

Labor Market Polarization and Inequality: A Roy Model Perspective*

Andrés Erosa[†]

Universidad Carlos III de Madrid

Gueorgui Kambourov[§]

University of Toronto

Luisa Fuster[‡]

Universidad Carlos III de Madrid

Richard Rogerson[¶]

Princeton University and NBER

October 2025

Abstract

We study the forces driving polarization and higher wage inequality since 1980 using a structural model of occupation choice in the tradition of Roy (1951). In our model, changes in relative occupational skill prices proxy for changes in relative demand for occupational labor services. Our analysis yields three main findings. First, although changes in skill prices have quantitatively important effects on employment shares and mean wages, they play essentially no role in accounting for the sharp rise in wage inequality. Second, changes in relative wages are driven by changes in higher order moments and do not reflect changes in relative demand. Third, changes in the variance of idiosyncratic within occupation productivity are the dominant factor behind the sharp rise in inequality.

JEL Classification: J22, J23, J24, J31.

Keywords: Inequality, occupational choice, polarization, Roy model, Frechet distribution.

*We thank seminar and conference participants at the Federal Reserve Bank of Minneapolis, IIES, Trinity College Dublin, Universidad Autónoma de Madrid, Università di Padova, Universitat de Barcelona, Universitat de les Illes Balears, University of Bristol, University of New Hampshire, Wharton, 2024 Armenian Economic Association Meeting, 2024 Barcelona Summer Forum, 2024 Society for Economic Dynamics Winter Meeting, and 2024 Workshop for Australasian Macroeconomic Studies. Erosa and Fuster acknowledge financial support by grants CEX2021-001181-M, and PID2023-151842NB-I00 funded by MICIU/AEI /10.13039/501100011033. 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 School for Policy and International Affairs, Princeton NJ 08544, USA. E-mail: rdr@Princeton.edu.

1 Introduction

During the last 15 years the framework that economists use to organize their thinking about inequality has shifted from the skill based framework of Katz and Murphy (1992) to the task based framework proposed by Acemoglu and Autor (2011). Central to this shift was an empirical finding documented by Acemoglu and Autor (2011): the polarization of US labor market outcomes along the dimensions of both wages and employment. Specifically, if one ranks occupations according to mean wages in 1980 on the x-axis, plots of employment changes and wage changes in the 1980-2007 period are both U-shaped.¹ This finding has been used to promote a simple and intuitive narrative about the causal forces shaping inequality. This narrative, which we will label as the *polarization narrative*, holds that the U-shaped pattern for mean wage changes reflects a U-shaped pattern for changes in relative demand for occupational services, with supply responses to these changes in relative demands generating the U-shaped pattern for employment changes. Much recent work on inequality fits within this narrative, including recent papers assessing the impact of automation, computers and international trade on the wage structure.²

Because the Roy (1951) model was specifically designed to study the connection between the occupational wage and employment structure and overall inequality, it is a natural framework for studying polarization and inequality. But viewed from the perspective of a standard Roy model, the link between the polarization evidence and the polarization narrative seems tenuous. In a Roy model, changes in relative demands for occupational services affect relative skill prices across occupations. But importantly, one of the key messages of the Roy model is that changes in mean wages need not reflect changes in skill prices.³ This is because changes in mean wages within an occupation can reflect changes in the shape of the distribution of wages for incumbent workers in that occupation, or entry and exit of workers in specific parts of the distribution. The goal of this paper is to assess the polarization of US labor market outcomes from the perspective of a simple Roy model.

Our analysis yields three main findings. First, in our model the changes in mean occupational wages and employment shares observed in the US data are not consistent with a story in which

¹We note that the phenomenon of employment polarization had been previously documented by Acemoglu (1999) for the US, Goos and Manning (2007) for the UK and Goos, Manning and Salomons (2009) for Europe.

²For an analysis of computers and international trade see, e.g., Burstein, Morales and Vogel (2019), for the effects of the China shock see, e.g., Basco, Liegey, Mestieri and Smagghue (2025) and for an analysis of automation see, e.g., Acemoglu and Restrepo (2022) as well as the earlier paper by Autor, Levy and Murnane (2003).

³Böhm, von Gaudecker and Schran (2024) use a Roy model to interpret changes in the occupational wage and employment structure in Germany between 1985 and 2010. They show that selection effects are quantitatively important, implying that there is a large gap between changes in mean wages and changes in skill prices.

changes in relative demands for occupational services are the dominant driving force behind changes in the occupational wage and employment structure. Second, while our analysis finds that there have been large changes in relative demands for occupational services, they are not well proxied by changes in mean occupational wages and they have contributed very little to overall inequality. Third, the dominant source of relative changes in mean occupational wages is an increase in the variance of within occupation wages.

We begin our analysis by documenting some basic facts in the US data on wages and employment at the occupational level. We start with a simple calculation showing that changes in mean wages and employment at the three digit occupational level are an important source of the overall increase in wage inequality in the US since 1980. This remains true if we focus on the smoothed changes documented in Acemoglu and Autor (2011) or Autor and Dorn (2013). If these changes in the occupational wage and employment structure are driven by changes in relative demand for occupational skills, it follows that changes in relative demand are an important source of increased inequality.

Next we document two additional properties of occupational wage distributions. First, we show that occupational wage distributions have large ranges of overlapping support, implying that the difference in mean wages across occupations may not be a good predictor of the wage change a given worker would receive if they were to choose to change occupations. A simple corollary is that one cannot assume that moving employment out of middle wage occupations will necessarily lead to greater wage dispersion. Second, we show that there are large changes in the variance of within occupation wages over time, and that high mean wage growth occupations tended to have larger increases in the ratio of mean to median wages, suggesting higher wage growth in the right tail of their wage distribution.

Motivated by these considerations, we develop a multi-occupation Roy model in order to study wage and employment polarization. Following Heckman and Sedlacek (1985) our model assumes log normal distributions for individual productivity, as well as idiosyncratic tastes for occupations. We calibrate our model to match occupational employment shares and detailed features of the within occupation wage distributions in 1980. For reasons of tractability, we assume that there are nine occupations.⁴ We then use this model to understand the forces that generate the sharp rise in wage inequality between 1980 and 2010.

⁴Heckman and Sedlacek (1985) only considered two sectors. Dix-Carneiro (2014) extends their framework to allow seven.

Our first exercise asks whether changes in skill prices alone, which capture changes in relative demand for occupational services, can account for the changes in employment shares and mean wages across occupations. Our answer to this question is no.⁵ Moreover, we find that the changes in skill prices that generate the best fit to the polarization facts generate almost none of the observed increase in wage inequality found in the data.

In our second exercise, we recalibrate all of the primitives of our model to match the employment and wage distributions in 2010. Our recalibrated model closely matches the occupational employment and wage distributions in 2010. We perform decompositions to quantify the role of specific forces in accounting for polarization and changes in inequality.

Four main findings emerge. First, changes in skill prices and productivity variances are the two quantitatively important changes between 1980 and 2010. That is, changes in occupation tastes and the correlation structure of occupational skills play little role. Second, changes in skill prices and productivity variances are both important factors in accounting for employment polarization. Third, changes in mean wages across occupations are driven by changes in productivity variances. In particular, mean wage changes due solely to changes in skill prices are negatively correlated with observed mean wage changes. Fourth, essentially all of the increase in inequality is due to changes in skill variances; if skill prices were the only change between 1980 and 2010 there would have been a modest decrease in inequality.

Notably, rising productivity variances generate substantial increases in the gap between mean and median wages within occupations. This result underscores a distinctive feature of the Roy model: the occupational selection mechanism censors low-productivity draws and concentrates employment among higher-productivity individuals. This property helps the model capture the critical role of sharp wage increases at the top percentiles of the highest-paying occupations for US inequality.

Because we study occupational wage and employment distributions that are thirty years apart, our analysis focuses on the long-run forces shaping polarization and inequality. Specifically, we examine how shifts in the relative demand for occupational services affect inequality across cohorts entering the labor market under different long-run conditions. This perspective necessarily abstracts from the short-run disruptions that occur when, for example, a mid-career assembly line worker loses a middle-income job and transitions into low-wage service employment. For analyses that emphasize

⁵A similar conclusion is offered by Mishel, Schmitt and Shierholz (2013) and Hunt and Nunn (2022), though their arguments are quite distinct from ours.

the short-run consequences of occupational or sectoral shifts, we refer readers to Kambourov and Manovskii (2009), Artuc, Chaudhuri and McLaren (2010), Dix-Carneiro (2014), as well as recent work by Huckfeldt (2022) and Carrillo-Tudela and Visschers (2023).

In addition to our substantive messages about occupations and inequality, our analysis also offers an important methodological message. Multi-occupation Roy models face a curse of dimensionality: as the number of occupations J increases, the number of parameters in the covariance matrix for occupation specific skills increases as J -squared. It is for this reason that we limit our quantitative analysis to a model with nine occupations. Recent work (see, e.g., Hsieh, Hurst, Jones and Klenow (2019)) has assumed skills are distributed according to independent Frechet distributions. While this assumption is very powerful in terms of tractability, we show that that it is ill-suited to the analysis of inequality.⁶

Our paper is related to two literatures: one that studies changes in wage inequality over time and the other that studies occupational patterns using Roy models. Each of these is too vast for us to summarize here, so we simply note a few papers that are particularly closely related.

Our analysis of changes in mean wages and employment patterns across occupations follows the important contributions of Acemoglu and Autor (2011) and Autor and Dorn (2013). Many papers have since engaged with these patterns to study inequality. Burstein et al. (2019) and Acemoglu and Restrepo (2022) both use occupational data on employment and mean wages by demographic group to understand how changes to primitives affect the wage structure across demographic groups. Comin, Danieli and Mestieri (2020) study the role of income growth in generating polarization. Differently from us, these papers do not explicitly address changes in within occupation wage distributions for the behavior of mean wages. Johnson and Keane (2013) develop a general equilibrium model to analyze the changes in the occupational wage and employment structure over the period 1968-1996. While they consider more dimensions of observable heterogeneity than we do, they implicitly assume that workers with the same observables will earn the same wage in a given occupation.

Our Roy model analysis can be viewed as an extension of Heckman and Sedlacek (1985) to a setting with many occupations. Bárány and Siegel (2018) use a Roy model with log normal productivity to study polarization since 1950. Their analysis focuses on the role of sectoral reallocation between agriculture, manufacturing and services and does not consider the role of changes in the

⁶The independent Frechet assumption has also been used extensively in the trade literature. Lind and Ramondo (2023) discuss the limitation of this assumption in the context of international trade and develop a tractable extension that allows for correlation.

distribution of skills. Caunedo, Jaume and Keller (2023) use a Roy model with Frechet shocks to study how capital embodied technical change affects employment and wages across occupations.

Grigsby (2022) also studies a many-occupation version of Heckman and Sedlacek (1985) to study the impact of changes in skill prices. Whereas we follow Heckman and Sedlacek (1985) and impose lognormal distributions, Grigsby (2022) allows for discrete productivity distributions with mass on a relatively small set of values and estimates them nonparametrically. A key difference between our analysis and that of Grigsby (2022) is that he studies the effect of fluctuations in skill prices over the business cycle on movements in aggregate employment and wages, whereas we focus on long run changes in inequality.

An outline of the paper follows. Section 2 documents key empirical patterns in wages and occupations, suggesting that an occupational perspective may offer valuable insights into changes in inequality over time. However, the evidence raises questions about how well the polarization narrative explains shifts in wage inequality. This discussion motivates our use of the Roy model. In Section 3, we introduce our multi-occupation extension of the Roy model and contrast its implications with those of the Roy-Frechet model. Section 4 outlines our calibration strategy and evaluates the model’s fit to data from 1980. Section 5 examines whether changes in relative demand alone can account for shifts in occupational wages, employment shares, and overall inequality. Section 6 presents our main decomposition exercises to analyze changes in occupational outcomes and inequality. Finally, Section 7 concludes.

2 An Occupational Perspective on Inequality

This section begins with a simple exercise to suggest that an occupational perspective may be useful in studying inequality. In particular, our calculation shows that the evolution of employment shares and mean wages across three digit occupations goes quite far in accounting for the overall rise in inequality since 1980.⁷ We connect this result with the literature on polarization by showing that the smoothed U-shape patterns of employment and wage changes documented by Autor and Dorn (2013) capture much of the action from our first exercise.

The remainder of this section documents two key features of within occupation wage distributions. First, there are large overlaps of wage distributions across three digit occupations. Specifically, even occupations with large mean wage differences have large regions of overlap in their wage

⁷Similar exercises were carried out by Goos and Manning (2007) for the UK and by Autor (2019) for the US.

distributions. Second, there have been large changes in the variance of within occupation wages, and in particular, occupations that experienced high mean wage growth between 1980 and 2010 tended to experience larger increases in the variance of log wages, driven by increases at the top of the distribution. These facts are an important source of caution for the polarization narrative as an explanation for increased inequality. We close by arguing that the evidence points toward a Roy model as the appropriate model for assessing the relationship between occupational changes and changes in inequality.

2.1 Data

Our analysis is based on the IPUMS-CPS files from the 1976-2015 Current Population Survey (CPS).⁸ 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.

In order to match individuals to specific occupations, we only use observations for individuals who 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 0.2% of the hourly wage distribution. We also drop observations that report business income or farm income. Our analysis focuses on male workers aged 25-61.

Our focus is on longer-run changes in inequality over the period 1976-2015. To increase sample sizes we pool cross-section observations over two ten-year intervals: 1976-1985 and 2006-2015. For ease of exposition, we will refer to the earlier pooled data as representing 1980 and to the later pooled data as representing 2010.

2.2 Occupational Changes and Inequality

For our benchmark sample of males aged 25-61, the cross-sectional variance of log wages increased from 0.30 in the 1976-1985 period to 0.49 in the 2006-2015 period, an increase of 0.19.

⁸The data and a detailed description can be found at <http://cps.ipums.org/cps/>. See Flood, King, Rodgers, Ruggles and Warren (2018).

We now carry out a simple calculation that supports the notion that an occupational perspective on inequality is valuable. Let $e_{j,t}$ be the employment share for the three digit occupation j in year t and let $w_{j,t}$ be the log of mean wages in three digit occupation j in year t .

Next, we compute the variance of log wages for the discrete distribution defined by the two vectors $w_{j,t}$ and $e_{j,t}$ for both the 1980 and 2010 cross-sections. The variance in 1980 is 0.067, and the variance in 2010 is 0.178. The change is 0.111 and represents 58 percent of the total increase of 0.19 in the variance of log wages over this time period.⁹ We conclude that understanding the evolution of employment shares and mean log wages across occupations would go quite far in helping us understand the large increase in wage inequality since 1980.

The 0.111 increase in the variance of log wages reflects changes in both employment shares and relative mean wages. If we hold the employment distribution fixed at its 1980 level and consider the mean wage distribution from 2010 we obtain a variance of log wages equal to 0.139. This represents roughly 65 percent of the overall change. When we hold the mean wage distribution fixed at its 1980 level and consider the employment distribution from the 2010 cross-section, we obtain a variance of log wages equal to 0.077. This represents roughly 9 percent of the overall change. There is also an interaction term. We conclude that changes in mean wages play a more important role in accounting for the sharp rise in inequality than changes in employment shares.

2.3 Polarization and Changes in Inequality

The previous calculation was based on the observed changes in occupational mean wages and employment shares. Acemoglu and Autor (2011) and Autor and Dorn (2013) showed that while there was much idiosyncratic variation in mean wages and employment levels across occupations, there was an underlying U-shaped pattern in which occupations in the middle part of the mean wage distribution experienced large declines in both employment shares and wages relative to both high and low mean wage occupations. These patterns reflect the phenomena of employment and wage polarization.

In this subsection, we assess the importance of this underlying polarization pattern in accounting for the change in inequality calculated in the previous subsection. That is, we repeat our previous calculation, but this time assuming that the 2010 distribution reflects the smoothed distribution of changes. Doing this, we find that the variance of log wages in 2010 is now 0.153 instead of 0.178 so


⁹Although the variance of this discrete distribution in 1980 is less than 25 percent of the overall variance of log wages, it accounts for almost 60 percent of changes in the overall variance of log wages.

occupational polarization accounts for 45 percent of the total increase in the variance of log wages between the 1980s and 2010s.

2.4 Issues With the Polarization Narrative

Two notions are central to the intuitive appeal of the simple polarization narrative. The first is that if we move workers from occupations with mean wages in the middle of the distribution to the two tails, then wage inequality will increase. The second is that movements of employment shares and relative mean wages in the same direction are evidence of changes in relative demand. While both of these notions have intuitive appeal, in this subsection, we present evidence to suggest that some caution is called for.

Figure 1: Within- and Between Occupation Wage Dispersion by 3 Digit Occupations, Men.




Notes: Each point represents a 3-digit occupation in the 1976-1985 (left panel) and 2006-2015 (right panel) time period. The scatter plots describe the relationship between the variance of log wages within an occupation and the occupational log mean wage. The size of the circle indicates the relative size of the occupation.

The first key piece of evidence that we present is that the variance of within occupation log wages is large relative to the variance of log mean occupational wages. The two panels of Figure 1 show the distribution of the within occupation variance of log wages as a function of log mean occupational wages in 1980 and 2010. Both panels show that the within-occupation variance of log wages is large relative to the variance of log mean occupational wages. In 1980, the mean log occupational wage was equal to 2.43, the variance of log mean occupational wages was 0.067, and the mean within-occupation variance of log wages was equal to 0.23. Considering the average within-occupation variance of 0.23, a two-standard deviation interval around the occupation with

the median log wage (which is 2.42) yields an interval of [1.46, 3.38] for log mean wages, which spans the entire support for log mean occupational wages. Comparing the two panels reveals that the mean within-occupation variance increased from 1980 to 2010, from 0.22 to 0.33, thereby strengthening this observation.

An important implication is that there are large overlaps in occupational wage distributions, even for occupations with very different mean wage levels. As a specific example of this general point, Figure 2 plots the densities of the within occupation wage distributions for the occupations at the 33rd and 67th percentiles.¹⁰ The figure illustrates the large area of overlap despite the large gap in mean wages.

Figure 2: Wage Distributions, 1980s, Occupations at the 33rd and 67th Percentiles.



Notes: This figure plots wage distributions for occupations at the 33rd and 67th percentiles using pooled data for the period 1976-1985. To increase the sample size, the 33rd percentile in this figure pools observations from the 32.5th to the 33.5th percentile and the 67th percentile pools observations from the 66.5th to the 67.5th percentile.

This evidence is directly relevant for assessing one piece of the logic associated with the polarization narrative. Specifically, if occupational wage distributions have large regions of overlap, it is no longer obvious that moving employment out of occupations in the middle of the mean wage distribution will lead to an increase in inequality.

A second issue is that because variances of within-occupation wage rates are large, changes in mean wage rates within an occupation may reflect changes in higher order moments rather than a shift in the overall distribution. Specifically, changes in either the right or left tail of the distribution

¹⁰To increase the sample size, the 33rd percentile in this figure pools observations from the 32.5th to the 33.5th percentile and the 67th percentile pools observations from the 66.5th to the 67.5th percentile.

might have important effects on the mean. This could happen either because the distribution of wages for a given set of workers changes or because of entry or exit of workers.

Figure 3: Within- and Between Occupation Dispersion in Wages: by 3d Occupations, Men.



Notes: Each point represents a 3-digit occupation. The scatter plots describe the relationship between the change in the variance of log wages within an occupation between the 1980s and 2010s and the log mean occupational wage in the 1980s (left panel) or the change in the log mean occupational wage between the 1980s and 2010 (right panel). 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 offers a first look at this issue. The left panel shows the change in the within-occupation variance of log wages between 1980 and 2010 as a function of mean wages in 1980. Notably, this panel shows that the within-occupation variance of log wages has increased substantially in many occupations. The right panel is a scatter plot of the change in the within occupation variance against the *change* in the occupation mean wage. The key feature of this figure is that occupations with higher growth in mean wages also tend to have larger increases in variance.

To pursue this point further, Figure 4 plots the change in the ratio of mean to median wages as a function of the mean occupation wage in 1980. It shows that high mean wage occupations tended to have large increases in the ratio of mean to median wages, suggesting that changes in the right tail of the distribution are playing an important role.

2.5 Summary

These facts about within occupation wage distributions lead us to conclude that the appropriate framework for assessing the connection between polarization, changes in relative demand for occupational services and inequality must address within occupation wage dispersion in addition to

Figure 4: Change in the Mean-to-Median Wage Ratios, 1980s to 2010s.



Notes: Each point represents a 3-digit occupation. The scatter plot describes the relationship between the change in the mean-to-median ratio in an occupation between the 1980s and 2010s and the log mean occupational wage in the 1980s. 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.

mean wage differences across occupations. The standard model for thinking about occupational wage distributions and occupational choice is the Roy model. In the rest of this paper we develop a many occupation version of the Roy model and use it to analyze the connection between changes in the occupational employment and wage structure and changes in inequality.

3 A Model of Occupations and Wage Inequality

In this section, we present a static version of a Roy (1951) model that we will use to investigate the connection between polarization and inequality. The first subsection introduces our baseline model, which features log normal distributions for individual heterogeneity. In the second subsection we consider a more parsimonious model, that we label as the Roy Frechet model. This model has recently been used by several researchers, but we argue that it is ill-suited to study the role of occupations for understanding (changes in) inequality.

3.1 Baseline Model

Our baseline model focuses on the labor supply decisions of a unit mass of workers indexed by i , that choose between J occupations denoted by j . As in the standard textbook Roy Model, each individual i is endowed with a vector $\vec{\epsilon} = (\epsilon_{i,1}, \epsilon_{i,2}, \dots, \epsilon_{i,J})$ that reflects their productivity in each

of the J occupations, along with a unit of time, and does not value leisure.¹¹ The distribution of idiosyncratic productivity in the population follows a lognormal distribution:

$$\ln(\epsilon_{i,1}, \epsilon_{i,2}, \dots, \epsilon_{i,J}) = \ln(\vec{\epsilon}) \sim N(\mu_\epsilon, \Sigma_\epsilon). \quad (1)$$

Each occupation j is characterized by a skill price, which we denote by w_j , and individuals make occupational choices taking skill prices w_j as given. The literature considers many different factors that might affect w_j , including automation and other forms of technical change, international trade, and income. For analyzing the pure role of changes in relative demand for occupational services, it is the relative change in skill prices that matters and not the source of the changes. For this reason we work in a partial equilibrium setting and do not model the underlying factors that influence skill prices.

If individual i chooses to work in occupation j they will receive earnings of $\epsilon_{i,j}w_j$. They use their earnings to purchase a final good, whose price is normalized to unity.

Differently than in standard textbook versions of the Roy model, but following Heckman and Sedlacek (1985), we assume idiosyncratic tastes for working in each occupation. In particular, if individual i works in occupation j , their effective consumption is given by actual consumption of the final good multiplied by $v_{i,j}$, where $v_{i,j}$ is independent across occupations and individuals and is drawn from a lognormal distribution: $\ln(v_{i,j}) \sim N(\mu_{v,j}, \sigma_{v,j}^2)$. As we show later, this feature is important in allowing the model to fit employment and wage moments.

An individual knows productivity and taste shock realizations as well as all prices prior to choosing an occupation. If individual i chooses to work in occupation j they receive utility $u(\epsilon_{i,j}w_jv_{i,j})$ where u is a strictly increasing function. It follows that individual i chooses the occupation j that maximizes:

$$V_{ij} = \epsilon_{i,j}w_jv_{i,j}$$

More formally, an individual decides to work in occupation j if

$$\ln V_{ij} = \ln \epsilon_{ij} + \ln w_j + \ln v_{i,j} \geq \ln V_{i,k} \text{ for all } k \neq j.$$

For future reference we close this section with a comment about interpreting the individual

¹¹See Erosa, Fuster, Kambourov and Rogerson (2022, 2024) for extensions of the Roy model that incorporate hours choices.

productivity differences in our model. Productivity differences imply differences in output, and the key point is that these output differences should be interpreted as reflecting a combination of differences in skill and the return to skill. To illustrate this point we consider the specification of production introduced by Heathcote, Perri and Violante (2020). Letting s denote a worker’s skill and γ parameterize technology, they assume that the output of a worker with skill s using technology γ is given by $\exp(\gamma s)$, so that the log of output for this worker will be given by γs . In this specification, the parameter γ influences the return to skill and so captures the possibility that a change in technology influences the return to skill. If s follows a log normal distribution then it follows that a change in the variance of individual output can be induced either by a change in technology, i.e., a change in γ , or a change in the variance of individual skills. Later in this paper we will consider changes in the variance of occupation specific productivity, and in view of this discussion it should be understood that these changes in variance may reflect both changes in worker skills and changes in the return to skill driven by changes in technology. The recent paper by Lochner, Park and Shin (2025) proposes a method for distinguishing between these two components.

3.2 The Roy Frechet Model

Our model economy essentially extends Heckman and Sedlacek (1985) to more than two occupations. A key difficulty in connecting a Roy model with many occupations to the data is that it brings a curse of dimensionality: The number of parameters grows at a quadratic rate with the number of occupations.¹²

A standard approach that has been adopted in much of the recent literature is to model occupation specific productivities as IID draws from a Frechet distribution. (See, e.g., Hsieh et al. (2019).)¹³ This reduces the number of parameters drastically and also allows for convenient analytical characterizations.¹⁴ In what follows we will refer to this as the Roy Frechet model. Despite its many advantages, we argue below that a key limitation of the Roy Frechet model in our context is that it cannot account for the key facts about occupational wage inequality.

The details of the Roy Frechet model are as follows. Productivity draws are iid across individuals and occupations and follow a Frechet distribution with shape parameter θ , scale parameter 1,

¹²Specifically, the variance-covariance matrix of skills has $J \times (J - 1)$ parameters.

¹³Lagakos and Waugh (2013) is an exception but they have only 2 occupations.

¹⁴The maximum among independent Frechet random variables with a common shape follows a Frechet distribution, making analytical derivations easy.

and location parameter 0. There is no idiosyncratic component to preferences, but we allow for occupation specific utility, denoted by v_j , that is common to all individuals. In what follows, we will refer to the v_j as occupation amenities. An individual who obtains consumption c and works in occupation j receives utility $u(cv_j)$ where u is an increasing function.

An individual will choose an occupation that yields the highest value of $V_{i,j}$, where $V_{i,j}$ is given by:

$$V_{i,j} = w_j v_j \epsilon_{i,j} = \tilde{w}_j \epsilon_{i,j}, \quad (2)$$

where $\tilde{w}_j = w_j v_j$ is the effective reward per efficiency unit in occupation j .

Proposition 1 characterizes optimal occupation choices and wage distributions in the economy. Proof of this proposition is found in Appendix A.

Proposition 1 *Optimal occupational choices in the Roy Frechet model lead to allocations with the following properties:*

1. (Employment shares.) *The fraction of individuals choosing occupation j is given by $e_j = \frac{\tilde{w}_j^\theta}{\sum_{j=1}^N \tilde{w}_j^\theta}$.*
2. (Wage distributions.) *The wage distribution in occupation j follows a Frechet distribution with scale parameter $\bar{w}_j$*

$$F_j(w) = P(w_j \epsilon_j < w \mid j \text{ chosen}) = \exp \left[- \left(\frac{w}{\bar{w}_j} \right)^{-\theta} \right], \text{ where } \bar{w}_j = \frac{1}{v_j} \left[\sum_{s=1}^J (\tilde{w}_s)^\theta \right]^{1/\theta}.$$

3. (Inequality.) *If there is no heterogeneity in occupational amenities, the wage distribution does not vary across occupations. If there is heterogeneity in occupational amenities, the ratio of mean wages between two occupations is inversely related to the amenity ratio. Moreover, the distribution of wages do not vary across occupations when normalized by the mean wage in each occupation.*

Proposition 1 has important implications regarding the usefulness of the Roy Frechet model for studying inequality. If there is no heterogeneity in occupational amenities, wages would be identically distributed across all occupations. This would in turn imply that the distribution of wages in each occupation is identical to the aggregate wage distribution. In particular, mean wages

would be identical across all occupations and there would be no hierarchy of occupations with regard to mean wages. Both of these properties are strongly contradicted by the data.

Allowing for heterogeneity in amenities allows the model to generate heterogeneity in wage distributions across occupations, and in particular, would generate a hierarchy across occupations in terms of mean wages. But, although occupational wage distributions will differ, all occupations would have the same distribution when each occupational wage distribution is normalized by its mean or median. One implication of this is that the variance of log wages should be identical across occupations, a property that is also strongly contradicted by the data.

In summary, Proposition 1 implies that the Roy Frechet model fails to capture the most basic facts about inequality and occupations. Having established this we next argue that this model is also not well suited to assess the essence of the polarization narrative. The key feature that characterizes polarization is that relative employment and relative wages move in the same direction across occupations. The standard narrative in this literature is that changes in skill prices due to changes in demand are the source of these movements. But Proposition 1 states that while changes in relative skill prices and changes in relative employment are positively correlated across occupations, they will have no effect on relative mean wages. That is, changes in relative mean wages across occupations are not a measure of changes in relative demand for occupational services.

Proposition 1 implies that the only way to affect relative mean wages across occupation is via changes in relative amenity values across occupations. Because relative wages and relative amenity values are negatively correlated, the mean wage in a given occupation will increase only if the amenity value decreases. But this will in turn lead to a decrease in relative employment.

Generating a positive relationship between changes in relative employment and wages across occupations could be achieved with an appropriate combination of changes in skill prices and amenity values, but would require a systematic negative correlation between changes in skill prices and amenity values. While logically possible, we do not view this as a compelling explanation.

We conclude that despite the tractability offered by the Frechet Roy model, it is ill-suited as a framework for understanding the connection between occupational changes and changes in inequality.

4 Calibration

In this section we describe our calibration procedure, report the calibrated parameter values, and assess the fit of the calibrated model.

4.1 Number of Occupations

Our choice for the number of occupations, J , reflects a trade-off between the desire to have a rich occupational structure and feasibility, given the curse of dimensionality mentioned earlier. This leads us to set $J = 9$. This choice requires that we aggregate the data at the 3 digit level to 9 occupations. To do this we rank the 3 digit occupations by mean wage for our 1980 cross-section and then create 9 occupation bins, each of which contains the same number of 3 digit occupations.

It is useful to confirm that the properties that we documented in Section 2 at the 3 digit level continue to hold for our 9 occupation specification. If we focus on the discrete distribution of employment shares and mean occupation wages, we find that the change in the implied variance of log wages increases from 0.063 in 1980 to 0.142 in 2010, an increase of 0.079. This corresponds to roughly 42 percent of the overall change of 0.19 in the variance of log wages over this time period. Not surprisingly, the 9 occupation specification implies both a smaller overall variance and a smaller increase in variance of log wages relative to the case in which we use the 3 digit disaggregation. But importantly, it remains the case that changes in mean occupational wages and employment shares are an important driver of the overall increase in the variance of log wages.

Table 1: Changes in Occupational Wage and Employment Structure: 1976-1985 to 2006-2015.

Description	Oc1	Oc2	Oc3	Oc4	Oc5	Oc6	Oc7	Oc8	Oc9
Δ Emp. Shares	0.033	0.009	0.019	-0.029	-0.014	0.004	-0.015	-0.013	0.005
$\Delta\%$ Mean Wage	-4.4	0.7	-2.3	-13.9	-8.1	4.2	20.8	12.7	32.3
Δ Var. Log Wage	-0.012	0.132	0.103	0.067	0.114	0.145	0.246	0.144	0.208
Δ Mean to Median Wage	0.073	0.182	0.155	0.104	0.109	0.170	0.277	0.129	0.237

Notes: The table reports changes in employment shares, mean wages, and variance of log wages for 9 occupations in the US economies from 1976-1985 to 2006-15.

Table 1 displays the changes in mean occupational wages and employment shares between 1980

and 2010. It reveals a general U-shaped pattern for both the changes in employment shares and the changes in mean wages. In what follows we will refer to these changes in employment shares and mean wages as the polarization facts. The table also shows the change in the variance of within occupation log wages and the change in the ratio of mean to median wages. There are significant increases in the variance of log wages for almost all occupations, especially for the top ones. There are also large increases in the ratio of mean to median wages, and the correlation between the change in mean wages and the change in the mean to median ratio is 0.80. Notably, there is a sizeable increase in the ratio of mean to median wages in all occupations, including those occupations that experienced mean wage decreases.

4.2 Calibration Procedure

Under our parametric specifications and given our choice of $N = 9$, the model has 72 parameters: 45 parameters of the variance-covariance matrix for idiosyncratic productivity, 18 parameters for the distribution of occupational tastes (9 mean values and 9 variances), and 9 skill prices (wage rates). Recall that we normalize the price the final good to unity. Because the optimal occupational choice only depends on utility differences across occupations, we fix the mean value of tastes in occupation 5 to zero without any loss of generality. This leaves us with 71 parameters.

We choose parameter values by targeting 80 moments from the 1980 data aggregated to nine occupations: i) 7 moments of the wage distribution for each occupation (wage percentiles 5, 10, 25, 50, 75, 90, and 95); ii) 9 mean occupational wages, and iii) 8 employment shares.¹⁵ The calibration is performed by minimizing a loss function of the weighted sum of squared deviations between model statistics and targets.¹⁶

4.3 Calibration Results: Model Fit and Parameter Values

Table 2 presents the parameter values that are delivered by our calibration procedure and shows the ability of the calibrated model to match the targets for mean occupational wages and employment shares.

We begin by discussing the ability of the model to match the features of the data. The calibrated

¹⁵We do not target either the 1st or 99th percentiles. Acemoglu and Autor (2011) argue that one should not consider the two tails of the wage distribution when using CPS data on account of data issues. An additional consideration is that it is well known that one cannot match the upper tail of the wage distribution with a log normal distribution.

¹⁶The weights on moments involving wages are set to 1, and the weights on moments involving employment shares are set to 100. Deviations between wages in the model and data are expressed as percent deviations.

Table 2: Calibration Results.

Parameters					Calibration Results			
Occupation	w_j	$\sigma_{\epsilon_j}^2$	$\mu_{v,j}$	$\sigma_{v_j}^2$	Emp. Shares		Mean Wages	
					Data	Model	Data	Model
1	0.975	0.412	0.403	0.315	0.073	0.075	6.54	6.41
2	0.973	0.423	0.151	0.221	0.031	0.031	8.03	7.99
3	1.414	0.352	0.222	0.202	0.099	0.099	9.06	8.82
4	1.844	0.228	0.173	0.172	0.171	0.171	9.90	9.66
5	1.882	0.224	0	0.121	0.078	0.078	10.75	10.61
6	1.873	0.242	0.180	0.037	0.103	0.103	11.37	11.22
7	2.246	0.168	-0.233	0.177	0.132	0.132	12.49	12.24
8	2.145	0.306	-0.033	0.167	0.229	0.229	14.40	14.36
9	1.889	0.385	-0.209	0.056	0.082	0.081	17.27	17.18


Notes: The table reports parameter and calibration results for the baseline model calibrated to US data from 1976-1985. The first column reports the log of skill prices, the second column the variance of log skills, the third column the mean of the log of occupational tastes, and the fourth column the variance of log occupational tastes. The rest of the columns report the fit of the calibration targets regarding employment shares and mean wages by occupation.

model closely matches the employment shares, with the largest discrepancy between model and data being two-tenths of a percentage point. It also does a reasonable job of capturing the distribution of mean wages, though the fit is not as close for mean wages as it is for employment shares. As we show below, our model does a good job of matching the targeted percentiles that range from the 5th to the 95th percentile. But with lognormal distributions of productivity there is some tension between matching the percentiles from the 5th through the 95th and also matching the mean. Nonetheless, the deviations between model and data are relatively small. Mean wages are slightly lower in the model than the data for all occupations, but the discrepancy never exceeds three percent. In particular, our calibrated model replicates the hierarchical structure of mean wages across occupations and does a good job of matching the magnitude of mean wage differences across steps in the occupational hierarchy.

Figure 5 shows the ability of the model to match the within occupation wage distributions. The left panel shows the percentiles of the wage distribution for occupations 1 to 5, while the right panel does the same for occupations 5 to 9. We have used the same scale in both panels and report statistics for occupation 5 in both panels to facilitate comparisons across the two panels. The targeted percentiles are indicated with a circle. As noted above, our calibrated model matches these

targets quite closely. Importantly, the calibrated model captures the fact that wage distributions across occupations have large regions of overlap.

Figure 5: Fit of the Calibrated Model: Wage Percentiles.



Notes: The figure displays the occupational wage distributions in the model and the data targets from 1975-1985. The left panel shows the percentile of the wage distributions for occupations 1 to 5, while the right panel displays the wage distributions for occupations 5 to 9.

Next, we examine the parameter values that generate the close fit to the data. Table 2 shows the implied values for the skill prices, productivity variances, and mean and variance of the occupational taste shocks.

Although there is a general tendency for skill prices to increase as we move up the occupational hierarchy, they peak at occupation 7 and decrease quite substantially as we move to occupations 8 and 9. This property highlights the fact that in a Roy model, skill prices and mean wages need not have the same properties, a result that will figure prominently in our analysis later in the paper.

The variance of productivity draws varies quite substantially across occupations and exhibits a U-shaped pattern. Notably, there is not a strong connection between the observed within occupation variance of log wages and the unconditional variance of productivity shocks within an occupation.

Table B-1 in Appendix B shows the pattern of correlation for productivity draws across occu-

pations. All but two of the correlations are positive (ρ_{21} and ρ_{26}). Although the correlations are all positive, they tend to be somewhat modest in magnitude. Only 5 of the 36 independent values of the correlation matrix exceed 0.50, with the highest value being 0.69, and 17 of the values are between 0 and 0.20.

The calibration implies substantial heterogeneity in mean occupational tastes, ranging from 0.40 in occupation 1 to -0.21 in occupation 9. Mean taste values tend to decrease with the occupational skill price, a force that helps shift employment towards occupations with low wages. Conditional on other parameter values (e.g., the variance-covariance matrix of skills), the variation in tastes helps the model match the substantial within-occupational variation in wages while also matching the employment shares. This variation in tastes proves particularly important for matching the high wage variation in occupations at the bottom of the hierarchy.

The negative correlation between mean occupational tastes and mean occupational wages may seem counterintuitive if one thinks in terms of non-wage benefits. Our preferred interpretation of this feature of our calibration is that occupation tastes reflect the disutility associated with the effort required to be employed in a particular occupation. With this interpretation, our results indicate that high mean wage occupations tend to be more demanding. However, as a practical matter we show next that this feature of the calibration does not play a quantitatively important role.

To understand the role of preference heterogeneity we consider three restricted versions of our model. In the first version, we eliminate heterogeneity in preferences for occupations across individuals by setting $\sigma_{\nu_j} = 0$ for all j . In this specification, all individuals have the same tastes for occupations. In the second version, we allow for heterogeneity in tastes across individuals but eliminate mean differences in tastes across occupations by setting $\mu_{\nu_j} = 0$ for all j . In the third version, we eliminate all variation in preferences across occupations. Table 3 shows how the fit of the model varies across these restricted models.

Our baseline calibration (Column 1) features a value for the loss function of 0.132. When we eliminate the common component of preferences (Column 2) the loss function increases to 0.210, while when we eliminate individual heterogeneity (Column 3) the loss function increases to 0.329. When we eliminate all preference heterogeneity (Column 4) the loss function balloons to 1.632. These results show that the model requires at least one source of preference heterogeneity in order to get a reasonable match to both employment shares and within occupation wage distributions. We interpret these results to indicate that allowing for heterogeneity in mean tastes across occupations

Table 3: Loss Function in Alternative Models.

Model Spec.	(1)	(2)	(3)	(4)
Occ. Tastes	yes	no	yes	no
Het Tastes	yes	yes	no	no
Contribution to loss function				
Emp. Shares ($\times 100$)	0.00	0.02	0.00	0.87
Mean Wages	0.003	0.003	0.005	0.017
Wage Percentiles	0.128	0.191	0.323	0.741
Total	0.131	0.210	0.329	1.632

Notes: The table reports the loss function for three model economies calibrated to the same targets. The first column corresponds to the baseline model. The second column is the model with idiosyncratic occupational taste shocks that all have zero mean. The third column allows assumes no idiosyncratic component to occupational tastes but does allow for common mean effects. The fourth column allows for not occupational tastes. For each model, we report the contribution to the loss function by three categories of targets: i) employment shares (times 100), ii) mean occupational wages, and iii) occupational wage distributions.

is not critical in allowing the model to fit the targeted moments.

4.4 Discussion

Our goal is to understand the forces that shape long-run changes in inequality. In particular, we do not view our analysis as focusing on the short-run effects on inequality associated with mid-career workers who are displaced from middle-wage occupations and need to find employment elsewhere.¹⁷ Rather, we view our model as capturing the choices made by different cohorts of young workers that enter the labor market under changing economic conditions, as captured by the changes in skill prices.

This interpretation raises the issue that our data analysis focuses on cross-sectional inequality for workers aged 25-61, which need not accurately reflect the ex-ante options faced by young individuals making occupational choices. To address this issue we have also carried out an alternative calibration for which we constructed data moments for young (25-35) male workers. All of our main findings are robust to this alternative calibration strategy. Because the results are so similar,

¹⁷As noted in the introduction, both Dix-Carneiro (2014) and Grigsby (2022) focus on these effects.

all of the detailed results for this alternative calibration are contained in Appendix C.¹⁸

A second issue is that our calibration procedure relies entirely on cross-section data. As shown in Heckman and Honoré (1990), the Roy model cannot be identified non-parametrically from a single cross-section dataset on occupational choice and wage rates. Simply put, observed wages will pin down the product of the skill price and occupational productivity in the chosen occupation, and there is an infinite set of choices for any individual’s productivity in other occupations that is consistent with the individual choosing their observed occupation. Following Heckman and Sedlacek (1985) and other early contributions based on the Roy model, we deal with this issue by imposing that heterogeneity in productivity and tastes are described by lognormal distributions.

Consistent with our previously mentioned focus on long-run inequality, the key information that we require is the covariance of productivity across occupations for young workers from the perspective of making career choices at the time of labor market entry. Wage changes for mid-career occupational choices of displaced workers may not be informative about this. Given the challenges associated with obtaining non-parametric estimates of the appropriate covariance matrix, we feel that imposing lognormal distributions provides a useful first step in providing answers to the questions that we pose.

5 Skill Prices, Polarization and Changes in Inequality

In this section, we use our calibrated model to evaluate an extreme form of the polarization narrative. Specifically, we consider the case in which the only exogenous changes in the economy are changes in skill prices, which we take as proxies for changes in relative demand for occupational services. The question we ask is to what extent this narrative can account for the observed changes in the occupational wage and employment structure? Notably, these are the facts that have been used to motivate the polarization narrative.

To answer this question, we find the changes in skill prices that best match the polarization facts, that is, that provide the best fit to the changes in employment shares and mean occupation wages between 1980 and 2010. To implement this exercise we minimize a loss function which sums the squared deviations of changes in log mean occupation wages and employment shares expressed as percentages. Having chosen changes in skill prices to provide the best match to the polarization

¹⁸As we discuss in the Conclusion, it is of interest to explicitly consider a dynamic life cycle model and to estimate it using panel data. The small sample size of publicly available panel data sets in the US impose a severe constraint on this type of analysis.

facts, we ask two questions. First, does this generate a close fit to the polarization facts? And second, will it generate a large increase in inequality?

Before presenting the results of this exercise we note that if we had adopted the Roy Frechet model described in Section 3.2, there is no possibility that changes in skill prices alone could match the changes in mean occupational wages. Proposition 1 implies that mean occupational wages all change by the same amount, so that the Roy Frechet model is incapable of generating a U-shaped pattern of mean wage changes in response to changes in skill prices. In contrast, our baseline model allows for the possibility that changes in skill prices generate U-shaped responses in both employment shares and mean occupational wages.

Table 4: Pure Effect of Changing Skill Prices: 1980 to 2010.

Description	Oc1	Oc2	Oc3	Oc4	Oc5	Oc6	Oc7	Oc8	Oc9
Panel A: Data									
Δ Occ. Emp. Shares	0.031	0.009	0.019	-0.028	-0.014	0.004	-0.016	-0.013	0.006
% Δ Mean Wage	-4.4	0.7	-2.3	-13.9	-8.1	4.2	20.8	12.7	32.3
Δ Var Log Wage	-0.012	0.132	0.103	0.067	0.114	0.145	0.246	0.144	0.208
Δ Mean to Median Wage	0.073	0.182	0.155	0.104	0.109	0.170	0.277	0.129	0.237
Panel B: Minimizing Loss Function Over Emp. Share and Mean Wage Changes									
% Δ Skill Price	15.3	8.0	9.2	-4.5	-3.5	2.0	-1.3	-1.1	5.5
Δ Occ. Emp. Shares	0.025	0.007	0.019	-0.030	-0.014	0.005	-0.010	-0.012	0.010
% Δ Mean Wage	10.4	4.8	6.7	-2.9	-2.1	1.9	-0.6	0.2	3.2
Δ Var Log Wage	-0.0018	0.0010	0.0038	-0.0022	-0.0040	0.0008	-0.0006	-0.0027	0.0007
Δ Mean to Median Wage	-0.0019	-0.0012	0.0009	-0.0007	0.0044	-0.0013	-0.0016	-0.0039	0.0048

Notes: The table reports results from an experiment that changes skill prices in the baseline economy (calibrated to data from 1976-1985) to match changes in the occupational employment shares and mean wages between 1976-1985 and 2006-2015.

Results are shown in Table 4. The top panel displays changes in the data, and the second panel shows the results for the exercise that we just described. The first row in Panel B shows the changes in skill prices that deliver the best fit. Interestingly, the required changes in skill prices are U-shaped, though we note that the change in the bottom occupations is much larger than the change in the top occupations.

The second row of Panel B shows that the model is able to match the changes in employment

shares quite well when changes in skill prices are the only exogenous shift. The third row of Panel B shows the implied changes in mean occupational wages. The pattern for mean wage changes in the model exhibits a U-shaped pattern similar to the one for changes in skill prices, but with a shallower shape. The shallower shape reflects the fact that in our calibrated model, selection effects tend to dampen the direct effect of skill price changes.

The key finding is the large discrepancy between the results for mean wage changes in Panel B and the corresponding values from the data reported in Panel A. We conclude that there must have been some additional changes in fundamentals between 1980 and 2010 beyond changes in relative skill prices.

An important clue as to what this exercise might be missing is contained in the final two rows of Panel B. The fourth row reports the change in the variance of within occupation log wages and the fifth row reports the change in the mean to median wage. Comparing these outcomes with the values from the data in the top panel, we see that changes in skill prices generate changes in within-occupation variances that are often too small by more than an order of magnitude and changes in the ratio of mean to median wages that are in many cases too small by two orders of magnitude.

Next we turn to our second question: do changes in skill prices have an important effect on inequality? The fact that changes in skill prices are not the only quantitatively important driving force does not necessarily imply that changes in skill prices have a small effect on inequality. Nonetheless, it turns out that this is indeed the case: the increase in the variance of log wages in this exercise is 0.0042, which is roughly two percent of the overall change of 0.19 observed in our data set over this time period.

One of the questions that we have addressed in this section is whether changes in skill prices alone are able to account for the observed changes in employment shares and mean occupational wages. In the exercise just considered, we found the changes in skill prices that provided the best fit to the changes in both employment shares and mean wages. In Table B-2 in Appendix B we report the results from two other versions of this exercise: one where we choose changes in skill prices to only target employment share changes and another where we choose skill price changes to only target changes in log mean wages.¹⁹

The results for the case where we only target employment share changes look very similar to

¹⁹Because occupational shares necessarily sum to one, the first exercise only provides 8 targets for the 9 changes in skill prices. We eliminate the final degree of freedom by also targeting the change in the average wage in the US between the two periods.

the results for the exercise presented Table 4. This is not too surprising given that we previously found a reasonably close match to the employment share changes. The results for the case in which we target mean wage changes are quite different. In contrast to the results Table 4, this exercise closely matches the changes in log mean wages, but it dramatically misses the changes in employment shares. For example, in this exercise, occupation 9 experiences an increase of 18 percentage points, whereas in the data this increase was less than one percent. And occupations 4 and 5 experience losses that are more than three times larger than those found in the data. One implication of this result is that changes in mean wages could only be driven by changes in skill prices if there were changes in other primitives that had extremely large effects on employment shares without affecting wages.

In summary, we conclude that viewed through the lens of our model, a pure polarization narrative in which changes in relative demand result in changing skill prices is not able to reconcile the observed changes in employment shares and mean occupational wages. Moreover, if we infer changes in skill prices from choosing the best fit to the polarization facts, these changes in skill prices account for a very small component of the overall increase in wage inequality.²⁰

6 Sources of Inequality

The results in the previous section imply that changes in relative demands alone that manifest themselves as changes in relative skill prices are not sufficient to account for either the polarization facts or the changes in inequality observed in the data. In this section, we use our model to infer the changes in all of our model primitives that provide the best match to the employment and wage distributions in 2010. We then carry out counterfactuals to assess the relative importance of different driving forces.

6.1 Calibrating to Match the 2010 Data

As a first step we recalibrate all of the parameters of our model using the same procedure as previously, but now using data from 2010 instead of 1980. Details of the new calibration and its ability to fit the data are contained in Appendix B. Here we simply note that the model is again able to provide a good fit to the data. The loss function takes on a value of 0.109, which is

²⁰This property also holds for the exercise in which we only target changes in employment shares. When we only target changes in mean wages the exercise does generate a larger effect on inequality, as the variance of log wages increases by 0.069. However, this effect is driven by the very large expansion of occupation 9, an effect which is highly counterfactual.

comparable to the loss function when we calibrated the model using data from 1980. (Refer to the calibration results in Table B-3 and Figure B-1 in Appendix B).

When reporting the change in primitives between 1980 and 2010 there is one degree of freedom that requires a normalization. The reason is that changes in mean occupational productivity are isomorphic to changes in skill prices. In our calibration we normalized the mean of log productivity draws to zero for all nine occupations. But changes in the variance of log productivity will affect the mean level of productivity even if one holds the mean of log productivity equal to zero. In what follows we will keep mean occupational productivity constant by considering mean preserving changes in log variances.

6.2 Changes in Model Primitives

We begin by highlighting the key changes in primitives across the two calibrations. Table 5 displays the results.

Table 5: Change in Calibrated Parameter Values Using 2006-2015 Data.

Parameter	Oc1	Oc2	Oc3	Oc4	Oc5	Oc6	Oc7	Oc8	Oc9
Δ Var Log Prod.	0.001	0.205	0.161	0.115	0.127	0.158	0.247	0.270	0.634
Δ Prices (%)	3.6	-19.1	-12.5	-22.7	-18.0	-9.0	-18.4	-12.2	-20.8
Δ Mean Log Tastes	-0.008	0.029	-0.022	0.017	0	-0.005	-0.008	0.010	0.009
Δ Var Log Tastes	-0.073	0.005	0.034	0.021	-0.002	-0.004	-0.036	-0.042	-0.002

Notes: The table reports the change in the variance of log skills, skill prices, and the mean and variance of occupational tastes when the model is calibrated to data from 2006-2015 instead of 1976-1985. Changes in the variances of log productivities are accompanied by changes in the mean of log productivity that leave the mean level of productivity unchanged.

Perhaps not surprising given the results from the previous section, the first row indicates that all of the occupations except for Occupation 1 exhibit very large increases in the variance of idiosyncratic productivity. Recalling that the mean of the variances in the earlier calibration was roughly 0.17, the change in variances imply more than a doubling on average. This change is particularly high for occupation 9, which shows an increase of 0.634, more than twice as large as the second-highest value.

The second row shows that our calibration generates large but uneven changes in skill prices, all of which are negative with the exception of occupation 1. Notably, while occupation 9 experiences the largest increase in mean occupational wages, it also experiences one of the largest decrease in its skill price (-20.8%). Only occupation 4 has a larger decrease.


The third and fourth rows report the changes in the parameters characterizing the distribution of occupational tastes. The main message is that these values are relatively stable across the two calibrations. Tables B-4 and B-5 in Appendix B show the new matrix of correlations for productivity shocks as well as the difference between the values for 1980 and 2010. As we show below, changes in the matrix of correlations and parameters characterizing preference heterogeneity are of second-order importance.

We are particularly interested in the ability of our model to capture three sets of changes between 1980 and 2010: the overall change in inequality, as measured by the variance of log wages, the polarization facts, as captured by the changes in employment shares and mean occupational wages, and the change in within occupation wage distributions. Our calibration procedure explicitly targets employment shares and mean wages but does not explicitly target the overall variance of log wages. In the data, the variance of log wages increases from 0.30 in 1980 to 0.49 in 2010, an increase of 0.19. In our calibrated model the variance of log wages increases from 0.27 in 1980 to 0.44 in 2010, an increase of 0.17. In both 1980 and 2010 our calibrated model slightly underpredicts the variance of log wages. The reason for this is that our model does not fully account for the two extreme ends of the wage distribution. Nonetheless, our calibrated model does capture roughly 90 percent of the change in the variance of log wages. As noted earlier, our model matches the percentiles of within-occupation wage distributions in both 1980 and 2010, and as a result also matches the changes in within-occupation wage distributions.

Figure 6 shows the extent to which our model captures the polarization facts. The left panel of this figure shows a scatter plot of changes in employment shares in the data versus the model, along with a 45 degree line. Our calibration procedure explicitly targets employment shares and hits these targets quite closely, so as a result the points in this figure closely follow the 45 degree line.

The right panel of Figure 6 shows the analogous plot for changes in occupational mean wages. While the points do follow the 45 degree line, there is more dispersion around the 45 degree line than was the case for changes in employment shares. To understand the reason for this it is useful to recall our earlier discussion of the model when calibrated to data for 1980. There we noted

Figure 6: Fit of the Polarization Facts: Theory and Data.



Notes: The figure compares the changes in employment shares and mean occupational wages in the model and data between 1970s and 2010s. The model changes are computed by comparing stats from the 1970s and 2010s calibrations.

that although mean occupational wages in the calibrated model fit the data quite well, there were discrepancies as large as two to three percent for the 1980 calibration. For the 2010 calibration some of the discrepancies were a bit larger. It follows that changes in occupational mean wages between the two calibrated models will potentially differ by as much as seven or eight percent for some occupations. In fact, the right panel of Figure 6 reveals that there are discrepancies of this magnitude for a couple of occupations. Nonetheless, the fact that the points do follow the 45 degree line indicates that the calibrated model does a good job of capturing the quantitative properties of the polarization facts.

6.3 Counterfactuals

In the previous section we showed that our fully calibrated model does a good job of tracking both the overall change in inequality as well as the polarization facts. In this subsection we carry out counterfactual exercises to understand the relative importance of the different changes in model primitives between 1980 and 2010 for these outcomes.

6.3.1 The Joint Role of Skill Prices and Idiosyncratic Productivity Variance

In our first counterfactual, we allow skill prices and productivity variances to assume their 2010 values but hold all preference parameters and the productivity correlation matrix fixed at their 1980 values. This amounts to changing 18 of the 72 calibrated values. The goal of this counterfactual is to show that these are the two quantitatively important changes for the outcomes of particular interest.

Table 6: The Overall Wage Distribution.

Description	p5	p10	p25	p50	p75	p90	p95	p99
Percent Changes in Wages at Various Percentiles of the Overall Wage Distribution								
Data	-19.7	-20.4	-18.7	-12.9	0.7	15.7	25.7	113.7
Full Model	-21.1	-20.2	-16.2	-8.1	2.9	15.8	25.6	48.0
Restricted Model	-24.4	-22.5	-18.4	-10.3	0.5	13.4	23.3	46.9
Wage Inequality in 2010: Wages at Various Percentiles Relative to Median Wages								
Data	0.35	0.44	0.65	1	1.55	2.34	3.07	7.53
Full Model	0.36	0.44	0.65	1	1.57	2.40	3.16	5.34
Restricted Model	0.35	0.44	0.65	1	1.57	2.41	3.18	5.43


Notes: The table compares moments of the overall wage distribution with those from the full calibration and the restricted model in which we only consider the changes in skill prices and the variance of idiosyncratic productivity.

In this counterfactual, the variance of log wages increases from 0.27 in 1980 to 0.45 in 2010, an increase of 0.18. This increase is modestly larger than the 0.17 increase in the fully recalibrated model, implying that the other changes in primitives have a very modest negative effect on the change in inequality. More generally, Table 6 shows that both the full recalibrated model as well as the restricted model with only changes in skill prices and the variance of idiosyncratic productivity each do a good job of capturing the properties of the wage distribution in 2010 as well as the changes in the wage distribution between 1980 and 2010, with the lone qualification being that the 2010 calibration fails to generate the particularly large increase in the 99th percentile.

To assess the extent to which this counterfactual captures the polarization facts, Figure 7 presents the same information as our earlier Figure 6, but now only allowing for changes in skill prices and the variance of idiosyncratic productivity. Comparing these two figures, the basic mes-

sage is that this counterfactual also does a good job of capturing the polarization facts.²¹

Figure 7: Fit of the Polarization Facts: The Combined Role of Skill Prices and Productivity Variance.



Notes: The figure compares the changes in employment shares and mean occupational wages in the model and data between 1970s and 2010s. The model changes are computed by simulating the changes in skill prices and productivity variances between the 1970s and 2010s model economies, keeping all other parameters fixed.

6.3.2 The Individual Roles of Skill Prices and Productivity Variance

Having established that changes in skill prices and productivity variances are the dominant changes to fundamentals, in this section we examine the relative importance of these two changes. Table 7 presents the results of counterfactual exercises when each of these changes are introduced individually, holding all other primitives fixed at their 1980 values. The exercise where we change the values of skill prices to their 2010 values but hold all other primitives fixed at their 1980s values is similar in spirit to the exercise in Section 5, but differs in the method used to infer the changes in skill prices. In Section 5 we inferred changes in skill prices by trying to match various targets assuming no changes in other primitives, whereas here we take the changes in skill prices from the full recalibration to data from the 2010s.

We highlight five properties of Table 7. First, skill prices have large effects on both the employ-

²¹Table B-6 and Figure B-2 in the Appendix provide a fuller set of results for this counterfactual.

Table 7: The Role of Changes in Skill Prices and Productivity Variances.

Description	Oc1	Oc2	Oc3	Oc4	Oc5	Oc6	Oc7	Oc8	Oc9
Changes in Employment Shares									
Data	0.033	0.009	0.019	-0.029	-0.014	0.004	-0.015	-0.013	0.005
Skill Prices Only	0.040	-0.003	0.007	-0.045	-0.012	0.030	-0.017	0.015	-0.015
Prod. Var. Only	0.003	0.008	0.008	-0.002	0.002	0.002	-0.012	-0.027	0.019
Percent Changes in Mean Wages									
Data	-4.4	0.7	-2.3	-13.9	-8.1	4.2	20.8	12.7	32.3
Skill Prices Only	-3.3	-17.2	-13.0	-20.9	-17.8	-11.9	-18.1	-13.0	-17.5
Prod. Var. Only	-6.8	17.7	12.2	8.9	10.0	13.9	29.9	22.7	55.4
Change in Variance of Log Wages									
Data	-0.012	0.132	0.103	0.067	0.114	0.145	0.246	0.144	0.208
Skill Prices Only	0.000	-0.004	0.005	-0.002	0.000	0.002	-0.002	0.002	-0.004
Prod. Var. Only	0.018	0.111	0.098	0.088	0.089	0.101	0.172	0.158	0.214
Change in Mean to Median Wages									
Data	0.073	0.182	0.155	0.104	0.109	0.170	0.277	0.129	0.237
Skill Prices Only	0.001	-0.002	0.000	0.004	0.001	0.008	-0.002	0.001	-0.001
Prod. Var. Only	0.011	0.075	0.061	0.058	0.062	0.067	0.107	0.099	0.162

Notes: The table reports results from counterfactuals in which we allow only skill prices or log productivity variances to change between 1980 and 2010. When we change log productivity variances, we adjust the mean of log productivity so that mean productivity is unchanged.

ment distribution and mean occupational wages. The average absolute changes between 1980 and 2010 for employment shares and the percent change in mean wages in the data are 0.0155 and 11.0 respectively. When we change only skill prices, the corresponding values are 0.021 and 14.7.

Second, changing skill prices has virtually no impact on within occupation variances of log wages or the mean to median ratios. In the data the mean absolute change in the variance of log wages is 0.130, but when we change only skill prices the corresponding value is only 0.003. And the mean absolute change in the mean to median ratio is 0.160 in the data and only 0.003 when only skill prices change.

Third, the change in productivity variances has large effects on all four outcomes. The average absolute change in employment shares, the percent change in mean wages, the variance of log wages,

and the mean to median wage are 0.0092, 19.72, 0.12, and 0.078, respectively.²²

Fourth, changes in skill prices and changes in productivity variances tend to have opposing effects on mean wages and roughly orthogonal effects on employment shares. Specifically, the correlation for the effects of skill price and productivity variances on mean wages is -0.50, while for employment share effects the correlation is -0.076. Related to this last point, the correlation between mean wage changes in the data and the component of mean wage changes due to the change in productivity variances is 0.88, while the correlation between the mean wage changes in the data and the component of wage changes due to changes in skill prices is -0.48. The key message is that our model assigns a dominant role to changing variances as the driving force behind changes in mean wages. Put somewhat differently, our results imply that one should not interpret changes in mean wages as reflecting changes in relative demands.

Figure 8 serves to highlight that both changes play a role in accounting for the polarization facts. The top panel shows the contribution of changes in skill prices only to the polarization facts, while the bottom panel does the same for changes in productivity variances. The broad message that emerges is that changes in skill prices only are able to track a substantial portion of the changes in employment shares but account for little of changes in mean wages.²³ In contrast, mean preserving changes in the variances of log productivity alone are able to account for much of the differences in mean wage changes and less of the differences in employment share changes.


Finally, and similar to what we found in Section 5, it is again true that changes in skill prices alone explain little of the overall change in inequality. In fact, if skill prices were the only primitive that changes from 1980 to 2010, there would have been a modest *decrease* in the variance of log wages, falling from 0.27 to 0.26.

Three conclusions emerge from this analysis. First, changes in the occupational wage and employment structure requires considering the interaction between changes in skill prices and changes in productivity variances. Second, changes in relative wages are driven by changes in second moments and should not be interpreted as reflecting changes in relative demands. Third, changes in skill prices alone are effectively irrelevant as a source of changes in inequality as measured by the variance of log wages.

²²We note that the corresponding values for changes in preference parameters are 0.0071, 1.9, 0.005, and 0.031, consistent with our earlier message that our decomposition assigns a very small role to changing preference parameters.

²³This result echoes the finding in Section 5 that changes in skill prices alone were better able to account for changes in employment shares.

Figure 8: Fit of the Polarization Facts: Skill Price Changes vs. Productivity Variance Changes.



Notes: The figure compares the changes in employment shares and mean occupational wages in the model and data between 1980 and 2010. In Panel (a) only the skill prices are changed, while in Panel (b) only the skill distributions are changed.

7 Conclusion

There have been large changes in the occupational wage and employment structure since 1980. If we order occupations according to their mean wages in 1980, the changes in both employment

shares and mean wages exhibit a U-shaped pattern. The literature has coined the term labor market polarization to describe this feature of the data. This process of labor market polarization was accompanied by a large increase in inequality as measured by the cross-sectional variance of log wages.

We use a structural model of occupation choice in the tradition of Roy (1951) to understand the driving forces behind polarization and the sharp rise in inequality, and specifically the connection between polarization and the rise in inequality. In our model, changes in relative occupational skill prices proxy for changes in the relative demand for the labor services provided by different occupations. We were especially interested in assessing the extent to which changes in occupational employment shares and mean wages can be explained purely by changes in skill prices across occupations.

Our analysis yields three main findings. First, although changes in skill prices play a quantitatively important role in shaping the changes in employment shares and mean wages across occupations, they play essentially no role in accounting for the sharp rise in inequality since 1980. Second, changes in relative mean wages across occupations are a very poor proxy for changes in relative demand for occupational services. Third, changes in the variance of idiosyncratic within-occupation productivity is the dominant factor behind the sharp rise in inequality.

We close by noting some important directions for future research. Our analysis interpreted cross-section data through the lens of a static Roy model with nine occupations. Our results are robust to considering cross-section data for all workers or only younger workers. A richer analysis will allow for dynamic considerations and make use of panel data. Keane and Wolpin (1997) is an early example of a dynamic model of occupational choice, though they only allowed for two occupations. Data availability poses an important constraint on the feasibility of such an analysis in the context of the US. Publicly available panel data sets do not have sufficiently large samples to allow one to disaggregate beyond a few occupations. Extending this analysis to other countries may be useful, though many of these data sets do not allow one to compute wage rates since they do not include data on hours.

We have argued that Roy models with many occupations are the appropriate framework for addressing the issues posed in this paper. One challenge for this research agenda is that Roy models with large numbers of occupations lead to a curse of dimensionality in terms of parameterizing the correlations of individual productivity across occupations. Recent work has shown that this curse of dimensionality can be avoided by assuming that productivity draws are iid draws from Frechet

distributions with a common shape parameter. We show that such a specification is not useful for asking questions about inequality in the context of occupational choice. We thus build a theory with lognormal productivity and occupational tastes in which occupational returns and the sorting of workers are affected by both factors. Our theory essentially extends the 2-occupation model of Heckman and Sedlacek (1985, 1990) to nine occupations. Additional progress on connecting large Roy models to the data will also be important for the future development of this literature.

References

- ACEMOGLU, D., “Changes in Unemployment and Wage Inequality: An Alternative Theory and Some Evidence,” *American Economic Review* 89 (1999), 1259–78.
- ACEMOGLU, D. AND D. AUTOR, “Skills, Tasks and Technologies: Implications for Employment and Earnings,” in O. Ashenfelter and D. Card, eds., *Handbook of Labor Economics* volume 4B (Amsterdam: Elsevier, 2011), 1043–1171.
- ACEMOGLU, D. AND P. RESTREPO, “Tasks, Automation, and the Rise in US Wage Inequality,” *Econometrica* 90 (2022), 1973–2016.
- ARTUC, E., S. CHAUDHURI AND J. MCLAREN, “Trade Shocks and Labor Adjustment: A Structural Empirical Approach,” *American Economic Review* 100 (2010), 1008–1045.
- AUTOR, D., “Work of the Past, Work of the Future,” *AEA Papers and Proceedings* 109 (2019), 1–32.
- 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.
- AUTOR, D., F. LEVY AND R. J. MURNANE, “The Skill Content of Recent Technological Change: An Empirical Exploration,” *Quarterly Journal of Economics* 118 (2003), 1279–1333.
- BÁRÁNY, Z. L. AND C. SIEGEL, “Job Polarization and Structural Change,” *American Economic Journal: Macroeconomics* 10 (2018), 57–89.
- BASCO, S., M. LIEGEY, M. MESTIERI AND G. SMAGGHUE, “The Effect of Import Competition Across Occupations,” *Journal of International Economics* (2025), Forthcoming.
- BÖHM, M., H.-M. VON GAUDECKER AND F. SCHRAN, “Occupation Growth, Skill Prices, and Wage Inequality,” *Journal of Labor Economics* 42 (2024), 201–243.
- BURSTEIN, A., E. MORALES AND J. VOGEL, “Changes in Between-Group Inequality: Computers, Occupations, and International Trade,” *American Economic Journal: Macroeconomics* 11 (2019), 348–400.
- CARRILLO-TUDELA, C. AND L. VISSCHERS, “Unemployment and Endogenous Reallocation Over the Business Cycle,” *Econometrica* 91 (2023), 1119–1153.

- CAUNEDO, J., D. JAUME AND E. KELLER, “Occupational Exposure to Capital-Embodied Technical Change,” *American Economic Review* 6 (2023), 1642–1685.
- COMIN, D., A. DANIELI AND M. MESTIERI, “Income-Driven Labor-Market Polarization,” Mimeo, Dartmouth College, 2020.
- DIX-CARNEIRO, R., “Trade Liberalization and Labor Market Dynamics,” *Econometrica* 82 (2014), 825–885.
- EROSA, A., L. FUSTER, G. KAMBOUROV AND R. ROGERSON, “Hours, Occupations, and Gender Differences in Labor Market Outcomes,” *American Economic Journal: Macroeconomics* 14 (2022), 543–590.
- , “Wage and Earnings Inequality Between and Within Occupations: The Role of Labor Supply,” *Journal of Monetary Economics* 145 (2024), 103553.
- 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).
- GOOS, M. AND A. MANNING, “Lousy and Lovely Jobs: The Rising Polarization of Work in Britain,” *Review of Economics and Statistics* 89 (2007), 118–133.
- GOOS, M., A. MANNING AND A. SALOMONS, “Job Polarization in Europe,” *American Economic Review* 89 (2009), 58–63.
- GRIGSBY, J., “Skill Heterogeneity and Aggregate Labor Market Dynamics,” Working paper #30052, NBER, 2022.
- HEATHCOTE, J., F. PERRI AND G. L. VIOLANTE, “The Rise of US Earnings Inequality: Does the Cycle Drive the Trend?,” *Review of Economic Dynamics* 37 (2020), 181–204.
- HECKMAN, J. AND B. E. HONORÉ, “The Empirical Content of the Roy Model,” *Econometrica* 58 (1990), 1121–1149.
- HECKMAN, J. AND G. SEDLACEK, “Heterogeneity, Aggregation, and Market Wage Functions: An Empirical Model of Self-selection in the Labor Market,” *Journal of Political Economy* 93 (1985), 1077–1125.

- , “Self-selection and the Distribution of Hourly Wages,” *Journal of Labor Economics* 8 (1990), S329–363.
- HSIEH, C.-T., E. HURST, C. I. JONES AND P. J. KLENOW, “The Allocation of Talent and U.S. Economic Growth,” *Econometrica* 85 (2019), 1439–1474.
- HUCKFELDT, C., “Understanding the Scarring Effect of Recessions,” *American Economic Review* 112 (2022), 1273–1310.
- HUNT, J. AND R. NUNN, “Has U.S. Employment Really Polarized? A Critical Reappraisal,” *Labour Economics* 75 (2022), 102117.
- JOHNSON, M. AND M. P. KEANE, “A Dynamic Equilibrium Model of the US Wage Structure, 1968–1996,” *Journal of Labor Economics* 31 (2013), 1–49.
- KAMBOUROV, G. AND I. MANOVSKII, “Occupational Mobility and Wage Inequality,” *Review of Economic Studies* 76 (2009), 731–759.
- KATZ, L. F. AND K. M. MURPHY, “Changes in Relative Wages, 1963–1987: Supply and Demand Factors,” *The Quarterly Journal of Economics* 107 (1992), 35–78.
- KEANE, M. P. AND K. I. WOLPIN, “The Career Decisions of Young Men,” *Journal of Political Economy* 105 (June 1997), 473–522.
- LAGAKOS, D. AND M. E. WAUGH, “Selection, Agriculture, and Cross-Country Productivity Differences,” *American Economic Review* 103 (2013), 948–980.
- LIND, N. AND N. RAMONDO, “Trade with Correlation,” *American Economic Review* 113 (2023), 317–353.
- LOCHNER, L., Y. PARK AND Y. SHIN, “The Evolution of Unobserved Skill Returns in the US: A New Approach Using Panel Data,” Working paper #33407, NBER, 2025.
- MISHEL, L., J. SCHMITT AND H. SHIERHOLZ, “Don’t Blame the Robots: Assessing the Job Polarization Explanation of Growing Wage Inequality,” Working paper, Economic Policy Institute, 2013.
- ROY, A. D., “Some Thoughts on the Distribution of Earnings,” *Oxford Economic Papers* 3 (1951), 135–146.

A Roy Frechet Model

A.1 Proof of Proposition 1

We have assumed skill draws are made from a multivariate (IID) Frechet distribution

$$F(\epsilon_1, \dots, \epsilon_J) = \exp \left[- \sum_{j=1}^J \epsilon_j^{-\theta} \right], \quad (\text{A-1})$$

where θ is the shape parameter (scale has been set to 1, and location parameter to 0).

The wage among workers choosing occupation j is $w_j \epsilon_j$. Individual choose occupation j when the return to this occupation is higher than for all other occupations $u_j = w_j \epsilon_j v_j = \max_s w_s \epsilon_s v_s$, so that $\epsilon_s \leq \epsilon_j \frac{w_j v_j}{w_s v_s} = \epsilon_j \frac{\tilde{w}_j}{\tilde{w}_s} \forall s$.

The employment share in occupation j is equal to the probability a worker choosing occupation j .

$$\begin{aligned} P(\text{choose } j) &= \Pi_{s=1}^J P(u_s \leq u_j) = \Pi_{s=1}^J P(\epsilon_s \leq \epsilon_j \frac{w_j v_j}{w_s v_s}) = \\ &= \int_0^\infty \theta \epsilon_j^{-1-\theta} \exp[-\epsilon_j^{-\theta}] \Pi_{s \neq j} P(\epsilon_s \leq \epsilon_j \frac{w_j v_j}{w_s v_s}) d\epsilon_j \\ &= \int_0^\infty \theta \epsilon_j^{-1-\theta} \exp[-\epsilon_j^{-\theta}] \exp \left[- \sum_{s \neq j} \left(\epsilon_j \frac{\tilde{w}_j}{\tilde{w}_s} \right)^{-\theta} \right] d\epsilon_j, \\ &= \int_0^\infty \theta \epsilon_j^{-1-\theta} \exp \left[- \sum_{s=1}^J \left(\epsilon_j \frac{\tilde{w}_j}{\tilde{w}_s} \right)^{-\theta} \right] d\epsilon_j \\ &= \int_0^\infty \theta \epsilon_j^{-1-\theta} \exp \left[- \epsilon_j^{-\theta} \sum_{s=1}^J \left(\frac{\tilde{w}_j}{\tilde{w}_s} \right)^{-\theta} \right] d\epsilon_j \\ &= \frac{\exp \left[- \epsilon_j^{-\theta} \sum_{s=1}^J \left(\frac{\tilde{w}_j}{\tilde{w}_s} \right)^{-\theta} \right]}{\sum_{s=1}^J \left(\frac{\tilde{w}_j}{\tilde{w}_s} \right)^{-\theta}} \Big|_0^\infty \\ &= \frac{1}{\sum_{s=1}^J \left(\frac{\tilde{w}_j}{\tilde{w}_s} \right)^{-\theta}} = \frac{\tilde{w}_j^\theta}{\sum_{s=1}^J \tilde{w}_s^\theta} = e_j \end{aligned} \quad (\text{A-2})$$

The CDF for skills in occupation j (ϵ_j) among workers choosing occupation j is:

$$\begin{aligned} P(\epsilon_j < x | j \text{ chosen}) &= \Pi_{s=1}^J P \left(\epsilon_s < x \frac{\tilde{w}_j}{\tilde{w}_s} \right) \\ &= F \left(x \frac{\tilde{w}_j}{\tilde{w}_1}, \dots, x \frac{\tilde{w}_j}{\tilde{w}_J} \right) = \exp \left[- \sum_{s=1}^J \left(x \frac{\tilde{w}_j}{\tilde{w}_s} \right)^{-\theta} \right] = \exp[-m_j x^{-\theta}], \\ \text{where } m_j &= \sum_{s=1}^J \left(\frac{\tilde{w}_s}{\tilde{w}_j} \right)^\theta = 1/e_j, \end{aligned} \quad (\text{A-3})$$

which corresponds to the CDF of the Frechet with scale parameter $(m_j)^{\frac{1}{\theta}}$:

$$P(\epsilon_j < x | j \text{ chosen}) = \exp \left[- \left(\frac{x}{m_j^{1/\theta}} \right)^{-\theta} \right] \quad (\text{A-4})$$

Hence, when the returns to occupation j increase ($\tilde{w}_j \uparrow$), the employment share $e_j \uparrow$, and the scale parameter $m_j \downarrow$, implying that the skill distribution shifts down ($\exp[-m_j x^{-\theta}] \uparrow$).

Using (A-4), the CDF of wages in occupation j is:

$$P(w_j \epsilon_j < w | j \text{ chosen}) = P(\epsilon_j < \frac{w}{w_j} | j \text{ chosen}) = \exp \left[- \left(\frac{w}{w_j (m_j)^{\frac{1}{\theta}}} \right)^{-\theta} \right], \quad (\text{A-5})$$

which, using (A-3), yields the Frechet distribution with scale parameter

$$w_j (m_j)^{\frac{1}{\theta}} = w_j \left[\sum_{s=1}^J \left(\frac{\tilde{w}_s}{\tilde{w}_j} \right)^{\theta} \right]^{\frac{1}{\theta}} = \frac{1}{v_j} \left[\sum_{s=1}^J (\tilde{w}_s)^{\theta} \right]^{\frac{1}{\theta}} = \bar{w}_j \quad (\text{A-6})$$

If v_j does not vary across occupations, all occupations will have the same wage distribution since the scale parameter $\bar{w}_j$ is constant across occupations. If v_j varies across occupations, inversion of the CDF of the wage distribution implies that the percentile q of the wage distribution in occupation j is given by

$$w_j(q) = - \frac{\bar{w}_j}{[\ln q]^{\frac{1}{\theta}}}, \quad (\text{A-7})$$

so that the wage ratio at percentiles q_1 and q_2 does not depend on j . Occupational wage inequality (normalized by mean, median, or any percentile) is constant across occupations.

A.2 Changes in the occupational structure

If the taste for occupation j increases ($v_j \uparrow$), the employment share of occupation j rises. The scale parameter in occupation j decreases so that the mean wage in j decreases relative to other occupations. There is no change in wage inequality within occupations (as measured by the coefficient of variation of any two percentiles of the wage distribution). If the wage rate in occupation j increases ($w_j \uparrow$), the employment share in occupation j rises, and the scale parameter $\bar{w}_j$ increases proportionally for all occupations, implying that the wage distribution shifts right proportionally *in all occupations* (see equations (A-2), (A-5), and (A-6)).

B Figures and Tables

Table B-1: Matrix of Skill Correlations: Baseline Economy (1976-1985).

Skill	Oc1	Oc2	Oc3	Oc4	Oc5	Oc6	Oc7	Oc8	Oc9
Oc1	1.0000	-0.1922	0.0589	0.6234	0.0295	0.1656	0.0255	0.0009	0.0192
Oc2	-0.1922	1.0000	0.0698	0.4380	0.0624	-0.0033	0.1722	0.2471	0.0490
Oc3	0.0589	0.0698	1.0000	0.0585	0.5517	0.3154	0.1188	0.0242	0.1365
Oc4	0.6234	0.4380	0.0585	1.0000	0.2615	0.3912	0.3164	0.1404	0.2371
Oc5	0.0295	0.0624	0.5517	0.2615	1.0000	0.5968	0.4348	0.5563	0.2332
Oc6	0.1656	-0.0033	0.3154	0.3912	0.5968	1.0000	0.6865	0.5713	0.1317
Oc7	0.0255	0.1722	0.1188	0.3164	0.4348	0.6865	1.0000	0.6771	0.3612
Oc8	0.0009	0.2471	0.0242	0.1404	0.5563	0.5713	0.6771	1.0000	0.1503
Oc9	0.0192	0.0490	0.1365	0.2371	0.2332	0.1317	0.3612	0.1503	1.0000

Notes: The table reports the correlation between skills implied by the calibrated variance-covariance matrix in the baseline economy.

Table B-2: Pure Effect of Changing Skill Prices.

Description	Oc1	Oc2	Oc3	Oc4	Oc5	Oc6	Oc7	Oc8	Oc9
Panel A: Data									
Δ Occ. Emp. Shares	0.031	0.009	0.019	-0.028	-0.014	0.004	-0.016	-0.013	0.006
% Δ Mean Wage	-4.4	0.7	-2.3	-13.9	-8.1	4.2	20.8	12.7	32.3
Δ Var Log Wage	-0.012	0.132	0.103	0.067	0.114	0.145	0.246	0.144	0.208
Δ Mean to Median Wage	0.073	0.182	0.155	0.104	0.109	0.170	0.277	0.129	0.237
Panel B: Targeting Employment Share Changes									
% Δ Skill Price	26.4	17.2	16.6	2.3	2.7	8.6	3.8	5.1	10.6
Δ Occ. Emp. Shares	0.033	0.009	0.019	-0.029	-0.014	0.004	-0.015	-0.013	0.005
% Δ Mean Wage	19.7	12.7	13.9	3.7	4.2	8.7	4.8	6.7	9.1
Δ Var Log Wage	-0.0010	-0.0005	0.0030	-0.0013	-0.0037	0.0011	-0.0008	-0.0033	0.0007
Δ Mean to Median Wage	-0.0024	-0.0013	0.0008	0.0008	0.0048	-0.0001	-0.0001	-0.0034	0.0032
Panel C: Targeting Mean Wage Changes									
% Δ Skill Price	-14.9	-5.8	-10.0	-23.0	-14.6	1.3	24.6	15.1	75.6
Δ Occ. Emp. Shares	-0.029	-0.009	-0.036	-0.110	-0.048	-0.022	0.053	0.022	0.180
% Δ Mean Wage	-4.4	0.6	-2.3	-13.9	-8.1	4.2	20.8	12.8	32.3
Δ Var Log Wage	-0.0213	0.0011	-0.0040	-0.0057	-0.0043	-0.0003	0.0039	0.0024	0.0317
Δ Mean to Median Wage	0.0001	0.0025	-0.0008	0.0046	0.0088	0.0068	0.0071	0.0007	0.0257

Notes: The table reports results from three experiments in which skill prices are changed in the baseline economy calibrated to data from 1976-1985.

Table B-3: Calibration Results: 2006-2015 Targets.

Parameters					Calibration Results			
Occupation	w_j	$\sigma_{\epsilon_j}^2$	$\mu_{v,j}$	$\sigma_{v_j}^2$	Employ. Data	Shares Model	Mean Wages Data	Wages Model
1	1.006	0.413	0.395	0.242	0.106	0.105	6.251	6.153
2	0.663	0.628	0.180	0.226	0.039	0.040	8.085	7.714
3	1.199	0.513	0.200	0.236	0.118	0.117	8.845	8.605
4	1.526	0.343	0.190	0.193	0.143	0.143	8.524	8.321
5	1.625	0.351	0	0.119	0.065	0.065	9.883	9.604
6	1.700	0.400	0.176	0.033	0.108	0.107	11.847	11.257
7	1.924	0.415	-0.241	0.141	0.117	0.118	15.097	14.143
8	1.881	0.576	-0.023	0.124	0.217	0.217	16.226	15.640
9	1.340	1.018	-0.200	0.054	0.087	0.088	22.835	22.493

Notes: The table reports parameter and calibration results for the baseline model calibrated to US data from 2006-2015. The first column reports the log of occupational TFP (occupational prices), the second column the variance of log skills, the third column the mean of the log of occupational tastes, and the fourth column the variance of log occupational tastes. The rest of the column reports the fit of the calibration targets regarding employment shares and mean wages by occupation.

Table B-4: Matrix of Skill Correlations: 2006-2015 Economy.

Skill	Oc1	Oc2	Oc3	Oc4	Oc5	Oc6	Oc7	Oc8	Oc9
Oc1	1.0000	-0.1580	0.0283	0.6013	0.01750	0.0994	0.0135	-0.0009	0.0069
Oc2	-0.1580	1.0000	0.0663	0.5158	0.0213	0.0266	0.1059	0.0807	0.0297
Oc3	0.0283	0.0663	1.0000	0.0528	0.4803	0.2221	0.1090	0.0050	0.0682
Oc4	0.6013	0.5158	0.0528	1.0000	0.2776	0.3385	0.2258	0.0492	0.2674
Oc5	0.0175	0.0213	0.4803	0.2776	1.0000	0.7258	0.2535	0.6112	0.2786
Oc6	0.0994	0.0266	0.2221	0.3385	0.7258	1.0000	0.5527	0.7276	0.2466
Oc7	0.0135	0.1059	0.1090	0.2258	0.2535	0.5527	1.0000	0.4866	0.4099
Oc8	-0.0009	0.0807	0.0050	0.0492	0.6112	0.7276	0.4866	1.0000	0.1004
Oc9	0.0069	0.0297	0.0682	0.2674	0.2786	0.2466	0.4099	0.1004	1.0000

Notes: The table reports the correlation between skills implied by the calibrated variance-covariance matrix in the baseline economy.

Table B-5: Change in Skill Correlations: 2006-2015 vs. 1976-1985 Economy.

Skill	Oc1	Oc2	Oc3	Oc4	Oc5	Oc6	Oc7	Oc8	Oc9
Oc1	0	0.0342	-0.0306	-0.0221	-0.0120	-0.0662	-0.0120	-0.0018	-0.0123
Oc2	0.0342	0	-0.0035	0.0778	-0.0411	0.0299	-0.0663	-0.1664	-0.0193
Oc3	-0.0306	-0.0035	0	-0.0057	-0.0714	-0.0933	-0.0098	-0.0192	-0.0683
Oc4	-0.0221	0.0778	-0.0057	0	0.0161	-0.0527	-0.0906	-0.0912	0.0303
Oc5	-0.0120	-0.0411	-0.0714	0.0161	0	0.1290	-0.1813	0.0549	0.0454
Oc6	-0.0662	0.0299	-0.0933	-0.0527	0.1290	0	-0.1338	0.1563	0.1149
Oc7	-0.0120	-0.0663	-0.0098	-0.0906	-0.1813	-0.1338	0	-0.1905	0.0487
Oc8	-0.0018	-0.1664	-0.0192	-0.0912	0.0549	0.1563	-0.1905	0	-0.0499
Oc9	-0.0123	-0.0193	-0.0683	0.0303	0.0454	0.1149	0.0487	-0.0499	0


Notes: The table reports the changes in the correlation between skills in the economy calibrated to 2010 relative to the baseline economy calibrated to 1980.

Table B-6: Restricted vs Full Model.

Description	Oc1	Oc2	Oc3	Oc4	Oc5	Oc6	Oc7	Oc8	Oc9
Changes in Occupational Shares									
Restricted Model	0.042	0.005	0.012	-0.044	-0.008	0.026	-0.026	-0.015	0.008
Full Model	0.030	0.009	0.018	-0.028	-0.014	0.004	-0.014	-0.013	0.006
Data	0.033	0.009	0.019	-0.029	-0.014	0.004	-0.015	-0.013	0.005
Changes in Mean Wages (in %)									
Restricted Model	-8.3	-2.2	-2.9	-13.6	-7.9	0.1	7.67	6.5	28.8
Full Model	3.8	-3.8	-2.7	-14.7	-10.2	0.2	14.6	8.5	26.6
Data	-4.4	0.7	-2.3	-13.9	-8.1	4.2	20.8	12.7	32.3
Var Log Wages									
Restricted Model	0.015	0.104	0.101	0.084	0.092	0.104	0.171	0.155	0.205
Full Model	0.000	0.097	0.099	0.077	0.085	0.092	0.157	0.132	0.196
Data	-0.012	0.132	0.103	0.067	0.114	0.145	0.246	0.144	0.208


Notes: The table reports results from the restricted model (setting the variance of skills and the prices implied by the 2006-2015 calibration into the baseline economy calibrated to 1976-1985) and the full model (calibrated to the targets in 2006-2015).

Figure B-1: Wage Distribution: Calibration to US Data in 2006-2015.



Notes: The figure displays the occupational wage distributions in the model and the data targets from 2006-2015. The left panel shows the percentile of the wage distributions for occupations 1 to 5, while the right panel displays the wage distributions for occupations 5 to 9.


Figure B-2: Wages: The Restricted Model and the Data.




Notes: The figure compares data moments on the wage distribution with those generating when feeding into the baseline model the prices and the variance of skills obtained when calibrating the model to data from 2006-2015. The left panel shows the percentile of the wage distributions for occupations 1 to 5, while the right panel displays the wage distributions for occupations 5 to 9.

Figure B-3: Fit of the Polarization Facts: Skill Price Changes vs. Productivity Variance Changes.

Panel a: Only Skill Price Changes



Panel b: Only Prod. Var. Changes



Notes:

C Sensitivity Analysis: Age Group 25-35

Table C-1: Calibration Results, 1976-1985, Age Group 25-35.

Parameters					Calibration Results			
Occupation	w_j	$\sigma_{\epsilon_j}^2$	$\mu_{v,j}$	$\sigma_{v_j}^2$	Employ. Data	Shares Model	Mean Wages Data	Model
1	0.929	0.413	0.349	0.309	0.071	0.071	6.26	6.15
2	0.861	0.434	0.194	0.177	0.032	0.032	7.62	7.63
3	1.336	0.373	0.236	0.193	0.112	0.114	8.66	8.49
4	1.732	0.249	0.187	0.159	0.179	0.179	9.31	9.14
5	1.792	0.226	0	0.111	0.084	0.084	10.10	9.92
6	1.775	0.225	0.179	0.024	0.108	0.107	10.28	10.17
7	2.088	0.175	-0.260	0.196	0.124	0.124	11.44	11.26
8	2.012	0.258	-0.038	0.156	0.209	0.209	12.23	12.09
9	1.740	0.350	-0.193	0.060	0.081	0.080	14.33	14.26

The table reports parameter and calibration results for the baseline model calibrated to US data from 1976-1985 (targets for individuals aged 25-35). The first column reports the log of occupational skill prices, the second column the variance of log skills, the third column the mean of the log of occupational tastes, and the fourth column the variance of log occupational tastes. The rest of the columns reports the fit of the calibration targets regarding employment shares and mean wages by occupation.

Table C-2: Calibration Results, 2006-2015 Targets, Age Group 25-35.

Parameters					Calibration Results			
Occupation	A_j	$\sigma_{\epsilon_j}^2$	$\mu_{v,j}$	$\sigma_{v_j}^2$	Employ. Data	Shares Model	Mean Wages Data	Wages Model
1	1.010	0.363	0.421	0.247	0.130	0.129	5.673	5.561
2	0.679	0.542	0.192	0.205	0.046	0.044	7.213	6.838
3	1.177	0.460	0.236	0.252	0.138	0.138	7.746	7.544
4	1.452	0.316	0.183	0.192	0.142	0.143	7.692	7.492
5	1.516	0.339	0	0.137	0.066	0.066	8.961	8.643
6	1.583	0.354	0.188	0.039	0.104	0.105	9.967	9.481
7	1.755	0.362	-0.217	0.197	0.100	0.101	11.360	10.874
8	1.774	0.470	-0.031	0.129	0.195	0.195	13.119	12.669
9	1.360	0.695	-0.183	0.047	0.081	0.079	16.738	16.083

The table reports parameter and calibration results for the baseline model calibrated to US data from 2006-2015 (targets for individuals aged 25-35). The first column reports the log of occupational skill prices, the second column the variance of log skills, the third column the mean of the log of occupational tastes, and the fourth column the variance of log occupational tastes. The rest of the column reports the fit of the calibration targets regarding employment shares and mean wages by occupation.

Table C-3: Pure Effect of Changing Skill Prices, 1980 to 2010, Age Group 25-35.

Description	Oc1	Oc2	Oc3	Oc4	Oc5	Oc6	Oc7	Oc8	Oc9
Panel A: Data									
Δ Oc. Emp Shares	0.058	0.014	0.025	-0.037	-0.018	-0.004	-0.024	-0.014	-0.000
$\Delta\%$ Mean wage	-9.4	-5.3	-10.6	-17.4	-11.3	-3.1	-0.7	7.3	16.8
Δ Var log wage	-0.034	0.104	0.073	0.066	0.124	0.117	0.162	0.136	0.142
Δ Mean to Median Wage	0.042	0.174	0.136	0.099	0.139	0.135	0.174	0.126	0.187
Panel B: Minimizing Loss Function Over Emp. Share and Mean Wage Changes									
$\%\Delta$ Skill Price	25.3	7.9	5.6	-9.2	-9.2	-4.6	-9.3	-5.8	0.5
Δ Occ. Emp. Shares	0.054	0.012	0.023	-0.038	-0.019	-0.004	-0.021	-0.012	0.006
$\%\Delta$ Mean Wage	15.0	2.1	2.3	-7.7	-7.2	-3.3	-7.6	-4.2	-1.4
Δ Var Log Wage	-0.0035	-0.0026	0.0012	-0.0019	0.0012	-0.0008	-0.0017	-0.0021	0.0003
Δ Mean to Median Wage	0.0051	0.0101	0.0029	-0.0010	0.0020	0.0020	-0.0008	-0.0004	0.0020

Notes: The table reports results from an experiment in which skill prices are changed in the baseline economy, calibrated to data from 1976-1985, to match changes between 1976-1985 and 2006-2015 (targets for individuals aged 25-35).

Table C-4: The Role of Changes in Skill Prices and Productivity Variances, Age Group 25-35.

Description	Oc1	Oc2	Oc3	Oc4	Oc5	Oc6	Oc7	Oc8	Oc9
Changes in Employment Shares									
Data	0.058	0.014	0.025	-0.037	-0.018	-0.004	-0.024	-0.014	-0.000
Skill Prices Only	0.045	0.004	0.012	-0.043	-0.017	0.017	-0.026	0.017	-0.009
Prod. Var. Only	-0.002	0.005	0.001	-0.003	0.000	0.004	-0.005	-0.016	0.018
Percent Changes in Mean Wages									
Data	-9.4	-5.3	-10.6	-17.4	-11.3	-3.1	-0.7	7.3	16.8
Skill Prices Only	-1.7	-13.8	-12.3	-20.2	-18.5	-13.3	-20.4	-13.2	-17.0
Prod. Var. Only	-10.0	8.5	5.1	5.0	12.8	13.7	24.2	21.2	31.5
Change in Variance of Log Wages									
Data	-0.034	0.104	0.073	0.066	0.124	0.117	0.162	0.136	0.142
Skill Prices Only	-0.002	-0.005	0.000	-0.002	-0.001	0.000	-0.001	0.001	-0.002
Prod. Var. Only	-0.012	0.056	0.053	0.050	0.074	0.073	0.125	0.119	0.130
Change in Mean to Median Wages									
Data	0.042	0.174	0.136	0.099	0.139	0.135	0.174	0.126	0.187
Skill Prices Only	0.004	0.006	0.001	0.003	0.001	0.002	-0.001	0.002	-0.003
Prod. Var. Only	-0.007	0.040	0.035	0.035	0.043	0.050	0.083	0.079	0.090

Notes: The table reports results from counterfactuals in which we allow only skill prices or log productivity variances to change between 1980 and 2010 (targets for individuals aged 25-35). When we change log productivity variances we adjust the mean of log productivity so that mean productivity is unchanged.