NBER WORKING PAPER SERIES

WAGE AND EARNINGS INEQUALITY BETWEEN AND WITHIN OCCUPATIONS: THE ROLE OF LABOR SUPPLY

Andrés Erosa
Luisa Fuster
Gueorgui Kambourov
Richard Rogerson

Working Paper 31665
http://www.nber.org/papers/w31665

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
September 2023

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.

At least one co-author has disclosed additional relationships of potential relevance for this research. Further information is available online at http://www.nber.org/papers/w31665

Wage and Earnings Inequality Between and Within Occupations: The Role of Labor Supply
Andrés Erosa, Luisa Fuster, Gueorgui Kambourov, and Richard Rogerson
NBER Working Paper No. 31665
September 2023
JEL No. E20,J2,J3

ABSTRACT

We document systematic differences in wage and earnings inequality between and within occupations and show that these differences are intimately related to systematic differences in labor supply across occupations. We then develop a variant of a Roy model in which earnings are a non-linear function of hours, with the extent of this non-linearity differing across occupations. In our theory, the interplay between heterogeneity in tastes for leisure and occupational differences in non-linearities affects the sorting of workers. Moreover, this interplay is crucial to account for the facts on the distributions of hours, wages, and earnings within and across occupations.

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

1 Introduction

Recent empirical work on inequality emphasizes the value of approaching the data from the perspective of occupations (see, e.g., Autor and Dorn (2013)). The natural theoretical framework associated with this approach is the Roy (1951) model. In this paper we argue that textbook versions of the Roy model fail to account for some key features of inequality at the occupational level. In particular, whereas textbook versions of the Roy model do not distinguish between wage and earnings inequality, we document large and systematic differences between wage and earnings inequality across and within occupations driven by systematic differences in the distribution of hours worked across occupations. We develop a new variant of the Roy model that can account for the key quantitative patterns that we document.

The starting point for our analysis is to document the systematic and substantial differences in wage and earnings inequality both within and across occupations. Although related, we emphasize that wage and earnings inequality are of independent interest: wage inequality is more relevant when studying inequality of opportunity, but earnings inequality is more relevant when studying inequality in wealth or consumption. We first document that occupation differences in mean earnings are substantially larger than occupation differences in mean wages, and that within occupation variances of log earnings are much greater than within occupation variances of log wages. We next document that these gaps are systematically related to mean occupation wages: the gap between mean earnings and mean wages is increasing in occupational mean wages, and the gap between the within occupation variance of log earnings and log wages is decreasing in occupational mean wages.

Wages and earnings are intimately connected via hours worked. We next examine differences in the distribution of hours worked across occupations and show that these differences reflect the differences in wage and earnings inequality.[1] Specifically, we first show that there is a strong positive correlation between mean hours worked in an occupation and the mean wage in an occupation. And second, we document a strong negative correlation between the variance of log hours in an occupation and mean wages in an occupation. Combining these two facts yields a novel finding about labor supply and occupational choice: a strong negative correlation between mean hours in an occupation and the variance of log hours in an occupation.

Whereas textbook treatments of time allocation abstract from occupational choice, and textbook versions of Roy models take the time allocated to market work as given, the facts that we document suggest important interactions between time allocation and occupational choice. The second part of our paper develops a model that integrates time allocation decisions into an otherwise standard model of occupational choice. Our model features three occupations, heterogeneous tastes for leisure, and lognormal distributions over idiosyncratic taste and productivities. Building on the earlier work of Cogan (1981) and the more recent work of French (2005) and Prescott, Rogerson and Wallenius (2009), we assume a non-linearity in the mapping from individual hours worked to the supply of efficiency units of labor. But following Goldin (2014) and Erosa et al. (2022), a key innovation is that the

[1] The relationship between wage and earnings inequality is also affected by the correlation between hours and wages. We find that the cross-sectional correlation between wages and hours is modestly positive and very similar across occupations.

extent of this non-linearity differs across occupations.

We show theoretically that our model generates an intimate connection between occupational choice and time allocation: holding all else constant, an individual's choice of hours is affected by their occupational choice, and an increased taste for leisure by an individual will influence their occupational choice. These implications are not present if one assumes that earnings are linear in hours. Relative to a model with linear budget sets, our model implies that heterogeneous tastes for leisure influence both the allocation of individuals across occupations and the distribution of wages and earnings. These effects feature prominently in our quantitative analysis.

The third part of our paper carries out a quantitative analysis of our model. We calibrate the model and show that it can match key facts about wage and hours inequality between and within occupations. In particular, our model does a better job of accounting for these facts than a linear Roy model that allows for correlations between tastes for leisure and occupation-specific productivities. Our calibration procedure does not target the joint distribution between wages and hours and so does not target properties of the earnings distribution. Our model also does a much better job of matching the facts about earnings inequality between and within occupations than a linear Roy model.

A distinguishing prediction of our theory is that the choice of occupation affects optimal work hours. An implication of this effect is that workers with a low taste for leisure are more likely to sort into occupations with higher non-linearities, thus amplifying mean earnings differences between occupations. Although we find the choice of occupation has a sizeable impact on hours worked, differences in the sorting of workers account for most of the differences in mean occupational hours. The interaction between preference heterogeneity and occupational differences in non-linearities plays a key role in allowing our model to jointly account for the facts on earnings, hours, and wages.

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.[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 Goldin (2014) and Erosa et al. (2022). Like Goldin (2014), we study how occupational differences in the extent of non-linearities affect labor supply choices. But whereas Goldin (2014) focused on how this might qualitatively affect occupational choices of men and women, we calibrate our model and study the quantitative effects on the overall distributions of wages, earnings and hours. Our analysis shares

[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.

many features with that of Erosa et al. (2022), but pursues issues not addressed by them.[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 across occupations, and so did not address the systematic patterns in wage and earnings inequality within and across occupations. Second, the theoretical analysis in this paper identifies a channel through which differences in the non-linearity of earnings across occupations will generate differences in the dispersion of hours across occupations. Third, we calibrate the non-linearities in a three occupation model using cross-sectional moments, whereas Erosa et al. (2022) imposed values based on combining disparate pieces of information from various external studies. Fourth, we quantify the relative magnitude of occupation and selection effects in accounting for differences in hours distributions across occupations.

An outline of the paper follows. Section 2 presents our key empirical findings on wages, earnings and hours dispersion within and across occupations. Section 3 presents the simple benchmark model and Section 4 highlights the novel implications of our non-linear Roy model for occupational choice, hours and wages. Section 5 presents our calibration procedure and results and illustrates the quantitative aspects of the model's mechanisms. Section 6 examines the model's ability to account for features of earnings inequality in addition to wage inequality. Section 7 presents additional evidence relevant for a key prediction of our model, and Section 8 concludes.

2 Empirical Facts

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

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 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%

[4]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).

of the hourly wage distribution. Appendix A provides a detailed description of variables and sample restrictions.

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

2.2 Wage and Earnings Inequality Within and Across Occupations

We begin our empirical analysis by documenting substantial differences between wage and earnings inequality between and within occupations.

Comparing across occupations, the variance of log mean wages is 0.14 and the variance of log mean earnings is 0.22, implying that the variance of log mean earnings is about 55% larger than the variance of log mean wages across occupations. This difference is not driven by outliers in the tails; the interquartile range of the distribution of mean log earnings is around 30% higher than that of the distribution of mean log wages (0.65 versus 0.51).

Comparing within occupations, we find that the average within-occupation variance of log wages is 0.31 and the average within occupation variance of log earnings is 0.52, so that the within occupation variance of log earnings is about 70% larger than the within occupation variance of log wages.

Figure 1: Gaps in Earnings and Wages Across and Within Occupations, CPS 1976-2015: by 3-Digit Occupations.

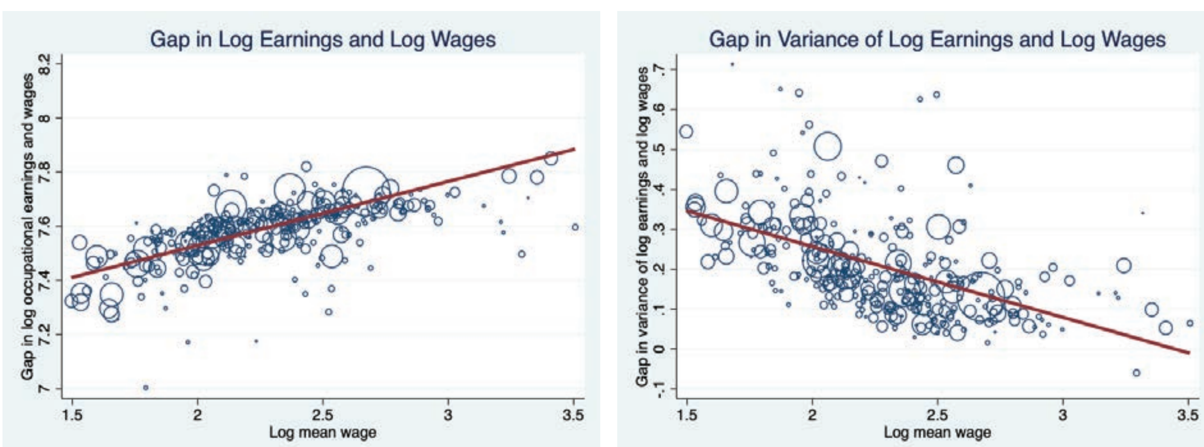

Notes: Each point represents a 3-digit occupation in the 1976-2015 time period. The left panel plots the gap in log mean earnings and log mean wages as a function of log mean occupation wage. The right panel plots the gap in the within occupation variance of log earnings and log wages as a function of log mean occupation wages. The size of the circle indicates the relative size of the occupation while the solid red lines are the fitted weighted regression lines.

Having documented large quantitative differences in inequality in wages and earnings both across and within occupations, we next document a systematic pattern in the relationship between these gaps and mean occupation wages. Figure 1 displays the results. In the left panel of Figure 1 we plot the gap between log mean earnings and log mean wages as a function of log mean occupational wages. In the right panel of Figure 1 we plot the gap between the within occupation variance of log earnings and the within occupation variance of log wages as a function of log mean occupation wages. The left panel reveals a robust positive relationship: occupations with high mean wages tend to exhibit a larger gap in mean log earnings and mean log wages. The right panel reveals a robust negative relationship: occupations with high mean wages tend to exhibit smaller gaps between the within occupation variance of log earnings and the within occupation variance of log wages.[6] Both panels include a linear regression line.[7]

Because hours generate gaps between wages and earnings, these patterns are suggestive about the inequality in hours within and across occupations. We turn to this issue in the next subsection.

2.3 Hours Inequality Within and Across Occupations

Motivated by the previous findings, in this subsection we investigate the variation in hours worked within and across occupations. Differences in log mean hours across occupations are large: log mean hours range from less than 7.4 to more than 7.7. This difference across occupations is similar to the aggregate differences observed between the US and countries in Western Europe. Differences in the within occupation variance of log hours are also large, varying from less than 0.05 to more than 0.30.

Next we document a systematic component to the variation in these two values. First, note that the left panel in Figure 1−a higher earnings-wage gap for occupations with higher mean wages−implies that mean wages are highly positively correlated with mean hours in an occupation. A more novel finding is that high mean wages are also associated with lower variance of log hours: Figure 2 shows a robust negative relationship between the within occupation variance of log hours and log mean wages.[8]

Not surprisingly given the two patterns documented above, one obtains a robust negative relationship between log mean hours and the within occupation variance of log hours. Figure 3 illustrates this pattern.[9] To the best of our knowledge this is a novel fact.[10] For this reason it is of interest to examine its robustness. Online Appendices D through F show that this negative relationship holds when we split the sample by gender (men, women), education (non-college, college), age (22-35, 36-49, and 50-64), time periods (ten-year period between 1976 and 2015), and full-time and part-time workers.[11] Furthermore, Online Appendix G

[6]This downward sloping relationship is driven by a downward sloping relationship between the variance in log earnings and log mean wages. There is no statistically significant relationship between the variance of log wages and log mean wages.

[7]In panel (a) the coefficient on log mean wage is 0.17 with a standard error of 0.01. In panel (b) the coefficient on log mean wage is -0.17 with a standard error of 0.02.

[8]The regression coefficient on log mean wages is -0.11 with a standard error of 0.01.

[9]The regression coefficient on the variance log hours is -1.07 with a standard error of 0.05.

[10]Online Appendix B shows that the main pattern remains unchanged if we were to use other measures of dispersion in hours as the coefficient of variation in annual hours or the 95/5 ratio and the 90/10 ratio of annual hours in an occupation. Online Appendix C shows that a similar pattern emerges when considering usual hours per week or weeks worked.

[11]One might conjecture that the pattern documented above is entirely driven by different propensities for

Figure 2: Mean Wages and Dispersion in Hours, CPS 1976-2015: by 3-Digit Occupations.

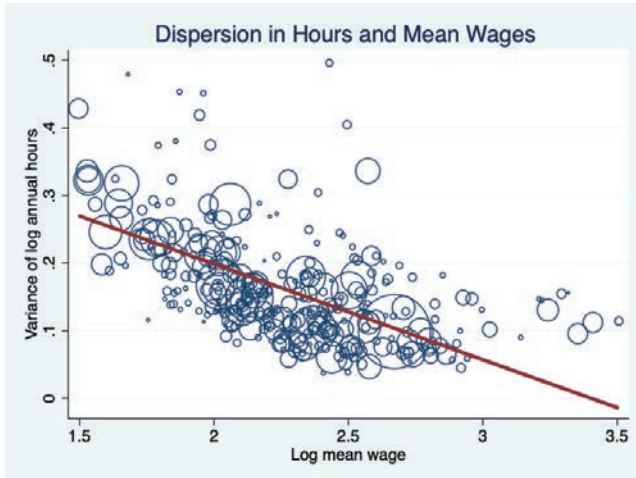

Notes: Each point represents a 3-digit occupation in the 1976-2015 time period. The scatter plot describes the relationship between the variance of log annual hours and log mean annual hours. The size of the circle indicates the relative size of the occupation while the solid red lines are the fitted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

Figure 3: Log Mean Annual Hours vs. the Variance of Log Annual Hours, CPS 1976-2015: by 3-Digit Occupations.

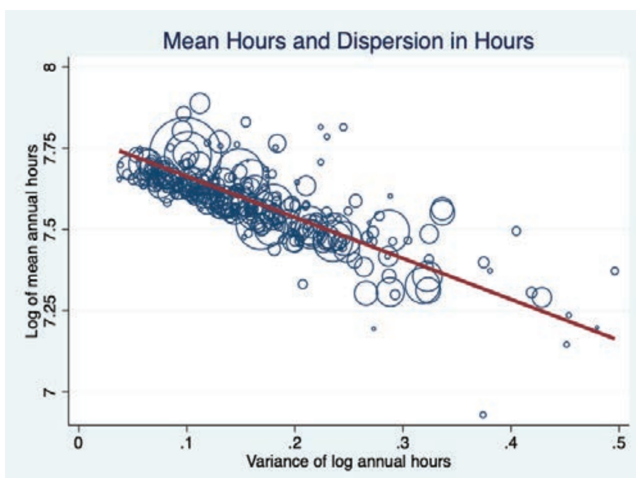

Notes: Each point represents a 3-digit occupation in the 1976-2015 time period. The scatter plot describes the relationship between log mean annual hours and the variance of log annual hours. The straight solid red line represents a linear regression, weighted by the relative size of each occupation,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

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 F shows that as occupation mean hours worked increases, there is shows that the position of an occupation in the mean-dispersion space is a relatively stable characteristic of an occupation over time.

2.4 Aggregating the Data to Three Occupations

Our theoretical and quantitative analysis focuses on a three-occupation model to maximize transparency and best highlight the forces at work. Connecting our model with the data requires that we aggregate three digit occupations into three broad groups. Because we focus on labor supply, we do this by ranking occupations by their mean hours and partitioning them into three groups of equal size based on person-level weights. We compute the log of mean hours, the variance of log hours, the log of mean wages, the variance of log wages, the log of mean earnings, and the variance of log earnings in each of the three occupational groups as a weighted average of the values for all occupations within each group.[12,13]

Table 1: Data Moments, CPS, 1976-2015.

	Emp. share	Log mean h	Var log h	Log mean w	Var log w	Log mean e	Var log e
H	1/3	7.70	0.10	2.61	0.33	10.32	0.46
M	1/3	7.59	0.15	2.28	0.28	9.87	0.48
L	1/3	7.46	0.24	1.93	0.29	9.41	0.60
gap H-M	0	0.114	-0.047	0.334	0.053	0.450	-0.012
gap L-M	0	-0.134	0.093	-0.346	0.013	-0.47	0.122

Notes: We rank occupations by their mean hours and partition them into three groups of equal size based on person-level weights. We denote the high, medium, and low mean hours occupations by H, M, and L, respectively. We compute the log of mean hours, the variance of log hours, the log of mean wages, the variance of log wages, the log of mean earnings, and the variance of log earnings in each of the three occupational groups as a weighted average of the values for all occupations within each group.

The moments of interest for our three "representative occupations," each accounting for the same employment share, are reported in Table 1. We denote the high, medium, and low mean hours occupations by H, M, and L, respectively. We briefly highlight the following patterns, all consistent with the facts documented at the 3-digit occupational level. First,

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).

[12] Appendix A provides more detail.

[13] Table H-1 in Online Appendix H provides an alternative way of computing the data moments. Instead of reporting the data moments for a representative occupation in an occupational group, we compute the data moments using all observations in that particular occupational group. By definition the reported means in both cases are the same, but the variances of log hours, log wages, and log earnings are slightly higher than those reported in Table 1.

occupations differ substantially in mean wages, mean earnings, mean hours, and in the within-occupation variances of log earnings, log wages, and log hours. Differences in within-occupation variances of log earnings are much larger than differences in the variance of log wages (0.14 versus 0.04). As a result, the gaps between the variance of log earnings and log wages vary substantially across occupations, ranging from 0.13 in occupation H to 0.31 in occupation L. Second, the gap between log mean earnings and log mean wages increases as we move to higher log mean wage occupations; whereas the gap in log mean wages between occupations H and L is 0.68, the gap in log mean earnings between these two occupations is more than 20 log points higher, at 0.91. Third, the gap between the within-occupation variance in log earnings and the within-occupation variance of log wages decreases as log mean wages increase. Fourth, the within-occupation variance of log hours decreases as the log mean wage increases falling from 0.24 in occupation L to only 0.10 in occupation H.

3 A Model of Occupational Choice and Time Allocation

In this section we present a generalized version of a three occupation Roy (1951) model. Our model generalizes a standard textbook Roy model along three dimensions. First, we incorporate an hours margin. This is essential in order to account for the distinction between wage and earnings inequality. Second, we allow for heterogeneity in tastes for leisure. Heterogeneity in hours is central to our analysis. Observables account for a small share of the cross-sectional variance in hours, so consistent with the literature, we rely on heterogeneity in preferences as an important source of hours dispersion.[14] And third, we allow for non-linearities in the mapping from hours to efficiency units, with this non-linearity potentially varying across occupations.

3.1 A Generalized Three Occupation Roy Model

There is a continuum of individuals of unit mass, indexed by i. Individual i has 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 $\gamma > 0$.

Two features of these preferences are worth noting. First, as is standard in the macro literature, we assume offsetting income and substitution effects. Second, the marginal utility of leisure 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 even as ϕ becomes very small, and seems a reasonable property to impose on preferences in a model that features preference

[14]See, for example the discussion in Bick, Blandin and Rogerson (2022).

heterogeneity. More specifically, the elasticity of the marginal utility of leisure with respect to hours of work is increasing in hours worked.[15]

There are three occupations, denoted by $j \in \{H, M, L\}$. Each occupation is a technology to produce the single final good of the economy according to a linear technology in efficiency units of labor:

$$Y_j = E_j, \tag{2}$$

where Y_j is the output from occupation j, E_j is the aggregate input of efficiency units of labor to occupation j, and we have implicitly chosen units so that TFP in each occupation is normalized to unity. As in a standard Roy model, each individual is endowed with a triplet of occupational specific productivities (a_{iH}, a_{iM}, a_{iL}). Heterogeneity across individuals is described by the four-tuple $(a_{iH}, a_{iM}, a_{iL}, \phi_i)$, which we assume is drawn from a log-normal distribution.

A novel feature of our framework is to allow for a non-linear mapping between hours and efficiency units, with the extent of the non-linearity potentially varying across occupations. In particular, if individual i supplies h_{ij} hours to occupation j, it will provide e_{ij} efficiency units of labor to occupation j, with e_{ij} given by:

$$e_{ij} = a_{ij} h_{ij}^{1+\theta_j}, \tag{3}$$

where $\theta_j \geq 0$ for $j \in \{H, M, L\}$. Without loss of generality we will impose that $\theta_H \geq \theta_M \geq \theta_L$.

In what follows we focus on two special cases of this model. The first special case imposes that $\theta_H > \theta_M > \theta_L \geq 0$ and will be referred to as the *non-linear* Roy model. Importantly, this specification imposes not only the presence of non-linearities in the earnings function but also that there is *heterogeneity* in the non-linearities across occupations. For many of the properties that we stress, the key issue is not whether the θ_j are positive, but rather whether they are heterogeneous. The second special case imposes that $\theta_j = 0$ for $j \in \{H, M, L\}$, and will be referred to as the *linear* Roy model. The key message from our quantitative work is that the non-linear Roy model is better able than the linear Roy model to account for the patterns documented in the previous section.

3.2 Equilibrium

We study a competitive equilibrium for this economy. The economy features four markets: one for the final good and one for efficiency units of labor in each of the three occupations. If we normalize the price of the final good to unity, the linearity of production functions with TFP normalized to one implies that the competitive equilibrium price of one efficiency unit of labor will equal unity in all occupations. Solving for an equilibrium then amounts to solving the individual's optimization problem and aggregating the solution across the distribution of individuals.

Given our normalizations, in equilibrium, an individual's earnings will be identical to their output, and their wage rate, defined as earnings divided by hours will be output divided by

[15]The elasticity of utility from leisure 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 implies a larger percentage decrease in hours of leisure.

hours. A simple implication of our earnings functions is that holding occupation constant, an individual's wage rate is increasing in their hours, with elasticity equal to θ_j. Despite the non-linear effect of hours on wages, it remains true in our model that for a given individual, the log of earnings is equal to the sum of log wages and log hours.

The decision problem of an individual characterized by the vector $\{a_{i,H}, a_{i,M}, a_{i,L}, \phi_i\}$ is:

$$\max_{c_i, \{h_{ij}\}_{j=H,M,L}} \left\{ \ln c_i + \phi_i \frac{\left(T - \sum_{j=H,M,L} h_{i,j}\right)^{1-\gamma}}{1-\gamma} \right\}$$

$$\text{subject to } c_i = \sum_{j=H,M,L} a_{ij} h_{ij}^{1+\theta_j}, \quad \sum_{j=H,M,L}^{3} h_{ij} \leq T, \; h_{ij} \geq 0.$$

This problem can be reformulated as a two-stage problem. In the first stage, individuals choose optimal hours conditional on an occupational choice, and in the second stage, they choose the optimal occupation taking these hours choices as given. In what follows we abstract from the set of measure zero of individuals who are indifferent between two or more occupations.

4 Qualitative Properties

In this section we present analytic results to highlight the differing forces in the non-linear and linear Roy models defined in the last section, in each case noting the significance of these differences for hours inequality within and between occupations. In the following sections we assess the quantitative importance of these differences.

We begin by analyzing the first stage decision for hours conditional on an occupational choice. Substituting the budget equation into the objective function and rearranging the first order condition for h_{ij} gives:[16]

$$\frac{1+\theta_j}{\phi_i} = h_{ij}(T - h_{ij})^{-\gamma} \equiv g(h_{ij}). \tag{4}$$

The function g defined in equation (4) is strictly increasing and convex (i.e., $g', g'' > 0$), has $g(0) = 0$ and satisfies $\lim_{h \to T} g(T) = \infty$. Thus, given $\frac{1+\theta_j}{\phi_i}$ there exist a unique value of h_{ij} that satisfies the first order condition.

Three results follow immediately from equation (4). First, h_{ij} is independent of skills $a_{i,j}$. This reflects our assumption on preferences that income and substitution effects are offsetting. Second, each of the h_{ij} are decreasing in the value of ϕ_i. Third, holding ϕ_i constant, a higher value of θ_j implies a higher value for h_{ij}. We state this result as Property 1.

Property 1: In the non-linear Roy model, hours for a given individual depend on occupational choice. In particular, $h_{iH} > h_{iM} > h_{iL}$. In the linear Roy model, hours for a given individual are independent of occupational choice, i.e., $h_{iH} = h_{iM} = h_{iL}$.[17]

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

[17] As we will see throughout this section, heterogeneity in the θ_j is key to this result. If the θ_j are all positive but equal then occupational choice has no impact on hours.

Because hours affect both wages and earnings, this property identifies a force in the non-linear model that leads to both wage differences across occupations and heterogeneity in the gap between wages and earnings across occupations.

Next, we establish that the non-linear model has implications for the *within* occupation hours distribution. In particular, we show that holding the distribution of ϕ_i constant across occupations, the dispersion in log hours decreases with θ_j. To show this, differentiate equation (4) 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. The elasticity, denoted by $\varepsilon_{h_{ij},\phi_i}$, decreases with hours worked:

$$\varepsilon_{h_{ij},\phi_i} = \frac{dh_{ij}}{h_{ij}} \Big/ \frac{d\phi_i}{\phi_i} = -\frac{1}{1+\gamma\frac{h_{ij}}{T-h_{ij}}}. \tag{5}$$

While the θ_j do not directly appear in expression (5), they influence the elasticity through their effect on the optimal choice of h_{ij}. It follows that holding the distribution of ϕ_i constant across occupations, $\theta_H > \theta_M > \theta_L$ implies that occupation H exhibits the highest mean hours and lowest dispersion of log hours and occupation L has the lowest mean hours and highest dispersion of log hours. Figure 4 illustrates graphically this result. In contrast, the linear model predicts that if the distribution of ϕ_i is the same across occupations, then the distribution of hours will be identical in all occupations.[18] We summarize this as Property 2.

Figure 4: Choice of h, Conditional on Occupation.

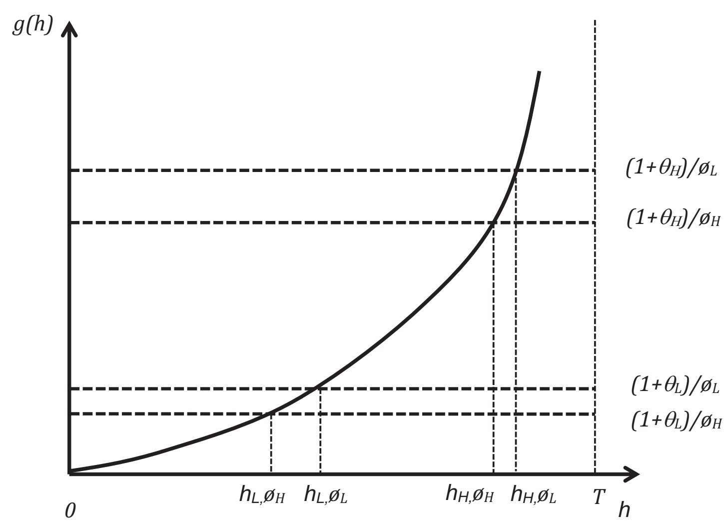

[18]It is again the case that this result also holds in a model with non-linear earnings if the non-linearity is the same across occupations.

Property 2: Holding the distribution of ϕ_i within an occupation constant, differences in θ_j across occupations generate a negative relationship between the mean and variance of log hours across occupations.

This property identifies a force in the non-linear model that will generate heterogeneity across occupations in the gap between the variance of log earnings and log wages.

The first two properties have highlighted the direct effect of heterogeneity in the θ_j on heterogeneity in hours. The next property will focus on indirect effects of heterogeneity in the θ_j that operate via selection effects. To pursue this, we examine the second stage decision in which an individual chooses an optimal occupation given the values of the h_{ij}. For concreteness, here we focus on the decision of whether to choose occupation H. Individual i compares the utilities of working in each occupation and decides to work in occupation H if and only if the following inequalities hold:

$$\ln\left(a_{iH}h_{iH}^{1+\theta_H}\right) + \phi_i \frac{(T-h_{iH})^{1-\gamma}}{1-\gamma} > \ln\left(a_{ij}h_{ij}^{1+\theta_j}\right) + \phi_i \frac{(T-h_{ij})^{1-\gamma}}{1-\gamma}, \text{ for } j = L, M \tag{6}$$

where h_{ij} are the solutions to equation (4) for $j = \{H, M, L\}$. Recalling that h_{ij} depends only on ϕ_i and not on a_{ij}, this expression can be re-arranged as:

$$\ln\left(\frac{a_{iH}}{a_{ij}}\right) > z_j(\phi_i), \text{ for } j = L, M \tag{7}$$

where

$$z_j(\phi_i) = -(1+\theta_H)\ln(h_{iH}) + (1+\theta_j)\ln(h_{ij}) + \phi_i\left[\frac{(T-h_{ij})^{1-\gamma}}{1-\gamma} - \frac{(T-h_{iH})^{1-\gamma}}{1-\gamma}\right].$$

Two results follow from these expressions. First, in the linear model both of the $z_j(\phi_i)$ are equal to zero, and we obtain the standard result from a Roy model without an hours decision: an individual chooses occupation H if and only if they have comparative advantage in occupation H.[19] In the non-linear model, recalling that $h_{iH} > h_{iM} > h_{iL}$ and applying the envelope theorem, it follows that $z_j'(\phi_i) > 0$ for both $j = M, L$. Holding all else constant, it follows that a higher value of ϕ_i increases the comparative advantage threshold for choosing occupation H, thereby making it less likely that the individual will choose this occupation.[20] We summarize this as Property 3.

Property 3: In the linear model, occupational choice is completely determined by the a_{ij}; in the non-linear model occupational choice is determined by both the a_{ij} and ϕ_i.

As noted earlier, ϕ_i is an important determinant of hours. Selection into occupations based on ϕ_i will influence the distribution of hours across occupations, which will in turn impact wage and earnings inequality within and across occupations.

[19] This result also holds as long as the θ_j are all equal, again highlighting the role of heterogeneity in the θ_j.

[20] A similar line of argument implies that high ϕ_i individuals are more likely to choose occupation L.

Property 3 shows that the distinct selection effects in the non-linear model have implications for how the distribution of ϕ_i might vary across occupations. It is thus of interest to consider the effects of differences in the distribution of the ϕ_i across occupations. The last property of the non-linear model that we highlight is that differences in the mean of ϕ_i across occupations will generate differences in the within occupation dispersion in hours. To see this, note that another implication of equation (5) is that a proportional decrease in ϕ_i 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. We summarize this as Property 4.

Property 4: In the non-linear model a proportional decrease in the distribution of ϕ_i within an occupation leads to an increase in mean hours and a decrease in the variance of log hours within the occupation.

This is significant for wage and earnings inequality for the same reasons as Property 2.

Selection effects may also affect the relative magnitude of the dispersion of $\log \phi_i$ across the three occupations. This effect will depend on the joint distribution of skills and tastes for leisure, and we cannot say anything about it at a general level. In our quantitative work, we find that selection serves to generate both a lower mean and variance for $\log \phi_i$ in occupation H than in occupations M and L, so that both effects play a role.

In closing this section we emphasize an implication that will influence the quantitative work that follows. In the linear model occupational choice depends solely on the a_{ij} and the choice of hours depends solely on ϕ_i. An immediate implication is that if ϕ_i is uncorrelated with the a_{ij}, then the hours distribution will be the same in all occupations. Put somewhat differently, generating differences in hours distributions across occupations in the linear model relies on the correlation structure between preferences and comparative advantage. In contrast, the analysis of this section has argued that the heterogeneity in the θ_j in the non-linear model can generate differences in hours distributions across occupations even if ϕ_i is uncorrelated with the a_{ij}.

5 Model Calibration

In this section we describe our procedure for calibrating the linear and non-linear Roy models and use the calibrated models to highlight the quantitative importance of the channels emphasized in the previous section.

5.1 Calibration Strategy

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. In particular, 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 T, γ and the θ_j. To see why, note that given values for T, γ and the θ_j, one can always find a value of ϕ_i and a value for a_{ij} in an individual's observed occupation to match their observations for hours and wages. One can then choose a value for the a_{ij} values in the other two occupations (e.g., set them equal to zero) that makes the observed occupational choice optimal. If hours and wages are perfectly matched then so are earnings. An important implication is that one can impose $\theta_H = \theta_M = \theta_L = 0$ and still match any pattern of hours and wages within and across occupations. Put somewhat differently, absent parametric restrictions, allowing for heterogeneous non-linearities across occupations does not improve the model's ability to account for the cross-sectional facts.

In this paper we follow a parametric approach, and have made the relatively standard assumption that individual heterogeneity in $(a_{iH}, a_{iM}, a_{iL}, \phi_i)$ is jointly lognormally distributed. As a practical matter, imposing log-normality in a standard Roy model (i.e., one without an hours margin) does not constrain its ability to do a good job of matching first and second moments for the distribution of log wages within and across occupations. Our first exercise will assess the extent to which this result generalizes when we also consider hours distributions. In particular, we ask whether our linear Roy model with lognormal distributions can account for the cross-sectional properties of both the wage and hours distributions within and across occupations.

As noted at the end of the previous section, if one imposes that tastes for leisure are not correlated with occupational productivities, the linear Roy model implies that the hours distribution will be identical across all occupations. That is, the sole channel through which the linear model can match differences in hours distributions across occupations is via the correlations between ϕ_i and the a_{ij}. Our second exercise seeks to assess the extent to which our non-linear Roy model with lognormal distributions offers an alternative channel that can account for the differences in hours distributions across occupations. To do this we impose that ϕ_i is uncorrelated with all of the a_{ij} and assess the extent to which the non-linear model can account for the cross-sectional distributions of wages and hours across occupations.

These two exercises are not nested; while the second exercise frees up three parameters that are restricted in the first exercise (i.e., the θ_j), it also restricts three of the free parameters in the first exercise (the correlations between ϕ_i and the a_{ij}). One question of interest will be the extent to which one of the models is better able to match the targets.

A second question of interest concerns the ability of the two models to match untargeted moments. In particular, in the introduction we stressed the different quantitative properties of inequality in wages and earnings across and within occupations. Our exercise separately targets the cross-sectional properties of hours and wage distributions across occupations, and so does not explicitly target any joint moments of the hours-wage distributions. As a result, we do not explicitly target any moments of the occupational earnings distributions. We will be especially interested in how well each of the models can match the differences between the moments for wages and earnings.

5.2 Calibration Details

With the parametric restriction to lognormality, the linear model has 16 parameters: the value of γ, the value of T, and 14 distributional parameters (two for the distribution of ϕ_i (mean and standard deviation), nine for the joint distribution of the a_{ij} (three means, three standard deviations and three correlations corresponding to each of the possible pairs of abilities), and three for the correlation between the a_{ij} and ϕ_i. As noted above, our restricted version of the non-linear model in which the correlations between the a_{ij} and ϕ_i are set to zero eliminates 3 parameters, but also adds 3 new parameters (the three θ_j), so also features 16 parameters.

Interpreting our model as predicting annual outcomes averaged over the life cycle we set $T = 5200$. We set the value of $\gamma = 4$.[21] This leaves 14 parameters for each model. For both models we find the values of the 14 remaining parameters which best fit the 14 moments listed in Table 2. Consistent with our earlier discussion, the list of targeted moments consists of cross-sectional moments of the occupational hours and wage distributions (for each occupation, the log of mean hours, the variance of log hours, the log of mean wages, and the variance of log wages), as well as the occupational employment shares (only two of which are independent). To find the best fit, we minimize a loss function that sums the square of the residuals between moments in the data and model.

Our calibration procedure uses 14 moments to identify 14 parameters in each of the two models. In some exercises of this nature one can point to particularly strong relationships between individual moments and individual parameters. This is not the case in our context, in the sense that changes in individual parameters tend to have substantial effects on several moments. The importance of selection effects makes this result somewhat intuitive. For example, increasing θ_H in the non-linear model tends to increase employment in occupation H, creating an intuitive link between the employment in occupation H and the value of θ_H. One might also expect that increasing θ_H would lead to an increase in average hours in occupation H, but this is not necessarily the case. Increasing the employment in occupation H affects the distribution of individual characteristics for workers in occupation H. These composition effects can both produce a substantial increase in the variance of hours in H and a substantial decrease in the mean wage in H.

This property of our model is perhaps not too surprising given the properties of standard two-occupation Roy models without an hours choice. In that model it is well known that occupational choices are affected by all the parameters characterizing the idiosyncratic productivity distribution, and that changes in the occupational employment distribution will typically affect all moments of the wage distribution.

For the reasons just discussed it is not possible to offer any simple intuition about which moments are largely determining which parameters. To provide some information on the mapping from parameters to moments, in Online Appendix I we present results from an exercise in which we consider small perturbations of each of the 14 individual parameters and evaluate the effect on each of the 14 moments used in our estimation exercise.

5.3 Calibration Results

In this section we assess the extent to which our two models are able to match the targeted moments. Parameter values and calibration results are shown in the top and bottom panels of Table 2, with the bottom row displaying the values of the loss function. The two models generate quite similar properties for the productivity distributions. In both cases the three

[21]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).

productivities are all very highly correlated, and the variance is decreasing as we move from H to M to L. The linear model requires a slightly larger value for the variance of tastes.

Table 2: Calibration of Non-linear and Linear Economies.

Description	Parameter	Non-linear	Linear
non-linearity H	θ_H	0.4490	0.0
non-linearity M	θ_M	0.3576	0.0
non-linearity L	θ_L	0.2673	0.0
corr (a_H, ϕ)	$\rho_{a_H,\phi}$	0.0	-0.0970
corr (a_M, ϕ)	$\rho_{a_M,\phi}$	0.0	-0.0009
corr (a_L, ϕ)	$\rho_{a_L,\phi}$	0.0	0.1563
corr (a_H, a_M)	ρ_{a_H,a_M}	0.9863	0.9507
corr (a_H, a_L)	ρ_{a_H,a_L}	0.9392	0.8527
corr (a_M, a_L)	ρ_{a_M,a_L}	0.9779	0.9077
mean ab occ. H	μ_{a_H}	-1.3631	1.9280
mean ab occ. M	μ_{a_M}	-1.3190	1.9914
mean ab occ. L	μ_{a_L}	-0.6888	1.9218
var ab occ. H	$\sigma^2_{a_H}$	0.4199	0.5077
var ab occ. M	$\sigma^2_{a_M}$	0.3532	0.3759
var ab occ. L	$\sigma^2_{a_L}$	0.2929	0.2631
mean taste for leisure	μ_ϕ	25.0072	24.7033
var taste for leisure	σ^2_ϕ	1.6371	1.6970
Target	**Data**	**Non-linear**	**Linear**
log mean hours occ. H	7.705	7.707	7.717
log mean hours occ. M	7.590	7.591	7.591
log mean hours occ. L	7.456	7.454	7.447
log mean wages occ. H	2.611	2.611	2.610
log mean wages occ. M	2.277	2.276	2.272
log mean wages occ. L	1.931	1.931	1.933
share of emp. occ. H	0.333	0.333	0.324
share of emp. occ. M	0.333	0.333	0.337
var log hours occ. L	0.239	0.238	0.223
var log wages occ. H	0.334	0.332	0.339
var log wages occ. M	0.281	0.287	0.291
var log wages occ. L	0.294	0.290	0.257
var log hours occ. H	0.099	0.100	0.114
var log hours occ. M	0.146	0.147	0.156
Loss Function $\times$ $(\mathbf{10^{-5}})$	–	6.47	246.7

Two messages emerge from this table. First, as is evident from both the detailed list of values for the 14 moments and the values of the loss functions, both models do a very good job overall of accounting for the cross-sectional wage and hours moments. But second, the non-linear model offers a closer fit to each of the moments, and in fact has a value of the loss function that is lower by a factor of almost 40 (6.4×10^{-5} relative to 246.7×10^{-5}).

It is useful to focus on how the two theories account for the gaps in outcomes across occupations, presented in Table 3. In what follows we measure gaps relative to occupation M. We begin with properties of the hours distribution. The gaps in mean hours in the data for occupations H and L are 0.114 and -0.137 respectively. The non-linear model matches those gaps almost exactly while the linear model slightly overpredicts both gaps. The gaps in the variance of log hours in the data are -0.047 and 0.093 in occupations H and L respectively. Again, the non-linear model matches those gaps almost exactly while the linear model slightly underpredicts both.

Table 3: Occupational Differences.

	Data	Non-linear	Linear
Log Mean Hours			
Occ H - Occ M	0.114	0.116	0.126
Occ L - Occ M	-0.134	-0.137	-0.144
Var Log Hours			
Occ H - Occ M	-0.047	-0.047	-0.042
Occ L - Occ M	0.093	0.091	0.067
Log Mean Wages			
Occ H - Occ M	0.334	0.335	0.338
Occ L - Occ M	-0.346	-0.345	-0.339
Var Log Wages			
Occ H - Occ M	0.053	0.045	0.048
Occ L - Occ M	0.013	0.003	-0.034
Emp shares			
Occ H	0.333	0.333	0.324
Occ M	0.333	0.333	0.337
Occ L	0.333	0.334	0.340

Notes: The table reports statistics on occupational differences in the distribution of hours and wages in the data, non-linear model, and linear model. The last panel reports employment shares for each of the three occupations.

Next we turn to occupational gaps in properties of the wage distributions. The occupational gaps in mean wages in the data are 0.344 and -0.346 in occupation H and L respectively. Both models match the data quite well, with the non-linear model exhibiting an almost perfect fit. The gaps in the variance of log wages in the data are 0.053 and -0.013 for occupations H and L respectively. These values are equal to 0.045 and 0.003 in the non-linear model and 0.048 and −0.034 in the linear model. Overall, even though both models match the data fairly well, the mean absolute discrepancy of the non-linear model is smaller.

Lastly, while the non-linear model perfectly matches the targets for employment shares, the linear model slightly understates the fraction of workers in occupation L (0.324 instead of 0.333).

5.4 Mechanisms

Section 4 highlighted the differing mechanisms through which the two models generate differences in hours distributions across occupations. In the linear model the key parameters are the correlations between tastes and productivities (the ρ_{a_{ij},ϕ_i}), while in the non-linear model it is heterogeneity in the θ_j. Consistent with this intuition, Table 2 shows that in the linear model the correlation between ϕ_i and a_{ij} is modestly negative for $j = H$, modestly positive for $j = L$, and effectively zero for $j = M$. This pattern serves to increase mean hours in occupation H and decrease mean hours in occupation L relative to the economy wide average.

For the non-linear model, Table 2 shows that θ_j increases from 0.27 in occupation L to 0.45 in occupation H. The mean value of the θ_j is equal to 0.35. This mean value is close to the value of 0.4 that was estimated by Aaronson and French (2004). To our knowledge our paper is the first to use a structural model to infer occupational gaps in non-linearity. Our estimated gap between θ_L and θ_H of 0.18 is only about half as large as the gap of 0.40 used by Erosa et al. (2022) in their benchmark calibration of a two occupation model.[22]

In Section 4 we highlighted that heterogeneity in the θ_j affect hours through both direct and indirect channels, with the indirect channel reflecting the effect on selection in ϕ_i. Using our calibrated model we can shed some light on the relative importance of these effects. To explore this we set $\sigma_\phi = 0$ and re-solve the model. In this specification the only source of hours heterogeneity is due to occupational choice. In the data, the gaps in the log of mean hours between H and M and between M and L are 0.114 and 0.134 respectively, and our calibrated model matches these very closely (0.116 and 0.137). But when we set $\sigma_\phi = 0$, these gaps are reduced to 0.018 and 0.020 respectively. We conclude that roughly 85 percent of the differences in mean hours across occupations in the non-linear model are due to the selection of individuals based on tastes. Consistent with this we find that selection has a significant impact on the mean of ϕ_i across occupations. Specifically, the population mean of ϕ_i is 0.575 while the profile for the mean value of ϕ_i across occupations H, M, and L is $(0.18, 0.58, 0.97)$

Selection effects also have a substantial impact on the variance of $\log \phi_i$ across occupations: the population variance for $\log \phi_i$ is 1.63 while the equilibrium profile for the variance of log ϕ_i across occupations H, M, and L is $(1.38, 1.48, 1.70)$. In our qualitative analysis we noted that the dispersion in hours within an occupation is influenced by both the mean and dispersion of ϕ_i. Variation in the within occupation dispersion in $\log \phi_i$ directly generates variation in the within occupation dispersion in hours, whereas variation in the mean of ϕ_i operates indirectly via its effect on the elasticity of hours with respect to ϕ_i, as seen in

[22] We note that our loss function is fairly flat with respect to increases in the value of this gap, so we do not view our calibration results as ruling out larger values.

equation (5). Both of these channels are quantitatively relevant. To establish this we hold occupational choices fixed, and shift the distribution of ϕ_i for individuals in occupation H proportionately so that mean hours of these workers is equal to mean hours of workers in occupation M. In this counterfactual occupations H and M have the same mean hours but differ in the value of the variance of $\log \phi_i$. We then compute the fraction of the original difference in the variance of log hours between occupations H and M that is still accounted for in this counterfactual. When comparing occupations H and M we find that dispersion in $\log \phi_i$ accounts for roughly 27% of the gap in the variance of log hours across occupations. When we repeat the exercise comparing occupations H and L the answer is 42%.[23] We conclude that more than half of the gaps in the variance of log hours are driven by differences in mean hours, via the elasticity effect reflected in equation (5). Consistent with this, we find large differences in the mean elasticity for individuals across occupations: the mean values for occupations H, M, and L are -0.27, -0.31, and -0.36 respectively.

Property 3 in Section 4 noted that occupational choice in the non-linear model is influenced by both ϕ_i and the a_{ij}, whereas in the linear model it is only influenced by the a_{ij}. To assess the role of preferences in shaping occupational choice we eliminate preference heterogeneity in the non-linear model and re-solve the model. We find that almost fifteen percent of workers will choose a different occupation.[24]

6 Implications for Earnings Inequality

A key objective of our model building exercise is to have a model of occupational choice that can simultaneously account for the quantitative patterns of inequality in both wages and earnings. In the previous section we showed that our estimated linear and non-linear models both do a good job of accounting for the patterns of wage inequality within and across occupations. In this section we compare how the two models fare with regard to the facts about earnings inequality, and in particular the differences between wage and earnings inequality. Our key finding is that the non-linear model does a good job of accounting for these untargeted moments, and in particular, does much better than the linear model.

Results are reported in Table 4. The top two panels report results for the mean and variance of log earnings across occupations. Both models match the values for mean earnings from the data quite closely, though the non-linear model does a slightly better job of matching the gaps across occupations. The non-linear model does a substantially better job of matching the variances, as the linear model misses the values for occupation M by more than twenty percent and occupation L by more than thirty percent.

Next we examine the extent to which the two models help us understand the differences in the within occupation variance of wages and earnings. These values are presented in the third panel of Table 4. Consistent with the results just reported for variances of log earnings, we see that the non-linear model matches the data quite closely whereas the linear model

[23] Recall from an earlier calculation that the differences in mean hours are about 85% due to differences in the ϕ_i distributions, with the remainder due to the direct effect of θ_j on hours.

[24] Specifically, the switching rates are roughly 10%, 18% and 15% for workers initially in occupations H, M, and L respectively. The vast majority of switches out of H are to M, and the vast majority of switches out of L are to M.

Table 4: Non-targeted Dimensions.

	Data	Non-linear	Linear
Mean Log Earnings			
Occ H	10.322	10.339	10.313
Occ M	9.872	9.889	9.834
Occ L	9.407	9.420	9.336
Log Earn Gap H-M	0.449	0.450	0.479
Log Earn Gap L-M	-0.466	-0.469	-0.497
Var Log Earnings			
Occ H	0.464	0.480	0.413
Occ M	0.476	0.486	0.379
Occ L	0.598	0.621	0.381
Var log earn- Var log wages			
Occ H	0.130	0.147	0.074
Occ M	0.195	0.199	0.088
Occ L	0.304	0.331	0.123
Corr of log hours and log wages			
Occ H	0.075	0.130	-0.102
Occ M	0.115	0.127	-0.161
Occ L	0.120	0.177	-0.208

Notes: The table reports (non-targeted) statistics on mean log earnings, variance of log earnings, gap between the variance of log earnings and the variance of log wages, and the correlation of log hours and log wages. The columns correspond to the data, the non-linear model, and the linear model.

misses substantially in all cases, and in particular by more than fifty percent for both the M and L occupations.

The fourth panel of Table 4 helps us to understand the reason for the different results across models. It reports the cross-sectional correlation of wages and hours in each of our three occupations, in the data and the two calibrated models. In the data and the non-linear model these correlations are modestly positive, whereas in the linear model they are modestly negative. Non-linearities are key to this result: in the non-linear model the earnings function generates a positive correlation between hours and wages holding individual productivity fixed, whereas in the linear model wages are independent of hours worked holding productivity constant. To show that this effect plays a key role we compute the within occupation correlations between hours and productivity (i.e, the value of a_{ij} in the individual's observed occupation) in the non-linear model. In contrast to the modestly positive correlations between wages and hours, we find modestly negative correlations between individual productivity and hours. The profile of correlations between productivity and hours is -0.12, -0.13, and -0.07 for occupations H, M, and L respectively. The average of these values is -0.11, which is quite similar to the average value of -0.15 for the linear model. We conclude that non-linearities play a fundamental role in providing a better account of the data.

7 Suggestive Evidence for Occupation Effects on Hours

A key distinguishing feature of our non-linear model relative to the linear model is the presence of a causal effect of occupation on hours of work. In this section we present some suggestive evidence supportive of this effect.

It is intuitive that data on occupation switchers might be useful in assessing the significance of occupation effects on hours. In the current context this presents two challenges. First, because our static model framed occupational choice as a career choice, it implicitly assumed that there is no switching between occupations. Second, it is important that the occupational change is not associated with a change in tastes for leisure. We propose that examining young occupational switchers offers a plausible way to deal with both of these challenges. In particular, consistent with models that feature learning, occupational switches of young individuals could be viewed as part of a process in which they are learning about their productivities in order to choose an optimal career. Jovanovic (1979) is an early model that stresses learning about productivity as a source of worker turnover, while Papageorgiou (2014) studies this in the context of a model of occupational choice. From this perspective, occupational switches of young individuals are not driven by changes to tastes, and reflect individuals choosing among careers, which is consistent with how we view our model. While not definitive, we think this evidence is suggestive about the existence of occupation effects.

Table 5: Changes in Log Hours upon an Occupational Switch.

	Data	Non-linear	Linear
H to M	-0.034	-0.019	0.000
H to L	-0.079	-0.046	0.000
M to H	0.042	0.019	0.000
M to L	-0.026	-0.023	0.000
L to H	0.095	0.040	0.000
L to M	0.087	0.024	0.000

Notes: Data from the 1990 Survey of Income and Program Participation. Moments are computed for young workers between the ages of 22 and 35.

We implement this exercise using data from the Survey of Income and Program Participation (SIPP). We use this data rather than the CPS because it offers a longer panel component with a still relatively large sample size. Details of the data sample and procedure are described in Appendix A. The first column of Table 5 reports the results from the data for all six of the potential occupational switches. Interestingly, the signs of all six changes are consistent with the prediction of our non-linear model.

To gauge the extent to which the magnitudes are in line with our model predictions we carry out the following simple exercise. If individuals are acquiring information about the relative benefits of different careers then it seems plausible that switchers may be those who are closest to the boundary of indifference between occupations. With this in mind we compute the implied change in hours in our calibrated model for all individuals whose benefit measured in terms of consumption is less than 5% of equilibrium consumption. The resulting values are reported in the second column of Table 5, and the values in the nonlinear model are quantitatively similar to those observed in the data.

A more definitive analysis of this issue would require that we extend our model to a dynamic setting and explicitly model the reasons for occupational mobility. We leave such an exercise for future work. But we view the above evidence on signs and magnitudes to be suggestive evidence in support of the presence of occupational effects on hours.

8 Conclusion

Recent work on inequality has found it useful to view the data from the perspective of occupations. In this paper we have argued that systematic differences in labor supply across occupations have important effects for the measurement of inequality between and within occupations. In particular, we document large and systematic differences in wage and earnings inequality between and within occupations, and show that they are intimately related to differences in the distribution of hours worked across occupations. In particular, high mean wage occupations tend to have higher mean hours and lower variance of log hours. These patterns amplify differences in mean earnings across occupations, and create greater variance in log earnings in low mean wage occupations.

We develop a variant of standard Roy model that can account for these patterns. Building on the work of Goldin (2014) and Erosa et al. (2022), the key feature of our model is that individual earnings are a non-linear function of individual hours worked, and the extent of this non-linearity varies across occupations. We show that this model is better able to account for the patterns in the data than a model with linear earnings. A distinctive implication of our model relative to a standard textbook Roy model is that occupational choice is determined both by comparative advantage as well as tastes for leisure. This interaction between preferences and occupation choice has important quantitative effects on how individuals sort across occupations and the distribution of hours worked within and across occupations.

Our model has viewed occupational choice as a career choice. An important direction for future work in this area is to extend our analysis to include life cycle dynamics. This would allow us to distinguish between static and dynamic source of non-linearities in earnings, and provide a framework in which data on occupation switchers can be used to identify occupation fixed effects.

References

Aaronson, D. and E. French, "The Effect of Part Time Work on Wages: Evidence from Social Security Rules," *Journal of Labor Economics* 22 (2004), 329–352.

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.

Bick, A., A. Blandin and R. Rogerson, "Hours and Wages," *Quarterly Journal of Economics* 137 (2022), 1901–1962.

Chetty, R., "Bounds on Elasticities With Optimization Frictions: A Synthesis of Micro and Macro Evidence on Labor Supply," *Econometrica* 80 (2012), 969–1018.

Cogan, J., "Fixed Costs and Labor Supply," *Econometrica* 49 (1981), 945–963.

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).

Jovanovic, B., "Job Matching and the Theory of Turnover," *Journal of Political Economy* 87 (October 1979), 972–990.

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

Papageorgiou, T., "Learning Your Comparative Advantages," *Review of Economic Studies* 81 (2014), 1263–1295.

Prescott, E. C., R. Rogerson and J. Wallenius, "Lifetime Aggregate Labor Supply with Endogenous Workweek Length," *Review of Economic Dynamics* 12 (2009), 23–36.

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.

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).[25]

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.[26] 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, (log) mean earnings, the standard deviation of log hours, the standard deviation of log wages, and the standard deviation of log earnings 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 three occupational 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 three groups that are equal in size, based on person-level weights. We then compute (log) mean hours, (log) mean wages, (log) mean earnings, the standard deviation of log hours, the standard deviation of log wages, and the standard deviation of log earnings in each occupation, using person-level weights in the analysis. Finally, for each of the three sectors, we report the averages of these moments, using the relative share of individuals in each occupation in that sector.

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

[26] 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 three occupational groups: Occupation H, M or L. We identify occupational switchers 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 restrict the analysis to young workers between the ages of 22 and 35. We find that individuals that stay in their current occupation do not experience on average any change in their hours worked. Occupational switchers, however, on average experience significant changes in their hours worked.

ONLINE APPENDICES

B 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 variance 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 (Figure B-1) or the 95/5 and 90/10 ratio of annual hours in an occupation (Figure B-2).

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

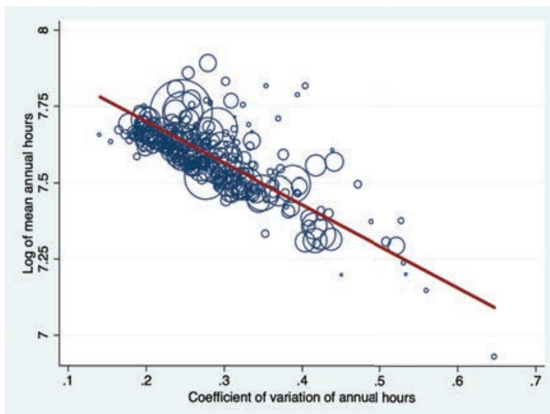

Notes: Each point represents a 3-digit occupation in the 1976-2015 time period. The straight solid red line represents a linear regression, weighted by the relative size of each occupation, while the size of the circle indicates the relative size of the occupation.

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

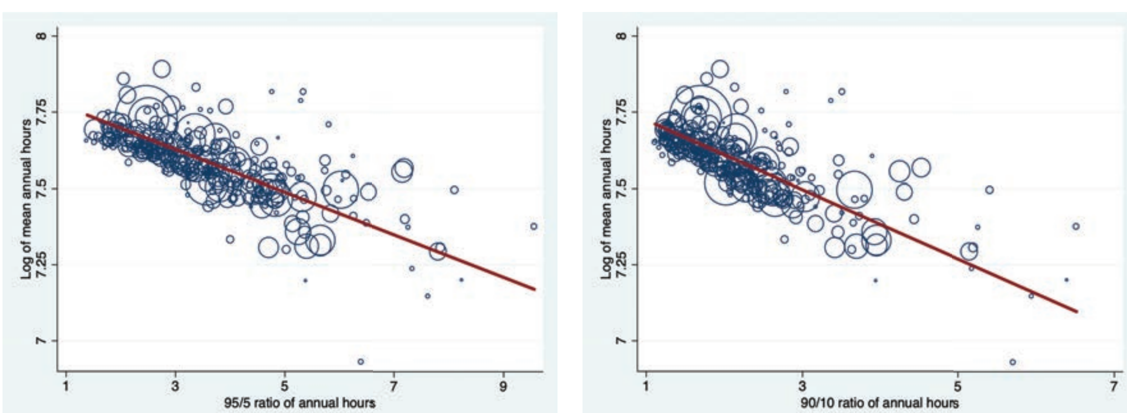

Notes: Each point represents a 3-digit occupation in the 1976-2015 time period. The straight solid red line represents a linear regression, weighted by the relative size of each occupation, while the size of the circle indicates the relative size of the occupation.

C Intensive and Extensive Margin: 1976-2015

Figure C-1 shows the relationship between log mean weeks and the variance 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: Log Mean Weeks vs. the Variance of Log Weeks, CPS, 1976-2015, by 3-Digit Occupations.

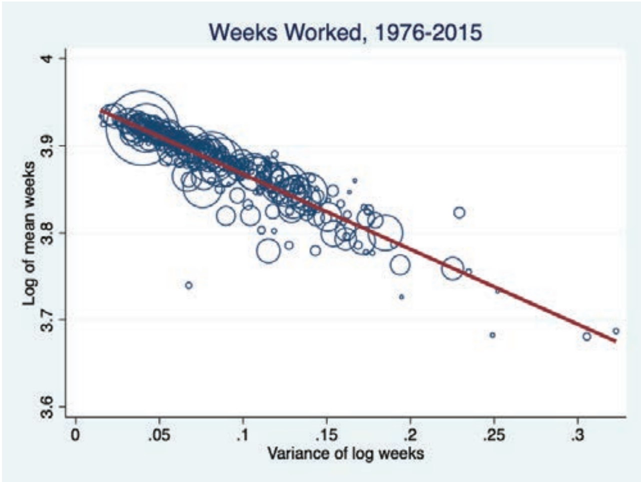

Notes: Each point represents a 3-digit occupation in the 1976-2015 time period. The straight solid red line represents a linear regression, weighted by the relative size of each occupation, while the size of the circle indicates the relative size of the occupation.

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

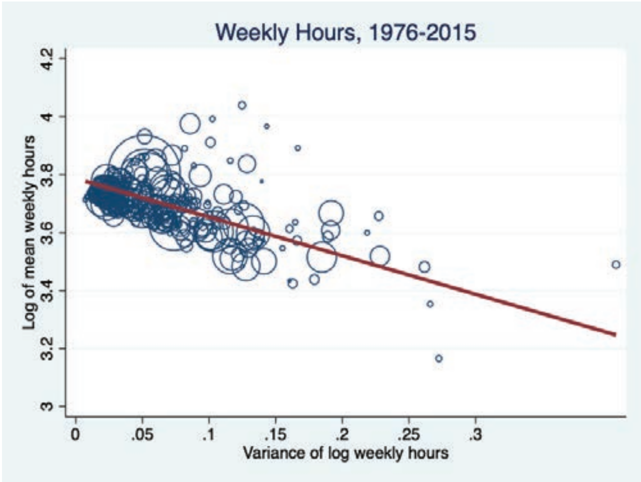

Notes: Each point represents a 3-digit occupation in the 1976-2015 time period. The straight solid red line represents a linear regression, weighted by the relative size of each occupation, while the size of the circle indicates the relative size of the occupation.

Figure C-2 shows the relationship between log mean weekly hours and the variance 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 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-3: Variance of Log Weekly Hours vs. Variance of Log Weeks, CPS, 1976-2015, by 3-Digit Occupations.

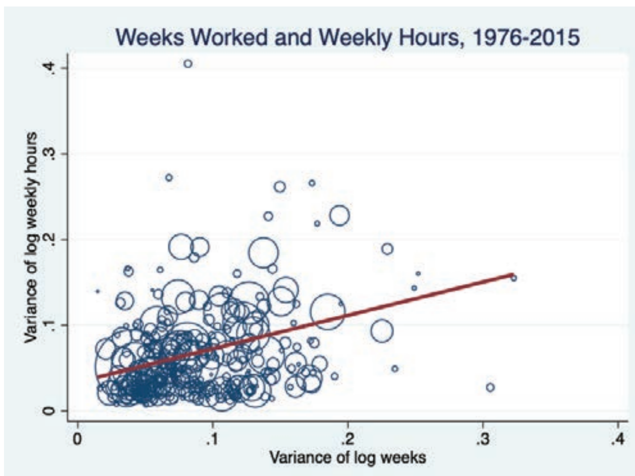

Notes: Each point represents a 3-digit occupation in the 1976-2015 time period. The straight solid red line represents a linear regression, weighted by the relative size of each occupation, while the size of the circle indicates the relative size of the occupation.

D Age, Gender, and Education Groups: 1976-2015

Figure D-1 illustrates that the main pattern observed in the data−of a negative relationship between mean hours and the dispersion in hours in an occupation−holds when we split the sample by gender (men, women), education (non-college, college), and age (22-35, 36-49, and 50-64).

Figure D-1: Log Mean Annual Hours vs. the Variance of Log Annual Hours, CPS, 1976-2015, 3-Digit Occupations: by Gender, Education, and Age.

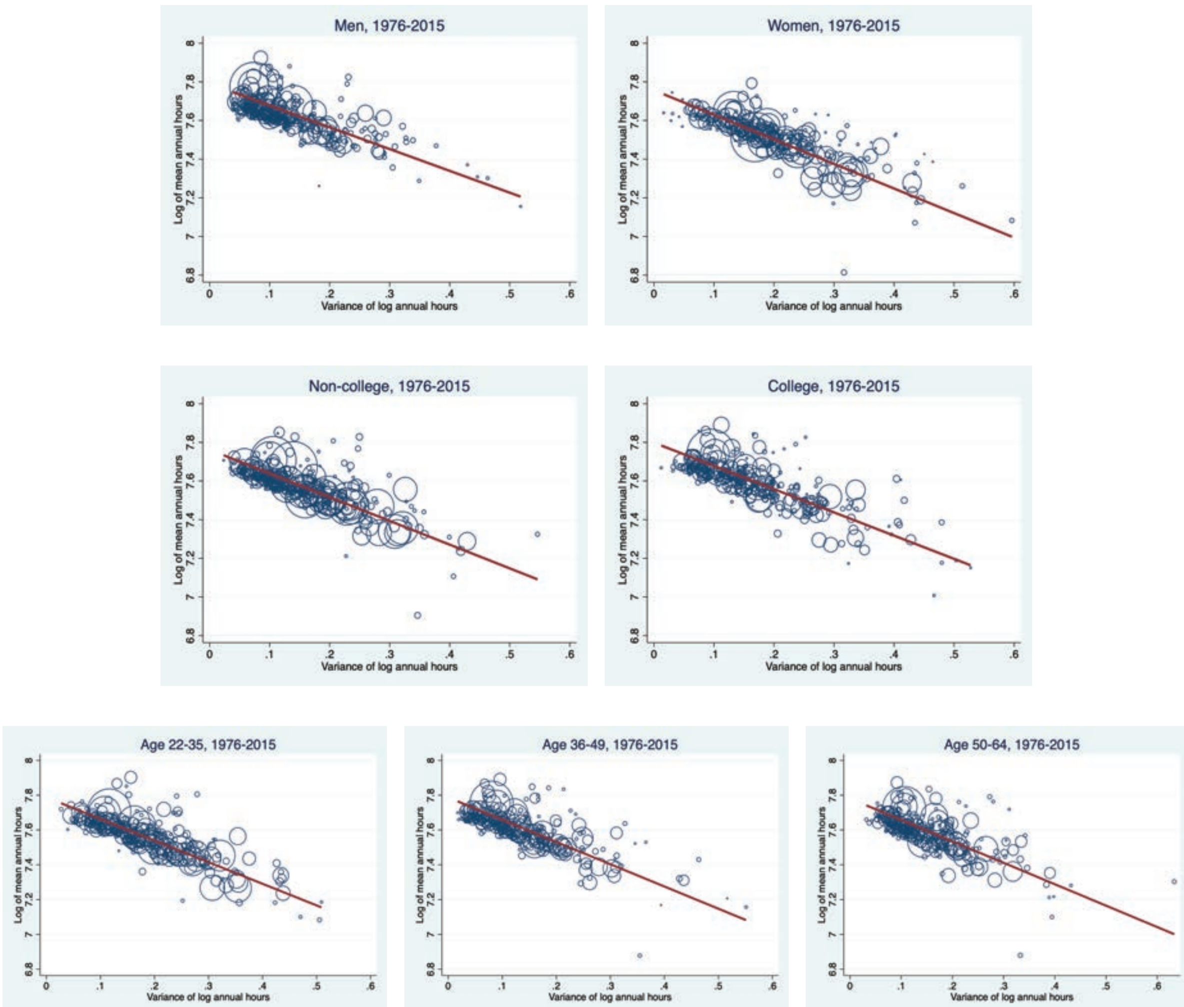

Notes: Each point represents a 3-digit occupation in the 1976-2015 time period. The straight solid red line represents a linear regression, weighted by the relative size of each occupation, while the size of the circle indicates the relative size of the occupation.

E Time Periods: 1976-2015

Figure E-1 illustrates that the main pattern observed in the data−of a negative relationship between mean hours and the dispersion in hours in an occupation−holds if we consider each successive ten year period between 1976 and 2015.

Figure E-1: Log Mean Annual Hours vs. the Variance of Log Annual Hours, CPS, 1976-2015, 3-Digit Occupations: by Different Time Periods.

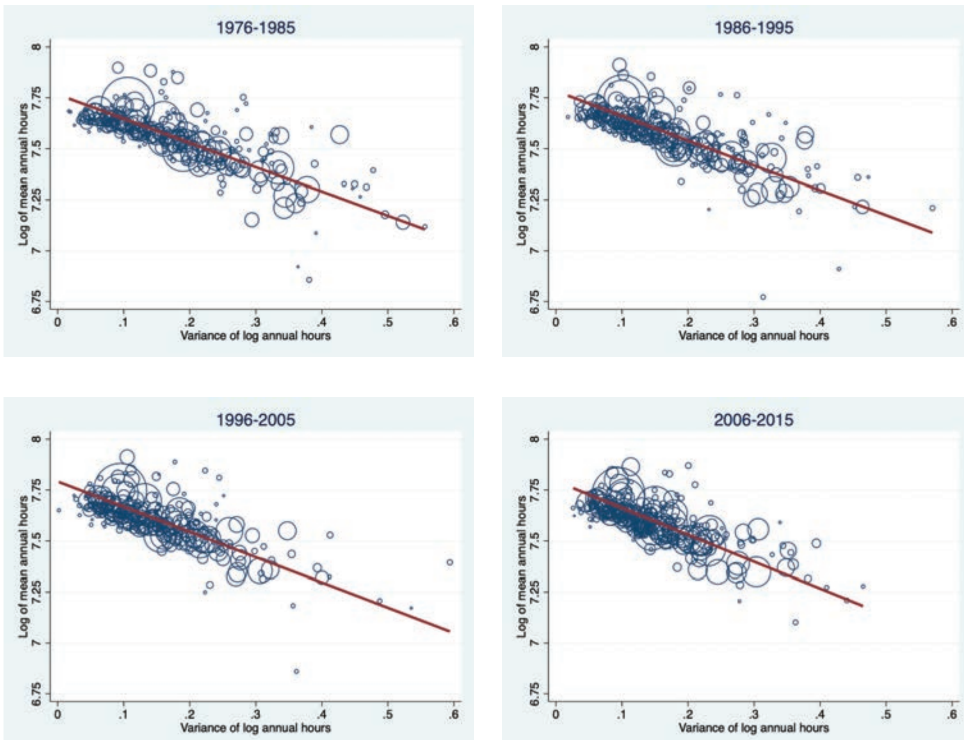

Notes: Each point represents a 3-digit occupation in the given 10-year time period. The scatter plot describes the relationship between the log of mean annual hours worked and the variance of log annual hours in a given occupation. The straight solid red line represents a linear regression, weighted by the relative size of each occupation, while the size of the circle indicates the relative size of the occupation.

F Part-Time vs. Full-Time Workers: 1976-2015

Figure F-1 illustrates that the main pattern observed in the data—of a negative relationship between mean hours and the dispersion in hours in an occupation—holds if we consider part-time (less than or equal to 34 hours per week) and full-time (more the 34 hours per week) workers.

Figure F-1: Log Mean Annual Hours vs. the Variance of Log Annual Hours, CPS, 1976-2015, 3-Digit Occupations: Part-Time vs. Full-Time.

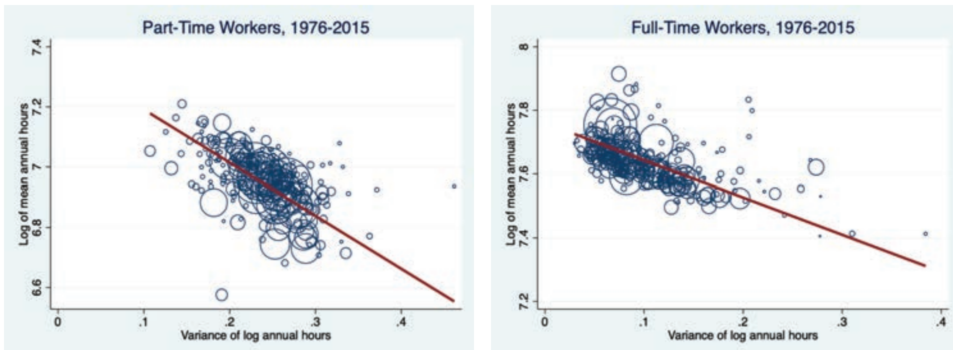

Notes: Each point represents a 3-digit occupation in the 1976-2015 time period. The straight solid red line represents a linear regression, weighted by the relative size of each occupation, while the size of the circle indicates the relative size of the occupation.

Further, 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 F-2: Fraction Working "Short" and "Long" Hours, CPS, 1976-2015: by 3-Digit Occupations.

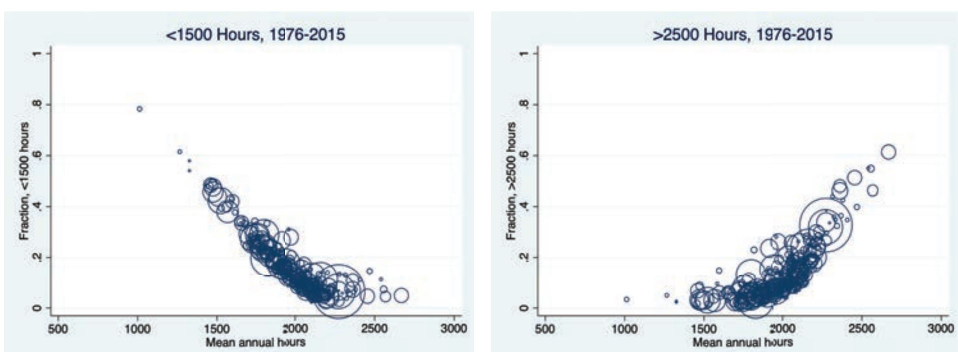

Notes: Each point represents a 3-digit occupation in the 1976-2015 time period. The size of the circle indicates the relative size of the occupation.

Figure F-2 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.

G Occupational Hours over Time: 1976-2015

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 G-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 variance of log annual hours, as reported in Figure G-2. The plots for the changes in the variance over time exhibit more dispersion around the 45 degree line, but this is to be expected if the estimate of the variance of hours within an occupation is noisier than the estimate of mean hours.

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

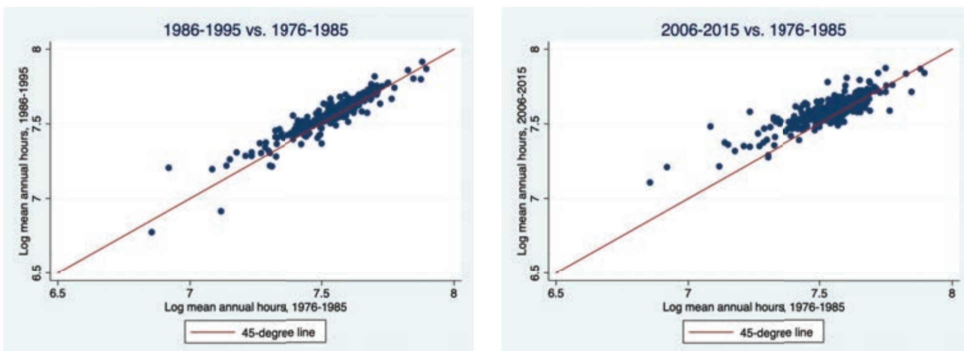

Notes: Each point represents a 3-digit occupation in a given 10-year time period. The scatter plot describes the change in log of mean annual hours in a given occupation over time, relative to the 1976-1985 period.

Figure G-2: Variance of Log Annual Hours, Over Time: by 3-Digit Occupations.

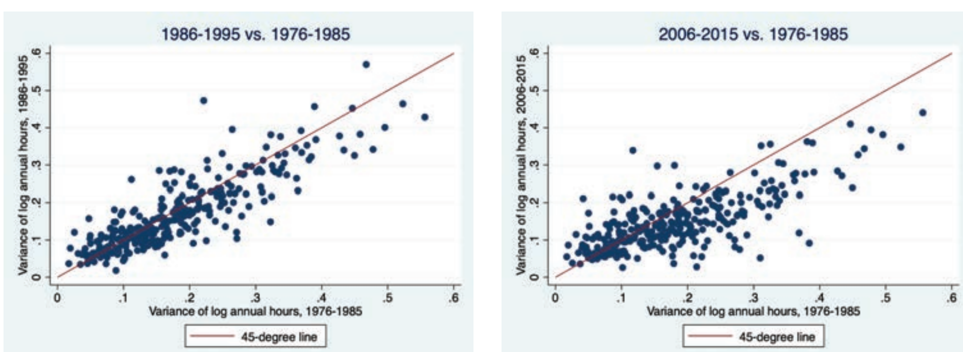

Notes: Each point represents a 3-digit occupation in a given 10-year time period. The scatter plot describes the change in the variance of log annual hours in a given occupation over time, relative to the 1976-1985 period.

H Data Moments in an Occupational Sector, an Alternative Computation

We provide an alternative way of computing the data moments reported in Table 1. Instead of reporting the data moments for a representative occupation in an occupational group, we compute the data moments using all observations in that particular occupational group. By definition the reported means in both cases are the same, but the variances of log hours, log wages, and log earnings are slightly higher than those reported in Table 1.

Table H-1: Data Moments, CPS, 1976-2015.

	Emp. share	Log mean h	Var log h	Log mean w	Var log w	Log mean e	Var log e
H	1/3	7.70	0.10	2.61	0.41	10.32	0.55
M	1/3	7.59	0.15	2.28	0.34	9.87	0.56
L	1/3	7.46	0.25	1.93	0.36	9.41	0.71
gap H-M	0	0.11	-0.05	0.33	0.07	0.45	-0.01
gap L-M	0	-0.13	0.10	-0.35	0.02	-0.46	0.15

Notes: We rank occupations by their mean hours and partition them into three groups of equal size based on person-level weights. We denote the high, medium, and low mean hours occupations by H, M, and L, respectively. We compute the log of mean hours, the variance of log hours, the log of mean wages, the variance of log wages, the log of mean earnings, and the variance of log earnings in each of the three occupational groups.

I Elasticity of Targets to Parameter Changes

Tables I-1 and I-2 report the elasticity of the calibration targets to a 1% change in each of the parameters in the linear and nonlinear models.

Table I-1: Elasticity to Parameter Changes: Non-linear Model.

Target	μ_{a_H}	μ_{a_M}	μ_{a_L}	μ_ϕ	$\sigma^2_{a_H}$	$\sigma^2_{a_M}$	$\sigma^2_{a_L}$	σ^2_ϕ	ρ_{a_H,a_M}	ρ_{a_H,a_L}	ρ_{a_M,a_L}	θ_H	θ_M	θ_L
Emp H	-11.55	10.41	0.38	-0.14	3.11	-3.08	-1.57	-0.04	-1.38	-4.02	7.05	29.68	-22.03	-1.24
Emp M	10.88	-20.34	5.26	-0.04	-1.11	5.72	-1.15	-0.05	-6.75	6.92	-5.42	-27.51	41.44	-14.98
Emp L	0.68	9.86	-5.61	0.18	-1.99	-2.63	2.70	0.09	8.09	-2.89	-1.64	-2.21	-19.28	16.14
Mean h H	0.06	-0.06	-0.01	-0.02	-0.12	-0.05	0.01	0.01	0.15	0.05	-0.06	-0.11	0.09	0.01
Mean h M	0.09	-0.01	-0.03	-0.02	0.09	0.00	-0.07	0.00	-0.12	0.00	0.11	-0.31	0.10	0.02
Mean h L	0.01	0.09	-0.06	-0.02	-0.03	0.07	0.13	-0.02	0.12	-0.05	-0.19	-0.06	-0.28	0.23
Var h H	-3.25	2.02	0.22	0.21	6.45	2.94	-0.61	0.54	-7.04	-1.94	2.17	4.69	-3.26	-0.31
Var h M	-0.79	-0.95	0.87	0.35	-1.80	1.81	3.26	1.21	1.38	0.14	-4.79	4.12	-0.61	-0.95
Var h L	-0.15	-1.86	1.24	0.31	0.30	-0.57	-1.85	1.29	-1.76	0.57	2.44	0.53	4.85	-3.68
Mean w H	0.47	-0.92	-0.06	-0.02	-0.44	-0.28	0.19	0.01	0.89	0.65	-1.04	-0.96	2.02	0.19
Mean w M	0.38	-0.20	-0.39	0.00	0.95	1.33	-0.78	0.01	-1.68	-0.40	0.85	-1.24	0.67	1.18
Mean w L	0.04	0.74	-0.83	-0.03	-0.14	0.43	2.38	-0.04	0.79	-0.59	-2.24	-0.14	-1.56	2.28
Var w H	-0.53	1.03	0.17	-0.07	2.13	0.82	-0.16	-0.02	-1.92	-2.48	3.02	1.27	-1.50	-0.77
Var w M	0.22	-1.96	0.32	0.12	1.76	4.44	2.77	0.14	-4.22	0.52	-3.05	-1.06	3.45	-1.02
Var w L	-0.07	0.35	-0.18	-0.06	-0.21	0.63	4.21	-0.07	0.49	-1.20	-2.00	0.03	-0.88	0.52

Notes: The table reports the elasticity of each targeted moment to a 1% change around the calibrated parameter value.

Table I-2: Elasticity to Parameter Changes: Linear Model.

Target	μ_{a_H}	μ_{a_M}	μ_{a_L}	μ_ϕ	$\sigma^2_{a_H}$	$\sigma^2_{a_M}$	$\sigma^2_{a_L}$	σ^2_ϕ	ρ_{a_H,a_M}	ρ_{a_H,a_L}	ρ_{a_M,a_L}	$\rho_{a_H,\phi}$	$\rho_{a_M,\phi}$	$\rho_{a_L,\phi}$
Emp H	8.38	-6.49	-2.19	0.00	1.43	-1.18	-0.51	0.00	-4.42	-1.27	2.25	0.000	0.000	0.000
Emp M	-6.07	11.86	-5.18	0.00	-0.43	1.92	-0.21	0.00	-5.45	2.30	-1.56	0.000	0.000	0.000
Emp L	-1.96	-5.57	7.21	0.00	-0.93	-0.77	0.69	0.00	9.61	-1.06	-0.59	0.000	0.000	0.000
Mean h H	-0.07	0.05	0.03	-0.01	-0.07	0.00	0.00	0.00	0.28	0.04	-0.04	0.012	0.000	0.003
Mean h M	-0.08	0.00	0.06	-0.01	0.04	0.01	-0.01	0.00	-0.12	0.00	0.03	-0.010	0.000	0.010
Mean h L	-0.03	-0.05	0.08	-0.01	0.00	0.00	0.04	0.00	0.02	-0.05	-0.04	-0.004	0.000	-0.016
Var h H	1.90	-0.70	-0.47	0.18	2.51	0.37	0.06	1.34	6.37	-1.74	0.69	-0.411	0.002	-0.148
Var h M	1.12	0.42	-1.43	0.17	-0.13	0.15	0.32	1.34	7.27	-0.25	-1.02	0.058	-0.003	-0.325
Var h L	0.29	0.69	-1.29	0.17	0.08	-0.05	-0.11	1.30	-8.48	0.38	0.68	-0.004	0.001	0.071
Mean w H	-0.06	0.56	0.35	0.00	0.05	-0.06	0.03	0.00	0.86	0.23	-0.38	0.000	0.000	0.000
Mean w M	-0.31	0.49	0.59	0.00	0.18	0.66	-0.14	0.00	-1.88	-0.21	0.20	0.000	0.000	0.000
Mean w L	-0.16	-0.07	1.25	0.00	0.01	0.09	0.76	0.00	0.87	-0.29	-0.53	0.000	0.000	0.000
Var w H	0.60	0.24	-1.08	0.00	1.11	0.51	0.35	0.00	0.09	-0.88	1.41	0.000	0.000	0.000
Var w M	-0.20	0.58	-0.47	0.00	0.11	1.38	0.91	0.00	-2.77	0.00	-0.07	0.000	0.000	0.000
Var w L	-0.09	0.37	-0.07	0.00	0.20	0.28	1.65	0.00	-12.22	-0.39	. -0.49	0.000	0.000	. 0.000

Notes: The table reports the elasticity of each targeted moment to a 1% change around the calibrated parameter value.