



Contents lists available at ScienceDirect

Review of Economic Dynamics

www.elsevier.com/locate/red


Involuntary unemployment and the business cycle [☆]

 Lawrence J. Christiano ^{a,*}, Mathias Trabandt ^{b,c,*}, Karl Walentin ^{d,*}

^a Northwestern University, Department of Economics, 2001 Sheridan Road, Evanston, IL 60208, USA

^b Freie Universität Berlin, School of Business and Economics, Chair of Macroeconomics, Boltzmannstrasse 20, 14195 Berlin, Germany

^c German Institute for Economic Research (DIW) and Halle Institute for Economic Research (IWH), Germany

^d Sveriges Riksbank, Research Division, 103 37 Stockholm, Sweden

ARTICLE INFO

Article history:

Received 10 September 2018

Received in revised form 11 May 2020

Available online 22 May 2020

JEL classification:

E2

E3

E5

J2

J6

Keywords:

DSGE

Unemployment

Labor force participation

Business cycles

Monetary policy

Bayesian estimation

ABSTRACT

Can a model with limited labor market insurance explain standard macro and labor market data *jointly*? We construct a monetary model in which: i) the unemployed are worse off than the employed, i.e. unemployment is involuntary and ii) the labor force participation rate varies with the business cycle. To illustrate key features of our model, we start with the simplest possible framework. We then integrate the model into a medium-sized DSGE model and show that the resulting model does as well as existing models at accounting for the response of standard macroeconomic variables to monetary policy shocks and two technology shocks. In addition, the model does well at accounting for the response of the labor force and unemployment rate to these three shocks.

© 2021 The Authors. Published by Elsevier Inc. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Can a model with limited labor market insurance explain standard macro and labor market data *jointly*? To answer this question, we construct a monetary model in which the unemployed are worse off than the employed, i.e. unemployment is involuntary and the labor force participation rate varies with the business cycle. We investigate whether the resulting model fits standard real and nominal macro data and unemployment and labor force participation data in response to monetary policy and technology shocks.¹

[☆] Many people and seminar participants provided us with excellent comments for which we are very grateful. The views expressed in this paper are those of the authors and do not necessarily reflect those of Sveriges Riksbank.

* Corresponding authors.

E-mail addresses: l-christiano@northwestern.edu (L.J. Christiano), mathias.trabandt@gmail.com (M. Trabandt), karl.walentin@riksbank.se (K. Walentin).

¹ We are interested in a monetary environment since it allows us to study the general equilibrium repercussions between e.g. unemployment, inflation and nominal interest rates. In addition, monetary models such as Christiano et al. (2005, CEE) and Altig et al. (2011, ACEL) have proved to be useful to account for VAR-based evidence for real and nominal variables in response to monetary as well as technology shocks. The model features developed in CEE and ACEL have become standard ingredients in modern business cycle models, see e.g. Smets and Wouters (2003, 2007) and many others. Integrating our model of unemployment into such an environment therefore provides a useful empirical test for our approach to the labor market in general.

<https://doi.org/10.1016/j.red.2020.05.003>

1094-2025/© 2021 The Authors. Published by Elsevier Inc. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

Recently, the unemployment rate and the labor force participation rate have been discussed prominently in the light of the Great Recession. A shortcoming of standard monetary dynamic stochastic general equilibrium (DSGE) models is that they are silent about these important variables. Work has begun on the task of introducing unemployment into monetary DSGE models. The Diamond-Mortensen-Pissarides search and matching approach of unemployment represents a leading framework and has been integrated into monetary models by a number of authors.²

However, the approaches taken to date have several important shortcomings. First, they assume the existence of perfect consumption insurance against labor market outcomes, so that consumption is the same for employed and non-employed workers.³ With this kind of insurance, a worker is delighted to be unemployed because it is an opportunity to enjoy leisure without a drop in consumption.⁴ In other words, unemployment in these models is voluntary rather than involuntary. Second, it is generally assumed that labor force participation is constant and exogenous. This assumption was challenged during the recovery after the Great Recession, when policymakers saw that much of the initial drop in unemployment reflected a puzzling drop in labor force participation. This increased demand from policymakers for models that capture the business cycle dynamics in labor force participation (for further discussion, see Christiano et al. (2015)).

To remedy these limitations, we pursue an approach to model the labor market that has not been used in the monetary DSGE literature. Our approach follows the work of Hopenhayn and Nicolini (1997) and others, in which finding a job requires exerting a privately observed effort.⁵ In this type of environment, the higher utility enjoyed by employed workers is necessary for people to have the incentive to search for and keep jobs.⁶ Moreover, our approach implies that workers take an optimal decision whether or not to join the labor force. In other words, the labor force participation margin in our framework responds endogenously to business cycle shocks.

We define unemployment the way it is defined by the agencies that collect the data. To be officially unemployed a person must assert that she (i) has recently taken concrete steps to secure employment and (ii) is currently available for work.⁷ To capture (i) we assume that people who wish to be employed must undertake a costly effort. Our model has the implication that a person who asserts (i) and (ii) enjoys more utility if she finds a job than if she does not, i.e., unemployment is involuntary. Empirical evidence appears to be consistent with the notion that unemployment is in practice more of a burden than a blessing.⁸ For example, Chetty and Looney (2007) and Gruber (1997) find that U.S. workers suffer roughly a 10 percent drop in consumption when they lose their job. Also, there is a substantial literature which purports to find evidence that insurance against labor market outcomes is imperfect. An early example is Cochrane (1991). These observations motivate our third defining characteristic of unemployment: (iii) a person looking for work is worse off if they fail to find a job than if they find one.⁹

To highlight the mechanisms in our model, we first introduce it into the simplest possible framework. In our model, workers gather into “households” for the purpose of partially ensuring themselves against bad labor market outcomes. We regard the “household” as a label or stand-in for all the various market and non-market arrangements that actual workers have for dealing with idiosyncratic labor market outcomes.¹⁰ In line with this view of the household, workers are assumed to have no access to loan markets, while households have access to complete markets.

Each worker experiences a privately observed shock that determines its aversion to work. Workers that experience a sufficiently high aversion to work stay out of the labor force. The other workers join the labor force and are employed with

² Examples include Blanchard and Galí (2010), Campolmi and Gnocchi (2016) Christiano et al. (2008), Christiano et al. (2011b), Christoffel et al. (2009a), Christoffel and Kuester (2008), Christoffel et al. (2009b), den Haan et al. (2000), Gertler et al. (2008), Goshenny (2009), Krause et al. (2008), Lechthaler et al. (2010), Sala et al. (2008), Sveen and Weinke (2008, 2009), Thomas (2011), Trigari (2009) and Walsh (2005).

³ Gornemann et al. (2016) model unemployment in a search and matching framework allowing for household heterogeneity in terms of wealth in an incomplete-market setting, i.e. emphasizing self-insurance against unemployment while abstracting from the labor force participation decision. The main focus of Gornemann et al. (2016) is on the distributional effects of monetary policy.

⁴ The drop in utility reflects that models typically assume preferences that are additively separable in consumption and labor or that have the King et al. (1988) form. Examples include all papers cited in footnote 2.

⁵ An early paper that considers unobserved effort is Shavell and Weiss (1979). Our approach is also closely related to the efficiency wage literature, as in Alexopoulos (2004). The present paper is also related to Landais et al. (2012) who study the cyclical nature of optimal unemployment insurance in a real model with imperfect labor market insurance. However, the authors barely spell out the macro implications of their approach. In contrast to these authors, we study the implications of limited labor market insurance in a monetary model and, more importantly, examine the ability of the approach to explain actual macro- and labor market data in response to technology and monetary policy shocks quantitatively.

⁶ Lack of perfect insurance in practice probably reflects other factors too, such as adverse selection. Alternatively, Kocherlakota (1996) explores lack of commitment as a rationale for incomplete insurance. Lack of perfect insurance is not necessary for the unemployed to be worse off than the employed (see Rogerson and Wright, 1988).

⁷ See the Bureau of Labor Statistics website, http://www.bls.gov/cps/cps_htgm.htm#unemployed, for an extended discussion of the definition of unemployment, including the survey questions used to determine a household's employment status.

⁸ There is a substantial sociological literature that associates unemployment with an increased likelihood of suicide and domestic violence.

⁹ Although the great majority of monetary DSGE models that we know of fail (iii), they do not fail (ii). In these models there are workers who are not employed and who would say ‘yes’ in response to the question, ‘are you currently available for work?’. Although such people in effect declare their willingness to take an action that reduces utility, they would in fact do so. This is because they are members of a large household insurance pool. They obey the household's instruction that they value a job according to the value assigned by the household, not themselves. In these models everything about the individual worker is observable to the household, and it is implicitly assumed that the household has the technology necessary to enforce verifiable behavior. In our environment - and we suspect this is true in practice - the presence of private information makes it impossible to enforce a labor market allocation that does not completely reflect the preferences of the individual worker.

¹⁰ Alternative labels in this regard would be “a zero profit insurance company”, “the government”, “a social planner” or “a representative agent”.

a probability that is an increasing function of a privately observed effort. The only thing about a worker that is observed is whether or not it is employed. Although consumption insurance is desirable in our environment, perfect insurance is not feasible because everyone would claim high work aversion and stay out of the labor force.

For simplicity we suppose the wage rate is determined competitively so that firms and households take it as given.¹¹ Firms face no search frictions and hire workers up to the point where marginal costs and benefits are equated. But it is important to note that our modeling approach in principle could encompass search frictions for firms and wage bargaining, and that the friction that we emphasize – workers have to make a job finding effort which is unobservable – might well be viewed as a complement to the currently dominating paradigm. At this point it is worth emphasizing that unemployment in our model is purely frictional. It is not generated by unions or other factors pushing up the general wage level to a point where supply exceeds demand. However, note that our environment is flexible enough to allow for market power in the labor market, as will be the case in the estimated medium-sized DSGE model that we present in section 3.

Although individual workers face uncertainty, households are sufficiently large that there is no uncertainty at the household level. Once the household sets incentives by allocating more consumption to employed workers than to non-employed workers, it knows exactly how many workers will find work. The household takes the wage rate as given and adjusts employment incentives until the marginal cost (in terms of foregone leisure and reduced consumption insurance) of additional market work equals the marginal benefit. The firm and household first order necessary conditions of optimization are sufficient to determine the equilibrium wage rate. It turns out that our environment has a simple representative agent formulation, in which the representative household has an indirect utility function that is a function only of market consumption and labor.

Our theory of unemployment has interesting implications for the optimal variation of labor market insurance over the business cycle. In a boom more labor is demanded by firms. To satisfy the higher demand, the household provides workers with more incentives to look for work by raising consumption for the employed, c_t^w , relative to consumption of the non-employed, c_t^{nw} . Conversely, in a recession, the consumption premium falls and thus the replacement ratio, c_t^{nw}/c_t^w , increases. Thus, our model implies a procyclical consumption premium – or equivalently – a countercyclical replacement ratio. Put differently, optimal labor market insurance is countercyclical in our model.

Next, we introduce our model of unemployment into a medium-sized monetary DSGE model that has been fit to data. In particular, we work with a version of the model proposed in Christiano et al. (2005) (CEE). In this model there is monopoly power in the setting of wages, there are wage setting frictions, capital accumulation and other features.¹² We estimate and evaluate our model using the Bayesian version of the impulse response matching procedure proposed in Christiano et al. (2011a) (CTW). The impulse response methodology has proved useful in the basic model formulation stage of model construction, and this is why we use it here. The three shocks we consider are the same ones as in ACEL. In particular, we consider VAR-based estimates of the impulse responses of macroeconomic variables to a monetary policy shock, a neutral technology shock and an investment-specific technology shock. Our model can match the impulse responses of standard macro variables as well as the standard model, i.e. the model in CEE and ACEL. However, our model also does a good job matching the responses of the labor force and unemployment to the three shocks.

Our paper emphasizes the importance of labor supply for the dynamics of unemployment and the labor force and is thereby related to Galí (2011). In his model, the presence of unemployment rests on the assumption of market power in the labor market. By contrast, in our model unemployment reflects frictions that are necessary for people to find jobs and does not require monopoly power. Further, Galí (2011) assumes i) that available jobs can be found without effort and ii) the presence of perfect labor market insurance which implies that the employed have lower utility than the non-employed, i.e. unemployment is voluntary from an individual workers perspective. Finally, Galí's theory of unemployment with standard preferences implies a drop of labor supply in response to an expansionary monetary policy shock.¹³ The drop in labor supply is counterfactual, according to our VAR-based evidence. We estimate the standard model that contains Galí's theory of unemployment with and without imposing data for unemployment and the labor force. In both cases, our model of involuntary unemployment outperforms the standard model in terms of data fit.

These results highlight another important implication of our work. In particular, it is in general not sufficient to account only for the response of employment or total hours to be able to draw conclusions about the unemployment rate. In particular, when the standard model is estimated without data on unemployment and the labor force, the fit of total hours of the model is in fact very good. By contrast, the implications of the model for unemployment and the labor force are counterfactual. Conversely, when the standard model is estimated on unemployment and labor force data too, the fit of these two variables indeed improves somewhat. However the improvement of fit comes at the cost of not fitting total hours well. In other words, the standard model provides an example that it is not straightforward to account for the dynamics of unemployment and labor force participation *jointly* with other standard macroeconomic variables. By contrast, our model does a good job in this regard.

¹¹ One interpretation of our environment is that job markets occur on Lucas-Phelps-Prescott type islands. Effort is required to reach those islands, but a person who arrives at the island finds a perfectly competitive labor market. For recent work that uses a metaphor of this type, see Veracierto (2008).

¹² The model of wage setting is the one proposed in Ercog et al. (2000).

¹³ This drop in labor supply, or the labor force, is induced by the positive wealth effect. Galí (2011) and Galí et al. (2011) show that changes to the household utility function that offset wealth effects reduce the counterfactual implications of the standard model for the labor force.

Finally, our model of unemployment has several interesting microeconomic implications. As mentioned above, the consumption premium is procyclical while the replacement ratio is countercyclical. Studies of the cross-sectional variance of log household consumption are a potential source of evidence on the cyclical behavior of the premium. Evidence in Heathcote et al. (2010) suggests indeed that the dispersion in log household non-durable consumption decreased in the 1980, 2001 and 2007 recessions. Thus, the observed cross sectional dispersion of consumption across households lends support to our model's implication that the consumption premium is procyclical. Another indication that the replacement ratio may indeed be countercyclical is the fact that the duration of unemployment benefits is routinely extended in recessions, e.g. in the U.S. during the Great Recession. Second, our model predicts that high unemployment in recessions reflects the procyclicality of effort in job search. There is some evidence that supports this implication of the model. Data from the Bureau of Labor Statistics suggests that the number of “discouraged workers” jumped 70 percent from 2008Q1 to 2009Q1. The number of discouraged workers is only a tiny fraction of the labor force. However, to the extent that the sentiments of discouraged workers are shared by workers more generally, a jump in the number of discouraged workers could be a signal of a general decline in job search intensity in recessions.

The organization of the paper is as follows. The next section lays out our basic model of limited labor market insurance. Section 3 proceed by integrating our model into the medium-sized DSGE (CEE) model. After that, in section 4, we describe our estimation method. Section 5 reports the estimation results for our model. Moreover, section 6 discusses some microeconomic implications and examines evidence that provides tentative support for the model. The paper ends with concluding remarks.

2. Model of limited labor market insurance

We begin by describing the physical environment of a typical worker. Workers are subject to two sources of privately observed idiosyncratic uncertainty: a shock to work aversion, l , and the uncertainty of finding employment for workers that participate in the labor market. In this environment, there is a need for insurance, but insurance cannot be perfect because of the presence of asymmetric information. With standard separable preferences in consumption and leisure, under perfect insurance all workers would enjoy the same level of consumption, regardless of their realized value of l and of whether or not they find employment. Under this first-best insurance arrangement, workers would have no incentive to participate in the labor market and if they did, they would then have no incentive to exert effort in finding work. Instead, we consider the optimal insurance arrangement in the presence of asymmetric information. The optimal insurance contract balances the trade-off between incentive and insurance provision.

Under the insurance arrangement, workers band together into large households. Individual workers have no access to credit or insurance markets other than through their arrangements with the household. In part, we view the household construct as a stand-in for the market and non-market arrangements that actual workers use to insure against idiosyncratic labor market experiences. In part, we are following Andolfatto (1996) and Merz (1995), in using the household construct as a technical device to prevent the appearance of difficult-to-model wealth dispersion among workers. Households have sufficiently many members that there is no idiosyncratic household-level labor market uncertainty. The environment is sufficiently simple that we can obtain an analytic representation for the equally weighted utility of all the workers in a household. This utility function corresponds to the preferences in a representative agent formulation of our economy. At the end of this section, we discuss some important implications of our basic model structure.

2.1. Workers

The economy consists of a continuum of households. In turn, each household consists of a continuum of workers. A worker can either work, or not.¹⁴ At the start of the period, each worker draws a privately observed idiosyncratic shock, l , from a stochastic process with support, $[0, 1]$.¹⁵ We assume the stochastic process for l exhibits dependence over time and that its invariant distribution is uniform. A worker's realized value of l determines her utility cost of working:

$$\zeta (1 + \sigma_L) l^{\sigma_L}. \quad (2.1)$$

The parameters ζ and $\sigma_L \geq 0$ are common to all workers. In (2.1) we have structured the utility cost of employment so that σ_L affects its variance in the cross section and not its mean.¹⁶

After drawing l , a worker decides whether or not to participate in the labor force. In case a worker chooses non-participation, her utility is simply:

$$\ln(c_t^{nw} - bC_{t-1}). \quad (2.2)$$

¹⁴ In assuming that labor is indivisible, we follow Hansen (1985) and Rogerson (1988). The indivisible labor assumption has attracted substantial attention recently. See, for example, Mulligan (2001), and Krusell et al. (2008, 2011). The labor indivisibility assumption is consistent with the fact that most variation in total hours worked over the business cycle reflects variations in number of people employed, rather than in hours per worker.

¹⁵ A recent paper which emphasizes a richer pattern of idiosyncracies at the individual firm and household level is Brown et al. (2015).

¹⁶ To see this, note that $\int_0^1 (1 + \sigma_L) l^{\sigma_L} dl = 1$ and $\int_0^1 [(1 + \sigma_L) l^{\sigma_L} - 1]^2 dl = \frac{\sigma_L^2}{1+2\sigma_L}$.

The term bC_{t-1} reflects habit persistence in consumption at the household level which the worker takes as given. A non-participating worker does not experience any disutility from work or from exerting effort to find a job.

The probability that a worker which participates in the labor market finds work is $p(e_{l,t}; \tilde{\eta}_t)$ where $e_{l,t} \geq 0$ is a privately observed level of effort expended by the worker. Let:

$$\tilde{p}(e_{l,t}; \tilde{\eta}_t) = \tilde{\eta}_t + ae_{l,t} \quad (2.3)$$

where $a > 0$. The sign of a implies that the marginal product of effort is non-negative. Further,

$$\tilde{\eta}_t = \eta + \mathcal{M}(\bar{m}_t/\bar{m}_{t-1}) \quad (2.4)$$

where $\eta < 0$. We discuss the negative sign of η below. The function $\mathcal{M}(\bar{m}_t/\bar{m}_{t-1})$ reflects the impact of aggregate economic conditions – in particular the change of the aggregate labor force \bar{m}_t/\bar{m}_{t-1} – on the worker's probability to find work. We will discuss details about the function \mathcal{M} in subsection 3.4 and in subsection B.6 in the technical appendix.¹⁷

In order to preserve analytic tractability, we assume the following piecewise linear specification for the probability of finding work:

$$p(e_{l,t}; \tilde{\eta}_t) = \begin{cases} 1 & \tilde{p}(e_{l,t}; \tilde{\eta}_t) > 1 \\ \tilde{p}(e_{l,t}; \tilde{\eta}_t) & 0 \leq \tilde{p}(e_{l,t}; \tilde{\eta}_t) \leq 1 \\ 0 & \tilde{p}(e_{l,t}; \tilde{\eta}_t) < 0 \end{cases} \quad (2.5)$$

A worker whose work aversion is l and which participates in the labor market and exerts effort $e_{l,t}$ enjoys the following expected utility:

$$p(e_{l,t}; \tilde{\eta}_t) \overbrace{\left[\ln(c_t^w - bC_{t-1}) - \varsigma(1 + \sigma_L)l^{\sigma_L} - \frac{1}{2}e_{l,t}^2 \right]}^{\text{ex post utility of worker that joins labor force and finds a job}} \\ + (1 - p(e_{l,t}; \tilde{\eta}_t)) \overbrace{\left[\ln(c_t^{nw} - bC_{t-1}) - \frac{1}{2}e_{l,t}^2 \right]}^{\text{ex post utility of worker that joins labor force and fails to find a job}} \quad (2.6)$$

Here, $e_{l,t}^2/2$ is the utility cost associated with effort. In (2.6), c_t^w and c_t^{nw} denote the consumption of employed and non-employed workers, respectively. These are outside the control of a worker and are determined in equilibrium given the arrangements which we describe below. In addition, $\tilde{\eta}_t$ is also outside the control of a worker. Our notation reflects that in our environment, an individual worker's consumption can only be dependent on its current employment status because this is the only worker characteristic that is publicly observed.¹⁸

We now characterize the effort and labor force participation decisions of the worker. Because workers' work aversion type and effort choice are private information, their effort and labor force decisions are privately optimal conditional on c_t^{nw} and c_t^w . In particular, the worker decides its level of effort and labor force participation by comparing the magnitude of (2.2) with the maximized value of (2.6). In the case of indifference, we assume the worker chooses non-participation.

2.2. Characterizing worker behavior

As described above, the worker takes the replacement ratio $r_t \equiv c_t^{nw}/c_t^w < 1$ as given. The workers's utility of participating in the labor market, minus the utility, $\ln(c_t^{nw} - bC_{t-1})$, of non-participation is given by:

$$\max_{e_{l,t} \geq 0} f(e_{l,t}), \quad f(e_{l,t}) \equiv p(e_{l,t}; \tilde{\eta}_t) \left[\ln\left(\frac{c_t^w - bC_{t-1}}{c_t^{nw} - bC_{t-1}}\right) - \varsigma(1 + \sigma_L)l^{\sigma_L} \right] - \frac{1}{2}e_{l,t}^2.$$

Define $\tilde{r}_t = \frac{c_t^{nw} - bC_{t-1}}{c_t^w - bC_{t-1}}$ and note the distinction between this expression and the replacement ratio, r_t if $b > 0$. Then, the difference in utility can be expressed as follows:

$$\max_{e_{l,t} \geq 0} f(e_{l,t}), \quad f(e_{l,t}) \equiv p(e_{l,t}; \tilde{\eta}_t) \left[\ln(1/\tilde{r}_t) - \varsigma(1 + \sigma_L)l^{\sigma_L} \right] - \frac{1}{2}e_{l,t}^2. \quad (2.7)$$

We suppose that if more than one value of $e_{l,t}$ solves (2.7), then the worker chooses the smaller of the two. The worker chooses non-participation if the maximized value of (2.7) is smaller than, or equal to, zero.

¹⁷ The technical appendix is available at the following URL: http://sites.google.com/site/mathiastrabandt/home/downloads/CTWinvoluntary_techapp.pdf.

¹⁸ For example, we do not allow worker consumption allocations to depend upon the employment history or the history of worker reports of l . We make this assumption to preserve tractability. It would be interesting to investigate whether the results are very sensitive to our assumption that consumption is not allowed to depend on the individual history. We suspect that if the history of past reports were publicly known, then the difference between discounted utility when household types and labor effort are public or private would narrow (see, e.g., Atkeson and Lucas, 1995).

2.2.1. Optimal effort

It is convenient to consider a version of (2.7) in which the sign restriction on $e_{l,t} \geq 0$ is ignored and $p(e_{l,t}; \tilde{\eta}_t)$ in (2.7) is replaced with the linear function, $\tilde{p}(e_{l,t}; \tilde{\eta}_t)$ (see (2.3))¹⁹:

$$\max_{e_{l,t}} \tilde{f}(e_{l,t}; \tilde{\eta}_t, \tilde{r}_t), \quad \tilde{f}(e_{l,t}; \tilde{\eta}_t, \tilde{r}_t) \equiv \tilde{p}(e_{l,t}; \tilde{\eta}_t) [\ln(1/\tilde{r}_t) - \varsigma(1 + \sigma_L)l^{\sigma_L}] - \frac{1}{2}e_{l,t}^2. \tag{2.8}$$

The function, \tilde{f} , is quadratic with negative second derivative, and so the unique value of $e_{l,t}$ that solves the above problem is the one that sets the derivative of \tilde{f} to zero:

$$\tilde{e}_{l,t} = a [\ln(1/\tilde{r}_t) - \varsigma(1 + \sigma_L)l^{\sigma_L}]. \tag{2.9}$$

Substituting this expression into (2.8), we obtain:

$$\tilde{f}(\tilde{e}_{l,t}; \tilde{\eta}_t) = \frac{\tilde{e}_{l,t}}{2} \left[\frac{2}{a}\tilde{\eta}_t + \tilde{e}_{l,t} \right], \tag{2.10}$$

2.2.2. Optimal participation

There exists a unique $0 < l < 1$ such that the object in square brackets in (2.10) is zero. That value of l is the labor force participation rate, which we denote by m_t and which solves:

$$a [\ln(1/\tilde{r}_t) - \varsigma(1 + \sigma_L)m_t^{\sigma_L}] = -\frac{2}{a}\tilde{\eta}_t. \tag{2.11}$$

We can rewrite this indifference condition as:

$$\ln\left(\frac{c_t^w - bC_{t-1}}{c_t^{nw} - bC_{t-1}}\right) = \varsigma(1 + \sigma_L)m_t^{\sigma_L} - \frac{2}{a^2}\tilde{\eta}_t, \tag{2.12}$$

which implies that the utility of the marginal labor force participant from the extra consumption from working equals the disutility of working and searching for a job. Equation (2.11) can also be written as:

$$m_t = \left[\frac{\ln(1/\tilde{r}_t) + \frac{2}{a^2}\tilde{\eta}_t}{\varsigma(1 + \sigma_L)} \right]^{\frac{1}{\sigma_L}}. \tag{2.13}$$

Note that for all $l \geq m_t$ such that $\tilde{e}_{l,t} \geq 0$, $\tilde{f}(\tilde{e}_{l,t}; \tilde{\eta}_t, \tilde{r}_t) \leq 0$ and for all $l < m_t$, $\tilde{f}(\tilde{e}_{l,t}; \tilde{\eta}_t, \tilde{r}_t) > 0$. Furthermore, we impose the following restriction, which ensures an interior solution for the labor force participation rate (see section A.2.1 in the technical appendix for details):

$$a \ln(1/\tilde{r}_t) > -\frac{2}{a}\tilde{\eta}_t > a [\ln(1/\tilde{r}_t) - \varsigma(1 + \sigma_L)]. \tag{2.14}$$

We can then summarize our findings in the form of the following proposition:

Proposition 2.1. *Suppose that (2.14) is satisfied and the worker's objective is described in (2.8), with \tilde{r}_t taken as given by the worker. Let m_t be as defined by (2.13). Then workers with $m_t \leq l \leq 1$ choose non-participation and workers with $l < m_t$ and $\tilde{e}_{l,t} \geq 0$ choose participation. For those who choose participation, their effort level is given by (2.9).*

The previous proposition was derived under the assumption that the workers's objective is (2.8). We use the results based on (2.8) to understand the case of (2.7), i.e. without imposing $e_{l,t} \geq 0$ or linearity of $p(e_{l,t}; \tilde{\eta}_t)$. One can show that there is a largest value of l , denoted \dot{l}_t , such that for all $l \leq \dot{l}_t$, the constraint, $p(e_{l,t}; \tilde{\eta}_t) \leq 1$ is binding. In other words, there is a share of workers \dot{l}_t that has $p(e_{l,t}; \tilde{\eta}_t) = 1$. The cutoff, \dot{l}_t , solves:

$$p(e_{l,t}; \tilde{\eta}_t) = \tilde{\eta}_t + a^2 [\ln(1/\tilde{r}_t) - \varsigma(1 + \sigma_L)\dot{l}_t^{\sigma_L}] = 1,$$

or after making use of (2.11) to substitute out $\ln(1/\tilde{r}_t)$ and re-arranging:

$$\dot{l}_t = \left[m_t^{\sigma_L} - \frac{1 + \tilde{\eta}_t}{\varsigma(1 + \sigma_L)a^2} \right]^{\frac{1}{\sigma_L}}. \tag{2.15}$$

¹⁹ Considering the unconstrained case first will be helpful to understand more easily the constrained case, i.e. $e_{l,t} \geq 0$ and $0 \leq p(e_{l,t}; \tilde{\eta}_t) \leq 1$ which we characterize below.

2.3. Equilibrium conditions

Before deriving the representative household's utility, $u(C_t, h_t)$, as a function of household aggregate employment, h_t , and household aggregate consumption, C_t , it is useful to derive a few helpful equilibrium conditions. The number of employed workers, h_t , is, using our uniform distribution assumption of workers across l^{20} :

$$h_t = \int_0^{m_t} p(e_{l,t}; \tilde{\eta}_t) dl = \int_0^{\dot{l}_t} 1 dl + \int_{\dot{l}_t}^{m_t} \tilde{p}(e_{l,t}; \tilde{\eta}_t) dl \quad (2.16)$$

$$= -\tilde{\eta}_t m_t + a^2 \zeta \sigma_L m_t^{\sigma_L+1} + \dot{l}_t \left[1 + \tilde{\eta}_t + a^2 \zeta (1 + \sigma_L) \left(-m_t^{\sigma_L} + \frac{\dot{l}_t^{\sigma_L}}{1 + \sigma_L} \right) \right].$$

Combining the latter equation with expression (2.15) to substitute out for $1 + \tilde{\eta}_t$ and re-arranging yields:

$$h_t = -\tilde{\eta}_t m_t + a^2 \zeta \sigma_L (m_t^{\sigma_L+1} - \dot{l}_t^{\sigma_L+1}) \quad (2.17)$$

Note that one can use (2.15) to rewrite this expression for employment as:

$$h_t = -\tilde{\eta}_t m_t + a^2 \zeta \sigma_L \left(m_t^{\sigma_L+1} - \left[m_t^{\sigma_L} - \frac{1 + \tilde{\eta}_t}{\zeta (1 + \sigma_L) a^2} \right]^{\frac{\sigma_L+1}{\sigma_L}} \right) \equiv Q(m_t; \tilde{\eta}_t)$$

or

$$m_t = Q^{-1}(h_t; \tilde{\eta}_t), \quad (2.18)$$

where Q^{-1} is the inverse function of Q . We show in the technical appendix section A.8 that the mapping between h_t and m_t is unique. As we will see below, expression (2.18) will be useful when computing the household utility function.

Equation (2.17) can also be re-arranged to obtain the following expression for the unemployment rate:

$$u_t = \frac{m_t - h_t}{m_t} = \frac{m_t + \tilde{\eta}_t m_t - a^2 \zeta \sigma_L (m_t^{\sigma_L+1} - \dot{l}_t^{\sigma_L+1})}{m_t} = 1 + \tilde{\eta}_t - a^2 \zeta \sigma_L \frac{(m_t^{\sigma_L+1} - \dot{l}_t^{\sigma_L+1})}{m_t} \quad (2.19)$$

Suppose the household decides to send a measure, h_t , of workers to work and to consume C_t . The household that has chosen a level of employment, h_t , must set the labor force, m_t , to the level indicated by (2.18). To ensure that a measure, m_t , of workers has the incentive to enter the labor force requires setting the consumption premium as indicated by (2.11). Expression (2.11) determines the ratio of the consumption of employed and not employed workers. Given this ratio, the household's resource constraint,

$$h_t c_t^w + (1 - h_t) c_t^{nw} = C_t, \quad (2.20)$$

determines the level of c_t^w and c_t^{nw} . Solving this expression for c_t^{nw} yields

$$c_t^{nw} = \frac{r_t C_t}{h_t + (1 - h_t) r_t}. \quad (2.21)$$

The consumption of employed workers can then be obtained by using $c_t^w = c_t^{nw} / r_t$.

Note that in our environment there is no reason to describe a household optimization problem for selecting c_t^{nw} or c_t^w since there is only one setting for these variables that satisfies the resource constraint, (2.20), and the labor force participation constraint, (2.12).

2.4. Household utility function

The equally weighted utility of the workers within the household is given by:

$$u(C_t, h_t, m_t; C_{t-1}, \tilde{\eta}_t) = \int_0^{m_t} \left(p(e_{l,t}; \tilde{\eta}_t) \left[\ln(c_t^w - bC_{t-1}) - \zeta (1 + \sigma_L) l^{\sigma_L} \right] + (1 - p(e_{l,t}; \tilde{\eta}_t)) \left[\ln(c_t^{nw} - bC_{t-1}) - \frac{1}{2} e_{l,t}^2 \right] \right) dl + (1 - m_t) \ln(c_t^{nw} - bC_{t-1})$$

²⁰ See the technical appendix section A.7 for the intermediate steps in the derivation of the following equation.

We wish to express this as a function of C_t and h_t only, given C_{t-1} and $\tilde{\eta}_t$, using the results in the previous section. The derivation of the following expression for household utility is described in the technical appendix sections A.3–A.8:

$$u(C_t, h_t; C_{t-1}, \tilde{\eta}_t) = \ln(C_t - bC_{t-1}) - z(h_t; \tilde{\eta}_t), \quad (2.22)$$

where

$$\begin{aligned} z(h_t; \tilde{\eta}_t) = & \ln\left(h_t \left[\frac{1}{\tilde{r}_t} - 1\right] + 1\right) - \alpha_1 \tilde{\eta}_t m_t^{\sigma_L+1} - \alpha_2 m_t^{2\sigma_L+1} \\ & - \left[\alpha_3 (1 + \tilde{\eta}_t) m_t^{\sigma_L} - \alpha_2^{2\sigma_L} m_t + \alpha_4 (1 + \tilde{\eta}_t)^2\right] \times \left[m_t^{\sigma_L} - \frac{1 + \tilde{\eta}_t}{\zeta (1 + \sigma_L) a^2}\right]^{\frac{1}{\sigma_L}} \\ \tilde{r}_t = & e^{-\left[\zeta(1+\sigma_L)m^{\sigma_L} - \frac{2}{a^2}\tilde{\eta}_t\right]} \end{aligned} \quad (2.23)$$

and m_t is a function of h_t given by equation (2.18), $\tilde{\eta}_t$ is given by equation (2.4) and the expression for \tilde{r}_t can be obtained by rearranging (2.13).²¹

A notable feature of (2.22) is that consumption enters the household's utility function in the same way that it enters the individual worker's utility function. Moreover, consumption and employment are separable in utility.

It is useful to define a measure of the curvature of the function z in the neighborhood of the steady state:

$$\sigma_z \equiv \frac{z_{hh}h}{z_h} \quad (2.24)$$

Here z_h (z_{hh}) denotes the first (second) derivative of z with respect to h , evaluated in steady state. Note that $1/\sigma_z$ is the consumption-compensated elasticity of household labor supply in steady state. In our environment, all changes in labor supply occur on the extensive margin, so the empirical counterpart to $1/\sigma_z$ is the extensive-margin labor supply.

2.5. Implications of our basic model structure

We now briefly discuss expression (2.22) as well as implications of our basic model structure.

First, note that the derivation of the household utility function, (2.22), involves no explicit maximization problem even though the resulting insurance arrangement is optimal given our information assumption. This is because the household labor force participation and resource constraints, (2.11) and (2.20), are sufficient to determine c_t^w and c_t^{nw} conditional on h_t and C_t .

Second, we can see from (2.22) that our model is likely to be characterized by a particular observational equivalence property. To see this, note that although the agents in our model are in fact heterogeneous, C_t and h_t are chosen as if the economy were populated by a representative agent with the utility function specified in (2.22). A model such as the one in Clarida et al. (1999, henceforth CGG) which specifies representative agent utility as the sum of the log of consumption and a separable disutility of labor term is indistinguishable from our model, as long as data on the labor force and unemployment are not used. This is particularly obvious if, as is the case here, we only study the linearized dynamics of the model about the steady state. In this case, the only properties of a model's utility function that are used are its second order derivative properties in the nonstochastic steady state.

Third, our model and the standard CGG model are distinguished by the following two features: i) our model addresses a larger set of time series than the standard model does and ii) in our model the representative agent's utility function is a reduced form object. With respect to the utility function, its properties are determined by i) the details of the technology of job search, and ii) the cross-sectional variation in preferences with regard to attitudes about market work. As a result, the basic structure of the utility function in our model can in principle be informed by time use surveys and studies of job search.²²

Fourth, we gain insight into the determinants of the unemployment rate in the model by re-stating (2.19):

$$u_t = 1 + \tilde{\eta}_t - a^2 \zeta \sigma_L \frac{\left(m_t^{\sigma_L+1} - i_t^{\sigma_L+1}\right)}{m_t}. \quad (2.25)$$

Ceteris paribus, a rise in the labor force, m_t , is associated with a fall in the unemployment rate, u_t . To generate the fall in the unemployment rate, the rise in employment must be larger than the rise in the labor force. The greater rise in employment reflects that an increase in the labor force requires raising employment incentives, and this generates an increase in search intensity.

²¹ In the technical appendix, in equation (A.34), we express equation (2.23) in a way that is more useful for computational purposes.

²² A similar point was made by Benhabib et al. (1991). They argue that a representative agent utility function of consumption and labor should be interpreted as a reduced form object, after non-market consumption and labor activities have been maximized out. From this perspective, construction of the representative agent's utility function can in principle be guided by surveys of how time in the home is used.

Fifth, our theory of unemployment implies a procyclical consumption premium – or equivalently – a countercyclical replacement ratio. So see this, consider, for simplicity, a version of equation (2.12) without habit ($b = 0$) and constant $\tilde{\eta}_t = \tilde{\eta}$:

$$\ln \left(\frac{c_t^w}{c_t^{nw}} \right) = \varsigma (1 + \sigma_L) m_t^{\sigma_L} - \frac{2}{a^2} \tilde{\eta}. \quad (2.26)$$

This equation shows that when the labor force, m_t , increases in a boom, the equilibrium consumption premium, c_t^w/c_t^{nw} increases. The boom results in more labor demanded by firms. In order to satisfy the higher demand, the household provides workers with more incentives to look for work by raising consumption for the employed, c_t^w , relative to consumption of the non-employed, c_t^{nw} . Conversely, in a recession, the consumption premium falls and thus the replacement ratio $r_t = c_t^{nw}/c_t^w$ increases.²³ In other words, our model implies that workers are provided with more insurance in a recession, i.e. optimal labor market insurance is countercyclical.

3. Limited labor market insurance in a medium-sized DSGE model

Next, we show how we embed our model of limited labor market insurance into an otherwise standard medium-sized New Keynesian DSGE framework as e.g. CEE or Smets and Wouters (2003, 2007).

3.1. Final and intermediate goods

A final good is produced by a competitive, representative firm using a continuum of inputs as follows:

$$Y_t = \left[\int_0^1 Y_{i,t}^{\frac{1}{\lambda_f}} di \right]^{\lambda_f}, \quad 1 < \lambda_f. \quad (3.1)$$

The i^{th} intermediate good is produced by a monopolist with the following production function:

$$Y_{i,t} = (z_t H_{i,t})^{1-\alpha} K_{i,t}^\alpha - \phi_t, \quad (3.2)$$

where $K_{i,t}$ denotes capital services used for production by the i^{th} intermediate good producer. Also, $\log(z_t)$ is a technology shock whose first difference has a positive mean and ϕ_t denotes a fixed production cost which we will discuss below. The economy has two sources of growth: the positive drift in $\log(z_t)$ and a positive drift in $\log(\Psi_t)$, where Ψ_t is the state of an investment-specific technology shock discussed below. Let z_t^+ be defined as $z_t^+ = \Psi_t^{\frac{\alpha}{1-\alpha}} z_t$. Along a non-stochastic steady state growth path, Y_t/z_t^+ and $Y_{i,t}/z_t^+$ converge to constants. The two shocks, z_t and Ψ_t , are specified to be unit root processes in order to be consistent with the assumptions we use in our VAR analysis to identify the dynamic response of the economy to neutral and capital-embodied technology shocks. The two shocks have the following time series representations:

$$\Delta \log z_t = \mu_z + \varepsilon_t^n, \quad E(\varepsilon_t^n)^2 = (\sigma_n)^2 \quad (3.3)$$

$$\Delta \log \Psi_t = \mu_\psi + \rho_\psi \Delta \log \Psi_{t-1} + \varepsilon_t^\psi, \quad E(\varepsilon_t^\psi)^2 = (\sigma_\psi)^2. \quad (3.4)$$

Our assumption that the neutral technology shock follows a random walk matches closely the finding in Smets and Wouters (2007) who estimate $\log z_t$ to be highly autocorrelated. The direct empirical analysis of Prescott (1986) also supports the notion that $\log z_t$ is a random walk.

In (3.2), $H_{i,t}$ denotes homogeneous labor services hired by the i^{th} intermediate good producer. Intermediate good firms must borrow the wage bill in advance of production, so that one unit of labor costs is given by $W_t R_t$ where R_t denotes the gross nominal rate of interest. Intermediate good firms are subject to Calvo price-setting frictions. With probability ξ_p the intermediate good firm cannot reoptimize its price, in which case it is assumed to set its price according to the following rule:

$$P_{i,t} = \bar{\pi} P_{i,t-1}, \quad (3.5)$$

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

²³ The estimated model discussed in the next section features habit formation and a time-varying $\tilde{\eta}_t$. We verified numerically that in the estimated model, the consumption premium is pro-cyclical too.

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

where s_t denotes the marginal cost of production, denominated in units of the homogeneous good. s_t is a function only of the costs of capital and labor, and is described in the technical appendix, section B.11.1. In the firm's discounted profits, $\beta^j v_{t+j}$ is the multiplier on the households's nominal period $t + j$ budget constraint. The equilibrium conditions associated with this optimization problem are reported in section B.11.1 of the technical appendix.

We suppose that the homogeneous labor hired by intermediate good producers is itself 'produced' by competitive labor contractors. Labor contractors produce homogeneous labor by aggregating different types of specialized labor, $j \in (0, 1)$, as follows:

$$H_t = \left[\int_0^1 (h_{t,j})^{\frac{1}{\lambda_w}} dj \right]^{\lambda_w}, \quad 1 < \lambda_w. \quad (3.6)$$

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

$$W_{j,t} = W_t \left(\frac{H_t}{h_{t,j}} \right)^{\frac{\lambda_w - 1}{\lambda_w}}. \quad (3.7)$$

Equation (3.7) is the demand curve for the j^{th} type of labor.

3.2. Worker and household preferences

We integrate the model of unemployment in the previous section into the Erceg et al. (2000) (EHL) model of sticky wages that is commonly used in empirically relevant DSGE models. Each type, $j \in [0, 1]$, of labor is assumed to be supplied by a particular household. The j^{th} household resembles the single representative household in the previous section, with one exception. The exception is that the unit measure of workers in the j^{th} household is only able to supply the j^{th} type of labor service. Each worker in the j^{th} household has the utility cost of working, (2.1), and the technology for job finding, (2.5). The preference and job finding technology parameters are the same across households.

Let $c_{j,t}^{nw}$ and $c_{j,t}^w$ denote the consumption levels allocated by the j^{th} household to non-employed and employed workers within the household. Although households all enjoy the same level of consumption, C_t , for reasons described momentarily each household experiences a different level of employment, $h_{j,t}$. Because employment across households is different, each type j household chooses a different way to balance the trade-off between the need for consumption insurance and the need to provide work incentives. For the j^{th} type of household with high $h_{j,t}$, the premium of consumption for employed workers to non-employed workers must be high. Accordingly, the incentive constraint is given by (2.12) which we repeat here for convenience:

$$\ln \left(\frac{c_{j,t}^w - bC_{t-1}}{c_{j,t}^{nw} - bC_{t-1}} \right) = \zeta (1 + \sigma_L) m_{j,t}^{\sigma_L} - \frac{2}{a^2} \tilde{\eta}_t$$

where $m_{j,t}$ solves the analog of (2.17):

$$h_{j,t} = -\tilde{\eta}_t m_{j,t} + a^2 \zeta \sigma_L \left(m_{j,t}^{\sigma_L + 1} - i_{j,t}^{\sigma_L + 1} \right) \quad (3.8)$$

and

$$i_{j,t}^{\sigma_L} = m_{j,t}^{\sigma_L} - \frac{1 + \tilde{\eta}_t}{\zeta (1 + \sigma_L) a^2}. \quad (3.9)$$

Consider the j^{th} household that enjoys a level of household consumption and employment, C_t and $h_{j,t}$, respectively. Note that given (2.22) from the previous section, the j^{th} household's discounted utility is given by:

$$E_0 \sum_{t=0}^{\infty} \beta^t \left[\ln (C_t - bC_{t-1}) - z(h_{j,t}; \tilde{\eta}_t) \right]. \quad (3.10)$$

Note that the utility function is additively separable, like the utility functions assumed for the workers. Additive separability is convenient because perfect consumption insurance at the level of households implies that consumption is not indexed by labor type, j .

3.3. The household problem

The j^{th} household is the monopoly supplier of the j^{th} type of labor service. The household understands that when it arranges work incentives for its workers so that employment is $h_{j,t}$, then the nominal wage $W_{j,t}$ takes on the value implied by the demand for its type of labor, (3.7). The household therefore faces the standard monopoly problem of selecting $W_{j,t}$ to optimize the welfare, (3.10), of its member workers. It does so subject to the requirement that it satisfy the demand for labor, (3.7), in each period. We follow EHL in supposing that the household experiences Calvo-style frictions in its choice of $W_{j,t}$. In particular, with probability $1 - \xi_w$ the j^{th} household has the opportunity to reoptimize its wage rate. With the complementary probability, the household must set its wage rate according to the following rule:

$$W_{j,t} = \tilde{\pi}_{w,t} W_{j,t-1} \quad (3.11)$$

$$\tilde{\pi}_{w,t} = (\pi_{t-1})^{\kappa_w} (\tilde{\pi})^{(1-\kappa_w)} \mu_{z^+}, \quad (3.12)$$

where $\kappa_w \in [0, 1]$. Note that in a non-stochastic steady state, non-optimizing households raise their real wage at the rate of growth of the economy. Because optimizing households also do this in steady state, it follows that in the steady state, the wage of each type of household is the same.

In principle, the presence of wage setting frictions implies that households have idiosyncratic levels of wealth and, hence, consumption. However, we follow EHL in supposing that each household has access to perfect consumption insurance. At the level of the household, there is no private information about consumption or employment. The private information and associated incentive problems all exist among the workers inside a household. Because of the additive separability of the household utility function, perfect consumption insurance at the level of households implies equal consumption across households. We have used this property of the equilibrium to simplify our notation and not include a subscript, j , on the j^{th} household's consumption. Of course, we hasten to add that although consumption is equated across households, it is not constant across workers.

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

$$P_t \left(C_t + \frac{1}{\Psi_t} I_t \right) + B_{t+1} \leq W_{t,j} h_{t,j} + X_t^k \bar{K}_t + R_{t-1} B_t + a_{jt}. \quad (3.13)$$

Here, B_{t+1} denotes the quantity of risk-free bonds purchased by the worker, R_t denotes the gross nominal interest rate on bonds purchased in period $t - 1$ which pay off in period t , and a_{jt} denotes the payments and receipts associated with the insurance on the timing of wage reoptimization. Also, P_t denotes the aggregate price level and I_t denotes the quantity of investment goods purchased for augmenting the beginning-of-period $t + 1$ stock of physical capital, \bar{K}_{t+1} . The price of investment goods is P_t/Ψ_t , where Ψ_t is the unit root process with positive drift specified in (3.4). This is our way of capturing the trend decline in the relative price of investment goods.²⁴

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

$$\bar{K}_{t+1} = (1 - \delta) \bar{K}_t + \left(1 - S \left(\frac{I_t}{I_{t-1}} \right) \right) I_t. \quad (3.14)$$

Here, S is a convex function which we discuss below.

For each unit of \bar{K}_{t+1} acquired in period t , the household receives X_{t+1}^k in net cash payments in period $t + 1$,

$$X_{t+1}^k = u_{t+1}^k P_{t+1} r_{t+1}^k - \frac{P_{t+1}}{\Psi_{t+1}} a(u_{t+1}^k), \quad (3.15)$$

where u_t^k denotes the rate of utilization of capital. The first term in (3.15) is the gross nominal period $t + 1$ rental income from a unit of \bar{K}_{t+1} . The household supply of capital services in period $t + 1$ is:

$$K_{t+1} = u_{t+1}^k \bar{K}_{t+1}.$$

It is the services of capital that intermediate good producers rent and use in their production functions, (3.2). The second term to the right of the equality in (3.15) represents the cost of capital utilization, $a(u_{t+1}^k) P_{t+1}/\Psi_{t+1}$ which we discuss below.

The household's problem is to select sequences, $\{C_t, I_t, u_t^k, W_{j,t}, B_{t+1}, \bar{K}_{t+1}\}$, to maximize (3.10) subject to (3.7), (3.11), (3.12), (3.13), (3.14), (3.15) and the mechanism determining when wages can be reoptimized. The equilibrium conditions associated with this maximization problem are standard, and so appear in section B.11.2 of the technical appendix.

²⁴ We suppose that there is an underlying technology for converting final goods, Y_t , one-to-one into C_t and one to Ψ_t into investment goods. These technologies are operated by competitive firms which equate price to marginal cost. The marginal cost of C_t with this technology is P_t and the marginal cost of I_t is P_t/Ψ_t . We avoid a full description of this environment so as to not clutter the presentation, and simply impose these properties of equilibrium on the household budget constraint.

3.4. Aggregate resource constraint, monetary policy and functional forms

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

$$Y_t = G_t + C_t + \tilde{I}_t. \quad (3.16)$$

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

$$\tilde{I}_t = \frac{1}{\Psi_t} \left(I_t + a \left(u_t^k \right) \bar{K}_t \right). \quad (3.17)$$

As discussed above, the investment goods, I_t , are used by the households to add to the physical stock of capital, \bar{K}_t , according to (3.14). The remaining investment goods are used to cover maintenance costs, $a \left(u_t^k \right) \bar{K}_t$, arising from capital utilization, u_t^k . Finally, Ψ_t in (3.17) denotes the unit root investment specific technology shock with positive drift discussed after (3.2).

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

$$\ln \left(\frac{R_t}{R} \right) = \rho_R \ln \left(\frac{R_{t-1}}{R} \right) + (1 - \rho_R) \left[r_\pi \ln \left(\frac{\pi_t}{\pi} \right) + r_y \ln \left(\frac{gdp_t}{gdp} \right) \right] + \frac{\sigma_R \varepsilon_{R,t}}{400}, \quad (3.18)$$

where $\varepsilon_{R,t}$ is an iid monetary policy shock with unit variance and σ_R scales the effective variance of monetary policy shocks. As in CEE and ACEL, we assume that period t realizations of ε_R are not included in the period t information set of workers and firms, so that the only variable that is contemporaneously affected by the monetary policy shock is the nominal interest rate.

Let gdp_t denote scaled real GDP defined as:

$$gdp_t = \frac{G_t + C_t + I_t / \Psi_t}{z_t^+}, \quad (3.19)$$

and gdp denote the nonstochastic steady state value of gdp_t .

We adopt the following specification for government spending, G_t , and the fixed cost of production, ϕ_t , in response to technology shocks. To guarantee balanced growth in the nonstochastic steady state, we require that each element in $[\phi_t, G_t]$ grows at the same rate as z_t^+ in steady state. Following Christiano et al. (2011a), we assume:

$$[\phi_t, G_t]' = [\phi, G]' \Omega_t. \quad (3.20)$$

Here, Ω_t is defined as follows:

$$\Omega_t = (z_{t-1}^+)^{\theta} (\Omega_{t-1})^{1-\theta}, \quad (3.21)$$

where $0 < \theta \leq 1$ is a parameter to be estimated. With this specification, Ω_t / z_t^+ converges to a constant in nonstochastic steady state. When θ is close to zero, Ω_t is virtually unresponsive in the short-run to an innovation in either of the two technology shocks, a feature that we find attractive on *a priori* grounds. Given the specification of the exogenous processes in the model, Y_t / z_t^+ , C_t / z_t^+ and $I_t / (\Psi_t z_t^+)$ converge to constants in nonstochastic steady state.

In terms of fiscal policy, we assume that lump-sum transfers balance the government budget.

Finally, we assume the following functional forms. For the capacity utilization cost function a we adopt the functional form:

$$a(u_t^k) = \sigma_a \sigma_b (u_t^k)^2 / 2 + \sigma_b (1 - \sigma_a) u_t^k + \sigma_b (\sigma_a / 2 - 1),$$

where σ_a and σ_b are the parameters of this function. For a given value of σ_a we select σ_b so that the steady state value of u_t^k is unity. The object, σ_a , is a parameter to be estimated.

We assume that the investment adjustment cost function takes the form:

$$S(I_t / I_{t-1}) = \frac{1}{2} \left\{ \exp \left[\sqrt{S''} (I_t / I_{t-1} - \mu_{z^+} \mu_\Psi) \right] + \exp \left[-\sqrt{S''} (I_t / I_{t-1} - \mu_{z^+} \mu_\Psi) \right] - 2 \right\}.$$

Here, μ_{z^+} and μ_Ψ denote the unconditional growth rates of z_t^+ and Ψ_t . The value of I_t / I_{t-1} in nonstochastic steady state is $(\mu_{z^+} \mu_\Psi)$. In addition, S'' denotes the second derivative of $S(\cdot)$, evaluated at steady state. The object, S'' , is a parameter to be estimated. It is straightforward to verify that $S(\mu_{z^+} \mu_\Psi) = S'(\mu_{z^+} \mu_\Psi) = 0$.

We assume the following functional form for the impact of aggregate economic conditions on the worker's probability to find a job:

$$\mathcal{M}(\bar{m}_t / \bar{m}_{t-1}) = 100\omega (\bar{m}_t / \bar{m}_{t-1} - 1).$$

We will estimate the parameter ω using a standard normal prior. That is, we are agnostic about the sign of ω . Recall that $\tilde{\eta}_t = \eta + \mathcal{M}(\tilde{m}_t/\tilde{m}_{t-1})$ and $p(e_{1,t}; \tilde{\eta}_t) = \tilde{\eta}_t + ae_{1,t}$. That is if, for example, $\omega < 0$, then this implies that an inflow of workers into the labor force reduces the probability of a worker to find a job. Importantly, it is the rate of change of the labor force that reduces the job finding probability of a worker. Intuitively, one might think about this as a bottleneck-type access to the labor market. When the labor force grows rapidly, many workers get ‘stuck’ in the process of finding work. In spirit of the various adjustment cost specifications in estimated medium-sized New Keynesian models, according to our specification, it is not the level of the labor force but its rate of change that affects the probability of a worker to find a job. In effect, if $\omega < 0$, our specification implies a more gradual adjustment of the labor force in response to shocks in line with model specifications with labor force adjustment costs as in e.g. Erceg and Levin (2014) and Christiano et al. (2015). We will elaborate more on the effects of our functional form assumption for the impact of aggregate economic conditions on the worker’s probability to find a job for the quantitative properties of our medium-sized New Keynesian model below when we discuss our model estimation results.

3.5. Aggregate labor force and unemployment in our model

We now derive our model’s implications for unemployment and the labor market. At the level of the j^{th} household, unemployment and the labor force are defined in the same way as in the previous section, except that the endogenous variables now have a j subscript (the parameters and shocks are the same across households). Thus, the j^{th} household’s labor force, $m_{j,t}$, and total employment, $h_{j,t}$, are related by (3.8) and $\dot{l}_{j,t}$ is given by (2.15). Log-linearizing these expressions gives:

$$\begin{aligned} h\hat{h}_{j,t} &= -\tilde{\eta}m(\hat{\eta}_t + \hat{m}_{j,t}) + (\sigma_L + 1)a^2\zeta\sigma_L(m^{\sigma_L+1}\hat{m}_{j,t} - \dot{l}^{\sigma_L+1}\hat{l}_{j,t}) \\ \sigma_L\dot{l}^{\sigma_L}\hat{l}_{j,t} &= \sigma_L m^{\sigma_L}\hat{m}_{j,t} - \frac{\tilde{\eta}}{\zeta(1+\sigma_L)a^2}\hat{\eta}_t \end{aligned} \quad (3.22)$$

Variables without subscript denote steady state values in the j^{th} household. Because we have made assumptions which guarantee that each household is identical in steady state, we drop the j subscripts from all steady state labor market variables (see the discussion after (3.11)).

Aggregate household hours and the labor force are defined as follows:

$$\hat{h}_t = \int_0^1 \hat{h}_{j,t} dj, \quad \hat{m}_t \equiv \int_0^1 \hat{m}_{j,t} dj, \quad \hat{l}_t \equiv \int_0^1 \hat{l}_{j,t} dj.$$

Using the fact that, to first order, type j wage deviations from the aggregate wage cancel, we obtain:

$$\hat{h}_t = \hat{H}_t. \quad (3.23)$$

See section B.7 in the technical appendix for a derivation. That is, to a first order approximation, the percent deviation of aggregate household hours from steady state coincides with the percent deviation of aggregate homogeneous hours from steady state. Integrating (3.22) over all j and substituting for \hat{l}_t yields:

$$h\hat{h}_t = \underbrace{\left(-\tilde{\eta}m + (\sigma_L + 1)a^2\zeta\sigma_L(m - \dot{l})m^{\sigma_L}\right)}_{>0}\hat{m}_t - \underbrace{\tilde{\eta}[m - \dot{l}]}_{>0}\hat{\eta}_t,$$

where $\hat{\eta}_t = \frac{\tilde{\eta}_t - \tilde{\eta}}{\tilde{\eta}}$. Aggregate unemployment is defined as $u_t \equiv \frac{m_t - h_t}{m_t}$ so that $du_t = \frac{h}{m}(\hat{m}_t - \hat{h}_t)$ where du_t denotes the deviation of unemployment from its steady state value, not the percent deviation.

3.6. The standard model

We derive the utility function used in the standard model as a special case of the household utility function in our involuntary unemployment model. In part, we do this to ensure consistency across models. In part, we do this as a way of emphasizing that we interpret the labor input in the utility function in the standard model as corresponding to the number of people working, not, say, the hours worked of a representative person. With our interpretation, the curvature of the labor disutility function corresponds to the (consumption compensated) elasticity with which people enter or leave the labor force in response to a change in the wage rate. In particular, this curvature does not correspond to the elasticity with which the typical person adjusts the quantity of hours worked in response to a wage change. Empirically, the latter elasticity is estimated to be small and it is fixed at zero in the model.

Another advantage of deriving the standard model from ours is that it puts us in position to exploit an insight by Galí (2011). In particular, Galí (2011) shows that the standard model already has a theory of unemployment implicit in it. The

monopoly power assumed by EHL has the consequence that wages are on average higher than what they would be under competition. The number of workers for which the wage is greater than the cost of work exceeds the number of people employed. Galí suggests defining this excess of workers as ‘unemployed’. The implied unemployment rate and labor force represent a natural benchmark to compare with our model.

Notably, deriving an unemployment rate and labor force in the standard model does not introduce any new parameters. Moreover, there is no change in the equilibrium conditions that determine non-labor market variables. Galí’s insight in effect simply adds a block recursive system of two equations to the standard DSGE model which determine the size of the labor force and unemployment. Although the unemployment rate derived in this way does not satisfy all the criteria for unemployment that we described in the introduction, it nevertheless provides a natural benchmark for comparison with our model. An extensive comparison of the economics of our approach to unemployment versus the approach implicit in the standard model appears in the appendix to this paper.

We suppose that in the standard model, the household has full information about its member workers and that workers which join the labor force automatically receive a job without having to expend any effort. As in the previous subsections, we suppose that corresponding to each type j of labor, there is a unit measure of workers which gather together into a household. At the beginning of each period, each worker draws a random variable, l , from a uniform distribution with support, $[0, 1]$. The random variable, l , determines a worker’s aversion to work according to (2.1). The fact that no effort is needed to find a job implies $m_{t,j} = h_{t,j}$. Workers with $l \leq h_{t,j}$ work and workers with $h_{t,j} \leq l \leq 1$ take leisure. The type j household allocation problem is to maximize the utility of its member workers with respect to consumption for non-working workers, $c_{t,j}^{nw}$, and consumption of working workers, $c_{t,j}^w$, subject to (2.20), and the given values of $h_{t,j}$ and C_t . In Lagrangian form, the problem is:

$$u(C_t - bC_{t-1}, h_{j,t}) = \max_{c_{t,j}^w, c_{t,j}^{nw}} \int_0^{h_{t,j}} \left[\ln(c_{t,j}^w - bC_{t-1}) - \varsigma(1 + \sigma_L)l^{\sigma_L} \right] dl \\ + \int_{h_{t,j}}^1 \ln(c_{t,j}^{nw} - bC_{t-1}) dl + \lambda_{j,t} \left[C_t - h_{t,j}c_{t,j}^w - (1 - h_{t,j})c_{t,j}^{nw} \right].$$

Here, $\lambda_{j,t} > 0$ denotes the multiplier on the resource constraint. The first order conditions imply $c_{t,j}^w = c_{t,j}^{nw} = C_t$. Imposing this result and evaluating the integral, we find:

$$u(C_t - bC_{t-1}, h_{j,t}) = \ln(C_t - bC_{t-1}) - \varsigma h_{t,j}^{1+\sigma_L}. \quad (3.24)$$

The problem of the household is identical to what it is in section 3.3, with the sole exception that the utility function, (3.10), is replaced by (3.24).

A type j worker that draws work aversion index l is defined to be unemployed if the following two conditions are satisfied:

$$(a) l > h_{j,t}, \quad (b) v_t W_{j,t} > \varsigma(1 + \sigma_L)l^{\sigma_L}. \quad (3.25)$$

Here, v_t denotes the multiplier on the budget constraint, (3.13), in the Lagrangian representation of the household optimization problem. Expression (a) in (3.25) simply says that to be unemployed, the worker must not be employed. Expression (b) in (3.25) determines whether a non-employed worker is unemployed or not in the labor force. The object on the left of the inequality in (b) is the value assigned by the household to the wage, $W_{j,t}$. The object on the right of (b) is the fixed cost of going to work for the l^{th} worker. Galí (2010) suggests defining workers with l satisfying (3.25) as unemployed. This approach to unemployment does not satisfy properties (i) and (iii) in the introduction. The approach does not meet the official definition of unemployment because no one is exercising effort to find a job. In addition, the existence of perfect consumption insurance implies that unemployed workers enjoy higher utility than employed workers.

We use (3.25) to define the labor force, m_t , in the standard model. With m_t and aggregate employment, h_t , we obtain the unemployment rate as follows $u_t = \frac{m_t - h_t}{m_t}$ or, after linearization about steady state $du_t = \frac{h}{m}(\hat{m}_t - \hat{h}_t)$. Here, $h < m$ because of the presence of monopoly power. The object, \hat{h}_t may be obtained from (3.23) and the solution to the standard model. We now discuss the computation of the aggregate labor force, m_t . We have $m_t \equiv \int_0^1 m_{j,t} dj$ where $m_{j,t}$ is the labor force associated with the j^{th} type of labor and is defined by enforcing (b) in (3.25) at equality. After linearization $\hat{m}_t \equiv \int_0^1 \hat{m}_{j,t} dj$. We compute $\hat{m}_{j,t}$ by linearizing the equation that defines $\hat{m}_{j,t}$. After scaling (3.25), we obtain

$$\psi_t \bar{w}_t \hat{w}_{j,t} = \varsigma(1 + \sigma_L) m_{j,t}^{\sigma_L}, \quad (3.26)$$

where $\psi_t \equiv v_t P_t z_t^+$, $\bar{w}_t \equiv \frac{W_t}{z_t^+ P_t}$, $\hat{w}_{j,t} \equiv \frac{W_{j,t}}{W_t}$. Log-linearizing (3.26) around the steady state and integrating the result over all $j \in (0, 1)$:

$$\hat{\psi}_t + \hat{w}_t + \int_0^1 \hat{w}_{j,t} dj = \sigma_L \hat{m}_t.$$

From the result in section B.7 in the technical appendix, the integral in the above expression is zero, so that:

$$\hat{m}_t = \frac{\hat{\psi}_t + \hat{w}_t}{\sigma_L}.$$

4. Estimation strategy

We estimate the parameters of the model in the previous section using the impulse response matching approach applied by Rotemberg and Woodford (1997), CEE, ACEL and other papers. We apply the Bayesian version of that method proposed in CTW. Specifically, our estimation machinery makes use of priors and posteriors, as well as the marginal likelihood as a measure of model fit in our impulse response function matching estimation. The advantage of the Bayesian impulse response matching estimation approach that we use is transparency and focus.²⁵ The transparency reflects that the estimation strategy has a simple graphical representation, involving objects - impulse response functions - about which economists have often very strong intuition. The advantage of focus comes from the possibility of studying the empirical properties of a model without having to specify a full set of shocks.

Impulse response matching estimation is often very useful when crafting new models with new transmission channels since the estimation procedure is very transparent and allows the researcher to focus on the particular new model features and new transmission mechanisms when taking the model to the data. Given that our paper is about constructing a new labor market model, we find the impulse response matching procedure particularly attractive.

To promote comparability of results across the two papers and to simplify the discussion here, we use the impulse response functions and associated probability intervals estimated using the 14 variable, 2 lag vector autoregression (VAR) estimated in CTW.²⁶ Here, we consider the response of 11 variables to three shocks: the monetary policy shock, $\varepsilon_{R,t}$ in equation (3.18), the neutral technology shock, ε_t in equation (3.3), and the investment specific shock, ε_t^Ψ in equation (3.4).²⁷ Nine of the eleven variables whose responses we consider are the standard macroeconomic variables displayed in Figs. 1–3. The other two variables are the unemployment rate and the labor force which are shown in Fig. 4. The VAR is estimated using quarterly, seasonally adjusted data covering the period 1952Q1 to 2008Q4.

The assumptions that allow us to identify the effects of our three shocks are the ones implemented in ACEL and Fisher (2006). To identify the monetary policy shock we suppose all variables aside from the nominal rate of interest are unaffected contemporaneously by the policy shock.²⁸ We make two assumptions to identify the dynamic response to the technology shocks: (i) the only shocks that affect labor productivity in the long run are the two technology shocks and (ii) the only shock that affects the price of investment relative to consumption is the innovation to the investment specific shock.²⁹ We emphasize that our medium-sized model structure is such that it is in line with the identifying assumptions for the monetary policy shock as well as for the two types technology shocks in the VAR. That is, the timing of monetary policy shocks as well as the long-run effects of technology shocks coincide in the model and the VAR.

Let $\hat{\psi}$ denote the vector of impulse responses used in the analysis here. Since we consider 15 lags in the impulses, there are in principle 3 (i.e., the number of shocks) times 11 (number of variables) times 15 (number of lags) = 495 elements in $\hat{\psi}$. However, we do not include in $\hat{\psi}$ the 10 contemporaneous responses to the monetary policy shock that are required to be zero by our monetary policy identifying assumption. Taking the latter into account, the vector $\hat{\psi}$ has 485 elements. To conduct a Bayesian analysis, we require a likelihood function for our data, $\hat{\psi}$. For this, we use an approximation based on asymptotic sampling theory. In particular, when the number of observations, T , is large, we have

$$\sqrt{T} \left(\hat{\psi} - \psi(\theta_0) \right) \overset{a}{\sim} N(0, W(\theta_0, \zeta_0)). \quad (4.1)$$

Here, θ_0 and ζ_0 are the parameters of the model that generated the data, evaluated at their true values. The parameter vector, θ_0 , is the set of parameters that is explicit in our model, while ζ_0 contains the parameters of stochastic processes

²⁵ Another advantage of the impulse response matching estimation compared to full information estimation is that the former does not require the underlying data to be normally distributed while the latter does.

²⁶ See CTW for a sensitivity analysis with respect to the lag length of the VAR. Further, see the technical appendix in CTW for details about the data.

²⁷ The VAR in CTW also includes data on vacancies, job findings and job separations, but these variables do not appear in the models in this paper and so we do not include their impulse responses in the analysis.

²⁸ This timing assumption has a long history in macroeconomics. The basic idea is that when a monetary policy shock occurs, many spending and other economic decisions are already under way, so that the economy's response to the shock occurs with a delay. Lengthy discussions of this assumption appear in Bernanke and Blinder (1992) and an even earlier discussion is Friedman (1959)'s testimony before Congress: "Monetary and fiscal policy is rather like a water tap that you turn on now and that then only starts to run 6, 9, 12, 16 months from now". A very recent discussion appears in Jarociński and Karadi (2020). They display some direct evidence on the idea that monetary policy shocks have no immediate impact on aggregate variables. They report (see their Fig. 2), that the contemporaneous response of output to an exogenous shock to monetary policy shock is numerically zero. Their identification strategy, based on high-frequency methods, does not require that response to be zero.

²⁹ Details of our strategy for computing impulse response functions imposing the shock identification are discussed in ACEL.

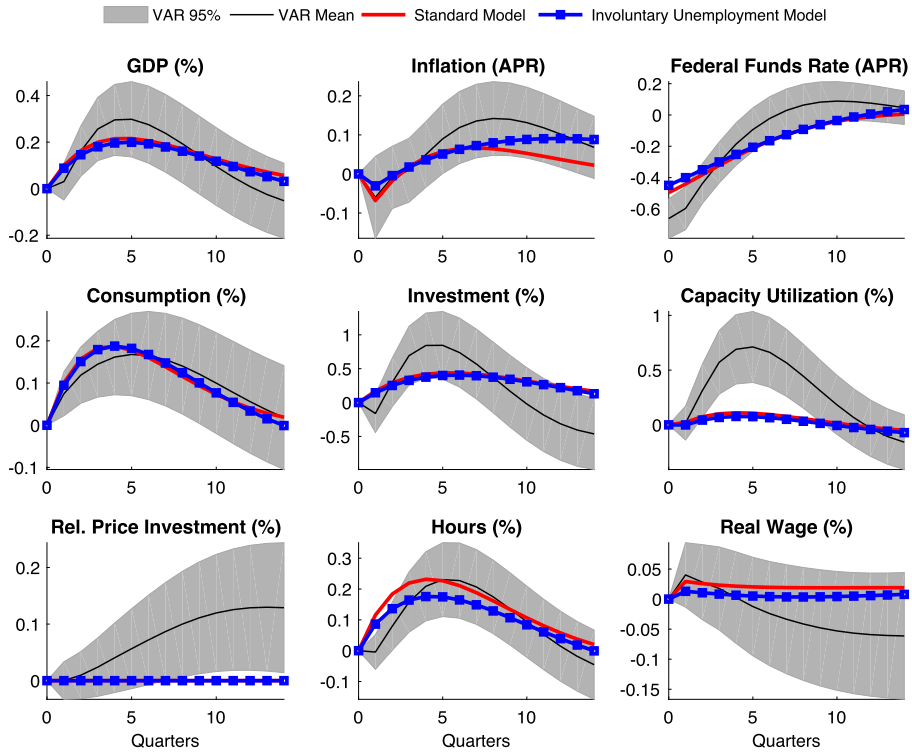


Fig. 1. Dynamic responses to a monetary policy shock.

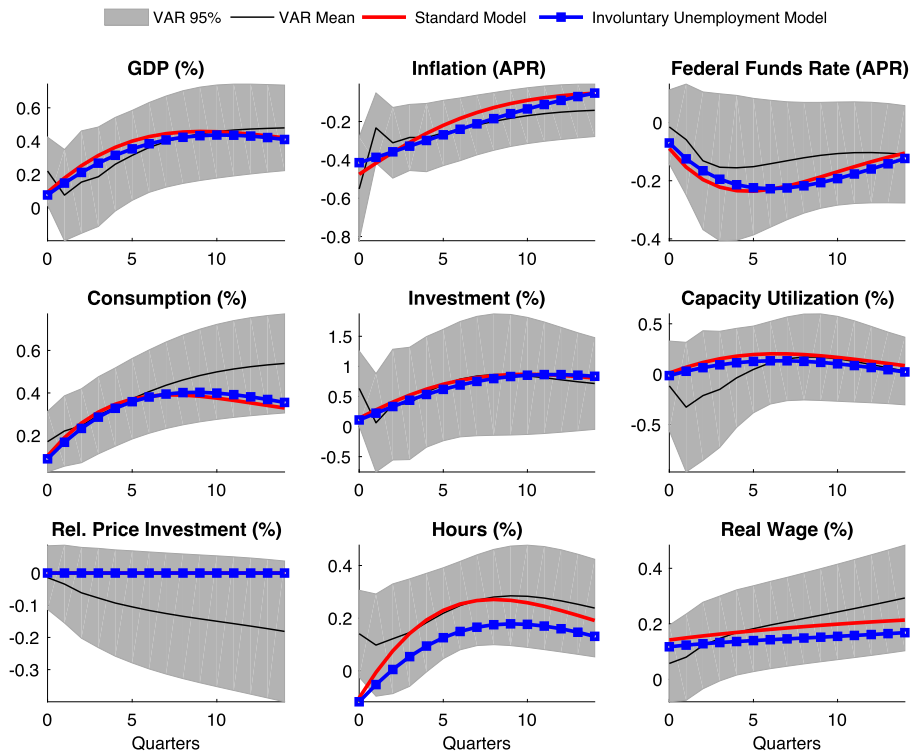


Fig. 2. Dynamic responses to a neutral technology shock.

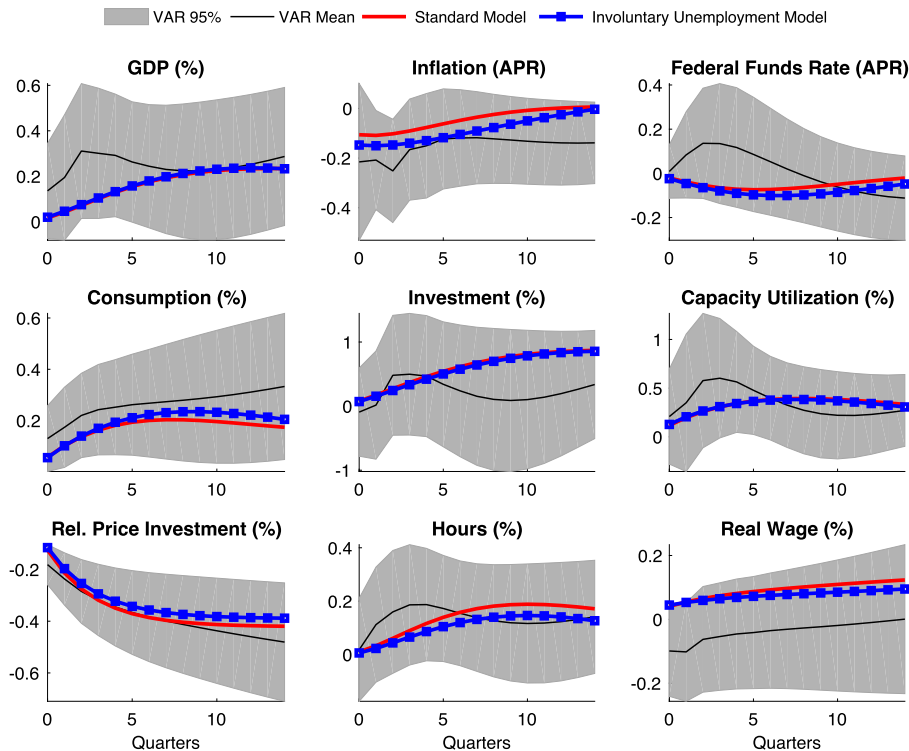


Fig. 3. Dynamic responses to an investment-specific technology shock.

not included in the analysis. In (4.1), $W(\theta_0, \zeta_0)$ is the asymptotic sampling variance of $\hat{\psi}$, which - as indicated by the notation - is a function of all model parameters. We find it convenient to express (4.1) in the following form:

$$\hat{\psi} \stackrel{a}{\sim} N(\psi(\theta_0), V), \tag{4.2}$$

where $V \equiv \frac{W(\theta_0, \zeta_0)}{T}$. For simplicity our notation does not make the dependence of V on θ_0, ζ_0 and T explicit. We treat V as though it were known. In practice, we work with a consistent estimator of V in our analysis (for details, see CTW). That estimator is a diagonal matrix with only the variances along the diagonal. An advantage of this diagonality property is that our estimator has a simple graphical representation.³⁰

We treat the following object as the likelihood of the data, $\hat{\psi}$, conditional on the model parameters, θ :

$$f(\hat{\psi}|\theta, V) = \left(\frac{1}{2\pi}\right)^{\frac{N}{2}} |V|^{-\frac{1}{2}} \exp\left[-\frac{1}{2}(\hat{\psi} - \psi(\theta))' V^{-1}(\hat{\psi} - \psi(\theta))\right]. \tag{4.3}$$

The Bayesian posterior of θ conditional on $\hat{\psi}$ and V is:

$$f(\theta|\hat{\psi}, V) = \frac{f(\hat{\psi}|\theta, V) p(\theta)}{f(\hat{\psi}|V)}, \tag{4.4}$$

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

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

As usual, the mode of the posterior distribution of θ can be computed by simply maximizing the value of the numerator in (4.4), since the denominator is not a function of θ . The marginal density of $\hat{\psi}$ is required when we want an overall measure

³⁰ The diagonal matrix V makes the estimator particularly transparent. Specifically, the estimator then corresponds to selecting the set of estimated parameters such that the model impulse responses lie inside a confidence tunnel around the estimated VAR impulses. If we instead had allowed for non-diagonal terms in V , then the estimator aims not just to put the model impulses inside a confidence tunnel about the VAR point estimates, but it is also concerned about the pattern of discrepancies across different impulse responses. In addition, by using a diagonal variance-covariance matrix V we ensure maximum comparability with the existing literature, including Christiano et al. (2005) and ACEL who also imposed this structure on V .

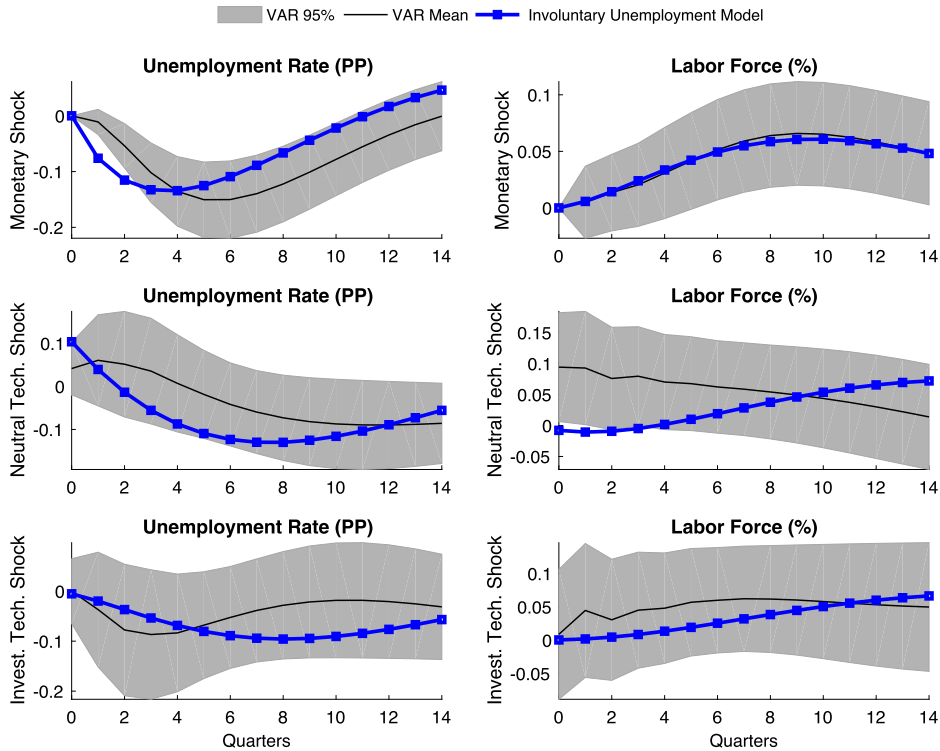


Fig. 4. Dynamic responses of unemployment and labor force to three shocks.

Table 1
Non-estimated parameters in medium-sized model.

Parameter	Value	Description
δ	0.025	Depreciation rate
β	0.99678	Discount factor
π	1.00625	Gross inflation rate
η_g	0.2	Government consumption to GDP ratio
κ_w	1	Wage indexation
λ_w	1.01	Wage markup
ξ_w	0.75	Wage stickiness
$400\ln\mu_{z^+}$	1.7	Annual output per capita growth rate
$400\ln\mu_{z^+}\mu_\psi$	2.9	Annual investment per capita growth rate

of the fit of our model and when we want to report the shape of the posterior marginal distribution of individual elements in θ . We do this using the MCMC algorithm.

5. Estimation results for the medium-sized model

The first subsection below discusses model parameter values. We then show that our model of involuntary unemployment does well at accounting for the dynamics of unemployment and the labor force. Fortunately, the model is able to do this without compromising its ability to account for the dynamics of standard macroeconomic variables.

5.1. Parameters

Parameters whose values are set a priori are listed in Table 1. We found that when we estimated the parameters, κ_w and λ_w , the estimator drove them to their boundaries. This is why we simply set λ_w to a value near unity and we set $\kappa_w = 1$. The steady state value of inflation (a parameter in the monetary policy rule and the price and wage updating equations), the steady state government consumption to output ratio, and the growth rate of investment-specific technology were chosen to coincide with their corresponding sample means in our data set.³¹ The growth rate of neutral technology was chosen

³¹ In our model, the relative price of investment goods represents a direct observation of the technology shock for producing investment goods.

Table 2
Priors and posteriors of parameters in the estimated medium-sized model.

Parameter		Prior		Posterior Mode [2.5% 97.5%]	
		Distribution [bounds]	Mode [2.5% 97.5%]	Standard Model	Involuntary Unemp. Model
<i>Price Setting Parameters</i>					
Price Stickiness	ξ_p	Beta [0, 1]	0.67 [0.45 0.83]	0.616 [0.55 0.71]	0.727 [0.67 0.78]
Price Markup	λ_f	Gamma [1.001, ∞]	1.19 [1.01 1.40]	1.230 [1.10 1.36]	1.399 [1.29 1.54]
<i>Monetary Authority Parameters</i>					
Taylor Rule: Int. Smoothing	ρ_R	Beta [0, 1]	0.76 [0.37 0.93]	0.873 [0.82 0.90]	0.890 [0.85 0.91]
Taylor Rule: Inflation Coef.	r_π	Gamma [1.001, ∞]	1.68 [1.41 2.00]	1.395 [1.19 1.65]	1.414 [1.19 1.69]
Taylor Rule: GDP Coef.	r_y	Gamma [0, ∞]	0.07 [0.02 0.21]	0.077 [0.03 0.14]	0.113 [0.05 0.18]
<i>Preference Parameters</i>					
Consumption Habit	b	Beta [0, 1]	0.75 [0.64 0.83]	0.761 [0.72 0.79]	0.776 [0.74 0.80]
Inverse Labor Supply Elast.	σ_z	Gamma [0, ∞]	0.26 [0.13 0.52]	0.165 [0.08 0.23]	0.334 [0.17 0.43]
Replacement Ratio	c^{nw}/c^w	Beta [0, 1]	0.75 [0.69 0.79]	–	0.797 [0.76 0.82]
Labor Force Impact on $p(e, \tilde{\eta})$	ω	Normal [$-\infty$, ∞]	0.0 [–1.96 1.96]	–	–0.533 [–0.74 –0.38]
<i>Technology Parameters</i>					
Capital Share	α	Beta [0, 1]	0.32 [0.28 0.37]	0.31 [0.25 0.33]	0.270 [0.24 0.31]
Technology Diffusion	θ	Beta [0, 1]	0.50 [0.12 0.86]	0.052 [0.01 0.80]	0.015 [0.01 0.05]
Capacity Adj. Costs Curv.	σ_a	Gamma [0, ∞]	0.31 [0.09 1.22]	0.462 [0.21 0.56]	0.256 [0.10 0.59]
Investment Adj. Costs Curv.	S''	Gamma [0, ∞]	7.50 [4.56 12.29]	11.56 [8.46 14.92]	15.72 [11.46 18.78]
<i>Shocks</i>					
Autocorr. Invest. Tech.	ρ_ψ	Beta [0, 1]	0.78 [0.53 0.91]	0.703 [0.54 0.77]	0.704 [0.59 0.82]
Std.Dev. Neutral Tech. Shock	σ_n	Inv. Gamma [0, ∞]	0.06 [0.04 0.44]	0.211 [0.18 0.25]	0.194 [0.17 0.23]
Std.Dev. Invest. Tech. Shock	σ_ψ	Inv. Gamma [0, ∞]	0.06 [0.04 0.44]	0.125 [0.09 0.17]	0.115 [0.08 0.15]
Std.Dev. Monetary Shock	σ_R	Inv. Gamma [0, ∞]	0.22 [0.14 1.49]	0.496 [0.41 0.60]	0.449 [0.37 0.53]

so that, conditional on the growth rate of investment-specific technology, the steady state growth rate of output in the model coincides with the corresponding sample average in the data. We set $\xi_w = 0.75$, so that the model implies wages are reoptimized once a year on average. We did not estimate this parameter because we found that it is difficult to separately identify the value of ξ_w and the curvature of household labor disutility.

The parameters for which we report priors and posteriors are listed in Table 2. The posterior mode and parameter distributions are based on a standard MCMC algorithm with a total of 2.5 million draws based on 10 chains. We use the first 20 percent of draws for burn-in. The acceptance rates are about 0.25 in each chain. We report results for two estimation exercises. In the first exercise we estimate the standard DSGE model discussed in section 3.6. In this exercise we only use the impulse responses of standard macroeconomic variables in the likelihood criterion, (4.3). In particular, we do not include the impulse responses of the unemployment rate or the labor force when we estimate the standard DSGE model.³² Results based on this exercise appear under the heading, ‘standard model’. In the second exercise we estimate our model with involuntary unemployment and we report those results under the heading, ‘involuntary unemployment model’.

We make several observations about the parameters listed in Table 2. First, the results in the last two columns are relatively similar. This reflects that the two models (i) are observationally equivalent relative to the impulse responses of standard macroeconomic variables and (ii) no substantial adjustments to the parameters are required for the involuntary unemployment model to fit the unemployment and labor force data.

³² Subsection 5.3 in the appendix discusses the implications when unemployment and the labor force are included in the estimation in the standard model.

Table 3
Medium-sized model steady state at posterior mode for parameters.

Variable	Standard Model	Involuntary Unemp. Model	Description
$p_k k/y$	8.765	7.665	Capital to GDP ratio (quarterly)
c/y	0.519	0.554	Consumption to GDP ratio
i/y	0.281	0.246	Investment to GDP ratio
$H = h$	0.628	0.628	Steady state labor input
c^{nw}/c^w	1.000	0.797	Replacement ratio
R	1.014	1.014	Gross nominal interest rate (quarterly)
R^{real}	1.0075	1.0075	Gross real interest rate (quarterly)
u	0.059	0.055	Unemployment rate
m	–	0.665	Labor force (involuntary unemployment model)
l^*	0.668	–	Labor force (standard model)
\bar{l}	–	0.504	Share of workers with $p(e; \bar{\eta}) = 1$
ζ	1.936	0.609	Slope, labor disutility
σ_L	0.165	4.287	Curvature, labor disutility
η	–	–0.467	Intercept, $p(e; \bar{\eta})$
a	–	1.170	Slope, $p(e; \bar{\eta})$

In the estimation of the involuntary unemployment model, we calculate the four parameters η , a , σ_L , ζ endogenously so as to set the following objects exogenously: h , m , $\frac{z_{hh}h}{z_h} = \sigma_z^{\text{target}}$, $r = c^w/c^{nw} = r^{\text{target}}$. See section C.2 in the technical appendix for more details. We set $h = 0.628$ and $m = 0.665$ which yields a steady state unemployment rate of 0.055. We estimate the values of σ_z^{target} and r^{target} . The resulting values for η , a , σ_L , ζ at the estimated posterior mode of the estimated parameters are provided in Table 3.

In the estimation of the standard model, we apply an analogous treatment to worker parameter values. In particular, throughout the estimation we fix the steady state level of hours worked, $h = 0.628$ and calculate the value of the parameter ζ endogenously. The resulting value for ζ at the estimated posterior mode of the estimated parameters is provided in Table 3. Similar to the involuntary unemployment model, we also estimate the value of $\frac{z_{hh}h}{z_h} = \sigma_z^{\text{target}}$ in the standard model. Note, however, that given the preference specification of the standard model, $\sigma_z = \sigma_L$ (evaluate (2.24) using (3.24)).

Turning to the parameter values themselves, note first that the degree of price stickiness, ξ_p , is modest. The implied time between price reoptimizations is a little less than 3 quarters in the standard model and a little less than 4 quarters in the involuntary unemployment model. The amount of information in the likelihood, (4.3), about the value of ξ_p is reasonably large in both models. The posterior standard deviation is roughly an order of magnitude smaller than the prior standard deviation and the posterior probability interval is half the length of the prior probability interval. Generally, the amount of information in the likelihood about all the parameters is large in this sense. An exception to this pattern is the coefficient on inflation in the Taylor rule, r_π . There appears to be relatively little information about this parameter in the likelihood. Note that σ_z is estimated to be quite small, implying a consumption-compensated labor supply elasticity for the household of around 6 in the standard model and 3 in the involuntary unemployment model. Such high elasticities would be regarded as empirically implausible if it was interpreted as the elasticity of supply of hours by an individual worker. But this elasticity should instead be interpreted as the elasticity at the household level incorporating all individual workers who are part of the household. Put differently, in our model of involuntary unemployment, the elasticity of labor supply is quite large at the household level as an outcome of our limited labor market insurance arrangement. By contrast the elasticity of labor supply is quite small at the individual worker level, consistent with micro evidence. For further discussion of this distinction between measures of individual and household-level labor supply elasticities, see section 2.3, in CTW.

The consumption replacement ratio, $r = c^{nw}/c^w$, is a novel feature of our model, that does not appear in standard monetary DSGE models. The replacement ratio is estimated to be roughly 80 percent. This is close to the value used for calibration by Landais et al. (2018). It is higher than the estimates of Hamermesh (1982) but somewhat lower than the empirical estimate of 90 percent reported by Chetty and Looney (2007) and Gruber (1997) mentioned in the introduction. Also, our consumption replacement ratio appears to be higher than the number reported for developed countries in OECD (2006). However, the replacement ratios reported by OECD pertain to income, rather than consumption.³³ So, they are likely to underestimate the consumption concept relevant for us.

Table 2 reveals that the estimated value of the parameter governing the impact of the labor force on the probability of a worker to find a job is $\omega = -0.533$. Recall that we use a standard normal prior, i.e. we are agnostic about the sign of ω . The data is informative about the sign of ω being negative. Note that $\omega < 0$ implies that an inflow of workers into the labor force reduces the probability of a worker to find a job. The negative value of ω helps the model to account for the slow response of the labor force to shocks, see Figs. 1–3 which are discussed in the next subsection. To examine the sensitivity of our results we re-estimated the involuntary unemployment model with ω set to zero. Technical appendix Table A.2

³³ The income replacement ratio for the US is reported to be 54 percent in Table 3.2, which can be found at <http://www.oecd.org/dataoecd/28/9/36965805.pdf>.

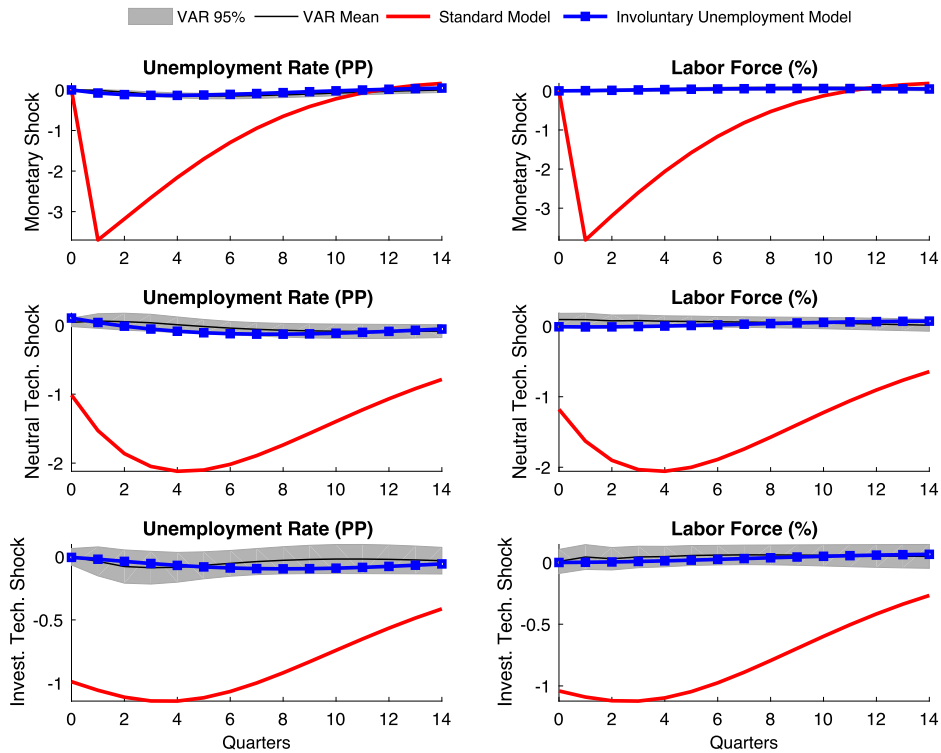


Fig. 5. Dynamic responses of unemployment and labor force to three shocks.

contains the corresponding estimated parameters and technical appendix Figs. 1 through 4 show the corresponding impulse responses. The impulse responses show that the model fit deteriorates when ω is set to zero instead of being set to its estimated value of $\omega = -0.533$. The principal impact of ω is on the response of the labor market variables and inflation to a monetary shock. The qualitative implications of the model remain unchanged, including the positive hump-shaped response of the labor force to an expansionary monetary policy shock. The effects of $\omega < 0$ are to reduce the magnitude of the peak response that occurs when $\omega = 0$ and to delay the timing of that peak. When $\omega = 0$, the response of inflation to a monetary shock is qualitatively similar, but quantitatively much weaker than in the data. Setting ω to its estimated value has the effect of making the VAR and model impulse responses of inflation virtually the same. The marginal data density of the involuntary unemployment model with $\omega = 0$ is about 60 log points lower than the baseline involuntary unemployment model with $\omega = -0.533$.

Finally, Table 3 reports steady state properties of the estimated standard model as well as the estimated baseline involuntary unemployment model, both evaluated at the posterior mode of the parameters.

5.2. Impulse response functions of macroeconomic variables

Figs. 1–3 display the results of the indicated macroeconomic variables to our three shocks. In each case, the solid black line is the point estimate of the dynamic response to a one standard deviation shock generated by our estimated VAR. The gray area is an estimate of the corresponding 95% probability interval.³⁴ Our estimation strategy selects a model parameterization that places the model-implied impulse response functions as close as possible to the center of the gray area, while not suffering too much of a penalty from the priors. The estimation criterion is less concerned about reproducing VAR-based impulse response functions where the gray areas are the widest.

The thick solid line and the line with solid squares in the figures display the impulse responses of the standard model and the involuntary unemployment model, respectively, at the posterior mean of the parameters. Note in Figs. 1–3 that in many cases only one of these two lines is visible. Moreover, in cases where a distinction between the two lines can be discerned, they are nevertheless very close. This reflects that the two models account roughly equally well for the impulse responses to the three shocks. This is a key result. Expanding the standard model to include unemployment and the labor

³⁴ We compute the probability interval as follows. We simulate 2,500 sets of impulse response functions by generating an equal number of artificial data sets, each of length T , using the VAR estimated from the data. Here, T denotes the number of observations in our actual data set. We compute the standard deviations of the artificial impulse response functions. The gray areas in Figs. 1–5 are the estimated impulse response functions plus and minus 1.96 times the corresponding standard deviation.

force does not produce a deterioration in the model's ability to account for the estimated dynamic responses of standard macroeconomic variables to monetary policy and technology shocks.

Consider Fig. 1, which displays the response of standard macroeconomic variables to an expansionary monetary policy shock (i.e., a negative shock to $\varepsilon_{R,t}$). Note how the model captures the slow response of inflation. Indeed, the model even captures the 'price puzzle' phenomenon, according to which inflation moves in the 'wrong' direction initially. This apparently perverse initial response of inflation is interpreted by the model as reflecting the reduction in labor costs associated with the cut in the nominal rate of interest.³⁵ It is interesting that the slow response of inflation is accounted for with a fairly modest degree of wage and price-setting frictions. The model captures the response of output and consumption to a monetary policy shock reasonably well. However, there is a substantial miss on capacity utilization. Also, the model apparently does not have the flexibility to capture the relatively sharp rise and fall in the investment response, although the model responses lie inside the gray area. The relatively large estimate of the curvature in the investment adjustment cost function, S'' , reflects that to allow a greater response of investment to a monetary policy shock would cause the model's prediction of investment to lie outside the gray area in the initial and later quarters. These findings for monetary policy shocks are broadly similar to those reported in CEE, ACEL and CTW.

Fig. 2 displays the response of standard macroeconomic variables to a neutral technology shock. Note that the models do reasonably well at reproducing the empirically estimated responses. The dynamic response of inflation is particularly notable. The estimation results in ACEL suggest that the sharp and precisely estimated drop in inflation in response to a neutral technology shock is difficult to reproduce in a model like the standard monetary DSGE model. In describing this problem for their model, ACEL express a concern that the failure reflects a deeper problem with sticky price models. Perhaps the emphasis on price and wage setting frictions, largely motivated by the inertial response of inflation to a monetary shock, is shown to be misguided by the evidence that inflation responds rapidly to technology shocks.³⁶ Our results suggest a far more mundane possibility. There are two differences between our model and the one in ACEL which allow it to reproduce the response of inflation to a technology shock more or less exactly without hampering its ability to account for the slow response of inflation to a monetary policy shock. First, as discussed above, in our model there is no indexation of prices to lagged inflation. ACEL follows CEE in supposing that when firms cannot optimize their price, they index it fully to lagged aggregate inflation. The position of our model on price indexation is a key reason why we can account for the rapid fall in inflation after a neutral technology shock while ACEL cannot. We suspect that our way of treating indexation is a step in the right direction from the point of view of microeconomic data. Micro observations suggest that individual prices do not change for extended periods of time. A second distinction between our model and the one in ACEL is that we specify the neutral technology shock to be a random walk (see (3.3)), while in ACEL the growth rate of the estimated technology shock is highly autocorrelated. In ACEL, a technology shock triggers a strong wealth effect which stimulates a surge in demand that places upward pressure on marginal cost and thus inflation.³⁷

Fig. 3 displays dynamic responses of macroeconomic variables to an investment-specific shock. The evidence indicates that the two models, parameterized at their posterior means, do well in accounting for these responses.

5.3. Impulse response functions of unemployment and the labor force

Fig. 4 displays the response of unemployment and the labor force to our three shocks. The key thing to note is that the model has no difficulty accounting for the pattern of responses. The probability bands are large, but the point estimates suggest that unemployment falls about 0.2 percentage points and the labor force rises a small amount after an expansionary monetary policy shock. The model roughly reproduces this pattern. In the case of each response, the model generates opposing movements in the labor force and the unemployment rate. This appears to be consistent with the evidence.

As discussed in section 3.6 above, Galí (2011) points out that the standard model has implicit in it a theory of unemployment and the labor force. Fig. 5 adds the implications of the standard model for these variables to the impulses displayed in Fig. 4 when data for unemployment and the labor force are not part of the dataset used in the standard model. Note that the impulses implied by the standard model are so large that they distort the scale in Fig. 5. Consider, for example, the first panel of graphs in the figure, which pertain to the monetary policy shock. The standard model predicts a massive fall in the labor force after an expansionary monetary policy shock. The reason is that the rise in aggregate consumption (see Fig. 1) reduces the value of work by reducing v_t in (3.25). The resulting sharp drop in labor supply strongly contradicts our VAR-based evidence which suggests a small rise. Given the standard model's prediction for the labor force, it is not surprising that the model massively over-predicts the fall in the unemployment rate after a monetary expansion.

³⁵ For a defense, based on firm-level data, of the existence of this 'working capital' channel of monetary policy, see Barth and Ramey (2001).

³⁶ The concern is reinforced by the fact that an alternative approach, one based on information imperfections and minimal price/wage setting frictions, seems like a natural one for explaining the puzzle of the slow response of inflation to monetary policy shocks and the quick response to technology shocks (see Maćkowiak and Wiederholt, 2009; Mendes, 2009, and Paciello, 2011). Dupor et al. (2009) suggest more modest changes in the model structure to accommodate the inflation puzzle.

³⁷ An additional, important, factor accounting for the damped response of inflation to a monetary policy shock (indeed, the perverse initial 'price puzzle' phenomenon) is the assumption that firms must borrow in advance to pay for their variable production costs. But, this model feature is present in both our model and ACEL as well as CEE.

An alternative approach to deduce the implications of the standard model for unemployment and the labor force is to impose the corresponding VAR impulse responses in the estimation. When doing so, the unemployment rate and labor force responses are virtually flat after the monetary shock, which is counterfactual with respect to the VAR evidence. Basically, the estimation procedure selects parameters such that the unemployment rate and labor force do not fall as much as displayed in Fig. 5.³⁸ However, selecting parameters in the standard model to basically shut down the responses of unemployment and the labor force comes at a heavy cost: the fit of all other macroeconomic data deteriorates sharply in this case. See Figs. A.1 to A.4 in the appendix and Table A.1 in the technical appendix for the estimated parameter values. Therefore, our model of involuntary unemployment outperforms the standard model when both models face the same dataset including unemployment and the labor force. Quantitatively, the log data density at the posterior mode for our involuntary unemployment model is about 200 log points higher than the one for the standard model. See appendix section A.1 for an in-depth discussion of the underlying estimation results for the standard model in this case.

The failure of the standard model raises a puzzle. Why does our involuntary unemployment model do so well at accounting for the unemployment rate and the labor force? The puzzle is interesting because the two models have household utility functions with comparable form (recall the indirect utility function in our model, eq. (3.10)). One might imagine that our model would have the same problem with wealth effects. In fact, it does not have the same problem because there is a connection in our model between the labor force and employment that does not exist in the standard model. In our model, the increased consumption premium from holding a job that occurs in response to an expansionary monetary policy shock simultaneously encourages workers to search for work more intensely, and to substitute into the labor force.

The standard model's prediction for the response of the unemployment rate and the labor force to neutral and investment-specific technology shocks is also strongly counterfactual. The problem is always the same, and reflects the operation of wealth effects on labor supply.

The problems in Fig. 5 with the standard model motivate Galí (2011) and Galí et al. (2011) to modify the worker utility function in the standard model in ways that reduce wealth effects on labor. Our view is that modifying the utility function is not the right way to go as microdata evidence indicate substantial and immediate wealth effects on labor supply. For example, Cesarini et al. (2017) find that winning a lottery prize reduces individual labor earnings already within a year, with very persistent effects thereafter. Earnings fall by approximately 1.1 percent of the prize amount per year. Somewhat stronger effects are documented by Imbens et al. (2001) who find that the discounted value of earnings falls by 11% in response to an exogenous increase in unearned income, i.e. a lottery prize. There is also macroeconomic evidence indicating that wealth effects on labor supply are non-negligible. Specifically, Mertens and Ravn (2011) use VAR evidence to show that the economy contracts in response to anticipated tax cuts. They then account for the VAR evidence by using a DSGE model with standard wealth preferences. In effect, our involuntary unemployment model represents an alternative strategy for dealing with these wealth effects. Our model has the added advantage of being consistent with all three characteristics (i)-(iii) of unemployment described in the introduction.

6. Further evidence in favor of our model

Our model of unemployment has several interesting microeconomic implications that deserve closer attention. The model implies that the consumption premium of employed workers over the non-employed, c_t^w/c_t^{nw} , is procyclical or, equivalently, the replacement ratio, c_t^{nw}/c_t^w , is countercyclical. Although Chetty and Looney (2007) and Gruber (1997) report that there is a premium on average, we cannot infer anything about the cyclicity of the premium from the evidence they present. Studies of the cross section variance of log worker consumption are a potential source of evidence on the cyclical behavior of the premium. To see this, let V_t denote the variance of log worker consumption in the period t cross section in our model³⁹:

$$V_t = (1 - h_t) h_t \left(\log \left(\frac{c_t^w}{c_t^{nw}} \right) \right)^2.$$

According to this expression, the model posits two countervailing forces on the cross-sectional dispersion of consumption, V_t , in a recession. First, for a given distribution of the population across employed and non-employed workers (i.e., holding h_t fixed), a decrease in the consumption premium leads to a decrease in consumption dispersion in a recession. Second, holding the consumption premium fixed, consumption dispersion increases as people move from employment to non-employment with the fall in h_t .⁴⁰ These observations suggest that (i) if V_t is observed to drop in recessions, this is evidence in favor of the model's prediction that the consumption premium is procyclical and (ii) if V_t is observed to stay

³⁸ Note that the standard model is not able to generate a rise in the labor force after an expansionary monetary policy shock. Thus, the "best" response possible in that model is a zero labor force response to the monetary shock.

³⁹ Note that the formula for V_t corresponds to the model developed in section 2 of this paper. For the model with capital developed in section 3, the relevant formula is more complicated as it requires a non-trivial aggregation across households that supply different types of labor services. To see how we derived the formula in the text, note that the cross-sectional mean of log household consumption is $E_t = h_t \log(c_t^w) + (1 - h_t) \log(c_t^{nw})$ so that $V_t = h_t (\log c_t^w - E_t)^2 + (1 - h_t) (\log c_t^{nw} - E_t)^2 = h_t (1 - h_t) (\log c_t^w - \log c_t^{nw})^2$.

⁴⁰ This statement assumes that the empirically relevant case applies, i.e. $h_t > 1/2$.

constant or rise in recessions then we cannot conclude anything about the cyclicity of the consumption premium. Evidence in Heathcote et al. (2010) suggests that the US was in case (i) in three of the previous five recessions.⁴¹ In particular, they show that the dispersion in log worker non-durable consumption decreased in the 1980, 2001 and 2007 recessions.⁴² We conclude tentatively that the observed cross-sectional dispersion of consumption across workers lends support to our model's implication that the consumption premium is procyclical. In addition, the fact that the duration of unemployment benefits routinely are extended in recessions (e.g. in the US) is an indication that the income premium is procyclical empirically.

Another interesting implication of the model is its prediction that high unemployment in recessions reflects the procyclicality of effort in job search. There is some evidence that supports this implication of the model. The Bureau of Labor Statistics (2009) constructs a measure of the number of discouraged workers. These are people who are available to work and have looked for work in the past 12 months, but are not currently looking because they believe no jobs are available. This statistic has only been gathered since 1994, and so it covers just two recessions. However, in both the recessions for which we have data, the number of discouraged workers increased substantially. For example, the number of discouraged workers jumped 70 percent from 2008Q1 to 2009Q1. In fact, the number of discouraged workers is only a tiny fraction of the labor force. However, to the extent that the sentiments of discouraged workers are shared by workers more generally, a jump in the number of discouraged workers could be a signal of a general decline in job search intensity in recessions. But, this is an issue that demands a more careful investigation.

Interestingly, Shimer (2004) reports evidence that search effort may be acyclical or even countercyclical. In his work, the number of different search methods that the unemployed use are counted at different stages of the business cycle. We interpret Shimer's finding as reflecting an *extensive* margin of search, i.e. how many alternative search methods are being used. By contrast, our model emphasizes the *intensive* margin of job search, i.e. how intensely one particular method of search is being used by the unemployed. Therefore, our model is not necessarily at odds with the evidence provided by Shimer.

7. Concluding remarks

We constructed a model in which workers must make an effort to find work. Because effort is privately observed, perfect insurance against labor market outcomes is not feasible. To ensure that people have an incentive to find work, workers that find jobs must be better off than people who do not work. With additively separable utility, this translates into the proposition that employed workers have higher consumption than the non-employed. We integrate our model of unemployment into a standard monetary DSGE model and find that the model's ability to account for standard macroeconomic variables is not diminished. At the same time, the new model appears to account well for the dynamics of variables like unemployment and the labor force.

The theory of unemployment developed here has interesting implications for the optimal variation of labor market insurance over the business cycle. In a boom more labor is demanded by firms. To satisfy the higher demand, workers are provided with more incentives to look for work by raising consumption for the employed, c_t^w , relative to consumption of the non-employed, c_t^{nw} . Conversely, in a recession, the consumption premium falls and thus the replacement ratio, c_t^{nw}/c_t^w , increases. Thus, our model implies a procyclical consumption premium – or equivalently – a countercyclical replacement ratio. Put differently, optimal labor market insurance is countercyclical in our model.

The empirical results highlight an important implication of our work. In particular, it is in general not sufficient to account for the response of only employment or total hours worked to be able to draw conclusions about the unemployment rate. In particular, when the standard model is estimated without data on unemployment and the labor force, the fit of total hours of the model is in fact very good. By contrast, the implications of the model for unemployment and the labor force are counterfactual. Conversely, when the standard model is estimated on unemployment and labor force data too, the fit of these two variables improves somewhat. However this improvement of fit comes at the cost of not fitting total hours well. In other words, the standard model provides an example that it is not straightforward to account for the dynamics of unemployment and labor force participation *jointly* with other standard macroeconomic variables. By contrast, our model does a good job in this regard.

We leave it to future research to quantify the various ways in which the new model may contribute to policy analysis. In part, we hope that the model is useful simply because labor market data are of interest in their own right. But, we expect the model to be useful even when labor market data are not the central variables of concern. An important input into policy analysis is the estimation of 'latent variables' such as the output gap and the efficient, or 'natural', rate of interest. Other important inputs into policy analysis are forecasts of inflation and output. By allowing one to systematically integrate labor market information into the usual macroeconomic dataset, our model can be expected to provide more precise forecasts, as well as better estimates of latent variables.⁴³ We also believe, in line with Veracierto (2008), that confronting models with labor market data such as unemployment and the labor force provides an important test for any business cycle model.

⁴¹ Of course, we cannot rule out that the drop in V_t in recessions has nothing to do with the mechanism in our model but rather reflects some other source of heterogeneity in the data.

⁴² A similar observation was made about the 2007 recession in Parker and Vissing-Jorgensen (2009).

⁴³ For an elaboration on this point, see Basistha and Startz (2004) and Christiano et al. (2011a).

Appendix A. Relationship of our work to Galí (2011)

In this section, we discuss the relationship of our work to Galí (2011) beyond those remarks made in the introduction and in section 3.6. Our paper emphasizes labor supply in its explanation of the dynamics of unemployment and the labor force. Galí adopts a similar perspective. To better explain our model, it is useful to compare its properties with those of Galí's model. Galí demonstrates that with a modest reinterpretation of variables, the standard DSGE model already contains a theory of unemployment. In particular, one can define the unemployed as the difference between the number of people actually working and the number of people that would be working if the marginal cost of work were equated to the wage rate. This difference is positive and fluctuating in the standard DSGE model because of the presence of wage-setting frictions and monopoly power. In effect, unemployment is a symptom of social inefficiency. People inflict unemployment upon themselves in the quest for monopoly profits. By contrast, in our model unemployment reflects frictions that are necessary for people to find jobs. The existence of unemployment does not require monopoly power. This point is dramatized by the fact that we introduce our model in the CGG framework, in which wages are set in competitive labor markets. At the same time, the logic of our model does create a positive relationship between monopoly power and unemployment. In our model, the employment contraction resulting from an increase in the monopoly power of unions produces a reduction in the incentives for workers to work. Workers' response to the reduced incentives is to allocate less effort to search, implying higher unemployment. So, our model shares the prediction of Galí's model that unemployment should be higher in economies with more union monopoly power. However, our model has additional implications that could differentiate it from Galí's. Ours implies that in economies with more union power both the labor force and the consumption premium for employed workers over non-employed workers are reduced. Galí's model predicts that with more union monopoly power, the labor force will be larger. The exact amount by which the labor force increases depends on the strength of wealth effects on leisure.

Other important differences between our model of unemployment and Galí's is that the latter fails to satisfy characteristics (i) and (iii) in the introduction. Galí's model assumes, as most of the related literature, that the available jobs can be found without effort. Because the model does not satisfy (i), unemployment does not meet the official U.S. definition of unemployment. In addition, the presence of perfect insurance in Galí's model implies that the employed have lower utility than the non-employed, violating (iii).

There are more differences between ours and Galí's theory unemployment. In standard DSGE models, labor supply plays little role in the dynamics of standard macro variables like consumption, output, investment, inflation and the interest rate. The reason is that the presence of wage setting frictions reduces the importance of labor supply. This is why the New Keynesian literature has been relatively unconcerned about all the old puzzles about income effects on labor and labor supply elasticities that were a central concern in the real business cycle literature. However, we show that these problems are back in full force if one adopts Galí's theory of unemployment. This is because labor supply corresponds to the labor force in that theory. To see how this brings back the old problems, we study the standard DSGE model's predictions for unemployment and the labor force in the wake of an expansionary monetary policy shock. Because that model predicts a rise in consumption, the model also predicts a decline in labor supply, as the income effect associated with increased consumption produces a fall in the value of work. The drop in labor supply is counterfactual, according to our VAR-based evidence. In addition, the large drop in the labor force leads to an counterfactually large drop in unemployment in the wake of an expansionary monetary policy shock.

Galí (2011) and Galí et al. (2011) show that changes to the worker utility function that offset wealth effects reduce the counterfactual implications of the standard model for the labor force. In effect, our paper proposes a different strategy. We preserve the additively separable utility function that is standard in monetary DSGE models, and our model nevertheless does not display the labor force problems in the standard DSGE model. This is because in our model the labor force and employment have a strong tendency to comove. In our model, the rise in employment in the wake of an expansionary monetary policy shock is accomplished by increasing people's incentives to work. The additional incentives not only encourage already active workers to intensify their job search, but also to shift into the labor force. More generally, the analysis highlights the fact that modeling unemployment requires thinking carefully about the determinants of the labor force.⁴⁴

A.1. Estimating the standard model on unemployment and labor force

In this section, we complement the discussion in section 5.3 when the standard model is also estimated on data for the unemployment rate and the labor force. In this case, the dataset used in the estimation of our involuntary unemployment model and the standard model is identical. Interestingly, there are four parameters that take on very different values at the posterior mode compared to the parameter estimates reported in Table 2 when the standard model is not estimated on unemployment and the labor force. These parameters are, the inverse labor supply elasticity, $\sigma_z = \sigma_L$, the steady state gross wage markup, λ^w , the curvature of capacity adjustment costs, σ_a and the Taylor rule coefficient, r_π . All other parameters listed in Tables 1 and 2 are affected only relatively little when the additional labor market data are taken on board in the estimation of the standard model. See technical appendix Table A.1 for the details.

⁴⁴ Our argument complements the argument in Krusell et al. (2011), who also stress the importance of understanding employment, unemployment and the labor force.

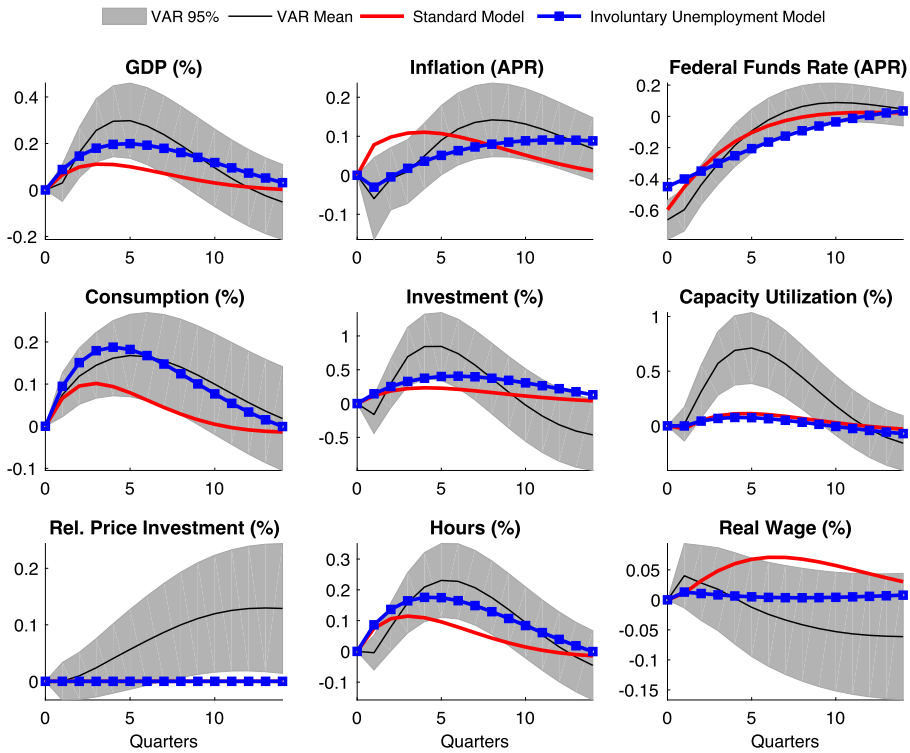


Fig. A.1. Dynamic responses to monetary policy shock when unemployment rate and labor force data are included in estimation of standard model.

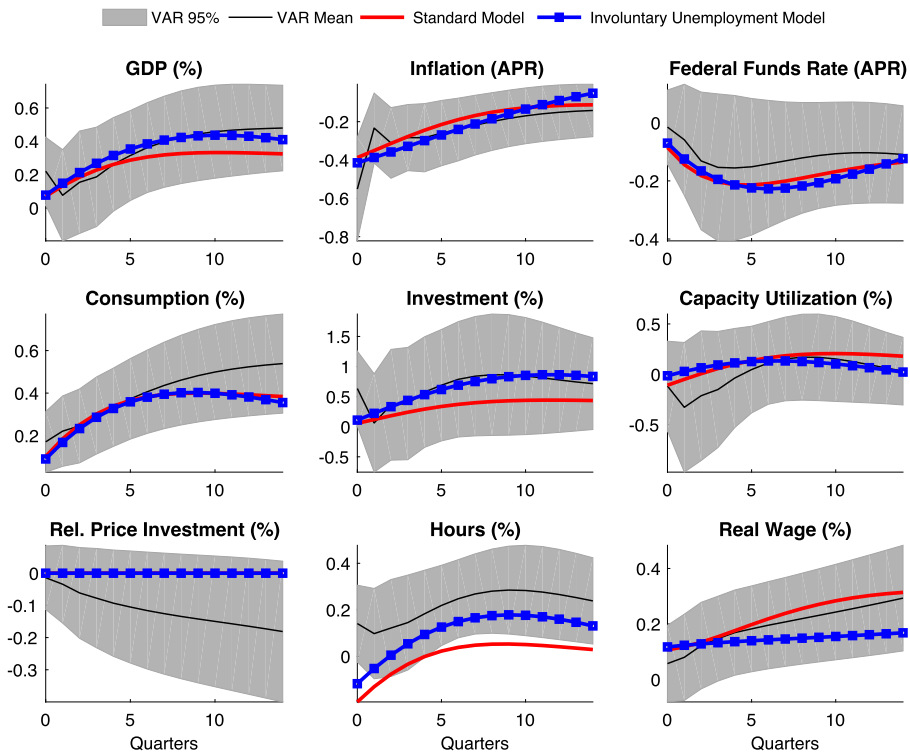


Fig. A.2. Dynamic responses to neutral technology shock when unemployment rate and labor force data are included in estimation of standard model.

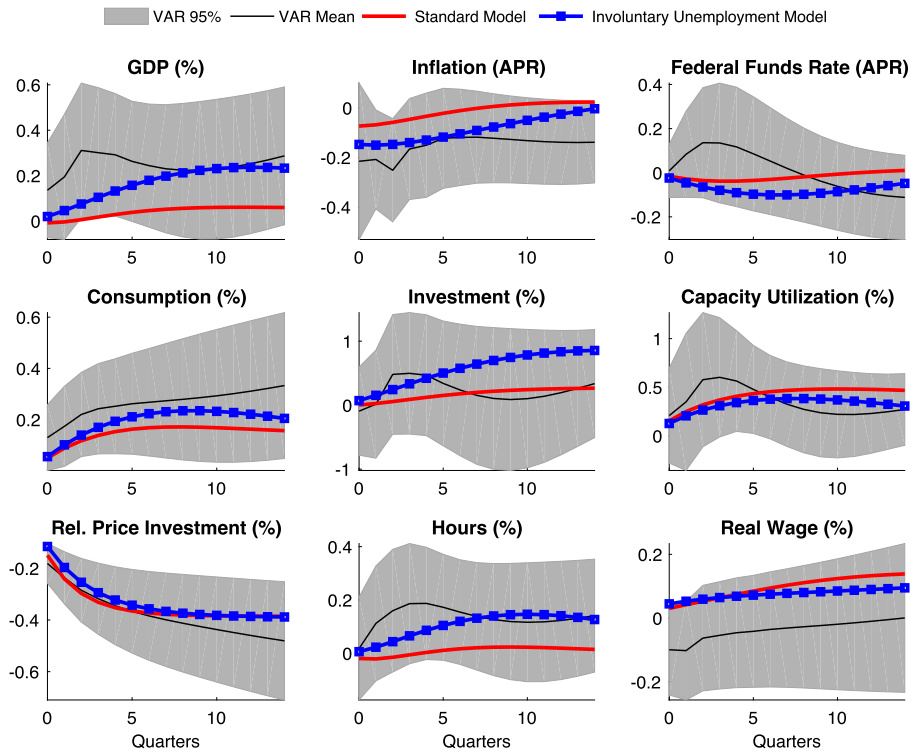


Fig. A.3. Dynamic responses to investment-specific technology shock when unemployment rate and labor force data are included in estimation of standard model.

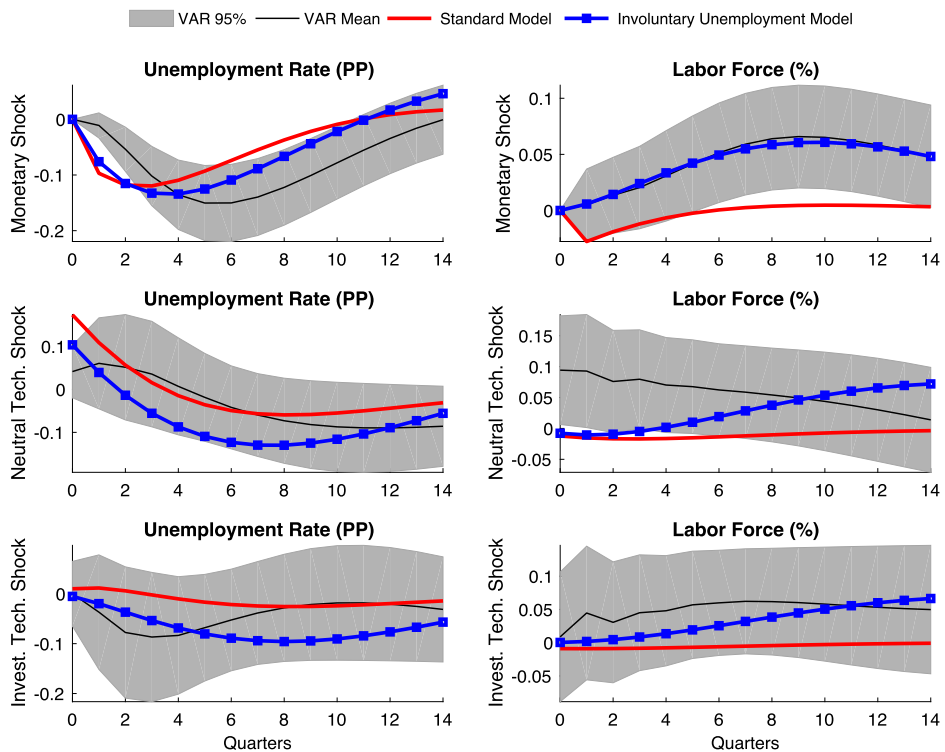


Fig. A.4. Dynamic responses of labor market variables to three shocks when unemployment rate and labor force data are included the estimation of standard model.

For convenience, let's repeat the equation from section 5.3 that determines the reaction of the labor force in the standard model, $\hat{l}_t^* = \frac{\hat{\psi}_{z+,t} + \hat{w}_t}{\sigma_L}$, where \hat{l}_t^* , $\hat{\psi}_{z+,t}$ and \hat{w}_t denote the labor force, marginal utility of consumption and the real wage, respectively. In the wake of an expansionary monetary policy shock, marginal utility of consumption falls much more than the real wage increases. Thus the labor force falls in the standard model while it rises according to the VAR. The only way the standard model can come close to the VAR responses is to drive $\sigma_z = \sigma_L$ to infinity and thereby shut down the response of the labor force. Setting $\sigma_z = \sigma_L$ to infinity, however, implies a zero labor supply elasticity and will therefore be harmful to the model in replicating the VAR responses for e.g. total hours. Thus, the estimation needs to balance the “miss” of the model for the labor force and e.g. total hours. It does so by selecting a posterior mode of $\sigma_z = \sigma_L = 18.12$ which is much higher than the value of about 0.165 reported in Table 2. Note that a value of $\sigma_z = \sigma_L$ as high as 18.18 relative to 0.165 generates a steady state unemployment rate close to zero when all other parameters are held fixed. In other words, the labor supply curve becomes essentially vertical. To enable maximum comparability with the model versions estimated in Table 2, we impose the same steady state unemployment rate of 5.5 percent in this experiment too. To do so, we need to set the gross wage markup $\lambda^w = 2.79$ at the posterior mode. The higher values of $\sigma_z = \sigma_L$ and λ^w imply that marginal costs rise much more steeply in response to e.g. an expansionary monetary policy shock. To at least partly offset this, the estimation wants to select a much higher steady state gross price markup $\lambda^f \gg 2$ which creates numerical issues in the estimation so that we have set λ^f to the estimated value of 1.23 in the estimated baseline standard model. Further, to at least partly offset the surge in marginal cost, the estimation selects a lower curvature of capacity adjustment costs of $\sigma_a = 0.02$, compared to Table 2. Appendix figures A1 to A4 show the responses of the model to the two technology shocks and to the monetary shock. Indeed, the standard model now delivers a worse fit for the standard macro variables. Still, the fit for unemployment and the labor force is not satisfactory. In terms of fit, the log data density at the posterior mode of our involuntary unemployment model is about 200 log points higher than for the standard model. Overall, it turns out that our model outperforms the standard model when both models face the same dataset including unemployment and the labor force.

Appendix B. Supplementary material

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.red.2020.05.003>.

References

- Altig, David, Christiano, Lawrence, Eichenbaum, Martin, Linde, Jesper, 2011. Firm-specific capital, nominal rigidities and the business cycle. *Review of Economic Dynamics* 14 (2), 225–247.
- Alexopoulos, Michelle, 2004. Unemployment and the business cycle. *Journal of Monetary Economics* 51 (2), 257–298.
- Andolfatto, David, 1996. Business cycles and labor market search. *The American Economic Review* 86, 112–132.
- Atkeson, Andrew, Lucas Jr., Robert E., 1995. Efficiency and equality in a simple model of efficient unemployment insurance. *Journal of Economic Theory* 66 (1), 64–88.
- Barth III, Marvin J., Ramey, Valerie A., 2001. The cost channel of monetary transmission. *NBER Macroeconomics Annual* 16, 199–240.
- Basistha, Arabinda, Startz, Richard, 2004. Measuring the NAIRU with Reduced Uncertainty: a Multiple Indicator-Common Component Approach. University of Washington. Manuscript.
- Benhabib, Jess, Rogerson, Richard, Wright, Randall, 1991. Homework in macroeconomics: worker production and aggregate fluctuations. *Journal of Political Economy* 99 (6), 1166–1187.
- Bernanke, Ben S., Blinder, Alan S., 1992. The federal funds rate and the channels of monetary transmission. *The American Economic Review* 82 (4), 901–921.
- Blanchard, Olivier, Galí, Jordi, 2010. Labor markets and monetary policy: a new Keynesian model with unemployment. *American Economic Journal: Macroeconomics* 2 (2), 1–30.
- Brown, Alessio J.G., Merkl, Christian, Snower, Dennis J., 2015. An incentive theory of matching. *Macroeconomic Dynamics* 19 (3), 643–668.
- Bureau of Labor Statistics, 2009. Ranks of Discouraged Workers and Others Marginally Attached to the Labor Force Rise During Recession, *Issues in Labor Statistics*. U.S. Department of Labor, U.S. Bureau of Labor Statistics. Summary 09-04/April, available at: <http://www.bls.gov/opub/ils/pdf/opbils74.pdf>.
- Campolmi, Alessia, Gnocchi, Stefano, 2016. Labor market participation, unemployment and monetary policy. *Journal of Monetary Economics* 79 (C), 17–29.
- Cesarini, David, Lindqvist, Erik, Notowidigdo, Matthew J., Östling, Robert, 2017. The effect of wealth on individual and worker labor supply: evidence from Swedish lotteries. *The American Economic Review* 107 (12), 3917–3946.
- Chetty, Raj, Looney, Adam, 2007. Income risk and the benefits of social insurance: evidence from Indonesia and the United States. In: *NBER Chapters*. In: *Fiscal Policy and Management in East Asia*, NBER-EASE, vol. 16. National Bureau of Economic Research, pp. 99–121.
- Christiano, Lawrence J., Eichenbaum, Martin, Evans, Charles L., 2005. Nominal rigidities and the dynamic effects of a shock to monetary policy. *Journal of Political Economy* 113 (1), 1–45.
- Christiano, Lawrence J., Eichenbaum, Martin S., Trabandt, Mathias, 2015. Understanding the great recession. *American Economic Journal: Macroeconomics* 7 (1), 110–167.
- Christiano, Lawrence J., Ilut, Cosmin, Motto, Roberto, Rostagno, Massimo, 2008. Monetary Policy and Stock Market Boom-bust Cycles. *European Central Bank Working Paper No. 955*.
- Christiano, Lawrence J., Trabandt, Mathias, Walentin, Karl, 2011a. DSGE models for monetary policy analysis. In: Friedman, B.M., Woodford, M. (Eds.), *Handbook of Monetary Economics*, vol. 3A.
- Christiano, Lawrence J., Trabandt, Mathias, Walentin, Karl, 2011b. Introducing financial frictions and unemployment into a small open economy model. *Journal of Economic Dynamics and Control* 35 (12), 1999–2041.
- Christoffel, Kai, Kuester, Keith, 2008. Resuscitating the wage channel in models with unemployment fluctuations. *Journal of Monetary Economics* 55, 865–887.
- Christoffel, Kai, Kuester, Keith, Linzert, Tobias, 2009a. The role of labor markets for euro area monetary policy. *European Economic Review* 53, 908–936.
- Christoffel, Kai, Costain, James, de Walque, Gregory, Kuester, Keith, Linzert, Tobias, Millard, Stephen, Pierrard, Olivier, 2009b. Inflation Dynamics with Labour Market Matching: Assessing Alternative Specifications. Working Paper No. 09-6. National Bank of Belgium.

- Clarida, Richard, Galí, Jordi, Gertler, Mark, 1999. The science of monetary policy: a new Keynesian perspective. *Journal of Economic Literature* XXXVII, 1661–1707.
- Cochrane, John H., 1991. A simple test of consumption insurance. *Journal of Political Economy* 99 (5), 957–976.
- den Haan, Wouter J., Ramey, Garey, Watson, J., 2000. Job destruction and propagation of shocks. *The American Economic Review* 90 (3), 482–498.
- Dupor, Bill, Han, Jing, Tsai, Yi-Chan, 2009. What do technology shocks tell us about the new Keynesian paradigm? *Journal of Monetary Economics* 56, 560–569.
- Erceg, Christopher J., Dale Henderson, W., Levin, Andrew T., 2000. Optimal monetary policy with staggered wage and price contracts. *Journal of Monetary Economics* 46, 281–313.
- Erceg, Christopher J., Levin, Andrew T., 2014. Labor force participation and monetary policy in the wake of the great recession. *Journal of Money, Credit, and Banking* 46 (S2), 3–49. Blackwell Publishing.
- Fisher, Jonas D.M., 2006. The dynamic effects of neutral and investment-specific technology shocks. *Journal of Political Economy* 114 (3), 413–452.
- Friedman, Milton, 1959. U.S. congress, joint economic committee, hearings, employment. In: *Growth, and Price Levels, Part 4. 86th Cong., 1st sess., 1959*, pp. 615–616.
- Galí, Jordi, 2011. The return of the wage Phillips curve. *Journal of the European Economic Association* 9 (3), 436–461.
- Galí, Jordi, Smets, Frank, Wouters, Raf, 2011. Unemployment in an estimated new Keynesian model. NBER Macroeconomics Annual 26.
- Gertler, Mark, Sala, Luca, Trigari, Antonella, 2008. An estimated monetary DSGE model with unemployment and staggered nominal wage bargaining. *Journal of Money, Credit, and Banking* 40 (8), 1713–1764.
- Gornemann, Nils, Kuester, Keith, Nakajima, Makoto, 2016. Doves for the Rich, Hawks for the Poor? Distributional Consequences of Monetary Policy. *International Finance Discussion Papers* 1167. Board of Governors of the Federal Reserve System.
- Grosheeny, Nicolas, 2009. Evaluating a Monetary Business Cycle Model with Unemployment for the Euro Area. National Bank of Belgium Working Paper No. 173.
- Gruber, Jonathan, 1997. The consumption smoothing benefits of unemployment insurance. *The American Economic Review* 87 (1), 192–205.
- Hamermesh, Daniel S., 1982. Social insurance and consumption: an empirical inquiry. *The American Economic Review* 72 (1), 101–113.
- Hansen, Gary D., 1985. Indivisible labor and the business cycle. *Journal of Monetary Economics* 16 (3), 309–327.
- Heathcote, Jonathan, Perri, Fabrizio, Violante, Giovanni L., 2010. Inequality in Times of Crisis: Lessons from the Past and a First Look at the Current Recession. University of Minnesota. Manuscript.
- Hopenhayn, Hugo, Nicolini, Juan Pablo, 1997. Optimal unemployment insurance. *Journal of Political Economy* 105 (2), 412–438.
- Imbens, Guido W., Rubin, Donald B., Sacerdote, Bruce I., 2001. Estimating the effect of unearned income on labor earnings, savings, and consumption: evidence from a survey of lottery players. *The American Economic Review* 91 (4), 778–794.
- Jarociński, Marek, Karadi, Peter, 2020. Deconstructing monetary policy surprises—the role of information shocks. *American Economic Journal: Macroeconomics* 12 (2), 1–43. <https://doi.org/10.1257/mac.20180090>.
- King, Robert G., Plosser, Charles I., Rebelo, Sergio T., 1988. Production, growth and business cycles: I. The basic neoclassical model. *Journal of Monetary Economics* 21 (2–3), 195–232.
- Kocherlakota, Narayana, 1996. Implications of efficient risk sharing without commitment. *The Review of Economic Studies* 63, 595–609.
- Krusell, Per, Mukoyama, Toshihiko, Rogerson, Richard, Sahin, Aysegul, 2008. Aggregate implications of indivisible labor, incomplete markets and labor market frictions. *Journal of Monetary Economics* 55, 961–979.
- Krusell, Per, Mukoyama, Toshihiko, Rogerson, Richard, Sahin, Aysegul, 2011. A three state model of worker flows in general equilibrium. *Journal of Economic Theory* 146 (3), 1107–1133.
- Krause, Michael U., Lopez-Salido, David, Lubik, Thomas A., 2008. Inflation dynamics with search frictions: a structural econometric analysis. *Journal of Monetary Economics* 55, 892–916.
- Landais, Camille, Michailat, Pascal, Saez, Emmanuel, 2012. Optimal Unemployment Insurance over the Business Cycle. NBER Working Papers 16526.
- Lechthaler, Wolfgang, Merkl, Christian, Snower, Dennis, 2010. Monetary persistence and the labor market: a new perspective. *Journal of Economic Dynamics and Control* 34 (5), 968–983.
- Maćkowiak, Bartosz, Wiederholt, Mirko, 2009. Optimal sticky prices under rational inattention. *The American Economic Review* 99 (3), 769–803.
- Mendes, Rhys, 2009. Information, Expectations, and the Business Cycle. Research Department, Bank of Canada. Manuscript.
- Mertens, Karel, Ravn, Morten O., 2011. Understanding the aggregate effects of anticipated and unanticipated tax policy shocks. *Review of Economic Dynamics* 14 (1), 27–54.
- Merz, Monika, 1995. Search in the labor market and the real business cycle. *Journal of Monetary Economics* 36, 269–300.
- Mulligan, Casey, 2001. Aggregate implications of indivisible labor. *Advances in Macroeconomics* 1 (1), 4.
- OECD, 2006. OECD Employment Outlook 2006: Boosting Jobs and Incomes. Organization for Economic Cooperation and Development, Renouf Pub Co Ltd. ISBN 92-64-02384-4.
- Paciello, Luigi, 2011. Does inflation adjust faster to aggregate technology shocks than to monetary policy shocks? *Journal of Money, Credit, and Banking* 43, 1663–1684.
- Parker, Jonathan, Vissing-Jorgensen, A., 2009. Who bears aggregate fluctuations and how? *The American Economic Review* 99 (2), 399–405.
- Prescott, Edward, 1986. Theory ahead of business cycle measurement. *Quarterly Review - Federal Reserve Bank of Minneapolis*, 9–22.
- Rogerson, Richard, 1988. Indivisible labor, lotteries and equilibrium. *Journal of Monetary Economics* 21 (1), 3–16.
- Rogerson, Richard, Wright, Randall, 1988. Involuntary unemployment in economies with efficient risk sharing. *Journal of Monetary Economics* 22, 501–515.
- Rotemberg, Julio, Woodford, Michael, 1997. An optimization-based econometric framework for the evaluation of monetary policy. In: Bernanke, B., Rotemberg, J. (Eds.), *NBER, Macroeconomics Annual*. MIT Press, Cambridge, MA.
- Sala, Luca, Söderström, Ulf, Trigari, Antonella, 2008. Monetary policy under uncertainty in an estimated model with labor market frictions. *Journal of Monetary Economics* 55, 983–1006.
- Shavell, Steven, Weiss, Laurence, 1979. The optimal payment of unemployment insurance over time. *Journal of Political Economy* 87 (6), 1347–1362.
- Shimer, Robert, 2004. Search Intensity. Manuscript.
- Smets, Frank, Wouters, Raf, 2003. An estimated dynamic stochastic general equilibrium model of the euro area. *Journal of the European Economic Association* 1 (5), 1123–1175.
- Smets, Frank, Wouters, Raf, 2007. Shocks and frictions in US business cycles. *The American Economic Review* 97 (3), 586–606.
- Svein, Tommy, Weinke, Lutz, 2008. New Keynesian perspectives on labor market dynamics. *Journal of Monetary Economics* 55, 921–930.
- Svein, Tommy, Weinke, Lutz, 2009. Inflation and labor market dynamics revisited. *Journal of Monetary Economics* 56, 1096–1100.
- Thomas, Carlos, 2011. Search frictions, real rigidities and inflation dynamics. *Journal of Money, Credit, and Banking* 43, 1131–1164.
- Trigari, Antonella, 2009. Equilibrium unemployment, job flows, and inflation dynamics. *Journal of Money, Credit, and Banking* 41 (1), 1–30.
- Veracierto, Marcelo, 2008. On the cyclical behavior of employment, unemployment and labor force participation. *Journal of Monetary Economics* 55 (6), 1143–1157.
- Walsh, Carl, 2005. Labor market search, sticky prices, and interest rate policies. *Review of Economic Dynamics* 8, 829–849.