

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

Yongsung Chang
Seoul National University

Sun-Bin Kim
Yonsei University

Kyooho Kwon
Korea Development Institute

Richard Rogerson
Princeton University
NBER

August 19, 2018

Abstract

We study business cycle fluctuations in heterogeneous-agent general equilibrium models that feature both intensive and extensive margins of labor supply. A nonconvexity in the mapping between time devoted to work and labor services combined with idiosyncratic shocks generates operative 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 salient features of the empirical distribution of hours worked, including how individuals transit within this distribution. We then study how the various specifications influence labor supply responses to aggregate technology shocks. We ask to what extent our predictions for business cycle fluctuations are affected by abstracting from the intensive margin and instead assuming that adjustment occurs only along the extensive margin. We find that abstracting from intensive margin adjustment can have large effects on the volatility of aggregate hours even if fluctuations along the intensive margin are small.

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

*Correspondence: Chang: School of Economics, Seoul National University, yongsung.chang@gmail.com; Kim: Department of Economics, Yonsei University, sunbin.kim@yonsei.ac.kr; Kwon: Korea Development Institute, kwonkh98@gmail.com; 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 SED Cyprus meeting. This work is supported by grants from the National Research Foundation of Korea funded by the Korean Government (NRF-2014S1A5A2A01011108).

1. Introduction

Early representative household models used to study business cycle fluctuations (e.g., Kydland and Prescott (1982)) assumed a household that smoothly adjusts labor supply. An immediate controversy revolved around the fact that these models were calibrated so as to give the representative household a Frisch labor supply elasticity larger than those estimated on micro data by researchers such as MaCurdy (1981), Browning et al (1984), and Altonji (1986). Heckman (1984) argued that this controversy was somewhat misdirected given that the extensive margin was the dominant margin of adjustment during US business cycles and that both the business cycle models and the micro estimation exercises abstracted from this margin.

It is now understood that the labor supply elasticity of a representative household should capture adjustment along both the intensive and extensive margins among heterogeneous agents. However, since the study of Hansen (1985), it has become common for macroeconomic analyses of business cycles to assume that *all* labor supply adjustment occurs along the extensive margin when deriving the implications of heterogeneous agent models for the aggregate labor supply. The usual motivation for this assumption is that between two thirds and three quarters of business cycle fluctuations in aggregate hours are due to changes in employment rather than hours per worker.

A natural question, but one that has not been asked in the literature, is whether abstracting from the intensive margin is a (relatively) harmless simplification for understanding aggregate labor supply in business cycle settings. This is the question that we address in this paper. Our main finding is that including an empirically reasonable channel of choice along the intensive margin can have important quantitative implications for labor market fluctuations, even if there is relatively little action along the intensive margin.

To answer this question we embed the nonconvex production model of Prescott et al. (2009) into the incomplete markets indivisible labor framework of Chang and Kim (2007). We first assess the ability of this model 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 the model is parsimonious, it is able to account for many stylized facts in the data.

We note two novel aspects to our steady state calibration exercise. First, we argue that the cross-sectional distribution of hours of work can serve as useful information regarding the extent of heterogeneity. Interestingly, based on this measure, we need a degree of heterogeneity that is roughly double the amount captured by standard estimates of idiosyncratic productivity shocks. Because the extent of heterogeneity has important implications, the development of simple procedures for assessing the empirically relevant amount of heterogeneity within an aggregate model is important. Second, to our knowledge, this is the first aggregate analysis to explicitly address how individuals transition not only between employment and non-employment but also within the hours worked distribution.

Having developed an empirically reasonable model that allows for adjustment along the intensive and extensive margin, we study the response of the model to business cycle shocks. To maximize comparability with the previous literature, we focus on the case in which aggregate shocks to productivity are the (only) driving force behind business cycles. Of particular interest is the comparison of models with an operative intensive margin to a model that abstracts from such adjustment. For this purpose we will consider the extensive-margin-only model of Chang and Kim (2007) as our benchmark model. Importantly, we assume that both models are calibrated to match the same aggregate steady-state moments and are subject to the same aggregate shocks. A key question is whether the model that abstracts from the intensive margin is a good approximation to the behavior of the model that features an operative intensive margin. That is, we assess the extent to which abstracting from the intensive margin is a harmless simplification.

Intuition and previous work both suggest that the answer to this question depends on the underlying primitives of the economy, notably the willingness of individuals to substitute hours intertemporally and the extent of heterogeneity. For example, if individuals are not very willing to substitute labor intertemporally then intuitively there will be very

little adjustment on the intensive margin and one might reasonably conjecture that it can be ignored. Because there remains some disagreement about the value of this curvature parameter we consider a set of values consistent with the range of estimates found in the literature. Existing work also suggests that the extent of heterogeneity matters. For example, Prescott et al. (2009) showed that when there is no heterogeneity, all adjustment takes place along the extensive margin even though the intensive margin is available. For this reason we also consider specifications that feature different degrees of heterogeneity.

For each specification of curvature and heterogeneity, we calibrate the models to the same aggregate (steady-state) targets and then compare the business-cycle fluctuations of our model with that of the benchmark model in which the intensive margin is exogenously restricted to be non-operative. Our main finding is that abstracting from the intensive margin can significantly distort inference regarding the volatility of aggregate hours. Moreover, the direction of the distortion depends on the underlying primitives. Surprisingly, even if variation along the intensive margin is very small, explicit modeling of the intensive margin can have a large impact on the volatility of aggregate hours. In particular, the presence of an intensive margin dampens the effect of increased heterogeneity on aggregate hours volatility in an important way. We conclude that abstracting from the intensive margin is a potentially serious issue for analyses that seek to understand the effect of heterogeneity on the magnitude of aggregate fluctuations.

Our model also has important implications for understanding the magnitude of fluctuations in efficiency units of labor. From the perspective of understanding the role of labor supply in accounting for fluctuations in output, this is the key dimension. An important feature of our model is that adjustment along the intensive and extensive margins can matter quite differently for the volatility of hours versus efficiency units. In particular, efficiency units respond much more to fluctuations along the intensive margin than along the extensive margin. This implies that when adjustment becomes relatively more important along the intensive margin, the relationship between changes in aggregate hours and changes in aggregate efficiency units also changes. We show that the distinction between

fluctuations in hours and efficiency units is quantitatively significant.

Our analysis offers an important message regarding intensive and extensive margin elasticities and their determinants. While it is intuitive to think that the intensive margin elasticity is determined by the curvature parameter whereas the extensive margin elasticity is determined by the properties of heterogeneity, we show that intensive and extensive elasticities are not independent. That is, heterogeneity matters for the extent of adjustment along the intensive margin, and curvature matters for the extent of adjustment along the extensive margin.

Our analysis provides a mapping from the specification of the underlying primitives of heterogeneity and curvature in preferences over hours of work into the curvature parameter in a representative-agent model that would generate the same volatility in aggregate hours. Relative to a model that only features an extensive margin, our model produces a much smaller range of curvature values for the stand-in household, with values of Frisch elasticity in the range between one and two. But as noted earlier, this is distinct from the mapping that would target the volatility of efficiency units of labor.

Our paper is related to several in the literature. Kydland and Prescott (1991) and Cho and Cooley (1994) are early business cycle models that feature adjustment along both the intensive and extensive margin, but both do this in the context of a representative household in which there is no heterogeneity in hours worked. Several papers have studied labor supply in incomplete markets models. An early example is Pijoan-Mas (2006), though he assumed all adjustment occurs along the intensive margin and did not study business cycles. Chang and Kim (2006, 2007), study business cycles but assume that all adjustment occurs along the extensive margin. Rogerson and Wallenius (2009) features adjustment along both the intensive and extensive margin but features a limited degree of heterogeneity, no uncertainty and does not study business cycles. Heathcote et al. (2014) study a model that features multiple dimensions of heterogeneity, progressive taxation and a flexible approach to risk-sharing, but they assume all adjustment occurs along the intensive margin and do not study business cycles. The paper that is probably closest to ours is Erosa et al. (2016).

Like us, they build a model that features heterogeneity and incomplete markets and allow 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, and calibrate the model to evaluate the labor supply elasticity of working age males in a partial equilibrium context. While their framework is richer, our somewhat more abstract model is more tractable and lends itself more readily to general equilibrium analyses. But, most importantly, we investigate very different questions: whereas they focus on calibrating their model to evaluate its aggregate labor supply elasticity, we focus on the role that explicit modeling of the intensive margin plays in shaping model predictions for aggregate labor market fluctuations. We therefore view these two pieces of work as complementary.

The paper is organized as follows: Section 2 specifies the benchmark model and the model with an operative intensive margin. Section 3 calibrates the different model specifications and Section 4 considers the steady state properties of the various specifications. Section 5 studies the business cycle properties of the calibrated models. Section 6 considers an alternative model of the intensive margin in which individuals face a discrete choice problem between full-time and part-time work. Section 7 concludes.

2. Models

In this section we describe the two models that we compare in our quantitative business cycle analysis. The benchmark model features adjustment only along the extensive margin, and is identical to the model in Chang and Kim (2007). The other model that we consider extends Chang and Kim (2007) in the spirit of Prescott et al. (2009) to also allow for adjustment along the intensive margin.

2.1. Benchmark Model: Extensive Margin of Adjustment Only

There is a unit measure of ex-ante identical infinitely lived individuals. Each individual has preferences over streams of consumption (c_{it}) and hours of work (h_{it}) given by:

$$\sum_{t=0}^{\infty} \beta^t [\log(c_{it}) - B \frac{h_{it}^{1+1/\gamma}}{1+1/\gamma}]$$

where $0 < \beta < 1$, $B > 0$ and $\gamma > 0$. Each individual is endowed with a unit of time in each period.

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 aggregate productivity shocks (Z_t):

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

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

$$\ln Z_{t+1} = \rho_Z \ln Z_t + \varepsilon_{Zt}, \quad \varepsilon_{Zt} \sim N(0, \sigma_Z^2).$$

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

Individuals are also subject to idiosyncratic productivity shocks, denoted by z_{it} . The stochastic evolution of z_{it} is described by the same transition probability distribution function for all individuals, $\pi_z(z'|z) = \Pr(z_{it+1} \leq z' | z_{it} = z)$, but realizations are i.i.d. across individuals. In our quantitative work we will also assume that z_{it} follows an AR(1) process in logs:

$$\ln z_{it+1} = \rho_z \ln z_{it} + \varepsilon_{zt}, \quad \varepsilon_{zt} \sim N(0, \sigma_z^2).$$

While this is the only source of idiosyncratic shocks in our model, our quantitative work will interpret this shock more broadly to represent all shocks that affect the relative return to market work and so influence the incentives to work.

In this benchmark model, individuals are restricted to choose h_{it} from the set $\{0, \bar{h}\}$ where $0 < \bar{h} \leq 1$. In this case the value of γ is not separately identified from B and \bar{h} , but we include this parameter in the current specification since the next subsection will extend the model to a setting in which h_{it} can take on any values in the interval $[0, 1]$.

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 z_{it} devotes h_{it} units of time to market work, the resulting labor earnings are $w_t z_{it} h_{it}$.

2.2. Extended Model: Extensive and Intensive Margins of Adjustment

In this section we describe an extension to the previous model that does not restrict the set of values for h_{it} that an individual can choose from and that gives rise to adjustment along both the intensive and extensive margins. In particular, we assume 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 z_{it} devotes h_{it} units of time to market work, this will generate $z_{it} \cdot g(h_{it})$ units of labor services. Following Prescott et al. (2009) and Rogerson and Wallenius (2009), we assume that $g(\cdot)$ takes the

following simple form¹:

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

where $0 < \hat{h} < 1$. All other aspects of the environment are as in the benchmark model.

It is important to note that extending the pure indivisible labor model in this fashion in order to endogenously generate adjustment along the intensive and extensive margins necessarily affects the model's implications for the relationship between aggregate hours and aggregate efficiency units of labor input. In particular, holding z_{it} constant, there is a larger increase in efficiency units, if a given change in hours occurs along the intensive margin than from having the same increase in hours resulting from adjustment along the extensive margin. The reason for this is that the former does not involve the “startup” penalty implicit in the \hat{h} term. For this reason we will be interested in studying how incorporating an intensive margin affects not only aggregate hours but also aggregate efficiency units of labor.

2.3. Equilibrium

We formulate equilibrium recursively. Here we formulate equilibrium for the extended model. The benchmark model is equivalent to the extended model in which \hat{h} is set to 0 and the choice of h is restricted. The individual state variables are beginning of period assets (a) and current idiosyncratic productivity (z), and the aggregate state variables will be the current aggregate productivity shock (Z) and a measure μ over the individual state variables (a, z). Prices are functions of the aggregate state: $w(Z, \mu)$ and $r(Z, \mu)$, and the equilibrium law of motion for μ is given by $\mu' = T(Z, \mu)$.

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

¹French (2005) considers the alternative specification of $g(h) = h^\theta$ where $\theta > 1$. Two differences are worth noting. First, this alternative allows for the possibility that individuals might work very few hours, whereas our specification implies that there will be no one working less than \hat{h} hours. Second, this alternative assumes that the nonconvexity is uniform throughout the hours distribution, whereas our specification assumes that the nonconvexity is most severe at low hours of work.

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

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

1. Individuals optimize:

Given factor prices, individual decision rules solve value function.

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

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

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

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

4. Factor markets clear:

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

5. Individual and aggregate behaviors are consistent :

$$\mu'(a^0, z^0) = \int_{a^0, z^0} \left\{ \int_{\mathcal{B}_a, \mathcal{B}_z} \mathbf{1}[a' = a'(a, z; Z, \mu)] d\pi_z(z'|z) d\mu \right\} da' dz'$$

for all $a^0 \subset \mathcal{B}_a$, $z^0 \subset \mathcal{B}_z$.

3. Calibration

As noted in the introduction, one of our main objectives is to examine whether the benchmark model with only an extensive margin reasonably captures the fluctuations in aggregate labor market variables in the model that also features an operative intensive margin of adjustment. That is, we want to ask whether a modeler might reasonably choose to work with the (simpler) benchmark model even though he or she recognizes that there is indeed some adjustment along the intensive margin in the data.

One would expect that the answer to this question may well depend on the values of some key parameters. Two special cases suggest features of the economy that will likely be important. As one extreme case, Prescott et al. (2009) show in their model that if all workers have identical productivity at each point in time, all adjustment will occur along the extensive margin even though the model explicitly allows for adjustment along the intensive margin. In this case there would be no loss in generality in considering the model that restricts adjustment to the extensive margin. Rogerson and Wallenius (2009) show that this result does not hold, if individuals have time varying productivity. The second extreme case corresponds to the limiting case in which γ becomes close to 0. In this case, holding the marginal utility of wealth constant, changes in an individual's productivity will have no effect on his or her hours of work. But as long as γ is positive, variation in individual productivity will lead to variation in hours worked.

Motivated by the above discussion, we will consider economies that differ in terms of the extent of heterogeneity and in the value of γ . One 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, σ_z , and this is the approach that we follow. If there were definitive estimates for σ_z and γ , it would perhaps be sufficient to just carry out our exercise for these specific values. However, we do not think that this is the case—keeping in mind that we want to think of z as reflecting all sources of variation in the relative return to working. Independently of this, examining a range of values is still of interest to get a clearer understanding of the underlying economics.

A reasonable lower bound on the size of the z shocks is provided by the literature that estimates idiosyncratic shocks to wages.² A sizable 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. (2008). 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 $\rho_z = 0.975$ and $\sigma_z = 0.165$.³

As we document later on, this specification will tend to generate too little cross-sectional variation in hours worked. In view of this we also consider specifications with larger values of σ_z . In particular, we assume that σ_z takes on values in the set $\{0.165, 0.2475, 0.33\}$. The largest value in this set tends to generate cross-sectional dispersion in hours worked that slightly exceeds that found in the data and so seems like a reasonable upper bound. For our baseline calibration, we will consider the intermediate value of 0.2475.

Motivation for the set of values considered for γ comes from the large literature that has sought to estimate this parameter. Chetty (2012) argues that an empirically reasonable value for this elasticity parameter is in the range between 0.4 and 0.5. Using a very different method, Pistaferri (2003) found an estimate around .75. Rogerson and Wallenius (2013, 2016) provide evidence for values of γ that are 1 or larger, and as discussed in Keane and Rogerson (2012), there are additional factors that Chetty abstracts from that would suggest higher values. For this reason, we think values of γ as high as 1 are still quite plausible, and consider 1.5 to be a very generous upper bound. We therefore assume that γ takes on values in the set $\{0.25, 0.5, 1, 1.5\}$.

Given values for σ_z and γ , we now describe how we calibrate the remaining parameters. 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 Cobb-Douglas technology parameter α

²This is reasonable as long as other idiosyncratic shocks (e.g., shocks to preferences or home production) 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. The values just reported correspond to the estimates in Floden and Linde (2001).

is set to 0.64 and the depreciation rate, δ , is set to .025, effectively targeting labor's share of income and the investment to output ratio. For the aggregate technology shocks we set $\rho_Z = 0.95$, and $\sigma_Z = 0.007$. These values are set independently of γ and σ_z .

There are four additional parameters to calibrate: β , B , \bar{a} , and \bar{h} (\hat{h} in the extended model). We set $\bar{a} = 0$ in all cases. Given values for γ and σ_z , we choose the values of the remaining three parameters to match three moments in the deterministic steady state equilibrium: a (quarterly) rate of return to capital of 1%, an employment rate of 70%, and average hours (conditional on working) equal to 1/3.⁴ In the benchmark model, this last target directly implies that $\bar{h} = 1/3$, while in the extended model the value of \hat{h} is adjusted so that optimally chosen hours have this average. In the benchmark model, there is no dispersion around this mean value, whereas in the extended model there will be dispersion of hours around this mean.

Note that we are following the standard practice in the business cycle literature of recalibrating the model as we vary the values of γ and σ_z so that each specification matches the same aggregate (steady-state) target moments. This is important given that we want to consider how predictions would be affected, if the researcher were to start with the benchmark model that abstracts from choice along the intensive margin rather than the more general model. In doing so, we want the researcher to continue to match the same targets.⁵

Table 1 summarizes the parameter values that are held constant across specifications as well as those that vary across specifications. As noted earlier, the value of γ is irrelevant for the benchmark economy, so we have set $\gamma = 1$ for each specification in the benchmark economy.

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.

⁵In some circumstances one might also be interested in a change in one parameter holding all other parameters constant as a way to develop intuition for how a given parameter affects model predictions. This distinction will be relevant when we discuss the extent to which some of our findings across specifications might seem counterintuitive.

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 is to examine how different values for γ and σ_z influence the model's ability to account for these features in the data.

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

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".

hours in the CPS and PSID so that average annual hours is the same as in the economy with $\gamma = 1$ and $\sigma_z = .165$.⁸

The cross-sectional standard deviations of hours in the PSID and CPS are fairly comparable: 0.42 in the former and 0.45 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.⁹ The standard deviations for the twelve model specifications range from .294 to .450. Consistent with intuition, the hours dispersion in the model is increasing in both γ and σ_z . When $\sigma_z = .165$, the model cannot generate sufficient dispersion in hours even with a relatively large value of γ . However, when $\sigma_z = .33$ 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 $\sigma_z = .33$ as a reasonable upper bound on the extent of heterogeneity in our model.

As noted earlier, we interpret the idiosyncratic shock z as a stand-in for all sources of variation that affect the relative return to market work. While it is possible to measure the component of this shock that corresponds to market opportunities, it is much less straightforward to measure other components. The close connection between the standard deviation of the idiosyncratic shocks and the cross-sectional dispersion in hours worked suggests that this dispersion can be used as a way to identify the standard deviation of our single idiosyncratic shock. The above exercise suggests that $\sigma_z = 0.33$ is an empirically relevant case to study.

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 three model specifications (three σ_z 's 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

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

⁹The tendency for reported hours to exhibit considerable bunching at certain values (e.g., 35 or 40 for usual weekly hours) suggests an effect that may go in the opposing direction.

see that the specification with $\gamma = 1$ and $\sigma_z = .33$ does a much better job of tracking the empirical hours distributions than does the specification with the lower value of σ_z . 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 again 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 $\sigma_z = .2475$). We only report statistics for one of the model specifications because it turns out that these statistics are very similar across all twelve 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 evidenced by the fact that the model has a 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 .77 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 quartiles of the hours worked distribution (and non-employment) from the PSID and three model specifications: $\gamma = 1$ and the three values of σ_z .

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 quartile this value exceeds 60%. Second, conditional on working in both periods and switching quartiles, the transition probabilities are monotone decreasing in the distance of the destination quartile from the originating quartile. Third, individuals who adjust along the extensive margin between years t and $t + 1$ are disproportionately from the lowest quartile 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 quartile 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 quartile in year $t + 1$.

Next we consider how the model fares in terms of reproducing these features. As Table 4 shows, all model specifications generate considerable persistence, though somewhat less as σ_z increases. The average of the diagonal elements for those with positive hours in both periods is 55.6% in the data. For the three model specifications the corresponding values are 52.1, 49.2 and 43.2 as σ_z runs from .165 to .33. While the model specifications come close to matching the persistence in the highest quartile, the lowest quartile displays quite a bit more persistence in the data than in the model—55.8% versus values between 36.9% and 38.4%. 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 $\sigma_z = .33$ the model has the same monotonicity property found in the data, whereas for the other specifications 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 each of the model specifications shown in the table, more than eighty percent of all transitions along the extensive margin involve workers who are in the lowest quartile of the hours distribution.

Overall, we take the fact that the model can match the salient features of transitions within the hours distributions to suggest that our parsimonious representation of individual heterogeneity is empirically reasonable.

4.3. 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. In particular, when $\sigma_z = .165$ 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, $\sigma_z = .33$, we obtain similar patterns qualitatively, but perhaps not surprisingly, the model generates more dispersion in earnings than is found in the data. In this section we document these properties.

Given that we calibrate our model using employment data from the PSID, we prefer 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.¹⁰

Table 5 reports the Gini coefficients for both the wealth and earnings distributions for our 15 different model specifications, as well as their corresponding values in the PSID and SCF.

A few patterns are evident. First, given a value for σ_z , the Gini coefficients for both the

¹⁰This is not the case, if one focuses on the extreme upper tail of the wealth distribution, as the PSID under-samples the upper extremes of the wealth distribution. However, given our emphasis on labor supply this extreme upper tail is probably not of major concern.

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 5 is that the quantitative importance of these effects are quite small. While moving from $\gamma = 1$ to the extensive margin only case does generate modest changes in both Gini coefficients, the effect of changes in γ inside the interval $[\.25, 1.5]$ is of second order importance for each measure.

It is also intuitive that higher values of σ_z 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_z 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 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. To look a bit deeper into the nature of the wealth and earnings distributions, the bottom part of Table 5 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 model specifications with $\gamma = 1$.

The basic message from these measures 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 $\sigma_z = .165$ does a very good job of capturing the earnings shares, whereas the specification with $\sigma_z = .33$ generates excessive concentration

of earnings in the upper quintile of the wealth distribution. However, the specification with the higher value of σ_z 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. 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 Cycle Analysis

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). As is standard, we take logs and then HP filter (with the smoothing parameter of 1,600) the simulated series before computing summary statistics.

Table 6 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. 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 models lead to different propagation of these shocks.¹¹

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.

5.1. Properties of the Extensive Margin Only Economies

It is useful to first summarize the key patterns for the benchmark model that features only an extensive margin. The main pattern is that as σ_z increases from .165 to .33 the volatility of employment and aggregate hours falls very dramatically, from .93 to .30. The intuition for this result is simple and well-known. In steady state, the optimal decision rules define a curve that divides the individual state space into two distinct regions, one with working as the optimal decision and the other with not working as the optimal decision. A positive aggregate productivity shock increases the wage rate and encourages workers who were previously slightly below the work threshold to engage in market work. The size of the employment response depends on the density of individuals near the boundary, and this in turn is influenced by the extent of heterogeneity. A larger value of σ_z tends to spread out the distribution of workers across the state space, thereby lowering this density and decreasing the employment response.

While the qualitative pattern just noted is intuitive, we want to emphasize the magnitude of the effects. Moving from $\sigma_z = .165$ to $\sigma_z = .33$ reduces volatility in aggregate hours

¹¹A distinct issue that we do not pursue here is whether alternative model specifications might lead to different measurement of the shocks.

by a factor of three. Whereas the economy with $\sigma_x = .165$ accounts for roughly two thirds of the volatility of aggregate hours relative to output, the $\sigma_z = .33$ specification accounts for less than a third of this measure.

It is also of interest to contrast fluctuations in hours with fluctuations in efficiency units of labor. Two features are relevant. First, although fluctuations in efficiency units are also decreasing in the value of σ_z , the magnitude of the decrease is much smaller. Second, for the two smaller values of σ_z , the volatility in efficiency units is less than the volatility in hours, but for $\sigma_z = 0.33$ the reverse is true. To understand these patterns, it is important to know who the marginal workers are. If the fluctuations in employment are dominated by individuals with idiosyncratic productivity that is lower than the mean idiosyncratic productivity of employed individuals, then hours will fluctuate more than efficiency units. But, it is also possible that shocks induce transitions among “high-wealth / high-productivity” individuals. As σ_z increases more of the marginal individuals are of this type, thereby flipping the relative volatility of hours and efficiency units. As σ_z increases and the extent of heterogeneity increases, the gap between the volatility in hours and efficiency units becomes larger.

5.2. Properties of the Intensive and Extensive Margin Economies

Next we look at how business cycle features vary across the twelve specifications that feature adjustment along both the intensive and extensive margin. We note first that the property just noted for the benchmark economy—that increases in σ_z lead to lower volatility for aggregate output and aggregate hours—also holds for each value of γ . That is, if we fix γ and increase σ_z , we see that volatility in aggregate output and aggregate hours decreases in all cases. In particular, allowing for an intensive margin does not overturn this feature qualitatively.

Next we summarize the patterns as we increase γ holding σ_z fixed. Intuitively, one would expect that as γ increases we would see an increase in the extent of fluctuations

along the intensive margin. This intuition is confirmed by the table: for each value of σ_z we observe that higher values of γ lead to greater volatility of hours per worker. The magnitude of this effect is substantial: moving from $\gamma = 0.5$ to $\gamma = 1$ can almost double the magnitude of fluctuations along the intensive margin.

One might also intuitively expect that greater fluctuations along the intensive margin will lead to a reduction in fluctuations along the extensive margin, as these two margins are substitutes in terms of ways to adjust overall labor input. While this intuition is confirmed for the case of $\sigma_z = 0.165$, we see that it largely fails for the other two values of σ_z . To first approximation, the other two cases reveal little to no change in the magnitude of fluctuations along the extensive margin as we vary the value of γ .

Similarly, one might also conjecture that an increase in γ holding σ_z constant makes it less costly to vary aggregate hours over the business cycle and so should lead to an increase in the volatility of aggregate hours. Looking at Table 6 we see that this effect is confirmed for the two higher values of σ_z but that the reverse holds true for the case of $\sigma_z = 0.165$.

Why does our basic intuition seem to not hold consistently across these specifications? We think this is an important question, and highlights a key feature of our model. In particular, we suspect that the most basic intuition comes from a setting in which hours and employment are perfect substitutes in production (i.e., only total hours matters) and there is a representative household with preferences over each of employment and hours per worker that feature curvature. Cho and Cooley (1994) is an example of this setting. In this setting, less curvature on hours per worker will lead to greater fluctuations along the intensive margin, less fluctuations along the extensive margin, but greater fluctuations in terms of aggregate hours. Importantly, this kind of model does not offer a micro foundation for adjustment along the two margins.

In contrast, our model does provide micro foundations for adjustment along both margins. The two key elements of this micro foundation are heterogeneity among individuals and a non-linear mapping from hours into efficiency units, both of which destroy the assumption of perfect substitutes for the two margins in production. Specifically, if workers

are heterogeneous, then it matters whether an additional hour of total work comes from existing workers or from bringing in a new worker. And because of the nonlinear mapping from hours per worker into efficiency units, even if the potential added worker has the same productivity as existing workers, it is not the same to bring in a new worker to work one hour as it is to have an existing worker work one more hour.¹²

This discussion highlights, even more so than in the extensive margin only economy, the need to distinguish between aggregate hours and aggregate efficiency units of labor. In fact, when we examine the relationship between the value of γ and the volatility of efficiency units of labor holding σ_z fixed, we see that the relationship is basically increasing.

Our analysis also offers an important message concerning the determinants of fluctuations along the intensive and extensive margins. In particular, we start from the notion, 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_z . 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 .

While the results in Table 6 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_z , 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 6. Specifically, starting from the specification in which $\gamma = 0.5$ and $\sigma_z = .165$, we see that increasing γ to 1.5 leads to more than a doubling of the response along the intensive margin. However, we also get a doubling of the response along the intensive margin if we keep the value of γ

¹²Beyond these interactions, recall also that we recalibrate the parameters for each of our model economies, and in particular, the parameter \hat{h} , which affects the nature of the nonlinearity in the mapping from hours to efficiency units. As Table 1 shows, this value is strongly affected by changes in γ .

fixed at 0.50 but we instead increase the value of σ_z to .33. 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 $\sigma_z = .165$, we note that moving to the specification with $\sigma_z = .33$ leads to roughly a fifty percent reduction in fluctuations along the extensive margin. However, if we instead kept the value of σ_z 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 twenty percent. It follows that one should not think of “intensive margin elasticities” and “extensive margin elasticities” as two independent and primitive parameters.

5.3. Is the Extensive Margin Only Economy a Good Approximation?

We now ask to what extent the extensive margin only economy accurately captures labor market fluctuations in the more general models that include an operative intensive margin, i.e., to what extent does it matter that one abstracts from adjustment along the intensive margin. Before turning to specific results, it is instructive to return to the extreme special cases that we mentioned earlier. Specifically, we argued that either little heterogeneity or a very low value of γ might suggest that it is very innocuous to abstract from adjustment along the intensive margin.

With this in mind, consider the extended model with $\sigma_z = .165$ and $\gamma = .25$. In the calibration section we argued that both of these were conservative lower bounds and hence are a relevant case to consider. Table 6 reveals a striking result: whereas the benchmark model generates a volatility of aggregate hours equal to 1.04 when $\sigma_z = .165$, the extended model with $\sigma_z = .165$ and $\gamma = .25$ only generates a volatility of aggregate hours equal to .89, a reduction of roughly fifteen percent. In contrast, the volatility of efficiency units of labor is roughly fifteen percent larger than in the model with an intensive margin. In each case, it seems that the extensive margin only economy is not a close approximation.

It is important to emphasize that there is a striking difference between the volatility of aggregate hours and aggregate efficiency units in the two economies, despite the fact that in the extended model with $\gamma = 0.25$ and $\sigma_z = 0.165$, the volatility of hours per worker is only .03 and is only about one-thirtieth the size of fluctuations in employment. Put somewhat differently, looking at the business cycle fluctuations in the economy with an operative intensive margin, it would seem tempting to conclude that it is seemingly harmless to abstract from fluctuations along the intensive margin given how small they are. The key message here is that even if there is little volatility along the intensive margin, abstracting from it has important implications precisely because a model that has an active intensive margin will necessarily contain features that are important for the nature of labor market fluctuations.

If we require the model to generate the same amount of cross-sectional dispersion in hours as found in the data, we are led to focus on the case in which $\sigma_z = .33$. As noted before, in the benchmark model there is a dramatic decrease in the volatility of aggregate hours when we move from $\sigma_z = .165$ to $\sigma_z = .33$. When $\sigma_z = .165$, allowing for an active intensive margin significantly decreased the extent of hours volatility, and the decrease became larger as the value of γ was increased. We now observe the opposite pattern when $\sigma_z = .33$. That is, allowing for an intensive margin serves to increase the volatility of aggregate hours relative to the extensive margin only case, and the extent of the increase is increasing in the value of γ . It follows that the effect of increased heterogeneity on the volatility of aggregate hours is muted considerably by allowing for an intensive margin. For example, whereas moving from $\sigma_z = .165$ to $\sigma_z = .33$ in the benchmark model lowers the standard deviation of aggregate hours by .65, the analogous drop when $\gamma = .5$ is only .4. If we abstract from the intensive margin we overstate the decrease in volatility.

Our analysis also allows us to construct a mapping from the values of γ and σ_z into the implied curvature parameter for a representative household that would generate the same fluctuations in hours worked. In particular, we consider a representative household model

in which the household has period utility function given by:

$$\log c - B \frac{h^{1+1/\hat{\gamma}}}{1 + 1/\hat{\gamma}}$$

We calibrate the value of B so that hours worked in steady state are 1/3 and consider the same aggregate shock process as before. We then ask what value for $\hat{\gamma}$ provides the same volatility in aggregate hours as the model with heterogeneity and incomplete markets. Table 7 shows the results for both the general model and the model with only an extensive margin. The first point to note is that the range of implied values for $\hat{\gamma}$ is much smaller once one includes an intensive margin. Whereas the range is from .75 to 4.46 in the benchmark model with only an extensive margin, the range is basically from 1 to 3 for the models that include an intensive margin, even though we consider a wide range of values for the individual curvature parameter. Moreover, if we were to restrict attention to $\sigma_z = .33$ on the grounds that it generates an empirically reasonable amount of heterogeneity in hours, and also consider values of γ that lie between .5 and 1, we see that the implied representative agent elasticity parameter is quite tightly pinned down and is around 1. Even with γ as large as 1.5 we see that the representative agent elasticity is still only as high as 1.5

6. An Alternative Model of the Intensive Margin

One of the findings from model simulations in the previous section is that fluctuations along the intensive margin tend to be quite small relative to what we see in the data, even when viewed relative to fluctuations along the extensive margin. In our model, choice along the intensive margin is continuous and choice along the extensive margin is discrete. Recent work by Borowczyk-Martins and Lale (2018) has argued that a significant fraction of volatility along the intensive margin reflects movements between full-time and part-time work and so might also reflect discrete choice.

These two observations lead us to consider an alternative model of the intensive margin

in this section. In particular, we consider an extension to the extensive margin only model that simply adds a third choice for hours that we will interpret as part-time work and briefly summarize its properties.¹³

6.1. Model and Calibration

The model is the same as our previous benchmark model with two notable exceptions. First, as just noted, we now assume that the choice of h_t for an individual is restricted to the set $\{0, h_p, h_f\}$ where h_p and h_f will represent part-time and full-time hours, respectively. Second, we assume that efficiency units per unit of labor supply are less for part-time work. As noted above, this model features only discrete choices and so we will refer to it as the discrete choice model.

We modify the calibration procedure as follows. We assume that part-time work yields thirty percent fewer efficiency units per unit of time than does full time work, and hence that there is a thirty percent wage penalty for part time work. We assume that $\gamma = 1$ and choose B as a disutility of working parameter. As before we target the same average hours conditional on working (0.33) and the aggregate employment rate (70%) in the steady state. We now add as an additional target that the full-time employment is 80% of total employment. These three targets allow us to determine the values of B , h_p and h_f . Assuming $\sigma_z = .2475$ the resulting values are $h_p = .237$ and $h_f = .357$. This hours ratio of $\frac{h_p}{h_f} = 0.67$ implies 27 hours per week for part-time workers when the average hours of full-time workers is 43.

6.2. Steady State Results

Here we focus on one particular aspect of the calibrated model: transitions between the three work choices. We compute transition probabilities from the PSID based on using

¹³Chang et al. (2011) study a model of households in which each member must choose between not working, part-time work and full-time work, but do not study business cycles.

35 hours as the cutoff for defining part-time and full-time work. Results are presented in Table 8 for the case in which $\sigma_z = 0.2475$. We also report the transition probabilities of our model with the intensive margin (for the case of $\gamma = 1$ and $\sigma_z = 0.2475$) using the same cutoffs for full- and part-time work as we used in the PSID.

The two models imply very similar transition probabilities. There are two main issues with respect to matching the transition probabilities in the data, both of which were noted earlier. First, our model has slightly less persistence in the employment and non-employment state, implying that we have somewhat too many transitions between employment and non-employment. Second, we do not generate sufficient persistence in the part-time state.

6.3. Business Cycle Results

The bottom part of Table 6 shows the business cycle results for the discrete choice economy. Although not shown in the table, the model generates a procyclical series for the fraction of individuals working full-time and a countercyclical series for the fraction working part-time, consistent with the data. The fraction of individuals working part-time experiences inflows and outflows: when a positive productivity shock happens it moves some individuals currently working part-time into full-time employment, but also moves some individuals from non-employment into part-time employment. But the former effect dominates.

Because we have implicitly set $\gamma = 1$ in our calibration of the discrete choice model it is best to focus on this case when comparing with the previous model results in Table 6. A few patterns are worth noting. First, whereas the fluctuations along the intensive margin become larger in the continuous choice model as σ_z increases, here we find the opposite. The logic for this is essentially the same as the logic for the decreasing volatility of overall employment with respect to σ_z in the benchmark model. In the discrete choice model, fluctuations along the intensive margin result from individuals who cross between regions as boundaries shift in the individual state space. As cross-sectional heterogeneity increases,

there are fewer individuals close to the boundaries and so fewer transitions. In the discrete choice model the volatility of intensive and extensive margins move in the same direction as σ_z changes, whereas in the continuous choice model they move in opposing directions.

Second, one of our main results in the previous analysis—that allowing for an intensive margin dampens the effect of increased heterogeneity on fluctuations—continues to hold in the discrete choice model despite the fact that the intensive and extensive margins move together. Key to this is the fact that there are multiple boundaries in this model. As some individuals are pushed further from the boundary for full-time work, they are pushed closer to the boundary for part-time work.

Third, this same result holds for fluctuations in efficiency units as well: as the extent of heterogeneity increases, the discrete choice model experiences less of a reduction in volatility of efficiency units of labor than does the benchmark economy.

To what extent does the benchmark model offer a good approximation to fluctuations in this economy? Table 6 reveals that for the lowest value of σ_z the two models generate quite similar aggregate fluctuations in terms of output, hours and efficiency units. But for the two larger values of σ_z there are significant discrepancies along all three dimensions. While this alternative model of the intensive margin differs both qualitatively and quantitatively along some dimensions, we conclude that the basic message remains intact.

7. Conclusion

Recent advances in modeling aggregate labor supply have emphasized the importance of accounting for adjustment along the intensive and extensive margins. 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 nature of idiosyncratic shocks that individuals face, which in turn influences the extent of cross-sectional 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 aggregate technology shocks. Three key findings emerge. First, abstracting from adjustment along the intensive margin may be misleading about aggregate responses in both hours and output even if there are relatively small fluctuations along the intensive margin. Second, 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. Third, the model with both intensive and extensive margins of adjustment produces much less sensitivity in aggregate volatility in response to changes in the underlying primitives. We view our aggregation results as suggesting a representative household with a Frisch elasticity slightly larger than unity.

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] Borowczyk-Martins, D., and E. Lale (2018): “Employment Adjustment and Part-Time Work: Lessons from the United States and the United Kingdom.” Forthcoming, *American Economic Journal: Macroeconomics*.
- [6] Browning, M., A. Deaton and M. Irish (1985), “A Profitable Approach to Labor Supply and Commodity Demands over the Life-Cycle,” *Econometrica* 53, 503-44.
- [7] Card, D., (1994), “Intertemporal Labor Supply: An Assessment,” in *Advances in Econometrics*, edited by Chris Sims, Cambridge University Press, 49-81.
- [8] 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.
- [9] Chang, Yongsung, and Sun-Bin Kim (2007), “Heterogeneity and Aggregation: Implications for Labor-Market Fluctuations,” *American Economic Review*, 97(5), 1939-1956.

- [10] 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.
- [11] Chetty, R. (2012), “Bounds on Elasticities with Optimization Frictions: A Synthesis of Micro and Macro Evidence on Labor Supply.” *Econometrica* 80(3): 969-1018.
- [12] Cho, J.-O. and Cooley, Thomas F. (1994), “Employment and hours over the business cycle.”, *Journal of Economic Dynamics and Control* 18, 411-432.
- [13] Diaz-Gimenez, J., V. Quadrini and V. Rios-Rull (1997), “Dimensions of Inequality: Facts on the US Distributions of Earnings, Income and Wealth,” Quarterly Review of Federal Reserve Bank of Minneapolis Spring, 3-21.
- [14] Erosa, A., L. Fuster and G. Kambourov (2016) “Towards a Micro-Founded Theory of Aggregate Labor Supply,” *Review of Economic Studies* 83(3), 1001-1039.
- [15] 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.
- [16] French, Eric (2005), “The Effects of Health, Wealth, and Wages on Labor Supply and Retirement Behavior,” *Review of Economic Studies*, 72(2), 395-427.
- [17] Hansen, G. (1985), “Indivisible Labor and the Business Cycle,” *Journal of Monetary Economics* 16, 309-327.
- [18] Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L. Violante (2008), “Insurance and Opportunities: A Welfare Analysis of Labor Market Risk,” *Journal of Monetary Economics*, 55, 501–525.
- [19] Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L. Violante (2014), “Consumption and Labor Supply With Partial Insurance: An Analytical Approach,” *American Economic Review*, 104(7) , 2075-2126.

- [20] Heckman, J., (1984), “Comments on the Ashenfelter and Kydland Papers,” *Carnegie Rochester Conference Series on Public Policy*, 21, 209-224.
- [21] Huggett, M. (1996), “Wealth Distribution in Life-Cycle Economics,” *Journal of Monetary Economics* 38, 469-94.
- [22] Keane, M. and R. Rogerson (2012) “Macro Labor Supply Elasticities: A Reassessment of Conventional Wisdom,” *Journal of Economic Literature*, 50, 464-476.
- [23] Krusell, Per, and Anthony A. Smith (1998), “Income and Wealth Heterogeneity in the Macroeconomy”, *Journal of Political Economy*, 89(6), 1059-1085.
- [24] Kydland, Finn E. and Edward C. Prescott (1982), “Time to build and aggregate fluctuations,” *Econometrica*, 50, 1345-79.
- [25] Kydland, Finn E. and Edward C. Prescott (1991), “Hours and employment variation in business cycle theory,” *Economic Theory*, 1, 63-81.
- [26] MaCurdy, Thomas E. (1981), “An Empirical Model of Labor Supply in a Life-Cycle Setting,” *Journal of Political Economy*, 106(5), 867-96.
- [27] Pijoan-Mas, J. (2006), “Precautionary Savings or Working Longer Hours,” *Review of Economic Dynamics* 9(2), 326-352.
- [28] Pistaferri, L. (2003), “Anticipated and Unanticipated Wage Changes, Wage Risk and Intertemporal Labor Supply.” *Journal of Labor Economics* 21, 729-754.
- [29] Prescott, Edward, Richard Rogerson, and Johanna Wallenius (2009), “Lifetime aggregate labor supply with endogenous workweek length,” *Review of Economic Dynamics* 12, 23-36.
- [30] Rogerson, Richard, and Wallenius, Johanna (2009), “Micro and Macro Elasticities in a Life Cycle Model with Taxes,” *Journal of Economic Theory* 144, 2277-2293.

- [31] Rogerson, Richard, and Wallenius, Johanna (2014), “Nonconvexities, Retirement and the Elasticity of Labor Supply.” *American Economic Review* 103, 1445-1462.
- [32] Rogerson, Richard, and Wallenius, Johanna (2016), “Retirement, Home Production and Labor Supply Elasticities,” *Journal of Monetary Economics* 78 (2016), 23-34.

Table 1: Calibrated Parameter Values

Common Across All Models					
σ	α	δ	ρ_Z	σ_Z	ρ_z
1.0	0.64	0.025	0.95	0.007	0.975
		σ_z	β	B	\hat{h}
Ext		0.165	0.97855	18.4	-
		0.2475	0.9726	15.0	-
		0.33	0.96795	11.2	-
$\gamma = 0.25$		0.165	0.9772	1248	0.26
		0.2475	0.9705	990	0.252
		0.33	0.9654	710	0.24
$\gamma = 0.5$		0.165	0.97688	85.0	0.211
		0.2475	0.97	69.4	0.2
		0.33	0.96472	53.0	0.18
$\gamma = 1.0$		0.165	0.97656	18.9	0.151
		0.2475	0.9693	15.8	0.136
		0.33	0.96373	12.3	0.114
$\gamma = 1.5$		0.165	0.97628	11.0	0.115
		0.2475	0.9688	9.4	0.098
		0.33	0.963	7.3	0.076
Discrete		0.165	0.97749	18.3	-
		0.2475	0.9711	15.8	-
		0.33	0.9664	12.43	-

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

Table 2: Cross-Sectional Standard Deviation of Annual Hours

PSID: 0.42, CPS: 0.45				
Models				
	$\gamma = 0.25$	$\gamma = 0.5$	$\gamma = 1$	$\gamma = 1.5$
$\sigma_z = 0.165$	0.294	0.301	0.317	0.332
$\sigma_z = 0.2475$	0.308	0.311	0.340	0.371
$\sigma_z = 0.33$	0.317	0.338	0.397	0.450

Notes: Annual hours in the CPS and PSID are normalized so that average annual hours is the same as that in the economy with $\gamma = 1.0$ and $\sigma_x = .165$.

Table 3: Distribution of Employment Transitions

Emp _{t-1}	Emp _t	PSID	Model
1	1	.767	.729
1	0	.044	.053
0	1	.032	.053
0	0	.155	.165

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

Table 4: Annual Hours Transition

		<i>PSID</i>				
		$t + 1$				
		Not Work	1st	2nd	3rd	4th
t	Not Work	81.2	15.3	1.2	1.5	0.9
	1st	16.1	55.8	14.8	8.7	4.7
	2nd	3.0	15.6	53.4	20.7	7.2
	3rd	2.2	8.2	24.9	46.5	19.2
	4th	1.8	4.7	7.5	19.4	66.6
		<i>Model: $\gamma = 1, \sigma_z = 0.2475$</i>				
		$t + 1$				
		Not Work	1st	2nd	3rd	4th
t	Not Work	75.7	20.2	2.4	1.1	0.5
	1st	22.5	36.9	19.7	12.4	8.4
	2nd	2.6	16.2	44.2	28.1	21.7
	3rd	1.2	8.8	13.1	55.0	21.8
	4th	1.2	14.0	6.8	17.4	60.5
		<i>Model: $\gamma = 1, \sigma_z = 0.165$</i>				
		$t + 1$				
		Not Work	1st	2nd	3rd	4th
t	Not Work	75.5	20.4	1.8	1.0	1.2
	1st	23.2	38.4	14.4	10.0	13.9
	2nd	2.4	17.1	27.3	39.8	13.3
	3rd	0.7	5.9	6.2	76.5	11.4
	4th	2.3	19.9	3.5	7.9	66.3
		<i>Model: $\gamma = 1, \sigma_z = 0.33$</i>				
		$t + 1$				
		Not Work	1st	2nd	3rd	4th
t	Not Work	76.3	19.8	2.8	0.8	0.2
	1st	22.2	37.3	24.4	10.7	5.3
	2nd	3.5	20.7	36.8	28.7	10.2
	3rd	1.2	9.9	19.3	40.7	28.7
	4th	0.6	8.8	9.9	22.7	57.9

Notes: Transition probability matrix of annual hours worked across 5 groups including zero hours ('Not Work').

Table 5: Wealth and Earnings

<u>Gini Coefficients: Earnings</u>					
PSID=0.53, SCF=0.63					
	Ext	$\gamma = 0.25$	$\gamma = 0.5$	$\gamma = 1.0$	$\gamma = 1.5$
$\sigma_z = 0.165$	0.55	0.58	0.59	0.59	0.60
$\sigma_z = 0.2475$	0.66	0.69	0.70	0.71	0.72
$\sigma_z = 0.33$	0.75	0.77	0.78	0.79	0.80

<u>Gini Coefficients: Wealth</u>					
PSID=0.76, SCF=0.78					
	Ext	$\gamma = 0.25$	$\gamma = 0.5$	$\gamma = 1.0$	$\gamma = 1.5$
$\sigma_z = 0.165$	0.68	0.71	0.71	0.71	0.72
$\sigma_z = 0.2475$	0.72	0.73	0.74	0.74	0.75
$\sigma_z = 0.33$	0.76	0.77	0.78	0.78	0.79

<u>Wealth Share by Quintile</u>					
	I	II	III	IV	V
PSID	-0.52	0.50	5.06	18.74	76.22
SCF	-0.39	1.74	5.72	13.43	79.49
$\sigma_z = 0.165$	0.00	0.68	6.24	20.23	72.85
$\sigma_z = 0.2475$	0.04	1.05	5.24	16.46	77.21
$\sigma_z = 0.33$	0.09	0.96	3.90	12.89	82.16

<u>Earnings Share by Wealth Quintile</u>					
	I	II	III	IV	V
PSID	7.51	11.31	18.72	24.21	38.23
SCF	7.05	14.50	16.48	20.76	41.21
$\sigma_z = 0.165$	7.65	10.22	16.36	23.34	42.43
$\sigma_z = 0.2475$	2.67	7.27	12.48	21.68	55.91
$\sigma_z = 0.33$	1.15	4.21	8.71	18.32	67.62

Notes: Wealth and earnings shares by quintile are based on the model with $\gamma = 1$.

Table 6: Business Cycle Statistics

		σ_Y	σ_H	$\frac{\sigma_H}{\sigma_Y}$	σ_E	σ_h	σ_L	σ_w
Data	BLS	2.01	1.80	.89	1.51	.48	-	.98
Ext	$\sigma_z = .165$	1.63	1.04	.63	1.04	-	.96	.89
	$\sigma_z = .2475$	1.39	.62	.44	.62	-	.60	.98
	$\sigma_z = .33$	1.27	.39	.31	.39	-	.59	1.06
$\gamma = 0.25$	$\sigma_z = .165$	1.72	.89	.51	.86	.03	1.10	.85
	$\sigma_z = .2475$	1.52	.49	.32	.44	.05	.81	.93
	$\sigma_z = .33$	1.47	.39	.26	.39	.07	.90	1.00
$\gamma = 0.5$	$\sigma_z = .165$	1.73	.86	.50	.81	.06	1.09	.84
	$\sigma_z = .2475$	1.57	.54	.35	.46	.09	.89	.92
	$\sigma_z = .33$	1.52	.46	.30	.36	.12	.99	.99
$\gamma = 1.0$	$\sigma_z = .165$	1.74	.79	.45	.69	.12	1.09	.83
	$\sigma_z = .2475$	1.62	.59	.37	.45	.16	.99	.90
	$\sigma_z = .33$	1.61	.56	.35	.38	.21	1.14	.98
$\gamma = 1.5$	$\sigma_z = .165$	1.76	.78	.44	.63	.16	1.12	.82
	$\sigma_z = .2475$	1.67	.64	.38	.44	.22	1.08	.89
	$\sigma_z = .33$	1.68	.62	.37	.37	.29	1.27	.98
Discrete	$\sigma_z = .165$	1.69	.92	.55	.61	.37	1.03	.85
	$\sigma_z = .2475$	1.51	.74	.49	.52	.29	.78	.94
	$\sigma_z = .33$	1.38	.54	.39	.43	.16	.72	1.02

Note: Total hours (H) = Employment (E) \times Hours per worker (h). The wage rate is denoted by w . Data moments are based on the Labor Productivity and Costs (LPC) of the non-farm business sector (1959:I-2017:IV) from the Bureau of Labor Statistics (BLS). The aggregate efficiency unit of labor is denoted by L .

Table 7: Implied Frisch elasticity for a Stand-in Household Model

	$\sigma_z = .165$	$\sigma_z = .2475$	$\sigma_z = .33$
Ext	4.46	1.48	0.75
$\gamma = 0.25$	2.98	1.03	0.75
$\gamma = 0.5$	2.76	1.19	0.94
$\gamma = 1.0$	2.30	1.37	1.26
$\gamma = 1.5$	2.24	1.56	1.48
Discrete	3.22	2.03	1.19

Note: The Frisch elasticity of labor supply from a representative-agent model that generates the same volatility of total hours (σ_H).

Table 8: Transition between Full-Time and Part-Time


<i>PSID</i>				
		$t + 1$		
		Not Work	Part	Full
t	Not Work	81.2	14.4	4.4
	Part	18.4	51.0	30.5
	Full	2.6	7.7	89.7

Model: Divisible Hours				
		$t + 1$		
		Not Work	Part	Full
t	Not Work	75.9	20.7	3.4
	Part	21.9	39.0	39.1
	Full	1.6	13.2	85.2

Model: Discrete Choice of Hours				
		$t + 1$		
		Not Work	Part	Full
t	Not Work	74.9	19.8	5.2
	Part	22.7	35.1	42.2
	Full	2.3	19.8	85.6

Notes: The statistics from the models are based on $\gamma = 1.0$ and $\sigma_x = .2475$.

Figure 1: Distribution of Annual Hours Worked



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