NBER WORKING PAPER SERIES

LABOR SUPPLY AND OCCUPATIONAL CHOICE

Andrés Erosa
Luisa Fuster
Gueorgui Kambourov
Richard Rogerson

Working Paper 30492
http://www.nber.org/papers/w30492

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
September 2022

We thank numerous conference and seminar participants for their comments. Fuster acknowledges financial support from MCIN/ AEI/10.13039/501100011033/ grant #PID2019-107614GB-I00. Kambourov acknowledges financial support from the Social Sciences and Humanities Research Council of Canada. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2022 by Andrés Erosa, Luisa Fuster, Gueorgui Kambourov, and Richard Rogerson. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Labor Supply and Occupational Choice
Andrés Erosa, Luisa Fuster, Gueorgui Kambourov, and Richard Rogerson
NBER Working Paper No. 30492
September 2022
JEL No. J22,J24,J31

ABSTRACT

We document a robust negative relationship between mean annual hours in an occupation and the dispersion of annual hours within that occupation. We study a unified model of occupational choice and labor supply that features heterogeneity across occupations in the return to working additional hours and show that it can match the key features of the data both qualitatively and quantitatively. Occupational choice in our model is shaped both by selection on comparative advantage and selection on tastes for leisure. Our quantitative work finds that the dominant source of differences in hours across occupations is selection on tastes for leisure.

Andrés Erosa
Universidad Carlos III de Madrid
Department of Economics
Calle Madrid 126, 28903
Getafe, Madrid
Spain
aerosa@eco.uc3m.es

Luisa Fuster
Universidad Carlos III de Madrid
Department of Economics
Calle Madrid 126, 28903
Getafe, Madrid
Spain
lfuster@eco.uc3m.es

Gueorgui Kambourov
Department of Economics
University of Toronto
150 St. George Street
Toronto, ON M5S 3G7
CANADA
g.kambourov@utoronto.ca

Richard Rogerson
Department of Economics
JRR Building Room 292
Princeton University
Princeton, NJ 08544
and NBER
rdr@princeton.edu

Labor Supply and Occupational Choice*

Andrés Erosa[†]
Universidad Carlos III de Madrid

Luisa Fuster[‡]
Universidad Carlos III de Madrid

Gueorgui Kambourov[§]
University of Toronto

Richard Rogerson[¶]
Princeton University

Abstract

We document a robust negative relationship between mean annual hours in an occupation and the dispersion of annual hours within that occupation. We study a unified model of occupational choice and labor supply that features heterogeneity across occupations in the return to working additional hours and show that it can match the key features of the data both qualitatively and quantitatively. Occupational choice in our model is shaped both by selection on comparative advantage and selection on tastes for leisure. Our quantitative work finds that the dominant source of differences in hours across occupations is selection on tastes for leisure.

JEL Classification: J22, J24, J31.

Keywords: Labor supply, occupational choice, inequality.

1 Introduction

Two classic topics in labor economics are labor supply and occupational choice. Interestingly, textbook treatments of labor supply abstract from occupational choice, and textbook

*This version: September 14, 2022. We thank numerous conference and seminar participants for their comments. Some of the material in this paper appears in our earlier NBER Working Paper (Erosa, Fuster, Kambourov and Rogerson (2017)), but was subsequently removed and does not appear in the published version of that paper (Erosa, Fuster, Kambourov and Rogerson (2022)). Fuster acknowledges financial support from MCIN/ AEI/10.13039/501100011033/ grant #PID2019-107614GB-I00. Kambourov acknowledges financial support from the Social Sciences and Humanities Research Council of Canada.

[†]Universidad Carlos III de Madrid, Department of Economics, Calle Madrid 126, 28903 Getafe, Madrid, Spain. E-mail: aerosa@eco.uc3m.es.

[‡]Universidad Carlos III de Madrid, Department of Economics, Calle Madrid 126, 28903 Getafe, Madrid, Spain. E-mail: lfuster@eco.uc3m.es.

[§]University of Toronto, Department of Economics, 150 St. George St., Toronto ON, M5S 3G7 Canada. E-mail: g.kambourov@utoronto.ca.

[¶]Princeton University, Department of Economics and Woodrow Wilson School, Princeton NJ 08544, USA. E-mail: rdr@Princeton.edu.

treatments of occupational choice typically take the time allocated to market work as given. In this paper we argue that it is important to consider these two topics jointly. To this end we make three contributions.

Our first contribution is to document a robust empirical relationship between the first and second moments of the hours worked distribution and occupational choice. Using data on annual hours of work at the 3-digit occupational level from the Current Population Survey (CPS) over the period 1976-2015, we document a robust negative correlation between mean hours worked in an occupation and the standard deviation of log hours worked within an occupation. This result is robust to controlling for demographic characteristics such as age, education and gender, and also holds across subperiods.

Goldin (2014) argues that an important difference across occupations is the extent of non-linearity in the relationship between hours and earnings.[1] Our second contribution is to demonstrate that heterogeneity across occupations along this dimension serves as the basis for a parsimonious explanation for the distribution of both hours and wages across occupations. We establish this result both analytically and quantitatively.

The core of our model is a standard two occupation Roy (1951) model that includes a time allocation decision. We extend it along two dimensions: heterogeneous preferences over leisure, and earnings that are convex in hours worked, with the extent of this convexity varying across occupations. In our model, hours worked differ across occupations both because of occupation effects and selection effects. The occupation effect reflects the fact that a given individual chooses to work longer hours when choosing the occupation with the more convex hours-earnings relationship. The selection effect captures the fact that differences in the extent of convexity also lead individuals to sort based on tastes for leisure: holding all else constant an individual with lower taste for leisure is more likely to choose the occupation with the more pronounced convexity. Selection on preferences distinguishes our model from a standard Roy model in which comparative advantage (i.e., relative productivity) is the sole determinant of occupational choice. Selection on preferences reflects an interaction between the heterogeneity in the convexity of the hours-earnings relationship across occupations and heterogeneity in tastes for work across individuals. Preference heterogeneity alone does not give rise to selection on preferences if the convexity is the same across occupations.

Our third contribution is to quantify the relative importance of occupation and selection effects for the distribution of hours across occupations. We find that the selection effect is dominant, and this result holds for the entire range of differences in convexity that we consider. We also show that selection on preferences is quantitatively important in shaping occupational choices.

Our paper is related to two classic literatures in labor economics: time allocation and occupational choice. Each of these is too vast for us to attempt any meaningful survey. We also relate to a literature that studies labor supply when earnings are a non-linear function of hours. Rosen (1976) and Moffitt (1984) are early examples of empirical studies that incorporate a non-linear hours-earnings function and emphasize its role for labor supply

[1]See the discussion in Erosa et al. (2022) for a summary of evidence on this point.

responses.[2] Our specification of this non-linearity follows French (2005).[3] The distinguishing feature of our paper is to embed this feature into a model of occupational choice in which the non-linearity varies across occupations.

Our paper is most closely related to Erosa et al. (2022). Both papers cover broadly similar terrain but their analyses are largely complementary.[4] First, while Erosa et al. (2022) documented the differences in mean hours worked across 3 digit occupations, it did not document the robust relationship between mean hours and the dispersion in hours that is the focus of this paper. Second, the theoretical analysis in this paper identifies a channel through which differences in the convexity of earnings across occupations will affect the dispersion in hours. Third, we show that our model captures the key patterns in the data for a wide range of differences in the extent of non-linearities across occupations. Fourth, we quantify the relative magnitude of occupation and selection effects. Lastly, we also examine the model's implications for differences between wage and earnings inequality across occupations.

An outline of the paper follows. Section 2 presents our key empirical findings on occupational choice and hours of work. Section 3 presents the simple benchmark model and illustrates its ability to qualitatively account for the salient facts about hours of work across occupations. Section 4 presents our quantitative analysis, and Section 5 concludes.

2 Empirical Facts

In this section we document systematic differences in the distribution of hours worked across three digit occupations, and their relation to differences in the distribution of wages.

2.1 Data

Our analysis is based on the IPUMS-CPS files from the 1976-2015 Current Population Survey (CPS).[5] The CPS provides information on number of weeks worked, usual hours per week, and annual wage and salary income. We construct annual hours as the product of weeks worked and usual weekly hours, and hourly wages are constructed by dividing wage and salary income in a calendar year by annual hours worked in that year. Nominal wages are converted to real wages using the CPI, with 1983 used as the benchmark year. We use the occupational classification provided in Autor and Dorn (2013) to construct consistent occupational codes for the 1976-2015 period.

We restrict our sample to individuals between the ages of 22 and 64. In order to match individuals to specific occupations, we only use observations for individuals that report

[2] See Barzel (1973) and Rosen (1978) regarding the general notion of wages that depend on hours. There is a large literature starting with Hausman (1985) on econometric estimation of models with nonlinear budget sets.

[3] Hornstein and Prescott (1993) show that this specification is the competitive equilibrium outcome of an economy with a production technology in which hours of work and number of workers are imperfect substitutes.

[4] As noted earlier, some of the material in the current paper was contained in Erosa et al. (2017), which is an early version of Erosa et al. (2022), but does not appear in the published version.

[5] The data and a detailed description can be found at http://cps.ipums.org/cps/. See Flood, King, Rodgers, Ruggles and Warren (2018).

having a single employer during the survey year. We drop observations with annual hours less than 250 or greater than 4500, or with a real hourly wage in the top and bottom 0.1% of the hourly wage distribution. Online Appendix A provides a detailed description of variables and sample restrictions.

We pool data across all years and bin the observations on annual hours and wages by occupation, and compute four values for each occupation: mean hours, the standard deviation of log hours, mean wages, and the standard deviation of log wages.

2.2 Occupational Differences in the Distribution of Hours Worked

Figure 1: The Log of Mean Annual Hours vs. the Standard Deviation of Log Annual Hours, CPS 1976-2015: by 3-Digit Occupations.

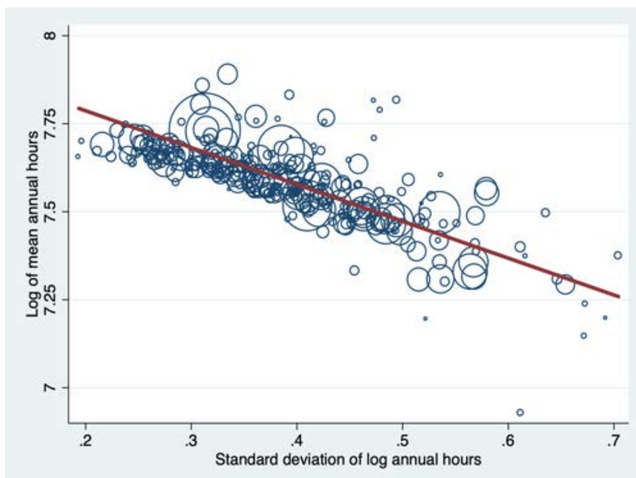

Notes: Each point represents a 3-digit occupation in the 1976-2015 time period. The scatter plot describes the relationship between log mean annual hours and the standard deviation of log annual hours. The straight solid red line represents a linear regression, weighted by the relative size of each occupation, while the size of the circle indicates the relative size of the occupation.

Figure 1 displays the key pattern for hours worked across occupations. The vertical axis measures the log of mean annual hours worked in an occupation while the horizontal axis measures the standard deviation of log annual hours in that same occupation. The straight line represents a linear regression, weighted by the relative size of each occupation, with the size of each circle indicating relative size. The figure illustrates a significant negative relationship between the mean hours worked in an occupation and the standard deviation of log hours worked in that occupation.[6]

Differences in mean hours worked across occupations are large; log mean hours range from 7.7 in the top-left part of the figure to 7.5 in the middle of the figure. This difference across occupations is similar to the large aggregate differences observed between the US and

[6]Online Appendix B shows that the main pattern remains unchanged if we were to use other measures of dispersion in hours such as the coefficient of variation in annual hours or the 95/5 ratio and the 90/10 ratio of annual hours in an occupation.

several countries in Western Europe. Differences in the standard deviation in log hours are also large, ranging from around 0.25 in the top left to around 0.45 in the middle of the figure.

Figure 1 pooled the data across all individuals and years. Importantly, this pattern holds when we disaggregate along many dimensions of interest. Figure 2 presents results when we split the sample by gender (men, women), education (non-college, college), and age (22-35, 36-49, and 50-64).[7]

Figure 2: The Log of Mean Annual Hours vs. the Standard Deviation of Log Annual Hours, CPS, 1976-2015, 3-Digit Occupations: by Gender, Education, and Age.

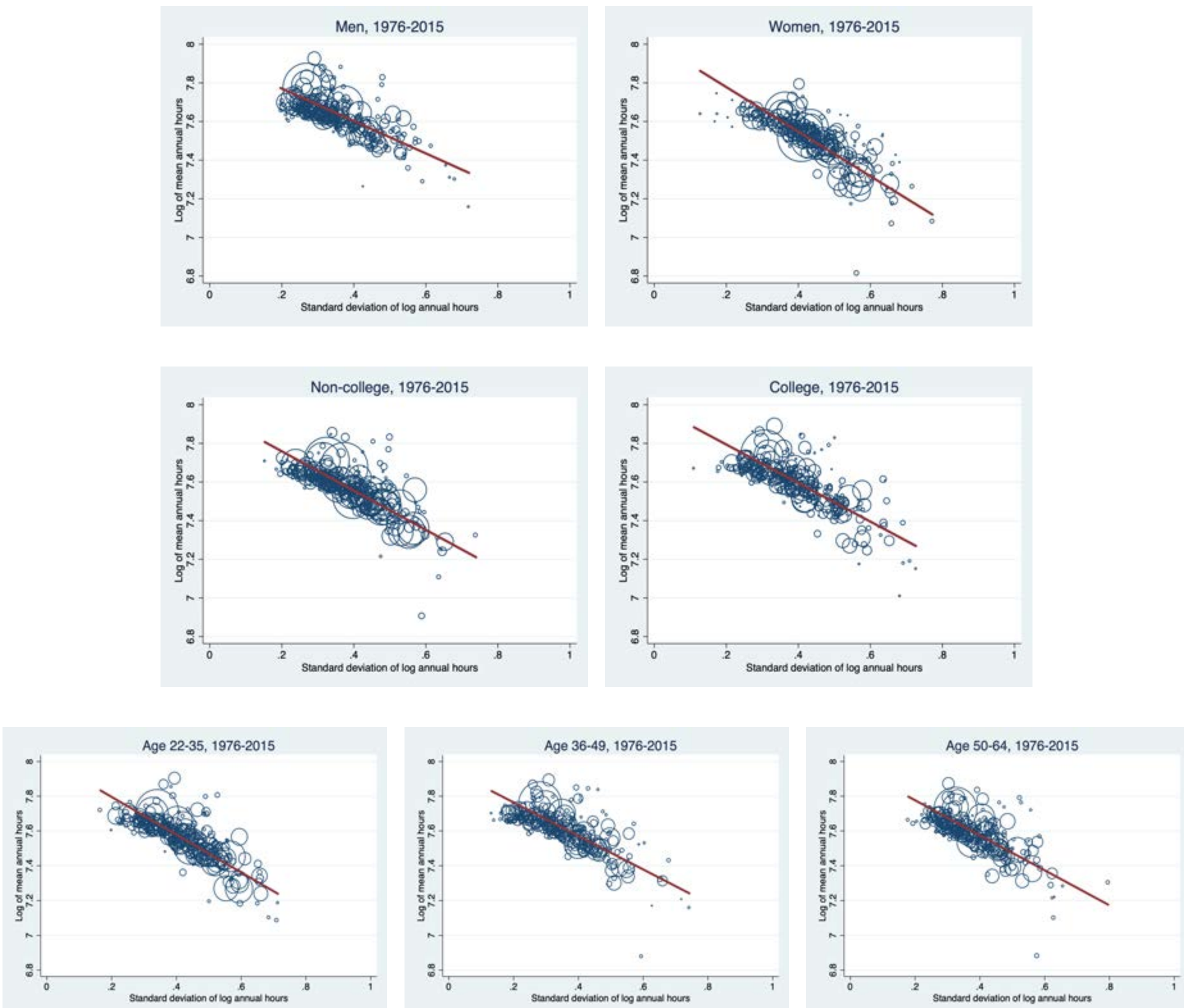

Notes: Each point represents a 3-digit occupation in the 1976-2015 time period. The straight solid red line represents a linear regression, weighted by the relative size of each occupation, while the size of the circle indicates the relative size of the occupation.

Figure 3 shows that the same pattern appears if we consider each successive ten year period between 1976 and 2015. Furthermore, Online Appendix E shows that the position

[7]Online Appendix C shows that a similar pattern emerges if we separately consider usual hours per week and weeks worked.

of an occupation in the mean-dispersion space is a relatively stable characteristic of an occupation over time.

Figure 3: The Log of Mean Annual Hours vs. the Standard Deviation of Log Annual Hours, CPS, 1976-2015, 3-Digit Occupations: by Different Time Periods.

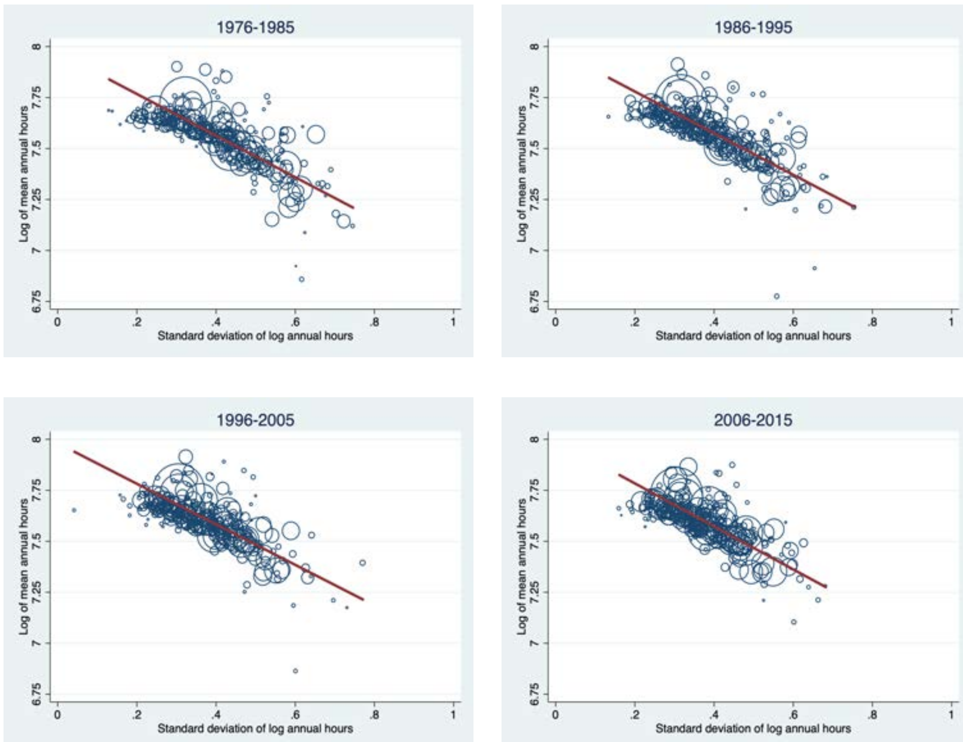

Notes: Each point represents a 3-digit occupation in the given 10-year time period. The scatter plot describes the relationship between the log of mean annual hours worked and the variance of log annual hours in a given occupation. The size of the circle indicates the relative size of the occupation.

One might conjecture that the pattern documented above is entirely driven by different propensities for part-time work across occupations; moving individuals from part-time to full-time would tend to mechanically increase mean hours and decrease the dispersion in hours. While this is part of what is going on in the data it is not everything. Online Appendix D shows that as occupation mean hours worked increases, there is both a large decline in the fraction working "short" hours (less than 1500) and a large increase in the fraction working "long" hours (more than 2500).

2.3 Hours and Wages

Following Roy (1951), models of occupational choice often focus on differences in the level and dispersion of wages across occupations. It is of interest to document the relationship between these variables and hours worked across occupations. Figure 4 provides scatter plots for each of the log of mean wages and the standard deviation of log wages against the log of mean hours, along with a linear regression line. Two patterns emerge. First, hourly wages tend to be higher in occupations with higher mean hours. Second, while there is a positive relationship between wage dispersion and mean hours, the variation in wage dispersion is much smaller than the variation in hours dispersion.

Figure 4: Occupational Hours and the Mean and Dispersion in Occupational Hourly Wages, CPS 1976-2015: by 3-Digit Occupations.

Notes: Each point represents a 3-digit occupation in the 1976-2015 time period. The scatter plot in the left panel describes the relationship between log mean annual hours and log mean hourly wages. The scatter plot in the right panel describes the relationship between log mean annual hours and the standard deviation of log hourly wages. The size of the circle indicates the relative size of the occupation while the solid red lines are the fitted weighted regression lines.

2.4 Aggregating the Data to Two Occupations

Our theoretical and quantitative analysis focuses on a two occupation model to maximize transparency and best highlight the key forces at work. For this reason it is of interest to aggregate the data to two occupation groupings. To do this we rank occupations by the level of mean hours, and separate them into two groups that are equal in size based on person-level weights. Having sorted individual observations into two occupation bins, we compute the log of mean hours, the standard deviation of log hours, the log of mean wages, and the standard deviation of log wages in each of the two occupations as a weighted average of values for occupations within each group.[8]

The moments of interest for the two aggregated occupations are reported in the first column of Table 1. We denote the high and low mean hours occupations by H and L respectively. By construction, each occupation accounts for 50 percent of total employment. The patterns reflected in this table mirror those found at the 3-digit level: there are large differences in mean hours, the standard deviation of log hours and mean wages; mean hours are negatively related to the standard deviation of log hours and positively related to wages; and there is only a small difference in the standard deviation of wages across the two occupations.

[8]Online Appendix A provides more detail.

3 A Model of Occupational Choice and Time Allocation

We study a simplified version of the model in Erosa et al. (2022). We highlight the presence of both occupation and selection effects for the distribution of hours worked across occupations and show that the model contains channels that can generate the pattern in Figure 1.

3.1 Model

A continuum of individuals, indexed by i, have preferences over consumption (c_i) and leisure ($T - h_i$) given by:

$$\ln c_i + \phi_i \frac{(T - h_i)^{1-\gamma}}{1 - \gamma}, \tag{1}$$

where T is the endowment of discretionary time, h_i is hours of work for individual i, ϕ_i is an individual specific preference parameter and $\gamma > 0$. As is standard in the macro literature, we assume offsetting income and substitution effects. Another important property of this utility function is that the marginal disutility of work approaches infinity at finite hours, i.e., as h tends to T. This serves to keep individuals away from the corner solution of zero leisure as ϕ becomes very small, and seems a reasonable property to impose on preferences in a model that features preference heterogeneity. More specifically, the elasticity of the marginal disutility of work with respect to hours of work is increasing in hours worked.[9]

There are two occupations, denoted by $j \in \{H, L\}$, and each individual is endowed with a pair of occupational specific productivities (a_{iH}, a_{iL}). The mapping from hours worked to efficiency units in occupation j for individual i depends both on the idiosyncratic productivity of the individual and an occupation specific non-linearity:

$$e_{ij} = a_{ij} h_{ij}^{1+\theta_j}. \tag{2}$$

If $\theta_j = 0$ for $j \in \{H, L\}$ the supply of efficiency units is linear in hours worked, and our model is the Roy model extended to include heterogeneity in leisure and an endogenous hours choice. We allow for the θ_j to be positive, and most importantly, to differ across occupations. In what follows we assume that $\theta_H > \theta_L$. Without loss of generality, the wage rate per efficiency unit is normalized to unity for both occupations.

In the quantitative exercise later in the paper we will interpret our static model to reflect outcomes for an individual averaged over their working life conditional on an occupational choice.[10] With this interpretation, the non-linear effect of hours on efficiency units reflects both static and dynamic effects. A part-time wage penalty would be an example of a static effect, while learning by doing that depends on hours would be an example of a dynamic effect.

[9]The elasticity of utility with respect to leisure is independent of the level of leisure, but as hours of work increase, a given percentage increase in hours of work imply a larger percentage decrease in hours of leisure.

[10]With this interpretation we are implicitly abstracting from occupational changes. Occupational mobility necessarily decreases as one considers a coarser partition of occupations. Aggregating 3-digit occupations into two groups serves to minimize the role of switching.

3.2 Properties

We solve the individual labor supply problem in two stages. The first stage solves for optimal hours conditional on occupational choice, denoted by h_{ij}, and the second stage solves for the optimal occupational choice.

The FOC for an interior solution for the first stage choice of h_{ij} is[11]:

$$\frac{1+\theta_j}{\phi_i} = h_{ij}(T-h_{ij})^{-\gamma} \equiv g(h_{ij}) \tag{3}$$

for $j \in \{H, L\}$. The function g defined in equation (3) is strictly increasing and convex (i.e., $g', g'' > 0$).

Three properties follow from this first order condition. First, our assumption of offsetting income and substitution effects implies that h_{ij} is independent of occupational productivity a_{ij}. Second, and intuitively, since higher values of ϕ_i indicate a higher preference for leisure, each h_{ij} is decreasing in ϕ_i. Third, for a given individual i, $\theta_H > \theta_L$ implies that $h_{iH} > h_{iL}$. This third property says that a given individual works more if they choose occupation H rather than occupation L and reflects an occupation effect on hours of work.

Next we consider the second stage choice of occupation. Individual i chooses to work in occupation H if the following inequality holds

$$\ln\left(a_{iH}h_{iH}^{1+\theta_H}\right) + \phi_i\frac{(T-h_{iH})^{1-\gamma}}{1-\gamma} > \ln\left(a_{iL}h_{iL}^{1+\theta_L}\right) + \phi_i\frac{(T-h_{iL})^{1-\gamma}}{1-\gamma}, \tag{4}$$

where h_{iH} and h_{iL} are the solutions to equation (3). Recalling that h_{ij} depends only on ϕ_i and not on a_{ij}, this expression can be re-arranged as:

$$\ln\left(\frac{a_{iH}}{a_{iL}}\right) > z(\phi_i) \equiv -(1+\theta_H)\ln(h_{iH}) + (1+\theta_L)\ln(h_{iL}) + \phi_i\left[\frac{(T-h_{iL})^{1-\gamma}}{1-\gamma} - \frac{(T-h_{iH})^{1-\gamma}}{1-\gamma}\right], \tag{5}$$

where $z'(\phi_i) > 0$ follows directly from the Envelope Theorem and $h_{iH} > h_{iL}$. Thus, individual i chooses to work in occupation H if the log of their skill ratio $\frac{a_{iH}}{a_{iL}}$ is higher than $z(\phi_i)$.

An important implication is that occupational choice is determined both by comparative advantage and tastes for leisure; holding a_{iH} and a_{iL} constant an individual is less likely to choose occupation H as ϕ_i increases. This result requires $\theta_H > \theta_L$; in a standard Roy model extended to include a leisure choice (i.e., our model with $\theta_H = \theta_L = 0$), occupational choice is solely determined by comparative advantage even if tastes for leisure are heterogeneous.[12]

Next we establish that holding the distribution of ϕ constant across occupations, the dispersion in log hours is less in occupation H than in occupation L. To show this, differentiate equation (3) to get:

$$-\frac{d\phi_i}{\phi_i}\frac{1+\theta_j}{\phi_i} = dh_{ij}\ (T-h_{ij})^{-\gamma} + dh_{ij}\ h_{ij}\gamma(T-h_{ij})^{-\gamma-1}$$

which gives a formula for the elasticity of optimal h_{ij} to ϕ_i, which we denote by $\varepsilon_{h_{ij},\phi_i}$:

$$\varepsilon_{h_{ij},\phi_i} = \frac{dh_{ij}}{h_{ij}}\Big/\frac{d\phi_i}{\phi_i} = -\frac{1}{1+\gamma\frac{h_{ij}}{T-h_{ij}}}. \tag{6}$$

[11] Our utility function implies that the solution for leisure will always be interior.

[12] This remains true if earnings are non-linear but $\theta_H = \theta_L$.

While the θ_j do not directly appear in expression (6), they influence the elasticity through their effect on the optimal choice of h_{ij}. It follows that holding the distribution of ϕ constant across occupations, $\theta_H > \theta_L$ implies higher mean hours and lower dispersion of log hours in occupation H. Another implication of equation (6) is that a proportional decrease in ϕ for all individuals in a given occupation (i.e., a leftward shift of the density for $\log \phi$) will lead to less dispersion in log hours.

As noted earlier, the distribution of ϕ across occupations will be determined by how individuals select across occupations. We previously showed that holding all else constant, individuals with lower values of ϕ are more likely to choose occupation H. This will tend to increase mean hours in occupation H and reduce the dispersion of log hours. Selection may also affect the relative magnitude of dispersion of $\log \phi$ across the two occupations. This effect will depend on many details of the distribution of heterogeneity and we cannot say anything at a general level about it. In our quantitative work we find that selection serves to generate less dispersion in $\log \phi$ in occupation H than in occupation L, but that this channel plays a small role relative to differences in the mean value of $\log \phi$.

3.3 Summary

We highlight three key properties of the model. First, the model generates an interaction between desired labor supply and occupational choice. Second, the heterogeneity in θ_j across occupations generates differences in mean hours worked across occupations via two effects: an occupation effect and a selection effect. Third, the model has forces that can generate a negative correlation between the mean and standard deviation of log hours across occupations.

4 Quantitative Assessment

This section assesses the extent to which the forces captured in our model can quantitatively account for the empirical patterns documented earlier, specifically those displayed in the first column of Table 1.[13]

4.1 Quantitative Exercise

To motivate our quantitative exercise it is useful to first recall the well-known result from Heckman and Honoré (1990). They show that the standard Roy model cannot be rejected non-parametrically using cross-section data on wages. This result is trivially extended to the case in which we include data on hours. If we do not impose parametric assumptions on the distribution of individual heterogeneity, our model can perfectly match any given cross-sectional distribution of hours, wages and occupational choice for any values of γ and the θ_j. To see why, note that given values for γ and the θ_j, and observations on hours,

[13] We noted earlier that we will interpret our model as reflecting annual outcomes averaged over a worker's lifetime. The cross-sectional statistics in the first column of Table 1 need not reflect cross-sectional differences in average lifetime outcomes across workers if individuals experience significant transitory shocks to hours and wages over their lifetime. We address this issue in our sensitivity analysis.

wages and occupational choice for a given individual, one can always find a value of ϕ_i and a value for a_{ij} in their observed occupation to match the observations for hours and wages. One can then choose a value for productivity in the other occupation (e.g., set it equal to zero) that makes the observed occupation choice optimal. An important implication is that one can impose $\theta_H = \theta_L$ and still match any pattern of hours and wages within and across occupations.

This non-parametric procedure places no restrictions on the distribution of individual heterogeneity, and in particular places no restrictions on the correlation between tastes for leisure and occupational specific productivities. Conversely, if one imposes that tastes for leisure are not correlated with occupational productivities, then with $\theta_H = \theta_L$ the hours distribution will be the same across occupations and the model cannot match the differences in hours distributions found in the data.

In what follows we adopt a parametric approach in which we make the relatively standard assumption that individual heterogeneity in (a_{iH}, a_{iL}, ϕ_i) is joint lognormally distributed. In order to highlight the potential of heterogeneity in the θ_j to account for the patterns in the hours data, our exercise imposes a zero correlation between ϕ_i and the a_{ij}. Given this restriction, our model has 11 parameters: the two values of the θ_j, the value of γ, the value of T, and 7 distributional parameters (two for the distribution of ϕ_i (mean and standard deviation), and five for the joint distribution of the a_{ij} (two means, two standard deviations and the correlation).)

Interpreting our model as predicting annual outcomes averaged over the life cycle we set $T = 5200$. We set the value of $\gamma = 4$.[14] Definitive estimates of the θ_j do not exist; the summary of evidence in Erosa et al. (2022) suggests that a reasonable range for θ_L is $[0, 0.4]$ and that the gap between θ_H and θ_L is probably no larger than 0.5. In view of this, and because many of our findings are not contingent on specific values for the θ_j, we present results for a range of combination of values for the θ_j. In particular, in what follows we will fix the value for θ_L and consider a range of values for θ_H.

For each combination of values for the two θ_j, we find the values of the 7 distributional parameters that best fit the 9 moments in the first column of Table 1 (the employment share of occupation H, the log of mean hours in each occupation, the standard deviation of log hours in each occupation, the log of mean wages in each occupation and the standard deviation of log wages in each occupation) by minimizing a loss function that is the sum of squared residuals. Because the number of targets exceeds the number of parameters being chosen, there is no presumption that the model can provide a good fit.

4.2 Results

Our main finding is robust to varying θ_L between 0 and 0.4, so in the interests of space we only report results here for a single value of θ_L, which following Erosa et al. (2022) we set to 0.2.[15] Table 1 shows the fit for each of the moments.[16] For each moment we also report the

[14]If on average leisure is 60% of total discretionary time and market work is the remaining 40%, the corresponding intertemporal elasticity of substitution for work along the intensive margin evaluated at these averages is roughly .4, in line with standard values assumed in the literature. See, e.g., Chetty (2012).

[15]Results for other values of θ_L are provided in Online Appendix G.

[16]Values for the underlying parameter estimates are contained in Online Appendix G.

difference between the two occupations in order to highlight the model's ability to account for the differences across occupations.

Table 1: Benchmark Results.

	Data	$\theta_H = .2$	$\theta_H = .3$	$\theta_H = .4$	$\theta_H = .5$	$\theta_H = .6$
Log Mean Hours						
Occ H	7.68	7.58	7.68	7.69	7.69	7.70
Occ L	7.49	7.58	7.49	7.48	7.48	7.47
Occ H−Occ L	0.19	0.00	0.20	0.20	0.21	0.23
Std Log Hours						
Occ H	0.33	0.39	0.33	0.33	0.33	0.32
Occ L	0.46	0.39	0.46	0.46	0.46	0.46
Occ H−Occ L	−0.14	0.00	−0.13	−0.13	−0.13	−0.14
Log Mean Wages						
Occ H	2.54	2.54	2.54	2.54	2.54	2.54
Occ L	2.02	2.02	2.02	2.02	2.02	2.03
Occ H−Occ L	0.53	0.53	0.53	0.52	0.52	0.51
Std Log Wages						
Occ H	0.56	0.56	0.56	0.57	0.59	0.60
Occ L	0.54	0.54	0.54	0.53	0.52	0.52
Occ H−Occ L	0.02	0.02	0.02	0.04	0.07	0.08
Emp Occ H	0.50	0.50	0.50	0.50	0.50	0.50
Loss Function$\times$ ($\mathbf{10^{-5}}$)	–	2696	8.2	46.0	144.7	293.3

We start with the results in the second column, corresponding to $\theta_H = \theta_L$. By assumption, this specification cannot generate differences in the hours distribution across occupations. The table shows that it can match the employment shares and the differences in wage distributions across occupations, which is essentially a confirmation in our context that the Roy model is capable of accounting for the first and second moment properties of occupational wage distributions even if we restrict attention to log normal distributions.[17] Consistent with the failure to match the properties of the hours distributions, the loss function attains a very high value in this case.

The key message from the other columns in Table 1 is that allowing for even a relatively small difference between the θ_j allows the model to capture the properties of both the wage and hours distributions. While the loss function is minimized at $\theta_H = 0.3$, the broader

[17]We note that because the two θ_j are not equal to zero, our model is not truly identical to the standard Roy model.

message is that the loss function is relatively flat over the entire range of values of θ_H that are greater than or equal to 0.3 and decreases by more than an order of magnitude when moving from $\theta_H = 0.2$ to any of the other values in $[0.3, 0.5]$.[18]

4.3 Occupation vs Selection Effects

Section 3 highlighted two distinct channels through which our model generates differences in hours across occupations: an occupation effect and a selection effect. Our estimation procedure targeted the occupational gap in mean hours and the results in Table 1 show that all of the economies with $\theta_H \in [0.3, 0.6]$ do a good job of accounting for this gap. In this subsection we assess the relative magnitude of the occupation and selection effects in our model. We also present evidence that speaks to the magnitude of the occupation effect in the data.

The empirical evidence that we present on the magnitude of the occupation effect uses data from the Survey of Income and Program Participation (SIPP) to document how hours of work change for individuals who move across our two aggregated occupation bins.[19] We find that average annual hours for individuals that move from Occupation H to Occupation L decline by 7.4 percent, while they increase by 6.0 percent for individuals that move from Occupation L to Occupation H.

Equation (3) implies that the difference between h_{iH} and h_{iL} depends on ϕ_i. To compare the magnitudes from the SIPP with our model we make the assumption that switchers are most likely those individuals who are close to the boundary in terms of their occupational choice decision.[20] We then compute the average change in hours that would result if these individuals were forced to switch occupations.[21]

We define close to the boundary to be the 5 percent of individuals in each occupation that have the smallest utility differences associated with their occupational choice.[22] When $\theta_H = 0.3$ the model implied change in mean hours is -2.4% for individuals moving from Occupation H to Occupation L, and $+2.5\%$ for individuals moving from Occupation L to Occupation H. When $\theta_H = 0.5$ the values are -6.7% and $+7.0\%$. Given that our loss function was relatively flat over the interval of θ_H values in $[0.3, 0.5]$, we interpret this evidence to imply that our model can match the key features of the wage and hours distributions with an empirically reasonable magnitude for the occupation effect.[23]

Recalling that the gap in log mean hours across occupations is 0.19, the preceding calculation implies that selection effects are the dominant source of differences in hours across occupations in our model. To further illustrate this, for each of the estimated economies in Table 1 we eliminate the heterogeneity in tastes for leisure (i.e., we set $\sigma_\phi = 0$) and solve

[18] When we considered a finer grid (distance 0.01 between gridpoints) we found that the minimum still occurred at $\theta_H = 0.30$.

[19] Details are provided in Online Appendix A.

[20] More formally, this would be true if we thought that occupation switching was the result of idiosycratic shocks to individual productivies or to occupation specific productivity shocks.

[21] As a practical matter, it turns out that the change in hours for switchers does not vary that much across the population, so that restricting attention to the group close to the boundary is not of particular significance.

[22] Changing this threshold has little effect on the magnitudes we report.

[23] One could view this moment as a useful way to provide additional discipline on the value of θ_H.

for optimal decisions holding all other parameters fixed. When we do this for the $\theta_H = 0.3$ economy, the gap in log mean hours across occupations shrinks from 0.20 to 0.02.[24] While the selection effect is the dominant effect, it is important to remember that the selection effect is only operative if there is an occupation effect; if $\theta_H = \theta_L$ then occupational choice is independent of tastes for work.[25]

Eliminating heterogeneity in preferences also allows us to gauge the importance of selection on preferences for occupational choice. When we do this for the $\theta_H = 0.3$ economy, the net change in the size of the H occupation is roughly 2 percent, but about 8% and 13% of workers in occupations H and L would choose a different occupation.

Next we consider the forces that shape the differences in the dispersion of hours across occupations. In our theoretical analysis we argued that an increase in mean hours due to either a higher value of θ or to a leftward shift in the distribution of $\log \phi$ would lead to lower dispersion in log hours. But we also noted that selection effects could give rise to differences in the dispersion in $\log \phi$ across occupations. Figure G-1 in Online Appendix G shows the densities for the distribution of $\log \phi$ in each of the two occupations as well as for the aggregate.[26] Both occupation specific distributions closely resemble log normal distributions, with the distribution for occupation H featuring both a lower mean and a lower standard deviation. Here we carry out a simple exercise to show that the dominant force behind the differences in the standard deviation in log hours reported in Table 1 is the difference in mean hours.

To do this, for each economy we hold fixed the set of individuals in occupation H and then scale the distribution of ϕ proportionately so as to make mean hours in occupation H the same as in occupation L. Importantly, this leaves the standard deviation of $\log \phi$ unchanged. We then compute the increase in the standard deviation of log hours in occupation H as a result of this scaling. This increase accounts for roughly 70 percent of the log difference between the standard deviation of hours across occupations when $\theta_H = 0.3$, and 73% when $\theta_H = 0.5$. We conclude that the mean hours channel highlighted by our theoretical analysis is the dominant factor behind the difference in the dispersion of hours across occupations.

4.4 Wage and Earnings Inequality Across Occupations

In the data, occupational gaps in earnings are substantially larger than occupational gaps in wages. For example, the gap in log median earnings is more than one third larger than the gap in log median wages (0.66 versus 0.49). A standard Roy model without a labor supply decision can match differences in wages or differences in earnings but cannot do both simultaneously. Whereas wage inequality may be of primary interest in some contexts, earnings inequality may be more relevant in others. For example, wage differences may be

[24]For the $\theta_H = 0.5$ economy the gap in log mean hours across occupations decreases from 0.21 to 0.04.

[25]The presence of an occupation effect is also relevant for comparing our model against alternatives. In particular, we previously pointed out that one could set $\theta_H = \theta_L$ and match the data on hours, wages and occupation if the distribution of productivity and preference parameters are estimated non-parametrically. But this specification implies that hours of work for a given individual are independent of occupational choice, and switching occupations in response to changes in productivity or skill prices would have no implications for changes in hours.

[26]Online Appendix G also shows the densities for comparative advantage.

more relevant when studying inequality of opportunity, but earnings inequality may be more relevant when studying wealth inequality or consumption. More generally, understanding the role of labor supply is likely to be relevant for thinking about the welfare implications of inequality.

Our model makes predictions for both wages and earnings, and does a good job of capturing the gaps in both wage and earnings distributions across occupations. In particular, when $\theta_H = 0.3$ the occupational gaps in median log earnings and median log wages are 0.72 and 0.51. Our model also does a good job of accounting for the levels of both wage and earnings inequality in the overall economy.

Because the dominant source of the occupational gap in hours is selection on tastes for leisure, it follows that heterogeneity in tastes for leisure is an important element for a theory of earnings differences across occupations. Put differently, if we eliminate heterogeneity in tastes for leisure in our calibrated models, our model predicts that occupational gaps in wages and earnings differ by very little relative to the data.

4.5 Sensitivity

We previously mentioned that our findings are robust to alternative choices for the value of θ_L. Here we briefly mention two additional sensitivity exercises that are summarized in the online appendix. The first exercise assumes that the relationship between earnings and hours becomes linear beyond some level of hours $\bar{h}$. In particular, Online Appendix G contains results for the case in which $\bar{h} = 2100$. The second exercise adjusts the two dispersion moments in Table 1 to control for the fact that some cross-sectional dispersion reflects transitory fluctuations that average out over time for a given individual. In both cases our main messages are unaffected.

5 Conclusion

We document a novel fact about hours of work across three digit occupations in the US: a robust negative correlation between mean individual hours of work and the standard deviation of log individual hours. We show that an extended version of the standard Roy model offers a parsimonious explanation for differences in both hours and wage distributions across occupations. The success of the model relies on the interaction of two features: heterogeneity in preferences across individuals and heterogeneity in the convexity of the hours-earnings profile across occupations. In a standard Roy model, occupational choice is driven solely by selection on comparative advantage. Our analysis also attributes an important role to a novel channel that reflects selection based on preferences. The connection between hours and occupational choice that we document could also prove important for understanding labor supply and earnings inequality in many contexts that feature changes or differences in the occupational distribution of employment.

References

AUTOR, D. AND D. DORN, "The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market," *American Economic Review* 103 (2013), 1553–1597.

BARZEL, Y., "The Determination of Daily Hours and Wages," *Quarterly Journal of Economics* 87 (1973), 220–238.

EROSA, A., L. FUSTER, G. KAMBOUROV AND R. ROGERSON, "Hours, Occupations, and Gender Differences in Labor Market Outcomes," Working paper 23636, NBER, 2017.

———, "Hours, Occupations, and Gender Differences in Labor Market Outcomes," *American Economic Journal: Macroeconomics* 14 (July 2022), 543–590.

FLOOD, S., M. KING, R. RODGERS, S. RUGGLES AND R. J. WARREN, *Integrated Public Use Microdata Series, Current Population Survey: Version 6.0 [dataset]* (Minneapolis, MN: IPUMS, 2018).

FRENCH, E., "The Effects of Health, Wealth, and Wages on Labor Supply and Retirement Behavior," *Review of Economic Studies* 72 (2005), 395–427.

GOLDIN, C., "A Grand Gender Convergence: Its Last Chapter," *American Economic Review* 104 (2014), 1091–1119.

HAUSMAN, J., "The Econometrics of Nonlinear Budget Sets," *Econometrica* 53 (1985), 1255–1282.

HECKMAN, J. J. AND B. E. HONORÉ, "The Empirical Content of the Roy Model," *Econometrica* 58 (1990), 1121–1149.

HORNSTEIN, A. AND E. C. PRESCOTT, "The Firm and the Plant in General Equilibrium Theory," in R. Becker, M. Boldrin, R. Jones and W. Thomson, eds., *General Equilibrium, Growth, and Trade II: The Legacy of Lionel McKenzie* (San Diego: Academic Press, 1993).

MOFFITT, R., "The Estimation of a Joint Wage-Hours Labor Supply Model," *Journal of Labor Economics* 2 (1984), 550–566.

ROSEN, H., "Taxes in a Labor Supply Model with Joint Wage-Hours Determination," *Econometrica* 44 (1976), 485–507.

ROSEN, S., "The Supply of Work Schedules and Employment," in *Work Time and Employment* (Washington, DC: National Commission for Manpower Policy, 1978).

ROY, A. D., "Some Thoughts on the Distribution of Earnings," *Oxford Economic Papers* 3 (1951), 135–146.

ONLINE APPENDICES

A Data Description

A.1 Current Population Survey, 1976-2015

Our analysis is based on the IPUMS-CPS files from the 1976-2015 Current Population Survey (CPS).[27]

Annual hours and hourly wages. Total annual hours worked last year are constructed using the variables (i) weeks worked last year (WKSWORK1) and (ii) usual hours worked per week last year (UHRSWORK). Hourly wages last year are constructed using the variables (i) wage and salary income for the previous calendar year (INCWAGE) and (ii) total hours worked last year constructed above. Real hourly wages are obtained using the CPI index(=100 in 1982/84).

Consistent 1976-2015 occupational classification. The occupational classification has changed four times over the period 1976-2015.[28] We use the occupational classification provided in Autor and Dorn (2013) to construct consistent occupational codes for the 1976-2015 period.

Sample restrictions.

- *Age.* 22-64.
- *Time Period.* Our benchmark results use the pooled data from 1976-2015. In some of the analysis we also use four 10-year periods: (1) 1976-1985; (2) 1986-1995; (3) 1996-2005; (4) 2006-2015.
- *Annual Hours.* Drop observations with annual hours of less than 250 or more than 4500.
- *Real Hourly Wages.* Drop observations with a real hourly wage in the top and bottom 0.1% of the hourly wage distribution.
- *Number of Observations in an Occupation.* Use observations from occupations with at least 30 observations.
- *One Employer per Year.* Some individuals might have worked in multiple occupations during the survey year. To address this, for each survey year, we focus on individuals who report having had a single employer (NUMEMPS variable).

Aggregate moments. The aggregate moments are computed as follows. We compute (log) mean hours, (log) mean wages, the standard deviation of log hours, and the standard deviation of log wages in each occupation, using person-level weights in the analysis. Then, we report the averages of these moments across all occupations, using the relative share of individuals in each occupation.

Moments for the linear and nonlinear sectors. We compute mean hours in each occupation, using person-level weights, rank all occupations by the level of mean hours, and separate them into two groups that are equal in size, based on person-level weights. We then compute (log) mean hours, (log) mean wages, the standard deviation of log hours, and the standard deviation of log wages in each occupation, using person-level weights in the analysis. Finally, for each of the two sectors, we report the averages of these moments, using the relative share of individuals in each occupation in that sector.

[27]The data and a detailed description can be found at http://cps.ipums.org/cps/.

[28]Specifically, the 1970 census classification scheme is used 1971-1982, the 1980 census classification scheme is used for 1983-1991, the 1990 census classification scheme is used for 1992-2002, and the 2000 census classification scheme is used for 2003-2015.

A.2 Survey of Income and Program Participation

We use the 1990 Survey of Income and Program Participation (SIPP) dataset, that runs from 1989 until 1992, to compute changes in hours worked for occupational switchers (and stayers). The SIPP data is monthly; however, individuals are interviewed every four months when they provide information on each of the months. Based on the analysis in the paper, and as described in Online Appendix A.1, individuals belong to one of two occupational groups: Occupation H or occupation L. We identify occupational switchers, from occupation H to occupation L or vice versa, at the monthly level between months $t-1$ to t and look at the average change in hours worked between months $t+1$, $t+2$, and $t+3$ and months $t-2$, $t-3$, and $t-4$, controlling for the fact that the only occupational switch during this period is between months $t-1$ and t.

We find that individuals that stay in their current occupation do not experience on average any change in their hours worked. In particular, those that remain in occupation H and L experience a 0% and -0.2% change, respectively. Occupational switchers, however, on average experience significant changes in their hours worked. Workers that switch from occupation H to L see their hours worked decrease by 7.4% while those that switch from occupation L to H experience an increase of 6%.

B Sensitivity with Respect to Various Measures of Dispersion in Annual Hours across Occupations: 1976-2015

We provide sensitivity of the main pattern observed in the data − of a negative relationship between mean hours and the dispersion in hours in an occupation − with respect to various measures of dispersion in annual hours. In the main text we used the standard deviation in log hours as the preferred measure of dispersion. However, the main pattern remains unchanged if we were to use the coefficient of variation in annual hours (B-1) or the 95/5 and 90/10 ratio of annual hours in an occupation (B-2).

Figure B-1: The Log of Mean Annual Hours vs. the Coefficient of Variation of Annual Hours, CPS 1976-2015: by 3-Digit Occupations.

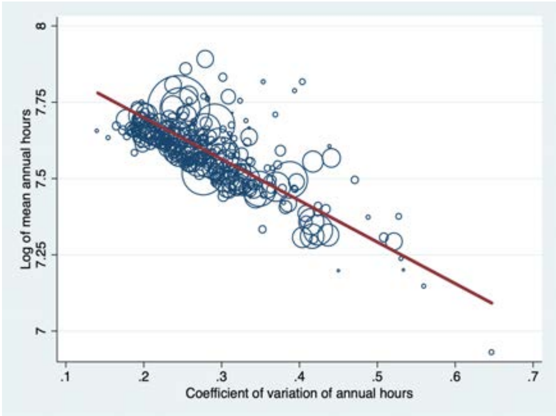

Notes: Each point represents a 3-digit occupation in the 1976-2015 time period. The size of the circle indicates the relative size of the occupation.

Figure B-2: The Log of Mean Annual Hours vs. 95/5 and 90/10 Ratio of Annual Hours, CPS 1976-2015: by 3-Digit Occupations.

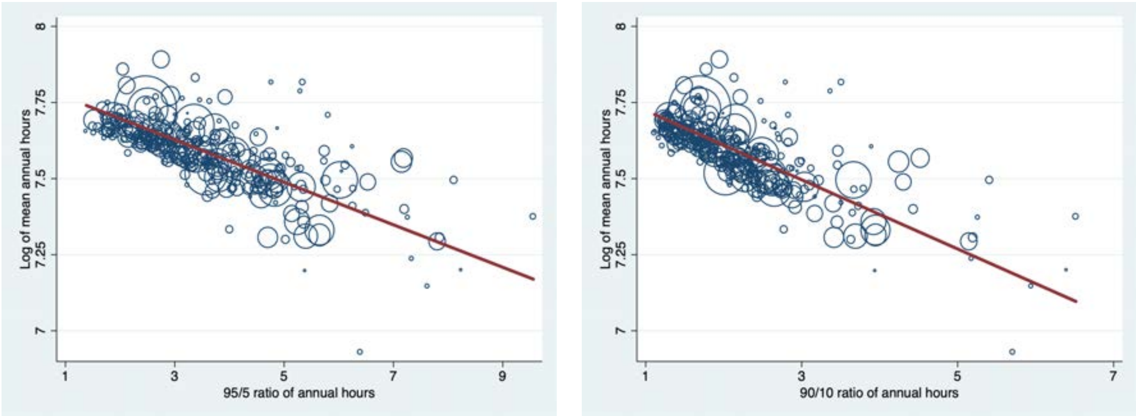

Notes: Each point represents a 3-digit occupation in the 1976-2015 time period. The size of the circle indicates the relative size of the occupation.

C Intensive and Extensive Margin: 1976-2015

Figure C-1 shows the relationship between log mean weeks and the standard deviation of log weeks in an occupation. The variable used here is number of week worked last year. Similarly to what we observed for total annual hours worked, there is a negative relationship between the mean number of weeks and the dispersion in weeks worked in a given occupation.

Figure C-1: The Log of Mean Weeks vs. the Standard Deviation of Log Weeks, CPS, 1976-2015, by 3-Digit Occupations.

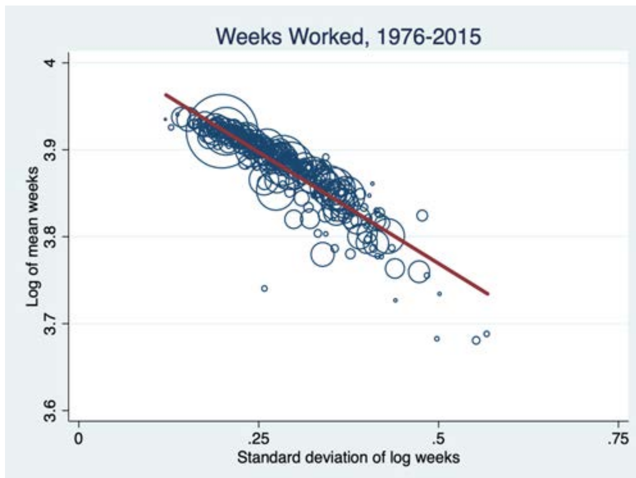

Notes: Each point represents a 3-digit occupation in the 1976-2015 time period. The size of the circle indicates the relative size of the occupation.

Figure C-2 shows the relationship between log mean weekly hours and the standard deviation of log weekly hours in an occupation. The variable used here is usual hours per week last year. Similarly to what we observed for total annual hours worked, there is a negative relationship between the mean hours per week and the dispersion in hours per week in a given occupation.

Finally, we analyze the correlation between weeks worked and weekly hours in occupational hours. Is it the case that occupations that have a high mean and low dispersion in number of weeks worked also exhibit a high mean and low dispersion in usual hours worked in a week? Figure C-3 shows that the correlation between the dispersion in log weekly hours and the dispersion in log weeks in an occupation is positive. This implies that, taking into account the facts described above, some occupations exhibit a high mean and low dispersion in their hours both in terms of usual weekly hours and weeks worked while other occupations exhibit a low mean and high dispersion in their hours both in terms of usual weekly hours and weeks worked.

Figure C-2: The Log of Mean Weekly Hours vs. the Standard Deviation of Log Weekly Hours (Usual Hours per Week Last Year), CPS, 1976-2015, by 3-Digit Occupations.

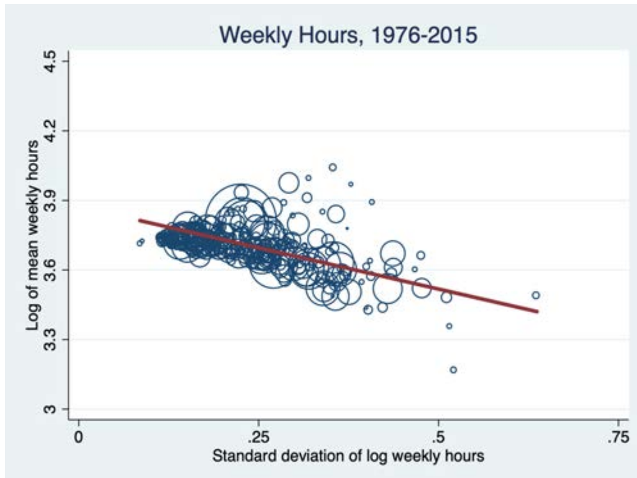

Notes: Each point represents a 3-digit occupation in the 1976-2015 time period. The size of the circle indicates the relative size of the occupation.

Figure C-3: The Standard Deviation of Log Weekly Hours vs. the Standard Deviation of Log Weeks, CPS, 1976-2015, by 3-Digit Occupations.

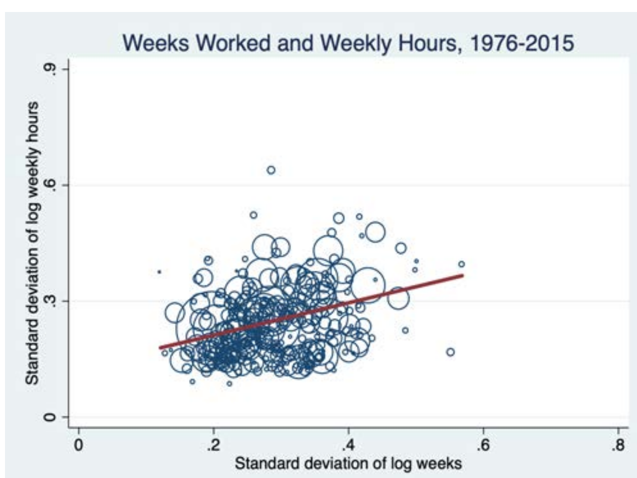

Notes: Each point represents a 3-digit occupation in the 1976-2015 time period. The size of the circle indicates the relative size of the occupation.

D The Distribution of Hours Worked across Occupations: 1976-2015

For each occupation over the 1976-2015 period, we compute the fraction of people working less than 1500 annual hours (30 hours per week) and more than 2500 annual hours (50 hours per week).

Figure D-1: Fraction Working "Short" and "Long" Hours, CPS, 1976-2015: by 3-Digit Occupations.

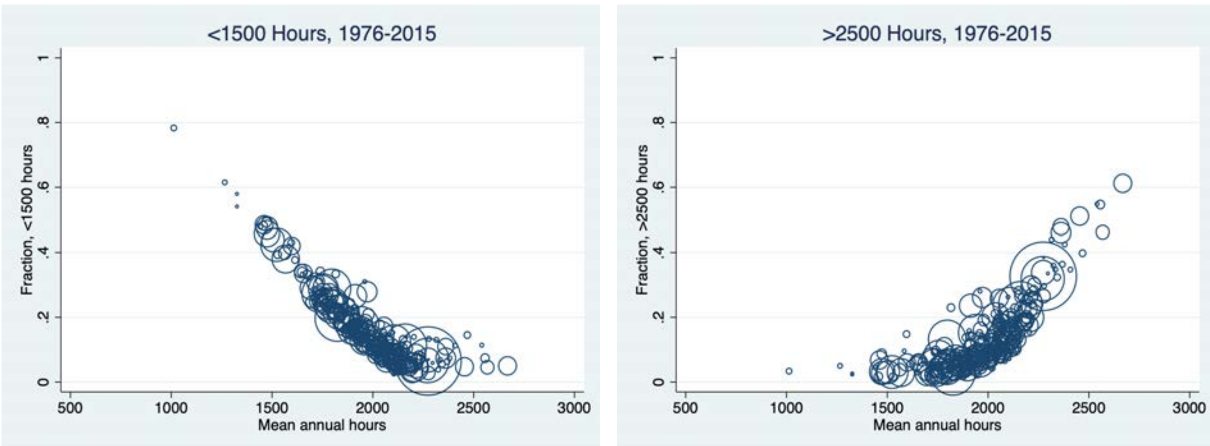

Notes: Each point represents a 3-digit occupation in the 1976-2015 time period. The size of the circle indicates the relative size of the occupation.

Figure D-1 reports the results. The horizontal axis displays the mean annual hours worked in a particular occupation. As we move along the x-axis from occupations with low mean hours towards the occupations with high mean hours, the fraction working "short" hours (less than 1500 annual hours) declines while the fraction working "long" hours (more than 2500 annual hours) increases. This indicates that as the level of mean hours worked in an occupation increases, the entire distribution of hours worked shifts to the right.

E Occupational Hours over Time: 1976-2015

Figure 3 in the paper illustrates that the pattern observed over the 1976-2015 time period is also observed in each of the 10-year periods from 1976 until 2015. Furthermore, there are no major changes over time in the mean and dispersion in hours in occupations, implying that the position of an occupation in the mean-dispersion space is a somewhat fixed characteristic of an occupation. Figure E-1 shows the log mean annual hours for each occupation in 1986-1995 and 2006-2015 relative to the initial 1976-1985 time period. Although there are some changes, mostly towards higher mean hours, 30 years later most occupations still line up closely along the 45-degree line. A similar pattern emerges for the changes in the standard deviation of log annual hours, as reported in Figure E-2. The plots for the changes in standard deviation over time exhibit more dispersion around the 45 degree line, but this is to be expected if the estimate of the standard deviation of hours within an occupation is noisier than the estimate of mean hours.

Figure E-1: Log of Mean Annual Hours, Over Time: by 3-Digit Occupations.

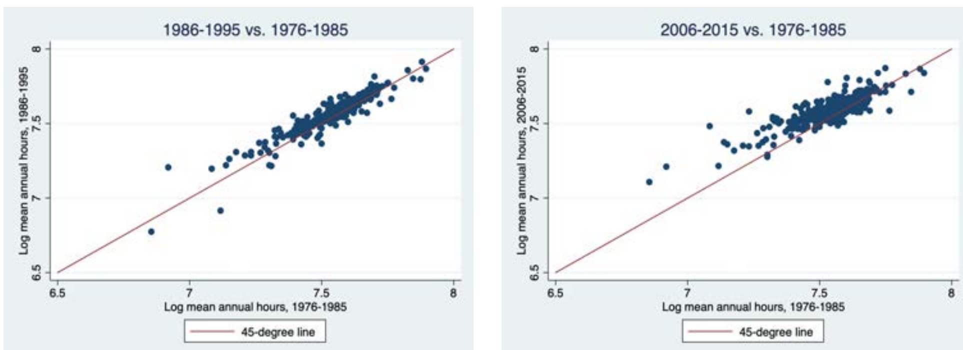

Notes: Each point represents a 3-digit occupation in a given 10-year time period. The scatter plot describes the change in log of mean annual hours in a given occupation over time, relative to the 1976-1985 period.

Figure E-2: Standard Deviation of Log Annual Hours, Over Time: by 3-Digit Occupations.

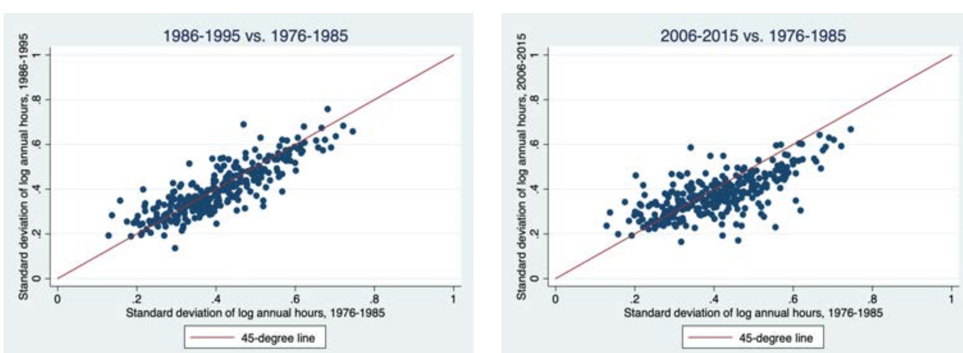

Notes: Each point represents a 3-digit occupation in a given 10-year time period. The scatter plot describes the change in the standard deviation of log annual hours in a given occupation over time, relative to the 1976-1985 period.

F Occupational Distribution of Workers: 1976-2015

We study the distribution of workers across occupations, based either on the occupational mean hours or on the dispersion in occupational hours. A useful approach is to construct the complementary cumulative distribution function in period t, $\overline{F}_{m,t}(x)$, defined as the probability that X − either log mean hours or the standard deviation of log hours − will take a value greater than x. Figure F-1 plots the two distributions.

Figure F-1: Complementary Cumulative Distributions, 1976-2015: by 3-Digit Occupations.

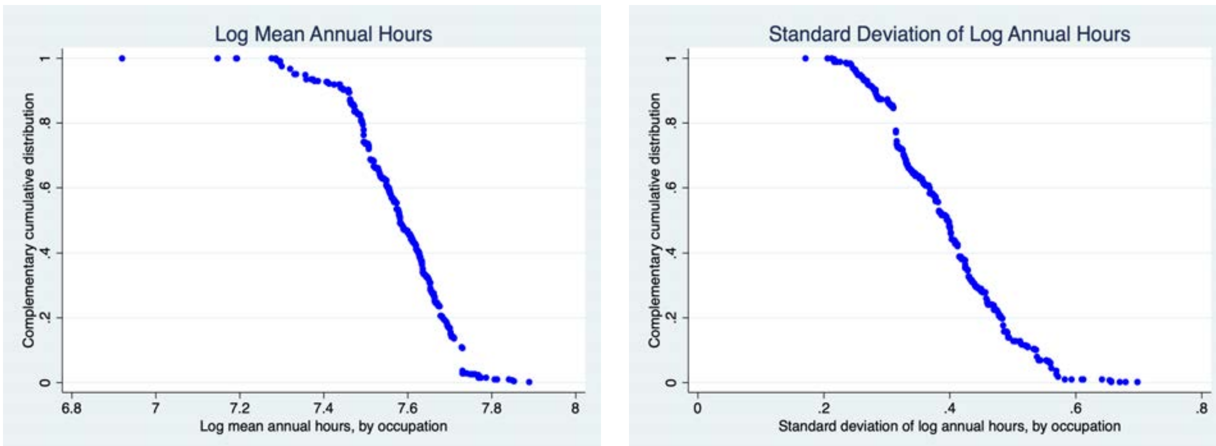

Notes: The scatter plots show the 1976-2015 complementary cumulative distributions over occupations in terms of the log of mean annual hours in an occupation (left panel) or the standard deviation of log annual hours in an occupation (right panel).

G Calibration Results

Figure G-1: Distributions of log ϕ.

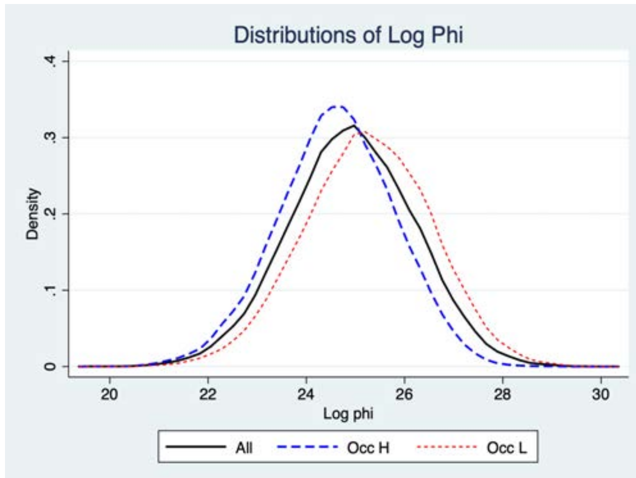

Notes: The figure plots the distributions of log ϕ for the whole economy, occupation H, and occupation L.

Figure G-2: Distributions of log $\left(\frac{a_H}{a_L}\right)$.

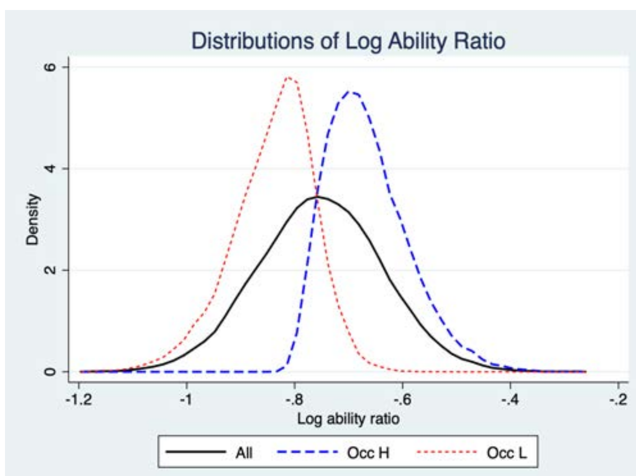

Notes: The figure plots the distributions of the log ability ratio, log $\left(\frac{a_H}{a_L}\right)$, for the whole economy, occupation H, and occupation L.

Table G-1: Calibration of Economies Differing on θ_H.

Parameters	$\theta_H = 0.2$	$\theta_H = 0.3$	$\theta_H = 0.4$	$\theta_H = 0.5$	$\theta_H = 0.6$
μ_{a_H}	0.5649	-0.1764	-0.9863	-1.8022	-2.6138
μ_{a_L}	0.5660	-0.1833	-0.9998	-1.8216	-2.6384
μ_ϕ	0.5014	0.5131	0.5488	0.5775	0.6095
$\sigma^2_{a_H}$	0.4120	0.3943	0.4267	0.4650	0.4936
$\sigma^2_{a_L}$	0.3191	0.3247	0.2904	0.2657	0.2581
σ^2_ϕ	1.4807	1.6334	1.6038	1.5881	1.5570
ρ_{a_H,a_L}	0.9785	0.9861	0.9324	0.8244	0.6807

Notes: The table shows the calibration results for economies that vary in θ_H. All economies featured in the table have $\theta_L = 0.2$. The calibration searches for 7 parameters that minimize the sum of the square deviations between 9 statistics in the model and in the data.

Table G-2: Calibration of Economies Differing on θ_H.

Targets	$\theta_H = 0.2$	$\theta_H = 0.3$	$\theta_H = 0.4$	$\theta_H = 0.5$	$\theta_H = 0.6$	Data
E_m^{NL}	0.50	0.50	0.50	0.50	0.50	0.50
$\ln \overline{h_H}$	7.58	7.68	7.69	7.69	7.70	7.68
$\ln \overline{h_L}$	7.58	7.49	7.48	7.48	7.47	7.49
$\ln \overline{w_H}$	2.54	2.54	2.54	2.54	2.54	2.54
$\ln \overline{w_L}$	2.02	2.02	2.02	2.02	2.03	2.02
$sd(\ln h_H)$	0.39	0.33	0.33	0.33	0.32	0.33
$sd(\ln h_L)$	0.39	0.46	0.46	0.46	0.46	0.46
$sd(\ln w_H)$	0.56	0.56	0.57	0.59	0.60	0.56
$sd(\ln w_L)$	0.54	0.54	0.53	0.52	0.51	0.54

Notes: All economies featured in the table have $\theta_L = 0.2$. The table shows the calibration results for economies that vary in θ_H. The calibration searches for 7 parameters that minimize the sum of the square deviations between 9 statistics in the model and in the data.

Table G-3: Calibration of Economies Differing on (θ_H, θ_L).

Parameters	$\theta_H = 0.3$ $\theta_L = 0.3$	$\theta_H = 0.4$ $\theta_L = 0.3$	$\theta_H = 0.5$ $\theta_L = 0.3$	$\theta_H = 0.4$ $\theta_L = 0.4$	$\theta_H = 0.5$ $\theta_L = 0.4$	$\theta_H = 0.6$ $\theta_L = 0.4$
μ_{a_H}	-0.1820	-0.9278	-1.7363	-0.9498	-1.6784	-2.4881
μ_{a_L}	-0.1810	-0.9347	-1.7499	-0.9491	-1.6854	-2.5017
μ_ϕ	0.5711	0.5903	0.6235	0.6443	0.6621	0.6927
$\sigma^2_{a_H}$	0.3995	0.3832	0.4118	0.4045	0.3685	0.3953
$\sigma^2_{a_L}$	0.3126	0.3164	0.2829	0.2905	0.3053	0.2703
σ^2_ϕ	1.4939	1.6317	1.6076	1.4988	1.6323	1.6130
ρ_{a_H,a_L}	0.9805	0.9853	0.9298	0.9659	0.9847	0.9236

Notes: The table shows the calibration results for economies that vary in θ_H. The first three economies feature $\theta_L = 0.3$ and the last three economies $\theta_L = 0.4$. The calibration searches for 7 parameters that minimize the sum of the square deviations between 9 statistics in the model and in the data.

Table G-4: Calibration of Economies Differing on (θ_H, θ_L).

Targets	$\theta_H = 0.3$ $\theta_L = 0.3$	$\theta_H = 0.4$ $\theta_L = 0.3$	$\theta_H = 0.5$ $\theta_L = 0.3$	$\theta_H = 0.4$ $\theta_L = 0.4$	$\theta_H = 0.5$ $\theta_L = 0.4$	$\theta_H = 0.6$ $\theta_L = 0.4$	Data
E_m^{NL}	0.50	0.50	0.50	0.50	0.50	0.50	0.50
$\ln \overline{h_H}$	7.58	7.68	7.68	7.58	7.68	7.68	7.68
$\ln \overline{h_L}$	7.58	7.49	7.48	7.59	7.49	7.49	7.49
$\ln \overline{w_H}$	2.54	2.54	2.54	2.55	2.54	2.54	2.54
$\ln \overline{w_L}$	2.02	2.02	2.02	2.02	2.02	2.02	2.02
$sd(\ln h_H)$	0.39	0.33	0.33	0.39	0.33	0.33	0.33
$sd(\ln h_L)$	0.40	0.46	0.46	0.40	0.46	0.46	0.46
$sd(\ln w_H)$	0.56	0.56	0.57	0.57	0.56	0.57	0.56
$sd(\ln w_L)$	0.54	0.54	0.53	0.53	0.54	0.53	0.54

Notes: The table shows the calibration results for economies that vary in θ_H. The first three economies feature $\theta_L = 0.3$ and the last three economies $\theta_L = 0.4$. The calibration searches for 7 parameters that minimize the sum of the square deviations between 9 statistics in the model and in the data.

Table G-5: Calibration of Economies with $\bar{h} < 5200$ and $(\theta_H, \theta_L) = (0.3, 0.2)$

Parameters	Baseline	$\bar{h} = 2500$	$\bar{h} = 2300$	$\bar{h} = 2100$
μ_{a_H}	-0.1764	-0.1656	-0.1576	-0.1426
μ_{a_L}	-0.1833	-0.1750	-0.1696	-0.1578
μ_ϕ	0.5131	0.4748	0.4582	0.4328
$\sigma^2_{a_H}$	0.3943	0.3943	0.3946	0.3929
$\sigma^2_{a_L}$	0.3247	0.3275	0.3296	0.3339
σ^2_ϕ	1.6334	1.7117	1.7391	1.7787
ρ_{a_H,a_L}	0.9861	0.9876	0.9884	0.9908

Notes: The table shows the calibration results for economies with $\bar{h} < 5200$. All economies in the table feature $(\theta_H, \theta_L) = (0.3, 0.2)$. The first economy corresponds to the baseline economy with $\bar{h} = 5200$. In the other economies $\bar{h} = \{2100, 2300, 2500\}$ in both occupations. The calibration searches for 7 parameters that minimize the sum of the square deviations between 9 statistics in the model and in the data.

Table G-6: Calibration of Economies with $\bar{h} < 5200$.

Parameters	Baseline	2500	2300	2100	Data
E_m^{NL}	0.50	0.50	0.50	0.50	0.50
$\ln \overline{h_H}$	7.68	7.67	7.66	7.66	7.68
$\ln \overline{h_L}$	7.49	7.50	7.50	7.50	7.49
$\ln \overline{w_H}$	2.54	2.54	2.54	2.54	2.54
$\ln \overline{w_L}$	2.02	2.02	2.02	2.02	2.02
$sd(\ln h_H)$	0.33	0.32	0.31	0.31	0.33
$sd(\ln h_L)$	0.46	0.47	0.47	0.48	0.46
$sd(\ln w_H)$	0.56	0.56	0.56	0.56	0.56
$sd(\ln w_L)$	0.54	0.54	0.54	0.54	0.54

Notes: The table shows the calibration results for economies with $\bar{h} < 5200$. All economies in the table feature $(\theta_H, \theta_L) = (0.3, 0.2)$. The first economy corresponds to the baseline economy with $\bar{h} = 5200$. In the other economies $\bar{h} = \{2100, 2300, 2500\}$ in both occupations. The calibration searches for 7 parameters that minimize the sum of the square deviations between 9 statistics in the model and in the data.

Table G-7: Calibration of Economies with Lower Dispersion in Hours.

Parameters	$\theta_H = 0.2$	$\theta_H = 0.3$	$\theta_H = 0.4$	$\theta_H = 0.5$	$\theta_H = 0.6$
μ_{a_H}	0.5560	-0.1647	-0.9535	-1.7486	-2.5477
μ_{a_L}	0.5571	-0.1687	-0.9613	-1.7593	-2.5617
μ_{ϕ}	0.4965	0.5172	0.5542	0.5858	0.6178
$\sigma^2_{a_H}$	0.4165	0.3842	0.4017	0.4220	0.4447
$\sigma^2_{a_L}$	0.3185	0.3452	0.3205	0.2981	0.2797
σ^2_{ϕ}	0.7390	0.8526	0.8270	0.8068	0.7896
ρ_{a_H,a_L}	0.9755	0.9964	0.9812	0.9468	0.8874

Notes: The table shows the calibration results for economies that vary in θ_H. The calibration targets 2/3 of the standard deviation of hours in the cross sectional data. All economies feature $\theta_L = 0.2$ and $\bar{h} = 5200$. The calibration searches for 7 parameters that minimize the sum of the square deviations between 9 statistics in the model and in the data.

Table G-8: Calibration of Economies with Lower Dispersion in Hours.

Targets	$\theta_H = 0.2$	$\theta_H = 0.3$	$\theta_H = 0.4$	$\theta_H = 0.5$	$\theta_H = 0.6$	Data
E_m^{NL}	0.50	0.50	0.50	0.50	0.50	0.50
$\ln \overline{h_H}$	7.58	7.68	7.68	7.68	7.69	7.68
$\ln \overline{h_L}$	7.58	7.49	7.49	7.49	7.48	7.49
$\ln \overline{w_H}$	2.54	2.54	2.54	2.54	2.54	2.54
$\ln \overline{w_L}$	2.02	2.02	2.02	2.02	2.02	2.02
$sd(\ln h_H)$	0.26	0.23	0.23	0.22	0.22	0.22
$sd(\ln h_L)$	0.26	0.31	0.30	0.30	0.30	0.31
$sd(\ln w_H)$	0.56	0.55	0.56	0.57	0.58	0.56
$sd(\ln w_L)$	0.54	0.55	0.54	0.53	0.52	0.54

Notes: The table shows the calibration results for economies that vary in θ_H. The calibration targets 2/3 of the standard deviation of hours in the cross sectional data. All economies feature $\theta_L = 0.2$ and $\bar{h} = 5200$. The calibration searches for 7 parameters that minimize the sum of the square deviations between 9 statistics in the model and in the data.