

**CROSS-SECTIONAL
AND
AGGREGATE LABOR SUPPLY**

Yongsung Chang
Sun-Bin Kim
Kyoocho Kwon
Richard Rogerson

October 2019

The Institute of Social and Economic Research
Osaka University
6-1 Mihogaoka, Ibaraki, Osaka 567-0047, Japan

Cross-Sectional and Aggregate Labor Supply*

Yongsung Chang
Seoul National University

Sun-Bin Kim
Yonsei University

Kyooho Kwon
Korea Development Institute

Richard Rogerson
Princeton University
NBER

October 9, 2019

Abstract

Standard heterogeneous agent macro models that highlight idiosyncratic productivity shocks do not generate the near zero cross-sectional correlation between hours and wages found in the data. We ask whether matching this moment matters for business cycle properties of these models. To do this we explore two extensions of the model in Chang et al. (2019) that can match this empirical cross-section correlation. One of these departs from the assumption of balanced growth preferences. The other introduces an idiosyncratic shock to the opportunity cost of market work that is highly correlated with the shock to market productivity. While both extensions can match the empirical correlation, they have large and opposing effects on the cyclical volatility of the labor market. We conclude that the cross-sectional moment is important for business cycle analysis and that more work is needed to distinguish the potential mechanisms that can generate it.

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

*We thank Hyun-Tae Kim for excellent research assistance. This work is supported by grants from the National Research Foundation of Korea funded by the Korean Government (NRF-2016S1A5A2A03926178). Part of this work has been accomplished while Sun-Bin Kim was visiting the Institute of Social and Economic Research at Osaka University, who is grateful for their hospitality. Correspondence: Chang: yohg@snu.ac.kr; Kim: sunbin.kim@yonsei.ac.kr; Kwon: kwonkh98@gmail.com; Rogerson: rdr@princeton.edu.

1. Introduction

The modern business cycle literature emphasizes the need to construct models that contain an explicit description of agents' objective functions and the constraints that they face, as well as a notion of equilibrium. One of the benefits of following this approach was the promise that such models could be connected to micro data to provide discipline regarding functional forms and parameter values. The early models in this literature, such as the seminal work of Kydland and Prescott (1982), assumed a representative household, and thus precluded any possibility to connect the model with cross-sectional micro data. But in the last two decades, advances in computational methods have allowed researchers to develop aggregate business cycle models that feature heterogeneous households, thus allowing researchers to study aggregate phenomena while simultaneously connecting with rich micro data sets. In this paper we take advantage of these advances to address a somewhat long standing tension between “micro” and “macro” models of labor supply related to the relative magnitude of income and substitution effects.

Dating back at least to the work of Lucas and Rapping (1969), macroeconomists have argued that aggregate time series data on hours of work and real wages suggest offsetting income and substitution effects. In particular, it is argued that because average hours per individual have displayed little or no trend at the same time that average real wages have increased by several multiples, that the real wage must have no effect on desired labor supply, i.e., that income and substitution effects must offset each other.

Because this condition—offsetting income and substitution effects—also allows one to find a balanced growth path equilibrium in the one sector growth model, the macro literature typically refers to preferences with this property as “balanced growth preferences”.¹

¹A recent paper by Boppart and Krusell (2016) argues that hours have been trending down and that one can have a balanced growth path with this property if income effects dominate. Ohanian et al. (2008) argue that most of the decline in hours can be attributed to changes in tax and transfer policies without departing from the assumption of offsetting income and substitution effects.

Preferences with this property are a standard feature of modern business cycle analyses. At the same time, labor supply analyses based on micro cross-sectional data often conclude that income and substitution effects are not perfectly offsetting, instead finding evidence for specifications in which income effects dominate substitution effects. See for example, the analyses in Pijoan-Mas (2006) and Heathcote et al. (2014).

In this paper we revisit this tension in a state of the art heterogeneous agent business cycle model. The starting point for our analysis is the recent model of Chang et al. (2019). This paper extends the heterogeneous agent-indivisible labor models of Chang and Kim (2006, 2007) to allow for active adjustment along both the intensive and extensive margins, implying that in equilibrium there are non-degenerate distributions for both hours worked and wages for workers in the cross-section. As is common in much of the heterogeneous agent macro literature, Chang et al. (2019) assumes a single dimension of exogenous heterogeneity, namely idiosyncratic productivity shocks. However, assuming balanced growth preferences in this environment leads to the prediction of a large positive correlation between hours and wages in the cross-section. This correlation is at odds with the data—which reveals a correlation near zero—and is intimately related to the tension between “micro” and “macro” studies of labor supply noted above. As we discuss in detail below, one common practice to achieve a near zero correlation between wages and hours in this setting is to adopt preferences under which income effects dominates substitution effects. This is the route taken by both Pijoan-Mas (2006) and Heathcote et al. (2014).

It is perhaps not surprising that a model in which idiosyncratic productivity shocks are the sole source of randomness generates a cross-sectional correlation between wages and hours that is at odds with the data. For this reason we extend the model of Chang et al. to include an additional source of randomness—idiosyncratic shocks to the opportunity cost to market work, which we model as a shock to preferences. This shock could proxy

for either a shock to the value of leisure or to home productivity. Although we cannot measure this shock directly, we infer the size of this shock by calibrating it to match the cross-sectional dispersion in hours worked. Adding a source of idiosyncratic randomness that is uncorrelated with productivity does serve to reduce the cross-sectional correlation between wages and hours worked, but it remains large and positive relative to the data.

Starting from this benchmark with two uncorrelated sources of idiosyncratic shocks we develop two extensions that can generate a cross-sectional correlation between hours and wages that is near zero. The first extension maintains a zero correlation between the two idiosyncratic shocks but departs from balanced growth preferences. The second extension maintains balanced growth preferences but allows the two idiosyncratic shocks to be correlated. We show that if the two shocks are sufficiently highly positively correlated one can generate a cross-sectional correlation near zero. Intuitively, these results suggest that there is a continuum of specifications that can generate a near zero cross-sectional correlation between wages and hours, with there being a trade-off between a more positive correlation between the two shocks and the extent to which the income effect dominates the substitution effect.

Having isolated specifications that are consistent with the cross-sectional relationship between wages and hours, we next ask whether matching this moment is important for business cycle properties. To answer this question we expose the benchmark model as well as both extensions just described to the same stochastic process for aggregate productivity shocks and compare the implied business cycle properties across the three models. Both extensions imply first order effects for the magnitude of labor market fluctuations, implying that matching the cross-sectional correlation between wages and hours is indeed important. But most importantly, the two extensions imply affect the magnitude of labor market fluctuations in opposite directions: whereas maintaining balanced growth preferences but imposing a high positive correlation between the two shocks implies significantly

larger fluctuations, maintaining a zero correlation between the shocks and departing from balanced growth preferences results in significantly smaller fluctuations. Comparing the two extensions, we find that labor market fluctuations are more than twice as large in the case with balanced growth preferences and correlated shocks.

The correlation of the two shocks offers a simple economic interpretation regarding the elasticity of aggregate labor supply. A high correlation between the two shocks (productivities in the market and non-market activities) leads to a weak cross-sectional comparative advantage in market work relative to non-market work, whereas a negative correlation between the two shocks represents an economy with a strong comparative advantage. The economy with strong comparative advantage exhibits much less movement of labor between the market and home.

As noted above, our paper is closely related to several earlier papers, including Chang and Kim (2006, 2007), Pijoan-Mas (2006), Heathcote et al. (2014), Erosa et al. (2016) and Chang et al. (2019). Here we note the key differences not already noted. Both Pijoan-Mas (2006) and Heathcote et al. (2014) assume that all hours adjustment occurs along the intensive margin, and neither studies business cycle fluctuations. Pijoan-Mas (2006) features a single idiosyncratic shock. While Heathcote et al. (2014) allow for two idiosyncratic shocks, they only study the case in which they are assumed to be uncorrelated. Like us, Erosa et al. (2016) study a model that features labor adjustment along the intensive and extensive margin as well as with multiple sources of idiosyncratic shocks. Their model is richer than ours along several dimensions but they do not focus on the cross-sectional correlation between wages and hours and its significance for aggregate fluctuations.

An outline of the paper follows. In Section 2, we develop our benchmark model. We start by introducing the model from Chang et al (2019) where the cross-sectional heterogeneity solely arises from the realization of idiosyncratic productivity shocks. We

refer to this as a one-shock model. A calibrated version of this model implies a large positive cross-sectional correlation between wages and hours. We then extend the Chang et al. (2019) model to include a shock that is orthogonal to productivity and that affects hours. We refer to this as a two-shock model and identify it as our benchmark model. In Section 3, we extend the benchmark model in two different directions to match the cross-sectional correlation between hours and wage in the data. Section 4 compares the business-cycle statistics generated from the benchmark model as well as both extensions from Section 3. Section 5 concludes.

2. Aggregate Labor Supply in a Heterogeneous Agent Model

We develop a benchmark model for simultaneously studying labor supply at both the aggregate level and in the cross-section. We proceed in two steps. We begin by describing the one-shock model of Chang et al. (2019), which is a generalization of the standard heterogeneous agent-incomplete markets model that has become one of the benchmark models in macroeconomics. Notably, this model has a single source of exogenous randomness at the individual level, namely productivity shocks. Having introduced this model we note its inability to account for the cross-sectional correlation between wages and hours of work found in the data. This failure motivates us to extend the model to allow for a second idiosyncratic shock that influences hours of work and is orthogonal to productivity (two-shock model in Section 3). While this extension generates a cross-sectional correlation between wages and hours that is closer to that found in the data, the gap between model and data remains very substantial.

2.1. One-Shock Model: Chang et al. (2019)

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.

Individuals are 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, but realizations are iid across individuals. In particular, we will assume that z_{it} follows an AR(1) process in logs:

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

More generally, one might want to consider a specification which features both permanent and multiple random components with varying persistence, as well as a deterministic time varying components. Our more parsimonious specification is attractive in part for its tractability, but it is also able to capture elements of such a more general specification. If ρ_z is large, implying a large amount of persistence, then in a short panel some of the cross-sectional variation would be identified as permanent fixed effects. The mean-reversion process would imply a certain amount of deterministic dynamics. While we could feasibly consider more complicated processes, we follow much of the literature in choosing this simple specification as we feel it facilitates transparency.

In order to generate adjustment along both intensive and extensive margins we assume a nonconvexity in the mapping from time devoted to work to the resulting efficiency units of labor. We view this non-convexity as capturing factors such 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 $z_{it}g(h_{it})$ units of labor services. Following Prescott et al. (2009) and Rogerson and Wallenius (2009), we assume that $g(\cdot)$ takes the following simple form:²

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

where $0 < \hat{h} < 1$.

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

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

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

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

$$a_t \geq \underline{a}$$

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 $w_t z_{it} g(h_{it})$. The period budget equation for an individual is:

$$c_{it} + a_{it+1} = w_t z_{it} g(h_{it}) + (1 + r_t) a_{it}$$

²French (2005) considers an alternative and smoother specification in which the wage per unit of time is smoothly increasing in the number of hours worked.

In this section we focus on a steady state equilibrium, in which the two prices w_t and r_t are constant, as is the distribution of individual state variables (a_{it}, z_{it}) .

2.2. Calibration

In this subsection we describe a simple calibration procedure that can be used to assign all of the model's parameters. As is standard in the business cycle literature, we set a time period equal to one quarter. Our calibration procedure will not pin down a value for the preference parameter γ ; instead this parameter will be set up front and then the calibration procedure will identify all of the other parameters given this choice.

A sizable literature has estimated processes for idiosyncratic wage shocks using annual data, typically for males, and we use estimates from this literature to parameterize our idiosyncratic productivity shock process. Papers from this literature include 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, we set $\rho_z = 0.975$ and $\sigma_z = 0.165$.³ The Cobb-Douglas parameter α and the depreciation rate are set to standard values in the literature: α is set to 0.64 and δ is set to .025, effectively targeting labor's share of income and the investment to output ratio. Note that none of these values are influenced by the choice of γ .

There are four additional parameters to calibrate: β , B , \underline{a} , and \hat{h} . We set $\underline{a} = 0$ independently of γ . Given a value for γ and our previous choices, we choose the values of the remaining three parameters to match three moments in the steady state equilibrium: a (quarterly) rate of return to capital of 1%, an employment rate of 70%, and average

³The specific values that we adopt correspond to the estimates in Floden and Linde (2001). Because these values are based on annual data, we need to convert them to quarterly values for our calibration. We choose values for the quarterly process such that when we estimate an annual process it delivers the appropriate values for persistence and the standard deviation of the innovations.

hours (conditional on working) equal to $1/3$.⁴

The steady state properties of the above model were analyzed extensively in Chang et al. (2019), so we refer the reader there for a more detailed analysis. For each of several values of γ , they examined the implications for the steady state distributions of hours worked, earnings and wealth, and the transition of individuals both between non-employment and employment and within the annual hours worked distribution. The main message was that for a wide range of values of γ , going from 0.25 to 1.5, the calibrated model did a reasonable job of accounting for the patterns in the data.

For concreteness and future reference, we reproduce the results from Chang et al. (2019) for the case in which $\gamma = 1$. Table 1 shows the full set of calibrated parameters and Table 2 shows the ability of the model to account for the cross-sectional dispersion in hours, earnings and wealth.⁵

Table 1: Calibration for One-Shock Model

α	β	γ	δ	B	ρ_z	σ_z	\hat{h}
0.36	0.97656	1.00	0.25	18.9	0.975	0.165	0.151

Table 2: Cross-Sectional Dispersion Measures

	Model	Data
S.D. of Annual Hours	0.32	0.45 (CPS)
Earnings Gini	0.59	0.63 (SCF)
Wealth Gini	0.71	0.78 (SCF)

Note that this one-shock model in which productivity shocks are calibrated to the data does not generate sufficient cross-sectional dispersion in hours worked relative to the data. We will return to this issue in the next section when we introduce an additional

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.

⁵When we compute the standard deviation of hours in the model we normalize the level of hours to match the mean level of annual hours in the data.

shock.⁶ The earnings Gini in the model is modestly smaller than in the data; additional dispersion in hours would potentially address this problem as well. While this model cannot generate the extreme concentration of wealth among the top 1% that we see in the data, it nonetheless is able to capture a large amount of the wealth dispersion as measured by the Gini coefficient. As noted earlier, Chang et al. (2019) reports on the ability of the model to account for the data along several other dimensions. For example, it does a reasonable job of capturing the flow of workers between employment and non-employment, as well as within the hours worked distribution.

2.3. Hours and Wages in the Cross-Section

Although the calibrated model does a reasonable job of capturing many salient features of the cross-sectional data, there is one dimension along which it performs poorly: the cross-sectional correlation between hours worked and wages. In the calibrated economy with $\gamma = 1$ this value is 0.78, but in the data this value is around zero. Controlling for age and education, we estimate that this correlation is equal to -0.04 using data from the PSID and equal to 0 using data from the SIPP.⁷

In considering this correlation in the data it is important to be aware of the potential effects of measurement error, since measurement error will tend to bias the correlation to zero. Heathcote et al. (2014) estimate the amount of measurement error in earnings and hours as part of their exercise and find that it does not create a sufficiently large negative bias to reconcile the model with the data. Given the raw correlations reported above, the presence of measurement error suggests that the “true” correlation is modestly positive. Given the model implied value of 0.78 for this correlation, it seems of second

⁶Chang et al. (2019) shows that one can match the standard deviation of hours worked in the data with a one-shock model by some combination of increasing the standard deviation of the shock innovations and increasing the magnitude of γ .

⁷Pijoan-Mas (2006) and Heathcote et al. (2014) also report values for this correlation that are close to zero.

order importance to consider the issue of whether the true correlation is 0, 0.1 or 0.2. For concreteness, in what follows we will take 0 as the target value for this correlation, since this represents the most challenging case.

Although the value for the cross-sectional correlation between wages and hours just reported was for the specific case of $\gamma = 1$, we note that the model implied correlation reflects a very robust mechanism. With so-called balanced growth preferences there is no intratemporal effect of wages on hours of work controlling for wealth. But, transitory fluctuations in individual productivity induce intertemporal substitution effects that generate a positive correlation between wages and hours controlling for wealth. So in fact, the strong positive relation between wages and hours in the cross-section is not specific to any particular value of γ . For example, if we instead chose $\gamma = 0.5$ or $\gamma = 1.5$ the implied values for this correlation would be 0.82 and 0.76 respectively.

2.4. A Two-Shock Benchmark Model

The model of Chang et al. (2019) assumed that market productivity was the only source of exogenous heterogeneity across individuals. It is natural to think that allowing for an additional source of randomness that influences hours and is orthogonal to productivity could lead to substantial improvement in terms of the model's ability to account for the cross-section correlation between hours and wages. In this subsection we extend the previous model to include a second dimension of exogenous cross-sectional heterogeneity, namely heterogeneity in the opportunity cost of market work.⁸

The only change relative to the previous model is that we will now write preferences as:

⁸Heathcote et al. (2014) also study a two shock model. While their model is richer than ours along some dimensions (e.g., they model life cycle choices), all of the labor supply adjustment in their model occurs along the intensive margin. Erosa et al. (2016) consider a life cycle model with multiple shocks and adjustment along the intensive and extensive model.

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

where all parameters are as before, and x_{it} is an idiosyncratic preference shock affecting the disutility of working. We assume that this shock follows the same AR(1) process in logs for all individuals, with realizations iid across individuals:

$$\log x_{it+1} = \rho_x \log x_{it} + \varepsilon_{xit+1}, \quad \varepsilon_{xit} \sim N(0, \sigma_x^2).$$

This shock will capture the role of non-wage factors in generating cross-sectional dispersion in hours of work. These non-wage factors could reflect heterogeneous preferences for leisure, or more generally heterogeneous value for time spent in home production. At this point we will impose that the two shocks are uncorrelated, and will refer to this economy with two uncorrelated idiosyncratic shocks as our benchmark model.

2.5. The Role of Comparative Advantage

It is useful to develop some intuition to better understand the results that will be presented later in the paper. In particular, consider an infinitely lived individual with period utility function:

$$\log(c_t) - \frac{1}{1+1/\gamma} x_t h_t^{1+1/\gamma}.$$

Suppose that the discount factor is β and period labor income is equal to $wh_t z_t$. Assume that w is constant over time, that the individual can borrow and lend at an interest rate equal to $(1/\beta) - 1$, and that the values of x and z vary over time in a deterministic fashion. Suppose also that there is a utility cost associated with participation in each period and that this cost is represented as $\bar{u}x_t$, so that the cost is scaled up by the preference shock x_t .

This utility maximization problem is not identical to our model. But it has some of the key features from our model—a nonconvexity that generates adjustment along both the intensive and extensive margins, and two sources of idiosyncratic variation—and permits an analytic characterization of the optimal decision rules for labor supply. For this reason it is useful in providing some valuable intuition about how the two sources of idiosyncratic variation affect labor supply decisions.

Deriving first order conditions for the individual’s optimal choices of consumption, hours and participation over time, straightforward manipulation yields that consumption will be constant over time, that the optimal choice for h_t conditional on working satisfies:

$$h_t = \left(\frac{w}{c}\right)^\gamma \left(\frac{z_t}{x_t}\right)^\gamma \quad (1)$$

and the decision about whether to work is determined by a reservation value for the ratio z_t/x_t .

A key property is that the ratio z_t/x_t serves as a sufficient statistic to characterize labor supply decisions along both the intensive and extensive margin. The ratio z_t/x_t can be interpreted as a measure of comparative advantage in market work. In contrast, absolute advantage in market work is captured by z_t , given that the wage per efficiency unit w is assumed to be constant. Importantly, this model says that labor supply decisions along both margins (intensive and extensive) are driven by comparative advantage in market work rather than absolute advantage in market work. In the special case in which z_t is always just a positive multiple of x_t , the value of z_t/x_t is constant and hours of work will be constant despite changes in the observed wage.

A key insight from the indivisible labor literature (see, for example, Chang and Kim (2006)) is that the responsiveness of the aggregate employment rate depends upon the density of individuals in the vicinity of the reservation wage. Loosely speaking, the more dispersion there is in the relevant individual state variables, the less responsive is the

extensive margin. The fact that z_t/x_t is a sufficient statistic for the extensive margin implies that the extent of dispersion in z_t/x_t is key to determining responsiveness along the extensive margin. Importantly, if z_t and x_t move in the same direction then the dispersion in z_t/x_t will tend to be less than the dispersion in z_t alone. Conversely, if x_t tends to move in the opposite direction of z_t then the dispersion in z_t/x_t will tend to be larger than the dispersion in z_t alone.

2.6. Calibration of the Two-Shock Model

Much of the calibration proceeds as before, but there are now two new parameters to be determined: ρ_x and σ_x . In what follows we will set the persistence parameter for the idiosyncratic preference shock equal to the same value that we used for the idiosyncratic productivity shock: $\rho_x = 0.975$.⁹ This leaves us with one additional parameter that needs to be calibrated, namely σ_x . To accomplish this we add one additional moment to our list of moments to be matched, as we now require that the model generates the same dispersion in annual hours worked as found in the data. Importantly, although it is not possible to directly measure the size of the second shock, we can discipline the size of this shock by requiring that the model generates the appropriate dispersion in hours of work. Table 3 shows the calibrated parameter values for this model. In what follows we will refer to this as our benchmark model.

Table 3: Calibration for Two-Shock Model (Benchmark)

α	β	γ	δ	B	ρ_z	σ_z	ρ_x	σ_x	\hat{h}
0.36	0.97678	1.00	0.25	20.6	0.975	0.165	0.975	0.103	0.133

⁹There is nothing special about the case in which $\rho_z = \rho_x$. We choose a highly persistent process based on the fact that empirical studies of micro data tend to find evidence for fixed effects, and in a short panel our idiosyncratic shock would at least partially be captured by fixed effects.

2.7. Steady State Properties of the Benchmark Model

The model with two dimensions of exogenous heterogeneity looks very similar to the benchmark model along the dimensions previously reported in Table 2, though by construction it now achieves the same dispersion of hours worked as found in the data. The two-shock model also performs quite similarly to the one shock model along the other dimensions that Chang et al. (2019) documented. In particular, the transition of individuals between employment and non-employment, as well as within the hours worked distribution capture some key patterns from the data, as Table 4 below shows.

Table 4: Annual Hours Transition

		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
		Model					
		$t + 1$					
		Not Work	1st	2nd	3rd	4th	5th
t	Not Work	72.49	17.64	6.53	2.39	0.81	0.14
	1st	21.12	34.43	24.53	12.60	5.61	1.72
	2nd	8.46	20.54	37.54	23.19	7.88	2.39
	3rd	3.24	11.03	17.67	37.28	24.68	6.09
	4th	1.69	6.99	7.57	19.31	39.95	24.48
	5th	0.49	3.57	3.83	5.37	20.49	66.25

Notes: Transition matrix of annual hours worked by quintile groups and non-employment ('Not Work').

Although the model captures the key patterns found in the empirical transition matrix, we note that overall, it does not exhibit as much persistence as in the data—relative to

the data the model has too many transitions from non-employment to employment, and too many transitions out of each hours quintile.¹⁰

But, most importantly for our purposes is the implication of the extended model for the cross-sectional correlation between wages and hours of work. This value is now equal to 0.53. Although the effect of adding an idiosyncratic shock orthogonal to individual productivity has the intuitive effect of decreasing this correlation relative to the model with only idiosyncratic productivity shocks, the quantitative magnitude of this effect is not sufficiently large to allow the model to match the data. We note again that this would remain true if we adopted a modestly positive correlation as the true target. In the next section we explore two extensions to this benchmark model that can help us to match this empirical correlation.

3. Two Extensions

In this section we generalize the benchmark model along two dimensions. First, we generalize the utility function to now read:

$$\sum_{t=0}^{\infty} \beta^t \left[\frac{c_{it}^{1-\sigma}}{1-\sigma} - Bx_{it} \frac{h_{it}^{1+1/\gamma}}{1+1/\gamma} \right].$$

Our benchmark model corresponds to the limiting case in which σ tends to one. Given this class of preferences, this limiting case is the only specification that is consistent with balanced growth, i.e., that delivers perfectly offsetting income and substitution effects. If σ is greater than one, then the income effect dominates the substitution effect, but if σ is less than one, then the substitution effect dominates the income effect.

The second generalization is to allow the two idiosyncratic shocks, z_{it} and x_{it} , to be correlated. In particular, we assume that the two innovations ε_{zt} and ε_{xt} are correlated

¹⁰Allowing for greater persistence of the two shocks would presumably improve the fit of the model along this dimension.

and denote the value of this correlation as ρ_{zx} . In this section we show that each of these departures individually can reconcile the model-implied cross-sectional correlation between wages and hours with the value from the data.

3.1. Extension I: Departure from Balanced Growth Preferences

In this subsection we ask whether maintaining $\rho_{zx} = 0$ as in the benchmark model, we can find a value of σ such that the model will match all of the previous targets as well as generate a cross-sectional correlation between wages and hours of approximately zero. Pijoan-Mas (2006) and Heathcote (2014) pursue this same strategy for targeting the cross-sectional correlation, though Pijoan-Mas (2006) had only productivity shocks and both papers assumed that all hours adjustment occurred along the intensive margin. Assuming $\gamma = 1$ as previously, it turns out that setting $\sigma = 2$ will achieve this in our model economy. This is slightly larger than the 1.45 value adopted by Pijoan-Mas (2006) or the 1.78 value adopted by Heathcote et al. (2014) to achieve the same goal—a near zero correlation.¹¹ Table 5 shows the full set of calibrated parameters.

Table 5: Calibration of Extension I Model

α	β	γ	δ	B	ρ_z	σ_z	ρ_x	σ_x	\hat{h}	σ
0.36	0.96946	1.00	0.25	77.0	0.975	0.165	0.975	0.144	0.144	2

This result that σ needs to be greater than one is qualitatively intuitive. We observed previously that when $\rho_{zx} = 0$, our two-shock model still generates a significantly positive correlation between hours and wages in the cross-section. When we deviate from balanced growth preferences, the effect on the cross-sectional correlation between wages and hours depends on the value of σ . If σ is less than one then substitution effects dominate income

¹¹As noted earlier, there are various other differences between our paper and these two which make could potentially impact the value of σ . For example, Heathcote et al (2014) study a life cycle model with insurable and uninsurable shocks. Pijoan-Mas (2006) has a different borrowing constraint but only productivity shocks. And both of these papers assume that all labor supply adjustment occurs along the intensive margin.

effects, thereby generating a positive correlation. But, if σ is greater than one, income effects dominate substitution effects and the effect on the correlation is negative.

In addition to hitting all of the targeted moments, this specification performs about the same as the benchmark model in terms of the other moments we have considered. (See Appendix A2).

3.2. Extension II: Departure from Uncorrelated Shocks

In this subsection we ask whether maintaining balanced growth preferences as in the benchmark model, we can find a value of ρ_{zx} such that the model will match all of the previous targets and also generate a cross-sectional correlation between wages and hours of roughly zero. Assuming $\gamma = 1$ as previously, it turns out that setting $\rho_{zx} = 0.9$ will achieve this. Table 6 shows the full set of calibrated parameters.

Table 6: Calibration of Extension II Model

α	β	γ	δ	B	ρ_z	σ_z	ρ_x	σ_x	\hat{h}	ρ_{zx}
0.36	0.9814	1.00	0.25	26.2	0.975	0.165	0.975	0.147	0.133	0.9

This result is also qualitatively intuitive. To see why recall equation (1). This equation illustrates that a positive correlation between z and x will serve to dampen the positive effect of wages on hours of work. Put somewhat differently, in our simple example it is comparative advantage that matters for market hours, and a high positive correlation between the two shocks tends to decrease the correlation between market wages and comparative advantage. Conversely, if the two shocks are negatively correlated then the correlation between market wages and comparative advantage in market work tends to be very high.

3.3. Discussion

The two previous subsections demonstrate two different extensions of the benchmark model that are able to generate a near zero correlation between wages and hours in the cross-section while still accounting for the cross-sectional dispersion in hours and earnings. Here we draw two generalizations that follow.

First, we introduced two new parameters relative to the benchmark model and showed that each of them individually could be adjusted to as to hit the desired targets. It necessarily follows that one could find a continuum of parameterizations in which the two new parameters are varied together such that all of the desired targets are hit. In the next section we will continue to focus on the two specifications in which only one of the parameters is varied in order to highlight the effects of the two distinct channels.

Second, although we have chosen to target a cross-sectional correlation of zero between wages and hours, it should be clear that targeting a modest positive value for this correlation would simply require a modestly smaller adjustment for the two parameters being varied. For example, we note that for Extension II, reducing the targeted correlation to 0.34 would reduce the required value of ρ_{zx} from 0.9 to 0.5. Table 7 provides information on the relationship between the correlation of the two shocks and the resulting cross-sectional correlation between wages and hours.¹²

Table 7: ρ_{zx} and $corr(w, h)$

	-0.90	-0.50	0	0.50	0.90
ρ_{zx}					
$corr(w, h)$	0.78	0.66	0.53	0.34	-0.004

¹²See Appendix Tables A1 and A2 for the calibrated parameters and other cross-sectional moments from the models with various ρ_{zx} .

4. Implications for Business Cycles

In this section we examine the extent to which the cross-sectional correlation between wages and hours is an important moment from the perspective of business cycle fluctuations, and if so, if it matters how one chooses to match this moment. That is, we examine the extent to which our three models (Benchmark, Extension I, and Extension II) differ in their implications for business cycles.

We pursue this by considering how all three models behave when confronted with the same process for aggregate productivity shocks. That is, we now assume that the aggregate production function is:

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

where Z_t is the aggregate productivity shock and is assumed to follow an AR(1) process in logs:

$$\log Z_{t+1} = \rho_Z \log Z_t + \varepsilon_{Zt+1}$$

and ε_{Zt} is assumed to be log normally distributed with mean zero and standard deviation σ_Z . We adopt commonly used parameter values for the technology shock process, $\rho_Z = 0.95$ and $\sigma_Z = 0.007$.

A recursive representation of equilibrium implies that the aggregate state variable will include the measure of individuals over idiosyncratic shocks and asset holdings. To solve for the equilibrium with such a high dimensional aggregate state variable we employ Krusell and Smith's (1998) "bounded rationality" method, which approximates the distribution of workers over individual states by a limited number of its moments. In particular, we assume that agents make use of the average asset holdings of the economy

as well as the aggregate technology shock Z .¹³ We generate 3,000 quarterly periods for a model economy. After dropping the first 1,000 observations, we take logs and apply a Hodrick-Prescott filter (with smoothing parameter 1,600 to be comparable to those from the data) to produce the business cycle statistics.

Table 8 displays the results for several business cycle statistics of interest. Apart from aggregate output we focus entirely on labor market statistics for the simple reason that the models all imply similar properties for the behavior of consumption and investment relative to output. We also include results from the Chang et al. (2019) specification that we described earlier, in which the idiosyncratic productivity shock is the sole source of idiosyncratic randomness.

Table 8: 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	0.89	1.51	0.48	–	0.98
Chang et al. (2019)	1.74	0.79	0.45	0.69	0.12	1.09	0.83
Benchmark ($\sigma = 1, \rho_{zx} = 0$)	1.65	0.72	0.44	0.58	0.17	0.93	0.85
Extension I ($\sigma = 2, \rho_{zx} = 0$)	1.46	0.39	0.28	0.44	0.09	0.72	0.66
Extension II ($\sigma = 1, \rho_{zx} = 0.9$)	1.92	0.93	0.48	0.95	0.11	1.47	0.84

Notation in this table is as follows: Y denotes aggregate output, H denotes aggregate hours, E denotes aggregate employment, h denotes hours per worker, L denotes aggregate efficiency units of labor, and w denotes average wage per hour worked. There are several messages to point out. First, we note that that introducing a second shock into the framework of Chang et al. (2019) has a modest but not insignificant effect on the extent of labor market fluctuations. Fluctuations in both aggregate hours and aggregate efficiency units are reduced by around ten percent relative to the one shock model of Chang et al. (2019). This is consistent with the intuition that we described previously. Adding an additional shock that is uncorrelated with the individual productivity shock effectively

¹³Note that because we have a continuum of agents, the joint distribution of idiosyncratic shocks (z, x) is necessarily stationary, which is why this distribution does not enter the aggregate state vector.

increases the amount of heterogeneity in the cross-section and as a result leads to fewer marginal individuals. This results in smaller fluctuations along the extensive margin, and smaller fluctuations in the aggregate, though the effect of smaller fluctuations along the extensive margin is at least partially offset by greater fluctuations along the intensive margin.

The second message from Table 8 is that modifying the benchmark model so as to account for the cross-sectional correlation between hours and wages has first order effects on the resulting business cycle statistics. When we move from the benchmark model to Extension I there is a substantial decrease in the size of labor market fluctuations. In this case the decrease occurs along both the intensive and the extensive margin. The intuition for this effect is straightforward: with $\sigma = 2$ the income effect dominates the substitution effect in terms of intratemporal effects of wages on labor supply. As a result, the intertemporal substitution effects are now dampened, and there is less response of labor supply to changes in wages. This effect operates on both the intensive and extensive margins.

Next consider moving from the benchmark model to Extension II. In this case we observe a first order increase in the magnitude of fluctuations in labor market variables. For example, the standard deviation of fluctuations in total hours increases from 0.72 to 0.93 and the standard deviation of efficiency units of labor increases from 0.93 to 1.47. These increases are accounted for by greater fluctuations along the extensive margin, as fluctuations along the intensive margin actually decrease somewhat.

The intuition for why we observe increased fluctuations in this case is straightforward and draws on our earlier heuristic analysis of the two shock model that suggested it is the ratio of the two shocks that matters. Fluctuations along the extensive margin are determined by the mass of individuals who are close to indifference between working and not working. Our earlier analysis argued that when the two shocks are negatively

correlated, the ratio of the two shocks has much larger variance than the z shock alone. And if the two shocks are positively correlated the effective shock has smaller dispersion than the z shock alone.

To better illustrate this intuition, Table 9 below repeats the business cycle analysis for several alternative values of ρ_{zx} .

Table 9: Effect of ρ_{zx} on Business Cycle Statistics

	σ_Y	σ_H	$\frac{\sigma_H}{\sigma_Y}$	σ_E	σ_h	σ_L	σ_w
$\rho_{zx} = -0.9$	1.59	0.66	0.41	0.47	0.22	0.83	0.87
$\rho_{zx} = -0.5$	1.62	0.68	0.42	0.51	0.19	0.87	0.86
$\rho_{zx} = 0$ (Benchmark)	1.65	0.72	0.44	0.58	0.17	0.93	0.85
$\rho_{zx} = 0.5$	1.72	0.77	0.45	0.68	0.14	1.04	0.83
$\rho_{zx} = 0.9$	1.92	0.93	0.48	0.95	0.11	1.47	0.84

Consistent with the intuition offered, we see that overall fluctuations are monotonically increasing as we move from $\rho_{zx} = -0.9$ to $\rho_{zx} = 0.9$. Note however, that fluctuations along the intensive margin are actually decreasing at the same time that overall fluctuations are increasing. The reason for this is that the intensive and extensive margins are substitutes from the perspective of adjusting overall labor input. When there are lots of marginal workers, i.e., a large mass of workers who are close to indifferent being working and not working, then it is relatively easy to adjust total labor input by adjusting along the extensive margin. This reduces the need to make adjustment along the intensive margin.

Returning to Table 8, the third, and we think most dramatic message is that these two different ways of reconciling the benchmark model with the cross-sectional evidence lead to dramatically different implications for business cycle fluctuations. If we set the correlation of the two idiosyncratic shocks to zero and depart from balanced growth preferences, the resulting fluctuations in labor market variables are less than half as large as they are in the case where we maintain balanced growth preferences but allow the two idiosyncratic shocks to be highly positively correlated.

5. Conclusion

Recent advances in modeling aggregate labor supply allow us to now study business cycles in models that also have a rich set of implications for cross-sectional relationships. In this paper we leverage these advances in order to address a tension between the implications of standard macro models that impose balanced growth preferences and the near zero cross-sectional correlation between wages and hours found in the data. Commonly used heterogeneous agent macro models imply a large positive value for this correlation.

We ask whether the failure to match this correlation is substantively important for inferring implied business cycle properties of these models. Our answer to this question is affirmative. But more importantly, we show that what really matters is the mechanism through which one achieves a realistic cross-sectional correlation between hours and wages. We consider two different mechanisms, and find that they have large and opposite effects on the magnitude of labor market fluctuations that result from a given aggregate shock. A key parameter that we highlight and that has received little attention in the existing literature is the correlation between idiosyncratic variation in market productivity and the opportunity cost of time. We think it is important for future work to pay more attention to this correlation—which reflects the cross-sectional comparative advantage between market and non-market activities—given its apparent importance for the properties of business cycles. We note one piece of recent work that speaks to this correlation. Boerma and Karabarbounis (2018) argue that market and home productivity are highly positively correlated.

References

- [1] Aiyagari, S. Rao. (1994) “Uninsured Idiosyncratic Risk and Aggregate Saving,” *Quarterly Journal of Economics* 109(3), 659–684.
- [2] Bewley, Truman. (1986) “Stationary Monetary Equilibrium With a Continuum of Independently Fluctuating Consumers,” in *Contributions to Mathematical Economics in Honor of Gerard Debreu*, 79–102.
- [3] Boerma, Job, and Karabarbounis, Loukas. (2018) “Inferring Inequality with Home Production,” University of Minnesota Working Paper.
- [4] Boppart, Timo and Krusell, Per. (2016) “Labor Supply in the Past, Present, and Future: a Balanced-Growth Perspective,” *NBER Working Paper* No. 22215.
- [5] Card, David. (1994) “Intertemporal Labor Supply: An Assessment,” in *Advances in Econometrics*, edited by Chris Sims, Cambridge University Press, 49–81.
- [6] Chang, Yongsung, and Kim, Sun-Bin. (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 Kim, Sun-Bin. (2007) “Heterogeneity and Aggregation: Implications for Labor-Market Fluctuations,” *American Economic Review* 97(5), 1939–1956.
- [8] Chang, Yongsung, Kim, Sun-Bin, Kwon, Kyoocho and Rogerson, Richard. (2019) “Individual and Aggregate Labor Supply in Heterogeneous Agent Economies with Intensive and Extensive Margins,” *International Economic Review* 60(1), 1–21.

- [9] Erosa, Andres, Fuster, Luisa, and Kambourov, Gueorgui. (2016) “Towards a Micro-Founded Theory of Aggregate Labor Supply,” *Review of Economic Studies* 83, 1001–39.
- [10] Floden, Martin, and Linde, Jesper. (2001) “Idiosyncratic Risk in the United States and Sweden: Is There a Role for Government Insurance?” *Review of Economic Dynamics* 4(2), 406–437.
- [11] French, Eric. (2005) “The Effects of Health, Wealth, and Wages on Labor Supply and Retirement Behavior,” *Review of Economic Studies* 72(2), 395–427.
- [12] Heathcote, Jonathan, Storesletten, Kjetil, and Violante, Giovanni. (2008) “Insurance and Opportunities: A Welfare Analysis of Labor Market Risk,” *Journal of Monetary Economics* 55, 501–525.
- [13] Heathcote, Jonathan, Storesletten, Kjetil, and Violante, Giovanni. (2014) “Consumption and Labor Supply with Partial Insurance: An Analytical Approach,” *American Economic Review* 104, 2075–126.
- [14] Hugget, Mark. (1996). “Wealth Distribution in Life-Cycle Economics,” *Journal of Monetary Economics* 38, 469–94.
- [15] Krusell, Per, and Smith, Anthony A. (1998) “Income and Wealth Heterogeneity in the Macroeconomy,” *Journal of Political Economy* 89(6), 1059–1085.
- [16] Kydland, Finn, and Prescott, Edward. (1982) “Time to build and aggregate fluctuations,” *Econometrica* 50, 1345–79.
- [17] Lucas, Robert, and Rapping, Leonard. (1969) “Real Wages, Employment, and Inflation,” *Journal of Political Economy* 77(5), 721–54.

- [18] Ohanian, Lee, Raffo, Andrea, and Rogerson, Richard. (2008) “Long-term Changes in Labor Supply and Taxes: Evidence from OECD Countries, 1956-2004,” *Journal of Monetary Economics* 55(8), 1353–1362.
- [19] Pijoan-Mas, Josep. (2006) “Precautionary Savings or Working Longer Hours,” *Review of Economic Dynamics* 9, 326–352.
- [20] Prescott, Edward, Rogerson, Richard, and Wallenius, Johanna. (2009) “Lifetime Aggregate Labor Supply with Endogenous Workweek Length,” *Review of Economic Dynamics* 12, 23–36.
- [21] Rogerson, Richard, and Wallenius, Johanna. (2009) “Micro and Macro Elasticities in a Life Cycle Model with Taxes,” *Journal of Economic Theory* 144, 2277–2293.

Appendix: Additional Statistics

Table A1: Calibration of Extension II Model

ρ_{zx}	β	B	\hat{h}	σ_x
-0.9	0.97537	19.0	0.128	0.0825
-0.5	0.976	19.2	0.130	0.09
0.5	0.97818	22.6	0.135	0.124
0.9	0.9814	26.2	0.139	0.147

Table A2: Wealth and Earnings

<u>Gini Coefficient: Wealth</u>					
PSID=0.76, SCF=0.78					
$\sigma = 2, \rho_{zx} = 0$	0.63				
$\rho_{zx} =$	-0.9	-0.5	0.0	0.5	0.9
	0.63	0.64	0.65	0.67	0.68
<u>Gini Coefficient: Earnings</u>					
PSID=0.53, SCF=0.63					
$\sigma = 2, \rho_{zx} = 0$	0.53				
$\rho_{zx} =$	-0.9	-0.5	0.0	0.5	0.9
	0.60	0.59	0.59	0.57	0.53
<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
$\sigma = 2, \rho_{zx} = 0$	0.07	2.19	9.80	24.77	63.17
$\rho_{zx} = -0.9$	0.16	2.66	9.51	23.80	63.88
$\rho_{zx} = 0.0$	0.08	2.10	8.75	23.20	65.87
$\rho_{zx} = 0.9$	0.04	1.16	7.54	22.51	68.76
<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
$\sigma = 2, \rho_{zx} = 0$	10.05	14.97	18.79	23.38	32.81
$\rho_{zx} = -0.9$	5.56	11.14	16.91	24.87	41.52
$\rho_{zx} = 0.0$	6.26	11.95	17.06	24.43	40.31
$\rho_{zx} = 0.9$	10.27	13.09	18.65	23.21	34.78