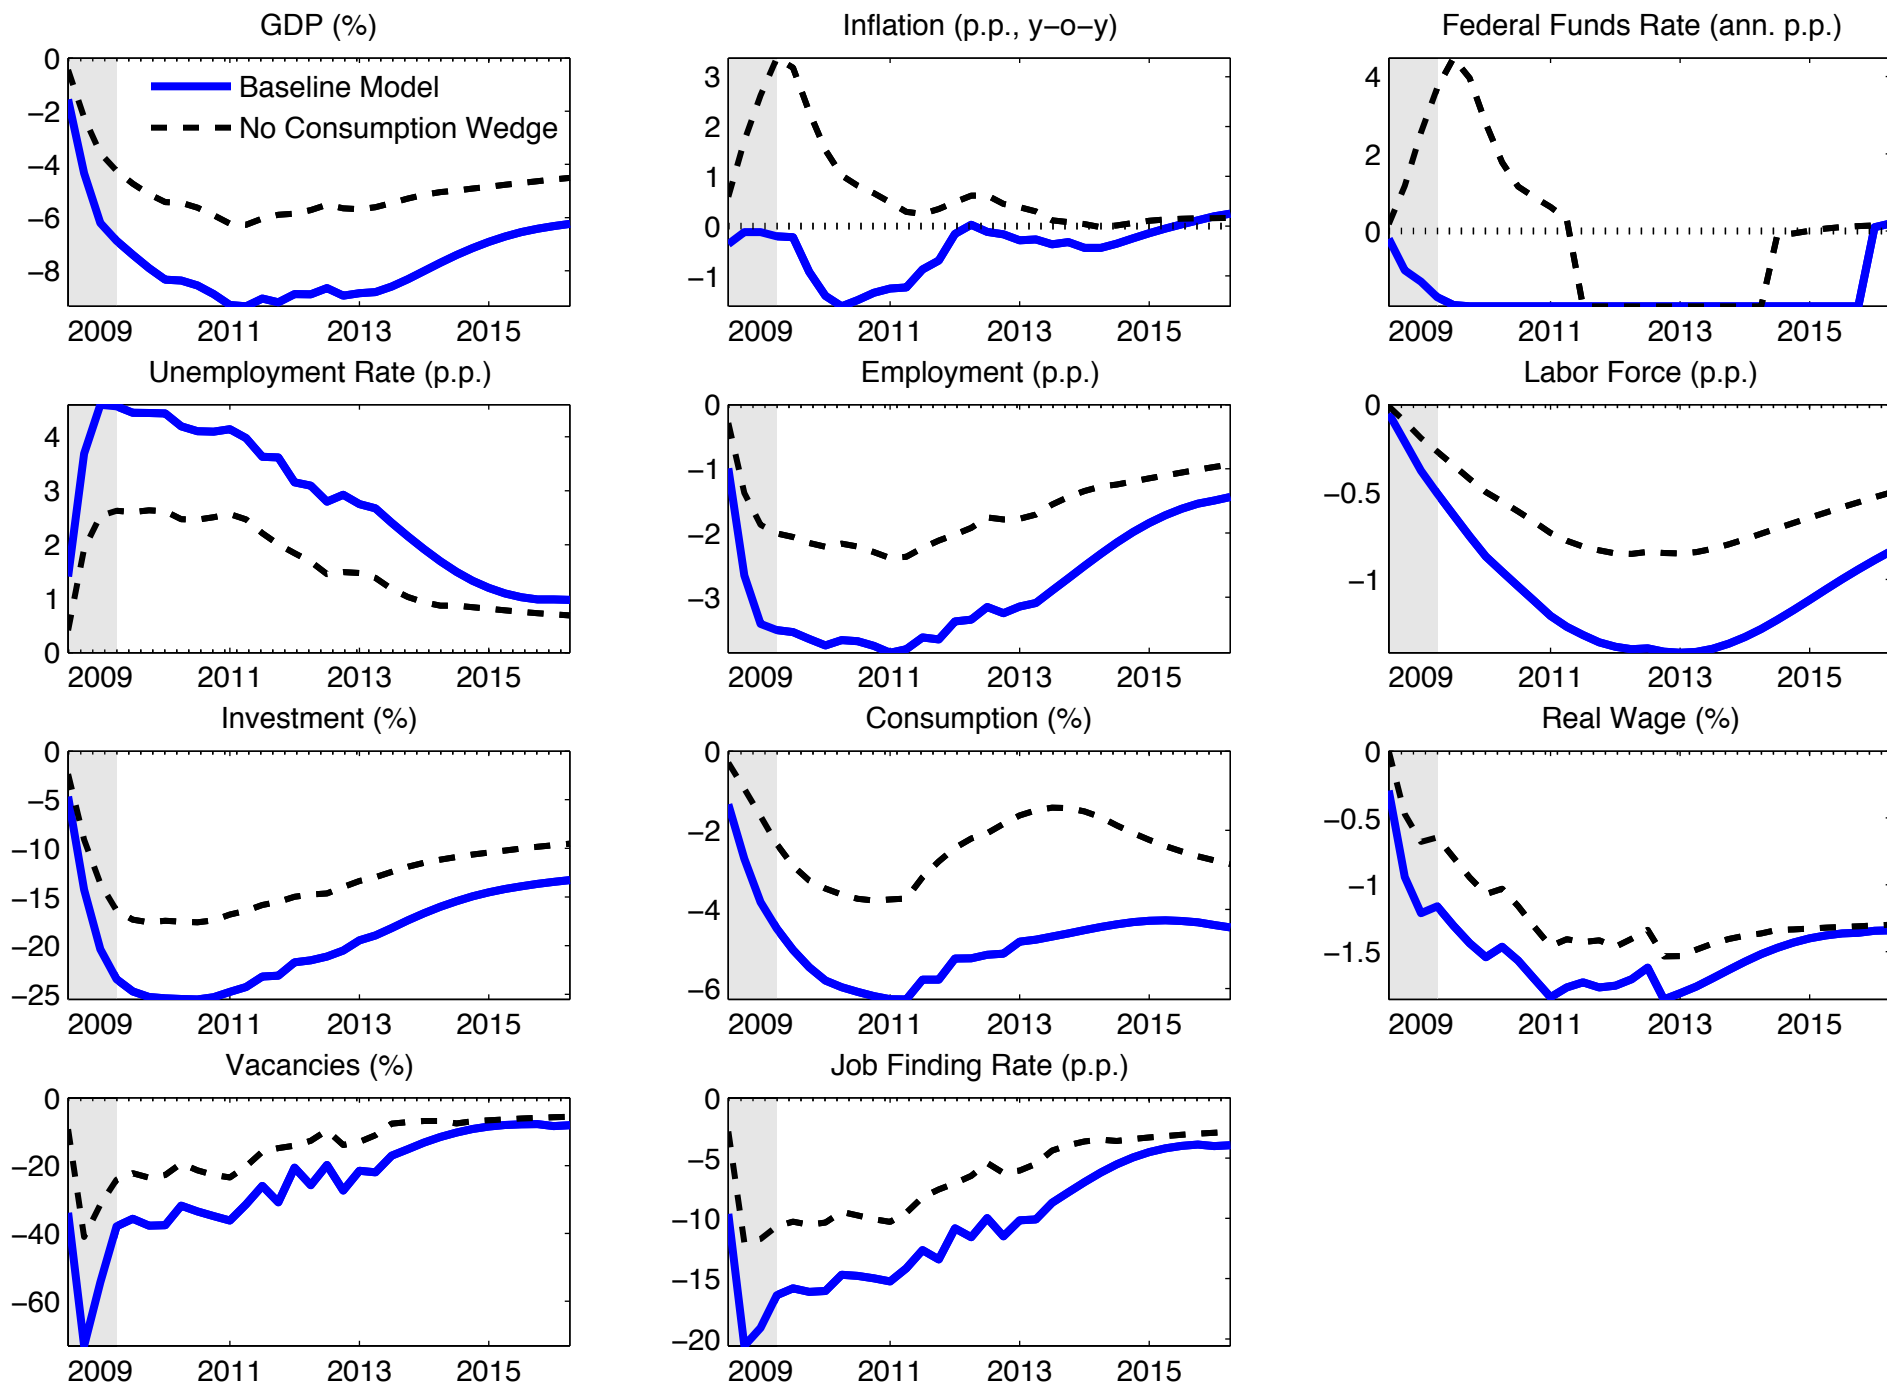


Figure 14: The U.S. Great Recession: Effects of Government Consumption and Investment




Figure 15: The U.S. Great Recession: Effects of Forward Guidance

