

What Do Participation Fluctuations Tell Us About Labor Supply Elasticities? *

Christian Haefke[#] and Michael Reiter[‡]
Institute for Advanced Studies, Vienna[§]

October 12, 2011

Abstract

In this paper we use information on the cyclical variation of labor market participation to learn about the aggregate labor supply elasticity. For this purpose, we extend the standard labor market matching model to allow for endogenous participation. A model that is calibrated to replicate the variability of unemployment and participation, and the negative correlation of unemployment and GDP, implies an aggregate labor supply elasticity along the extensive margin of around 0.3 for men and 0.5 for women. This is in line with recent micro-econometric estimates.

Key Words: Matching Models, Labor Market Participation, Labor Supply Elasticity.

JEL Classification: E24, E32, J21, J64.

*We thank Wouter den Haan, Tamas Papp and Étienne Wasmer for helpful discussions, and seminar participants in Barcelona, Linz, Oxford, Paris, Seoul, Stockholm and Vienna for many useful comments. Robert Shimer provided part of the raw CPS data.

[#]christian.haefke@ihs.ac.at

[‡]michael.reiter@ihs.ac.at

[§]Address for Correspondence: Stumpergasse 56, A-1060 Vienna, Austria.

1 Introduction

The elasticity of aggregate labor supply has important implications for many policy questions, such as optimal taxation. In the literature, there has been a tension between the elasticities estimated in micro-econometric studies, which tend to be relatively low, and the elasticities that are implicit in macroeconomic models of the business cycle. This debate is documented in Chetty et al. (2011, Table 1), and we reproduce the key numbers in our Table 3. While some of these numbers may be disputed, several recent papers have worked to reconcile the order of magnitude difference between estimates of the extensive margin labor supply elasticity based on micro-econometric identification strategies (0.28) and the Frisch labor supply elasticities typically needed in RBC style macro models (2.3). These papers are surveyed in Keane and Rogerson (2011, Section 2.4).

The purpose of our paper is to illustrate that there need not be a discrepancy between micro- and macro-elasticities, if one interprets the macro-evidence in the light of a standard search model of the labor market. We enrich the Mortensen-Pissarides model by a participation decision of workers. Then labor supply is reflected in labor market participation. The fluctuations of participation over the business cycle suggest a labor supply elasticity at the extensive margin in the range of 0.3 to 0.4 for men and 0.4 to 0.65 for women. This is well in the range of available micro-studies, although it is higher than what Chetty et al. (2011) suggest. The large fluctuations in employment, which are usually regarded as an indicator of high labor supply elasticity, come from changes in labor demand, reflected in the number of vacancies posted by firms. Variations in vacancies can be a consequence of wage rigidity, vacancy cost fluctuations, or other sources. The issue is, in our perspective, not micro vs. macro, but labor supply vs. labor demand. Labor supply is well measured by recent micro-econometric studies, while labor demand, similar to physical investment, is very volatile over the business cycle.

In a nutshell, our results are driven by the following facts. Over the business cycle, wages and job finding probabilities vary considerably. For example, the standard deviation of log average hourly earnings, relative to that of GDP, is 74 percent, and the wages of new jobs fluctuate even more, cf. Table 4. The job finding probability fluctuates much more than GDP. In contrast, the relative standard deviation of log participation is only 23 percent. A high labor supply elasticity (unity or above) would imply a much stronger response of participation. A very low labor supply elasticity (0.3 or lower) would imply an even lower variability of participation, or a smaller correlation of participation and GDP. These two statistics, the variability of participation and its correlation with GDP, identify the range of admissible labor supply elasticities.

Since our estimates are derived from a structural business cycle model, the results are only credible if the model is successful in explaining the dynamics of labor market aggregates over the business cycle. Earlier attempts to incorporate participation into a standard matching model (Veracierto 2008) generated a positive correlation between the unemployment rate and GDP, which is highly counterfactual. In contrast, we calibrate our model so as to generate the strong negative correlation between unemployment and GDP that is observed in the data. To make this possible, it is necessary to replicate the observed strong procyclicality of the job finding probability.

Apart from providing an estimate of the aggregate labor supply elasticity, our model provides a simple framework that combines endogenous labor supply with search and matching frictions in a way that is consistent with key labor market data. We find that endogenous participation helps in explaining the variability of employment over the cycle. It gives a greater magnification of technology shocks than the standard matching model.

In recent years, there have been a number of papers that model participation decisions in a job matching framework. In many respects, our model is similar to Garibaldi and Wasmer (2005), who study labor market policies in a model with endogenous participation. They focus on steady state analysis. Krusell et al. (2009) go a step further by introducing worker saving. We do not do this, in order to keep the model as simple as possible. Pries and Rogerson (2009) have emphasized the heterogeneity of expected durations in participation spells and implications for cross-country differences. Krusell et al. (2011) is complementary to our analysis. They focus mainly on explaining the gross flows between employment states and allow for both worker saving and changes in worker productivity. In independent work, Shimer (2011) shows that wage rigidity is helpful in explaining business cycle facts about participation. This is compatible with our findings below. Chang and Kim (2006) and Chang and Kim (2007) derive the labor supply elasticity from the cross-sectional distribution of labor productivity and capital, but they have no matching friction in their model. For a more detailed discussion of their work, see Section 4.2.3.

The outline of the paper is as follows. Section 2 presents the theoretical model and some analytical results. Section 3 discusses the data and calibration. Section 4 provides numerical results for the calibrated model and Section 5 concludes.

2 The model

2.1 Overview

Our aim is to stick as closely as possible to the standard Mortensen/Pissarides matching model, to focus on the issue of labor supply elasticity, and to keep the analysis transparent. In particular, our model does not aim at replicating the gross flows between employment states. We do therefore not model the stochastic processes of the individual productivity variables, which give rise to large flows between employment states, but rather keep individual productivities constant for each worker.

We also keep the model simple in other respects. We do not model endogenous search intensity, nor the intensive margin of labor supply, nor are there endogenous job separations in equilibrium. Most importantly, we abstract from worker saving and thereby from wealth effects on labor supply. This will be discussed in Section 2.5.1.

However, we have to deal with the problem pointed out in Shimer (2005) and Costain and Reiter (2008), called the “unemployment volatility puzzle”, that the standard matching model does not generate enough variation in vacancy posting, and therefore in labor market tightness and job finding probabilities. The strong procyclicality of those variables is an essential feature of the data, and plays a crucial role for our analysis. We achieve it in the model by a combination of two features.¹ First, we assume fluctuations in vacancy cost, which are potentially correlated with fluctuations in productivity. This captures an effect proposed by Pissarides (2009), and is consistent with micro-foundations via credit frictions proposed by Petrosky-Nadeau and Wasmer (2010). Second, we assume a certain degree of wage rigidity, but only to the extent that is consistent with the findings of Haefke et al. (2008). We do not claim that our formulation is a valid structural explanation of the unemployment volatility puzzle. We rather use it as a shortcut, because we think that the exact reasons for the strong fluctuations in vacancy posting are not essential for our results. In the same spirit, Krusell et al. (2011) take these fluctuations as exogenous.

¹The literature often deals with this problem by a calibration similar to Hagedorn and Manovskii (2008). An essential element of this calibration is a small job surplus. In our model, the average job surplus is not a free parameter, but is determined by the cross-sectional dispersion of market and home productivity, cf. the discussion in Haefke and Reiter (2006). We therefore cannot follow this calibration strategy.

2.2 Model Setup

2.2.1 Households

The economy is populated by a unit mass of one-person households, which we call workers. At any time t , workers are heterogeneous in four dimensions: their gender (male or female), which is indicated by the subscript g with $g \in \{m, f\}$; their idiosyncratic market productivity y_i ; their gender specific mean home productivity h_g ; and their employment status $s(i, t)$, which can be either employed, $s(i, t) = e$, or not employed, $s(i, t) = n$. The idiosyncratic component of market and home productivity differs across households, but it remains constant over the life of each household.

Men and women only differ in the cross-sectional distribution of market and home productivity. For each $g \in \{m, f\}$, the log of market productivity is distributed normally with mean $\overline{\log y_g}$ and standard deviation $\sigma_{y,g}$. We do not model home productivity as an independent variable, but we allow home and market productivity to be systematically related by

$$h_{i,t} = \left(\bar{h}_{g(i)} + h_y (y_i - e^{\overline{\log y_{g(i)}}}) \right) / \tilde{h}_t \quad (1a)$$

$$\log \tilde{h}_t = \rho_h \log \tilde{h}_{t-1} + \sigma_h \epsilon_{h,t} \quad (1b)$$

where $g(i)$ denotes the gender of worker i , \bar{h}_g is the gender-specific median of home productivity, and the parameter h_y measures the strength of the relationship between home and market productivity. The i.i.d. shock $\epsilon_{h,t}$ is standard normally distributed and can be interpreted as a shock to aggregate labor supply. It decreases average home productivity and leads to an inflow of workers into the labor market.

Employed workers decide whether to continue an employment relationship or not. Since the shocks in the model are sufficiently small, there are no endogenous separations in equilibrium. However, existing matches are separated exogenously with constant flow probability χ . For those who are not employed, the decision is whether or not to search for a job. Those who search find a job with probability p^W per period, which only depends on the aggregate state of the economy, not on any action by the worker. Non-employed agents who search for a job are called *unemployed* and by definition participate in the labor market. The mass of unemployed households in the economy at time t is denoted by U_t . Those who do not search for a job are called non-participants, or *out of the labor force*. We assume that there is no cost of switching between unemployment and non-participation, and a household's decision to participate is therefore independent of whether the household participated last period or not. In this sense, non-participation is not a

third (predetermined) labor market state, but only a decision variable of the household; we discuss this in Section 2.5.2.

All the non-employed agents engage in home production. We think of this utility as all kinds of returns to being at home, from taking care of your children or growing vegetables in your backyard to moonlighting. The utility flow of worker i then equals his or her respective home productivity $h_{i,t}$ unless the worker searches actively, in which case the utility flow reduces to $h_{i,t}(1 - \gamma_t)$. The variable γ_t measures the cost of job search, expressed as fraction of home production. Later we will calibrate this parameter by the time spent on searching. For reasons of numerical stability, we allow search costs to be random, being i.i.d. over time and across households. However, we think of the variability of γ as being very small.

Those who are employed do not enjoy any home production, but they earn a wage w . Following the bulk of the matching literature, we assume that wages are determined in each period by generalized Nash bargaining. The wage thus depends on the output of the match and on home productivity, among other things.

2.2.2 Production

Production takes place in one-worker firms using labor input only. Output of the firm in period t equals $Z_t y_i$, the product of aggregate productivity Z_t and of idiosyncratic market productivity y_i of worker i who is matched to the firm. Aggregate productivity follows an AR(1) in logs, such that

$$\log Z_t = \rho_Z \log Z_{t-1} + \sigma_Z \epsilon_{Z,t} \quad (2)$$

where $\epsilon_{Z,t}$ is a standard normal i.i.d. shock.

2.2.3 Job Creation

To form new employment relationships, workers must search and firms must post vacancies. All firms offer the same conditions (wage etc.) to a worker of a given characteristic, and therefore all vacancies are identical. Opening a vacancy comes at a flow cost κ_t , which varies over time

$$\kappa_t = \bar{\kappa} \tilde{\kappa}_t \quad (3a)$$

$$\log \tilde{\kappa}_t = \rho_\kappa \log \tilde{\kappa}_{t-1} + \sigma_\kappa \epsilon_{\kappa,t}. \quad (3b)$$

with $\epsilon_{\kappa,t}$ a standard normal i.i.d. shock. We allow for nonzero correlation $\rho_{Z,\kappa}$ between aggregate productivity and vacancy posting costs. For negative $\rho_{Z,\kappa}$ this setup has a similar effect as the fixed and variable vacancy cost components in Pissarides (2009).

Table 1: Timing Within a Period.

-
-
1. The aggregate productivity shock realizes.
 2. The idiosyncratic search costs realize.
 3. Participation decisions and separation decisions are taken; vacancies are created.
 4. Matching and exogenous separation shocks realize.
 5. Production, consumption etc. take place.
-

In line with the literature, we assume that the total number of new matches M is produced by a constant returns to scale Cobb-Douglas matching function

$$M = \mu U^\lambda V^{1-\lambda}. \quad (4)$$

Here U is the size of the unemployment pool (the non-employed who search), V is the total number of vacancies, μ is a scaling parameter and λ is the matching elasticity. Defining aggregate labor market tightness $\theta \equiv V/U$, we obtain the matching probabilities of the firm (p^F) and the worker (p^W) as

$$p_t^F \equiv \frac{M_t}{V_t} = q(\theta_t) \quad (5)$$

$$p_t^W \equiv \frac{M}{U} = \theta q(\theta_t), \quad (6)$$

where

$$q(\theta) \equiv \mu \theta^{-\lambda}. \quad (7)$$

2.2.4 Timing of Events

Within each period, events happen in the order presented in Table 1. Exogenous separation and matching takes place at the same time. This implies that a worker who gets exogenously separated cannot participate in this period's matching market, but has to wait (and be unemployed) for at least one period. Workers who are matched this period cannot be separated the same period. This timing serves to simplify the model setup because each agent only faces one binary shock per period (being separated or not when employed; or being matched or not when unemployed). Furthermore, people who decide to search will be matched (or not) the same period (i.e., without productivity shocks happening in the meantime) which ensures that every matched agent accepts the current job offer.

Employment E_t denotes the mass of employed workers after separation and matching shocks of period t have realized. Using the above timing assumptions, and anticipating

that there are no endogenous separations, the dynamics of employment is summarized by

$$E_{t+1} = (1 - \chi)E_t + p_{t+1}^W U_t. \quad (8)$$

Remember that U is not equal to non-employment ($1 - E$), but only contains those workers who decide to search.

2.3 Recursive Formulation

2.3.1 Participation and Separation

The value function at time t of an employed and of a non-employed agent of gender g and market productivity y are denoted by $V_{g,t}^E(y)$ and $V_{g,t}^N(y)$, respectively. Value functions of time t are defined after matching and separation shocks have realized. The time subscript t captures the effect of aggregate shocks on the value function.

Employed agents decide whether to continue working or whether to separate and leave the labor force. Standard techniques can then be used to show that $V_{g,t}^E(y)$ satisfies the recursive relationship

$$V_{g,t}^E(y) = W_{g,t}(y) + \beta \mathbb{E}_t \max \left\{ \begin{array}{ll} (1 - \chi)V_{g,t+1}^E(y) + \chi V_{g,t+1}^N(y); & \text{staying} \\ V_{g,t+1}^N(y) & \text{endog.sep.} \end{array} \right\} \quad (9)$$

where $W_{g,t}(y)$ denotes the wage of the worker with characteristics g and y . The non-employed decide whether to participate in the labor market as unemployed, or not to participate, therefore

$$V_{g,t}^N(y) = h_{g,t}(y) + \beta \mathbb{E}_t \max \left\{ \begin{array}{ll} p_{t+1}^W V_{g,t+1}^E(y) + (1 - p_{t+1}^W)V_{g,t+1}^N(y) - \beta^{-1}\gamma_{t+1}h_{g,t+1}(y); & \text{unemployed} \\ V_{g,t+1}^N(y) & \text{non-partic.} \end{array} \right\} \quad (10)$$

The value of the search cost at which the worker is indifferent between participation and non-participation, denoted by $\gamma_{g,t}^*(y)$, satisfies

$$h_{g,t}\gamma_{g,t}^*(y) \equiv p_t^W (V_{g,t}^E(y) - V_{g,t}^N(y)). \quad (11)$$

A worker employs a simple threshold strategy and participates if $\gamma \leq \gamma_{g,t}^*(y)$. Therefore (10) can be written as

$$V_{g,t}^N(y) = h_{g,t}(y) + \beta \mathbf{E}_t \left[V_{g,t+1}^N(y) + p_{t+1}^W (V_{g,t+1}^E(y) - V_{g,t+1}^N(y)) F_\gamma(\gamma_{g,t+1}^*(y)) - h_{g,t+1} \int_{-\infty}^{\gamma_{g,t+1}^*(y)} \gamma dF_\gamma(\gamma) \right], \quad (12)$$

where $F_\gamma(\gamma)$ denotes the distribution function of the search cost.

We think of the variation in γ as very small. In the limit, γ is a constant, and the participation decision is characterized by the threshold $y_{g,t}^*$, such that a household of gender g at time t participates if and only if $y \geq y_{g,t}^*$.

If the support of γ is bounded away from zero, it follows from (11) that a worker participates if and only if the value of employment is strictly greater than the value of non-employment. Since the idiosyncratic market productivity is constant for the worker, and if the fluctuations in aggregate productivity Z_t are sufficiently small, then the worker will never decide to quit, so there are no endogenous separations.

2.3.2 Job creation

The value of having a job filled with a worker of gender g and market productivity y is denoted by $J_{g,t}(y)$. Since total output is the product of aggregate productivity Z and individual productivity y , it satisfies

$$J_{g,t}(y) = Z_t y - W_{g,t}(y) + \beta(1 - \chi) \mathbf{E}_x \max \left\{ \begin{array}{ll} J_{g,t+1}(y); & \text{continuation} \\ 0 & \text{endog.sep.} \end{array} \right\}. \quad (13)$$

In order to create a vacancy, firms have to pay a vacancy posting cost of $\kappa(Z)$ per unit of time. Free entry assures that the expected net value of a vacancy is zero. Therefore, the expected surplus that a worker generates for the firm has to equal the expected hiring cost:

$$\kappa_t = p_t^F \overline{J(t)} \quad (14a)$$

$$\overline{J(t)} \equiv \frac{\sum_{g \in \{m,f\}} \left[\int_{y_{g,t}^*}^{\infty} J_{g,t}(y) dF_{g,t}^u(y) \right]}{\sum_{g \in \{m,f\}} \left[\int_{y_{g,t}^*}^{\infty} dF_{g,t}^u(y) \right]}. \quad (14b)$$

Here, $F_{g,t}^u(y)$ denotes the cross-sectional distribution of market productivity among unemployed. We see from (14) that, in order to compute the expected value of a vacancy, a firm needs to know the composition of the unemployment pool, represented by $F_{g,t}^u(y)$. This is an infinite-dimensional aggregate state variable, which is the reason why the analysis is substantially more complicated compared to the model without participation. In (14b), the integral is taken over all market productivities and over both sexes, which reflects the assumption of a single, non-segmented labor market.

2.3.3 Equilibrium Wage and Labor Market Tightness

Wages are determined by generalized Nash bargaining. The corresponding first order condition is

$$\alpha(Z)J_{g,t}(y) = (1 - \alpha(Z)) (V_{g,t}^E(y) - V_{g,t}^N(y)) \quad (15)$$

with bargaining weight $\alpha(Z)$ of the worker. Similar to κ , α is allowed to vary over the cycle as a function of Z , such that firms potentially grab a bigger part of the surplus in booms:

$$\alpha_t = \bar{\alpha} - \Delta_\alpha(Z - \bar{Z}), \quad (16)$$

for $\Delta_\alpha \geq 0$.

2.4 Some Analytical Insights

2.4.1 Aggregate Productivity and Participation

To obtain some insight into the properties of our model, consider the limiting case where the search cost γ is a constant. For a given aggregate labor productivity, there is a critical value y_c of market productivity where households are indifferent between participating and staying at home. The analysis is done separately for each gender $g \in \{m, f\}$, but in the following the gender subscript is omitted for notational convenience.

Participation for each gender is given by the mass of agents with idiosyncratic market productivity above the critical value, such that

$$P = \int_{\ln y_c}^{\infty} f(\ln y) d \ln y = 1 - F(\ln y_c)$$

where $f(\cdot)$ and $F(\cdot)$ denote the density and cumulative distribution function of log market productivity, respectively. The elasticity of participation with respect to aggregate productivity is then

$$\frac{d \ln P}{d \ln Z} = -f \frac{1}{P} \frac{d \ln y_c}{d \ln Z} = -\frac{f(\ln y_c)}{1 - F(\ln y_c)} \frac{d \ln y_c}{d \ln Z} \quad (17)$$

The response of participation to changes in aggregate productivity is determined by two factors: the change in the threshold value y_c , and the mass of unemployed households at the margin, relative to total participation.

The threshold value y_c of individual market productivity, where households start participating, satisfies

$$\frac{\gamma h}{p^W} = \frac{\beta}{1 - \beta + \chi\beta} \alpha (Z y_c - h) \quad (18)$$

γh is the per period utility reduction when searching, $1/p^W$ the expected number of periods spent searching. At the threshold this utility loss is offset by the premium of wages over the outside option, which is basically the worker's share of the flow surplus.

Taking the total derivative of the log of (18), and evaluating at $Z = 1$, we get (recall that λ is the elasticity of matches with respect to unemployment)

$$h_y \frac{y_c}{h} d \ln y_c - (1 - \lambda) d \ln \theta = \frac{1}{y_c - h} (y_c d \ln Z + (1 - h_y) y_c d \ln y_c). \quad (19)$$

Defining $\eta_y^h \equiv h_y \frac{y_c}{h}$, we can rearrange (19) to get

$$\frac{d \ln y_c}{d \ln Z} = -\frac{1}{1 - \eta_y^h} \left[1 + \frac{y_c - h}{y_c} (1 - \lambda) \frac{d \ln \theta}{d \ln Z} \right] \quad (20)$$

The response of participation to changes in productivity can then be derived to be:

$$\begin{aligned} \frac{d \ln P}{d \ln Z} &= -f \frac{1}{P} \frac{d \ln y_c}{d \ln Z} = -\frac{f(\ln y_c)}{1 - F(\ln y_c)} \frac{d \ln \ln y_c}{d \ln Z} \\ &= \frac{f(\ln y_c)}{1 - F(\ln y_c)} \frac{1}{1 - \eta_y^h} \left[1 + \frac{y_c - h}{y_c} (1 - \lambda) \frac{d \ln \theta}{d \ln Z} \right] \end{aligned} \quad (21)$$

i.e. the mass of agents at the participation threshold times the change in the threshold.

From (21), there are four factors determining the response of participation:

1. The cross-sectional dispersion of $\ln y$. The higher the variance, the lower is the density f at a given quantile of the distribution, i.e., for a given participation rate.
2. The average participation P . For a normal distribution, the ratio $\frac{f(\ln y_c)}{1 - F(\ln y_c)}$ increases monotonically in $\ln y_c$, which means that it decreases monotonically in the participation rate. For participation rates above 50 percent, the numerator decreases and the denominator increases with participation.
3. The relationship between market and home productivity h_y . A higher h_y causes larger changes in the threshold value y_c , therefore higher elasticity. Another interpretation of this is that a higher h_y makes the difference between market and home productivity less dispersed, therefore the density of this difference at the margin becomes larger.

4. The response of labor market tightness to productivity. A tighter labor market reduces effective search costs. Notice that this effect is scaled by $\frac{y_c - h}{y_c}$. If search costs are small anyway, h is close to y_c , and tightness has little effect on participation.

2.4.2 Labor Supply Elasticity

Rogerson and Wallenius (2009) stress that, in a dynamic framework, there is not just one number that measures labor supply elasticity. In the framework of our model, we define labor supply elasticity as the long-run reaction of labor market participation to a change in aggregate labor productivity, keeping labor demand conditions (i.e., the job finding probability) constant. With “long-run reaction”, we refer, technically speaking, to a comparative steady state analysis. In the context of our model, this should be interpreted as the reaction until a new flow equilibrium in the labor market is reached. It does not include wealth effects from permanent changes in productivity, because there are no wealth effects in our model. Neither does it include very long-run decisions of households, such as education or fertility choices.

To compare our definition to the estimates in the empirical micro literature, notice two things. First, the labor literature studies the elasticity of employment, rather than participation, with respect to wages. However, with constant job finding and separation rates, employment is a constant fraction of participation in steady state, and therefore the distinction employment vs. participation is irrelevant for our definition of the labor supply elasticity.

Second, labor supply elasticity is usually defined in terms of a response to wages, not to labor productivity. We prefer not to do this, for two reasons. First, the wage is the outcome of a bargaining process between workers and firms, not a parameter that workers take as given. Second, there is not one wage, but the bargained wage is influenced by both market and home productivity of the workers. In contrast, aggregate labor productivity is an exogenous parameter in steady state. In any case, the distinction between wages and labor productivity makes a small difference, because the elasticity of wages to productivity is close to unity in steady state.

Using the chosen definition, it is clear that the determinants of labor supply elasticity are described by the first three points in the list of Section 2.4.1. The fourth point, namely the response of labor market tightness to productivity, is a general equilibrium effect that reflects the labor demand decisions of firms. This effect is not included in our definition of labor supply.

2.5 Discussion

2.5.1 Worker Saving and Wealth Effects

In the model, workers do not save. This has at least two important consequences. First, workers do not engage in short-run intertemporal substitution between work in different time periods. This mechanism, which can lead to very high Frisch labor supply elasticities in models of indivisible labor, was stressed in recent work by Keane and Rogerson (2011), Ljungqvist and Sargent (2011), and Rogerson and Wallenius (2009). It is interesting that our model explains the volatility of labor supply only through the cross-sectional variation in the participation margin, similar to Chang and Kim (2006). This complements the recent paper by Krusell et al. (2011) who focus mainly on intertemporal substitution in a model of large households where they impose exogenous job finding and separation probabilities to study labor market flows.

Second, wage increases do not have wealth effects. We think that, within our model, this effect is not very big. One reason is that productivity changes have an autocorrelation of less than 0.9 quarterly, so that the wealth effect is clearly dominated by the substitution effect. Moreover, the wealth effect is not relevant for the participation decision of a single-earner non-employed household, because higher wages do not create a wealth effect if the household does not participate. The wealth effect only comes into play for endogenous quits, when the household has accumulated so much assets that it does not want to continue working, and for double earner household, where the wage of one earner affects the participation decision of the other earner (Faraglia 2003; Chang and Kim 2006). While those are important issues, we think that for business cycle purposes, a model without wealth effects is a useful benchmark. The empirical success in explaining the movements of participation supports this view.

2.5.2 Fixed costs

In our model, there are no fixed costs of switching between non-participation and participation. The question is whether there are such fixed costs, and whether they would affect the main results. We find fixed costs of switching hard to reconcile with the substantial gross flows between non-participation and participation. Using the CPS data described in Section 3.1, we find that each month, about 14 percent of the unemployed switch to non-participation, and about 2 percent of the non-participants switch to unemployment. Furthermore, we have experimented with fixed costs in earlier work, and found that it had little impact on the main results (Haefke and Reiter 2006, Footnote 4).

However, there are probably fixed costs of labor supply at a very long horizon. Those include, for example, the costs of getting a professional education, or costs related to fertility decisions. It is therefore important to interpret our results as labor supply elasticities at business cycle horizons, not necessarily at very long horizons. We cannot rule out higher labor supply elasticities in the long run, as was suggested, for example, by Rogerson and Wallenius (2009). Similarly, the microeconomic studies, to which we compare our results, are unlikely to pick up these very long-run elasticities.

2.5.3 Measurement Errors in Participation

Labor market participation is a somewhat vague concept, and difficult to measure. In practice, it is based on workers' answers to survey questions. It could be that the reported behavior is more sluggish than actual behavior. For example, some of the unemployed might always say they are looking for work, although in reality they are only looking in a boom, when conditions are good enough. Estimated changes in participation then underestimate true changes, and we underestimate the labor supply elasticity. This hypothesis can be tested. If it is true, then in a boom, there are fewer non-participants who are counted as unemployed, or more searchers who are counted as non-participant. In either case, it implies that the measured *NE*-hazard increases more strongly in a boom than the *UE*-hazard.

In the data, we find the opposite. The covariance between the unemployment rate and the log of the ratio of the *NE*-hazard and the *UE*-hazard equals 0.0860, which corresponds to a correlation of 0.0733. This means that in recessions, when unemployment is high, the *UE*-hazard drops even more than the *NE*-hazard. Therefore, although the *NE*-flows are substantial, the changes of unemployment over the business cycle are driven by changes in the *UE* more than by changes in the *NE*-hazard. More precisely, if unemployment is higher by one percentage point, the *UE*-hazard decreases by 8.6 percent compared to the *NE*-hazard, a quite strong effect. This serves as further evidence that the distinction between unemployment and non-participation is an important one.

3 Data and Calibration

3.1 Data

We consider data for the time period 1979–2007. All statistics (except the means) in table 5 are reported for the log of quarterly series which have been detrended using an HP filter

with smoothing parameter 10^5 . Over this time, labor productivity defined as GDP per worker fluctuated about 75% of GDP. Persistence measured by first order autocorrelation was high at 0.9.

We consider the entire population aged 16–99, male and female, rather than restricting ourselves to certain age groups as the previous papers in this literature (Faraglia 2003; Garibaldi and Wasmer 2005) have done. Even though the precise numbers we report are not identical to theirs, the qualitative picture remains the same. The overall participation rate averaged 65.8%, which can be broken down into 75.4% for men and 57% for women. Persistence was high at .89. Moderate procyclicality at a correlation of .57 with GDP and the low variability (23% of GDP overall, 19% for men and 35% for women) have so far often been used as arguments to abstract from the participation decision. Veracierto (2008) recently argued that a Mortensen-Pissarides style search and matching model cannot match these low participation fluctuations. For us, the low volatility of participation provides key information for identifying the labor supply elasticity.

The employment population ratio is highly procyclical, as indicated by a .85 correlation with GDP, and its variance over the business cycle is at 69% of GDP. Generating such strong employment variation has been a challenge for many RBC models. Endogenizing the participation decision is very helpful in this dimension, as discussed by Haefke and Reiter (2006) and Veracierto (2008).

The unemployment rate fluctuated substantially around its mean of 6.1%, with a standard deviation 7.4 times higher than that of GDP. As was pointed out by Shimer (2005) and Costain and Reiter (2008), standard matching models have a hard time explaining the high volatility, which will be a crucial issue in our calibration. The unemployment rate is also highly persistent and strongly countercyclical with a correlation of $-.85$ with GDP. Replicating this countercyclicality has been a major difficulty for previous models with search frictions and endogenous participation (Faraglia 2003; Tripier 2002; Veracierto 2008) or search intensity (Shimer 2004), because non-participants start searching at the onset of a boom but take some time to find a job. We will explain below that the problems of the high volatility and of the strongly negative correlation are closely linked. To solve those problems, it is crucial to capture the movements of the job finding probability. According to Robert Shimer’s data, it averaged 42.5% for monthly rates, is strongly procyclical and fluctuates 5.25 times as strongly as GDP.

A lot of attention has recently been given to wages. Using different approaches and data sets, several authors (Pissarides 2009; Haefke et al. 2008; Hagedorn and Manovskii 2010) find that the wage relevant for job creation fluctuates strongly with labor productivity. Similar to Michaillat (2011), our preferred estimate is 0.8 for the elasticity of wages with

respect to labor productivity. In addition to time series volatility, we also consider the cross sectional distribution of wages for men and women separately. Women’s wages are on average 25.7% below men’s wages. We capture the cross-sectional wage dispersion by considering average ratios of the median to the 10-th percentile for each sex.

3.2 Calibration

The model period is $1/48$ of a year, corresponding roughly to one week. The following parameters are set to conventional values, and are kept constant across different calibrations. The interest rate is set to 4% annually, so $\beta = 1/1.04^{1/48}$. For the elasticity of matches w.r.t. unemployment we follow Reiter (2008) and set $\lambda = 0.4$. This is well within the range of values that Petrongolo and Pissarides (2001, Table 3) report. With this value, the model does equally well in explaining the variability of unemployment and the variability of tightness. If one uses a higher value, as Shimer (2005) and Rudanko (2009) do, the model tends to underestimate the variability of unemployment, and overestimate that of tightness. Following most of the literature, we set the worker’s share of the surplus α equal to λ . We normalize the efficiency parameter μ in the matching function to 1, which implicitly scales the number of vacancies but otherwise has no real consequences.

A parameter we take directly from micro data is search cost. Average search costs $\bar{\gamma}$ are set, separately for each gender, to 6.56 percent of median home productivity. This reflects the average time spent on job search of the unemployed, as percent of a 36 hours workweek based on the American Time Use Survey for 2003–2006. As explained in Section 2.3.1, we assume a small i.i.d. variation in search costs, for numerical reasons. This is specified in Appendix A.

The autocorrelation parameters for the aggregate shock processes (2), (1b) and (3b) are set independently of the rest of the calibration. To obtain the autocorrelation of aggregate labor productivity, we identify it with average labor productivity in the data. We set $\rho_Z = 0.9970$ at weekly frequency such that the process of labor productivity, aggregated to quarterly frequency and HP filtered with weight 10^5 , has a first-order autocorrelation of 0.8974, which we find in the data. Since we have no data to estimate persistence in the first order process for vacancy posting costs (ρ_κ) and aggregate home productivity (ρ_h) we set both equal to ρ_Z for reasons of parsimony. Notice that our choice is close to the 0.91 estimated by Floden and Lindé (2001) for idiosyncratic productivity using U.S. data.

There remain 13 parameters which are set so as to make the model match 13 targets (8 steady state targets and 5 second moments) in the data. These parameters can vary between different calibrations. Table 2 reports the values for the baseline calibration, next

Table 2: Baseline Calibration

Parameters			Targets	
χ	Weekly Separation Rate	0.844%	Weekly Job Finding Prob.	12.92%
κ	Vacancy Posting Cost	3.94	Unemployment Rate	6.1%
\bar{h}_f	Home Productivity for Females	0.2647	Participation Rate for Females	57.1%
\bar{h}_m	Home Productivity for Males	0.2618	Participation Rate for Males	75.4%
$\overline{\log y_f}$	Average Productivity for Females	0.77	Wage Ratio female/male	0.7734
$\overline{\log y_m}$	Average Productivity for Males	1	Average Wage Male	1
$\sigma_{y,f}$	Std.Dev. Female Market Productivity	1.036	P50-P10 Ratio for Females	1.76
$\sigma_{y,m}$	Std.Dev. Male Market Productivity	1.065	P50-P10 Ratio for Males	2.06
σ_Z	Stdev Aggr. Prod. Shock, Weekly	0.0012	Quarterly Std.Dev of GDP	1.35
σ_h	Stdev. Labor Supply Shock, Weekly	0.0010	StdDev Partic. Rate, rel. to GDP	0.23
σ_κ	StdDev Vacancy Cost Shock, Weekly	0.0258	StdDev Employment, rel. to GDP	0.69
$\rho_{Z,\kappa}$	Corr. Agg. Prod. and Vac.Cost Shock	-0.181	Corr. Unemployment and GDP	-0.85
Δ_α	Change in Worker Bargaining Power	2.275	Elast. of New Wages to Aggr.Prod.	0.8

to the targets that they match. To give some intuition, the table pairs each parameter with a specific target that is prominently influenced by this parameter, although all parameters are determined simultaneously. For the calibration targets, we set the steady state wage of men to unity, which is just a normalization. The job finding probability was taken from Shimer (2007). He reports a monthly probability of 42.51%. We compute the weekly probability as $1 - (1 - 0.4251)^{1/4}$.

Apart from scaling GDP fluctuations, we have four second moments as calibration targets, namely the relative standard deviations of employment and participation, the correlation of unemployment with GDP, and the elasticity of average wage w.r.t. aggregate productivity. To match those, it turns out that the calibrated standard deviation of the labor supply shock is slightly higher than that of the productivity shock. The standard deviation of the vacancy cost shock is about 20 times higher than that of the productivity shock, and the two shocks are negatively correlated (-0.181). This is in line with the idea of Pissarides (2009), but notice that the required fluctuations in vacancy costs are large. The derivative of firms bargaining power w.r.t. aggregate productivity is 2.275. If productivity increases by 2.12 percent, which is the unconditional standard deviation of aggregate productivity in the baseline calibration, the bargaining power of firms increases from 0.6 to $0.6 + 2.275 \cdot 0.0212 = 0.6481$, a moderate increase.

Next to this *baseline calibration*, we also define a *one-shock calibration*. It is the same as the baseline, except for setting $\sigma_h = \sigma_\kappa = \rho_{Z,\kappa} = \Delta_\alpha = 0$. The only shock is to aggregate productivity, which is set so as to match the variability of GDP, and wages are perfectly

flexible. This does not change the steady state of the model, but we have dropped the calibration targets referring to second moments. This calibration comes closest to the standard version of the Mortensen/Pissarides model.

4 Results

4.1 Labor Market Fluctuations

Before going to the baseline calibration, we first consider the results for the one-shock calibration, which does not target the second moments. So it is lacking all the extensions that are necessary to generate a high labor market volatility. In particular, there are no labor supply or vacancy cost shocks, and wages are perfectly flexible. Results are shown in Table 7. The defects of this parameter combination are clearly visible. The relative standard deviation of unemployment, relative to GDP, is only 0.63, and unemployment is almost uncorrelated with GDP. The reasons for this negative result will become clear later. However, the model does very well in one dimension: the relative variability of participation is 0.22 for men and 0.35 for women, compared to 0.19 and 0.35 in the data. This is remarkable, because the calibration does not target those variables. Recall from Section 3.2 that home productivity is identical across households, conditional on gender. Participation is driven by the cross-sectional distribution of market productivity, which is identified by the cross-sectional distribution of observed wages. Interestingly, the model explains the volatility of female relative to male participation. Female participation fluctuates about 50 percent more. The reason for this is that the average female participation rate is lower, so that the participation threshold lies closer to the modus of the distribution, which gives a higher density at the threshold (cf. the discussion in Section 2.4.1). This result is consistent with the finding of Bargain et al. (2011) that women have significantly larger extensive margin labor supply elasticities in countries with low female labor force participation rates.

Table 6 displays the results for the baseline calibration. Notice first that the variability of participation is not substantially affected by the recalibration. It is mainly the cross-sectional distribution of market productivity, not the relative standard deviation of shocks, or wage rigidity, which drives this variability. With the baseline parameter values, the model matches a number of second moments quite well, apart from those that were used for the calibration. In particular, the correlation of participation with GDP is 0.69 in the model, for both men and women, while in the data it is 0.57 for men and 0.45 for women. The relative standard deviation of wages matches those for all workers in the data. One

dimension in which the model does not do so well is that it generates a too high volatility of the job finding probability. This is probably because the job separation rate is constant in the model, so that all the variability in unemployment comes from variations in the finding rate, while in the data part of it is explained by variations in separations.

Our calibration targets the strong negative correlation between the unemployment rate and GDP, in contrast to earlier attempts to incorporate participation into a standard matching model (Veracierto 2008). To understand how this works, it is useful to look at the impulse response functions in Figure 1. Panel a) gives the response to a one percent technology shock in the baseline calibration.² Numbers are percentage deviations from steady state. Notice that the responses of the unemployment rate and of the job finding probability have been scaled down by a factor of 10, to fit well into the same graph. We see that participation, and therefore unemployment, increases upon the technology shock, reflecting improved market opportunities over conditions at home. Despite of this, the shock generates a negative correlation between unemployment and GDP, because the job finding probability goes up strongly, so that the unemployed are quickly absorbed into employment. The strong effect of productivity on the job finding rate in this model is due to wage rigidity; in a one-shock calibration of the model, which suffers from the Shimer critique, it is well known that the increase in job finding is small (Shimer 2005; Costain and Reiter 2008). Then unemployment increases even more and stays above the steady state level for a long time. This is precisely what happens in the one-shock calibration, for which the results are shown in Panel b).

Next to the pro-cyclical job finding rate, two more aspects of the calibration are responsible for this success, as was already explained in Haefke and Reiter (2006). One is the calibration of the cross-sectional distribution of market productivity. If there is not enough heterogeneity across workers, too many are close to the participation threshold, such that the initial procyclical response of unemployment is bigger. To offset this effect, an even stronger response of the job finding probability is necessary. Another, quantitatively less important factor, is time aggregation. If the model is solved and simulated at a weekly rather than quarterly frequency, new entrants into the job market can be absorbed more quickly. If the weekly data are aggregated to monthly or quarterly values, the initial rise in unemployment becomes less important. This can be seen by comparing the quarterly correlations (ρ) with the monthly (ρ^m) and weekly (ρ^w) correlations reported in Tables 6–8.

Panels c) of Figure 1 displays the response to a ten percent vacancy cost shock. We see that the shock reduces the job finding rate, but does little to participation. In our

²For this exercise, we keep the vacancy cost constant, although it is correlated with technology.

calibration, the cost of searching is small, and it is therefore the comparison of home vs. market productivity, not the probability of finding a job, which mainly drives the participation decision (cf. the discussion in Section 2.4.1). Panel d) displays the response to a labor supply shock, which turns out to be qualitatively very different from a technology shock. The increase in participation leads to a decrease in the job finding rate because the composition of the unemployment pool deteriorates, which lowers the incentives to post vacancies.³ This causes a sharp increase in unemployment. The correlation between GDP and unemployment serves to discriminate between technology and labor supply shocks.

Notice that in all cases the response of participation is hump-shaped. This is because households close to the participation margin increase their probability of participating, as their critical level of the search cost declines. This temporarily increases the fraction of non-employed households who participate, while those who happen to find a job stay employed, and therefore continue participating.

4.2 Labor supply elasticity

The parameter values of the baseline calibration reported in Table 6 imply steady state labor supply elasticities of 0.403 for men and 0.646 for women. The values are well below the preferred estimates of Keane and Rogerson (2011). They are above the extensive-margin estimates cited in Chetty et al. (2011, Table 1), but we think they are well in the range of recent micro estimates. For example, van Soest et al. (2002) estimate an extensive margin elasticity of 0.72 for Dutch women. We cite this study, because it probably comes closest to estimating the labor supply elasticity as we have defined it. From those results, there is no discrepancy between micro and macro elasticities.

The question then is: which features of the data identify the elasticity, and how robust are our estimates? Are there other calibrations of the model that fit the data similarly well but imply different labor supply elasticities? The estimates depend on our calibration of the model. At least two concerns with this calibration come to our mind. First, the degree of wage rigidity is uncertain. Second, the assumption of constant home productivity for each gender appears arbitrary. We examine the two issues in turn.

³If market productivity increases, workers who were at the border between participation and non-participation now enter the labor market. Since their job surplus was small by definition, they reduce the average match surplus (cf. the discussion in Haefke and Reiter (2006)).

4.2.1 Wage Rigidity

The cyclical nature of wages is not easy to measure. Relevant for our purpose are the wages for new hires. Our one-shock calibration uses the estimates of Haefke et al. (2008), which are based on CPS data, but which come with substantial standard errors. Calibrating the model with a higher degree of wage rigidity would lead to substantially higher estimates of the labor supply elasticities. This is because smaller fluctuations in wages give less incentives to households to participate, therefore the observed fluctuations can only be explained by larger elasticities. A natural case is where fluctuations are explained exclusively by productivity shocks ($\sigma_h = \sigma_\kappa = 0$). In that case, we are left with only two parameters to calibrate (h_y and Δ_α). We use those to match the relative standard deviations of employment and participation. The results are reported in Table 8. The estimated elasticities are now 0.577 for men and 1.026 for women, but the model generates two counterfactual results: a very high correlation between participation and GDP, and a variability of wages that is clearly below what is observed in the data.

4.2.2 Varying h_y

Unfortunately, the covariance between home- and market productivity is not observed. Our baseline choice of $h_y = 0$ is natural, but arbitrary. Increasing h_y , that is, increasing this covariance, increases the labor supply elasticity, as was explained in Section 2.4. To be compatible with the observed variability of participation, this has to be compensated by a lower variability σ_h of the labor supply shock.

Table 9 summarizes the model results for varying values of h_y . The first column reflects a case that can be considered a natural upper limit for h_y . By setting $h_y = 0.236$, the variability of participation is matched without labor supply shocks, that means $\sigma_h = 0$. In the following columns, h_y gets reduced down to -0.6, at which the labor supply elasticities are similar to what Chetty et al. (2011) suggest. The column $h_y = 0$ refers to the baseline case. To explain the variability of participation in each case, we have to increase the size of the labor supply shocks. The size of the other shocks is not much affected. The last part of the table reports the statistics that vary most from one calibration to the next: the correlation of participation (male and female) with GDP, the variability of wages, and its correlation with GDP.

With $h_y = 0.236$, the elasticities are 0.515 for men and 0.866 for women, considerably higher than in the baseline case. However, this raises the correlation between GDP and participation to a value of 0.9, whereas it is 0.57 for men and 0.45 for women in the data. In the framework of our model, such a high labor supply elasticity is incompatible with

the data.

To match the correlation of total participation with GDP, we need a value of around $h_y = -0.15$. In this case, the elasticity is 0.33 for men and 0.55 for women. A negative h_y may appear implausible, but it need not be taken literally. If there is variation in home productivity h that is independent of the variation in market productivity y , it increases the cross-sectional variation of $y - h$, just as a negative h_y does. We do not allow for independent variation in h because we do not know how to identify it.

The above elasticities are still a bit higher than the micro elasticities reported in Chetty et al. (2011). To replicate those, one has to set h_y around -0.6 , as done in the last column of the Table. This requires a standard deviation of the labor supply shock that is more than twice as high as the productivity shock. Then the correlation of participation with GDP is below what we find in the data, and wages become more volatile than GDP. Hence it is hard to reconcile those very low labor supply elasticities with some second moments in the data. Our best estimates are therefore 0.33 for men and 0.55 for women.

4.2.3 Comparison with Chang and Kim (2006)

Our model is in many respects simpler than the one in Chang and Kim (2006), which includes worker saving, two-earner households, and a stochastic process for idiosyncratic productivity. Nonetheless, our estimated labor supply elasticities for the baseline case (0.403 and 0.646 for men and women, respectively) are close to the individual compensated labor supply elasticities reported in Chang and Kim (2006, Table 8), which are 0.41 for men and 0.78 for women in their baseline model. This suggests that our simple model has focused on the features that are most important for labor supply elasticities.

We go beyond Chang and Kim (2006) in two respects. First, we have stressed in Section 2.3 that it is the cross-sectional distribution of home vs. market productivity, not just of market productivity, that drives the labor supply elasticity. This has led us to a slight reduction in estimated elasticities. Second, we have introduced labor market matching into the model, such that employment is driven to a great extent by fluctuations in labor demand (vacancy posting), which are influenced by changes in vacancy cost and by wage rigidities. In this way we disconnect the labor supply elasticity from fluctuations in employment and unemployment.

5 Conclusion

We have extended a labor market model of the Mortensen/Pissarides type to allow for endogenous labor market participation. The model explains the business cycle properties of unemployment, labor market participation and wages well. Key to the empirical success of the model are a degree of wage rigidity consistent with Haefke et al. (2008), a vacancy cost shock that is negatively correlated with productivity shocks, and a careful calibration of the cross-sectional distribution of market productivity.

This model suggest that the labor supply elasticity is between 0.3 and 0.4 for men and between 0.4 and 0.65 for women, along the extensive margin. Adding an intensive-margin elasticity of 0.54, as suggested by Chetty et al. (2011), gives an overall elasticity close to unity. This elasticity should be interpreted as capturing the supply effects that are operative at business cycle frequencies. Elasticities at longer horizons might be even higher, as stressed by Keane and Rogerson (2011).

These elasticities are slightly higher than the estimates reported in Chetty et al. (2011), but well within the range of recent results in the micro-econometric literature. They are also compatible with the macroeconomic business-cycle evidence, because the high variability of employment and unemployment is not due to a very high elasticity of labor supply, but to the large fluctuations in vacancy creation. Those reflect the changing hiring incentives of firms, not the labor supply elasticity of households.

A Numerical implementation

The model cannot be solved analytically, and even an exact numerical solution poses important problems, for the following reason. Firms, when they decide on whether to post vacancies, have to consider the cross-sectional distribution of the unemployment pool, which determines the expected quality of the workers they are going to be matched with. Even though we only consider small aggregate fluctuations, such that only the households in a neighborhood of the participation threshold are affected in their participation decision, a complicated distribution dynamics arises in this region.

In the theoretical model of Section 2, each household has constant search cost γ . Given the aggregate state, the household either participates or not. The dynamics of participation then depends precisely on the density of the unemployed at the current participation margin. This density turns out to be hard to pin down exactly in the stochastic equilibrium. Numerically it is much better to smooth out the household decision in the following way. For each household, the search cost γ is uniformly distributed within $(\bar{\gamma} - \varepsilon, \bar{\gamma} + \varepsilon)$. If ε is

very small, only households very close to the participation margin will have a participation probability in the interior of $(0, 1)$. To model the distribution, we consider a grid of n types of households in that region, and for each type of household, the fraction that is unemployed is an aggregate state variable. To achieve high precision, we need n in the range of several hundreds. We solve the model by the methods described in Reiter (2009b) and Reiter (2009a).

The theoretical model is the limit of the smoothed model with $\varepsilon \rightarrow 0$. To see whether we can trust the numerical solution, we vary both n and ε . This has no significant effect on the results.

B Data Sources

All labor market series originate from the BLS website and are seasonally adjusted. Age is 16 years and older. The participation rate series (Id LNS11300000); the unemployment rate series (Id LNS14000000); the employment population rate series (Id LNS12300000).

The remaining series are all available via the FRED database. All are seasonally adjusted. Real GDP (2000=100) can be obtained via

<http://research.stlouisfed.org/fred2/series/GDPC1/106>.

Vacancies are measured by the help-wanted index

<http://research.stlouisfed.org/fred2/series/HELPWANT/10> (1987=100);

the wage series is average hourly earnings, total private industries

<http://research.stlouisfed.org/fred2/series/ahetpi/10>.

It is deflated using the private consumption expenditure price index

<http://research.stlouisfed.org/fred2/series/pcepi/21> (Base 2000).

Labor productivity is measured as output per worker:

<http://research.stlouisfed.org/fred2/series/OPHNFB> and was deflated using the GDP deflator.

When available, all series were downloaded at a monthly frequency and converted to quarterly frequency by averaging. GDP and output/worker are at quarterly frequency.

References

- Bargain, O., K. Orsini, and A. Peichl (2011). Labor supply elasticities in Europe and the US. *IZA Discussion Paper 5820*.

- Chang, Y. and S.-B. Kim (2006). From individual to aggregate labor supply: A quantitative analysis based on a heterogeneous agent macroeconomy. *International Economic Review* 47(1), 1–27.
- Chang, Y. and S.-B. Kim (2007). Heterogeneity and aggregation: Implications for labor-market fluctuations. *American Economic Review* 97(5), 1939–1956.
- Chetty, R., A. Guren, D. Manoli, and A. Weber (2011). Are micro and macro labor supply elasticities consistent? A review of evidence on the intensive and extensive margins. *American Economic Review* 101(May), 471–75.
- Costain, J. and M. Reiter (2008). Business cycles, unemployment insurance, and the calibration of matching models. *Journal of Economic Dynamics and Control* 32, 1120–1155.
- Faraglia, E. (2003). Search intensity, US gross worker flows and the business cycle. *Mimeo, NYU*.
- Floden, M. and J. Lindé (2001). Idiosyncratic risk in the united states and sweden: Is there a role for government insurance? *Review of Economic Dynamics* 4, 406–437.
- Garibaldi, P. and E. Wasmer (2005). Equilibrium search unemployment, endogenous participation and labor market flows. *Journal of the European Economic Association* 3, 851–882.
- Haefke, C. and M. Reiter (2006). Endogenous labor market participation and the business cycle. IZA DP 2029.
- Haefke, C., M. Sonntag, and T. van Rens (2008). Wage rigidity and job creation. IZA DP 3714.
- Hagedorn, M. and I. Manovskii (2008). The cyclical behavior of equilibrium unemployment and vacancies revisited. *American Economic Review* 98(4), 1692–1706.
- Hagedorn, M. and I. Manovskii (2010). Job selection and wages over the business cycle. *mimeo*.
- Keane, M. and R. Rogerson (2011). Reconciling micro and macro labor supply elasticities. forthcoming JEL.
- Krusell, P., T. Mukoyama, R. Rogerson, and A. Sahin (2009). A three state model of worker flows in general equilibrium. *NBER Working Paper 15251*.
- Krusell, P., T. Mukoyama, R. Rogerson, and A. Sahin (2011). Labor supply, frictions, and the business cycle. *mimeo*.

- Ljungqvist, L. and T. Sargent (2011). A labor supply elasticity accord? *American Economic Review* 101, 487–491.
- Michaillat, P. (2011). Do matching frictions explain unemployment? Not in bad times. *American Economic Review* 101, Forthcoming.
- Petrongolo, B. and C. Pissarides (2001). Looking into the black box: A survey of the matching function. *Journal of Economic Literature* 39, 390–431.
- Petrosky-Nadeau, N. and E. Wasmer (2010). The cyclical volatility of labor markets under frictional financial markets. *mimeo*.
- Pissarides, C. (2009). The unemployment volatility puzzle: Is wage stickiness the answer? *Econometrica* 77, 1339–1369.
- Pries, M. and R. Rogerson (2009, July). Search frictions and labor market participation. *European Economic Review* 53(5), 568–587.
- Reiter, M. (2008). Embodied technical change and the fluctuations of wages and unemployment. *Scandinavian Journal of Economics* 109(4), 695–721.
- Reiter, M. (2009a). Approximate and almost-exact aggregation in dynamic stochastic heterogeneous-agent models. Manuscript.
- Reiter, M. (2009b). Solving heterogenous agent models by projection and perturbation. *Journal of Economic Dynamics and Control* 33(3), 649–665.
- Rogerson, R. and J. Wallenius (2009). Micro and macro elasticities in a life cycle model with taxes. *Journal of Economic Theory* 144, 2277–2293.
- Rudanko, L. (2009). Labor market dynamics under long term wage contracting. *Journal of Monetary Economics* 56, 170–183.
- Shimer, R. (2004). Search intensity. *Manuscript*.
- Shimer, R. (2005). The cyclical behavior of equilibrium unemployment and vacancies. *American Economic Review* 95, 25–49.
- Shimer, R. (2007). Reassessing the ins and outs of unemployment. NBER Working Papers 13421, National Bureau of Economic Research, Inc.
- Shimer, R. (2011). Job search, labor force participation, and wage rigidities. Manuscript.
- Tripier, F. (2002). Can the labor market search model explain the fluctuations of allocations of time? *Economic Modelling* 21, 131–146.
- van Soest, A., M. Das, and X. Gong (2002). A structural labour supply model with flexible preferences. *Journal of Econometrics* 107, 345–374.

Veracierto, M. (2008). On the cyclical behavior of employment, unemployment and labor force participation. *Journal of Monetary Economics* 55, 1143–1157.

C Tables and Figures

Table 3: Micro vs. Macro Labor Supply Elasticities

		Intensive Margin	Extensive Margin	Aggregate Hours
Steady State	micro	0.33	0.26	0.59
(Hicksian)	macro	0.33	0.17	0.50
Intertemporal	micro	0.54	0.28	0.82
Substitution (Frisch)	macro	[0.54]	[2.30]	2.84

Note: reproduced from Chetty et al. (2011, Table 1).

Estimates in brackets show the values implied by the macro aggregate hours elasticity if the intensive Frisch elasticity is chosen to match the micro estimate of 0.54.

Table 4: Volatility of wages at business cycle frequencies

		BP filter		HP filter	
		Relative	Auto	Relative	Auto
		std. dev.	correl.	std. dev.	correl.
Aggregate wage	1951-2001	0.41	0.92	0.43	0.91
	1984-2006	0.85	0.92	0.84	0.93
CPS, all workers	1984-2006	0.44	0.91	0.67	0.92
CPS, new hires	1984-2006	0.68	0.80	1.09	0.71

Note: reproduced from Haefke et al. (2008, Table 2).

Table 5: Business Cycle Statistics for Agegroup 16–99, Years 1979 – 2007.

	Mean	$\frac{\sigma_x}{\sigma_{\text{GDP}}}$	ρ_{-1}	$\rho_{x,\text{GDP}}$
Participation	65.83	0.23	0.89	0.57
Partic. Men	75.44	0.19	0.78	0.57
Partic. Women	57.07	0.35	0.90	0.45
Employment	61.81	0.69	0.98	0.85
Unempl.Rate	6.13	7.42	0.98	-0.85
Monthly Job Finding Prob.	42.51	5.25	0.94	0.81
Monthly Sep. Prob	3.49	2.31	0.59	-0.43
Avg. Hourly Earnings	14.81	0.74	0.99	0.44

Note: $\frac{\sigma_x}{\sigma_{\text{GDP}}}$: standard deviation of x , rel. to GDP

ρ_{-1} : first order autocorrelation

$\rho_{x,\text{GDP}}$: correlation with GDP.

Table 6: Model Results, Baseline Model

Parameters	η_Z^{PM} : 0.403;	η_Z^{PF} : 0.646;	h_y : 0.000			
	σ_Z : 0.0012;	σ_κ : 0.0258;	σ_h : 0.0010			
	$\rho_{Z,\kappa}$: -0.181	Δ_α : 2.275				
Moments	Mean	$\frac{\sigma_x}{\sigma_{\text{GDP}}}$	ρ_{-1}	$\rho_{x,\text{GDP}}$	$\rho_{x,\text{GDP}}^m$	$\rho_{x,\text{GDP}}^w$
Partic. Men	75.4	0.18	0.98	0.69	0.67	0.65
Partic. Women	57.1	0.29	0.98	0.69	0.67	0.65
Employment	61.8	0.69	0.97	0.91	0.89	0.87
Unempl.Rate	6.1	8.42	0.97	-0.85	-0.84	-0.82
JobFindProbWeekly	12.9	9.53	0.90	0.86	0.84	0.82
AveWage	-	0.79	0.91	0.89	0.87	0.85
GDP	-	1.00	0.94	1.00	1.00	1.00
Labor Productivity	-	0.56	0.90	0.87	0.87	0.86

Note: η_Z^{PM}, η_Z^{PF} : labor supply elasticity male and female, respectively;

$\frac{\sigma_x}{\sigma_{\text{GDP}}}$: standard deviation of x , rel. to GDP; ρ_{-1} : first order autocorrelation;

$\rho_{x,\text{GDP}}, \rho_{x,\text{GDP}}^m, \rho_{x,\text{GDP}}^w$: correlation with GDP at quarterly, monthly and weekly frequency, respectively.

Table 7: Model Results, One-Shock Calibration

Parameters	η_Z^{PM} : 0.403;	η_Z^{PF} : 0.646;	h_y : 0.000			
	σ_Z : 0.0019;	σ_κ : 0.0000;	σ_h : 0.0000			
	$\rho_{Z,\kappa}$: 0.000	Δ_α : 0.000				
Moments	Mean	$\frac{\sigma_x}{\sigma_{GDP}}$	ρ_{-1}	$\rho_{x,GDP}$	$\rho_{x,GDP}^m$	$\rho_{x,GDP}^w$
Partic. Men	75.4	0.22	0.98	0.90	0.86	0.82
Partic. Women	57.1	0.35	0.98	0.90	0.86	0.82
Employment	61.8	0.30	0.99	0.84	0.76	0.62
Unempl.Rate	6.1	0.63	0.79	0.01	0.23	0.52
JobFindProbWeekly	12.9	0.42	0.99	0.87	0.78	0.58
AveWage	-	0.88	0.91	1.00	1.00	1.00
GDP	-	1.00	0.91	1.00	1.00	1.00
Labor Productivity	-	0.92	0.90	1.00	1.00	1.00

Note: cf. Table 6.

Table 8: Model Results, Rigid Wages

Parameters	η_Z^{PM} : 0.577;	η_Z^{PF} : 1.026;	h_y : 0.357			
	σ_Z : 0.0010;	σ_κ : 0.0000;	σ_h : 0.0000			
	$\rho_{Z,\kappa}$: 0.000	Δ_α : 6.847				
Moments	Mean	$\frac{\sigma_x}{\sigma_{GDP}}$	ρ_{-1}	$\rho_{x,GDP}$	$\rho_{x,GDP}^m$	$\rho_{x,GDP}^w$
Partic. Men	75.4	0.19	0.98	0.91	0.89	0.87
Partic. Women	57.1	0.28	0.98	0.92	0.90	0.88
Employment	61.8	0.69	0.97	0.98	0.97	0.95
Unempl.Rate	6.1	7.20	0.97	-0.99	-0.98	-0.96
JobFindProbWeekly	12.9	8.28	0.90	0.99	0.99	0.98
AveWage	-	0.11	0.95	1.00	0.99	0.99
GDP	-	1.00	0.94	1.00	1.00	1.00
Labor Productivity	-	0.47	0.90	0.99	0.98	0.97


Note: cf. Table 6.

Table 9: Model Results for Varying Values of h_y

h_y	0.236	0.000	-0.100	-0.150	-0.250	-0.500	-0.600
σ_h	0.0000	0.0010	0.0014	0.0015	0.0018	0.0024	0.0027
σ_Z	0.0012	0.0012	0.0012	0.0012	0.0012	0.0012	0.0012
σ_κ	0.0236	0.0258	0.0267	0.0271	0.0278	0.0291	0.0296
$\rho_{Z,\kappa}$	-0.182	-0.181	-0.176	-0.174	-0.170	-0.164	-0.161
η_Z^{PM}	0.515	0.403	0.357	0.334	0.289	0.172	0.121
η_Z^{PF}	0.866	0.646	0.577	0.546	0.491	0.381	0.345
$\rho_{PM,GDP}$	0.89	0.69	0.61	0.57	0.50	0.38	0.34
$\rho_{PF,GDP}$	0.89	0.69	0.61	0.58	0.52	0.42	0.38
σ_W/σ_{GDP}	0.63	0.79	0.87	0.92	1.01	1.27	1.38
$\rho_{W,GDP}$	0.94	0.89	0.86	0.85	0.83	0.79	0.78

Note: η_Z^{PM} , η_Z^{PF} : labor supply elasticity male and female, respectively;

$\frac{\sigma_x}{\sigma_{GDP}}$: standard deviation of x , rel. to GDP; $\rho_{x,GDP}$, correlation with GDP.



Note: The impulse responses of the unemployment rate and the job finding rate are scaled down by a factor of 10.

Figure 1: Impulse responses