

Individual and Aggregate Labor Supply in a Heterogeneous Agent Economy with Intensive and Extensive Margins*

Yongsung Chang
University of Rochester
Yonsei University

Sun-Bin Kim
Yonsei University

Kyooho Kwon
University of Rochester

Richard Rogerson
Princeton University
NBER

September 13, 2012

Abstract

We develop a heterogeneous-agent general equilibrium model that incorporates both intensive and extensive margins of labor supply. A nonconvexity in the mapping between time devoted to work and labor services distinguishes between extensive and intensive margins. We consider calibrated versions of this model that differ in the value of a key preference parameter for labor supply and the extent of heterogeneity. The model is able to capture the key features of the empirical hours worked distribution, including how individuals transit within this distribution. We then study how the various specifications influence labor supply responses to temporary shocks and permanent tax changes, with a particular focus on the intensive and extensive margin elasticities in response to these changes. We find important interactions between heterogeneity and the extent of curvature in preferences.

Keywords : Hours, Employment, Cross-section, Business Cycles

*Correspondence: Chang: Department of Economics, University of Rochester and Yonsei University, yongsung.chang@gmail.com; Kim: Department of Economics, Yonsei University, sunbin.kim@yonsei.ac.kr; Kwon: Department of Economics, University of Rochester. kkwon5@mail.rochester.edu; Rogerson: Department of Economics, Princeton University, rdr@princeton.edu. We thank Richard Blundell for comments as well as conference participants at Yonsei University, Institute for Fiscal Studies and the 2012 SED. Chang and Rogerson gratefully acknowledge the support from the WCU (R33-10005) Program through the Korea Research Foundation. Rogerson also thanks the NSF for financial support.

1. Introduction

Connecting individual and aggregate labor supply has long been controversial. An early controversy revolved around the fact that stand-in household models such as Kydland and Prescott (1982) assumed much larger Frisch labor supply elasticities than those estimated by researchers such as MaCurdy (1981) and Altonji (1986) from micro data. Heckman (1984) argued that this controversy was somewhat misdirected given that these studies abstracted from labor adjustment along the extensive margin. Subsequent work has addressed this deficiency, and led to a few general points of consensus. First, aggregate labor supply responses depend upon adjustment along both the intensive and extensive margins. Second, the individual preference parameter estimated by MaCurdy, Altonji and others is closely related to aggregate responses along the intensive (i.e., hours per worker) margin. And third, the extent of adjustment along the extensive margin is heavily influenced by the amount of heterogeneity, or more specifically, the mass of workers that are close to a “reservation wage curve”.

Chetty et al. (2011) argues that further headway in understanding labor supply requires connecting macro models with microeconomic estimates of elasticities for both intensive and extensive margins, in both steady state and business cycle settings. We agree with this call for future research, but note that this call for research points to a key void in the literature: a systematic analysis of how preference parameters and the extent of heterogeneity interact to affect the connection between individual and aggregate labor supply in aggregate models that allow for both intensive and extensive margins. Many papers that consider adjustment along the extensive margin abstract entirely from adjustment along the intensive margin, and models that feature both margins often abstract from heterogeneity.¹ This paper seeks to fill this void by developing and analyzing a model that features heterogeneity in

¹Hansen (1985) is the first of many papers that study business cycles in a setting with only adjustment along the extensive margin and no heterogeneity. Chang and Kim (2007) introduced heterogeneity but still did not allow for an intensive margin. Cho and Cooley (1994) and Kydland and Prescott (1991) for early studies that feature adjustment along both margins of labor supply in a general equilibrium framework, but without heterogeneity.

addition to adjustment along the intensive and extensive margins. Our model merges the framework of Chang and Kim (2006, 2007) with that of Prescott et al. (2009) and Rogerson and Wallenius (2009).

We consider several specifications of the model that differ in the extent of heterogeneity and the value of the preference parameter that dictates curvature in the disutility of work function. All of the specifications are calibrated so as to match the average employment rate and the average level of hours per worker. We assess the ability of these specifications to account for various steady state observations, including the distribution of hours of work across individuals, the transition of individuals in the hours of work distribution over time, and the distribution of labor earnings and wealth. Although our model is parsimonious, it is able to account for many of the salient features found in the data. There are two novel aspects to our steady state calibration exercise that we think are worth noting. First, we argue that the cross-sectional distribution of hours of work can serve as useful information regarding the extent of heterogeneity in the data. Interestingly, based on this measure, we need a degree of heterogeneity that is roughly double the amount captured by estimates of idiosyncratic wage shocks. As we note later on, it turns out that the extent of heterogeneity has important implications, so the development of simple procedures for assessing the appropriate degree of heterogeneity within an aggregate model is important. Second, to our knowledge ours is the first aggregate analysis to address how individuals transition within the hours worked distribution. Previous analyses have instead focused on how individuals transition between the states of employment and non-employment.

We then consider the implications of these different specifications for two standard types of analyses: business cycle fluctuations and steady state tax analyses. For ease of exposition and transparency we focus on specifications that are common in the literature. In the business cycle analysis we consider aggregate shocks to productivity as the driving force behind business cycles, and for the steady state tax analysis we consider a lump-sum transfer program financed by a proportional tax on labor earnings. Several findings emerge from our analysis. First, in contrast to the apparent consensus that has emerged,

we find that fluctuations along the intensive and extensive margin are both functions of the underlying curvature parameter in preferences and the extent of heterogeneity. That is, intensive and extensive elasticities are not independent of each other and so must be considered jointly. Second, we find that the extent of heterogeneity has very different effects on extensive margin responses in the two settings: in the business cycle setting we find that increased heterogeneity leads to less response along the extensive margin, whereas in the case of permanent changes to tax and transfer programs, we find that increased heterogeneity leads to more response along the extensive margin. Third, in the business cycle context we find that the presence of an intensive margin dampens the effects of increased heterogeneity in an important way. It follows that abstracting from the intensive margin is a serious issue for analyses that seek to understand the effect of heterogeneity on the magnitude of aggregate fluctuations. Fourth, we find that in the context of permanent tax changes, the preference parameter plays effectively no role in influencing the change in aggregate hours when the extent of heterogeneity is relatively low, but does play a role as the extent of heterogeneity increases. That is, the extent of heterogeneity interacts with the preference parameter in a quantitatively important manner.

Our paper is related to several in the literature. Relative to the business cycle analysis of Chang and Kim (2007), we add an intensive margin. Relative to the tax analysis of Rogerson and Wallenius (2009), we consider a richer environment in terms of heterogeneity, and allow for uncertainty and incomplete markets. The paper that is closest to ours is Erosa et al. (2011). Like us, they build a model that features heterogeneity and incomplete markets and allows for labor supply adjustment along the intensive and extensive margin. They adopt a life cycle structure, consider a richer environment in terms of sources of heterogeneity, make more effort to isolate a definitive calibration. We carry out a more extensive analysis of business cycle fluctuations and consider a wider range of specifications in order to assess the implications of variation in the key preference parameter and the extent of heterogeneity. We view these two pieces of work as complementary.

The paper is organized as follows: Section 2 specifies the model. Section 3 calibrates

the different specifications of the model economy and Section 4 considers the steady state properties of the various specifications. Section 5 studies the business cycle properties of the model, while Section 6 considers the effects of permanent tax changes. Section 7 concludes.

2. Model

The model is essentially the indivisible labor incomplete markets model of Chang and Kim (2007) extended in the spirit of Prescott et al. (2009) to allow for adjustment along the intensive margin. The details follow. There is a unit measure of ex-ante identical infinitely lived individuals. Each individual has preferences over streams of consumption (c_t) and hours of work (h_t) given by:

$$\sum_{t=0}^{\infty} \beta^t [\log c_t - B \frac{h_t^{1+1/\gamma}}{1+1/\gamma}]$$

where $0 < \beta < 1$, $B > 0$ and $\gamma > 0$.

There is an aggregate Cobb-Douglas production function that produces output using inputs of labor services (L_t) and capital services (K_t) and is subject to TFP shocks (λ_t):

$$Y_t = \lambda_t L_t^\alpha K_t^{1-\alpha}.$$

The aggregate productivity λ_t evolves with a transition probability distribution function $\pi_\lambda(\lambda'|\lambda) = \Pr(\lambda_{t+1} \leq \lambda' | \lambda_t = \lambda)$. In our quantitative analysis we will assume that λ_t follows an AR(1) process in logs:

$$\ln \lambda_{t+1} = \rho_\lambda \ln \lambda_t + \varepsilon_{\lambda t}, \quad \varepsilon_{\lambda t} \sim N(0, \sigma_\lambda^2).$$

Output can be used for either consumption or investment, and capital depreciates at rate δ .

Two features influence the mapping from time devoted to work to labor services. The first is that individuals are subject to idiosyncratic productivity shocks, denoted by x_t . The

stochastic evolution of x_t is described by the transition probability distribution function $\pi_x(x'|x) = \Pr(x_{t+1} \leq x' | x_t = x)$. In our quantitative work we will also assume that x_t follows an AR(1) process in logs:

$$\ln x_{t+1} = \rho_x \ln x_t + \varepsilon_{xt}, \quad \varepsilon_{xt} \sim N(0, \sigma_x^2).$$

The second feature is a non-convexity associated with such factors as set-up costs, supervisory time and/or the need to coordinate with other workers. If an individual with idiosyncratic productivity x_t devotes h_t units of time to market work, this will generate $x_t g(h_t)$ units of labor services. Following Prescott et al (2009) and Rogerson and Wallenius (2009), we assume that $g(\cdot)$ takes the following simple form::

$$g(h_t) = \max\{0, h_t - \bar{h}\}, \quad h_t \in [0, 1].$$

Following Bewley (1986), Huggett (1993), and Aiyagari (1994) we assume that markets are incomplete in the sense that there are no markets for insurance and the only asset is physical capital. Individuals trade claims to physical capital, and these claims are denoted by a . Additionally, there is an exogenous borrowing constraint that limits the amount of debt that an individual can acquire:

$$a_t \geq \bar{a}$$

In each period t there is a market for units of labor services, with price w_t , and a rental market for capital services, with price $r_t + \delta$, so that r_t is the rate of return to capital. When a worker of productivity x_t devotes h_t units of time to market work, the resulting labor earnings are $w_t x_t g(h_t)$. The government taxes labor earnings at a flat rate of τ and rebates all the revenue with a lump-sum transfer amount, Tr , to all households.

Our model assumes that all (exogenous) heterogeneity occurs along one dimension—that of productivity in market work. More generally, one could also imagine that individuals differ along a second dimension, which could be thought of as the return to non-market work, either in the form of differences in the value of leisure time or in the productivity of non-market work. From the perspective of market labor supply choices, what really

matters is the relative return to market work. While we could have introduced a second idiosyncratic shock, to maintain parsimony we have opted for a single idiosyncratic shock that we will interpret as a composite shock. It will be relevant to keep this in mind when we discuss calibration.

We formulate equilibrium recursively. The individual state variables are beginning of period assets (a) and current idiosyncratic productivity (x), and the aggregate state variables will be the current aggregate productivity shock (λ) and a measure μ over the individual state variables (a, x). Prices are functions of the aggregate state: $w(\lambda, \mu)$ and $r(\lambda, \mu)$, and the equilibrium law of motion for μ is given by $\mu' = T(\lambda, \mu)$.

The value function for a worker, denoted by V , is:

$$\begin{aligned}
 V(a, x; \lambda, \mu) &= \max_{c, a', h} \left\{ \log(c) - B \frac{h^{1+\frac{1}{\gamma}}}{1+\frac{1}{\gamma}} + \beta E [V(a', x'; \lambda', T(\lambda, \mu)) \mid x, \lambda] \right\} \\
 \text{s.t.} \quad &c = (1 - \tau)w(\lambda, \mu)x \max \{0, h - \bar{h}\} + (1 + r(\lambda, \mu))a - a' + Tr \\
 &c \geq 0, a' \geq \bar{a}, 0 \leq h \leq 1
 \end{aligned}$$

An equilibrium consists of a value function $V(a, x; \lambda, \mu)$, individual decision rules $c(a, x; \lambda, \mu)$, $a'(a, x; \lambda, \mu)$, $h(a, x; \lambda, \mu)$, aggregate inputs $\{K(\lambda, \mu), L(\lambda, \mu)\}$, factor prices $\{w(\lambda, \mu), r(\lambda, \mu)\}$, and a law of motion $T(\lambda, \mu)$ such that

1. Individuals optimize:

Given factor prices, individual decision rules solve value function.

2. The representative firm maximizes profits : For all (λ, μ)

$$\begin{aligned}
 w(\lambda, \mu) &= F_1(L(\lambda, \mu), K(\lambda, \mu), \lambda) \\
 r(\lambda, \mu) &= F_2(L(\lambda, \mu), K(\lambda, \mu), \lambda) - \delta
 \end{aligned}$$

3. The goods market clears : For all (λ, μ)

$$\int \{a' + c\} d\mu = Y + (1 - \delta)K$$

4. Factor markets clear :

$$\begin{aligned} L(\lambda, \mu) &= \int x g(h(a, x; \lambda, \mu)) d\mu \\ K(\lambda, \mu) &= \int a d\mu \end{aligned}$$

5. Government balances its budget :

$$\int \{\tau w(\lambda, \mu) x \max\{0, h - \bar{h}\}\} d\mu = Tr$$

6. Individual and aggregate behaviors are consistent :

$$\mu'(A^0, X^0) = \int_{A^0, X^0} \left\{ \int_{\mathcal{A}, \mathcal{X}} 1 [a' = a'(a, x; \lambda, \mu)] d\pi_x(x'|x) d\mu \right\} da' dx'$$

for all $A^0 \subset \mathcal{A}$, $X^0 \subset \mathcal{X}$.

3. Calibration

Our key objective is to examine how different values of the preference parameter γ and the amount of heterogeneity influence the magnitude and nature of responses in aggregate hours to business cycle shocks and permanent fiscal policy changes. A simple way to vary the amount of heterogeneity in the economy is to vary the standard deviation of the innovations to the idiosyncratic shock process, and this is the approach that we follow. Accordingly, we will consider specifications in which γ takes on values in the set $\{0.5, 1.0, 1.5\}$ and σ_x takes on values in the set $\{0.165, 0.2475, 0.330\}$. Motivation for these values is given below.

As is standard in the business cycle literature we assume that each period corresponds to one quarter. Many of our parameters are standard in the literature and so we set them to be in line with previous studies. In particular, the labor-income share, α , is set to 0.64, the depreciation rate, δ , is set to .025, and for the aggregate technology shocks we set $\rho_\lambda = 0.95$, and $\sigma_\lambda = 0.007$. For the benchmark economy, there is no tax and transfer: $\tau = 0$ and $Tr = 0$.

We noted earlier that our idiosyncratic shock is best thought of a composite shock that reflects the net effect of idiosyncratic shocks on the relative return to working in the market versus not working. A reasonable lower bound on the size of these shocks is provided by the literature that estimates idiosyncratic shocks to wages.² A sizeable literature has done this for prime age males, including, for example, Card (1994), Floden and Linde (2001), French (2005), Chang and Kim (2006), and Heathcote et al. (2007). While there is some variation across studies, the consensus is that these shocks are large and persistent. Guided by these empirical studies, for one of our specifications we set $\rho_x = 0.975$ and $\sigma_x = 0.165$.³ We assume that the individual and the aggregate shock processes are orthogonal.

As noted above, we view this as a reasonable lower bound on the extent of heterogeneity induced by idiosyncratic shocks. As an upper bound we consider a specification with the same persistence, i.e., $\rho_x = .975$, but double the standard deviation of the innovations so that $\sigma_x = .330$. As we document later, the reason we view this is a reasonable upper bound is that it generates a dispersion in hours worked that exceeds what is found in the data. We also consider one intermediate value ($\sigma_x = .2475$).

Motivation for the set of values considered for γ comes from the implied range of Frisch elasticities when we run a standard labor supply regression for workers with positive hours using micro data generated from the model. Recent work by Chetty (2010) argues that an empirically reasonable value for this elasticity is in the range of .40 – .50. As discussed in Keane and Rogerson (2011), there are additional factors that Chetty abstracts from that would suggest higher values. Our set of values for γ leads to a range of estimated Frisch elasticities that run from around .25 to 1.00, which we think of as a reasonable range.

There are four additional parameters to calibrate. For a given value of γ and σ_x , we choose values of the discount factor, β , disutility of work parameter, B , the borrowing limit, \bar{a} , and the fixed hours cost, \bar{h} , so that in the steady state the (quarterly) rate of

²This is reasonable as long as other idiosyncratic shocks are not perfectly negatively correlated with idiosyncratic wage shocks.

³Note that all of the papers previously cited estimated shocks based on annual data, so that our benchmark values need to be converted to annualized values when comparing them to the literature. Our values correspond to the estimates in Floden and Linde (2001).

return to capital is 1%, the employment rate is 70%, the borrowing limit is equal to two quarters' earnings of a worker with productivity equal to the mean of all employed workers, and average hours (conditional on working) is $1/3$.⁴

For purposes of comparison it is of interest for us to also have an economy that does not feature an intensive margin. We do this by fixing h to be $1/3$ for all workers, and removing the parameter \bar{h} from the calibration exercise. All other aspects of the calibration exercise remain the same. In particular, the value of B will be adjusted so that steady state employment rate in this economy will also equal 70%. This alternative economy is a version of Chang and Kim (2007) and we will denote it as the Extensive Economy, abbreviated as "Ext".⁵

Table 1 summarizes the parameter values that are held constant across specifications as well as those that vary across specifications.

4. Properties of Steady State

In this section we consider some of the properties of the steady state equilibrium. There are two main objectives of this section. The first is to demonstrate that although our model is highly stylized, it is able to capture many features of the heterogeneity in wealth, earnings and hours worked found in the data. The second objective is to examine how different values for γ and σ_x influence the model's ability to account for these features in the data.

⁴With a quarterly employment rate of 70%, the average annual employment rate in our model (i.e., fraction of individuals who work at least one quarter during a year) is 76.7%. This corresponds to the average annual employment rate in the PSID over the period 1968-1998 for household heads and spouses with ages between 18 and 65. We use a cutoff of 240 annual hours as the threshold for employment, i.e., we treat individuals with less than 240 annual hours worked as not employed.

⁵There is one exception to note. When $\sigma_x = .330$ the natural borrowing constraint for the economy with only an extensive margin is tighter than the borrowing constraint used previously, so in this case the borrowing limit is set to zero. The borrowing constraint is tighter in an economy without an intensive margin since it precludes people in low productivity states with high debt from working more hours as a way to generate more income in these states.

4.1. The Hours Worked Distribution

As noted in the calibration section, all of our model specifications are calibrated so as to generate the same fraction of people employed and the same average hours for workers conditional on employment. In this subsection we examine the extent to which the model can account for the distribution of hours worked among workers, and how this distribution varies across the various model specifications. We begin with Table 2, which reports standard deviations of the steady state distribution of annual hours of work, conditional on working. We compute this measure at the annual level because it is not available at the quarterly level in the PSID. Our sample is all household heads and spouses between the ages of 18 and 65 during the period 1968-1997. In the data there are some individuals who work very few hours during the year. We therefore classify a worker as employed if he or she works at least 240 hours per year, and treat those with less than 240 annual hours as having zero hours.⁶ In the model, an individual is classified as employed if he or she has positive hours for at least one quarter during the year. While for some of the subsequent analysis we will utilize the panel nature of the PSID, this feature is not essential to this calculation, and so as a robustness check we also include a measure based on the CPS.⁷ We normalize annual hours in the CPS and PSID so that average annual hours is the same as in the economy with $\gamma = 1$ and $\sigma_x = .165$.⁸

The standard deviations in the PSID and CPS are fairly comparable – .418 in the former and .454 in the latter. Acknowledging that there is likely to be some measurement error in hours, the true dispersion in hours will be less than indicated by these values. The standard deviations for the 9 model specifications range from .301 to .480. Consistent with intuition the hours dispersion is increasing in γ and in σ_x . When $\sigma_x = .165$, the model cannot generate sufficient dispersion in hours even with a relatively large value of γ .

⁶This adjustment affects relatively few individuals and our results are not very sensitive to variation in this cutoff.

⁷Annual Hours for the CPS data is obtained by multiplying “Average Hours per Week” and “Number of Weeks Worked”.

⁸Average annual hours do not vary much across model specifications, so we do not renormalize for comparison with each specification.

However, when $\sigma_x = .330$ the model is able to match the dispersion found in the PSID as long as γ is around 1.0 or larger. As noted earlier, this motivates our choice of $\sigma_x = .330$ as a reasonable upper bound on the extent of heterogeneity in our model.⁹

To examine the implications for the distribution of hours worked in somewhat more detail, we next look at the average hours worked at various percentiles of the hours distribution. Figure 1 plots these values for two of the model specifications (the high and low values of σ_x with $\gamma = 1$) as well as the corresponding values for the PSID.

Consistent with the message based on looking simply at the standard deviations, we see that the specification with $\gamma = 1$ and $\sigma_x = .330$ does a much better job of tracking the empirical hours distributions than does the specification with the lower value of σ_x . In fact, this specification tracks the distribution in the PSID quite well.

4.2. Employment and Hours Transitions

Having assessed the model's ability to account for the distribution of individuals between employment and non-employment as well as the distribution of hours among employed workers, we next examine the model's ability to account for the movement of individuals within the hours worked distribution, including transitions into and out of employment.

Our data on transitions comes from the PSID and so are based on annual measures. We begin by looking at transitions into and out of employment. Table 3 shows the distribution of individuals across different combinations of employment states in consecutive years for the PSID and one of our model specifications ($\gamma = 1$ and $\sigma_x = .165$). We only report statistics for one of the model specifications because it turns out that these statistics are virtually identical across all 9 specifications.

The model does a good job of accounting for transitions into and out of employment, though employment status is somewhat less persistent in the model than in the data, as

⁹Table 2 might lead one to consider even higher values for σ_x combined with smaller values for γ . However, as we show in a later subsection, this leads to an estimated Frisch elasticity that is well below .40, and so we do not consider this region of parameter space.

evidenced by the fact that the model has greater share of workers changing employment status across consecutive years than does the PSID. However, this difference is relatively small. For example, in the PSID, the persistence of employment (i.e., the probability of being employed next year conditional on being employed this year) is .74 whereas this probability is .73 in the model.

Next we consider the transition matrices for annual hours worked between years t and $t+1$. Table 4 reports the transition rates between quintiles of the hours worked distribution (and non-employment) from the PSID and the two model specifications in which $\gamma = 1$ and the two extreme values of σ_x . The transition matrices for different values of γ holding σ_x constant turn out to be quite similar, so in the interest of space we do not include them. As we will see, although changes in σ_x do produce some quantitative differences, most notably in the degree of persistence, the basic patterns are also not affected much by changes in σ_x .

We start by noting three prominent features in the data. First, annual hours of work exhibit a significant degree of persistence, especially for those workers with high hours of work. All of the diagonal elements for workers with positive hours are greater than 40%, and for the highest quintile this value exceeds 60%. Second, conditional on working in both periods and switching quintiles, the transition probabilities are monotone decreasing in the distance of the destination quintile from the originating quintile. Third, individuals who adjust along the extensive margin between years t and $t + 1$ are disproportionately from the lowest quintile of the hours distribution. Specifically, for those individuals who work in year t but not in $t + 1$, roughly two-thirds of them have annual hours of work in the lowest quintile in year t . Similarly, for those individuals who did not work in year t but did work in year $t + 1$, roughly three-quarters of them have annual hours in the lowest quintile in year $t + 1$.

Next we consider how the model fares in terms of reproducing this features. As the table shows, both model specifications also generate considerable persistence, though less in the economy with $\sigma_x = .33$. The average of the diagonal elements for those with positive

hours in both periods is 49.48 in the data, 46.13 in the economy with $\sigma_x = .165$ and 38.48 in the economy with $\sigma_x = .33$. While both model specifications come close to matching the persistence in the highest quintile, the lowest quintile displays quite a bit more persistence in the data than in either model—49.45 versus 32.33 and 32.66. One possible explanation for this is the existence of a group of workers in the data who desire part-time work on a more permanent basis than captured by the idiosyncratic shocks in our model.

For the most part the model also matches the second observation noted above. Specifically, for the specification with $\sigma_x = .33$ the model has the same monotonicity property found in the data, whereas for the other specification though there are a couple of values which violate the pattern.

Finally, the model does a good job of matching the nature of adjustment along the extensive margin. For both of the model specifications shown in the table, roughly three quarters of all transitions along the extensive margin involve workers who are in the lowest quintile of the hours distribution.

4.3. Micro Frisch Elasticity

Next we examine what one would conclude about the Frisch elasticity in the steady state of our model based on running a standard labor supply regression. To do this we generate an artificial panel data from our model economy (50000 workers with 100 quarters), and aggregate it to the annual level. As is standard in the empirical labor supply literature, and in the spirit of Altonji (1986), we run the following regression:

$$\log h_{it} = b_0 + b_w \log w_{it}^h + b_c \log c_{it} + \varepsilon_{it}$$

using all observations with positive hours, where the individual wage rate is defined as earnings divided by hours, $w_{it}^h = w_t x_{it} g(h_{it}) / h_{it}$, conditional on working ($h_{it} > 0$). The resulting parameter estimate b_w is the so-called micro Frisch labor supply elasticity. If a worker with preferences as in our model were to make optimal choices facing a linear budget constraint, the first order condition would imply that b_w is equal to γ , so that under

these circumstances this regression can uncover the true value of the preference parameter γ . As is known from previous work, a modification of the standard labor supply model that dispenses with a linear budget constraint will break the close link between preference parameter γ and observed labor supply elasticity b_w (see, e.g., Blomquist (1983), Moffit (1983), Rogerson and Wallenius (2009), Chang et al. (2011)). Table 5 reports the results of running this regression in the various specifications of our model.

While the estimated values of b_w are increasing in the value of the preference parameter γ , there is a substantial discrepancy between the two. In the economies with $\sigma_x = .165$ the value of b_w is slightly less than 60% as large as γ .¹⁰ When $\sigma_x = .330$ the estimated values of b_w are roughly two-thirds of the value of γ . The two key messages from this table are first, that in this setting with non-linear wages, the estimated value of b_w is substantially lower than the underlying preference parameter γ , and second, that the extent of this discrepancy is significantly affected by the value of σ_x . As we have noted earlier, to the extent that a reasonable lower bound for this elasticity is in the interval $.40 - .50$, we think that $.50$ is a reasonable lower bound for values of γ .

4.4. Distributions of Wealth and Earnings

Chang and Kim (2007) and An et al. (2009) have previously shown that a model with idiosyncratic productivity shocks calibrated to micro data, incomplete markets and indivisible labor captures many quantitative features of the wealth and earnings distribution. It turns out that adding an intensive margin to the analysis has little impact along this dimension. As a result, when $\sigma_x = .165$ our wealth and earnings distributions look very similar to those in An et al. (2009). When we consider the specification with a much greater degree of heterogeneity, $\sigma_x = .330$, we get similar patterns qualitatively, but the model generates too much dispersion in earnings. In this section we document these properties.

¹⁰This result is similar to the finding in Rogerson and Wallenius (2009), though their result was based on a slightly different labor supply regression and their calibration was somewhat more stylized.

Given that we calibrate our model using employment data from the PSID we think it is preferable to compare our model to data that is also based on the PSID. For this reason our primary source of information on the cross-sectional wealth and earnings measures are based on the 1984 PSID. As a robustness check we also report comparable figures for properties of the wealth distribution from the work of Diaz-Gimenez et al. (1997) that is based on the Survey of Consumer Finances (SCF). For the measures that we focus on the two data sets provide very similar answers, so this does not seem to be a major issue.¹¹

Table 6 reports the Gini coefficients for both the wealth and earnings distributions in the eight of the different model specifications that we consider, as well as their corresponding values in the PSID and SCF.

A few patterns are evident. First, given a value for σ_x , the Gini coefficients for both the wealth and earnings distributions are (weakly) increasing in the value of γ . (Note that the Extensive Only case can be thought of as the limiting case as γ goes to zero.) This effect is intuitive; a higher value of γ leads to greater intertemporal substitution of labor supply, so that individuals work more when productivity is high and less when productivity is low, thereby amplifying the direct effect of productivity on earnings. Given that individuals accumulate assets to smooth consumption in the face of fluctuations in earnings, greater fluctuations in earnings leads to greater dispersion in assets. Although the qualitative effects are intuitive, the main message from Table 3 is that the quantitative importance of these effects are quite small. While moving from $\gamma = 1.0$ to the extensive only case does generate modest changes in both Gini coefficients, the effect of changes in γ inside the interval $[\.5, 1.5]$ is of second order importance for each measure.

It is also intuitive that higher values of σ_x would similarly lead to increases in the Gini coefficients for both wealth and earnings distributions. However, in contrast to the previous case, changes in σ_x for a given value of γ do generate first order effects on both measures, with the effect on the earnings Gini being almost twice as large as the effect on

¹¹This is not the case if one focuses on the extreme upper tail of the wealth distribution, as it is well known that the PSID undersamples the upper extremes of the wealth distribution. However, as noted before, given our emphasis on labor supply, this extreme upper tail is not of primary concern.

the wealth Gini.

Comparing the values in the various model specifications with their counterparts in the data, all of the model specifications seem to capture much of dispersion in the wealth and earnings distributions. If anything, the models generate too much dispersion in earnings, especially in comparison to what is found in the PSID. To look a bit deeper into the nature of the wealth and earnings distributions, Table 7 shows the wealth and earnings shares by quintiles of the wealth distribution. Because variation in γ turns out to be not very important quantitatively in terms of these distributions, we only report results for the two model specifications with the extreme values of σ_x and with $\gamma = 1$.

The basic message from Table 7 is that in addition to doing a reasonable job of accounting for the absolute amount of dispersion in wealth and earnings as captured by the Gini coefficient, the model also does a good job of accounting for the shape of these distributions. The specification with $\sigma_x = .165$ does a very good job of capturing the earnings shares, whereas consistent with Table 2, the specification with $\sigma_x = .330$ generates excessive concentration of earnings in the upper quintile of the wealth distribution. However, the specification with the higher value of σ_x is better able to capture the amount of wealth concentrated in the upper quintile. Analyzing the wealth shares by quintiles of the wealth distribution hides the extreme concentration of wealth at the very top of the distribution. It is well known (see, for example, Diaz-Gimenez et al (1997)) that the model is not able to account for the concentration found in say the upper 1% of the wealth distribution. However, from the perspective of labor supply, accounting for the likes of Bill Gates is probably not of first order importance, and so we do not focus on the extreme upper part of the wealth distribution.

In summary, this subsection shows that all of the model specifications generate significant dispersion in earnings and wealth relative to the data. If anything, some of the model specifications generate too much dispersion. The nature of the dispersion is also empirically reasonable, in terms of matching earnings and wealth shares by quintiles of the wealth distribution. We conclude that adding an intensive margin of labor supply to the

previous analyses of Chang and Kim (2006, 2007) does not have first order effects on the earnings and wealth distributions.

5. Business Cycles

In this section we study the business cycle properties of our model economies. We solve the equilibrium of the model using the method proposed by Krusell and Smith (1998). Details are included in the appendix. In this section we focus on aggregate statistics, and since aggregate data is available at quarterly frequency, we also compute model statistics using quarterly data. As is standard, we take logs and then HP filter (with the smoothing parameter of 1,600) the simulated series before computing statistics.

Table 8 reports the properties of output and labor market variables from our models. Aggregate hours worked, employment (extensive margin), hours per worker (intensive margin) and aggregate efficiency units of labor services are denoted as H , E , h , and L respectively. By definition, $H = E \times h$. Because the behavior of consumption and investment is basically the same as in standard real business cycle style exercises we do not report statistics for these variables. When reporting statistics from the US data we report measures based on both the Establishment Survey (ES) and the Household Survey (HS). The statistics for the US economy are taken from Prescott and Cooley (1998).

We start by noting two properties from the data. First, total hours, whether measured by the household or the establishment survey, are almost as volatile as output. Second, fluctuations in total hours are dominated by fluctuations in employment. According to the Establishment Survey, employment is almost three times as volatile as hours per worker, whereas according to the household survey this ratio is roughly two. It is relevant to note that the intensive and extensive margins from the establishment and household surveys are not necessarily comparable. The establishment survey is based on payroll positions, whereas the household survey is based on individuals. To the extent that some individuals hold multiple jobs these surveys will provide different measures of employment and hours

volatility. We do not take a stand on which of these is “preferable”; we will simply use as a reference point for the fact that employment is two to three times as volatile as hours per worker.

Next we turn to the properties of fluctuations in the models. We emphasize that in all of these model economies the driving force behind aggregate fluctuations is identical, i.e., the parameters of the technology shock process are held constant across all specifications. Hence, to the extent that fluctuations are different in the various model economies, it is the result of how the different models lead to different propagation of these shocks. While our focus is on labor market variables, we first note that consistent with many previous exercises, the technology shocks that we feed into our model generate output fluctuations that are between two-thirds and three-quarters of observed fluctuations in output. We will say more about the nature of these differences below, when we examine the nature of labor market fluctuations in more detail.

Rather than systematically walking the reader through the results in Table 8, we think it is useful to organize our discussion around several messages that we want the reader to take away from the table. The first message concerns the determinants of fluctuations along the intensive and extensive margins. In particular, we start from the notions, implicit in much of the recent literature, that fluctuations along the intensive margin are determined mostly (if not exclusively) by the value of the preference parameter γ , whereas fluctuations along the extensive margin are determined mostly (if not exclusively) by the extent of heterogeneity, which in our model is reflected in the value of σ_x . Intuitively, fluctuations along the intensive margin are increasing in the value of γ , while fluctuations along the extensive margin are decreasing in the value of σ_x . Additionally, intuition based on simple models would suggest that fluctuations along the intensive margin are proportional to the value of γ .

While the results in Table 8 provide partial support for both of these notions, they also reveal that these notions reflect an oversimplification that is potentially very misleading from a quantitative perspective. For example, consistent with the first notion above, we see

that for a given value of σ_x , increases in γ lead to greater fluctuations along the intensive margin. However, the notion that γ is the dominant, let alone the exclusive factor that determines this response is strongly contradicted by the results in Table 8. Specifically, starting from the specification in which $\gamma = 0.5$ and $\sigma_x = .165$, we see that increasing γ to 1.5 leads to almost a tripling of the response along the intensive margin. However, this same effect occurs if we keep the value of γ fixed at 0.50 but we instead increase the value of σ_x to .330. To the best of our knowledge, we are the first to point out that the aggregate fluctuations along the intensive margin are affected in a quantitatively critical way by the amount of heterogeneity in the economy.

Similarly, starting from the specification $\gamma = .5$ and $\sigma_x = .165$, we note that moving to the specification with $\sigma_x = .33$ leads to roughly a 50% reduction in fluctuations along the extensive margin. However, if we instead kept the value of σ_x unchanged at .165 and instead increased the value of γ to 1.5, we would still have a decrease in fluctuations along the extensive margin of almost 20%.

The key message that emerges from this exercise is that fluctuations along the intensive and extensive margin are jointly determined in a quantitatively important fashion by both the value of the preference parameter γ and the extent of heterogeneity, as captured here by the value of σ_x . It follows that one should not think of “intensive margin elasticities” and “extensive margin elasticities” as two independent parameters.

A second and related message is that pure indivisible labor models are a poor guide for assessing the impact of heterogeneity on fluctuations in aggregate hours. Table 8 shows that the drop in the volatility of aggregate hours in the extensive only economies as σ_x increases is much greater than in the economies with $\gamma = 0.5$.

The third message concerns the implications of γ and σ_x for fluctuations in aggregate hours. We begin by looking at how the results vary as we change γ holding the value of σ_x fixed at .165. Keeping in mind that the aggregate shocks are constant across all specifications, and that each specification matches both the average employment rate and the average level of hours per worker, the first striking result is the huge variation in

the volatility of aggregate hours across these four specifications, ranging from .67 to .93. However, despite these large differences in the volatility of aggregate hours, Table 8 also shows that differences in the volatility of labor services are roughly an order of magnitude smaller, ranging from .85 to .89. Because of this, differences in output fluctuations are also relatively minor. Before we explain the reason for these results, we note that a key message is that the split of aggregate fluctuations into intensive and extensive margins has first order implications for fluctuations in total hours but not necessarily for output.

To understand the patterns just discussed, it is useful to start with the observation that as we increase γ holding σ_x constant, one would intuitively expect an increase of the elasticity of aggregate labor supply and therefore an increase in the volatility of hours. This intuition does show up in the results if we apply it to labor services, i.e., the volatility of labor services increases as we increase γ . This raises the question of why hours do not exhibit the same behavior? The key to understanding this discrepancy lies in understanding the role of heterogeneity in individual productivity and how this matters for aggregate fluctuations. To answer this question, we start with the following exercise. Consider the response of the economy with $\gamma = 1.5$ and $\sigma_x = .165$ to a positive aggregate shock, and consider how this response changes as we decrease γ . Intuitively, decreasing γ leads to smaller increases along the intensive margin. One piece of intuition for the results comes from asking what would happen to fluctuations in aggregate hours of work if we were to compensate for this decrease in fluctuations along the intensive margin with sufficiently greater fluctuations along the extensive margin so as to keep the total increase in labor services constant. To the extent that the employed are on average of higher productivity than the non-employed, when we substitute increases along the extensive margin for increases along the intensive margin, the hours are on average less productive. It follows that it would take a larger increase in hours along the extensive margin in order to perfectly compensate for the smaller increase in hours per worker. In fact, intuition suggests that the overall increase in labor services should be lower as we decrease γ . To see why, note that if the change in labor services is unchanged, then so is the change in the wage. But in

order to get an additional increase along the extensive margin when we lower γ will require a larger increase in the wage rate, i.e., a smaller increase in labor services.

The net prediction is that the increase in labor services should decrease as we decrease γ , while the change in hours is ambiguous since there are two opposing effects. However, to the extent that the change in the response of labor services is not too large, we would expect that the net effect is for the change in hours to increase. To pursue this a bit further, the extent to which labor services increase by a smaller amount is dictated by the density of individuals who are willing to move into employment in response to small increases in the wage. The more compressed is the distribution of individual heterogeneity, the more likely it is that a large mass of individuals are willing to enter employment in response to small increases in the wage.

With this intuition in mind, we next turn to the results for the case of $\sigma_x = .330$. While the basic economics are the same for the different values of σ_x , the higher value of σ_x is associated with a less compressed distribution of (the exogenous component of) individual heterogeneity. Based on the above intuition, we would expect this to show up as a more pronounced effect of γ on the magnitude of fluctuations in labor services. Indeed, we observe this in Table 8, with the range now being $.46 - .61$. And, given the larger effects on fluctuations in labor services, we also see larger differences in hours.

The other pattern that is evident is the decrease in fluctuations in both aggregate hours and labor services as we increase the value of σ_x for a given value of γ . This is consistent with the intuition expressed earlier—the greater dispersion in individual heterogeneity leads to lower responses along the extensive margin. Although this results in greater fluctuations along the intensive margin, the curvature in the disutility of work function implies that this effect only partially offsets the lower fluctuations on the extensive margin.

There are some additional interesting patterns across the specifications. For example, fixing $\sigma_x = .165$, we see that changes in the value of γ within the interval $[.5, 1.5]$ has virtually no impact on the magnitude of fluctuations in aggregate hours, though going to the extreme case of no adjustment along the intensive margin leads to a very large increase

in the standard deviation of aggregate hours. That is, for these specifications, a tripling of the willingness of individuals to substitute work effort over time, in the form of an increase in γ from .5 to 1.5, has virtually no impact on the actual fluctuations in aggregate hours of work. However, Table 8 shows that this result breaks down if we consider a specification with substantially more heterogeneity. Specifically, in the specifications with $\sigma_x = .330$ we see that increasing γ from .50 to 1.5 results in roughly a twenty percent increase in the size of fluctuations in aggregate hours.

Lastly, we assess the extent to which the various specifications are able to account for the fact that over the business cycle, employment fluctuates between two and three times as much as hours per worker. With the lower degree of heterogeneity, i.e., $\sigma_x = .165$, achieving this ratio would require a value of γ of at least 1.5, since this implies that employment is three times as volatile as hours per worker and this ratio is decreasing in γ . When $\sigma_x = .165$ and $\gamma = 0.50$, this ratio exceeds ten. With the higher degree of heterogeneity, this range is achieved for γ somewhat less than .50, since when $\gamma = .50$ and $\sigma_x = .330$, employment fluctuations are a bit less than twice those of hours per worker.

6. Taxes

In this section we consider the effect of steady state changes in fiscal policy in the various specifications. Given our focus on labor supply responses, we choose to focus on a single fiscal policy for which labor supply responses are known to be of first order importance: a proportional tax on labor earnings used to finance a lump-sum transfer. Specifically, in each of the models we solve for the effect on steady state of instituting a 10% proportional tax on labor earnings used to finance a lump-sum transfer ($\tau = 0.1$).

Recalling that in all of our economies the initial steady state equilibrium has the same values for both the employment to population ratio and the average hours of work per employed worker. Table 9 shows the effects on various labor market outcomes.

There are two patterns worth noting. One concerns the effects of increases in σ_x for a

fixed value of γ , while the second concerns the effects of increases in γ for a fixed value of σ_x . Before presenting these patterns it is useful to articulate two key mechanisms at work in steady state equilibrium that lay behind optimal individual labor supply decisions.

One key mechanism at work at the individual level in the steady state equilibrium of this economy is intertemporal substitution of labor supply. With complete markets and no uncertainty, individuals would arrange their labor supply so that they work only when productivity is above some threshold, and hours when working would be increasing in productivity above this threshold. Both the threshold for working versus not working and the slope of the hours profile reflect intertemporal substitution. Intuitively, by having a steeper slope to the profile of hours versus productivity the individual would generate more income while working and could therefore increase the productivity threshold for working. The intertemporal trade off between changes in hours of work while working and the fraction of periods in employment is influenced both by the preference parameter γ and the mass of states that are near the reservation productivity level.

However, the presence of uncertainty and incomplete markets adds another aspect to the analysis. Specifically, because the individual cannot perfectly insure against idiosyncratic shocks, optimal choices will entail an individual working in some low productivity states as a way to help smooth consumption. In these states, the individual's labor supply problem is effectively a static problem, and in a static setting we know that a tax and transfer system decreases the incentives to work. The lower the level of productivity, the greater is the income effect of the transfer payment, making it more likely that low productivity individuals choose to not work. Intuitively, it turns out that this phenomenon is more present in the economies that feature greater heterogeneity, since greater an increase in σ_x implies greater uncertainty, thereby making it harder for individuals to self-insure.

With these mechanisms in mind, we turn to the patterns in Table 9. Note first that increasing the value of σ_x for a fixed value of γ leads to a greater response along both the extensive and intensive margins, and therefore a greater response in aggregate hours. This pattern is in contrast to our findings in the business cycle context and at first pass

may seem counterintuitive. Specifically, in the previous section we argued that an increase in heterogeneity led to less mass in the neighborhood of the reservation wage curve and therefore a smaller response along the extensive margin. The second mechanism noted above—the presence of low productivity workers with positive hours—is key to understanding the employment response. More specifically, the greater the dispersion in idiosyncratic shocks creates a larger mass of individuals who are working despite having low productivity, thereby creating a large negative response in employment in response to an increase in the income transfer.¹²

Why is there also a greater decrease along the intensive margin as we increase the value of σ_x ? This response is basically what one would predict from intertemporal substitution considerations. A tax and transfer system basically has the government carrying out part of the intertemporal substitution that the individual wants to carry out, in the sense that the net transfer from the government is relatively high when individual productivity is low. As a result the individual will do less intertemporal substitution. This can happen along the extensive margin or the intensive margin. With more heterogeneity there is less scope for intertemporal substitution along the extensive margin, similar to what we found in the analysis of business cycles. As a result, we get a greater response along the intensive margin as we increase the value of σ_x .

The second pattern concerns the effect of an increase in the value of γ for a given value of σ_x . Rogerson and Wallenius (2009) effectively considered an experiment of this sort, and found that the first order effect was to increase the response along the intensive margin, decrease the response along the extensive margin, but to leave the change in aggregate hours relatively unchanged. Rogerson and Wallenius carried out their exercise in a setting with a relatively small degree of heterogeneity and no uncertainty. The results in Table 9 show that these results do not necessarily carry over to settings with higher degrees of

¹²The reader may wonder why this mechanism was not evident in the business cycle analysis. Key to understanding the difference is that the two driving forces are different. As noted above, these individuals effectively solve a static problem. In a static model with balanced growth preferences, an increase in the wage has no effect on labor supply, whereas an increase in the size of a tax and transfer program has a negative effect on labor supply.

heterogeneity.

In particular, when $\sigma_x = .165$ our results exactly mirror those found in Rogerson and Wallenius, in that we find that the change in aggregate hours is almost independent of the value of γ .¹³ Interestingly, however, for higher values of σ_x the aggregate response is not independent of γ as γ increases from .5 to 1.5. To the extent that increases in γ represent a “more elastic” labor supply, one would intuitively expect that increasing γ would lead to a greater response in aggregate hours, which is what happens as we increase γ from .5 to 1.5. Additionally, the size of this effect is increasing in the value of σ_x . These effects can be understood through the manner in which σ_x influences the relative efficacy of intertemporal substitution along the intensive and extensive margins. To see why it is useful to begin with an analysis of the $\sigma_x = .165$ case. When $\sigma_x = .165$, we see this tradeoff at work as we move to higher values of γ , in the sense that the greater decrease along the intensive margin is accompanied by a smaller decrease on the extensive margin. Moreover, the two effects are roughly offsetting. As we consider higher values of σ_x we see the same pattern for the intensive margin, in that the responsiveness of hours per worker increases as we increase γ . However, what differs as we consider higher values of σ_x is that the opposing effect on employment becomes much smaller. The reason that the employment response becomes smaller is precisely due to the effects that we have emphasized earlier: with a higher value of σ_x there is less mass of “states” near the reservation curve and as a result there is less scope for intertemporal substitution along the extensive margin. In fact, when $\sigma_x = .33$, we see that the response along the extensive margin is effectively the same for all three values of γ . One additional factor that influences the statistics in this case is composition effects on average hours associated with changes on the extensive margin. Specifically, since individuals who are on the margin of working versus not-working for purposes of intertemporal substitution are working fewer hours, when these individuals

¹³Note that we found a similar result for the case of business cycle fluctuations. Interestingly, for higher degrees of heterogeneity, the business cycle and steady state tax effects are of opposite signs, with increases in σ_x leading to smaller fluctuations in hours in the business cycle setting but greater changes in the steady state tax and transfer setting.

drop out there is an added increase in hours per worker for employed individuals. This further explains why the drop along the intensive margin increases as the employment drop decreases.

We note two additional points from Table 9. First, the employment response in the extensive only model is quite similar to the employment response in the $\gamma = .50$ economies. Second, it is striking that the variation in the response of aggregate labor services is much smaller than the variation in aggregate hours.

In summary, we again find that in general responses along both the intensive and extensive margin are influenced in a non-trivial way by both the preference parameter γ and the extent of heterogeneity as captured by the parameter σ_x . For relatively large values of σ_x we did find that the response along the extensive margin was basically unaffected by the value of γ . However, changes along the intensive margin are affected by both parameters.

7. Conclusion

Recent advances in modeling aggregate labor supply have emphasized the importance of accounting for adjustment along the intensive and extensive margins. Adjustment along the extensive margin has also been shown to depend on the extent of heterogeneity. In this paper we build a model in which individual labor supply adjusts along both the intensive and extensive margins in an environment that features heterogeneity and incomplete markets. We believe that this is the appropriate benchmark model for understanding the joint determination of adjustment along the two margins. We consider a family of specifications of this model that differ along two key dimensions: the value of the preference parameter that influences curvature of utility in hours of work, and the standard deviation of innovations in the idiosyncratic shock process, which in turn influences the extent of heterogeneity in the invariant distribution for idiosyncratic shocks.

We consider the ability of the various specifications of the model to account for key

features of employment and hours worked in the cross-section. We then use this model to consider labor supply responses to temporary shocks and permanent tax changes, along both the intensive and extensive margins. Three key findings emerge. First, extensive and intensive margin elasticities are jointly determined by both the preference parameter and the extent of heterogeneity. That is, one cannot speak of intensive and extensive margin elasticities as independent parameters of the economic environment. Second, the effect of increased heterogeneity on extensive margin elasticities is of opposite sign in the two contexts: in the business cycle setting increased heterogeneity leads to less adjustment along the extensive margin, whereas in the case of permanent changes in the size of the tax and transfer system we find that increased heterogeneity leads to more adjustment along the extensive margin. Third, in terms of fluctuations in aggregate hours, we find that abstracting from the intensive margin can be very misleading about the effect of heterogeneity. Fourth, although the various specifications that we consider generate very large differences in responses in aggregate hours, the associated responses in aggregate labor services is considerably smaller.

References

- [1] Aiyagari, S. Rao (1994): “Uninsured Idiosyncratic Risk and Aggregate Saving,” *Quarterly Journal of Economics*, 109(3), 659-684.
- [2] Altonji, Joseph G. (1986) ”Intertemporal Substitution in Labor Supply : Evidence from Micro Data.”, *Journal of Political Economy*, 94(3), s176-215.
- [3] An, Sungbae, Yongsung Chang, and Sun-Bin Kim (2009): “Can a Representative Agent Model Represent a Heterogeneous Agent Economy?” *American Economic Journal - Macroeconomics*, 1, 29-54.
- [4] Bewley, Truman (1986) ”Stationary Monetary Equilibrium with a Continuum of Independently Fluctuating Consumers,” in *Contributions to Mathematical Economics in Honor of Gerard Debreu*.
- [5] Blomquist, Soren (1983) ”The Effect of Income Taxation on the Labour Supply of Married Men in Sweden.” *Journal of Public Economics*, 22, 169-197
- [6] Chang, Yongsung, and Sun-Bin Kim (2006): “From Individual to Aggregate Labor Supply: A Quantitative Analysis Based on a Heterogeneous Agent Macroeconomy,” *International Economic Review*, 47(1), 1-27.
- [7] Chang, Yongsung, and Sun-Bin Kim (2007): “Heterogeneity and Aggregation: Implications for Labor-Market Fluctuations,” *American Economic Review*, 97(5), 1939-1956.
- [8] Chang, Yongsung, Sun-Bin Kim, Kyooho Kwon and Richard Rogerson (2011): “Interpreting Labor Supply Regressions in a Model of Full- and Part-Time Work,” *American Economic Review: Papers & Proceedings* , 101(3), 476-481.
- [9] Cho, J.-O. and R. Rogerson (1988) “Family Labor Supply and Aggregate Fluctuations,” *Journal of Monetary Economics* 21, 233-245.

- [10] Cho, J.-O. and Cooley, Thomas F. "Employment and hours over the business cycle", *Journal of Economic Dynamics and Control* 18 (1994), 411-432.
- [11] Cooley, Thomas F. *Frontiers of Business Cycle Research*, Princeton University Press 1998
- [12] Erosa, A., L. Fuster and G. Kambourov (2011) "Towards a Micro-Founded Theory of Aggregate Labor Supply," mimeo, University of Toronto.
- [13] Floden, Martin, and Jesper Linde, (2001) "Idiosyncratic Risk in the United States and Sweden: Is There a Role for Government Insurance?," *Review of Economic Dynamics*, 4 (2), 406-437.
- [14] French, Eric (2005) "The Effects of Health, Wealth, and Wages on Labor Supply and Retirement Behavior," *Review of Economic Studies*, 72(2), 395-427.
- [15] Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L. Violante (2008): "Insurance and Opportunities: A Welfare Analysis of Labor Market Risk," *Journal of Monetary Economics*, Vol. 55, pp. 501–525.
- [16] Hugget, M (1996) "Wealth Distribution in Life-Cycle Economics," *Journal of Monetary Economics* 38, 469-94.
- [17] Krusell, Per, and Anthony A. Smith (1998) "Income and Wealth Heterogeneity in the Macroeconomy", *Journal of Political Economy*, 89(6), 1059-1085.
- [18] Kydland, Finn E. and Edward C. Prescott (1982), "Time to build and aggregate fluctuations," *Econometrica*, 50, 1345-79.
- [19] Kydland, Finn E. and Edward C. Prescott (1991), "Hours and employment variation in business cycle theory," *Economic Theory*, 1, 63-81.
- [20] MaCurdy, Thomas E. (1981): "An Empirical Model of Labor Supply in a Life-Cycle Setting," *Journal of Political Economy*, 106(5), 867-96.

- [21] Moffitt, Robert (1984) “The Estimation of a Joint Wage-Hours Labor Supply Model,” *Journal of Labor Economics* 2, 550-566.
- [22] Rogerson, Richard (1988): “Indivisible Labor, Lotteries and Equilibrium,” *Journal of Monetary Economics*, 21(1), 3-16.
- [23] Rogerson, Richard, and Wallenius, Johanna (2009). “Micro and Macro Elasticities in a Life Cycle Model with Taxes,” *Journal of Economic Theory* 144, 2277-2293.
- [24] Prescott, Edward, Rogerson, Richard, and Wallenius, Johanna (2009). “Lifetime aggregate labor supply with endogenous workweek length,” *Review of Economic Dynamics* 12, 23-36.

Table 1: Calibrated Parameter Values

A. Values Held Constant Across Specifications					
α	δ	ρ_λ	σ_λ	ρ_x	
0.64	0.025	0.95	0.007	0.975	
B. Values That Vary Across Specifications					
		β	B	\bar{h}	\bar{a}
$\sigma_x = .165$	$\gamma = 1.5$.9770	10.72	.1132	-.6029
	$\gamma = 1.0$.9773	18.47	.1490	-.5012
	$\gamma = 0.5$.9776	82.70	.2099	-.3328
	Ext	.9792	18.05	-	-.4741
$\sigma_x = .2475$	$\gamma = 1.5$.9698	9.32	.0945	-1.0899
	$\gamma = 1.0$.9702	15.93	.1318	-.9075
	$\gamma = 0.5$.9708	70.67	.1968	-.6041
	Ext	.9734	14.75	-	-.8230
$\sigma_x = .330$	$\gamma = 1.5$.9640	8.03	0.070	-2.2845
	$\gamma = 1.0$.9645	13.50	.1074	-1.9170
	$\gamma = 0.5$.9654	58.63	.1777	-1.2876
	Ext	.9678	11.49	-	0.0

Notes: 'Ext' denotes the model specification with the extensive margin only.

Table 2: Standard Deviation of Annual Hours Conditional on Working

<u>Data</u>		<u>Models</u>								
		$\sigma_x = .165$			$\sigma_x = .2475$			$\sigma_x = .330$		
PSID	CPS	$\gamma=1.5$	$\gamma=1$	$\gamma=.5$	$\gamma=1.5$	$\gamma=1$	$\gamma=.5$	$\gamma = 1.5$	$\gamma = 1$	$\gamma = .5$
.418	.454	.337	.318	.301	.379	.347	.314	.480	.413	.342

Notes: Annual hours in the CPS and PSID are normalized so that average annual hours is the same as in the economy with $\gamma = 1$ and $\sigma_x = .165$.

Table 3: Distribution of Employment Transitions

E_{t-1}	E_t	Shares in PSID	Shares in Model
1	1	.740	.727
1	0	.046	.053
0	1	.032	.053
0	0	.182	.168

Notes: Employment status at time t (E_t), is denoted by 1 (working) or 0 (not working).

Table 4: Transition Probability of Hours (Annual)

PSID		$t + 1$					
		Not Work	1st	2nd	3rd	4th	5th
t	Not Work	83.57	12.25	1.69	0.91	0.99	0.60
	1st	21.08	49.45	14.91	6.15	5.29	3.12
	2nd	4.77	15.40	45.77	18.27	11.15	4.63
	3rd	2.81	6.75	19.77	46.24	17.88	6.54
	4th	2.26	5.14	10.63	19.91	42.42	19.64
	5th	1.81	3.41	4.69	6.80	19.77	63.52
	$\sigma_x = 0.165$						
t	Not Work	76.12	18.62	3.64	0.28	0.92	0.43
	1st	25.33	32.33	23.05	3.30	9.33	6.66
	2nd	4.48	15.11	35.75	20.77	16.88	7.02
	3rd	0.42	2.24	13.86	62.16	17.95	3.37
	4th	1.92	11.79	14.05	12.04	39.22	20.98
	5th	1.54	12.28	8.15	1.56	15.29	61.17
	$\sigma_x = 0.330$						
t	Not Work	76.96	18.32	3.85	0.65	0.17	0.04
	1st	25.30	32.66	25.75	10.43	3.89	1.97
	2nd	5.32	21.01	32.37	25.91	10.07	5.32
	3rd	1.44	9.97	19.25	30.65	26.11	12.58
	4th	0.54	5.24	9.49	19.89	40.82	24.03
	5th	0.26	4.98	7.49	12.42	18.93	55.91

Table 5: Estimated Values of b_w based on the labor supply regression.

	$\gamma = 1.5$	$\gamma = 1.0$	$\gamma = 0.5$
$\sigma_x = .165$.87	.57	.28
$\sigma_x = .2475$.93	.60	.29
$\sigma_x = .330$	1.02	.66	.32

Table 6: Gini Coefficients for Wealth and Earnings Distributions

	Data		Models							
	PSID	SCF	$\sigma_x = .165$				$\sigma_x = .330$			
			$\gamma=1.5$	$\gamma=1$	$\gamma=.5$	Ext	$\gamma=1.5$	$\gamma=1$	$\gamma=.5$	Ext
Wealth	.76	.78	.74	.74	.73	.70	.81	.80	.80	.78
Earnings	.53	.63	.64	.63	.62	.58	.84	.83	.82	.78

Table 7: Wealth and Earnings Shares by Wealth Quintiles

		1^{st}	2^{nd}	3^{rd}	4^{th}	5^{th}
Wealth Share	PSID	-0.52	0.50	5.06	18.74	76.22
	SCF	-0.39	1.74	5.72	13.43	79.49
	$\sigma_x = .165$	-1.18	-0.48	5.76	20.55	75.35
	$\sigma_x = .330$	-1.06	0.10	3.34	12.94	84.68
Earnings Share	PSID	7.51	11.31	18.72	24.21	38.23
	SCF	7.05	14.50	16.48	20.76	41.21
	$\sigma_x = .165$	6.24	14.00	17.85	23.69	38.21
	$\sigma_x = .330$	1.06	3.87	8.33	17.96	68.78

Table 8: Business Cycle Fluctuations

		σ_Y	σ_H	σ_E	σ_h	σ_L
Data	ES	1.72	1.69	1.41	.48	-
	HS	1.72	1.59	1.14	.63	-
$\gamma = 1.5$	$\sigma_x = .330$	1.28	.56	.26	.34	.61
	$\sigma_x = .2475$	1.35	.60	.35	.26	.71
	$\sigma_x = .165$	1.46	.68	.51	.17	.89
$\gamma = 1.0$	$\sigma_x = .330$	1.24	.49	.27	.24	.54
	$\sigma_x = .2475$	1.31	.54	.36	.19	.65
	$\sigma_x = .165$	1.44	.66	.54	.12	.87
$\gamma = 0.5$	$\sigma_x = .330$	1.19	.47	.30	.17	.46
	$\sigma_x = .2475$	1.27	.47	.36	.11	.57
	$\sigma_x = .165$	1.43	.67	.61	.06	.85
Ext	$\sigma_x = .330$	1.01	.30	.30		.16
	$\sigma_x = .2475$	1.13	.49	.49	0	.38
	$\sigma_x = .165$	1.39	.93	.93	0	.85

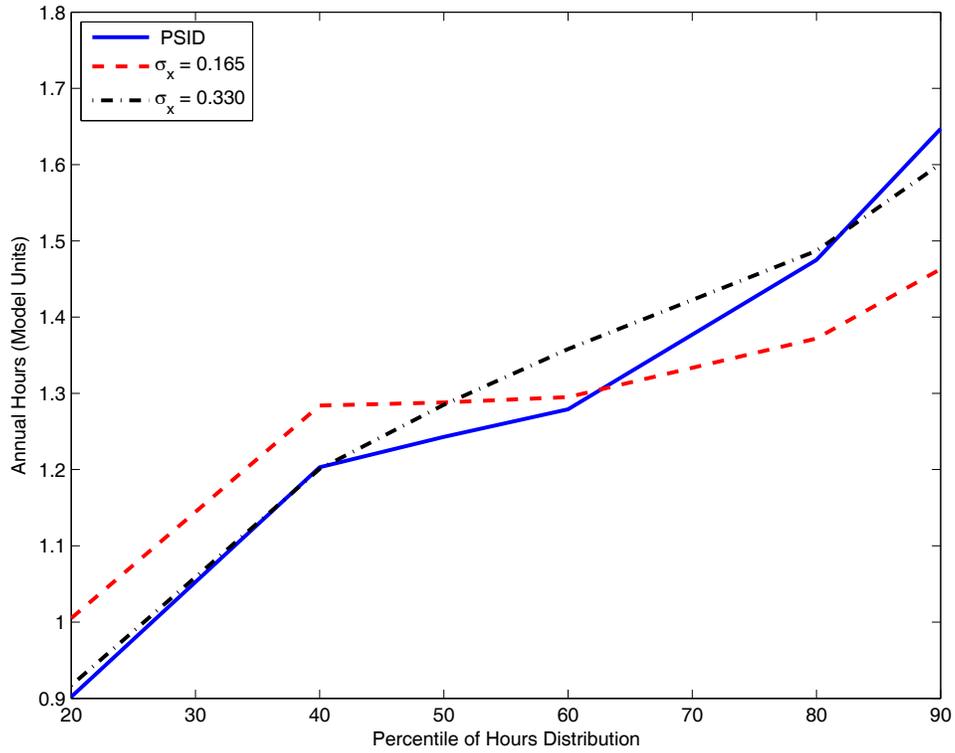
Note: Total hours (H) = Employment (E) \times Hours per worker (h). The variable L denotes labor hours in efficiency units. “ES” and “HS” denotes the data based on the Establishment Survey and Household Survey, respectively.

Table 9: Effects of Taxes on Steady State

		H	E	h	L
$\gamma = 1.5$	$\sigma_x = .330$.671	.774	.867	0.940
	$\sigma_x = .2475$.774	.843	.919	0.923
	$\sigma_x = .165$.866	.911	.951	0.921
$\gamma = 1.0$	$\sigma_x = .330$.695	.772	.900	0.944
	$\sigma_x = .2475$.788	.838	.941	0.932
	$\sigma_x = .165$.870	.901	.966	0.923
$\gamma = 0.5$	$\sigma_x = .330$.727	.770	.944	0.948
	$\sigma_x = .2475$.801	.827	.967	0.935
	$\sigma_x = .165$.870	.885	.983	0.929
Ext	$\sigma_x = .330$.722	.722	1.00	0.967
	$\sigma_x = .2475$.786	.828	1.00	0.951
	$\sigma_x = .165$.835	.835	1.00	0.929

Note: Total hours (H) = Employment (E) \times Hours per worker (h). The variable L denotes labor hours in efficiency units.

Figure 1: Distribution of Hours Worked



Notes: Annual hours in the PSID are normalized so that average annual hours is the same as in the economy with $\sigma_x = .165$.