

Hysteresis in Unemployment and Jobless Recoveries

Dmitry Plotnikov*

Abstract

This paper develops and estimates a general equilibrium rational expectations model with search and multiple equilibria where aggregate shocks have a permanent effect on the unemployment rate. If agents' wealth decreases, the unemployment rate increases for a potentially indefinite period. This makes unemployment rate dynamics path dependent as in Blanchard and Summers (1987). I argue that this feature explains the persistence of the unemployment rate in the U.S. after the Great Recession and over the entire postwar period.

JEL codes: E20, E30, E32

*Email: dplotnikov@imf.org. International Monetary Fund, 700 19th st, Washington, D.C., 20431. Tel: 202-623-8550. Fax: 202-589-8550.

1 Introduction

This paper addresses three questions. First, what caused the Great Recession of 2008? Was it a demand or a supply failure? Second, why has the unemployment rate been so persistent and why has it remained above 8 percent for more than 18 quarters since the official end of the recession? Third, how was this recession different from the other eleven post-war recessions? Can it be explained within a single framework together with the other postwar recessions?

Figure 1 presents employment dynamics for all postwar recessions including the Great Recession of 2008. The X-axis corresponds to the number of months from peak employment in the corresponding recession. The Y-axis measures the employment drop in percent from the peak. Zero is normalized to be peak employment in every recession.

Figure 1 shows that employment dynamics in postwar recessions varied significantly. In particular, prior to 1990 recoveries in employment were quick and took less than 8 quarters. After 1990 employment recoveries have been much more sluggish. For this reason economists dubbed the last three recoveries “jobless.”²

Several reasons behind employment sluggishness in the last three recoveries have been proposed in the literature. Jaimovich and Siu (2013) connect this phenomenon to the phenomenon of job polarization and argue that mechanization and technological advances contributed to disappearance of middle-income jobs during the last three recessions. Shimer (2012) emphasizes importance of wage rigidities. Berger (2013) suggests that a decrease in union power in the 1980s made it easier for firms to lay off less productive workers during the last recessions. Wiczer (2013) highlights importance of heterogeneity of unemployment spells across different occupations and sectors.

In this paper I focus on wealth changes as a reason behind the speed of an employment recovery. What I mean by this is that the last three “jobless” recessions and especially the Great Recession were associated with large wealth losses. For example, during the Great Recession, the real net worth of U.S. households and non-profit organizations declined by 20 percent within one year and the decline from the peak to the trough of the last recession was larger than 25 percent.

Table 1 suggests that wealth losses and slow employment recoveries are indeed related for all postwar recessions and not only for the Great Recession. For every postwar recession in the U.S. Table 1 indicates the number of months it took employment to recover to the pre-recession level (the second column) and the average quarterly growth rates of real wealth from the peak of employment to the end of each recession (the third column).³ I measure real wealth as net worth of households and nonprofit organizations (quarterly data) divided by CPI. Table 1 shows that the last three “jobless” recessions are associated with negative average wealth losses. In contrast, in all recessions prior to 1990 average wealth change was positive.

The link between large wealth losses and persistent decrease in employment is even more evident if one looks only at the recessions associated with large wealth losses such as financial crises and housing busts. Figure 2 presents employment dynamics during so called “Big Five” financial crises of the 20th century (Reinhart and Rogoff (2009)). As in Figure 1 the Y-axis measures the employment in percent from peak employment. Zero is normalized to be peak employment in every recession. In contrast to Figure 1, the X-axis corresponds to the number of years from the peak of employment.

In all of the recessions in Figure 2 employment sharply decreased in the beginning, was highly persistent and extremely slow to recover. For example, it took Finland over seventeen years to return to the pre-recession level of employment after the banking crisis in 1990.

To address the questions and observations raised above quantitatively, I develop and estimate a general equilibrium rational expectations model with multiple steady state equilibria.⁴ One implicit benefit for multiplicity of equilibria is that wealth in my model can change for both fundamental and non-fundamental reasons.⁵ In other words, if, for example, all agents rationally start believing that their real assets are worth

²Galí, Smets, and Wouters (2012) uses term “slow recoveries” to emphasize that there was no structural break in the relationship between U.S. GDP and employment in 1990. My model is entirely consistent with this evidence as it can generate both regular and jobless recoveries.

³Wealth data for the 1948 recession is not available.

⁴Since each steady state is associated with its own unique unemployment rate, my model does not contain a natural rate of unemployment. Instead many unemployment rate are sustainable as a long-run equilibria.

⁵Shiller (2000) among others argued that the run-up in housing prices prior to 2005 and the subsequent decline cannot be attributed to any change in fundamentals. Farmer, Nourry, and Venditti (2013) argue that most of the stock market movements are due to “animal spirits”.

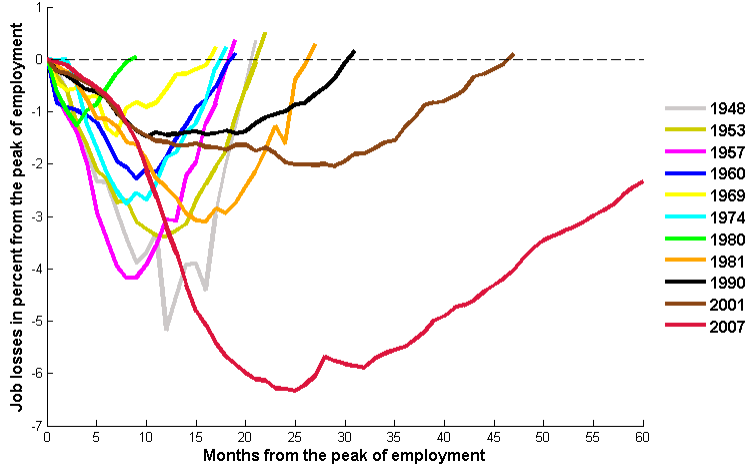


Figure 1: Employment dynamics in all post-war recessions in the United States. The X-axis corresponds to months from peak employment. The Y-axis measure the employment in percent from the peak. Zero is normalized to be peak employment in every recession. Source: Total Nonfarm Employment, BLS.

Recession	Employment recovery (months)	Average wealth change (%)
1953	8	1.20
1957	8	0.75
1960	6	1.12
1969	6	0.65
1974	6	0.71
1980	4	0.51
1981	9	0.75
1990	11	-0.15
2001	16	-0.13
2007	>66	-3.04

Table 1: Wealth losses and joblessness of recoveries. Wealth is calculated as Net Worth of Households and Nonprofit Organizations (quarterly data) divided by CPI. Source: FRED.

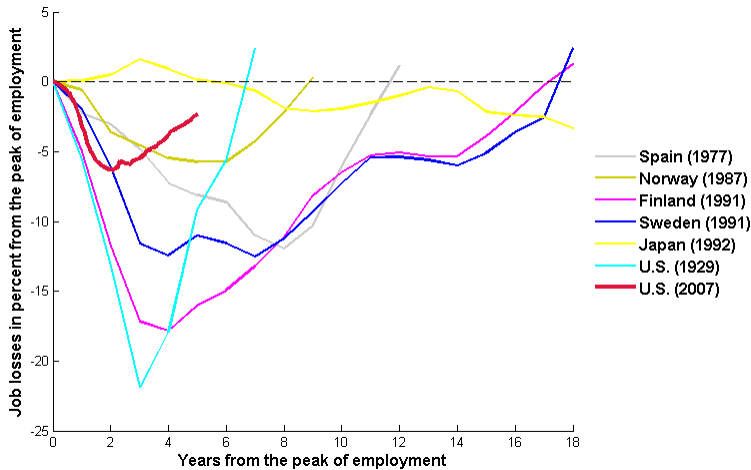


Figure 2: Total nonfarm employment dynamics during Big Five financial crises. The X-axis corresponds to years from peak employment. The Y-axis measures employment in percent from the peak. Zero is normalized to be peak employment in every recession. Source: IMF.

less than in the previous period, this will negatively affect everyone's real wealth, decrease aggregate demand and force firms to lay off workers and move the economy to a new equilibrium with a higher unemployment rate. If wealth does not recover for either fundamental or non-fundamental reasons, the economy can stay in the high unemployment rate equilibrium for a long time.

In my model, the unemployment rate exhibits hysteresis and is path-dependent in the sense of Blanchard and Summers (1986), Blanchard and Summers (1987).⁶ This result differs from the widely accepted assumption that demand shocks do not have long-run effects on the unemployment rate (see Blanchard and Quah (1989), Long and Plosser (1983), Nelson and Plosser (1982) among others).

I find in my estimated model, that demand and supply shocks are both important in understanding the dynamics of macroeconomic time series in the post-war period. I am also able to explain why the importance of persistent demand shocks has not been recognized in earlier papers: they have a permanent effect on the unemployment rate and the real wage, but not on real output and its components.

The core of my model is a real version of the incomplete factor market model introduced in Farmer (2006). In Farmer's model unemployed households and potential employers engage in a random job search and matching process. The labor market is incomplete because there are no markets for inputs to the matching technology. Because factor markets are incomplete, the system that describes general equilibrium has one more unknown than the total number of equations. As a result, the model has a continuum of steady state equilibria.⁷

Instead of adding an extra equation for the wage (such as the Nash-bargaining equation), I assume that both firms and workers take the wage as given, as in a standard RBC model. I pick a particular equilibrium by adding a separate belief function that specifies how forward-looking households form self-fulfilling expectations about their wealth. This equation determines the aggregate demand for goods which, in turn, pins down output.

My model differs from Farmer (2010), Farmer (2011a), Farmer (2011b), Farmer (2012b) in two important ways. First, Farmer closes a model of this kind by specifying self-fulfilling beliefs about the stock market. In my model, there is no analog of the stock market. Instead, I adapt the work of Friedman (1957) on adaptive expectations. In my model, households form expectations about their wealth that I model as permanent income. As in Farmer (2002), adaptive expectations do not violate rationality of agents because of the multiplicity of equilibria.⁸

Second, in my model capital is reproducible. This assumption makes it possible to analyze the behavior of consumption and investment separately. Since these two series behave very differently in the data, this modification to Farmer's model is an important step towards an empirically relevant theory.

The equations of my model pin down the dynamic path of the economy given the transversality condition plus the initial belief about permanent income. However, the model still has a continuum of steady states because permanent income in every steady state coincides with GDP. The adaptive expectations equation uniquely determines the path of the economy, but the steady state that the economy eventually converges to, depends on the initial belief about permanent income.

My model features two sources of fluctuations: i.i.d. changes to perceived permanent income (demand) and standard persistent productivity (supply) shocks. I show that, even when driven by i.i.d. self-fulfilling belief shocks, the model can produce a high degree of persistence in the unemployment rate. This contrasts with a standard RBC model where persistence in employment comes entirely from persistence in total factor productivity (TFP).

In my model TFP shocks have a qualitatively similar, although more protracted, effect on the economy than in a standard RBC model. By protracted I mean that the economy experiences permanent effects following a productivity shock. By qualitatively similar, I mean that a positive productivity shock causes a

⁶To my knowledge, very few papers argued in favor of this hypothesis. The major exception since 2000 is the paper by Ball (2009). This paper presents new empirical evidence from 20 countries that supports hysteresis in unemployment.

⁷The multiplicity of steady state equilibria is the key difference between the model in this paper and the previous generation of endogenous business cycles models such as Farmer and Guo (1994) or Benhabib and Farmer (1994). These models exhibit dynamic indeterminacy, but all dynamic paths converge to a single steady state. In contrast, the model I describe here has a continuum of steady states but a unique dynamic path that converges to each of them. In both types of models an independent expectations equation selects an equilibrium in every time period. For more details on comparison and evolution of the concept of indeterminacy see Farmer (2012a).

⁸Other papers that study when adaptive expectations can be rational include Evans and Honkapohja (1993), Evans and Honkapohja (2000).

large increase in real investment, a smaller increase in real output and a small increase in real consumption. TFP shocks act as a cyclical component in my model. As I show in Section 5, they play an important role in explaining postwar dynamics.

Together, demand and supply shocks can generate jobless recoveries, and a high persistent unemployment rate can coexist with positive real growth.

I use a Bayesian estimation procedure to estimate the parameters of my model.⁹ After I obtain estimates of all of the model parameters, I quantitatively describe the effect of each shock using impulse response functions. I show that together a TFP shock and a belief shock can reproduce the observed dynamics of the unemployment rate, output, consumption and investment for the entire post-war period, including the recovery from the Great Recession. By combining these two shocks in the right proportions, I am able to produce the characteristic business cycle dynamics of the first eight post-war recessions, as well as the jobless recoveries we have observed in the three recessions after 1990.

2 Model

This section builds a dynamic stochastic general equilibrium model. The distinguishing feature of the model is that households are not on their labor supply curve.¹⁰

My model deviates from existing models with multiple steady states in two dimensions. First, capital is reproducible. This lets me analyze the behavior of consumption and investment; two series that behave very differently in the data.

Second, because there is no analog of the stock market in my model, I do not close it with an exogenous sequence of asset prices as in Farmer (2012b). Instead, in Section 3, I adapt Friedman (1957) - I assume that households form adaptive expectations about their permanent income. As in Farmer (2002) because there is a multiplicity of equilibria, this assumption does not violate the rational expectations assumption.

A key advantage of closing the model with adaptive expectations is that it generates a feedback effect from the real economy to expectations of future income. Farmer (2012b) estimated a cointegrated VAR between the unemployment rate and the S&P 500 and found evidence of a significant negative short-run effect of increased unemployment on the stock market. This feedback channel is absent from previous incomplete factor models where the sequence of asset prices was assumed to be exogenous.

In my model, when GDP is high and the unemployment rate is low, agents revise up their perceived measure of permanent income. In response, they increase consumption demand, which in turn influences the real economy.

2.1 The household

There is a single representative household with a continuum of members. The household maximizes the expected discounted life-time utility of its members. All members of the household are perfectly insured against risks and thus they receive the same consumption value. Members derive utility from real consumption per household member and have no disutility from working.¹¹ I normalize the amount of time available to every member to one.

At every point in time t , a measure $l_t \in [0, 1]$ of members is employed and the rest $u_t = 1 - l_t$ are unemployed. The assumption that there is no disutility from working implies that members spend all available time searching for a job. This, in turn, implies that all household members are in the labor force. As a result all fluctuations in the equilibrium unemployment rate are due to agents dropping in and out of employment status and variation in the participation rate is absent in the model. Another implication of

⁹For a review of Bayesian methods see An and Schorfheide (2007). For applications see Fernandez-Villaverde and Rubio-Ramirez (2005a), Fernandez-Villaverde and Rubio-Ramirez (2005b).

¹⁰This assumption has both theoretical and empirical foundations. Theoretically, Farmer (2006) argues that the labor supply curve is missing because factor markets are incomplete. Kocherlakota (2012) refers to this feature as models of “incomplete labor markets.” Empirically, Kocherlakota (2012) argues that labor supply equation is inconsistent with real wage, consumption and employment dynamics after the 2008 recession. In the companion paper to the current one, Plotnikov (2012) argues that the labor supply assumption is the most problematic of all assumptions made in a standard business cycle model. Justiniano, Primiceri, and Tambalotti (2010) reach the same conclusion using a different dataset.

¹¹Evidence from microeconometrics studies (see Blundell and MaCurdy (1999)) supports a labor supply elasticity that is close to zero.

this assumption is that real consumption per household member corresponds to the real consumption per member of labor force in the data.

The one-period utility of the household is logarithmic:

$$u(c_t) = \log(c_t).$$

Each employed member of the household earns real wage w_t per period so that the real labor income of the household is $w_t L_t$. The household accumulates physical capital k_t that it rents out to firms for the gross real rental rate r_t per period. This leads to the standard budget constraint:

$$c_t + I_t = r_t k_t + w_t l_t, \quad (1)$$

where I_t is real investment in period t per member of the labor force. As usual, the price of the consumption good is equal to the price of the investment good and is normalized to unity because consumption and investment are perfect substitutes. Capital depreciates at the rate δ per period and evolves according to the expression:

$$I_t = k_{t+1} - (1 - \delta)k_t. \quad (2)$$

As in a standard RBC model the capital market is competitive: in equilibrium the gross real rental rate r_t is such that the amount of capital supplied by the household is equal to the quantity of capital demanded by a representative firm. In contrast, the labor market does not clear every period. Household members engage in search activities and become employed according to the following equation:

$$l_t = \tilde{q}_t h_t. \quad (3)$$

Equation (3) states that if a measure $h_t \in [0, 1]$ of agents searches for a job, $\tilde{q}_t h_t$ of them will find one, where \tilde{q}_t is determined in equilibrium from the labor matching technology (described below) and is taken as given by the household.

The household discounts future utility with discount factor β . It takes wages and rental rates and as given and maximizes discounted lifetime utility $\sum_{s=t}^{\infty} \beta^{s-t} u(c_s)$ subject to equations (1), (2), and (3). The solution to the household's maximization problem can be summarized by two equations. The first one is the standard Euler equation expressed in nominal terms:

$$\frac{1}{c_t} = \beta E_t \left[\frac{1}{c_{t+1}} (1 - \delta + r_{t+1}) \right]. \quad (4)$$

Second, no disutility from working trivially implies that all household members will search for a job in equilibrium:

$$h_t = 1. \quad (5)$$

2.2 Firms

Consumption and investment goods are produced by representative firms using a CES technology with capital and labor as inputs:

$$y_t = \left(a \cdot k_t^\rho + b \cdot s_t^\rho x_t^\rho \right)^{\frac{1}{\rho}}, \quad a + b = 1, \quad (6)$$

where y_t is real output, x_t is the amount of labor used in production of goods and s_t is a labor-augmenting TFP shock. The parameter ρ determines the elasticity of substitution between capital and labor:

$$\epsilon_{k,l} = \frac{1}{1 - \rho}. \quad (7)$$

The parameter ρ is less than one. The value $\rho = 0$ corresponds to the Cobb-Douglas production function, $\rho = -\infty$ corresponds to a Leontief technology and $\rho = 1$ is a linear technology. There are two main reasons why this specification is preferred to the standard Cobb-Douglas case. First, there is evidence in the literature (see, for example, Klump, McAdam, and Willman (2007)) that the elasticity between capital and labor is

significantly below unity and close to 0.5. Second, a Cobb-Douglas production function implies a constant labor income share. As I show in the companion paper Plotnikov (2012), the labor income share is strongly procyclical in the data. A CES production function helps to explain this observation.

A representative firm solves the following problem every period:

$$\max_{x_t, l_t, k_t} \left(a \cdot k_t^\rho + b \cdot s_t^\rho x_t^\rho \right)^{\frac{1}{\rho}} - w_t l_t - r_t k_t, \quad (8)$$

s.t.

$$x_t + v_t = l_t, \quad (9)$$

$$q_t v_t = l_t. \quad (10)$$

Equation (9) states that workers, l_t , hired by the firm can be allocated to one of two tasks: they can produce goods, x_t , or work in the human resources department, v_t . Both types of workers earn the same real wage, w_t , which the firm takes as given along with the interest rate, r_t .

The hiring process works as follows. A firm can hire as many job applicants as it wants at the real wage w_t ; however, not all applicants are suitable to work at the firm. For this reason the firm needs to screen potential employees using the human resources department, v_t . The variable q_t represents screening efficiency, which the firm takes as given.

Equation (10) states that the resulting number of employees at the firm, l_t , is equal to the number of people in the hiring/human resources department, v_t , augmented by the screening efficiency, q_t .¹² The screening efficiency of each worker, q_t , is determined in equilibrium by the aggregate matching technology. It depends on how many firms are looking to hire new workers at the same time. If many firms are searching, it is harder for each firm to find the necessary number of workers and q_t will be low. In contrast, if very few firms are searching, q_t will be high. In other words, equilibrium screening efficiency is a function of the aggregate unemployment rate.¹³

The firm's problem (8)–(10) can be reduced to the following:

$$\max_{l_t, k_t} \left(a \cdot k_t^\rho + b \cdot s_t^\rho l_t^\rho \Theta_t \right)^{\frac{1}{\rho}} - w_t l_t - r_t k_t,$$

where I define $\Theta_t = \left(1 - \frac{1}{q_t} \right)$ to be the externality that emerges from the labor search friction. Note that a high screening efficiency, q_t , implies a high value of Θ_t whereas a low value of q_t implies a low value of Θ_t .

The first order conditions for profit maximization of the firm are:

$$r_t = a \left(\frac{y_t}{k_t} \right)^{1-\rho}, \quad (11)$$

$$w_t = b s_t^\rho \Theta_t^\rho \left(\frac{y_t}{l_t} \right)^{1-\rho}. \quad (12)$$

These first order conditions have the standard intuition: the marginal product of capital (labor) is equal to the real rental rate (the real wage) respectively.¹⁴ This implies that firms always make zero profit in equilibrium, independent of externality level Θ_t .¹⁵

¹²Equation (10) implicitly makes a strong assumption that all workers are laid off in the end of each period. I intentionally make this assumption to show that the model can generate high persistence in the unemployment rate even without assuming that employment is a state variable. See beginning of the next subsection for further discussion.

¹³Equation (10) also implies that hiring costs are expressed in terms of labor units and not in terms of output as in standard labor search models. This assumption is not essential and is made for simplification.

¹⁴Notice that if $\rho = 0$ both equations reduce to the standard Cobb-Douglas case: income shares of both inputs are constant and equal to a for capital and b for labor.

¹⁵By definition, the income shares of capital and labor are:

2.3 Search in the labor market

The variables q_t , \tilde{q}_t and Θ_t are determined in equilibrium and depend on aggregate activity levels. To distinguish aggregate from individual variables, I put bars over variables to denote the aggregate level of activity. For example, l_t represents the number of people employed by an average representative firm (the firm's choice variable) whereas \bar{l}_t represents the aggregate level of employment (the individual firm does not have control over it). Because all firms and household members are the same, in equilibrium these two concepts coincide, but it is important to distinguish between them.

I assume that labor is rehired every period. This implies that the number of matches every period, m_t , is equal to the aggregate number of people employed, \bar{l}_t .¹⁶ I make this assumption for two reasons. First, it makes the dynamic system of equations of my model similar to a standard RBC model (see the next subsection for details). Second, assuming that all workers are rehired every period allows me to find a closed-form solution for Θ_t as a function of the aggregate employment level \bar{l}_t .¹⁷

The aggregate number of matches between firms and workers is determined according to the following Cobb-Douglas matching technology:

$$m_t = (\Gamma \bar{v}_t)^\theta \bar{h}_t^{1-\theta}, \quad (13)$$

where Γ is a scale parameter and θ is the elasticity of the matching function with respect to the aggregate number of recruiters, \bar{v}_t . As in a standard labor search model,

$$q_t = \frac{m_t}{\bar{v}_t}, \quad \tilde{q}_t = \frac{m_t}{\bar{h}_t}. \quad (14)$$

Equations(3), (5) and (10) along with definitions for q_t , \tilde{q}_t (Equations (14)) imply the following expressions for q_t and \tilde{q}_t :

$$q_t = \frac{\Gamma}{\bar{l}_t^{\frac{1-\theta}{\theta}}}, \quad \tilde{q}_t = \bar{l}_t. \quad (15)$$

One can also derive the following expression for the externality term Θ_t :

$$\Theta_t = \left(1 - \frac{1}{q_t}\right) = \left(1 - \frac{\bar{l}_t^{\frac{1-\theta}{\theta}}}{\Gamma}\right). \quad (16)$$

Notice that the externality term Θ_t is a decreasing function of the aggregate level of employment. Intuitively, since all firms search for workers in the same pool, it is harder to fill a vacancy if all firms exert higher search effort simultaneously.

This expression for Θ_t leads to the following aggregate production function for the economy:

$$\bar{y}_t = \left(a \bar{k}_t^\rho + b s_t^\rho \bar{l}_t^\rho \left(1 - \frac{\bar{l}_t^{\frac{1-\theta}{\theta}}}{\Gamma}\right)^\rho \right)^{\frac{1}{\rho}}. \quad (17)$$

$$\begin{aligned} (\text{Capital share})_t &= \frac{R_t k_t}{Y_t} \\ (\text{Labor share})_t &= \frac{W_t l_t}{Y_t} \end{aligned}$$

Notice that these shares are not constant in the general CES case (except for the Cobb-Douglas case, $\rho = 0$) and depend on the current level of income and both inputs. Additionally, the labor income share is affected by the externality in the labor market, Θ_t and the labor productivity s_t . This feature of the CES technology can partially explain procyclicality in the labor income share in the data (see previous section). On the other hand, this mechanism is not central to the paper and does not substantially improve the performance of a standard model.

¹⁶In contrast, the standard practice in the labor literature is to assume that employment is a state variable (see for example Rogerson, Shimer, and Wright (2005)).

¹⁷Even when the employment level is a state variable, the standard model generates a very low persistence of the unemployment rate. In contrast, my model produces a highly persistent unemployment rate even without this assumption. In general, making employment a state variable will not worsen performance of the model. Farmer (2011a) builds a similar model where employment is a state variable and shows that none of the results depend on this simplification.

In contrast to a standard Cobb-Douglas production technology, this production function has an inverted U-shape for a fixed \bar{k}_t and reaches peak at l^* , where

$$l^* = (\theta\Gamma)^{\frac{\theta}{1-\theta}}. \quad (18)$$

The level of employment l^* corresponds to the socially optimum level of employment. A social planner that maximizes the discounted stream of household's utility functions subject to the matching and production technologies would choose this employment level in every time-period. Intuitively, l^* does not depend on the capital stock or on time t , because employment is an intratemporal decision variable for the social planner. Later, I use Equation (18) to calibrate the scale parameter Γ in the estimation procedure.

2.4 Equilibrium

To sum up, the equilibrium of the model is represented by the following set of equations:

$$\frac{1}{c_t} = \beta E_t \left[\frac{1}{c_{t+1}} (1 - \delta + a \left(\frac{y_{t+1}}{k_{t+1}} \right)^{1-\rho}) \right], \quad (19)$$

$$c_t + I_t = y_t, \quad (20)$$

$$I_t = k_{t+1} - (1 - \delta)k_t, \quad (21)$$

$$y_t^\rho = ak_t^\rho + bs_t^\rho l_t^\rho \left(1 - \frac{l_t^{\frac{1-\theta}{\sigma}}}{\Gamma} \right)^\rho, \quad (22)$$

$$w_t = bs_t^\rho \left(1 - \frac{l_t^{\frac{1-\theta}{\sigma}}}{\Gamma} \right)^\rho \left(\frac{y_t}{l_t} \right)^{1-\rho}, \quad (23)$$

where I have substituted expressions for the rental rate (11) into the Euler equation (4) and the expression for the externality term Θ_t (Equation (16)) into the production function (Equation (6)) and the first order condition for labor for a firm (Equation (12)). Equations (19)–(21) are the same as in any standard RBC model. Equations (22)–(23) include the externality term $\Theta_t = \left(1 - \frac{l_t^{\frac{1-\theta}{\sigma}}}{\Gamma} \right)$ which is absent from the standard RBC model. Finally, my model does not have a labor-leisure equation.

Note that if the externality term Θ_t is absent, the model reduces to the standard RBC model with fixed labor supply. This is intentional: I constructed the model so that deviations from the standard model are minimal.

Assuming a standard autoregressive process for TFP in logs gives an additional equation:

$$s_t = s_{t-1}^\lambda \exp(\epsilon_t^p), \quad \epsilon_t^p \sim N(0, \sigma_\epsilon^2), \quad (24)$$

The parameter λ captures the persistence of labor productivity and ϵ_t^p represents the innovation to this process. The six equations (19)–(24) jointly determine the seven-dimensional vector of unknowns:

$$X_t = \left[s_t, k_{t+1}, c_t, I_t, y_t, l_t, w_t \right].$$

The model has one more unknown than the total number of equations because the market for search intensity of workers and vacancies posted by the firms is missing. This makes the model incomplete and, without an additional assumption to pin down the steady state, the system (19)–(24) has a continuum of steady state equilibria, each associated with a unique steady state employment level $l_{ss} \in (0, 1]$. Only one of these equilibria is socially efficient (see Equation (18)).

To see which variables are determined by the six equations (19)–(24) I evaluate the system at a steady state. Equation (24) gives the steady state productivity level $s_{ss} = 1$. The Euler equation (19) pins down the ratio of output to investment in wage units in every steady state,

$$\frac{y_{ss}}{I_{ss}} = \frac{1}{\delta} \left(\frac{\frac{1}{\beta} - (1 - \delta)}{a} \right)^{\frac{1}{1-\rho}},$$

and as a consequence of the national accounts identity (20) the ratio of consumption to output in wage units is determined by the expression,

$$\frac{c_{ss}}{y_{ss}} = 1 - \delta \left(\frac{a}{\frac{1}{\beta} - (1 - \delta)} \right)^{\frac{1}{1-\rho}}. \quad (25)$$

The existence of a multiplicity of steady state employment levels is not new in the labor search literature and is usually resolved by adding an extra equation to the system. This equation is *independent* of the rest of the model and is often represented by a wage equation that determines how the surplus is split between a matched firm and a matched worker. The best known assumptions used in the literature are the Nash-bargaining equation (originally proposed in the original Diamond-Mortensen-Pissarides framework) and the sticky wage assumption (see, for example, Shimer (2012), Blanchard and Gali (2008), Hall (2005)). Recently Farmer (2011b) proposed that the model should be closed by specifying an exogenous sequence of asset prices.

3 Closing the model

In this paper I extend Farmer (2011b). I add a belief function as a new independent fundamental that resolves the indeterminacy of dynamic equilibrium paths. The way I do this is new. By using adaptive expectations, I am able to explain how demand and supply shocks feedback into beliefs and influence the future path of unemployment, output, consumption, investment and the real wage.

In my model, consumption depends on permanent income.¹⁸ To incorporate this idea I adapt Friedman (1957). There, Friedman argued that consumption is proportional to permanent income:

$$c_t = \phi y_t^P, \quad (26)$$

and he assumed that expectations of permanent income are formed adaptively,

$$\frac{y_t^P}{w_t} = \left(\frac{y_{t-1}^P}{w_{t-1}} \right)^\chi \left(\frac{y_t}{w_t} \right)^{1-\chi} \exp(e_t^b) \quad e_t^b \sim N(0, \sigma_b^2). \quad (27)$$

In Equation (27), y_t^P is real permanent income (the perceived expected present value of all future and labor income) and $1 - \chi$ measures the speed of adjustment of permanent income to new information. The term e_t^b represents an independent shock to beliefs. This shock can be interpreted as a non-fundamental demand disturbance (“animal spirits”). I show later that it has a permanent effect on the unemployment rate. I refer to equations (26) and (27) as the permanent income hypothesis (PIH).

There are two differences between Friedman (1957) and Equations (26) and (27). First, the real wage w_t in Equation (27) is absent in Friedman’s work. Its presence in Equation (27) ensures existence of a balanced growth path in my model. Intuitively, the last period permanent income y_{t-1}^P and the current income y_t have to be expressed in the same units when they are weighted to obtain the current level of permanent income. The real wage is a natural example of such weight because the real wage grows with the same productivity time trend as real output. As a result ratio $\frac{y_t}{w_t}$ is absent of this time trend.

The second difference with Friedman’s work is that he obtained analogs of Equations (26) and (27) under the assumption that the household’s utility function is quadratic. When utility is not quadratic but the steady state is unique, Equations (26) and (27) are consistent with the rest of the model only if parameters ϕ and χ are complicated (and not constant) functions of the state variables. When there is a continuum of equilibria, by assuming that ϕ and χ are constant, I select a specific equilibrium where this is the case.

¹⁸Several recent papers pointed out the significance of wealth, in particular stock market and housing wealth, in the consumption decision. Ludvigson and Lettau (2004) pointed out that there is a low frequency relationship between consumption and asset wealth. Farmer (2012b) documented a similar relationship between consumption and S&P 500 index measured in wage units.

It is important to note that the belief function specified in Equations (26) and (27) does not violate the assumption of rational expectations. *In the case when there is a continuum of equilibria, the expectations equation selects an equilibrium in every time period.*

3.1 Steady state vs dynamic indeterminacy

To find when the adaptive expectations of permanent income are consistent with rational expectations, I evaluate Equations (26) and (27) at a steady state. Equation (27) implies that in every non-stochastic steady state, permanent income coincides with the economy's GDP, $y_{ss}^P = y_{ss}$. Because equation (26) has to hold in every steady state, it restricts the parameter ϕ to be the ratio of steady state consumption c_{ss} to the steady state level of income y_{ss} (see Equations (19)–(21)).

Thus, as a consequence of the rational expectations assumption, ϕ has to be the following function of the other parameters (see Equation (25)):¹⁹

$$\phi \equiv 1 - \delta \left(\frac{a}{\frac{1}{\beta} - (1 - \delta)} \right)^{\frac{1}{1-\rho}}. \quad (28)$$

Under condition (28), equations (26) and (27) are consistent with the system of equations (19)–(24). At the same time these two equations add no additional information about a steady state, and, as a consequence any unemployment rate can be a steady state equilibrium.

On the other hand, the belief function, in the form of the permanent income hypothesis, does resolve dynamic indeterminacy. By this I mean that Equations (26) and (27), together with the previously derived dynamic system of equations (19)–(21), pin down the model dynamics for a given set of initial conditions:

$$k_0 = \bar{k}_0, \quad (29)$$

$$s_0 = \bar{s}_0, \quad (30)$$

$$\lim_{T \rightarrow \infty} E_t \left(\beta^T \frac{k_T}{c_T} \right) = 0, \quad (31)$$

$$y_0^P = \bar{y}_0^P \quad (32)$$

The first three boundary conditions are standard and are implied by the initial value of capital, the initial level of labor augmenting technology process, and the transversality condition. The last condition is required because permanent income y_t^P is a new state variable.

Intuition behind the model dynamics is the following: the standard RBC model contains a unique saddle path and rational expectations govern evolution of the economy towards the steady state given initial conditions. My model contains a continuum of steady states with a unique saddle path associated with every steady state. As in the standard model, rational expectations govern evolution of the economy along every saddle path whereas the belief function specifies how the economy switches from one path to another every period.

In every steady state, permanent income coincides with current income: $y_{ss}^P = y_{ss}$. At the same time, initial permanent income, y_0^P , is required to pin down the model dynamics. This means that the model exhibits hysteresis as in Blanchard and Summers (1986), Blanchard and Summers (1987), and initial beliefs about permanent income y_0^P determine the steady state to which the model will converge. It also makes the model dynamics path-dependent.

4 Estimation

4.1 Overview

In this section I estimate the parameters of the model using Markov Chain Monte Carlo (MCMC). As an outcome of the procedure I obtain not only point estimates of all parameters of the model but also standard

¹⁹The necessity of this condition was first introduced by Muth (1961). There, the author showed that adaptive expectations about output are rational under specific restrictions on the speed of adjustment and on the output process, which, in his model, has to be ARIMA(0,1,1).

errors and the posterior distribution for each of the parameters.

I assume that I observe three series - the civilian unemployment rate, output in wage units and investment in wage units in quarterly data for the entire postwar period 1948:1 - 2011:4.²⁰ The model counterparts for these series are the variables $1 - l_t$, $\frac{y_t}{w_t}$ and $\frac{I_t}{w_t}$.

Bayesian estimation requires either that the number of shocks in the equations to be greater or equal to the number of observed series or that some of the variables are observed with measurement error. Because of this, I introduce a wedge ϵ_t^{LD} in Equation (23).

$$1 = b \cdot \exp(\epsilon_t^{LD}) s_t^\rho \left(1 - \frac{l_t^{1-\theta}}{\Gamma}\right)^\rho \left(\frac{y_t}{l_t}\right)^{1-\rho} \quad (33)$$

There are two reasons why inserting a wedge in Equation (23) is fitting. First, if $\rho = 0$, Equation (23) implies constant labor income share and that output in wage units is proportional to the employment rate. Because output in wage units and the employment rate are not exactly proportional in the data, the absence of this wedge creates a stochastic singularity. Second, in Plotnikov (2012) I found that the labor income share is highly countercyclical, although, not very volatile. Since $\rho = 0$ implies constant labor income share, inserting a wedge in Equation (23) is a way to evaluate how much of labor income share cyclicity can be explained by a CES production technology. In the same paper I argued that Equation (23), when $\rho = 0$, is the second most problematic equation based on the outcome of the accounting procedure of the type described in Chari, Kehoe, and McGrattan (2007).

The solution to the model is standard except for one detail. Because there is no unique steady state, I can choose a steady state around which to log-linearize the model. I picked the obvious candidate - the steady state associated with long-run statistical mean of the civilian unemployment rate in the data - 5.7%.²¹ This is the same steady state around which all DSGE models with unemployment are linearized. Thus the choice of this steady state eliminates potential differences due to the linear approximation between my model and any other DSGE model with unemployment.²²

After the equations of the model have been linearized, I solve them using the methods developed in Sims (2001). This brought the system to the form

$$X_t = F \cdot X_{t-1} + G\epsilon_t, \quad (34)$$

$$OBS_t = H \cdot X_t, \quad (35)$$

where $X_t = [\tilde{s}_t, \tilde{k}_{t+1}, \tilde{c}_t, \tilde{I}_t, \tilde{y}_t, \tilde{l}_t, \tilde{y}_t^p, E_t c_{t+1}, E_t y_{t+1}]$ is a vector of state variables (a tilde over a variable denotes a log deviation from the steady state), $OBS_t = [(\frac{y_t}{w_t})^{obs}, (\frac{I_t}{w_t})^{obs}, u_t^{obs}]$ is a vector of observables and $\epsilon_t = [\epsilon_t^p, \epsilon_t^b, \epsilon_t^{LD}]$ represents innovations to productivity and beliefs. The index t represents the time period and ranges from 1948:1 to 2011:4 - the maximum timespan available for the data series considered. The matrices F , G , and H are nonlinear functions of the underlying parameters of the model.

Equation (34) contains the estimated linearized policy functions and the functions that describe how expectations about the future are formed based on the state vector X_{t-1} . Equation (35) links the observed series to their model counterparts.

I constructed the likelihood for every set of parameters by treating the first equation as if it were the true data generating process. Since OBS_t contains fewer variables than X_t , in order to construct the likelihood of a set of parameters, I used the Kalman filter to construct forecasts of all the underlying variables, X_t , given the observed variables $\{OBS_s\}_{s=0}^t$.²³ For this, I used the algorithms that are part of the open source MATLAB package DYNARE (See Adjemian, Bastani, Juillard, Mihoubi, Perendia, Ratto, and Villemot (2011)).

²⁰The data was constructed as in Farmer (2010). In the next section I present summary statistics and plot these series.

²¹The results of the model are not sensitive to this choice.

²²Recall that, in contrast to standard models, this steady state does not correspond to the socially optimum steady state. The expression for the socially optimal steady state is given by Equation (18).

²³Since the matrix F always has a unit eigenvalue, I use the diffuse Kalman filter. Because the state space model (34) is non-stationary, the initial value for the filtering procedure matters. First introduced in Jong (1991) and later developed in Koopman (1997), Koopman and Durbin (2000), the diffuse Kalman filter addresses this problem.

Then, I combined prior distributions of all parameters with the maximized likelihood function to obtain a posterior likelihood. I computed posterior distributions numerically using the random walk Metropolis-Hastings algorithm (MCMC). I made 100,000 draws and kept 50,000 of them to ensure independence from the starting point. Further details on the computational procedure can be found in An and Schorfheide (2007).

4.2 Priors

To estimate the parameters of the model using Bayesian maximum likelihood, I need to specify priors for all of the parameters of the model. The priors I used are summarized in Table 2.

I chose the prior mean of the parameter a in the CES production function to be close to the calibration value in the Cobb-Douglas case.²⁴ Instead of estimating the parameter ρ directly, I estimate the elasticity between capital and labor $\epsilon_{k,l}$. Because of the one-to-one relationship between the two (Equation (7)), any statistic characterizing a posterior distribution of ρ can be inferred from a posterior distribution of $\epsilon_{k,l}$. Recent evidence in the literature (see for example Klump, McAdam, and Willman (2007)) indicates that the elasticity between capital and labor can be significantly below unity and is close to 0.5. Taking this into account I set the prior for $\epsilon_{k,l}$ to be as wide as possible with a mean of 0.5.

Parameter	Description	Distribution	Prior mean	Std. Dev.
a	\approx Capital share (= if $\rho = 0$)	beta	0.33	0.15
$\epsilon_{k,l}$	Elasticity b/w capital and labor	beta	0.50	0.25
δ	Capital depreciation	beta	0.03	0.015
β	Discount factor	Fixed	0.99	–
u^*	Optimal unemployment rate	Fixed	0.045	–
λ	Labor productivity persistence	beta	0.90	0.05
χ	Expectations persistence	beta	0.90	0.05
θ	Elasticity of the matching function	beta	0.5	0.25
σ^p	St.dev. of ϵ^p	Inv. Gamma	0.02	0.01
σ^b	St.dev. of ϵ^b	Inv. Gamma	0.02	0.01

Table 2: Prior distributions of parameters of the model

The discount factor β is fixed at 0.99. This parameter is not identified in the data. The choice of 0.99 corresponds to the long-run annual interest rate of 4 percent. I picked the value for the scale parameter Γ in the matching function so that the optimal unemployment rate, u^* , is 4.5 percent as in Hall (2011). Intuitively, the socially optimal unemployment rate should be lower than the long-term average and somewhat close to it.²⁵

For the rest of the parameters I chose somewhat wide priors. The prior mean for the persistence in the labor productivity process (λ) is 0.9. I picked the prior mean for the quarterly capital depreciation rate to be 0.03 which corresponds to an annual depreciation rate of 12 percent. The prior mean of the speed of adjustment of adaptive expectations is 0.8 - the same as the equilibrium persistence in Evans and Ramey (2006).

Estimates of the elasticity of the matching function (θ) range significantly in the literature - from 0.28 in Shimer (2005) to 0.54 in Mortensen and Nagypal (2007). Farmer (2012b) fixes θ at 0.5. Taking this into account I set the prior mean at 0.5 with a standard deviation of 0.25.

I picked the prior mean for the standard deviation of the productivity innovations to be 2 percent from the steady state - the standard calibration value for TFP process. I picked the same prior mean for the

²⁴Recall, that if the production technology is CES, the capital share is not constant and depends on the current physical level of output and capital.

²⁵Because l^* depends on both Γ and θ , every time I evaluate the likelihood for a given set of parameters, I choose Γ so that Equation (18) is satisfied for the given value of θ and the fixed value of u^* . The results of the model do not depend on the value of Γ as long as $\bar{l} < l^*$. This inequality is not restrictive and follows from the inverted U shape of the production function as a function of labor, l_t for a fixed capital stock (see Equation (17)). The peak of this function corresponds to the value l^* . The point at which the model is linearized, \bar{l} , should be on the increasing part of this function to ensure that increasing the number of workers leads to an increase in production.

standard deviation of belief shocks, σ^b , and the labor demand shock, σ^{LD} . For all of the parameters I chose prior standard deviations that are as large as possible, given the means and the natural limits of the parameters.²⁶

4.3 Posterior estimates

The outcome of the MCMC algorithm is presented in Table 3. For every parameter I report the prior mean, the posterior mean and the 90 percent posterior confidence interval centered around the posterior mode. All estimates are stable with respect to the choice of priors. Changing a prior mean or a standard deviation for one or several parameters does not change the posterior estimates.

The posterior means of all model parameters have reasonable values. Parameter a of the CES production function is estimated to be 0.4585. Recall that in a general CES case, $1 - a$ is not equal to the labor income share. To see that $a = 0.4585$ corresponds to a standard value of the labor income share, I calculate the labor share at the steady state that corresponds to the long-run statistical average of the unemployment rate, $\bar{u} = 5.7\%$. Solving for the steady state level of output \bar{Y} that corresponds to \bar{u} implies a standard value of the labor income share at this steady state of 0.65.²⁷

The estimated elasticity between capital and labor $\epsilon_{k,l}$ is equal to 0.92 and implies that the other parameter of the CES production function, ρ , equals -0.089 (see Table 3 and Equation (7)). Together the estimates for a and ρ imply that the U. S. economy can be well approximated by a Cobb-Douglas function.²⁸ This finding contrasts with the recent evidence by Klump, McAdam, and Willman (2007) who found that the elasticity between capital and labor is significantly less than unity in the U.S. aggregate data.

Parameter	Description	Posterior mean	$CI_{90\%}$
a	\approx Capital share (= if $\rho = 0$)	0.4585	[0.3929, 0.5212]
$\epsilon_{k,l}$	Elasticity b/w capital and labor	0.9209	[0.8804, 0.9611]
δ	Capital depreciation	0.0082	[0.0079, 0.0086]
β	Discount factor	–	–
u^*	Optimal unemployment rate	–	–
λ	Labor productivity persistence	0.9175	[0.8784, 0.9531]
χ	Expectations persistence	0.9513	[0.9223, 0.9820]
θ	Elasticity of the matching function	0.6284	[0.3680, 0.8920]
σ^p	St.dev. of ϵ^p	0.0156	[0.0141, 0.0172]
σ^b	St.dev. of ϵ^b	0.0082	[0.0076, 0.0089]
σ^{LD}	St.dev. of ϵ^{LD}	0.0245	[0.0227, 0.0263]
$\log L = 2101$	MCMC accept. rate 32.84%	100000 draws	50000 kept

Table 3: Posterior distributions of parameters of the model.

The capital depreciation rate is estimated to be approximately 3.3 percent per year. Even though depreciation rate values used in the literature are usually higher, my estimate coincides with Bureau of Economic Analysis estimates on the depreciation rate of physical capital (see Nadiri and Prucha (1997) for an overview).

These three parameters – a , ρ , δ – together with the discount factor β imply that households consume 84 percent of their permanent income (see Equation (25) for the formula). This number is in line with the original estimate of 90 percent in Friedman (1957).

The next two parameters in Table 3 are persistence in labor productivity process, λ and the speed of adjustment of adaptive expectations χ that have posterior means of 0.92 and 0.95. As expected, both parameters imply sufficient persistence in these two processes. At the same time both posterior means are far from unity. In contrast to the standard model where persistence in the simulated series comes from persistent shocks, my model does not require highly persistent shocks and belief shocks are, in fact, i.i.d.

²⁶A rule of thumb that ensures bell-shaped prior distributions with no positive probability mass on either end of the support requires the distance between the prior mean and the closest natural limit to be less than two standard deviations.

²⁷The labor share in this case can be calculated as $\frac{1-\bar{u}}{\bar{Y}}$, where \bar{Y} is the level of output in wage units that corresponds to \bar{u} .

²⁸Recall that Cobb-Douglas case corresponds to $\epsilon_{k,l} = 1$ and $\rho = 0$.

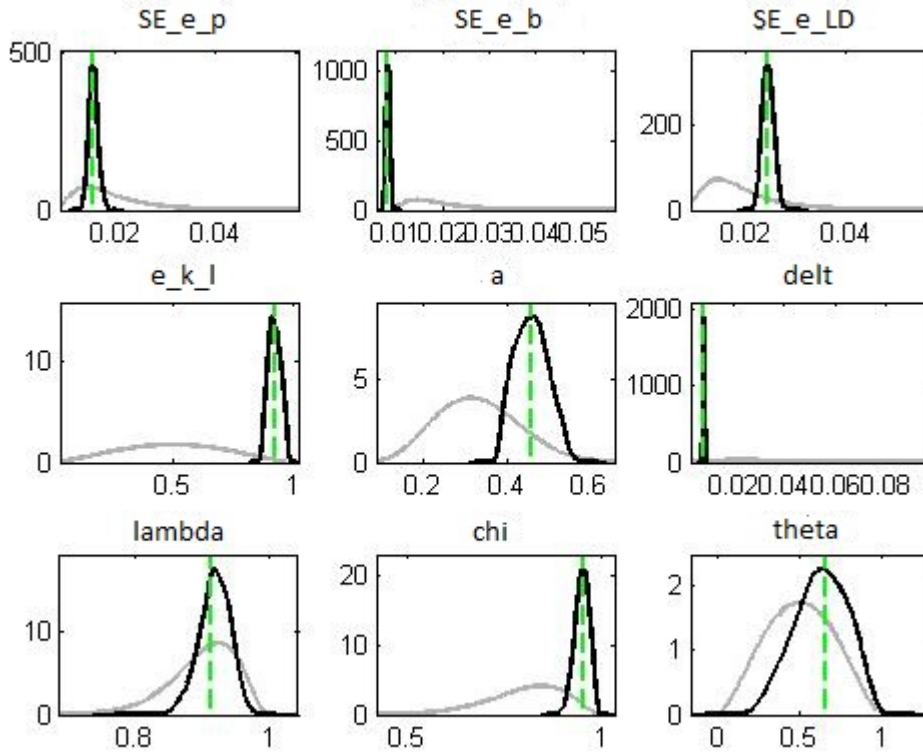


Figure 3: Posterior distributions. Grey line represents prior density. Dark line represents posterior density. Dashed line is estimated mode of the posterior density. SE stands for standard error.

Finally, the elasticity of the matching function (θ) is not well identified in the data. This is expected given the wide range of estimates of this parameter in the literature. The reported confidence interval contains values from 0.4 to 0.9.

The lower half of Table 3 presents estimates of the standard deviation of the realized shocks. All posterior standard deviations are relatively small. The standard error of the technology shock is estimated to be 1.6 percent, even lower than its standard calibration value of 2 percent. The standard deviation of the belief shocks is even smaller – less than 1 percent.

The low standard deviation for belief shocks means that non-fundamental disturbances do not need high volatility to explain the data. The low standard deviation for labor productivity innovations suggests that belief shocks explain a share of the observed volatility in the data that is usually attributed to TFP shocks.

In Figure 3, I plot posterior densities for all the estimated parameters, together with the prior distributions. This figure shows that all of the parameters are identified. If a parameter were to be weakly or non-identified, the difference between its prior and posterior densities would have been small. In contrast, all posterior densities, except for the one for θ , differ significantly from their respective priors, indicating that these parameters are identified. The posterior density for θ does not differ much from its prior distribution, which suggests that the data does not contain much information about this parameter.

5 Taking the model to the data

In this section I perform three exercises. First, I show that my model matches the stylized facts. Second, I explain the mechanism behind the model's performance and the contribution of each shock qualitatively. Finally, I present impulse response functions that clarify the effect of each shock on the economy.

5.1 Model performance

In this subsection I perform the following exercise. I use my model to simulate $R=10000$ Monte-Carlo samples of the same length as the length of the data sample ($T=256$ observations) using randomly drawn productivity and belief shocks. I report the Monte-Carlo average and 90 percent Monte-Carlo confidence interval centered around the average of each statistic of interest (such as the standard deviation, persistence of consumption, investment, etc).²⁹

Next, I compare each summary statistic of the data with its Monte-Carlo average.³⁰ The Monte-Carlo confidence interval represents the dispersion of a statistic across R Monte-Carlo samples. Fixing sample size and taking a sufficiently large R gives me a reasonable way to compare the similarity of the non-stationary series generated by the model with the non-stationary data.³¹

Note that even though there is no time trend in any of the data series in wage units and the unemployment rate, some of them still appear non-stationary (for example, the unemployment rate). Usually this degree of non-stationarity is filtered out using the HP-filter under the assumption of a unique steady state. I cannot apply a filter, because I do not assume a unique steady state and instead will explain the entire variation in the unemployment rate and the variable in wage units.

Table 4 presents the outcome of this exercise. I present statistics for four series: investment in wage units, consumption in wage units, output in wage units and the employment rate. All series are measured in log-deviations from their statistical means.³²

X_t	$Std(X_t)$		$Corr(X_t, X_{t-1})$		Unit Root?	
	Simulations	Data	Simulations	Data	Simulations	Data
	MC Avg. $CI_{90\%}$		MC Avg. $CI_{90\%}$		MC Share of unit root	p-value of unit root
$\frac{I_t}{w_t}$	14.468 [10.883,18.017]	12.363	0.904 [0.858,0.950]	0.904	0.0022	0.0001
$\frac{c_t}{w_t}$	5.589 [2.319,8.849]	5.648	0.981 [0.963,0.999]	0.988	0.7358	0.3407
$\frac{y_t}{w_t}$	6.007 [2.831, 9.173]	3.730	0.968 [0.942,0.997]	0.966	0.5167	0.0307
l_t	5.972 [2.784,9.149]	1.767	0.971 [0.942,0.997]	0.970	0.5518	0.0758

Table 4: Summary statistics of the simulated and the real data. All series are measured in log-deviations from their statistical means. MC=Monte-Carlo.

In the first two columns I compare the standard deviation in log-deviations for the simulated series and the actual data. The model not only generates standard deviations similar to those in the data, but also Monte-Carlo means that are very close to the actual standard deviations (for all series except the employment rate). This result is non-trivial because the standard deviation of a non-stationary series is highly history-dependent and the length of each sample is rather long.

As is in the data, simulated investment is the most volatile series among all four and the model captures this. The standard deviation of consumption in the data is lower than the standard deviation of investment and the Monte-Carlo mean matches it almost perfectly.

The model simulated GDP is more volatile – 6.007 – than the actual one – 3.730. Partially this is because GDP in the model consists only of consumption and investment whereas in the data GDP includes

²⁹The Monte-Carlo moments I report do not change once the number of simulated samples exceeds 5000. I take $R=10000$ to eliminate even small changes.

³⁰For example, I compare the standard deviation of investment to the Monte-Carlo average of the standard deviations of the simulated investment series.

³¹Because of nonstationarity I cannot take the limit $T \rightarrow \infty$, as many summary statistics do not converge. Moreover if some data series are non-stationary, statistics for these series will change over time. However as long as the sample size is fixed, statistics for the data and simulated series are comparable.

³²Since the long run statistical mean of the employment rate is 0.943, a one percent log-deviation for the employment rate is approximately a 1 percent deviation in terms of the labor force.

government expenditures. If I exclude government purchases from the measure of GDP in the actual data, the standard deviation of GDP increases from 3.730 to 5.204.³³

Because the estimated CES technology is very close to Cobb-Douglas, the employment rate is highly correlated with output. This implies that the standard deviation of the simulated employment rate is similar to the standard deviation of GDP - 5.972. The actual deviation of the employment rate - 1.767 - is outside the reported 90 percent confidence interval, although it still within a 95 percent confidence interval (not reported in Table 4).

The next two columns of Table 4 compare the persistence of the simulated series and the persistence of the actual data. As before I report the Monte-Carlo average of persistence along with the 90 percent confidence interval. The model matches persistence of all series almost exactly.

The model again correctly replicates the different behavior of the four series: as in the data, the investment series is the least persistent of all, whereas the persistence of the other series is as high as in the data and close to unity. This is again non-trivial because only the productivity process in the model is persistent with a relatively low persistence of 0.9175 (see Table 3) whereas belief shocks are i.i.d. Nevertheless the model can match very different persistence behavior in all four series.

In the last two columns I present the outcome of the unit-root test for the data and the simulated series.³⁴ For the actual data I report the p-value where the null is that the process has a unit root. The share of processes for which existence of a unit root is not rejected at the 5 percent level is reported in the simulation column.

In the data, the unit root hypothesis is strongly rejected for investment, strongly not rejected for the consumption series and both output and unemployment series are borderline cases.³⁵ The model reflects this observation - simulated investment is stationary in more than 99 percent of all cases, simulated consumption is a random walk in 73 percent of all simulations and both output and the employment rate are random walks slightly more than half of the time. These results are consistent with the p-values for the actual data. The higher the probability that the actual series is a random walk, the higher the share of simulated series that are random walks.

Based on the evidence I present in Table 4 I conclude that the model replicates key characteristics of the data. In the next subsection I address the mechanism behind this success.

5.2 The mechanism of the model

In this subsection I discuss the qualitative properties of the model-generated series. As before I pick parameters of the model to be equal to the posterior mean of the Bayesian estimation procedure (see Table 3).

The left column of Figure 4 presents typical dynamics of the model with randomly drawn technology and belief shocks. The right column of Figure 4 plots actual consumption and investment in wage units. On the bottom panel of the figure I plot the analog of output in the data - the sum of consumption and investment.³⁶ As in the previous subsection, all series are measured in log-deviations from their statistical means.

The left column of Figure 4 shows that the model qualitatively captures the data behavior: simulated consumption in wage units is smooth, persistent and drives medium term changes in output; whereas the investment series appears stationary and is much less persistent.

What is the mechanism behind the high serial correlation in the simulated data? Algebraically, the estimated policy rule always has an eigenvalue that is equal to unity.³⁷ This makes all the simulated series nonstationary. Economically, the simulated output and consumption series are persistent because the economy does not fluctuate around the unique steady state (as in a standard model), but constantly shifts

³³In Farmer and Plotnikov (2012) we explore effects of fiscal policy in this environment.

³⁴I used Dickey-Fuller test with no intercept and no trend since by construction means of all series are zero.

³⁵There are two reasons to believe that output in wage units is in fact a random walk. First, consumption and investment together constitute 81 percent of U.S. GDP. Since consumption is a random walk and investment is stationary GDP is likely to be a random walk as well. Second, since labor share in the U.S. data is close to being constant, this implies that the employment rate is roughly proportional to the output in wage units. Since existence of unit root in the unemployment rate is not rejected, it suggests that output in wage units has to have a unit root as well.

³⁶Consumption and investment constitute 81 percent of the U.S. GDP.

³⁷This is because Equations (26) and (27) are proportional to the rest of the equations in every steady state.

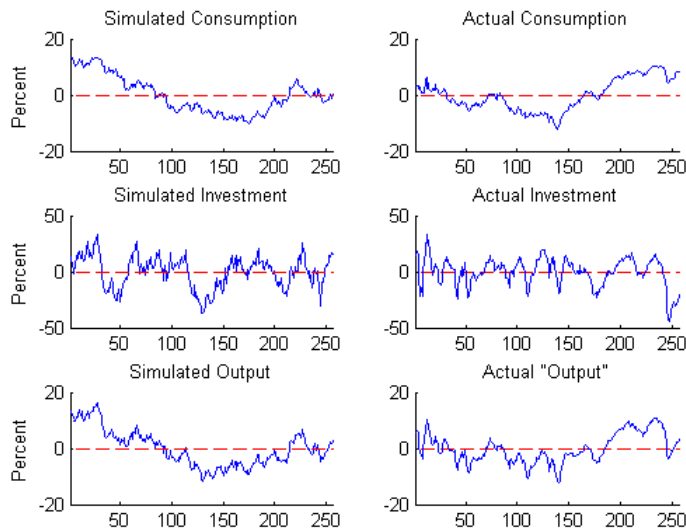


Figure 4: Simulated series in wage units (left column) and the actual data counterparts (right column). Output in quotation marks is defined as sum of consumption and investment. Scale of axis is the same for simulated and actual data.

from one steady state to another. This implies that the model economy exhibits hysteresis – the entire history of shocks determines the steady state to which the economy converges.

If all the simulated series are random walks, why is investment stationary (see Figure 4 center left and Table 4 last column)? In fact it is non-stationary, but the “degree of nonstationarity” is significantly lower for investment than for consumption and output in wage units. To see what I mean by this, consider the simulated series generated by the same innovations to beliefs and no productivity innovations, $\epsilon_t^p = 0$. These series are plotted in Figure 5 in the right column. The left column in Figure 5 is the same as the left column on Figure 4.

As expected, all the series in Figure 5, including investment, are highly persistent and satisfy the inequality (see Equations (25) and (28))

$$\frac{c_t}{y_t} \approx \phi \gg 1 - \phi \approx \frac{I_t}{y_t}. \quad (36)$$

Equations (25) and (28) imply that all the series in Figure 5 on the right are identical. This means that belief shocks are responsible for persistence in the data, but cannot replicate the high volatility of the investment series.

Because on average the share of consumption is approximately five times the share of investment in U.S. GDP, if one were to construct the investment series generated by the belief shocks in levels, it would appear almost constant. In other words, belief shocks imply very small changes in the investment series, but large changes in consumption and output.

What if beliefs shocks are shut down and only productivity shocks are used to generate data? The right column of Figure 6 presents the simulated series produced with the same productivity innovations as in the left column of Figure 6 (which is the same as the left column in Figure 4) and no belief shocks.

It follows from the right column of Figure 6 that TFP shocks have their largest effect on the investment series (the standard intuition from the real business cycle models applies) but almost no effect on consumption. This is because the adaptive expectations assumption (equation (27)) smooths the consumption series. As a result, all fluctuations in the investment series are transferred almost one to one to the output series.

If one were to construct the consumption and investment series implied by the right column of Figure 6, consumption would be almost flat, whereas investment would be similar to the one implied by the left column of Figure 6. This follows from the same logic as before: on average the share of consumption is

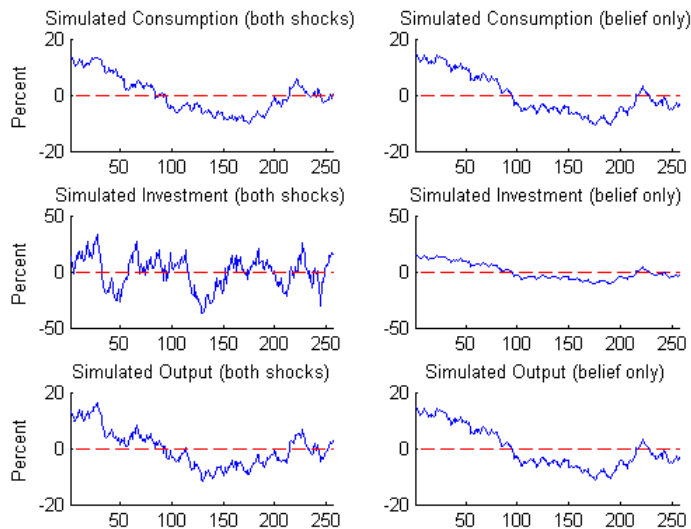


Figure 5: Simulated series with both shocks (left column) and simulated series with belief only (right column). The left column is the same as the left column of Figure 4. Scale of y-axis is the same for both columns.

approximately five times the share of investment in U.S. GDP. Because the fluctuations in consumption in levels, implied by productivity shocks alone (Figure 6, top right), are very small relative to the ones produced by both shocks (Figure 6, top left) these fluctuations result in an almost constant level of consumption.

Now consider how shocks to productivity and beliefs interact to produce the series in the left panel of Figure 4. Because belief shocks change the investment steady state only slightly, fluctuations in this series are dominated by productivity fluctuations. This results in an investment series that is statistically stationary. On the other hand, the consumption series is dominated by belief shocks and is very smooth. Combining the behavior of the consumption and investment series results in volatile and persistent output, where high persistence comes from consumption and volatility comes from the investment series - exactly as in the data.

5.3 Case study: the Great Recession

In this subsection I study how my model explains the output and employment dynamics during the Great Recession. I show how much of the decline in employment and output my model attributes to a decline in TFP and how much it attributes to the change in beliefs.

I compare my model with the standard RBC model and highlight the differences between the two models. I make and explain two points. First, my model amplifies and propagates standard TFP shocks. Second, when I add belief shocks implied by the data, my model can explain almost the entire persistent decline in employment after 2007.

To compare my model with a standard model, I use the standard RBC model with a separable linear disutility from labor (Hansen (1985)). The system of equations that describes the dynamic equilibrium in this model consists of Equations (19)-(24) with $\Gamma = +\infty$ with an additional labor-leisure condition:³⁸

$$\psi \cdot c_t = w_t \tag{37}$$

where ψ is the parameter that measures disutility from working.

To make the RBC model comparable to my model, I picked exactly the same parameter values as in my model (see Table 3). To calibrate parameter ψ I used Equation (37). Then, I linearized and solved the model using the methods described in Section 4.

³⁸Intuitively, $\Gamma = +\infty$ corresponds to an infinitely efficient matching function which means all matches are instant or, in other words, the labor market clears instantly.

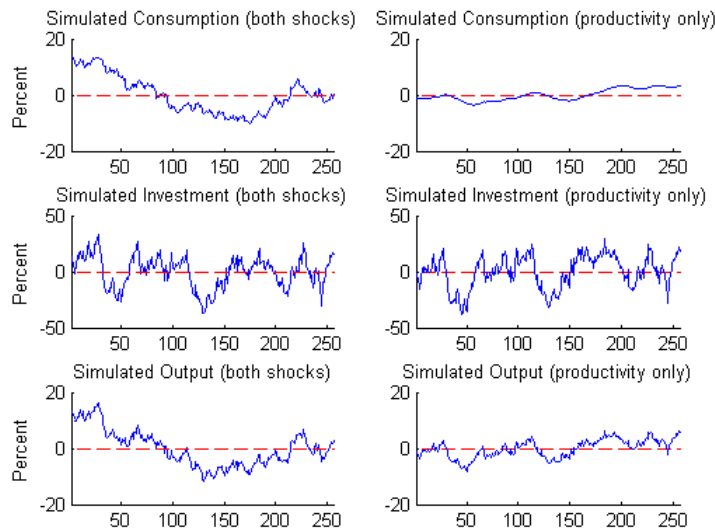


Figure 6: Simulated series with both shocks (left column) and simulated series with productivity shocks only (right column). The left column is the same as the left column of Figure 4. Scale of y-axis is the same for both columns.

To obtain the actual shock series, I downloaded raw quarterly TFP series from the San Francisco Federal Reserve website for the period from 2007:4 to 2011:4.³⁹ I adjusted these TFP series for the fact that the productivity process in both models is labor-augmenting. Then, assuming that these series come from the autoregressive process estimated in both models ($\log(s_t) = \hat{\lambda} \cdot \log(s_{t-1}) + \epsilon_t^p$, $\hat{\lambda} = 0.9175$ and $s_{2007:3} = 1$), I backed out the implied TFP innovations.

I assumed that the economy was at the steady state prior to 2007:3. Then, I plugged the implied TFP innovations ϵ_t^p in the linearized policy functions of both models. I plot the outcome of this experiment along with the actual employment rate and real output in Figure 7 (RBC model), Figure 8 (the benchmark model without belief shocks) and Figure 9 (the full model). Both the data series and the model-generated series are normalized to zero in 2007:4.⁴⁰

Figure 7 shows that the RBC model fails to explain the persistent decline in employment and real output since 2007. The RBC model predicts that both series should have recovered when the TFP series recovered - in the middle of 2010 (see Figure 7). This did not happen in the data. The RBC model also overestimates the decline in real output (Figure 7, the right panel).

When I plug in the same TFP series in the benchmark model, the effect is more persistent (see Figure 8). First, the model predicts a recovery only in the end of the sample. Second, the employment rate in the benchmark model is more persistent than in a standard model (see Figure 8, left). Finally, my model essentially matches the recovery in real output almost exactly (see Figure 8, right).

To obtain the belief shock innovations $\hat{\epsilon}_t^b$ implied by the TFP series, I rely on the benchmark model's response to belief shocks (see the next subsection). The result I use is that belief shocks do not affect any real variables except the real wage and the unemployment rate. To be more precise, a 1 percent negative shock to beliefs translates to a 1 percent permanent increase in the real wage and a 1 percent increase in the unemployment rate.

I then proceed as follows. First, I obtain the logarithm of the real wage series generated by the TFP series only, as described above. I subtract the logarithm of the real wage series implied by the TFP fluctuations alone from the logarithm of the actual real wage series. As I mentioned above, if this difference changes it

³⁹<http://www.frbsf.org/economic-research/total-factor-productivity-tfp/>

⁴⁰I defined the employment rate in the data to be one minus the civilian unemployment rate. Real output is defined as real U.S. GDP in chained 2005 U.S. dollars.

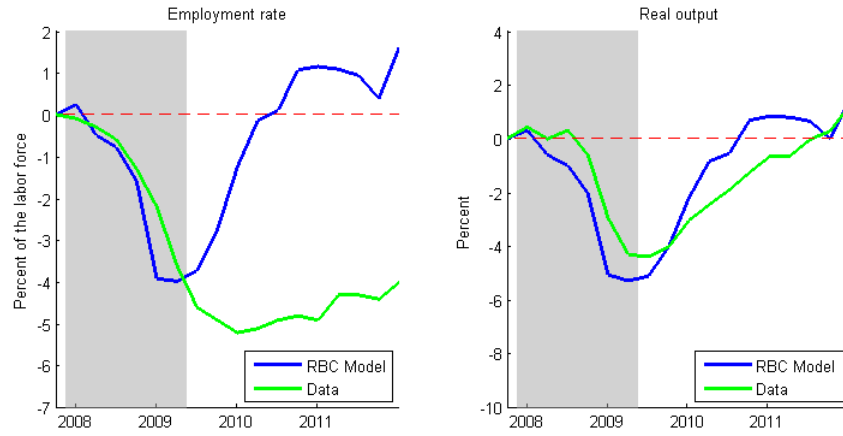


Figure 7: RBC model generated series and the actual data.

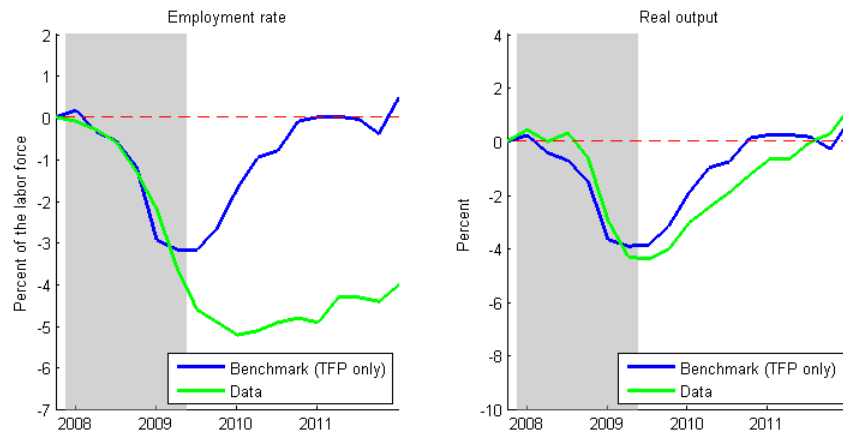


Figure 8: The benchmark model generated series using TFP shocks only against the actual data

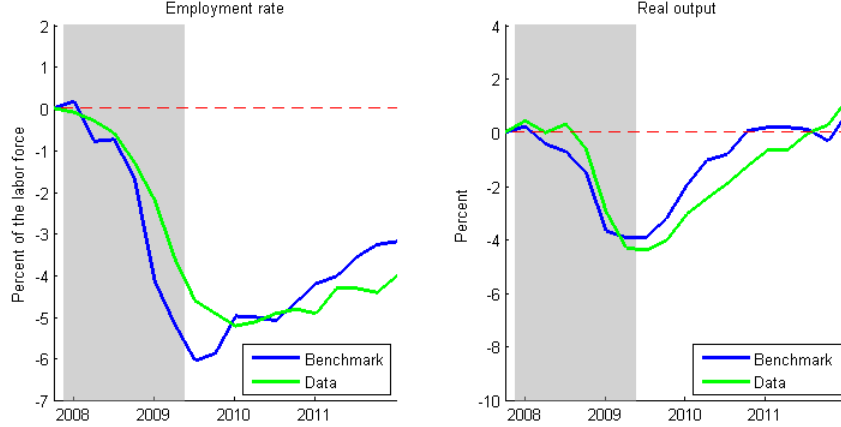


Figure 9: The benchmark model generated series using both TFP and belief shocks against the actual data.

can happen only because of the belief shock. Thus I can obtain $\hat{\epsilon}_t^b$ as

$$\hat{\epsilon}_t^b = \left(\log(w_t^{data}) - \log(w_t^{TFP}) \right) - \left(\log(w_{t-1}^{data}) - \log(w_{t-1}^{TFP}) \right).$$

If I plug these implied shock series in my model along with the TFP shocks, it can explain almost the entire decline in employment since 2007 (see Figure 9, left)⁴¹. Notice that real output series is almost unaffected by the belief shocks (compare Figure 9, right, and Figure 8, right).

5.4 Impulse response functions

In this subsection I explain the results in the previous subsection using impulse response functions (IRFs). First, I discuss why the multiple equilibrium model generates different dynamics in output from a standard model after a TFP shock as indicated on the left panels of Figures 8 and 9. Second, I show why TFP shocks generate more persistent unemployment rate after the 2008 crisis (Figure 8, left). I also point out that my model can generate rigid real wages as an equilibrium outcome. Finally, I explain how the full model with belief shocks can match the persistent unemployment rate after the 2008 crisis (see Figure 9, left).

First, I discuss why the dynamics for output and its components differ in my model and the RBC model. Figure 10 presents IRFs to a negative 1 percent TFP shock. This magnitude corresponds to approximately 60 percent of the standard deviation of productivity innovations. Because the productivity process is autocorrelated with persistence $\lambda=0.9175$, the effects of the initial shock disappear completely after approximately 45 quarters (see Figure 10, dotted line on both panels). The left panel of Figure 10 helps explain why the multiple equilibrium model produces a more protracted effect on real output, whereas the right panel focuses on the unemployment rate.

Qualitatively the effects of the TFP shock are standard with extra propagation and permanent effects. The left panel of Figure 10 shows that investment in my model is the most volatile component of GDP, as in the standard RBC model, whereas consumption is the least volatile series.

In contrast to the standard model, consumption changes only by less than 0.25 percent in absolute value.⁴² Consumption in my model drops less than in the standard model because of the adaptive belief function.

This result leads to a more persistent effect on output. Income decreases initially by the same amount in both models and, because consumption drops less, investment has to decrease more. Because agents are essentially pulling out capital from the economy, it takes longer for real output to recover. In the long-run the economy converges to the steady state with the same level of real output as before the initial impact, but slightly higher investment and slightly lower consumption.

⁴¹To be more precise, in Figure 9 I use output in wage units to identify belief shocks. Since real output is not affected by belief shocks, this procedure is identical to identifying belief shocks from the actual wages.

⁴²In the same situation consumption in the RBC model as calibrated in the previous section is almost 60 percent more sensitive.

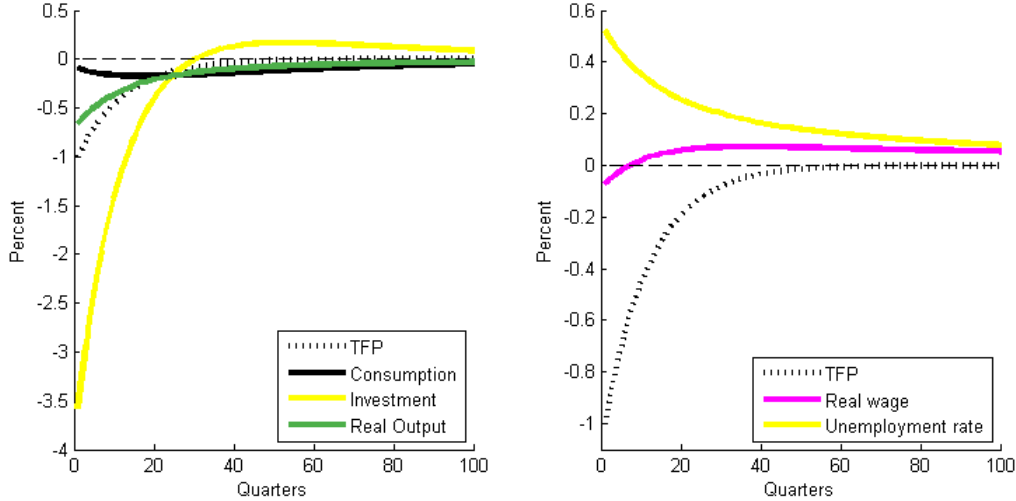


Figure 10: Impulse response functions to 1 percent negative TFP shock. The left panel presents the effects of this shock on real output and its components. The right panel presents the effects of the shock on the unemployment rate and the real wage. Effect on the unemployment rate is measured in percents of the labor force.

I now discuss why the unemployment rate is more persistent in my model even in the absence of belief shocks (see Figure 8, right). The right panel of Figure 10 explains why TFP shocks have a more persistent effect on the unemployment rate. Even though the TFP level is indistinguishable from its steady state value 50 quarters after the initial impact, the unemployment rate is still significantly affected by the shock even after 50 quarters. In particular, only 2/3 of the initial impact on the unemployment rate disappears after 50 quarters. After 75 periods 80 percent of the impact is absorbed and 10 percent is never absorbed and is permanent⁴³. This sharply contrasts with standard models of unemployment where a persistent response of variables to shocks comes mostly from persistence in the TFP process.

Intuitively, the effect of the TFP shock on the unemployment rate is persistent, because the real wage appears to be “rigid” with respect to TFP fluctuations (see the right panel of Figure 10), a point worth discussing in more detail. The overall change in the real wage is less than 0.2 percent. The real wage exhibits endogenous “rigidity” even though there are no artificial barriers that prevent wages from moving.

Why does the real wage appears rigid with respect to TFP fluctuations? The real wage in my model is the product of two terms:

$$w_t = MPL_t \cdot \Theta_t^\rho = bs_t^\rho \left(\frac{y_t}{l_t} \right)^{1-\rho} \cdot \left(1 - \frac{l_t^{1-\theta}}{\Gamma} \right)^\rho.$$

The first term, MPL_t , is the standard marginal product of labor that equals the real wage in the RBC model. This term decreases after the negative productivity shock as in the standard model. The second term, Θ_t^ρ , is the externality term. This term increases as the unemployment rate goes up because every firm has to spend less to hire more workers if required. In equilibrium, these two effects almost cancel each other, and the overall effect on the real wage is very small. Although the intuition behind the persistent effects of TFP shocks is similar to models with rigid wages, the difference is that real rigidity appears endogenously in my model.⁴⁴

Finally, I discuss why the model with added belief shocks can explain the entire increase in unemployment in 2008 and its slow recovery after that. Because the only variables that are significantly affected by a belief shock are the real wage and the unemployment rate, variables such as real output, investment and

⁴³In contrast to the rest of the IRFs reported, I report the effect on the unemployment rate in percent of the labor force, not as a percent of the steady state value.

⁴⁴This is in contrast to the literature where rigid wages are imposed as an assumption. See for example Shimer (2012), Hall (2005).

consumption change by a magnitude that is 20 times less than the initial shock magnitude, I do not explicitly plot IRFs for a one-time belief shock.

The effects of one-time negative 1 percent belief shock are simple: both the unemployment rate and the real wage increase by 1 percent and the increase is permanent. A negative belief shock moves the economy to a new steady state with a higher unemployment rate while keeping the ratio of consumption to output constant (see Equation (36).)

Why does the real wage go up while neither real output nor its components are affected by belief shocks? The first thing to notice is that belief shocks act as changes in consumption demand. Intuitively, the multiple equilibrium model is close to the standard model with fixed aggregate supply. In such a model, if aggregate demand drops, the price drops and there is no change in real quantities.

This is exactly what happens in my model as well. The only difference is that because aggregate consumption is expressed relative to real wages (see the belief function equation (26)-(27)) the price corresponds to the inverse of the real wage. Therefore, a price drop implies an increase in the real wage.

In my model the unemployment rate exhibits hysteresis. These dynamics explain the phenomenon that has been dubbed “jobless recoveries” – a high persistent unemployment rate and real growth at the same time. To see this, consider a negative belief shock. In response to a shock of this kind, unemployment will rise proportionally and stay high for a potentially infinite period of time. At the same time there will be no decrease in real output, consumption, or investment. If a positive productivity shock follows such a belief shock, the economy will converge back to a new steady state with the new higher unemployment rate.

In my model a negative belief shock causes an increase in real wages. This is what we observe in the data. Between 2006 and 2011, real hourly earnings increased by around 3 percent while hours worked and the employment rate dropped significantly (see Kocherlakota (2012)).

6 Conclusion

In this paper I constructed and estimated a general equilibrium rational expectations model that displays hysteresis in unemployment as in Blanchard and Summers (1987). My model can explain both “jobless” or post-1990 recoveries as well as postwar U.S recoveries prior to 1990.

My model has the potential to explain the prolonged increase in the unemployment rate that we observed in the U.S. after the Great Recession as a socially inefficient equilibrium outcome. Additionally, the model can reproduce stylistic business cycle dynamics in output, consumption, investment and real wages over the entire postwar period. In particular, in contrast to nearly all labor search models, in my model real wages are slightly countercyclical as in the data.

My model advances on two other branches of the literature. First, it expands second generation endogenous business cycle models to include the behavior of investment, consumption and real wages together with output and unemployment.⁴⁵ Second, it provides deeper micro-foundations for why real wages are rigid with respect to a TFP shock. In contrast to the existing rigid wages literature (see Hall (2005), Shimer (2005)), I obtain this property of the real wage without directly imposing an assumption of nominal or real rigidity.

My model is driven by two shocks: a TFP shock that explains fluctuations at business cycle frequencies and a belief shock that explains permanent shifts in the unemployment rate. Productivity shocks have similar transitory effects as in a standard model. But they also have long-lasting effects on the unemployment rate, produced by an endogenous propagation mechanism. Together, demand and supply shocks generate jobless recoveries, and a high persistent unemployment rate can coexist with positive real growth.

The model does not contradict the view that productivity shocks are a major source of economic fluctuations as in Nelson and Plosser (1982) and Long and Plosser (1983). Rather the model enriches this view by combining it with the idea that the unemployment rate displays the property of hysteresis.

I find that neither demand nor supply shocks alone can explain the observed dynamics of employment and GDP in the postwar data. With no demand disturbances one is likely to come across an “unemployment volatility puzzle” because the unemployment rate is a lot more volatile than the data generated by the model. With no TFP shocks the model cannot explain the observed volatility of investment.

What is a belief shock? Farmer (2012b) argues that the stock market crash of 2008 was similar in nature to what Keynes (1936) referred to as “animal spirits”. Stock market wealth and housing wealth both fell

⁴⁵See Farmer (2012a) for a recent review of endogenous business cycle literature

sharply in 2008. This drop in wealth, caused by self-fulfilling shock to beliefs about permanent income, can partially explain the Great Recession.

Acknowledgements

I wish to to thank Ariel Burstein, Francisco Buera, Andrew Atkeson, Pablo Fajgelbaum, Jang-Ting Guo, Lee Ohanian, Aaron Tornell and Pierre-Olivier Weill for their detailed comments as well as participants of the Monetary Economics Proseminar at UCLA, NBER Summer Institute and Society for Economic Dynamics, ICD and RES seminars at the IMF for their helpful feedback. I am especially thankful to Roger Farmer for his time and invaluable advice. All errors are my own.

References

- ADJEMIAN, S., H. BASTANI, M. JUILLARD, F. MIHOUBI, G. PERENDIA, M. RATTO, AND S. VILLEMOT (2011): “Dynare: Reference Manual, Version 4,” *Dynare Working Papers, CEPREMAP*, 1.
- AN, S., AND F. SCHORFHEIDE (2007): “Bayesian analysis of DSGE models,” *Econometric Reviews*, 6, 113–172.
- BALL, L. (2009): “Hysteresis in Unemployment: Old and New Evidence,” *NBER Working Papers 14818*.
- BENHABIB, J., AND R. E. A. FARMER (1994): “Indeterminacy and Increasing Returns,” *Journal of Economic Theory*, 63, 19–46.
- BERGER, D. (2013): “Countercyclical Restructuring and Jobless Recoveries,” *Northwestern mimeo*.
- BLANCHARD, O., AND J. GALI (2008): “Labor Markets and Monetary Policy: A New-Keynesian Model with Unemployment,” *MIT Department of Economics Working Paper*, No. 06-22.
- BLANCHARD, O. J., AND D. QUAH (1989): “The Dynamic Effects of Aggregate Demand and Supply Disturbances,” *The American Economic Review*, 79, 655–673.
- BLANCHARD, O. J., AND L. H. SUMMERS (1986): “Hysteresis and the European Unemployment Problem,” in *NBER Macroeconomics Annual*, vol. 1, pp. 15–90. National Bureau of Economic Research.
- (1987): “Hysteresis in Unemployment,” *European Economic Review*, 31, 288–295.
- BLUNDELL, R., AND T. H. MACURDY (1999): “Labor Supply: A Review of Alternative Approaches,” *Handbook of Labor Economics*, 3, 1559–1695.
- CHARI, V. V., P. KEHOE, AND E. MCGRATTAN (2007): “Business Cycle Accounting,” *Econometrica*, 75(3), 781–836.
- EVANS, G. W., AND S. HONKAPOHJA (1993): “Adaptive forecasts, hysteresis and endogenous fluctuations,” *Federal Reserve Bank of San Francisco Economic Review*, 1, 3–13.
- (2000): “Expectations and the Stability Problem for Optimal Monetary Policies,” University of Oregon, mimeo.
- EVANS, G. W., AND G. RAMEY (2006): “Adaptive Expectations, Underparameterization and the Lucas Critique,” *Journal of Monetary Economics*, 53, 249–264.
- FARMER, R., C. NOURRY, AND A. VENDITTI (2013): “The Inefficient Markets Hypothesis: Why Financial Markets Do Not Work Well in the Real World,” *NBER Working paper 18647*.
- FARMER, R. E. A. (2002): “Why Does the Data Reject the Lucas Critique?,” *Annales d’Economie et de Statistiques*, (67-68), 111–129.
- (2006): “Old Keynesian Economics,” *UCLA mimeo*.
- (2010): *Expectations, Employment and Prices*. Oxford University Press.
- (2011a): “Animal Spirits, Rational Bubbles and Unemployment in an Old-Keynesian Model,” *NBER WP 17137*.
- (2011b): “Confidence, Crashes and Animal Spirits,” *Economic Journal*.
- FARMER, R. E. A. (2012a): “The Evolution of Endogenous Business Cycles,” *CEPR Discussion Paper 9080*.
- (2012b): “The stock market crash of 2008 caused the Great Recession: Theory and evidence,” *Journal of Economic Dynamics & Control*, 36, 693–707.

- FARMER, R. E. A., AND J. T. GUO (1994): “Real Business Cycles and the Animal Spirits Hypothesis,” *Journal of Economic Theory*, 63, 42–73.
- FARMER, R. E. A., AND D. PLOTNIKOV (2012): “Fiscal policy in the model with hysteresis,” *UCLA mimeo*.
- FERNANDEZ-VILLAYERDE, J., AND J. RUBIO-RAMIREZ (2005a): “Comparing Dynamic Equilibrium Models to Data: A Bayesian Approach,” *Journal of Econometrics*, 123, 891–910.
- (2005b): “Estimating Dynamic Equilibrium Economies: Linear Versus Non-Linear Likelihood,” *Journal of Applied Econometrics*, 20, 891–910.
- FRIEDMAN, M. (1957): *A Theory of the Consumption Function*. Princeton University Press.
- GALÍ, J., F. SMETS, AND R. WOUTERS (2012): “Slow Recoveries: A Structural Interpretation,” *Journal of Money, Credit and Banking*, 44 (2), 9–30.
- HALL, R. E. (2005): “Employment Fluctuations with Equilibrium Wage Stickiness,” *American Economic Review*, 95(1), 50–65.
- HALL, R. E. (2011): “The Long Slump,” *American Economic Review*, 101(2), 431–69.
- HANSEN, G. D. (1985): “Indivisible Labor and the Business Cycle,” *Journal of Monetary Economics*, 16, 309–327.
- JAIMOVICH, N., AND H. SIU (2013): “The Trend is the Cycle: Job Polarization and Jobless Recoveries,” *Duke mimeo*.
- JONG, P. D. (1991): “The Diffuse Kalman Filter,” *The Annals of Statistics*, Volume 2, No 2, 1073–1083.
- JUSTINIANO, A., G. E. PRIMICERI, AND A. TAMBALOTTI (2010): “Investment Shocks and Business Cycles,” *Journal of Monetary Economics*, 57, 132–145.
- KEYNES, J. M. (1936): *The General Theory of Employment, Interest and Money*. MacMillan and Co.
- KLUMP, R., P. MCADAM, AND A. WILLMAN (2007): “Factor Substitution and Factor-Augmenting Technical Progress in the United States: a Normalized Supply-Side System Approach,” *The Review of Economics and Statistics*, 89(1), 183–192.
- KOCHERLAKOTA, N. (2012): “Incomplete Labor Markets,” *Research Papers of Minneapolis Federal Reserve*.
- KOOPMAN, S. J. (1997): “Exact Initial Kalman Filtering and Smoothing for Nonstationary Time Series Models,” *Journal of the American Statistical Association*, 92, No. 440 (Dec., 1997), 1630–1638.
- KOOPMAN, S. J., AND J. DURBIN (2000): “Fast Filtering And Smoothing For Multivariate State Space Models,” *Journal of Time Series Analysis*, 21, No. 3, 281–296.
- LONG, J. B., AND C. I. PLOSSER (1983): “Real Business Cycles,” *Journal of Political Economy*, 91(1), 39–69.
- LUDVIGSON, S. C., AND M. LETTAU (2004): “Understanding Trend and Cycle in Asset Values: Reevaluating the Wealth Effect on Consumption,” *American Economic Review*, 94(1), 277–299.
- MORTENSEN, D. T., AND E. NAGYPAL (2007): “More on Unemployment and Vacancy Fluctuations,” *Review of Economic Dynamics*, 10, 327–347.
- MUTH, J. F. (1961): “Rational Expectations and the Theory of Price Movements,” *Econometrica*, 29, 315–335.
- NADIRI, M. I., AND I. R. PRUCHA (1997): “Estimation of the Depreciation Rate of Physical and R&D Capital in the U.S. Total Manufacturing Sector,” *NBER Working Paper No. 4591*.

- NELSON, C. R., AND C. PLOSSER (1982): “Trends and random walks in macroeconomic time series: Some evidence and implications,” *Journal of Monetary Economics*, 10, 139–162.
- PLOTNIKOV, D. (2012): “Can an RBC model match medium-term business cycle fluctuations in the U.S. data?,” *UCLA working paper*.
- REINHART, C., AND K. ROGOFF (2009): *This Time is Different*. Princeton University Press.
- ROGERSON, R., R. SHIMER, AND R. WRIGHT (2005): “Search-Theoretic Models of the Labor Market: A Survey,” *Journal of Economic Literature*, 43, 959–988.
- SHILLER, R. J. (2000): *Irrational Exuberance*. Princeton University Press.
- SHIMER, R. (2005): “The Cyclical Behavior of Equilibrium Unemployment and Vacancies,” *American Economic Review*, 95(1), 25–49.
- SHIMER, R. (2012): “Wage Rigidities and Jobless Recoveries,” *Journal of Monetary Economics*, 59(Supplement), S65–S77.
- SIMS, C. A. (2001): “Solving Linear Rational Expectations Models,” *Journal of Computational Economics*, 20(1-2), 1–20.
- WICZER, D. (2013): “Long-Term Unemployment: Attached and Mismatched?,” *University of Minnesota Working Paper*.