

Aggregate Implications of Individual Labor Supply Heterogeneity *

José Mustre-del-Río

University of Rochester

First Version: September 30, 2010

This Version: November 12, 2010

Abstract

The aggregate labor supply elasticity plays a crucial role in understanding employment fluctuations, and the effect of taxes and government spending. This paper examines the Frisch elasticity at the extensive margin of labor supply in an economy consistent with the observed dispersion in average employment rates across individuals. A heterogeneous agent economy with indivisible labor is presented where agents differ in their disutility of labor and market skills. To impose quantitative discipline on the model, its key parameters are estimated via indirect inference. The elasticity of aggregate employment in the model is 0.4. A simple decomposition reveals that labor disutility differences, which capture the dispersion in average employment rates, are crucial for obtaining this quantitative result. In a version of the model with only skill differences across agents the elasticity is 1.2. In a version that only allows for labor disutility differences the elasticity is 0.5. These results suggest that the previous literature generates large aggregate labor supply elasticities by ignoring labor supply differences as captured by the dispersion in average employment rates across individuals.

*I would like to thank my advisor Mark Bilal for his endless advice, comments and support on this project. I would also like to thank Mark Aguiar, William Hawkins and Ronni Pavan for their useful comments and discussions along with participants at the Rochester Lunch Seminar. Finally, I am indebted to Yongsung Chang for his help and for providing useful computer code. All errors are my own.

1 Introduction

The labor supply elasticity plays a crucial role in understanding employment fluctuations over the business cycle and in evaluating the effect of taxes and government spending. Early business cycle models, (e.g. Lucas and Rapping, 1969), require a representative agent to have a large intertemporal substitution of leisure to be consistent with the observed movements in hours and wages. Similarly, Prescott (2004) postulates a large aggregate elasticity of labor supply when determining the effect of marginal labor tax rates on labor supply across countries and time. Meanwhile, estimates based on labor supply decisions over the life-cycle find elasticities that are positive but economically small.¹ More recently, work by Chang and Kim (2007); Rogerson and Wallenius (2009); Gourio and Noual (2006), and Erosa et al. (2010) argues that one can generate a large macro elasticity in spite of assuming a small elasticity at the micro level. In these papers, the large employment response to a wage change is crucially determined by differences across workers in the surplus that employment generates relative to non-employment; i.e. how different are their reservation wages in comparison to the market wage.

This paper measures the Frisch elasticity at the extensive margin of labor supply when individuals are ex-ante different in labor supply and skills, and hence heterogeneous in the surplus that employment generates for them. Motivated by observations from data on individuals (National Longitudinal Survey of the Youth–NLSY) that show large differences in average employment rates that do not project on wages, I develop a model that is consistent with these facts. The model is a heterogeneous agent economy with incomplete markets and indivisible labor supply, with two novel features. Firstly, agents differ in their disutility of labor and secondly, they differ in their market skills.

¹See for example Ghez and Becker (1975), Macurdy (1981), Altonji (1986), or Abowd and Card (1989).

To impose quantitative discipline on the model, I estimate its key parameters using data from the NLSY via indirect inference. Due to the NLSY's structure, I am able to estimate the model on quarterly data. In doing so, I circumvent the time-aggregation issues that arise when using annual surveys, such as the Panel Study of Income Dynamics (PSID), as discussed in Erosa et al. (2010).

The result of the paper is that once agents display a realistic amount of ex-ante heterogeneity in labor supply, as well as wages, a very large macro-level elasticity is no longer obtained through the extensive margin of labor supply. The implied aggregate labor supply elasticity of my model is 0.4. Further inspection of my model reveals that labor disutility differences across agents are essential in generating the low labor supply elasticity. In a version of my model with only ex-ante skill differences (in the spirit of Erosa et al., 2010), the implied elasticity is 1.2. Meanwhile, in a version of my model with only ex-ante labor disutility differences, the implied elasticity is 0.46. This version however, generates a counterfactual wealth effect on participation. Similar to Chang and Kim (2007), in this version of my model the wealthiest do not participate in the labor market as much as in the data. Once labor disutility and skills are both incorporated, the elasticity drops to 0.4. This model also generates a more realistic wealth effect on participation.

This paper is a natural direction in the literature that extends the neoclassical growth model to account for choices at the extensive margin of labor supply. In a representative agent model with indivisible labor, Hansen (1985); Rogerson (1988) are the first to show that individual and aggregate labor supply elasticities are unrelated. Because of the representative agent assumption, there is no heterogeneity in the value of non-market time and hence the Frisch elasticity at the extensive margin is infinite. Cho (1995) relaxes the representative agent assumption and allows for ex-post heterogeneity across agents in their market productivity. However, he maintains the complete markets assumption.

Chang and Kim (2006) go a step further and relax both the representative agent assumption and complete markets assumptions. Their model features ex-post heterogeneity as in Cho (1995), but no consumption insurance across agents. In their model, as in mine, the slope of the aggregate labor supply schedule is determined by the distribution of reservation wages. Their model does not allow for any ex-ante heterogeneity across agents and implies an elasticity at the extensive margin around 1.

Rogerson and Wallenius (2009), and Erosa et al. (2010) adopt a different approach and examine life-cycle models. In these models a non-linear mapping between hours of work and earnings plays a crucial role in providing the disconnect between micro and macro elasticities of labor supply. Both models allow for intensive and extensive margin adjustments. The work by Erosa et al. (2010) differs from Rogerson and Wallenius (2009) by allowing for incomplete markets and heterogeneous agents. In addition, Erosa et al. (2010) allow for ex-ante differences in skills across agents. Their model yields an aggregate labor supply elasticity of 1.27. This elasticity however, reflects both margins. They argue that the extensive margin explains between 54 to 60 percent of the aggregate labor supply response to a temporary wage change in their model. Hence, their implied Frisch elasticity at the extensive margin is between 0.69 and 0.76.

While these contributions allow for some heterogeneity across workers, the key dimension of heterogeneity they lack is in the value of non-market time. This dimension of heterogeneity matters greatly for the Frisch elasticity at the extensive margin and is crucial to capture the average employment rate differences observed in the data. In the NLSY, most individuals are typically employed and therefore display high average employment rates. Meanwhile, others are employed less frequently and display relatively low average employment rates. This suggests fewer individuals are located at the margin, than what is implied in the work of Chang and Kim (2007) or Erosa et al. (2010).

In a model without ex-ante differences in labor supply, all individuals are, on average, employed at the same frequency and thus display similar employment rates. In equilibrium, since everyone's willingness to work is roughly the same, the reservation wage distribution is dense around the market wage. Thus, for a small change in the wage rate there is a large aggregate labor supply response simply through individual extensive margin adjustments.

Conversely, in my model with ex-ante differences in labor supply, individuals differ in their average employment rates. Because of these labor supply differences, the reservation wage distribution implied by the estimated model is dispersed in a neighborhood around the equilibrium wage rate. As a result, for a small change in the wage rate there is a small aggregate labor supply response as few individuals change their employment decision due to the location of their reservation wage relative to the equilibrium wage.

Explaining these labor supply differences is beyond the scope of this paper. However, work by Bils et al. (2009) suggests that these differences may be driven by comparative advantage in market production relative to non-market production. In their model, individuals with low labor supply are those with high comparative advantage in non-market activity.

This paper proceeds as follows. The next section describes the model. Section 3 presents the NLSY sample used for the empirical analysis. Section 4 discusses the estimation procedure. Section 5 presents the results of the estimation procedure along with a discussion of the model's fit to the data. Section 6 presents the implied Frisch elasticity of the estimated model and the decomposition of this elasticity. Finally, section 7 concludes.

2 Model

The model economy is a heterogeneous agent model with incomplete markets and indivisible labor supply similar to the one considered by Chang and Kim (2007). Unlike their work, agents are both ex-ante and ex-post heterogeneous. As in Erosa et al. (2010), individuals are ex-ante heterogeneous in skills. In addition, agents are also ex-ante heterogeneous in labor disutility, which is the key distinguishing feature of this model from the rest of the literature.² These two new dimensions of heterogeneity allow the model to account for differences across workers in average employment rates and wages. As in Aiyagari (1994), individuals are ex-post different in wealth and labor productivity. The analysis is confined to a steady-state with no aggregate uncertainty.

2.1 Workers

The economy is populated by a continuum (measure one) of workers. Workers differ in terms of their time invariant disutility of labor $d_j \in \{d_1, d_2, \dots, d_M\}$ and market skills $s_i \in \{s_1, s_2, \dots, s_N\}$. They also differ in their idiosyncratic productivity x that evolves exogenously according to the stochastic process with transition probability function $\pi_x(x'|x) = Pr(x_{t+1} \leq x' | x_t = x)$. Workers have preferences over consumption c_t given by $\ln(c_t)$.

Workers can trade claims for physical capital a_t , which yields a rate of return r . Physical capital is the only asset available to workers (markets are incomplete) and they face a borrowing constraint $a_t \geq \bar{a}$ for all t as in Aiyagari (1994). Labor

²The two dimensions of heterogeneity across agents could alternatively be interpreted as market and non-market skills (as in Bils et al. 2009), leaving agents with a choice between working in the market or working at home. Because the data used to estimate the model's parameters has no information on non-market activities, it is not possible to distinguish between somebody valuing leisure more and working less in the market versus being more productive at home and working less in the market. Moreover, even if the data did include information on non-market work, the individual labor assumption in the model precludes the marginal decision between an hour of work in the market versus at home. Thus for expositional simplicity, the former interpretation is maintained.

supply is indivisible as in Hansen (1985); Rogerson (1988). When employed, a worker with skills s_i must supply \bar{h} units of labor and earns $w_t x_t s_i \bar{h}$, where w_t is the market wage rate per unit of effective labor $x_t s_i$.

The value function of an employed worker with market skills s_i , disutility of labor d_j , assets a , and idiosyncratic productivity x is:

$$V_{ij}^E(a, x) = \max_{a'} \ln(c) - d_j + \beta \mathbf{E}[\max\{V_{ij}^E(a', x'), V_{ij}^{NE}(a', x')\} | x] \quad (1)$$

subject to

$$\begin{aligned} c &= w x s_i \bar{h} + (1 + r)a - a' \\ a' &\geq -\bar{a}. \end{aligned}$$

A worker takes the wage w and the interest rate r as given. Meanwhile, the value function of a non-employed worker is defined as:

$$V_{ij}^{NE}(a, x) = \max_{a'} \ln(c) + \beta \mathbf{E}[\max\{V_{ij}^E(a', x'), V_{ij}^{NE}(a', x')\} | x] \quad (2)$$

subject to

$$\begin{aligned} c &= (1 + r)a - a' \\ a' &\geq -\bar{a} \end{aligned}$$

Finally, the labor supply decision of an individual with market skills s_i , disutility of labor d_j , assets a and idiosyncratic productivity x is thus characterized

by:

$$V_{ij}(a, x) = \max_{h \in \{0, \bar{h}\}} \{V_{ij}^E(a, x), V_{ij}^{NE}(a, x)\}. \quad (3)$$

Note that the reservation productivity $x_{ij}^*(a)$, the value of x such that the worker is indifferent between working and not working, is an increasing function of asset holdings a and labor disutility d , but decreasing in market skills s . Because workers face the same stochastic process for x , differences in labor disutility will lead to systematic differences in the frequency of employment across workers, as low d workers will have a wider range of acceptable x 's and thus will be employed more often, relative to high d workers. Conditional on being employed at the same productivity level, high skill workers will also systematically earn higher wages relative to low skill workers, through the scaling effect of s_i on effective wages wxs_i . It is thus through these two channels that the model will generate differences both in average employment and wages across workers. Meanwhile, the cross-sectional correlation between market skills and disutility of labor implicitly generates a cross-sectional correlation between average wages and employment.

The model abstracts from the intensive margin choice of labor supply and focuses on the Frisch elasticity at the extensive margin for several reasons. First, workers are rarely allowed to choose completely flexible work schedules or to supply a small number of hours. Second, a large fraction of hours fluctuations are accounted for by movements in and out of employment by workers (see for example Coleman, 1984; Heckman, 1984). Finally, Kimmel and Kniesner (1998) find that employment fluctuations account for three-fourths of wage-induced variation in labor hours. Erosa et al. (2010) document that the extensive margin accounts for about 54 percent of the aggregate labor supply response to the temporary wage change in their model.

Unlike Rogerson and Wallenius (2009), and Erosa et al. (2010), this model departs from the life-cycle. Instead, it follows the tradition of infinite horizon indivisible labor economies pioneered by Hansen (1985); Rogerson (1988) and continued by Chang and Kim (2006). As documented by Erosa et al. (2010), there is a pronounced life-cycle pattern in the labor supply behavior of men. By abstracting from life-cycle, this model misses the potentially elastic participation decisions of young individuals and those near retirement. This in principle, may lead my model to underestimate the true elasticity at the extensive margin. However, as shown in the next section, aggregate time-series of the extensive margin of employment for prime-age workers (ages 25-54) and all workers display very similar patterns. This suggests that the elastic employment responses of the young and old may not be important for characterizing the overall employment response.

2.2 Firms

There is a representative firm that takes capital K and effective units of labor L as inputs, and produces output Y according to a constant returns-to-scale Cobb-Douglas technology:

$$Y = F(K, L) = K^\alpha L^{1-\alpha} \quad (4)$$

Capital depreciates at a constant rate δ , while effective units of labor are measured as

$$L = \sum_{i=1}^N \sum_{j=1}^M \int h_{ij} x s_i d\mu_{ij}. \quad (5)$$

Here h_{ij} is the labor supply decision of a worker of type $s = s_i, d = d_j$; and $\mu_{ij} = \mu_{ij}(a, x)$ is the distribution of workers of type $s = s_i, d = d_j$ with assets a

and idiosyncratic productivity x . It is such that $\int d\mu_{ij} = p_{ij}$ and $\sum_{ij} p_{ij} = 1$, where p_{ij} denotes the proportion of workers with skills s_i and disutility of labor d_j .

2.3 Equilibrium

A steady-state equilibrium consists of a set of value functions $\{V_{ij}^E(a, x), V_{ij}^{NE}(a, x), V_{ij}(a, x)\}_{i=1, j=1}^{N, M}$, decision rules for consumption, asset holdings and labor supply, $\{c_{ij}(a, x), a'_{ij}(a, x), h_{ij}(a, x)\}_{i=1, j=1}^{N, M}$; aggregate inputs, K, L and factor prices w, r such that:

1. Individuals optimize: Given prices w and r , the individual decision rules $\{c_{ij}(a, x), a'_{ij}(a, x), h_{ij}(a, x)\}_{i=1, j=1}^{N, M}$ solve $\{V_{ij}^E(a, x), V_{ij}^{NE}(a, x), V_{ij}(a, x)\}_{i=1, j=1}^{N, M}$.
2. The representative firm maximizes profits:
 - $w = F_2(K, L)$
 - $r = F_1(K, L) - \delta$
3. The good market clears:

$$\sum_{i=1}^N \sum_{j=1}^M \int \{a'_{ij}(a, x) + c_{ij}(a, x)\} d\mu_{ij} = F(K, L) + (1 - \delta)K$$

4. Factor markets clear:

$$L = \sum_{i=1}^N \sum_{j=1}^M \int h_{ij}(a, x) x s_i d\mu_{ij}$$

$$K = \sum_{i=1}^N \sum_{j=1}^M \int a d\mu_{ij}$$

5. Individual and aggregate behaviors are consistent: For all $A^0 \subset \mathcal{A}$ and

$X^0 \subset \mathcal{X}$ and each i, j ,³

$$\mu_{ij}(A^0, X^0) = \int_{A^0, X^0} \left\{ \int_{\mathcal{A}, \mathcal{X}} \mathbf{1}_{a'=a'_{ij}(a,x)} d\pi_x(x'|x) d\mu_{ij} \right\} da' dx'$$

3 Data

The data used comes from the National Longitudinal Survey of Youth 1979 (NLSY79), survey years 1990 through 2000. The NLSY79 is a nationally representative sample of 12,686 young men and women who were 14-22 years old when first interviewed in 1979. Interviews were conducted annually through 1994 and biennially thereafter. Participants are asked questions regarding their family background, education, and work experience. Since average hourly wages and employment rates are the primary focus of this study, the NLSY is used as it consistently tracks workers' employment histories over several years. While individuals are not interviewed on a quarterly basis, it is possible to convert the data to a quarterly frequency as individuals are asked information both on jobs currently held and held since the last interview including calendar dates on when each of the jobs started and finished.

By using quarterly information on employment and wages, I am able to circumvent the bias introduced by time aggregation when using lower frequency data such as the PSID. This point is mentioned in Erosa et al. (2010), who argue that the wage rate obtained in the PSID as the ratio of annual earnings to annual hours is a noisy measure of the true returns to work faced by an individual during the year. This is because temporary low wage shocks will be unobserved in annual data if the individual chooses not to work during that portion of the year. However, their calibration strategy relies on aggregating quarterly model-generated data to match annual PSID observations. On the

³Let \mathcal{A} and \mathcal{X} denote the sets of all possible realizations of a and x , respectively.

contrary, my procedure is more consistent as it relies on comparing quarterly model-generated data with quarterly observations from the NLSY.

The drawback of using the NLSY is that respondents are fairly young when first interviewed in 1979. However, by the 1990 survey wave the youngest age that I observe individuals is 26. Conversely, the oldest age I observe respondents is 48. To gauge how representative this age group is for studying the responsiveness of labor force participation I compare the employment to population ratio for all workers and for workers ages 25-54. As can be seen from figure 1, the age group I consider has an overall higher participation rate. This difference is driven mostly by the very low participation rate of workers near or in retirement. If I project the HP filtered employment to population series of workers ages 25-54 on the HP filtered series of employment to population for all workers, the resulting coefficient is 1.04. This simple exercise suggests that workers in the age group I consider and workers from all age groups display very similar time-series movements at the extensive margin of employment.

I restrict my sample to the cross-sectional subsamples. Individuals must not be in the armed forces and not be attending school. In addition, over the 11 year period considered, individuals must have at least 22 quarters where employment status can be determined (either employed or non-employed).⁴ When employed, individuals must have data on both hours and wages earned. I ignore wage reports below \$1.00 in 1983 dollars, jobs where the individual works less than 30 hours a week and I censor hourly wage rates above \$500. Finally, I restrict my sample to individuals with positive average employment rates over the 11 year period considered. From the perspective of the model, individuals who never work can be either be extremely low skill workers, have a very high disutility of

⁴For those individuals with missing observations, they must have at least 22 valid quarters before the first missing observation as I ignore valid observations after the first missing observation. This is done for simplicity as adding valid observations after the first missing observation only increases my sample size by 3%. Moreover, the model simulation becomes more complicated as simulated data must replicate the observed frequency of valid observations after the first missing observation.

labor, or both. Because data is unavailable to distinguish between these three possibilities, they are excluded from the estimation procedure. Since these individuals are not marginal, in the sense of being near the margin between choosing employment or non-employment, their exclusion should not affect the implied elasticity at the extensive margin. The resulting sample consists of 220,199 observations from 5,082 individuals. Summary statistics appear in table 1.



Figure 1: Employment to population ratio, all workers (solid) and workers ages 25-54 (dash) from CPS.

4 Model Parametrization and Estimation

4.1 Parametrization

In this section I describe how the model is parametrized and the procedure used to estimate its key structural parameters. Details of how the steady-state equilibrium is computed appear in the appendix. To start, I define the unit

Table 1: Summary Statistics NLSY panel 1990-2000.

Variable	Mean	Standard Deviation
Long run employment rate	0.746	0.283
Long run log wage	1.994	0.518
Employment duration (in quarters)	15.152	11.833
Pr(E→N)	0.031	0.175
Pr(N→E)	0.122	0.328
Age	34.169	3.76
Male	0.491	0.500
White	0.808	0.394
Highest Grade Completed	13.52	2.54

Notes: Wages are in 1983 dollars. Cross-sectional correlation between long run employment rates and wages equals 0.397.

of time as one quarter. Individual productivity x follows an $AR(1)$ process: $\ln x' = (1 - \rho_x)\mu_x + \rho_x \ln x + \epsilon_x$, where μ_x is the unconditional mean of the process and $\epsilon_x \sim N(0, \sigma_x^2)$. As in Chang and Kim (2007), an employed individual spends one-third of discretionary time working, so $\bar{h} = \frac{1}{3}$. The capital-income share α is set to 0.36 while the depreciation rate δ is set to 2.5 percent. The discount factor β is chosen so that in equilibrium the quarterly rate of return on capital is 1 percent.

I assume that market skills and labor disutility can take on three values $\{s_1, s_2, s_3\}$ and $\{d_1, d_2, d_3\}$, yielding a total of 9 distinct worker types and hence 9 proportions p_{ij} to be determined.⁵ By normalization, I set the highest skill level s_1 to 1, while p_{33} is set so that $\sum_{ij} p_{ij} = 1$. Under these assumptions, there are a total of 16 structural parameters that must be estimated: $\Psi' = (s_2, s_3, d_1, d_2, d_3, p_{11}, p_{12}, p_{13}, p_{21}, p_{22}, p_{23}, p_{31}, p_{32}, \rho_x, \sigma_x, \mu_x)$. I turn to the procedure used to estimate these 16 parameters next.

⁵The choices for the number of skills and labor disutilities is primarily driven by computational concerns as adding more worker types increases the state-space and thus computational time significantly. Moreover, as will be seen in the next section, these modeling choices seem to fit the data well.

4.2 Estimation via Indirect Inference

Given the complicated structure of the model, rather than attempting to directly estimate the $k = 16$ structural parameters Ψ , I will instead estimate them using indirect inference.⁶ Indirect inference involves the use of an “auxiliary” statistical model that serves as a criterion to determine if actual data and model-generated data (given Ψ) are “close enough” in a sense that is formally defined below. I define the indirect inference estimator of Ψ , as the estimated value $\hat{\Psi}$ that is found when the estimated parameters of the auxiliary model obtained when using actual data and the estimated parameters of the auxiliary model obtained when using model-simulated data are close enough.

More formally, suppose that the observed data can be written as $\{y_{it}\}, i = 1, \dots, N; t = 1, \dots, T$, while data generated from the model can be written as $\{\tilde{y}_{it}(\Psi)\}, i = 1, \dots, N; t = 1, \dots, T$.

Next, suppose the auxiliary model is characterized by a vector of parameters Γ (of dimension $p \geq k$) that can be estimated using observed data as:

$$\hat{\Gamma} = \arg \max_{\Gamma} \mathcal{L}(y; \Gamma), \quad (6)$$

where $\mathcal{L}(y; \Gamma)$, is the likelihood function associated with the auxiliary model.

Meanwhile, the model can be simulated to generate M statistically independent data sets $\{\tilde{y}_{it}^m(\Psi)\}, m = 1, \dots, M$. As in the case with observed data, the auxiliary model can be estimated using each of the simulated data sets to obtain M estimated parameter vectors $\tilde{\Gamma}_m(\Psi)$, as:

$$\tilde{\Gamma}_m(\Psi) = \arg \max_{\Gamma} \mathcal{L}(y^m(\Psi); \Gamma), \quad (7)$$

⁶This method was first introduced by Smith (1990,1993) and extended by Gourieroux, Monfort, and Renault (1993) and Gallant and Tauchen (1996).

Finally, define the average of the estimated parameter vectors by $\tilde{\Gamma}(\Psi) = M^{-1} \sum_{m=1}^M \tilde{\Gamma}_m(\Psi)$. The criterion used to determine if the observed data and simulated data are “close enough” through the lens of the auxiliary model is the Wald approach to indirect inference that chooses Ψ to minimize the quadratic form in the vector $\hat{\Gamma} - \tilde{\Gamma}(\Psi)$:

$$\hat{\Psi}^{Wald} = \arg \min_{\Psi} (\hat{\Gamma} - \tilde{\Gamma}(\Psi))' W (\hat{\Gamma} - \tilde{\Gamma}(\Psi)) \quad (8)$$

where W is a positive definite “weighting” matrix.^{7,8}

Notice that accommodating sample restrictions and missing observations when estimating Ψ via indirect inference is straight forward. I simply apply the same sample restrictions and assumptions on missing observations across actual and simulated data sets. In the present context, each simulated data set consists of $I = 5082$ individuals contributing at most 44 quarters of data, as in the panel constructed from the NLSY. Because some individuals have fewer quarterly observations than others (or are never observed to be employed), I simply omit quarter observations in the simulated data so that the distribution of “quarter-counts” by individual in model-generated data is the same as in actual data.

⁷For the purposes of this paper set to the identity matrix \mathbf{I}_p ; more generally, the optimal weighting matrix is the inverse of the covariance matrix of the parameter vector $\hat{\Gamma}$ using observed data. Note that setting $W = \mathbf{I}_p$ only affects the efficiency of the estimated $\hat{\Psi}$, but not its consistency.

⁸In practice, I use a Nelder-Meade simplex algorithm in minimizing (8), as implemented in Press et al. (1992), and set $M=20$. As highlighted by Smith (2008), the usage of simulations inflates asymptotic standard errors by a factor of $(1 + M^{-1})^{1/2}$, and thus for $M \geq 10$, this factor is negligible.

4.3 The Auxiliary model

My choice for the auxiliary model is driven by two considerations: efficiency and computational complexity. From the perspective of efficiency it is important that the auxiliary model be flexible enough to provide a good description of the data. As stressed by Keane and Smith (2003), if the auxiliary model is correctly specified (in the sense that it provides a correct statistical description of the observed data), then the Wald approach to indirect inference is asymptotically equivalent to maximum likelihood, provided that M is sufficiently large. From the perspective of computational complexity, the auxiliary model should be one that can be estimated quickly as its parameters must be estimated M times for each choice of the structural parameters Ψ . Guided by these two considerations and following the related literature (Keane and Smith, 2003; Altonji et al. 2009), I use an auxiliary model that consists of a system of seemingly unrelated regressions (SUR) with 3 equations and 6 covariates common to both equations.⁹ The system is defined as:

$$E_{it} \cdot E_{it-1} = \gamma_0^E + \gamma_{ED}^E \ln(ED_{it-1} + 1) + \gamma_{ND}^E \ln(ND_{it-1} + 1) + \gamma_w^E w_{it-1}^* + \gamma_{\bar{e}}^E e_i + \gamma_{\bar{w}}^E w_i + \epsilon_{it}^E$$

$$E_{it} \cdot (1 - E_{it-1}) = \gamma_0^N + \gamma_{ED}^N \ln(ED_{it-1} + 1) + \gamma_{ND}^N \ln(ND_{it-1} + 1) + \gamma_w^N w_{it-1}^* + \gamma_{\bar{e}}^N e_i + \gamma_{\bar{w}}^N w_i + \epsilon_{it}^N$$

$$w_{it}^* = \gamma_0^w + \gamma_{ED}^w \ln(ED_{it-1} + 1) + \gamma_{ND}^w \ln(ND_{it-1} + 1) + \gamma_w^w w_{it-1}^* + \gamma_{\bar{e}}^w e_i + \gamma_{\bar{w}}^w w_i + \epsilon_{it}^w$$

or more compactly:

$$Y_{it} = Z_{it}\Gamma + \epsilon_{it} \tag{9}$$

⁹The choice of restricting the covariates to be the same in both equations is driven by computational simplicity as the SUR system can be estimated via equation-by-equation OLS.

where $\epsilon_{it} \sim N(0, \Sigma)$ and iid over i and t . The variable E_{it} denotes individual i 's employment status (1 or 0) in period t ; ED_{it-1} denotes the number of periods individual i has been continuously employed up to time $t-1$; ND_{it-1} denotes the number of periods individual i has been continuously non-employed up to time $t-1$; e_i is the individual's average employment rate; and w_i is the individual's average log hourly wage rate (conditional on being employed).¹⁰ The variable w_{it}^* represents the individual's log wage that is equal to 0 when non-employed and equals the observed wage otherwise.¹¹

The auxiliary model (9) is a variant of the auxiliary model used in Altonji et al. (2009). This is a natural starting point as the model they ultimately estimate via generalized indirect inference can be interpreted as a reduced-form version of my structural model. Unlike the previous literature, my auxiliary model includes terms that explicitly capture permanent differences across agents as embodied by their average employment rates and wages. By this dimension, the closest work is Guvenen and Smith (2010) who use average income as an explanatory variable in their auxiliary model that is then used to estimate a consumption-savings model.

Note that the system described in (9) consists of 24 parameters: 18 coefficients from the three equations and 6 unique elements in the covariance matrix Σ . Given that the identification of the two dimensions of heterogeneity (labor disutility and skills) precisely comes from cross-sectional variation in average employment and wages, it seems valuable for the purposes of calculating the Frisch elasticity at the extensive margin to discipline the estimation of model parameters by having model-generated data imitate these two distributions and their correlation. To this end, I also estimate from observed and simulated data the

¹⁰Both ED_{it} and ND_{it} are determined recursively as $ED_{it} = E_{it}(ED_{it-1} + 1)$ and $ND_{it} = (1 - E_{it})(ND_{it-1} + 1)$, respectively.

¹¹Setting the wage equal to 0 when non-employed is valid so long as this assumption is maintained both in the actual and simulated data. Alternatively, one could set the wage of non-employed workers to the sample mean, as suggested in Keane and Smith (2003).

means' (μ_e, μ_w) , standard deviations' (σ_e, σ_w) , and skewness' $(Skew_e, Skew_w)$ of the distributions of average employment rates and average wages along with their cross-sectional correlation ρ_{ew} . These additional parameters yield a total of 31 auxiliary parameters that are used to *indirectly infer* the 16 elements of Ψ . To this end, I define $\Pi' = (\Gamma, \mu_e, \mu_w, \sigma_e, \sigma_w, Skew_e, Skew_w, \rho_{ew})$ as the vector of parameters of the augmented auxiliary model and redefine the quadratic form to be minimized as:

$$\hat{\Psi}^{Wald} = \arg \min_{\Psi} (\hat{\Pi} - \tilde{\Pi}(\Psi))' W (\hat{\Pi} - \tilde{\Pi}(\Psi)) \quad (10)$$

5 Estimation Results

This section presents the estimation results. The estimated parameters of the model are presented first. Next, I discuss the goodness of fit of the model by showing how well it performs in replicated the motivating cross-sectional facts regarding average employment rates and wages.

5.1 Estimated Model Parameters

Table 5.1 presents the estimated values for Ψ , the vector of structural parameters of the model. Given that the highest skill level was normalized to 1, these estimates imply that in the model the lowest skill type is nearly 70 percent less productive in the market relative to the highest skill. In terms of the estimated labor disutility parameters, the results imply much larger variation. To illustrate this point, I consider a worker of type d_2 and a worker of type d_1 both with the highest skill level, s_1 , and assets equal to the economy-wide average. Under these circumstances, the reservation wage of the d_2 type worker (higher labor disutility) is 76 percent higher than the reservation wage of the d_1 type worker

(lower labor disutility). Doing the same exercise with a d_3 and a d_1 type worker, the reservation wage of the former is nearly 200 percent higher than that of the latter. This exercise over-exaggerates the true differences across workers in the model economy as the proportion of workers with skills equal to s_1 decreases with labor disutility. When setting the reservation wage of a d_2 type worker to the average reservation wage across all skill types (conditional on being type d_2), and analogously for a worker of type d_1 , the difference in reservation wages decreases to 19 percent. Likewise, the difference between d_3 and d_1 workers decreases to 47 percent.

The estimated persistence of the productivity process (0.935) is within the range for males and females that Chang and Kim (2006) estimate (0.948 and 0.925, respectively) using data from the PSID. The estimated standard deviation of the innovations to the productivity shock is much lower in comparison. I obtain an estimate of 0.12, while their estimates are 0.269 for males and 0.319 for females. This large difference is likely due to the fact that in my model wage differences across workers can be attributed to labor supply, skill and idiosyncratic differences. Meanwhile, in their model, wage differences are purely idiosyncratic (conditional on gender). Finally, the fact that individuals are disproportionately located along the diagonal of the matrix of disutility versus skill is expected given that the model must reproduce a positive correlation between employment (labor supply) and wages (skills).

Given the estimated parameter values from table 5.1, the aggregate steady-state employment rate of the model is 0.77. Table 3 presents the steady-state employment rates conditional on worker type. As expected given the utility specification, the model predicts fairly large employment rate differences across disutility types, and small differences within types. While the lowest disutility types are employed nearly all the time, the highest disutility types are employed roughly one third of the time. Hence, as in the data, most of the individuals

in the model display very high average employment rates, while a few others display comparatively low average employment rates.

Table 2: Estimated Parameter Values

Disutility of labor	Skills		
	$s_1 = 1.00^\dagger$	$s_2 = 0.557$ (0.002)	$s_3 = 0.319$ (0.0032)
$d_1 = 0.237$ (0.0033)	$p_{11} = 0.211$ (0.0117)	$p_{12} = 0.013$ (0.0125)	$p_{13} = 0.016$ (0.0074)
$d_2 = 0.644$ (0.0165)	$p_{21} = 0.0137$ (0.008)	$p_{22} = 0.471$ (0.0023)	$p_{23} = 0.017$ (0.010)
$d_3 = 1.922$ (0.008)	$p_{31} = 0.004$ (0.0038)	$p_{32} = 0.050$ (0.007)	$p_{33} = 0.203^\dagger$
	Value		
μ_x	1.583 (0.0051)		
ρ_x	0.936 (0.0034)		
σ_ϵ	0.120 (0.0068)		

Notes: \dagger by normalization. Asymptotic standard errors appear in parentheses.

Discount factor $\beta = 0.98751$, found from capital market clearing.

Table 3: Model steady-state employment rates, by worker type

Disutility of labor	Skills		
	s_1	s_2	s_3
d_1	0.99	0.99	0.99
d_2	0.88	0.89	0.90
d_3	0.29	0.30	0.30

Notes: Aggregate employment rate is 0.77.

5.2 Assessing the Model's Fit

In this subsection I discuss the goodness of fit of the model. Judging by the estimated results for the auxiliary model, it is unclear how well the model fits actual data.¹² Since it is ambiguous from these results if the differences are economically meaningful, I turn to the distributions of average employment and wages obtained from actual data and model-generated data next.

5.2.1 Employment and Wages

Figure 2 presents the distributions of average employment rates obtained from actual and model-generated data, while figure 3 presents the analogous distributions of average wages. Figure 4 presents the joint distributions of employment and wages from actual and model data. Most striking from figure 2 is how well the model matches the data distribution of employment rates. However, the model over-predicts the portion of individuals with average employment rates near 100%. This is likely why the model's overall employment rate is slightly higher than in the data (0.77 versus 0.745). In terms of the distribution of wages, the model also performs well. Relative to the data however, the distribution of wages in the model is less disperse and is skewed more to the right.

¹²Estimation results of the auxiliary model using actual data and model-generated data appear in tables 8 and 9 in appendix C.





Figure 2: Distribution of long run employment rates, data (top) and model (bottom).



Figure 3: Distribution of long run wage rates, data (top) and model (bottom).





Figure 4: Joint distributions of employment and wages, data (top) and model (bottom).

5.2.2 Hazard Rates

As can be seen from figures 5 and 6, the model is able to capture the negative duration dependence of both the hazard from employment to non-employment and the hazard from non-employment to employment. However, the model over-predicts the decline in both of these hazards for spells lasting at most 2 quarters. The reason for this result is purely compositional. In the model, flows from employment to non-employment occurring within the first 2 quarters of the duration of an employment spell, are disproportionately done by workers with the highest disutility of labor d_3 . Because these workers dislike market work so much, they engage in short lived employment spells, consistent with their low average employment rates. Likewise, flows from non-employment to employment occurring within the first 2 quarters of the duration of a non-employment spell also are disproportionately done by these same workers. Looking at the model's predicted hazards after 2 quarters (once most of the effect of type d_3 workers vanishes), the model performs better in replicating both the direction and level of both hazard rates.





Figure 5: Hazard rates from employment to non-employment, data (top) and model (bottom).

Note: Dashed lines represent 95% confidence interval of data.

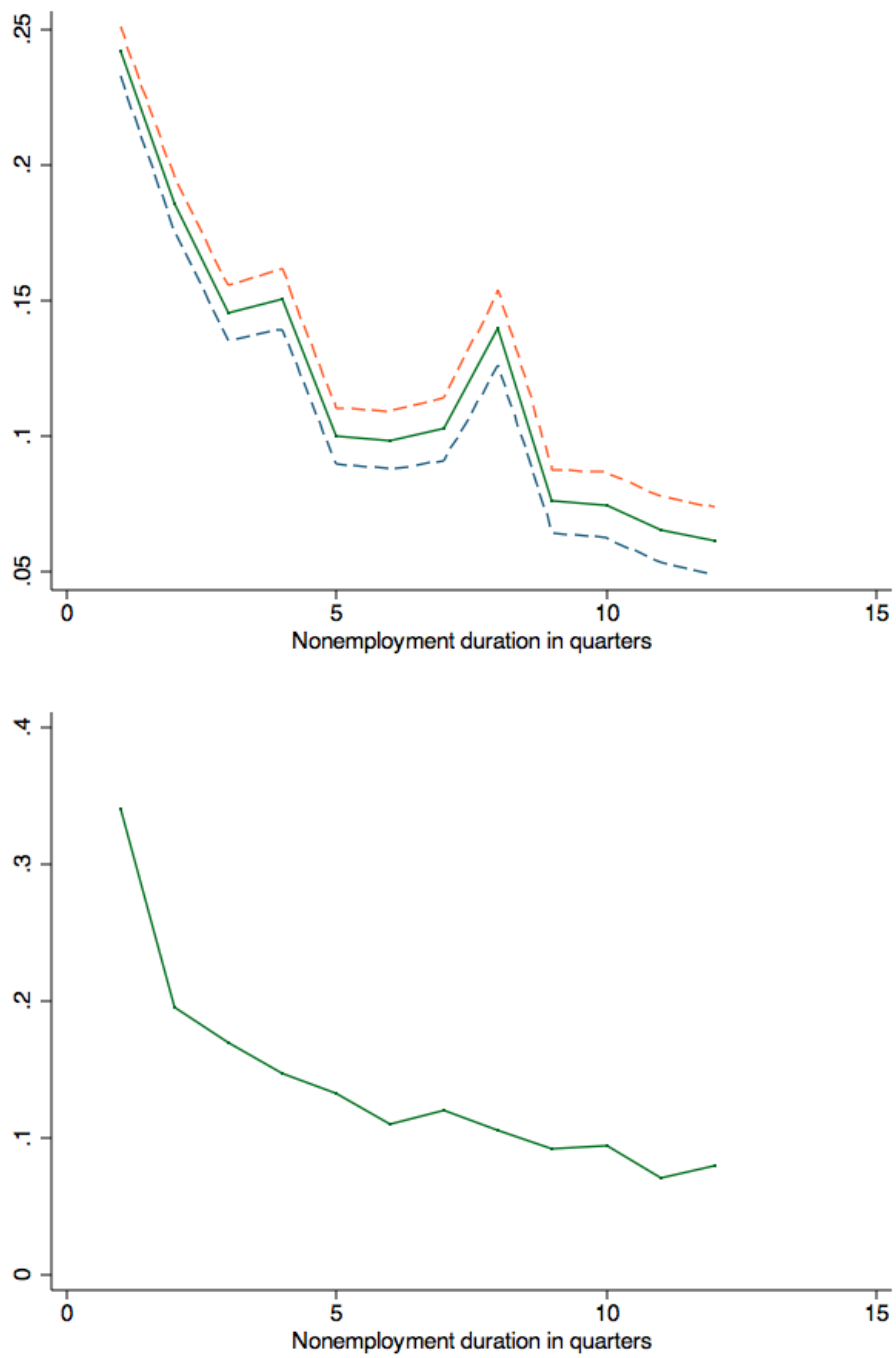


Figure 6: Hazard rates from non-employment to employment, data (top) and model (bottom).

Note: Dashed lines represent 95% confidence interval of data.

5.2.3 Wealth

As a final check on the model, I examine how well it replicates the cross-sectional wealth and earnings distribution observed in the data. In table 4 I present detailed statistics on wealth and earnings from the PSID, my model, and for comparison, Chang and Kim's (2007) model. As in Chang and Kim (2007), the category "PSID Primary Households" reflects households whose head is a high school graduate and whose age is between 35 and 55 as of 1983 (1984 survey). For each quintile group of wealth distribution, I calculate the wealth share, ratio of group average to economy-wide average, earnings share, and participation rate.

As can be seen in table 4, my model captures well earnings and wealth differences across quintiles. This is in spite of the fact that it was not estimated to match any of these features nor using data from the PSID. Both in the data and my model, the poorest 20 percent of families own almost nothing. In the data the richest 20 percent of families own nearly 58 percent of total wealth, while in my model they own nearly 56 percent of all wealth. Comparing my model to Chang and Kim (2007), the table shows that my model performs better in capturing the shares of wealth across all quintiles and in particular at the tails. While my model predicts essentially zero wealth for the poorest 20 percent, their model predicts negative wealth for this group. Meanwhile, their model over-predicts the wealth held by the richest 20 percent, while my model slightly under-predicts it. For the second through fourth quintiles my model also does a better job in matching these wealth shares. Most notably, for both the second and fourth quintiles, my model reduces the discrepancy in shares of wealth between model and data.

Finally, the key success of my model, which is absent in Chang and Kim (2007), is the predicted wealth effect on participation. In my model, because of the positive correlation between labor supply and skills, the wealthiest display fairly high participation rates. In the data, the fourth and fifth quintiles have labor market participation rates of 87 and 79 percent, respectively. In my model,

both quintiles have a participation rate of 70 percent. In Chang and Kim (2007), these quintiles participate at rates of 50 and 43 percent, respectively. Thus, my model does a considerably better job in capturing the labor supply decision of the wealthiest. More generally, except for the poorest 20 percent, my model does much better in replicating labor participation rates across all quintiles of wealth.

To summarize, my model replicates well the distributions of average employment rates and average wages as observed in the data. Moreover, it is also consistent with the negative duration dependence of both the hazard rate from employment to non-employment and vice-versa. A final check of the model's consistency shows that the model wealth distribution is consistent with salient features of the wealth distribution derived from the PSID. Most importantly, because of the positive correlation between labor supply and skills in the model, it generates a much more realistic wealth effect on participation.

Table 4: Summary Statistics of the Wealth Distribution

	Quintile				
	1	2	3	4	5
<i>PSID-primary households</i>					
Share of wealth	1.03	7.07	13.01	21.10	57.76
Group average/population average	0.05	0.36	0.64	1.06	2.97
Share of earnings	14.29	14.67	20.08	25.07	25.86
Participation rate	0.86	0.84	0.83	0.87	0.79
<i>Benchmark Model</i>					
Share of wealth	0.14	5.66	13.52	25.08	55.60
Group average/population average	0.01	0.28	0.68	1.25	2.78
Share of earnings	16.28	18.25	19.48	21.14	24.86
Participation rate	0.91	0.80	0.76	0.70	0.70
<i>Chang and Kim (2007)</i>					
Share of wealth	-2.46	3.27	12.21	26.05	60.93
Group average/population average	-0.12	0.16	0.61	1.30	3.08
Share of earnings	13.52	17.87	20.50	22.65	25.46
Participation rate	0.86	0.63	0.56	0.50	0.43

Notes: The PSID statistics reflect the family wealth and earnings in the 1984 survey as reported in Chang and Kim (2007).

6 Implications for the Frisch Elasticity at the Extensive Margin of Labor Supply

This section discusses the model's implications for the Frisch elasticity at the extensive margin of labor supply. First, it presents the baseline model's implied Frisch elasticity. Second, it presents a simple decomposition of this Frisch elasticity by considering two extreme cases of the baseline model.

6.1 Results for the Baseline Model

The implied Frisch elasticity at the extensive margin of labor supply of the model is 0.37. Note that this elasticity reflects no wealth effect as the entire wealth distribution is held constant. For comparison, Chang and Kim (2007) obtain an aggregate elasticity of 1.5, while Erosa et al. (2010) get an extensive margin elasticity of 0.69. Rogerson and Wallenius (2009) find elasticities ranging from 2.25 to 3.0. However, these elasticities reflect both intensive and extensive margins. The fact that the aggregate elasticity in my model is lower, precisely comes from constraining individual decisions in the model to match heterogeneity in individual employment rates in the data.

Table 5 presents the individual level employment elasticities by worker type. Again, these individual level elasticities reflect the percent change in participation (evaluated at the steady state participation rate for each worker type) given a one percent change in their steady state reservation wage. The results from table 5 show that the individual labor supply elasticity ranges from essentially zero to above 4. In the model, as in the data, a vast majority of the population is employed frequently and hence does not adjust their employment decision. Meanwhile, another portion of the population is employed less frequently and can adjust their labor supply more readily. However, because their contribution to overall employment is small, their elastic response is weighted less.

Table 5: Implied Elasticity from the steady-state reservation-wage distribution, by worker type and aggregate

Disutility of labor	Skills		
	s_1	s_2	s_3
d_1	0.05	0.05	0.05
d_2	0.5	0.47	0.45
d_3	4.6	4.4	5.2
Aggregate	0.37		

Notes: The numbers reflect the elasticity of the labor-market participation rate of each type (and overall) with respect to the reservation wage (evaluated at the steady-state) based on the steady-state reservation wage distribution.

6.2 The Role of Labor Supply Heterogeneity

In this subsection, I present results for two extreme cases of my model. I do this to understand whether labor supply or skill differences are the main reason for the low implied labor supply elasticity. In the first version of the model, agents only display ex-ante labor supply differences and have equal ex-ante market skills. In the second version of the model, agents only display ex-ante skill differences (akin to Erosa et al., 2010). For each case, I estimate the model using the same data and procedure as the baseline model and impose the corresponding restriction on skills or labor disutility. In both cases, the models are estimated to match a steady-state employment rate of 77 percent, which is what is obtained for the baseline model.

Table 6 presents the implied aggregate labor supply elasticities for each of the three models: labor disutility and skills (the baseline model); labor disutility only; and skills only. Each elasticity reflects a percentage change in the aggregate labor force participation rate (evaluated at a steady state rate of 77%, common to all models), given a percentage change in the steady state reservation wage holding the entire wealth distribution constant.

The first row reproduces the baseline aggregate elasticity of 0.37. What can be seen from the next two rows is that this low aggregate elasticity is overwhelmingly due to ex-ante labor supply differences. The model where skills are held constant produces an aggregate elasticity of 0.46. The model where labor disutility is held constant produces an aggregate elasticity of 1.2, only slightly below the elasticity of 1.5 recovered from Chang and Kim (2007). The key reason behind this result is that the model where labor disutility is held constant does a poor job in replicating the observed differences in average employment rates across workers. Figure 7 presents the distributions of average employment rates from the data (top), model with labor disutility differences (middle), and model with skill differences (bottom).

As can be seen from figure 7, the model with only skill differences produces a distribution of average employment rates which is dense near the steady-state employment rate. Because in the model these employment rate differences translate into reservation wages differences, the reservation wage distribution of this model is dense near a neighborhood of the steady-state wage rate. Hence, a large aggregate labor supply elasticity is recovered. Finally, table 7 shows another dimension where this model fails. Table 7 presents detailed statistics on wealth and earnings for each of the models and the PSID. The model with only skill differences under-performs, relative to the baseline model and model with labor supply differences, in reproducing a realistic wealth distribution and a realistic wealth effect on labor market participation. This further suggests that a model with ex-ante labor supply differences provides a closer description of actual data.

Conversely, the model with only labor supply differences is able to replicate a distribution of average employment rates similar to the one observed in the data. As a consequence, it produces a disperse reservation wage distribution and hence a low labor supply elasticity. Table 7 shows where the model with only labor disutility differences fails. As can be seen from the table, this model does not

have the same wealth effect on labor market participation as the model with both labor disutility and skill differences. In the model with only disutility differences, the richest 20 percent of the population work too little, while the poorest 20 percent work too much. This follows from the fact that in this model, the correlation between average employment and wages is negative. Individuals with a high disutility of labor work only when they receive high enough idiosyncratic productivity shocks and hence their average wage, conditional on employment, is counterfactually high. Because these individuals will have high asset holdings to finance their long non-employment spells, the counterfactual wealth effect on labor market participation is obtained.

Table 6: Aggregate Labor Supply Elasticity by Model

Model	Elasticity
Benchmark Model	0.37
Labor disutility only	0.46
Skills only	1.2

Notes: All elasticities are evaluated at a steady-state employment rate of 77%.




Figure 7: Distributions of average employment rates: data (top), model with labor disutility (middle), and model with skills (bottom).

Table 7: Summary Statistics of the Wealth Distribution

	Quintile				
	1	2	3	4	5
<i>PSID-primary households</i>					
Share of wealth	1.03	7.07	13.01	21.10	57.76
Group average/population average	0.05	0.36	0.64	1.06	2.97
Share of earnings	14.29	14.67	20.08	25.07	25.86
Participation rate	0.86	0.84	0.83	0.87	0.79
<i>Benchmark Model</i>					
Share of wealth	0.14	5.66	13.52	25.08	55.60
Group average/population average	0.01	0.28	0.68	1.25	2.78
Share of earnings	16.28	18.25	19.48	21.14	24.86
Participation rate	0.91	0.80	0.76	0.70	0.70
<i>Labor disutility only</i>					
Share of wealth	-0.27	4.44	12.48	25.19	58.17
Group average/population average	-0.01	0.22	0.62	1.26	2.92
Share of earnings	13.04	18.23	20.52	22.62	25.59
Participation rate	0.94	0.84	0.76	0.70	0.62
<i>Skills only</i>					
Share of wealth	-0.63	3.85	11.80	24.30	60.68
Group average/population average	-0.03	0.19	0.59	1.21	3.04
Share of earnings	15.81	18.35	18.83	21.39	25.63
Participation rate	0.96	0.83	0.75	0.70	0.61

Notes: The PSID statistics reflect the family wealth and earnings in the 1984 survey as reported in Chang and Kim (2007).

7 Conclusion

This paper examines the role of ex-ante heterogeneity across workers in determining the Frisch elasticity at the extensive margin of employment. Motivated by empirical observations from the NLSY that show large differences in average employment rates across individuals that do not project on wages, I develop a heterogeneous agent model with incomplete markets and indivisible labor supply to match these facts. The novel ingredients of the model are allowing agents to differ in their disutility of labor and market skills, both of which remain fixed across time. Unlike most of the previous literature, with Erosa et al. (2010) as an important exception, my model allows for a rich description of ex-ante heterogeneity (labor disutility and skills), and ex-post heterogeneity (idiosyncratic productivity shocks and assets) across agents. Rather than calibrating the model to match aggregate moments, I estimate the model's key micro-level parameters with indirect inference.

The main result of the paper can be summarized as follows. Once agents display a realistic amount of ex-ante heterogeneity in labor supply and skills a very large macro-level elasticity is no longer obtained through the extensive margin of labor supply. The implied aggregate labor supply elasticity of my model is 0.4. This result follows crucially from the labor supply differences observed in the data. In actual data, most individuals are typically employed and thus display high average employment rates, while few others display relatively low average employment rates.

A simple decomposition reveals the importance of these labor supply differences for the inferred Frisch elasticity of the extensive margin. In a version of my model with no ex-ante labor supply differences (akin to Erosa et al., 2010), the recovered elasticity is 1.2, nearly 3 times larger than the elasticity obtained from the baseline model. Meanwhile, in a version of my model with no ex-ante skill differences, the recovered elasticity is 0.46, only 15 percent larger than the elasticity

obtained from the baseline model. However, this version of my model violates the positive cross-sectional correlation between average employment rates and average wages that is observed in the data. In my baseline model with labor supply and skill differences, this correlation is positive. Moreover, because of this correlation, my model generates a realistic wealth effect on labor market participation, which is not found in Chang and Kim (2007).

Future research should consider allowing for some intensive margin adjustment (e.g. choice of hours conditional on being employed subject to some minimum requirement) as an extension of my model to verify that the results are not driven by the assumption of no intensive margin choice. Verifying that the hours choice by worker type is consistent with what is observed in the data is another important check of my model's consistency. Allowing for a distinction between men and women in the model is also a promising venue of research as the current model abstracts from the difference between the labor market participation decision of a married woman versus a single man. Work by Guner et al. (2008) shows that this distinction is very important. Finally, extending the model to allow for business cycle shocks is also a promising direction of research. The structure of the model can help quantify how much of the volatility of aggregate employment and wages is due to the employment response of each of the worker types over the business cycle. Obtaining answers to these questions will further our knowledge about the aggregate implications of individual level heterogeneity.

References

- Abowd, John M. and David Card**, “On the Covariance Structure of Earnings and Hours Change,” *Econometrica*, 1989, 57 (2), 411–45.
- Aiyagari, S. Rao**, “Uninsured Idiosyncratic Risk and Aggregate Saving,” *The Quarterly Journal of Economics*, 1994, 109 (3), 659–84.
- Altonji, Joseph, Anthony A. Smith, and Ivan Vidangos**, “Modeling Earnings Dynamics,” 2009. NBER Working Paper # 14743.
- Altonji, Joseph G.**, “Intertemporal Substitution in Labor Supply: Evidence from Micro Data,” *Journal of Political Economy*, 1986, 94 (3, part 2), S176–S215.
- Auerbach, Alan J. and Laurence J. Kotlikoff**, *Dynamic Fiscal Policy*, Cambridge, UK: Cambridge University Press, 1987.
- Bils, Mark, Yongsung Chang, and Sun-Bin Kim**, “Comparative Advantage and Unemployment,” 2009. Mimeo University of Rochester.
- Chang, Yongsung and Sun-Bin Kim**, “From Individual to Aggregate Labor Supply: A Quantitative Analysis Based on a Heterogeneous Agent Macroeconomy,” *International Economic Review*, 2006, 47 (1), 1–27.
- and —, “Heterogeneity and Aggregation: Implications for Labor-Market Fluctuations,” *American Economic Review*, 2007, 97 (5), 1939–56.
- Cho, Jang-Ok**, “Ex post heterogeneity and the business cycle,” *Journal of Economic Dynamics and Control*, 1995, 19, 533–51.
- Coleman, Thomas**, “Essays on Aggregate Labor Market Business Cycle Fluctuations.” PhD dissertation, University of Chicago 1984.

- Erosa, Andrés, Luisa Fuster, and Gueorgui Kambourov**, “A Micro-founded Theory of Aggregate Labor Supply with Non-linear Wages,” 2010. Mimeo University of Toronto.
- Gallant, A. Ronald and George Tauchen**, “Indirect Inference,” *Journal of Applied Econometrics*, 1993, 8, S85–S118.
- Ghez, Gilbert and Gary Becker**, *The Allocation of Time and Goods over the Life Cycle*, New York: National Bureau of Economic Research, Columbia University Press, 1975.
- Gourieroux, C, A Monfort, and E Renault**, “Indirect Inference,” *Journal of Applied Econometrics*, 1993, 8, S85–S118.
- Gourio, François and Pierre-Alexandre Noual**, “The Marginal Worker and The Aggregate Elasticity of Labor Supply,” 2006. Boston University Working Paper WP2006-009.
- Guner, Nezih, Remzi Kaygusuz, and Gustavo Ventura**, “Taxation, Aggregates and the Household,” 2008. IZA Discussion Paper No. 3318.
- Guvenen, Fatih and Anthony A. Smith**, “Inferring Labor Income Risk From Economic Choices: An Indirect Inference Approach,” 2010. Working paper.
- Hansen, Gary D.**, “Indivisible Labor and The Business Cycle,” *Journal of Monetary Economics*, 1985, 16, 309–27.
- Heckman, James**, “Comments on Ashenfelter and Kydland Papers,” *Carnegie-Rochester Conference Series on Public Policy*, 1984, 21, 209–24.
- Keane, Michael and Anthony A. Smith**, “Generalized Indirect Inference for Discrete Choice Models,” 2003. Working paper.

- Kimmel, Jean and Thomas J. Kniesner**, “New evidence on labor supply: Employment versus hours elasticities by sex and marital status,” *Journal of Monetary Economics*, 1998, *42*, 289–01.
- Kydland, Finn E. and Edward C. Prescott**, “Time to Build and Aggregate Fluctuations,” *Econometrica*, 1982.
- Lucas, Robert E. Jr. and Leonard Rapping**, “Real Wages, Employment and Inflation,” *Journal of Political Economy*, 1969.
- Macurdy, Thomas E.**, “An Empirical Model of Labor Supply in a Life-Cycle Setting,” *Journal of Political Economy*, 1981.
- Prescott, Edward C.**, “Why Do American Work So Much More Than Europeans?,” *Federal Reserve Bank of Minneapolis Quarterly Review*, 2004, *28* (1), 2–13.
- Press, William H., Saul A. Teukolsky, William T. Vetterling, and Brian P. Flannery**, *Numerical Recipes in C*, Cambridge: Cambridge University Press, 1992.
- Ríos-Rull, José-Víctor**, “Computation of Equilibria in Heterogenous-Agent Models,” in Ramon Marimon and Andrew Scott, eds., *Computational Methods for the Study of Dynamic Economies*, New York: Oxford University Press, 1999.
- Rogerson, Richard**, “Indivisible Labor, Lotteries and Equilibrium,” *Journal of Monetary Economics*, 1988, *21*, 3–16.
- **and Johanna Wallenius**, “Micro and Macro Elasticities in a Life Cycle Model with Taxes,” *Journal of Economic Theory*, 2009, *144* (6), 2277–92.
- Rupert, Peter, Richard Rogerson, and Randall Wright**, “Home Work

in Labor Economics: Household Production and Intertemporal Substitution,” *Journal of Monetary Economics*, 2000, 46 (3), 557–79.

Smith, Anthony A., “Estimating Nonlinear Time-Series Models using Simulated Vector Autoregressions,” *Journal of Applied Econometrics*, 1993, 8, S63–S84.

—, “Indirect Inference,” in Steven N. Durlauf and Lawrence E. Blume, eds., *The New Palgrave Dictionary of Economics*, 2nd ed., Palgrave Macmillan, 2008. The New Palgrave Dictionary of Economics Online. Palgrave Macmillan. 22 September 2010 <http://www.dictionaryofeconomics.com/article?id=pde2008_I000259> doi:10.1057/9780230226203.0778.

Smith, Anthony A. Jr., “Three Essays on the Solution and Estimation of Dynamic Macroeconomic Models.” PhD dissertation, Duke University 1990.

Tauchen, George., “Finite State Markov-Chain Approximations to Univariate and Vector Autoregressions,” *Economics Letters*, 1986.

Data Appendix

A Data

A.1 Linking Employers Across Survey Years

The NLSY allows the linking of an individual’s job reports across consecutive survey years. In linking reports across survey years I follow the method suggested in the NLSY technical Appendix # and use the variables defined as “*Previous job number at last interview #1-5*”.

A.2 Constructing Quarterly Employment Status

The NLSY79 provides variables containing the weekly employment status of each individual in the sample in their work history file. These variables are named “*Labor Force Status Week #*”, where $\#$ serves as a place holder for the week number in question. Each calendar week is assigned a number starting with 1 (corresponding to the first of January 1978), through 1531 (corresponding to the week starting with February 29th 2007). For each individual I construct their employment status for quarter q as follows:

1. For quarter q determine the week numbers \underline{w} and \bar{w} which correspond to the first and last weeks in the quarter.
2. For each week in $[\underline{w}, \bar{w}]$ check if the individual is employed (status code ≥ 100 or 3), non-employed (status code 2,4, or 5) or missing (status code 0 or 7).
3. If the individual is employed for at least 7 weeks in the quarter, set her quarterly employment status to employed. If the individual is not employed for at least 7 weeks, but has at least one week where her status is not missing, set her quarterly employment status to non-employed. Otherwise, set her status to missing.

A.3 Wages

I use the NLSY79 calculated hourly wage rates included in the variables “*Hours usually worked at current/most recent job*” and “*Hourly Rate of Pay Job #1-5*”. From 1979-1993 detailed information on the CPS or current/most recent employer is collected in the CPS section, while after 1993 the CPS employer is always the first job coded. Hence, for survey years 1979-1993 it is necessary to look at both sets of variables to obtain complete information on the CPS job.

If an individual reports wages in units other than hourly, the NLSY calculates an hourly wage rate based on the earnings reported, the unit in which they are reported and usual hours worked on the job. Nominal wages are deflated using the Consumer Price Index for all urban consumers and all items (CPI-U) which is seasonally adjusted. I impute missing wages using the average of wage from reports of the same job from other survey years. If only one report is available for a particular job, I impute the missing job using the average wage from all job reports from the same individual.

A.4 Hours

To identify hours worked in each job reported I combine the variables “*Hours usually worked at current/most recent job*” and “*Hours per week usually worked at Job # 1-5*”, deferring to the CPS report whenever the job coincides with the current/most recent employer.

A.5 Definition of Quarterly Wage

I define the quarterly wage rate as the hourly wage rate of the job the individual works at the most during the quarter in question. To do this, I look at the product of hours per week and weeks worked in the quarter to determine which job the individual worked at the most during the quarter.

B Computation of the Steady-State Equilibrium

The computational strategy used to compute the steady-state equilibrium of the model is an extension of the one used in Chang and Kim (2007), to take into account multiple worker types. As in Ríos-Rull (1999), my goal is to find the discount rate β which clears the capital market given an interest rate of 1%. To

do so, I solve for the invariant measure $\{\mu_{sd}(a, x)\}_{s=1, d=1}^{N_s, N_d}$ of workers across assets and productivity, given their type as follows:

0. Initialize guesses (or current estimates) for $\{s_1, \dots, s_{N_s}\}, \{d_1, \dots, d_{N_d}\}, \{p_{sd}\}_{s=1, d=1}^{N_s, N_d}, \sigma_x, \rho_x$.

1. Choose the grid points for asset holdings a and idiosyncratic productivity x . Denote the number of grids by N_a and N_x . I set $N_a = 1,666$ and $N_x = 10$. Asset holdings a are restricted to the range $[-2, 2000]$, where the average asset holdings are 13.7. The grid points on asset are not equally spaced; more points are assigned on the bottom of the asset range to better approximate the savings decisions of workers with lower assets. For idiosyncratic productivity, I construct a vector of length N_x , whose elements $\ln x_j$, are equally spaced on the interval $[-3\sigma_x/\sqrt{1-\rho_x^2}, +3\sigma_x/\sqrt{1-\rho_x^2}]$. I use Tauchen's (1986) algorithm to approximate the idiosyncratic productivity process using a transition matrix.
2. Given values for β , skills s , and labor disutilities d , I solve for the value functions $\{V_{sd}^E, V_{sd}^N, V_{sd}\}_{s=1, d=1}^{N_s, N_d}$ at each grid point of the individual states. By doing so I also obtain the optimal decision rules for asset holdings and labor supply for each worker type $\{a'_{sd}(a_i, x_j), h_{sd}(a_i, x_j)\}_{s=1, d=1}^{N_s, N_d}$. The value functions are found iteratively as follows:

- (a) Initialize the value functions $V_{sd}^E(a_i, x_j)$ and $V_{sd}^N(a_i, x_j)$ for all $i = 1, \dots, N_a, j = 1, \dots, N_x, s = 1, \dots, N_s$ and $d = 1, \dots, N_d$.
- (b) Obtain updated guesses of the value functions by evaluating the dis-

cretized versions

$$\begin{aligned}\tilde{V}_{sd}^E(a_i, x_j) &= \max_{a' \in \{a_1, \dots, a_{N_a}\}} \left\{ \ln \left(w\bar{h}s_s x_j + (1+r)a_i - a' \right) - d_d \right. \\ &\quad \left. + \beta \sum_{k=1}^{N_x} V_{sd}(a', x_j) \pi_x(x_k | x_j) \right\} \\ \tilde{V}_{sd}^N(a_i, x_j) &= \max_{a' \in \{a_1, \dots, a_{N_a}\}} \left\{ \ln \left((1+r)a_i - a' \right) \right. \\ &\quad \left. + \beta \sum_{k=1}^{N_x} V_{sd}(a', x_j) \pi_x(x_k | x_j) \right\}\end{aligned}$$

where $\pi_x(x'|x_j)$ is the transition probability of x_j to x' . Update $\tilde{V}_{sd}(a_i, x_j) = \max\{\tilde{V}_{sd}^E(a_i, x_j), \tilde{V}_{sd}^N(a_i, x_j)\}$.

- (c) If \tilde{V} and V are close enough for all grid points and for each s, d pair, then we have found the value functions. Otherwise, set $V_{sd}^E = \tilde{V}_{sd}^E$ for each s, d pair and all grid points (and similarly for V^N), and go back to step 2 (b).

3. Using $\{a'_{sd}(a_i, x_j)\}_{s=1, d=1}^{N_s, N_d}$ obtained from step 2 and $\pi_x(x'|x_j)$, obtain the time-invariant measures $\{\mu_{sd}^*(a_i, x_j)\}_{s=1, d=1}^{N_s, N_d}$ as follows:

- (a) Initialize the measures $\{\mu_{sd}(a_i, x_j)\}_{s=1, d=1}^{N_s, N_d}$, such that

$$\sum_{i=1}^{N_a} \sum_{j=1}^{N_x} \mu_{sd}(a_i, x_j) = p_{sd}, \text{ where } p_{sd} \text{ is the proportion of the population with skills } s = s_s \text{ and labor disutility } d = d_d.$$

- (b) Update each measure by evaluating a discretized version of (5) (for each (s, d) pair):

$$\mu'_{sd}(a_{i'}, x_{j'}) = \sum_{i=1}^{N_a} \sum_{j=1}^{N_x} \mathbf{1}_{a_{i'}=a'_{sd}(a_i, x_j)} \mu_{sd}(a_i, x_j) \pi_x(x_{j'} | x_j)$$

- (c) If μ'_{sd} and μ_{sd} are close enough for all grid points and each s, d pair, then we have found the time-invariant measure. Otherwise, set $\mu_{sd} = \mu'_{sd}$ and go back to step 3 (b).

4. Calculate the real interest rate as a function of β , $r(\beta) = \alpha \left(K(\beta), L(\beta) \right)^{1-\alpha} - \delta$, where

$$K(\beta) = \sum_{s=1}^{N_s} \sum_{d=1}^{N_d} \sum_{i=1}^{N_a} \sum_{j=1}^{N_x} a_i \mu_{sd}^*(a_i, x_j)$$

and

$$L(\beta) = \sum_{s=1}^{N_s} \sum_{d=1}^{N_d} \sum_{i=1}^{N_a} \sum_{j=1}^{N_x} s_s x_j h_{sd}(a_i, x_j) \mu_{sd}^*(a_i, x_j).$$

If $r(\beta)$ is close enough to the assumed value of the real interest rate, we have found the steady-state. Otherwise, choose another β and go back to step 2.

C Estimates from the Auxiliary Model

Table 8: Estimated Results for Auxiliary Model: Actual vs Model-Generated Data

Parameter	E equation		N equation		W equation	
	Actual	Model	Actual	Model	Actual	Model
γ_{ED}	.0465029 (.0000608)	0.05103 (0.00020)	-.0235211 (.0000608)	-0.02725 (0.00018)	-.0116675 (.0001098)	-0.01511 (0.00044)
γ_{ND}	-1.1350696 (.0000608)	-0.08960 (0.00057)	-.0068707 (.0000608)	-0.02165 (0.00030)	.0093461 (.0001098)	0.03886 (0.00041)
γ_w	.2419846 (.0008472)	0.26982 (0.00063)	-.0646542 (.0008472)	-0.06984 (0.00045)	.7662637 (.0015309)	0.81640 (0.00119)
$\gamma_{\bar{e}}$.1361421 (.0027194)	0.19271 (0.00111)	.14039 (.0027194)	0.11108 (0.00101)	.5001013 (.0049141)	0.49674 (0.00206)
$\gamma_{\bar{w}}$	-1.1660909 (.0013099)	-0.14918 (0.00122)	.0382842 (.0013099)	0.01124 (0.00069)	.1722257 (.0023671)	0.11485 (0.00116)
γ_0	.5522605 (.0025226)	0.39513 (0.00231)	-.0027554 (.0025226)	0.10051 (0.00103)	-.3351861 (.0045584)	-0.30821 (0.00156)
R^2	0.7972	0.80582	0.1177	0.11464	0.7985	0.82009

Notes: Standard Errors in parentheses. Model coefficients are averages over 50 simulations. Model standard errors are calculated from the distribution of each parameter over the 50 simulations.

Table 9: Estimated Results for Other Moments: Actual vs Model-Generated Data

Moment	Actual Data	Model Simulated Data
μ_e	.746	0.778 (0.0005)
μ_w	1.99	1.921 (0.001)
σ_e	.282584	0.2848 (0.0004)
σ_w	.5180253	0.430 (0.0007)
$Skewness_e$	-.9741168	-1.13 (0.005)
$Skewness_w$.1866516	0.056 (0.006)
$\rho(e, w)$	0.3965	0.31 (0.002)