The Allocation of Talent and U.S. Economic Growth

Chang-Tai Hsieh
Chicago Booth and NBER

Erik Hurst
Chicago Booth and NBER

Charles I. Jones
Stanford GSB and NBER

Peter J. Klenow *
Stanford University and NBER

February 16, 2012 – Version 0.8
Preliminary and Incomplete

Abstract

In 1960, 94 percent of doctors were white men, as were 96 percent of lawyers and 86 percent of managers. By 2008, these numbers had fallen to 63, 61, and 57 percent, respectively. Given that innate talent for these professions is unlikely to differ between men and women or between blacks and whites, the allocation of talent in 1960 suggests that a substantial pool of innately talented black men, black women, and white women were not pursuing their comparative advantage. This paper estimates the contribution to U.S. economic growth from the changing occupational allocation of white women, black men, and black women between 1960 and 2008. We find that the contribution is significant: 16 to 20 percent of growth over this period might be explained simply by the improved allocation of talent within the United States.

*We are grateful to seminar participants at Boston University, Boston College, Brown, Chicago, Harvard, MIT, UCLA, and USC for helpful comments, and to Huiyu Li and Gabriel Ulyssea for excellent research assistance.

1. Introduction

Fifty years ago, there were stark differences in the occupational distribution of white men versus women and blacks. For example, virtually all doctors, lawyers, engineers, and executives and managers in 1960 were white men: 94 percent of doctors, 96 percent of lawyers, 99 percent of engineers, and 86 percent of executives and managers. In contrast, 58 percent of white women were employed as nurses, teachers, sales clerks, secretaries, and food preparers; 54 percent of black men were employed as freight handlers, drivers, machine operators, and janitors. A vast literature has documented how these gaps have narrowed since then, particularly in high-skilled occupations.[1] By 2008, only 63 percent of doctors and 61 percent of lawyers were white men. Similarly, the share of women and blacks in skilled occupations increased from 2 percent in 1960 to 15 percent for women and 11 percent for black men by 2008.[2]

This paper measures the aggregate effect of the changes in the occupational distribution through the prism of a Roy (1951) model of occupational choice. We assume that every person is born with a range of talents across all possible occupations and chooses the occupation where she earns the highest returns. In this framework, differences in the occupational distribution between white men and women and blacks can arise if the distribution of talent differs between these groups; Rendall (2010), for example, shows that brawn intensive occupations (such as construction) in the U.S. are dominated by men. However, it seems unlikely that natural differences in ability between groups can explain much of the differences that we see in occupational choice and how it has changed over the last fifty years. Consider the world that Supreme Court Justice Sandra Day O'Connor faced when she graduated from Stanford Law School in 1952. Despite being ranked third in her class, the only private sector job she could get was as a legal secretary (Biskupic, 2006). Such barriers might explain why white men dominated the legal profession

[1]We will not attempt to survey this literature, but see Blau (1998), Goldin (1990), and Smith and Welch (1989) for assessments of this evidence.

[2]We define skilled occupations as as executives, managers, architects, engineers, computer scientists, mathematicians, scientists, doctors, and lawyers.

at that time. And the fact that private law firms are now more open to hiring talented female lawyers might explain why the share of women in the legal profession has increased dramatically over the last fifty years. Similarly, the Civil Rights movement of the 1960s is surely important in explaining the change in the occupational distribution of blacks over the last fifty years.[3]

To capture these forces, we depart from the standard Roy model in two dimensions. First, we allow for the possibility that each group faces different occupational frictions in the labor market. We model these frictions as a group-occupation specific "tax" on earnings that drives a wedge between the group's marginal product in the occupation and their take home pay. One interpretation of these "taxes" is that they represent preference-based discrimination as in Becker (1957). For example, one reason why private law firms would not hire Justice O'Connor is that the law firms' partners (or their customers) viewed the otherwise identical legal services provided by female lawyers as somehow less valuable.[4] Second, we allow for group-specific frictions in the acquisition of human capital. We model these frictions as a tax for each group and each occupation on the inputs into human capital production. These human capital frictions could represent the fact that some groups were restricted from elite higher education institutions, that black public schools are underfunded relative to white public schools, that there are differences in prenatal or early life health investments across groups, or that social forces steered certain groups towards certain occupations.[5]

In our augmented Roy model, all three forces — relative ability, labor market frictions, and human capital frictions — can generate gaps in the occupational distri-

[3]See Donohue and Heckman (1991) for an assessment of the effect of federal civil rights policy on the economic welfare of blacks.

[4]Consistent with the Becker (1957) interpretation of labor market frictions, Charles and Guryan (2008) show that relative black wages are lower in states where the marginal white person is more prejudiced (against blacks).

[5]Here is an incomplete list of the enormous literature on these forces. Karabel (2005) documents how Harvard, Princeton, and Yale systematically discriminated against blacks, women, and Jews in admissions until the late 1960s. Card and Krueger (1992) documents that public schools for blacks in the U.S. South in the 1950s were underfunded relative to schools for white children. See Chay, Guryan and Mazumder (2009) for evidence on the importance of improved access to health care for blacks. See Fernandez (2007) and Fernandez, Fogli and Olivetti (2004) on the role of social forces in women's occupational choice.

bution and relative wages between groups. To make progress analytically, we follow McFadden (1974) and Eaton and Kortum (2002) and assume that the distribution of talent follows an extreme value distribution. This assumption gives us two key results. First, we get a closed form expression relating the relative fraction of workers in different groups in an occupation to a composite of three forces: 1) the relative talent of the each group in the occupation, 2) the relative occupational frictions faced in the labor market and, 3) the relative friction in the production of human capital. We calculate this composite measure using data from the decadal U.S. Censuses and the American Community Surveys. We find that this composite measure increased dramatically for women and blacks in high-skilled occupations over the last 50 years, but was roughly unchanged in low skilled occupations.

Second, we get the result that the average wage gap between groups depends on a weighted average of the occupational frictions but is invariant across occupations. That is, our theory predicts that the wage gap will *not* be higher in an occupation where the group faces larger frictions. Intuitively, imagine that the barriers facing women in the legal profession decline. This increases the income of existing women lawyers, but it also induces less talented female lawyers to enter the legal profession. With an extreme value distribution, this quality dilution effect exactly offsets the direct effect of lower barriers on the average wage. The average wage of women rises overall but by the same amount in all occupations. Consistent with this prediction, we show that between 1960 and 2008, the relative wage of women in low-skilled occupations increased by almost exactly the same amount as that of women in high-skilled occupations.

Finally, we embed the Roy model in general equilibrium. This allows us to estimate the effect of a reduction in the barriers to occupational choice on aggregate productivity, wages, and labor force participation. In our baseline results, between 16 and 21 percent of aggregate wage growth between 1960 and 2008 is explained by a decline in occupational frictions and the resulting improved allocation of talent. Looking at the individual groups, the reduction in the frictions since 1960 boosts real wages by 44% for white women, 77% for black women, and 60% for black men, but *lowers* real wages by 7.4% for white men. In addition, about 40 percent of the rise in women's labor force participation is attributable to the decline in occupational frictions.

The paper proceeds as follows. The next section lays out the basic model of occupational choice. We then use this framework to measure the frictions in occupational choice between blacks and women versus white men in Section 3. Section 4 embeds the occupational choice framework into general equilibrium, allowing us to explore the macroeconomic consequences of misallocation in Section 5.

2. A Model of Occupational Sorting

The economy consists of a continuum of people working in N possible occupations, one of which is the home sector. Each person possesses an idiosyncratic ability in each occupation — some people are good economists while others are good nurses. The basic economic allocation to be determined in this economy is how to match workers with occupations.

2.1. People

Individuals are members of different groups, such as race and gender, indexed by g. A person with consumption c and leisure time $1-s$ gets utility

$$U = c^{\beta}(1-s) \tag{1}$$

where s represents time spent on schooling, and β parameterizes the tradeoff between consumption and schooling.

Each person chooses to work in an occupation indexed by i. A person's human capital is produced by combining time s and goods e. The production function for human capital in occupation i is

$$h(e,s;i) = s^{\phi_i}e^{\eta} = s^{\phi_i}e^{\eta}, \tag{2}$$

Occupations differ in how useful schooling is in generating human capital. This will give rise to differences in schooling by occupation, which in turn differences in wages across occupations.

We add two frictions here. The first friction affects human capital: a "tax" τ_{ig}^{h} is applied to the goods e invested in human capital, with the tax varying across occupations and groups. We think of this tax as representing forces that affect the cost of acquiring human capital. For example, τ_{ig}^{h} might represent discrimination against blacks or women in admission to universities, or differential allocation of resources to public schools attended by black vs. white children, or parental liquidity constraints that affect children's health and education.

The second friction we consider can be thought of as a friction in the labor market. A person in occupation i and group g is paid a wage equal to $(1-\tau_{ig}^{w})w_i$ where w_i denotes the wage per efficiency unit of labor paid by the firm. One interpretation of τ_{ig}^{w} is that it represents preference-based discrimination by the employer or customers as in Becker (1957).

Consumption is equal to labor income less expenditures on education, incorporating both taxes:

$$c = (1-\tau_{ig}^{w})w\epsilon h(e,s) - e(1+\tau_{ig}^{h}). \tag{3}$$

Note that pre-tax labor income is the product of the wage received per efficiency unit of labor, the idiosyncratic talent draw ϵ in the worker's chosen occupation, and the individual's acquired human capital h. We stress that these taxes have no real resource costs, and we will later assume they average out to zero across individuals in every occupation. The frictions will, however, affect occupational choices, human capital investments, and earnings.

For notational convenience, it will be useful to work with the following expressions: $\delta_{ig}^{h} \equiv (1+\tau_{ig}^{h})^{-\eta}$ and $\delta_{ig}^{w} \equiv 1-\tau_{ig}^{w}$. These δ's move in the opposite direction of the τ's.

2.2. Occupational Skills

Turning to the worker's idiosyncratic talent, we borrow from McFadden (1974)'s and Eaton and Kortum (2002)'s formulation of the discrete choice problem. We assume each person gets an iid skill draw ϵ_i from a Fréchet extreme value distribution for each occupation:

$$F_{ig}(\epsilon) = exp(-T_{ig}\epsilon^{-\theta}). \tag{4}$$

The parameter θ governs the dispersion of skills, with a higher value of θ corresponding to *smaller* dispersion. We assume that θ is common across occupations and groups. The parameter T_{ig}, however, can potentially differ. Across occupations, differences are obvious: talent is easy to come by in some occupations and scarce in others. In some occupations, it also seems reasonable to allow the distribution of talent to differ between men and women. For example, men may be relatively more endowed with physical strength, which is likely to be more valuable in occupations such as firefighting or construction.

2.3. Individual Optimization

The occupational choice problem for an individual can be described as follows. First, given an occupational choice, the occupational wage w_i, and idiosyncratic ability ϵ in that occupation, each individual chooses c, e, s to maximize utility:

$$U(\delta, w, \epsilon) = \max_{c,e,s} (1-s)c^{\beta} \;\; s.t. \;\; c = \delta w \epsilon s^{\phi_i} e^{\eta} - e \tag{5}$$

where $\delta \equiv \delta^h \, \delta^w$. Then, each individual chooses the occupation that maximizes his or her utility: $i^* = \arg\max_i U(\delta_i, w_i, \epsilon_i)$, taking $\{\delta_i, w_i, \epsilon_i\}$ as given.

We summarize the key results that follow from this setup in a series of propositions. Note that for many results, it is convenient to define a single discrimination parameter $\delta_{ig} \equiv \delta^h_{ig} \, \delta^w_{ig}$, which is the product of the human capital and wage discrimination parameters.

Proposition 1 (Individual Consumption and Schooling): *The solution to the indi-*

vidual's utility maximization problem, given an occupational choice, is

$$s_i^* = \frac{1}{1+\frac{1-\eta}{\beta\phi_i}}$$

$$e_{ig}^*(\epsilon) = \left(\eta\delta_{ig} w_i s_i^{\phi_i}\epsilon\right)^{\frac{1}{1-\eta}}$$

$$c_{ig}^*(\epsilon) = \bar{\eta}(\delta_{ig} w_i s_i^{\phi_i}\epsilon)^{\frac{1}{1-\eta}}, \ \ \bar{\eta} \equiv (1-\eta)\eta^{\frac{-\eta}{1-\eta}}$$

$$U(\delta_i, w_i, \epsilon_i) = \bar{\eta}^{\beta}(\tilde{\delta}_{ig}\epsilon_i)^{\frac{\beta}{1-\eta}}, \ \ \tilde{\delta}_{ig} \equiv \delta_{ig} w_i s_i^{\phi_i}(1-s_i)^{\frac{1-\eta}{\beta}}.$$

This result is an intermediate one, with the key piece coming in the last line describing the equation for U_{ig}. In particular, the individual's occupational choice problem then reduces to picking the occupation that delivers the highest value of $\tilde{\delta}_{ig}\epsilon_i$. The assumption that talent draws are iid and drawn from an extreme value distribution delivers the result that the highest utility can also be characterized by an extreme value distribution, a result reminiscent of those in McFadden (1974) and Eaton and Kortum (2002). The overall occupational share can then be obtained by aggregating the optimal choice across people, as we show in the next proposition.

Proposition 2 (Occupational Choice): *Let p_{ig} denote the fraction of people in group g that work in occupation i in equilibrium. Also, define m_g to be a sum of the economic forces that affect the occupational allocation for group g:*

$$m_g \equiv \left(\sum_{s=1}^{N} T_{sg}\tilde{\delta}_{sg}^{\theta}\right)^{\frac{1}{\theta}\cdot\frac{1}{1-\eta}} \quad \textit{where } \tilde{\delta}_{ig} \equiv \delta_{ig} w_i s_i^{\phi_i}(1-s_i)^{\frac{1-\eta}{\beta}}. \tag{6}$$

Aggregating across people, the solution to the individual's occupational choice problem leads to

$$p_{ig} = \frac{T_{ig}\tilde{\delta}_{ig}^{\theta}}{\sum_{s=1}^{N} T_{sg}\tilde{\delta}_{sg}^{\theta}} \tag{7}$$

The relative propensity of a group to work in occupation i is therefore

$$\frac{p_{ig}}{p_{i,wm}} = \frac{T_{ig}}{T_{i,wm}} \cdot \left(\frac{\delta_{ig}}{\delta_{i,wm}}\right)^{\theta} \left(\frac{m_g}{m_{wm}}\right)^{\theta(1-\eta)}. \tag{8}$$

The last equation of the proposition gives the key result for occupational choice. The relatively propensity of a group to work in an occupation (relative to white men) is the product of three terms: the relative mean talent in the occupation, the relative friction, and the relative m's. Shortly, we will show that this last term, the relative m's, is simply given by the aggregate wage gap between the groups.

Proposition 3 (Labor Supply): *The equilibrium labor supply by group g to occupation i is*

$$H_{ig} = \gamma\bar{\eta} q_g p_{ig} \cdot \frac{1}{\delta_{ig}^w w_i} \cdot (1-s_i)^{-1/\beta} \cdot m_g \tag{9}$$

where $\gamma \equiv \Gamma(1 - \frac{1}{\theta} \cdot \frac{1}{1-\eta})$ is related to the mean of the Fréchet distribution for abilities.

Equation (9) gives the equilibrium efficiency units of labor supplied to occupation i by group g. The first term in this equation captures the number of people working in the occupation; the remaining terms capture the "quality" of those people. For example, the second main term, $\frac{1}{\delta_{ig}^w w_i}$, is a selection effect: a higher wage per efficiency unit of labor attracts lower ability people to the occupation, other things equal. The third term captures the fact that occupations with higher schooling will have more human capital. Finally, the last term, m_g, captures a general equilibrium effect: the average quality of workers from group g going into all occupations depends on the average post-friction wages they face.

Proposition 4 (Occupational Wage Gaps): *Let $\overline{wage}_{ig}$ denote the average earnings in occupation i by group g. Its value in equilibrium is*

$$\overline{wage}_{ig} \equiv \frac{\delta_{ig}^w w_i H_{ig}}{q_g p_{ig}} = (1-s_i)^{-1/\beta} \gamma\bar{\eta}\, m_g. \tag{10}$$

Importantly, this implies that the occupational wage gap between any two groups is the same across all occupations. For example,

$$\frac{\overline{wage}_{ig}}{\overline{wage}_{i,wm}} = \left(\frac{\sum_s T_{sg} \tilde{\delta}_{sg}^{\theta}}{\sum_s T_{s,wm} \tilde{\delta}_{s,wm}^{\theta}} \right)^{\frac{1}{\theta} \cdot \frac{1}{1-\eta}} = \frac{m_g}{m_{wm}}. \tag{11}$$

The first equation of the proposition reveals that average earnings only differs across occupations because of the first term, $(1 - s_i)^{-1/\beta}$. Occupations in which schooling is especially productive (a high ϕ_i and therefore a high s_i) will have higher average earnings, and that is the only reason for earnings differences across occupations in the model. For example, occupations that face less discrimination or a better talent pool or a higher wage per efficiency unit do not yield higher average earnings. The reason is that each of these factors leads to lower quality (i.e. lower ϵ) workers entering those jobs. This composition effect exactly offsets the direct effect on earnings when the distribution of talent is Fréchet. This leads to the novel prediction, given in equation (11), that the earnings gap between two groups (men and women, for example, or blacks and whites) will be *constant* across occupations. We test this proposition in the empirical work that follows.

3. Empirically Evaluating the Occupational Sorting Model

3.1. Data

We use data from the 1960, 1970, 1980, 1990, and 2000 Decennial Censuses as well data from the 2006-2008 American Community Surveys (ACS) for all analysis in the paper. When using the 2006-2008 ACS data, we pool all the years together and treat them as one cross section.[6] We make only four restrictions to the raw data when constructing our analysis samples. First, we restrict the analysis to only include white men (wm), white women (ww), black men (bm) and black women (bw). These will be the four groups we analyze in the paper.[7] Second, we restrict the sample to include only individuals between the ages of 25 and 55 (inclusive). This restriction helps to focus our analysis on individuals after they finish schooling and prior to considering retirement. Third, we exclude individuals on active military duty. Fi-

[6]Henceforth, we refer to the pooled 2006-2008 sample as the 2008 sample. A full description of how we process the data, including all the relevant code, is available at http://faculty.chicagobooth.edu/erik.hurst/research/chad_data.html.

[7]We think an interesting extension would be to include Hispanics in the analysis. In 1960 and 1970, however, there are not enough Hispanics in the data to provide reliable estimates of occupational sorting. Such an analysis can be performed starting in 1980. We leave such an extension to future work.

nally, we exclude individuals who report their labor market status as being unemployed (i.e., not working but searching for work). Our model is not well suited to capture transitory movements into and out of employment. Appendix Table A1 reports the sample size for each of our six cross sections, including the fraction of the sample comprised of our four groups.[8]

A key to our analysis is to use the Census data to create a consistent set of occupations over time. We treat the home sector as a separate occupation. Anyone in our data who is not currently employed or who is employed but usually works less than ten hours per week is considered to be working exclusively in the home sector. Those who are employed but usually work between ten and thirty hours per week are classified as being part-time workers. We split the sampling weight of part-time workers equally between the home sector and the occupation to which they are working. Individuals working more than thirty hours per week are considered to be working full-time in an occupation outside of the home sector.

For our base analysis, we define the non-home occupations using the roughly 70 occupational sub-headings from the 1990 Census occupational classification system.[9] We use the 1990 occupation codes as the basis for our occupational definitions because the 1990 occupation codes are available in all Census and ACS years since 1960. We start our analysis in 1960, as this is the earliest year for which the 1990 occupational cross walk is available. Appendix Table A2 reports the 67 occupations we analyze in our main specification using the 1990 occupational subheadings. Example occupations include "Executives, Administrators, and Managers", "Engineers", "Natural Scientists", "Health Diagnostics", "Health Assessment", and "Lawyers and Judges". Appendix Table A3 gives a more detailed description of some of these occupational categories. For example, the "Health Diagnostics" occupation includes physicians, dentists, veterinarians, optometrists, and podiatrists, and the "Health Assessment and Treating" occupations include registered nurses, pharmacists, and dieticians. For short hand, we sometimes refer to these occupations as

[8]For all analysis in the paper, we weight our data using the sample weights available in the different surveys.

[9]http://usa.ipums.org/usa/volii/99occup.shtml.

doctors and nurses, respectively. The way the occupations are defined ensures that each of our occupational categories has positive mass in all years or our analysis.

As seen with the examples above, there is some heterogeneity within our 67 base occupational categories. To assess the importance of such heterogeneity, we perform a series of robustness exercises for many of our main empirical results where we use different levels of occupational aggregation. Specifically, in some robustness specifications, we use the roughly 340 occupations that are consistently defined (using the 1990 occupation codes) in 1980, 1990, 2000, and 2006-8. The reason we start this in 1980 is that the occupational classification system is roughly similar across the Censuses and ACS starting in 1980. We perform our main analysis using the 340 detailed occupation codes for the 1980–2008 period and show that the quantitative outcomes are very similar to what we get using our 67 base occupation codes for the same period. Additionally, we show that much of our quantitative results can be generated if we use only 20 broad occupation categories as opposed to the roughly 67 occupation codes in our base analysis. The 20 occupation categories we use for this robustness analysis are shown in Appendix Table A4. The 20 broad occupation categories include the same universe of 67 occupations just aggregated to broader categories. As we show throughout, our key empirical results come from the fact that women and blacks in recent periods are sorting with a more equal propensity relative to white men in a handful of high skilled occupations.

Our measures of earnings throughout the paper sum together the individual's labor, business, and farm income. The earnings measures in the Census are from the prior year. Implicitly we assume that individuals who are working in a given occupation in the survey year also worked in that same occupation during the prior year which corresponds to their income report. When measuring earnings, we only focus on those individuals who worked at least 48 weeks during the prior year and who had at least 1000 dollars of earnings (in year 2007 dollars). We define wage rate by dividing individual earnings from the prior year by the product of the weeks worked during the prior year and the reported current usual hours worked. When computing individual wage measures, we further the sample to those individuals that report that they usually work more than 30 hours per week.[10]

For a few of our empirical results, we need a measure of the average wages in the home sector. We impute average earnings for the home sector by extrapolating the relationship between average education and average earnings for the 66 non-home occupations taking into account group fixed effects. Using this year-specific relationship by group and the actual year-specific average education and group composition of participants in the home sector, we predict the average earnings of participants in the home sector. Our primary use of these imputed average earnings in the home sector is when we weight some of our estimates by the wage bill in each sector.

3.2. Occupational Sorting and Wage Gaps By Group

We begin our analysis by documenting the large amount of convergence in the occupational distribution between white men and the other groups over the last fifty years. To illustrate this fact, we create a simple occupational similarity index, Ψ_g which is defined as:

$$\Psi_g \equiv 1 - \frac{1}{2}\sum_{i=1}^{N} |p_{i,wm} - p_{ig}| \tag{12}$$

To construct the index, we compute the absolute value of the difference in the propensity for a given group to be in an occupation relative to the propensity that white men are in that occupation. We then sum these differences across all occupations. For ease of interpretation, we normalize the measure so that it runs between zero (no occupational overlap between the two groups) and 1 (complete occupational overlap between the two groups). When computing Ψ_g, we exclude the home sector. However, the broad patterns are very similar — particularly the index for white women — even when the home sector is included.

[10] In some census years, weeks worked during the prior year and usual hours worked are reported as categorical variables. In these instances, we use the midpoint of the range when computing the wage rate. See the full details of our data processing in the detailed online data appendix available on the author's web sites.

Panel A of Table 1 shows the measure of Ψ_g for white women, black men, and black women in 1960, 1980, and 2008. We also do the comparison for lower educated individuals (those with a high school degree and less) and higher educated individuals (those with more than a high school degree). Within the educational categories, for example, we compare the occupational distribution of lower educated white women to the occupational distribution of lower educated white men.

A few things are of note from Panel A of Table 1. First, each group experienced substantial occupational convergence relative to white men between 1960 and 2008. Second, the timing of the convergence occurred differentially across the groups. For example, occupational convergence occurred both during the 1960 and 1980 period and the 1980 and 2008 period for white women and black women. For black men, however, the bulk of the convergence occurred prior to 1960. Third, there are differences in the occupational convergence by educational attainment. This is seen particularly for white women. In 1960, there were substantial occupational differences both between high educated white women and high educated white men and between low educated white women and low educated white men. Specifically, low educated white men primarily worked in construction and manufacturing while low educated white women primarily worked as secretaries or in low skilled services like food service. High educated white men in 1960 were spread out across many high skilled occupations while high educated white women primarily worked as teachers and nurses. Between 1960 and 2008, however, the occupational similarity between higher educated white men and women converged dramatically while the occupational similarity between lower educated white men and women barely changed. Today, low skilled women still primarily work in services and office support occupations while low skilled men still primarily work in construction and manufacturing.

One of the strong predictions of our occupational sorting model is that the wage gaps relative to white men should be constant for a given group across all occupations. The reason for this is that an occupation that pays a high wage per unit of ability will attract less talented workers. As discussed above, this type of sorting is what makes the wage gap between two groups in a given occupation a poor mea-

Table 1: Occupational Similarity and Conditional Wage Gaps Relative to White Men

	Panel A: Occupational Similarity Index, Relative to White Men				
Race-Sex Group	1960	1980	2008	1980-1960 Difference	2008-1980 Difference
White Women: All	0.42	0.49	0.55	0.07	0.06
White Women: High Educated	0.38	0.49	0.59	0.10	0.12
White Women: Low Educated	0.40	0.43	0.46	0.02	0.04
Black Men: All	0.56	0.72	0.76	0.16	0.04
Black Men: High Educated	0.51	0.74	0.77	0.23	0.03
Black Men: Low Educated	0.59	0.75	0.75	0.16	0.00
Black Women: All	0.28	0.42	0.50	0.14	0.08
Black Women: High Educated	0.31	0.44	0.53	0.13	0.11
Black Women: Low Educated	0.27	0.41	0.44	0.14	0.03

	Panel B: Conditional Log Difference in Wages, Relative to White Men				
Race-Sex Group	1960	1980	2008	1980-1960 Difference	2008-1980 Difference
White Women: All	-0.57	-0.47	-0.26	0.10	0.21
White Women: High Educated	-0.50	-0.40	-0.24	0.10	0.16
White Women: Low Educated	-0.56	-0.47	-0.27	0.09	0.20
Black Men: All	-0.38	-0.22	-0.15	0.16	0.07
Black Men: High Educated	-0.29	-0.16	-0.18	0.13	-0.02
Black Men: Low Educated	-0.38	-0.23	-0.12	0.15	0.11
Black Women: All	-0.86	-0.48	-0.31	0.38	0.17
Black Women: High Educated	-0.62	-0.39	-0.31	0.23	0.08
Black Women: Low Educated	-0.88	-0.48	-0.28	0.40	0.20

Note: Panel A of the table reports our occupational similarity index for white women, black men, and black women relative to white men in 1960, 1980, and 2008. The occupational similarity index runs from zero (no overlap in the occupational distribution relative to white men) and one (complete overlap in the occupational distribution relative to white men). The index is computed separately for higher educated and lower educated individuals of the different groups. Panel B reports the difference in log wages between the groups and white men. The entries come from a regression of log wages on group dummies and controls for potential experience and hours worked per week. The regression only includes a sample of individuals working full time. See the text for additional details.

sure of any differential frictions or absolute advantage between the two groups in that occupation. There are, however, at least three reasons why the estimated wage gaps between groups will not be equated across all occupations. First, there is likely some measurement error in the occupational wage gap estimates due to small sample sizes for some groups in some occupations. Second, although we expect sorting will help offset the effect of differences in wages per ability on the average wage in an occupation, the exact offset due to sorting is a feature of the extreme value distribution. We would not get the complete offset if ability is not exactly distributed according to an extreme value distribution. Third, we focus on occupational sorting due to heterogeneity in ability, but some of the occupational sorting might be driven by other factors such as heterogeneity in preferences. High wage (per unit of ability) occupations might induce the entry of people with high disutility for an occupation rather than individuals with low ability in the occupation. All three forces will generate variation in wage gaps across occupations

We realize the model is highly stylized but the prediction with respect to the equivalence of wage gaps across occupations for a given occupation seems born out in the data at least when segmenting individuals by accumulated schooling. Panel B of Table 1 shows the estimated wage gap between white men and, respectively, white women, black men, and black women over time and by educational attainment. To obtain these estimates, we regress log wages of the individual on group dummies, a quadratic in potential experience, a polynomial in usual hours worked and our base specification occupation dummies. This regression is estimated only for those individuals who are currently working more than 30 hours per week and who worked at least 48 weeks during the prior year when earnings were measured. We estimated this regression separately for 1960, 1980, and the pooled 2006-2008 sample. The coefficients on the race-sex dummies are shown in the table and should be interpreted as being log deviations relative to white men. We also estimated the regression separately for individuals with 12 years of less of schooling and for individuals with more than 12 years of schooling.

As seen in panel B of Table 1, the wage gap for white women relative to white men is nearly identical by educational attainment in 1960, 1980, and 2008. For example, in 1960, low educated white women earned a wage that was 56 log points lower than low educated white men. The comparable number for high educated white women relative to high educated white men was 50 log points. Between 1960 and 2008, the relative wage of low educated white women narrowed by 29 log points. During the same time period, the relative wage of high educated white women narrowed by 26 log points. Despite the fact that the change in the relative occupational similarity was very different by educational attainment for white women (as seen in Panel A), the change in the relative wage gap was nearly identical by educational attainment for white women. According to our model, changes in the relative δ's for white women in high ϕ occupations would generate exactly this result.

Between 1960 and 1980, black men also had a relative wage gap that evolved nearly identically within a sample of lower educated individuals and a sample of higher educated individuals. After 1980, however, there was little change in relative occupational sorting for either high or low skilled black men and there was no change in the relative wage gap for high skilled black men. The wage gap for low skilled black men, however, continued to narrow after 1980. This may be due to the fact that there was a rapid decline in the labor market participation of low skilled black men during the last thirty years that was not random. As currently formulated, our model would not predict such results. However, as we discuss in Section 5, the change in labor market outcomes for black men between 1980 and 2008 do not affect our estimates of aggregate productivity gains in any way.

A further test of the plausibility of our framework is to examine occupation by occupation whether the change in the wage gap between two groups in that occupation is in any way related to the change in the relative propensities of the two groups to be in the occupation. Our model suggests that the two should be unrelated with no variation in the wage gaps. Figure 1 plots the (log) occupational wage gap for white women in 1980 against $p_{i,ww}/p_{i,wm}$.[11] The latter variable is the relative

[11] To compute the occupational wage gaps, we regress log wages on a quadratic in experience and controls for usual hours worked. Again, we only do this for full time workers. After taking the residuals from these regressions, we compute the average wage residual for each group in each occupation. The difference in the average wage residuals relative to white men is our measure of the relative wage gap in each occupation.

propensity of a white woman to work in a particular occupation relative to a white man. As an example, in 1980, white women were 65 times more likely than white men to work as a secretary, but only 0.14 times as likely to work as a lawyer. Given this enormous variation, the difference in the wage gaps between these two occupations is remarkably small. White women secretaries earned about 33 percent less than white men secretaries in 1980, while the gap was 41 percent less for lawyers. Fitting a regression line through the points in the figure shows that there is no relationship at all between the relative wage gap in the occupation and the propensity for white women to be in the occupation relative to white men in 1980.[12] The patterns in other years and for other groups were quite similar. Notice that, within the model, it is the relative propensity that pins down any potential frictions facing a group (relative to white men) in that occupation.[13]

Our productivity gains in the subsequent sections are based on the change in the occupational distribution over time. Figure 2 shows that the change in log $p_{i,ww}/p_{i,wm}$ between 1960 and 2008 is also uncorrelated with the change in the wage gap between white women and white men between 1960 and 2008. The relative fraction of white women who are doctors increased by 144 percent between 1960 and 2008. For nurses, in contrast, the relative fraction who are white women decreased by 52 percent. Yet the relative wage gap between white men and white women narrowed by between 20 and 30 log points in both occupations. Again, our model predicts that the change in the wage gaps should be uncorrelated with the change in the occupational sorting. This prediction is born out in the data.

[12]The coefficient on the log $p_{i,ww}/p_{i,wm}$ in a regression of the occupational wage gap on log $p_{i,ww}/p_{i,wm}$ was 0.002 with a standard error of 0.008 and an adjusted R-squared of essentially zero. For interpretation, the standard deviation of the independent variable was 1.96 and the mean of the dependent variable was -0.31. The regression was weighted using the number of individuals in the occupation across all groups.

[13]In additional work (not shown), we explored the relationship between occupational wage gaps and the average earnings of individuals in those occupations. On average, high income occupations tended to have larger wage gaps. This suggests that the extreme value distribution might not entirely correct for high income occupations. Nonetheless, the magnitude of this correlation was almost always small. For example, in 2006-2008, white working women had about a 3 percentage point larger wage gap relative to white men in response to a one-standard deviation increase in occupational log income. As seen from Table 1, the average wage gap was 26 percentage points.

Figure 1: Occupational Wage Gaps for White Women in 1980

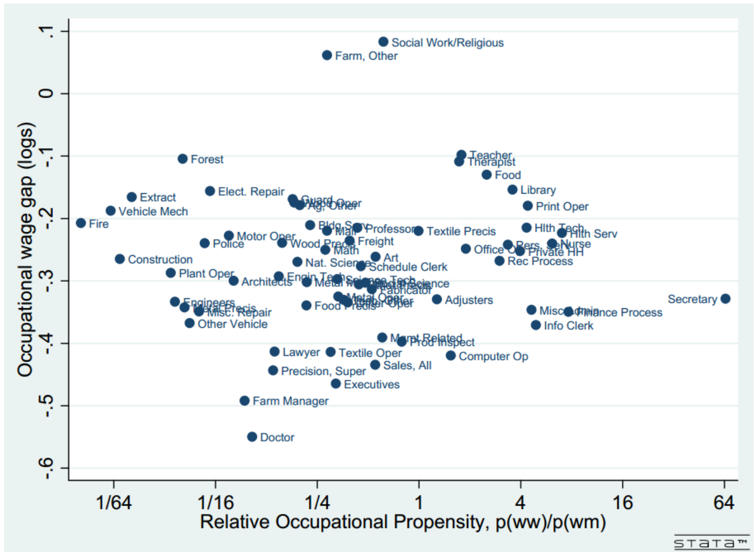

Note: The figure shows the relationship between the (log) occupational wage gap for white women compared to white men and the relative propensity to work in the occupation between white women and white men, $p_{i,ww}/p_{i,wm}$.

Figure 2: Change in Occupational Wage Gaps for White Women,1960–2008

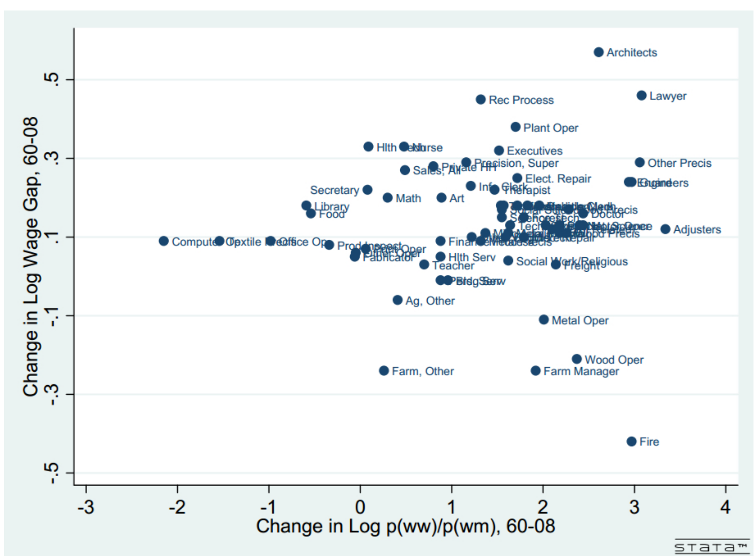

Note: The figure shows the relationship between the change in (log) occupational wage gap for white women compared to white men between 1960 and 2008 and the change in the relative propensity to work in the occupation between white women and white men, $p_{i,ww}/p_{i,wm}$, over the same time period.

3.3. Explaining Occupational Differences Across Groups Over Time

As seen from the previous section, we can use data on the difference in occupational propensities across groups as well as the average wage gaps to infer a composite measure of the occupation-specific relative frictions faced in the labor market between the two groups (τ^w_{ig}), the occupation specific relative frictions faced in the human capital market between the two groups (τ^h_{ig}), or the relative difference in occupation specific absolute advantage between the two groups (T_{ig}). Specifically, given equations (8) and (11), we can define the composite measure for each group (relative to white men) in each occupation as:

$$\hat{\delta}_{ig} \equiv \frac{\delta_{ig}}{\delta_{i,wm}} \left(\frac{T_{i,g}}{T_{i,wm}}\right)^{\frac{1}{\theta}} = \left(\frac{p_{ig}}{p_{i,wm}}\right)^{\frac{1}{\theta}} \left(\frac{\overline{\text{wage}}_g}{\overline{\text{wage}}_{wm}}\right)^{1-\eta}. \tag{13}$$

Aside from the p_{ig}'s and the average wage gap, we also need an estimate of θ and η to compute our estimates of $\hat{\delta}_{ig}$. The parameter θ is a key parameter that governs the dispersion of wages. Given the occupational choice model developed above, one can show that the dispersion of wages across people within an occupation-group obeys a Fréchet distribution with the shape parameter $\theta(1-\eta)$: the lower is this shape parameter, the *more* wage dispersion there is within an occupation. Wage dispersion therefore depends on the dispersion of talent (governed by $1/\theta$) and amplification from accumulating human capital via spending (governed by $1/(1-\eta)$). In particular, the coefficient of variation of wages within an occupation-group in our model satisfies

$$\frac{\text{Variance}}{\text{Mean}^2} = \frac{\Gamma(1-\frac{2}{\theta(1-\eta)})}{\left(\Gamma(1-\frac{1}{\theta(1-\eta)})\right)^2} - 1. \tag{14}$$

To estimate $\theta(1-\eta)$ in a given year, we first take residuals from a cross-sectional regression of log worker wages on 66 occupation dummies and 3 group dummies (one each for white women, black men, and black women). The wage is the hourly wage, and the sample includes both full-time and part-time workers. The occupation dummies capture the effect of schooling requirements (ϕ_i levels) on average wages in an occupation, and the group dummies absorb the wage gaps created by frictions (the average δ_{ig} across occupations for each group). We calculate the mean and variance across workers of the exponent of these wage residuals. We then solve equation (14) for the value of $\theta(1-\eta)$. Sampling error is trivial here because there are 300-400k observations per year for 1960 and 1970 and 2-3 million per year for 1980 onward. The point estimates for $\theta(1-\eta)$ hover around 3. They drift down over time, from 3.25 in 1960 to 2.84 in 2006-2008, as one would expect given rising wage inequality. For our baseline model, we use the simple average of the point estimates across years, namely $\theta(1-\eta) = 3.11$.[14]

The parameter η denotes the elasticity of human capital with respect to education spending. Related parameters have been discussed in the literature, for example by Manuelli and Seshadri (2005) and Erosa, Koreshkova and Restuccia (2010). In our model, η will equal the fraction of output spent on accumulating human capital in equilibrium, separate from time spent accumulating human capital. Absent any solid evidence on this parameter, we set $\eta = 1/4$ in our baseline and explore robustness to $\eta = 0$ and $\eta = 1/2$. In general, this parameter slightly affects the *level* of the δ_{ig} parameters, but not much else in the results.

Table 2 reports our estimates of $\hat{\delta}_{ig}$ for white women for a subset of our baseline occupations. The occupations we highlight in Table 2 represent five different types of occupations. First, we highlight the $\hat{\delta}_{ig}$ for the home sector. Second, we highlight high educated occupations for which white women were underrepresented in 1960. These occupations include executives, engineers, doctors, and lawyers. Third, we highlight high educated occupations like nursing and teachers for which white women were overrepresented relative to white men in 1960. Finally, we show lower educated occupations with a relatively large amount of white women in 1960 (e.g., secretaries and waitresses) and lower educated occupations with relatively few white women in 1960 (e.g., construction, firefighters, and vehicle mechanics).

Many interesting patterns emerge from Table 2. First, consider the results for

[14] When computing our counterfactuals in Section 5, we explore the robustness of results for $\theta(1-\eta)$ estimates that ranges from 2 to 15.

white women in the "home" occupation in 1960. Despite white women being 7 times as likely to work in the home sector as white men, we estimate $\hat{\delta}_{ig}$ for white women in the home sector to be just over one. This implies that white women in 1960 neither had an absolute advantage relative to white men in the home sector nor faced differential frictions in the home sector relative to white men. Two factors underlie our estimate of $\hat{\delta}_{ig}$ being equal to 1 in the home sector for white women. First, we are estimating that white women were choosing the home sector because they were facing either some relative frictions or relatively less absolute advantage in other sectors. The average amount of frictions or absolute advantage differences white women face in the other sectors shows up in the average wage gap between white women and white men. Given that white women earned roughly 57 percent less when working than white men, our model predicts that women should be much more likely to work in the home sector relative to white men all else equal. Second, how much more white women should be working in the home sector if the other sectors are less attractive for white women depends on θ. As the skill distribution becomes less dispersed (θ increasing), frictions in other sectors will push more women into the home sector. The reason for this is that the comparative advantage in a given occupation relative to another occupation gets stronger when θ is higher. Given our estimate of θ, the observed wage gap between white men and white women, and the relative propensity of each group to be in the home sector, we estimate that $\hat{\delta}_{ig}$ is roughly one for white women in the home sector.

Moreover, as seen from Table 2, $\hat{\delta}_{ig}$ is close to 1 for white women in the home sector in all years of our analysis. This suggests that the decline of white women working in the home sector relative to white men is not because women experienced a decline in their absolute advantage in the home sector relative to white men or a decline in the relative frictions in the home sector relative to white men. Instead, our results suggest that the decline of white women in the home sector is due to white women experiencing declining frictions or increased absolute advantage in other market sectors (as measured by the change in the wage gap). Additionally, as we show below, changes in the productivity of the home sector relative to the market sector for all groups will also contribute to women exiting the home sector

Table 2: Estimated $\hat{\delta}_{ig}$ for White Women

	Year						Difference 1980 vs. 1960	Difference 2008 vs. 1980
	1960	1970	1980	1990	2000	2008		
Home	1.05	0.99	1.03	1.02	0.98	1.02	-0.02	-0.01
Executives/Managers	0.40	0.42	0.53	0.66	0.70	0.73	0.13	0.20
Engineers	0.20	0.22	0.32	0.43	0.46	0.50	0.12	0.19
Doctors	0.31	0.33	0.41	0.53	0.62	0.70	0.10	0.29
Lawyers	0.27	0.30	0.44	0.57	0.65	0.71	0.17	0.28
Teachers, Non-Postsecondary	0.74	0.71	0.81	0.94	1.07	1.11	0.07	0.30
Nurses	0.94	0.96	1.09	1.23	1.29	1.33	0.15	0.25
Secretaries	1.41	1.41	1.92	2.07	1.82	1.82	0.51	-0.11
Food Prep/Service	0.80	0.82	0.88	0.89	0.90	0.89	0.08	0.01
Construction	0.18	0.23	0.26	0.31	0.34	0.35	0.08	0.08
Firefighting	0.15	0.17	0.23	0.31	0.36	0.38	0.09	0.15
Vehicle Mechanic	0.17	0.23	0.25	0.32	0.34	0.34	0.09	0.09

Note: Author's calculations based on equation (13) using Census data and imposing $\theta = 3.11$ and $\eta = 1/4$.

over this time period. In order to quantify this effect, we need the general equilibrium analysis formulated in the next section. To preview our results, we find that the changing productivity of the home sector relative to the market sector explains roughly 60 percent of the movement of white women out of the home sector. The remaining 40 percent is due to changes in the $\hat{\delta}_{ig}$ in the market sector.

The remainder of the 1960 results from Table 2 reinforce that the $\hat{\delta}_{ig}$'s for white women changed dramatically over time in certain occupations. For example, our estimates of $\hat{\delta}_{ig}$ for executives, lawyers, doctors, and engineers for white women in 1960 ranged from 0.20 to 0.40. In terms of our sorting model, the low relative participation of white women in these occupations in 1960 results in the low values of $\hat{\delta}_{ig}$. The model is attributing the low propensity for white women to work in these occupations as being the result of either white women facing some friction to work in these occupations (either in the human capital market or the labor market directly) or that white women have a lower absolute ability to work in these occupations. Interestingly, the $\hat{\delta}_{ig}$ for white women teachers is appreciably less than one in 1960. While white women were 1.7 times more likely than white men to work as teachers, this propensity is more than offset by the overall wage gap in 1960, where women earned about 0.57 times what men earned. If white women were not facing some friction or lower absolute advantage in the teacher occupation, our model predicts there should have been an even higher fraction of white women ending up as teachers in 1960.

Contrast this with secretaries in 1960. A white woman in 1960 was 24 times more likely to work as a secretary than was a white man. The model can only explain this enormous discrepancy by assigning a $\hat{\delta}_{ig}$ of 1.4 for white women secretaries. Thus the model interprets these data patterns as either white women secretaries had a larger absolute advantage or there was discrimination against white *men* being secretaries. For example, if there were discriminatory norms in 1960 preventing white men from being secretaries, the model treats this as akin to a subsidy for white women in this occupation relative to white men. Also in 1960, white women had very low $\hat{\delta}_{ig}$ in the construction, firefighting and vehicle mechanic professions.

For executives, lawyers, and doctors, the $\hat{\delta}_{ig}$'s approximately doubled, rising from around 0.3 or 0.4 to around 0.7 or 0.8 between 1960 and 2008. School teachers also saw a substantial increase in their average $\hat{\delta}_{ig}$ from 0.74 to a value exceeding one. While the $\hat{\delta}_{ig}$'s in many skilled professional occupations increased, the $\hat{\delta}_{ig}$ for low skilled occupations did not change that much. This is particularly true post 1980. For example, the estimated $\hat{\delta}_{ig}$ for white women barely changed (or fell) for secretaries, waitresses, construction workers, and vehicle mechanics between 1980 and 2008. Yet, the $\hat{\delta}_{ig}$'s for executives, engineers, doctors, lawyers, teachers, and nurses continued to rise sharply during this time period. These results are consistent with the results above showing that the convergence in occupational sorting propensities was primarily among high skilled individuals.

The $\hat{\delta}_{ig}$'s for black men and black women for these same select occupations are shown in Table 3. A similar overall pattern emerges, with the $\hat{\delta}_{ig}$'s being substantially less than one in general in 1960 and rising appreciably through 2008, though typically remaining below one, especially for the high-skilled occupations. Unlike for white and black women, almost the entire change in the $\hat{\delta}_{ig}$ for black men occurred prior to 1980.

Table 4 shows the mean, the standard deviation, and the standard deviation relative to the mean for $\hat{\delta}_{ig}$ for each group over time. As seen from Table 4, not only did

Table 3: Estimated $\hat{\delta}_{ig}$ for Black Men and Women

	Year			Difference 1980 vs. 1960	Difference 2008 vs. 1980
	1960	1980	2008		
Black Men					
Home	0.91	1.06	1.08	0.15	0.02
Executives/Managers	0.47	0.66	0.73	0.19	0.07
Engineers	0.35	0.60	0.70	0.25	0.10
Doctors	0.49	0.61	0.67	0.12	0.06
Lawyers	0.48	0.58	0.65	0.10	0.07
Teachers, Non-Postsecondary	0.72	0.77	0.89	0.05	0.12
Nurses	0.60	0.75	0.85	0.15	0.10
Food Prep/Service	0.96	0.95	0.97	-0.01	0.02
Construction	0.68	0.80	0.76	0.12	-0.04
Firefighting	0.49	0.73	0.80	0.23	0.07
Vehicle Mechanic	0.67	0.77	0.77	0.09	0.00
Black Women					
Home	0.82	0.99	0.95	0.17	-0.04
Executives/Managers	0.25	0.47	0.65	0.22	0.18
Engineers	0.14	0.29	0.42	0.15	0.13
Doctors	0.21	0.35	0.58	0.14	0.23
Lawyers	0.18	0.38	0.59	0.20	0.21
Teachers, Non-Postsecondary	0.60	0.81	1.02	0.21	0.21
Nurses	0.70	1.03	1.22	0.34	0.18
Food Prep/Service	0.72	0.93	0.87	0.21	-0.06
Construction	0.18	0.26	0.35	0.08	0.08
Firefighting	0.00	0.21	0.36	0.21	0.15
Vehicle Mechanic	0.00	0.27	0.34	0.27	0.07

Note: Author's calculations based on equation (13) using Census data and baseline parameter values.

Table 4: Summary Measures for $\hat{\delta}_{ig}$ for White Women, Black Men, and Black Women

	Year						Difference 2008 vs. 1960
	1960	1970	1980	1990	2000	2008	
White Women							
Average Delta	0.65	0.63	0.70	0.77	0.80	0.83	0.18
Standard Deviation of Delta	0.35	0.31	0.33	0.31	0.26	0.27	-0.08
Standard Deviation of Delta/Average Delta	0.53	0.49	0.46	0.40	0.33	0.32	-0.21
Black Men							
Average Delta	0.72	0.78	0.85	0.87	0.87	0.87	0.16
Standard Deviation of Delta	0.20	0.17	0.17	0.17	0.17	0.15	-0.06
Standard Deviation of Delta/Average Delta	0.28	0.22	0.20	0.20	0.19	0.17	-0.12
Black Women							
Average Delta	0.48	0.55	0.68	0.75	0.77	0.78	0.29
Standard Deviation of Delta	0.29	0.29	0.31	0.31	0.27	0.27	-0.02
Standard Deviation of Delta/Average Delta	0.59	0.52	0.46	0.41	0.35	0.34	-0.25

Note: Author's calculations based on equation (13) using Census data and baseline parameter values. We weight all the data using the share of that occupations income out of the total wage bill.

the mean $\hat{\delta}_{ig}$ increase for all groups over time, the standard deviation of $\hat{\delta}_{ig}$ fell over time. When computing the productivity gains from changes in the δ_{ig}'s in Section 5, it is the standard deviation of the δ's relative to the mean of the δ's that drives misallocation. As seen from Table 4, this statistic has fallen sharply for all groups over time.[15]

3.4. Summary

In this section, we have empirically explored some of the predictions of our occupational sorting model. First, we have shown that wage gaps within an occupation between groups are unrelated to the relative propensity of the groups to be in an occupation. The relative propensities for a group to be in an occupation is a composite measure of differences faced by the groups in occupational frictions (either δ^h_{ig} or δ^w_{ig}) or differences in absolute advantage between the groups in the occupation (T_{ig}). Second, we compute this composite measure ($\hat{\delta}_{ig}$) for white women, black men, and black women over time during the last 50 years. The big declines in the composite occurred primarily in high-skilled occupations for white women. This

[15]We showing the mean and standard deviations of the $\hat{\delta}_{ig}$'s, we weight each occupation by their share of income earned in that occupation out of the total wage bill.

is reflected in the fact that the occupational similarity between high-skilled white men and high-skilled white women have converged to a much greater extent than between low-skilled white men and low-skilled white women.

4. Closing the Model

In order to evaluate the macroeconomic consequences of the changing allocation of talent, we must aggregate across the different occupations in some way. We choose a relatively natural approach and show that our general results are robust to the way we aggregate.

In particular, assume the N occupations combine in a CES fashion to produce a single aggregate output Y according to

$$Y = \left(\sum_{i=1}^{N} (A_i H_i)^{\rho} \right)^{1/\rho} \tag{15}$$

where H_i denotes the total efficiency units of labor employed in occupation i and A_i is the exogenously-given productivity of the occupation.

The total efficiency units of labor in each occupation are given by

$$H_i = \sum_{g=1}^{G} q_g \int h_{ijg} \epsilon_{ijg} dj. \tag{16}$$

To understand this equation, start from the right. First, we integrate over all people j in group g, adding up their efficiency units, which are the product of their human capital and their idiosyncratic ability. Next, there are q_g people belonging to group g. Finally, we add up across all the groups.

That completes the setup of the model. We can now define an equilibrium and then start exploring the model's aggregate implications.

4.1. Equilibrium

A competitive equilibrium in this economy consists of individual choices $\{c, e, s\}$, an occupational choice by each person, total efficiency units of labor in each occupation H_i, final output Y, and an efficiency wage w_i in each occupation such that

1. Given an occupational choice, the occupational wage w_i, and idiosyncratic ability ϵ in that occupation, each individual chooses c, e, s to maximize utility:

$$U(\delta, w, \epsilon) = \max_{c,e,s} (1-s)c^{\beta} \;\; s.t. \;\; c = \delta w \epsilon s^{\phi_i} e^{\eta} - e \tag{17}$$

   where $\delta \equiv \delta^h \delta^w$.

2. Each individual chooses the occupation that maximizes his or her utility: $i^* = \arg\max_i U(\delta_i, w_i, \epsilon_i)$, taking $\{\delta_i, w_i, \epsilon_i\}$ as given.

3. A representative firm chooses labor input in each occupation, H_i, to maximize profits:

$$\max_{\{H_i\}} \left(\sum_{i=1}^{N} (A_i H_i)^{\rho} \right)^{1/\rho} - \sum_{i=1}^{N} w_i H_i \tag{18}$$

4. The occupational wage w_i clears the labor market for each occupation:

$$H_i = \sum_{g=1}^{G} q_g \int h_{ijg} \epsilon_{ijg} dj \tag{19}$$

5. Total output is given by the production function in equation (15).

6. "Revenue" associated with the distortions equals zero for each occupation.

The equations characterizing the general equilibrium are then given in the next result.

Proposition 5 (Solving the General Equilibrium): *The general equilibrium of the model is $\{p_{ig}, H_i^{supply}, H_i^{demand}, w_i\}$ and Y such that*

1. p_{ig} satisfies equation (7).

2. *H_i^{supply} aggregates the individual choices:*

$$H_i^{supply} = \gamma\bar{\eta}w_i^{\theta-1}(1-s_i)^{(\theta(1-\eta)-1)/\beta}s_i^{\theta\phi_i}\sum_g q_g(\delta_{ig}^h)^\theta(\delta_{ig}^w)^{\theta-1}m_g^{1-\theta(1-\eta)} \tag{20}$$

3. *H_i^{demand} satisfies firm profit maximization:*

$$H_i^{demand} = \left(\frac{A_i^\rho}{w_i}\right)^{\frac{1}{1-\rho}} Y \tag{21}$$

4. *w_i clears each occupational labor market: $H_i^{supply} = H_i^{demand}$.*

5. *Total output is given by the production function in equation (15).*

5. Estimating Productivity Gains from Changing Occupational Sorting

5.1. Parameter values and Exogenous Variables

The key parameters of the model — assumed to be constant over time — are η, θ, ρ, and β. We discussed the estimation and assumptions for η and θ above. The parameter ρ governs the elasticity of substitution among our 67 occupations in aggregating up to final output. We have little information on this parameter and choose $\rho = 2/3$ for our baseline value. We will explore robustness to a wide range of values for ρ shortly.

The parameter β is the geometric weight on consumption relative to time in an individual's utility function (1). As schooling trades off time for consumption, the model implies that wages must increase more steeply with schooling in equilibrium when β is lower. Workers must be more heavily compensated for sacrificing time to schooling the more they care about time relative to consumption. To be specific, the average wage of group g in occupation i is proportional to $(1-s_i)^{\frac{-1}{\beta}}$. If we take a log linear approximation around average schooling $\bar{s}$, then β is inversely related to the Mincerian return to schooling across occupations (call this return ψ): $\beta = (\psi(1-$

$\bar{s}))^{-1}$. We calculate s as years of schooling divided by a pre-work time endowment of 25 years, and find the Mincerian return ψ from a regression of log wages on average occupation schooling, with group dummies as controls. We then set $\beta = 0.693$, the simple average of the implied β values across years. This method allows the model to approximate the Mincerian return to schooling across occupations, which averages 12.7% across the six decades. For robustness we will also entertain a lower value of $\beta = 0.5$ and a higher value of $\beta = 0.8$.

As our model is static, we infer exogenous variables separately by year. In each year, we have $6N$ variables to be determined. For each of the $i = 1, \ldots, N$ occupations these are A_i, ϕ_i, and δ_{ig}, where g stands for white men, white women, black women, or black men. We also allow population shares of each group q_i to vary by year to match the data. Finally, we normalize average ability to be the same in each occupation-group, or $T_{ig} = 1$: differences in mean ability across occupations are isomorphic to differences in the production technology A_i. Across groups, we think the natural starting point is *no* differences in mean ability; this assumption will be relaxed in our robustness checks.

To identify the values of these $6N$ forcing variables, we match the following $6N$ moments in the data, decade by decade (numbers in parentheses denote the number of moments):

- $(4N-4)$ The fraction of people from each group working in each occupation, p_{ig}. (Less than 4N moments because the p_{ig} sum to one for each group.)
- (N) The average wage in each occupation.
- (N) Zero total revenue from the discrimination "tax" in each occupation.
- (3) Wage gaps between white men and each of our 3 other groups.
- (1) Average years of schooling in one occupation.

As discussed above, the δ_{ig} variables are easy to identify in the data given our setup. But recall that $\delta^h_{ig} \equiv (1+\tau^h_{ig})^(-\eta)$. From the data we currently have, we cannot separately identify the δ^h and δ^w components of δ. That is, we cannot distinguish between barriers to accumulating human capital and labor market barriers. We proceed by considering two polar cases. At one extreme, we assume all of the δ^w_{ig}'s are one, so that $\delta_{i,g}$'s solely reflect δ^h's. At the other extreme, we set all of the $\delta^h_{ig} = 1$ and assume the δ^w_{ig}'s are responsible for the δ's. i.e., we assume only human capital barriers (the δ^h case) or only labor market barriers (the δ^w case).

The A_i levels and the relative ϕ_i's across occupations involve the general equilibrium solution of the model.

Recall from equation (10) that wages are increasing in schooling across occupations. From Proposition 1, we know that schooling increases with ϕ_i. Thus we can infer from wages in each occupation the *relative* values of ϕ_i across occupations. But we cannot pin down the ϕ_i levels, as wage levels are also affected by the A_i productivity parameters. Thus we use the final moment – average years of schooling in one occupation – to determine the ϕ_i levels. We choose to match schooling in the lowest wage occupation, which is Farm Non-Managers. Calling this the min occupation, we set ϕ_{min} in a given year to match the observed average schooling among Farm Non-Managers in the same year: $\phi_{min} = \frac{1-\eta}{\beta} \frac{s_{min}}{1-s_{min}}$.

5.2. Productivity Gains

Given our model, parameter values, and the forcing variables we infer from the data, we can now answer one of the key questions of the paper: how much of overall earnings growth between 1960 and 2008 can be explained by the changing δ frictions?

In answering this question, the first thing to note is that output growth in our model is a weighted average of earnings growth in the market sector and in the home sector. Earnings growth in the market sector can be measured as real earnings growth in the census data. Deflating by the NIPA Personal Consumption Deflator, real earnings in the census data grew by 1.32 percent per year between 1960 and 2008.[16] For the home sector, we impute wages from the relationship between average education and average earnings across market sectors and from wage gaps by group in market sectors. (See the discussion in section 3.1. for additional details.)

[16]This might be lower than standard output growth measures because it is calculated solely from wages; for example, it omits employee benefits.

Taking a weighted average of the imputed wage in the home sector and the wage in the census data, we estimate that output (as defined by our model) grew by 1.47 percent per year between 1960 and 2008.

How much of this growth is accounted for by changing δ's, according to our model? We'd like to answer this question by holding the A's (productivity parameters by occupation), ϕ's (schooling parameters by occupation), and q's (group shares of the working population) constant over time and letting the δ's change. At which year's value should we hold the A's, ϕ's, and q's constant? We follow the standard approach in macroeconomics and use *chaining* to answer our question. That is, we compute growth between 1960 and 1970 allowing the δ's to change but holding the other parameters at their 1960 values. Then we compute growth between 1960 and 1970 from changing δ's holding the other parameters at their 1970 values. We take the geometric average of these two estimates of growth from changing δ's. We do the same for other decadal comparisons (1970 to1980 and so on) and cumulative the growth to arrive at an estimate for our entire sample from 1960–2008.

The results of this calculation are shown in Table 5. When the frictions are interpreted as occurring in human capital accumulation (the δ^h case), this calculation indicates that the change in occupational frictions contributed an average of 0.304 percentage points to growth per year. This would explain 20.6 percent of overall earnings growth over the last half century.

If we instead interpret the frictions as occurring in the labor market (the δ^w case), chain-weighted growth from changing δ^w's is 0.238 percent per year. According to this case, changing labor market frictions account for 16.2 percent of the cumulative earnings growth from 1960 to 2008.

A related calculation, perhaps more transparent, is to hold the δ's constant and calculate the hypothetical growth rate due to the changes in the A's, ϕ's, and q's. Figure 3 plots the results of this calculation for the δ^h case. The left panel considers the case when the occupational frictions are held constant at 1960 levels; the right panel presents the case when the δ^h's are kept at 2008 levels. Holding the δ^h's fixed at their 1960 level, output in 2008 would be 15.4 percent lower than it actually was. Holding the δ's fixed at their 2008 level, output in 1960 would be 12.7 percent higher

Table 5: Productivity Gains

	δ^h calibration	δ^w calibration
Average annual wage growth	*1.47*	
Chain-weighted wage growth due to changing frictions (annualized)	0.304	0.238
(Percent of total)	(20.6%)	(16.3%)
Robustness: No frictions in "brawny" occupations		
Chain-weighted wage growth due to changing frictions (annualized)	0.270	0.199
(Percent of total)	(18.3%)	(13.5%)

Note: Italicized entries in the table are data; non-italicized entries are results from the model. For the robustness panel in the table, we assume that there are no frictions for white women in occupations where physical strength is important. Instead, we allow $T_{i,ww}$ to change over time to match the occupational allocation for white women. For blacks in this case, we do allow for frictions, but also assume $T_{i,bw} = T_{i,ww}$.

Figure 3: Counterfactuals in the δ^h case

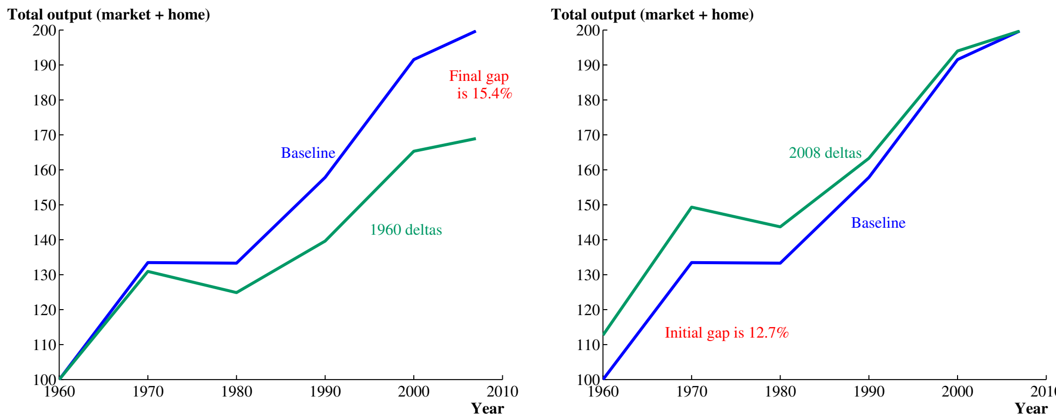

Note: The left panel shows the counterfactual path of output in the model if the δ^h's were kept at their 1960 values in every period. The right panel shows the counterfactual where the δ^h's are kept at their 2008 values.

than in the data.

Figure 4 presents similar estimates, this time for the δ^w case. Here, holding the δ^w's fixed at their 1960 level would result in output being 12.0 lower in 2008. Holding the δ^w's fixed at their 2008 level, output in 1960 would be 10.3 percent higher. Note that both models predict output growth would have been negative in the 1970s in the absence of the reduction in δ's for blacks and women in that decade. But over the entire sample, the bulk of growth is due to rising A's and ϕ's.

It is worth elaborating on the gains from the changing δ's. To this end, Figure 5 presents the mean and variance of δ^h over time for each group in the δ^h case. The left panel shows the average δ's rising for women and African-Americans, whereas the average for white males starts above one and falls (almost imperceptibly) toward one. According to the model, these δ's led white men to overinvest in human capital and blacks and women to underinvest in human capital (presuming $\eta > 0$ so that higher earnings induce more human capital investment). Over time these gaps in average δ's diminished, leading to a better allocation of human capital investment in 2008 than in 1960 and therefore some efficiency gains.

The right panel of Figure 5 shows that the δ's were also more dispersed across

Figure 4: Counterfactuals the δ^w case

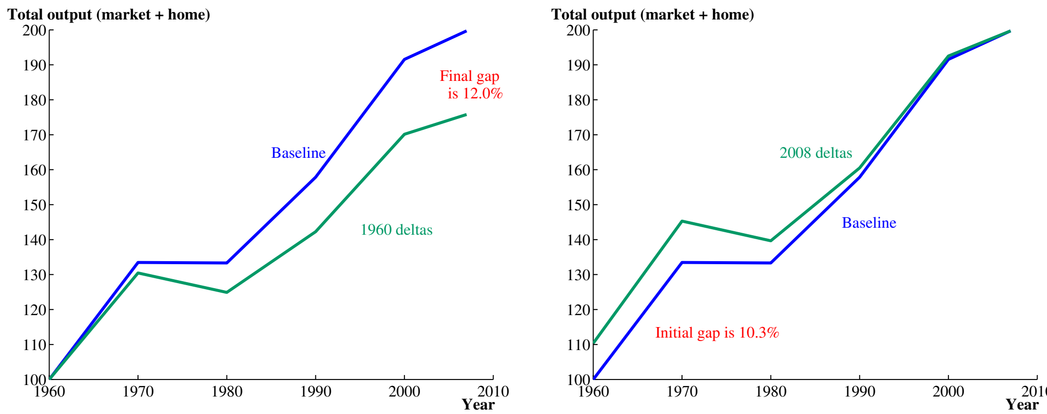

Note: The left panel shows the counterfactual path of output in the model if the δ^w's were kept at their 1960 values in every period. The right panel shows the counterfactual where the δ^w's are kept at their 2008 values.

occupations for blacks and women than for white men. It is this dispersion that leads to misallocation of talent across occupations. If there were no dispersion in the δ's across occupations for each group, there would be no misallocation of talent. All groups would have the same occupational distributions. The dispersion in the δ's leads to different occupational choices for each group – a misallocation if the distribution of talent is the same in each group. The falling variance of the log δ's leads to a better allocation of talent and hence some productivity growth.[17]

Could the productivity gains we estimate be inferred from a back-of-the-envelope calculation involving the wage gaps alone? In particular, suppose one takes white male wage growth as fixed, and calculates how much of overall wage growth comes from the faster growth of wages for the other groups. The answer is that faster wage growth for blacks and white women contributed 0.32 percentage points per year to overall wage growth from 1960 to 2008. This is strikingly close to our estimate of the productivity gain from changing δ^h's of 0.30 percent per year, and not that far away from the 0.23 percent per year growth in the δ^w case. But the similarity is more coin-

[17] Even for women, most of the decline in the variance can be seen between market occupations. But a notable portion of the overall decline for women comes from falling barriers in the market occupations relative to the home sector.

Figure 5: Means and Variances of the Frictions, δ^h Calibration

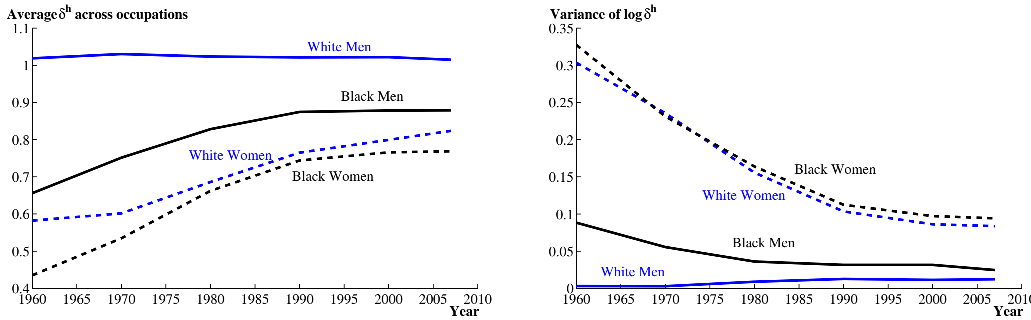

Note: The left panel shows the average level of δ^h by group, weighted by total earnings in each occupation. The right panel shows the variance of $\log \delta^h$, weighted in the same way.

cidental than fundamental. The reason is that white male wages are not exogenous with respect to the earnings of other groups, at least in the model.

As women and blacks enter previously white male-dominated occupations, this drives down the wages of white men remaining in those occupations. The wage gains to women and blacks, therefore, come (partly) at the expense of white males. In addition to this general equilibrium force, recall that we normalize the mean δ across groups to zero in each occupation. Thus rising δ's for women and blacks go hand in hand with falling δ's for white men. This is a direct force lowering wages of white men as barriers to blacks and women fall. It operates through human capital in the δ^h case and by hitting wages directly in the δ^w case (as δ's drive a wedge between wages and marginal products).

To drive home that wage gaps can be decoupled from productivity gains in the model, consider the δ^w case and $\eta = 0$. And imagine either all workers have the same ability or δ^w's are the same across occupations for each group. In this scenario, δ_w trends can entirely explain the path of wage gaps while contributing nothing to aggregate productivity growth. There are no gains from reallocation of workers across occupations (as everyone has the same talent or is in the right occupation already), and there are no gains from better allocation of human capital investments (as there are none). There is only redistribution of wages from white men to other groups.

The upshot is that model productivity gains cannot be gleaned from the wage gaps alone. The model's productivity gains can stray even farther from the back-of-the-envelope number under alternative parameter values (such as a lower η so that human capital investment is less important, and a higher θ so that talent is less disperse). We will see this in the next subsection.

5.3. Robustness

The final row in Table 5 considers the robustness of our productivity gains to relaxing the assumption that men and women draw from the same distribution of talent in all occupations. In particular, we consider the possibility that some occupations rely more on physical strength than others, and that this reliance might have changed because of technological progress. For this check, we go to the extreme in assuming that frictions for white women are completely absent from the set of occupations where physical strength is arguably important, including firefighters, police officers, and most of manufacturing.[18] That is, we estimate values for $T_{i,g}$ for white women that completely explain their observed allocation to these occupations for every period between 1960 and 2008. Our hypothesis going into this check was that most of the productivity gains were coming from the rising propensity for women to enter occupations like lawyers, doctors, scientists, professors, and managers, where physical strength is not important. Indeed, the results in Table 5 support this hypothesis. The amount of wage growth explained by changing frictions falls only slightly — for example, from 20.6% to 18.3% in the δ^h calibration.

How sensitive is the growth contribution of changing δ's to our chosen parameter choices? Tables 6 and 7 explore robustness to different parameter values. For each alternative set of parameter values, we the recalculate the δ_{ig}, A_i, and ϕ_i values so that the model continues to fit the occupation shares and wage gaps.

The first row checks sensitivity to the elasticity of substitution (ρ) between occupations in production. Under the δ^h case, the share of growth explained ranges from 16.6 percent when the occupations are almost Leontief ($\rho = -90$) to 23.6 per-

[18]These occupations are assigned based on Rendall (2010).

Table 6: Robustness Results: Percent of Growth Explained in the δ^h case

Base: ρ = 2/3	ρ = -90	ρ = -1	ρ = 1/3	ρ = 0.95
20.6%	16.6%	17.7%	19.3%	23.5%

Base: η = 0.25	η = 0.005	η = 0.01	η = 0.05	η = 0.1	η = 0.5
20.6%	15.5%	16.2%	18.8%	19.9%	19.3%

	η = 0.25		
θ(1-η) = 3.11	θ(1-η) = 2	θ(1-η) = 15	θ(1-η) = 15 with ρ = -90
20.6%	25.0%	16.4%	13.0%

Base: β = 0.693	β = 0.5	β = 0.8
20.6%	20.6%	20.6%

Note: Entries in the table represent the share of earnings growth that is explained by the changing δ^h's using the chaining approach. Each entry changes one of the parameter values relative to our baseline case.

cent when they are almost perfect substitutes ($\rho = 0.95$). This compares to 20.6 percent with our baseline value of $\rho = 2/3$. Outcomes are more sensitive under the δ^w case, with the share of growth explained by changing δ^w's going from 10.7 to 19.7 percent (vs. 16.2 percent baseline). Note that gains are increasing in substitutability.

The second row considers different values of the elasticity of human capital with respect to goods invested in human capital (η). The gains are generally increasing in η. In the δ^h case, the gains range from 15.5 percent to 20.6 percent as η rises from 0.005 to 0.25.[19] Gains are less sensitive to η under the δ^w case, rising from 15.7 percent to 16.4 percent as η goes from 0 to 0.5.

The third row indicates that the gains from changing δ's shrinks as θ rises. As $\theta(1-\eta)$ rises from 2 to 15, gains fall from 25.0 percent to 16.4 percent in the δ^h case, and from 17.9 percent to 13.1 percent in the δ^w case. When people are more similar in ability, lower barriers are required to explain why women and blacks were under-

[19]We must have $\eta > 0$ in the δ^h case as the only source of wage and occupation differences across groups is different human capital investments in this case.

Table 7: Robustness Results: Percent of Growth Explained in the δ^w case

Base: ρ = 2/3	ρ = -90	ρ = -1	ρ = 1/3	ρ = 0.95
16.2%	10.7%	12.3%	14.4%	19.7%

Base: η = 0.25	η = 0	η = 0.1	η = 0.5
16.2%	15.7%	15.9%	16.4%

	η = 0.25		
θ(1-η) = 3.11	θ(1-η) = 2	θ(1-η) = 15	θ(1-η) = 15 with ρ = -90
16.2%	17.9%	13.1%	4.5%

Base: β = 0.693	β = 0.5	β = 0.8
16.2%	16.2%	16.2%

Note: Entries in the table represent the share of earnings growth that is explained by the changing δ^w's using the chaining approach. Each entry changes one of the parameter values relative to our baseline case.

represented in high skill occupations. As the barriers drop, the women and blacks who replace previously privileged white males are not too far apart in terms of ability, limiting the gains from reallocation. The gains are smaller still if occupations are skill dispersion is narrow and occupations are complementary ($\theta(1-\eta) = 15$ and $\rho = -90$): 13.0 percent in the δ^h case and 4.9 percent in the δ^w case.

The final row of the robustness Tables shows that the results are insensitive to the weight placed on time vs. goods in utility (β).

5.4. Further Results

Here we describe a number of additional insights from the model.

In the Census data, the share of women working in the market rose from 32.9 percent in 1960 to 69.2 percent in 2008. One explanation is that women's market opportunities rose, say due to declining discrimination or better information. See Jones, Manuelli and McGrattan (2003), Albanesi and Olivetti (2009), and Fogli and Veldkamp (2011) for empirical analysis of these hypotheses. As Table 2 showed,

Table 8: Female Participation Rates

	δ^h calibration		δ^w calibration
Women's LF participation	*1960=0.329*		*2008=0.692*
Change, 1960 – 2008		*0.364*	
Change, Model	0.146		0.146
(Percent of total)	(40.1%)		(40.3%)

Note: Results are for white women and black women combined. Participation is defined as working in market occupations. The sampling weight of part-time workers is split evenly between the market sector and the home sector. Italicized entries in the table are data; non-italicized entries are results from the model.

the δ's rose in market occupations relative to the home sector for women. How much of the rising female labor-force participation rate can be traced to changing δ's? Table 8 provides the answer. Of the 36.3 percentage point increase, the changing $\delta's$ contributed 14.6 percentage points, or around 40 percent of the total increase. According to our model, the remaining 60 percent can be attributed to changes in technology such as the A's. This is in the spirit of the work by Greenwood, Seshadri and Yorukoglu (2005) on "engines of liberation".

As we report in Table 9, gaps in average years of schooling narrowed from 1960 to 2008 for all three groups vs. white males: by 0.41 years for white women, 1.81 years for black men, and 1.55 years for black women. If the δ's for blacks and women rose faster in occupations with above-average schooling, then the changing δ's contributed to this educational convergence. The Table indicates how much. For white women, the changing δ's account the trend and then some (0.59-0.62 years, vs. 0.41 in the data). For black men, falling frictions might have narrowed the schooling gap with white men by 0.62-0.63 years, about one-third of the convergence in the data. For black women, declining distortions might explain 70 percent (1.10/1.55) of their catch-up in schooling.

How much of the productivity gains reflect changes in the occupational frictions facing women vs. those facing blacks? Tables 10 and 11 provide the answer for δ^h

Table 9: Education Predictions, All Households (Age 25–55)

	1960 Level	2008 Level	Change 2008 - 1960	Change Relative to WM	Model Predictions τ_w	τ_h
White Men	11.11	13.47	2.35			
White Women	10.98	13.75	2.77	0.41	0.62	0.59
Black Men	8.56	12.73	4.17	1.81	0.63	0.62
Black Women	9.24	13.15	3.90	1.55	1.11	1.09

Note: Author's calculations using Census data.

and δ^w, respectively. The second column presents the overall wage growth for each time period. The third column replicates the estimates (already shown in Figures 3 and 4) of setting the δ's to their levels at the end of each period (1960–1980, 1980–2008, and 1960–2008 for Rows 1, 2, and 3). Take the δ^h case. Almost two-thirds (13.0 out of 20.6) of the total gains from reduced occupational frictions over the last fifty years can be explained by the changes facing white women. Falling frictions faced by blacks accounted for two-fifths of the gains. We expect

The share of gains associated with falling frictions for white women vs. blacks differs across the time periods. Again, consider the δ^h case. Blacks accounted for a larger share of the gains in the 1960s and 1970s than in later decades. From 1960 to 1980, reduced frictions for blacks account for a quarter of the overall gains from reduced frictions. From 1980 to 2008, reduced frictions for blacks account for less than one-tenth of the overall gains. This timing might link the gains for blacks to the Civil Rights movement of the 1960s.

What was the consequence of shifting occupational frictions for the wage growth of different groups? Tables 12 and 13 try to answer this question. The first column

Table 10: Contribution of Each Group to Total Earnings Growth, δ^h case

		Percent of Growth Explained			
Year	Base Model Growth	Setting All δ's to End Levels	Setting WW δ's to End Levels	Setting BM δ's to End Levels	Setting BW δ's to End Levels
1960-1980	34.7 percent	21.9%	9.2%	2.1%	3.3%
1980-2008	46.7 percent	19.7%	15.7%	0.6%	1.7%
1960-2008	97.6 percent	20.6%	13.0%	1.3%	2.4%

Note: Author's calculations using Census data and baseline parameter values.

Table 11: Contribution of Each Group to Total Earnings Growth, δ^w case

		Percent of Growth Explained			
Year	Base Model Growth	Setting All δ's to End Levels	Setting WW δ's to End Levels	Setting BM δ's to End Levels	Setting BW δ's to End Levels
1960-1980	33.3 percent	22.0%	14.4%	2.2%	3.7%
1980-2008	49.8 percent	12.0%	10.8%	0.4%	1.4%
1960-2008	99.7 percent	16.2%	12.3%	1.2%	2.4%

Note: Author's calculations using Census data and baseline parameter values.

Table 12: Group Changes in Wages, δ^h case

Group	Wage Growth	Percent Explained By Changing δ's: η = 0.25	Percent Explained By Changing δ's: η = 0.005
White Men	77.0 percent	-3.4%	-9.1%
White Women	126.3 percent	39.9%	34.9%
Black Men	143.0 percent	45.3%	41.6%
Black Women	143.0 percent	57.3%	53.7%

Note: Author's calculations using Census data and baseline parameter values.

presents the actual growth of real wages for the different groups from 1960 to 2008. Real wages increased by 77 percent for white men, 126 percent for white women, and 143 percent for both black men and black women. For brevity, consider the δ^h case. In the absence of the change in occupational frictions, the model says real wages for white men would have been 3 percent higher. Put differently, real income of white men declined due to the changing opportunities for blacks and women. But at the aggregate level, this loss was swamped by the wage gains for blacks and women. Almost 40 percent of the wage growth for white women was due to the change in occupational frictions. For blacks, around half of their earnings growth might be attributable to the increased opportunities they faced. The model explains the remainder of growth as resulting from changes in technology (A's) and skill requirements (ϕ's).

Tables 14 and 15 look at the regional dimension of the decline in frictions confronting blacks and women. Here, we assume that workers are immobile across regions. With this assumption, a decline in occupational frictions in the South relative to the North will increase average wages in the South relative to the North. From 1960 to 2008, wages in the South increased by 10 percent relative to wages in

Table 13: Group Changes in Wages, δ^w case

Group	Wage Growth	Percent Explained By Changing δ's: $\eta = 0.25$	$\eta = 0$
White Men	77.0 percent	-9.6%	-10.1%
White Women	126.3 percent	35.2%	33.6%
Black Men	143.0 percent	41.8%	41.0%
Black Women	143.0 percent	53.8%	52.6%

Note: Author's calculations using Census data and baseline parameter values.

the Northeast. In the δ^h case, about 5 percentage points of this convergence was due to reduced occupational frictions facing blacks and women in the South relative to the Northeast — with the bulk of the effect due to rising δ's for blacks.

From 1980 to 2008, we see a reversal of the North-South convergence, perhaps driven by the reverse migration of blacks to the U.S. South. Reverse migration is what one would expect to see if workers are responding to the improved labor market outcomes in the South by relocating to the South. In a long run with higher labor mobility, the main effect of declining occupational frictions for blacks in the South relative to the North might be to increase the number of blacks living in the South relative to the North. Persistent wage gaps might reflect skill differences between regions. Of course, to the extent mobility is costly even in the long run, frictions can contribute to wage gap differences across regions even in the long run.

Table 14: Contributions to Northeast - South Convergence, δ^h case

Year	Base Model Convergence in P. Points	Percentage Points of Growth Explained: Setting All δ's to End Levels	Setting BM and BW δ's to End Levels
1960-1980	20.7	2.8	3.3
1980-2008	-16.5	0.1	1.6
1960-2008	10.0	3.2	5.0

Note: Author's calculations using Census data and baseline parameter values.

Table 15: Contributions to Northeast - South Convergence, δ^w case

Year	Base Model Convergence in P. Points	Percentage Points of Growth Explained: Setting All δ's to End Levels	Setting BM and BW δ's to End Levels
1960-1980	20.7	1.9	2.3
1980-2008	-16.5	0.0	1.0
1960-2008	10.0	2.0	3.3

Note: Author's calculations using Census data and baseline parameter values.

5.5. Human Capital of Workers by Occupation

Using equation (9), the amount of human capital per worker — including innate ability — for group g in occupation i is given by

$$\frac{H_{ig}}{q_g p_{ig}} = \gamma \bar{\eta} \cdot \frac{1}{\delta^w_{ig} w_i} \cdot (1 - s_i)^{-1/\beta} \cdot m_g. \tag{22}$$

The amount of human capital per worker for a group *relative* to white men is therefore

$$\frac{H_{ig}/q_g p_{ig}}{H_{i,wm}/q_{wm} p_{i,wm}} = \frac{\delta^w_{i,wm}}{\delta^w_{ig}} \cdot \frac{\overline{\text{wage}}_g}{\overline{\text{wage}}_{wm}}. \tag{23}$$

That is, relative quality in an occupation is simply the wage gap divided by the occupational frictions.

Notice that in the δ^h case (where the δ^w variables are set to one), equation (23) implies that the amount of human capital per worker for a group relative to white men is *the same across all occupations*. In particular, relative quality is precisely equal to the wage gap.

Figure 6 shows the average amount of human capital per worker for white women for select occupations. These measures are shaped by several forces. First, there is the general rise in human capital over time. These forces are especially apparent for the δ^h case.

Second, however, are the substantial selection effects that occur as the δ's change, especially in the δ^w case. For example, human capital per worker among women doctors has declined over time in this version: in 1960, only the most able women became doctors, according to the model, while in 2008 far less able women have entered this profession, lowering the overall human capital per doctor among women. For teachers and managers, this same selection effect is roughly offset by the general rise in educational attainment, leading to a relatively stable amount of human capital per female worker.

For white men, the results are shown in Figure 7. In both cases, the general rise in human capital per worker is dominant. But the selection effects are also apparent in some occupations, such as school teachers, for example.

Figure 6: Human Capital per Worker for White Women

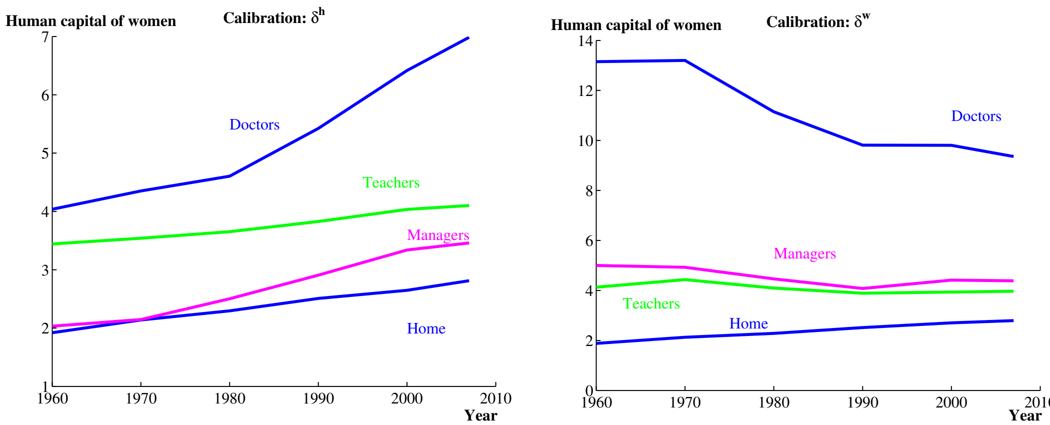

Note: Average quality (human capital and innate ability) in various occupations for white women, in the δ^h and δ^w cases. Computed using Census data, equation (22), and baseine parameter values.

Figure 7: Human Capital per Worker for White Men

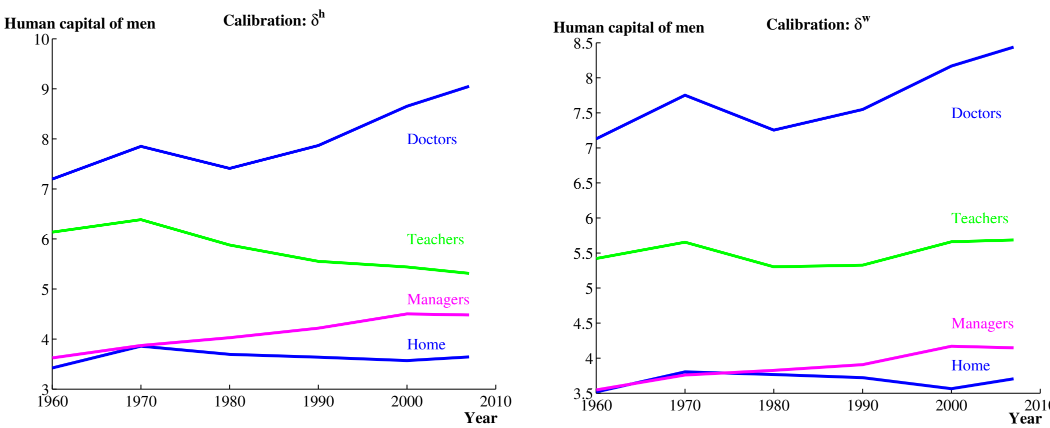

Note: Average quality (human capital and innate ability) in various occupations for white men, in the δ^h and δ^w cases. Computed using Census data, equation (22), and baseline parameter values.

Figure 8: Relative Human Capital per Worker, White Women vs. White Men

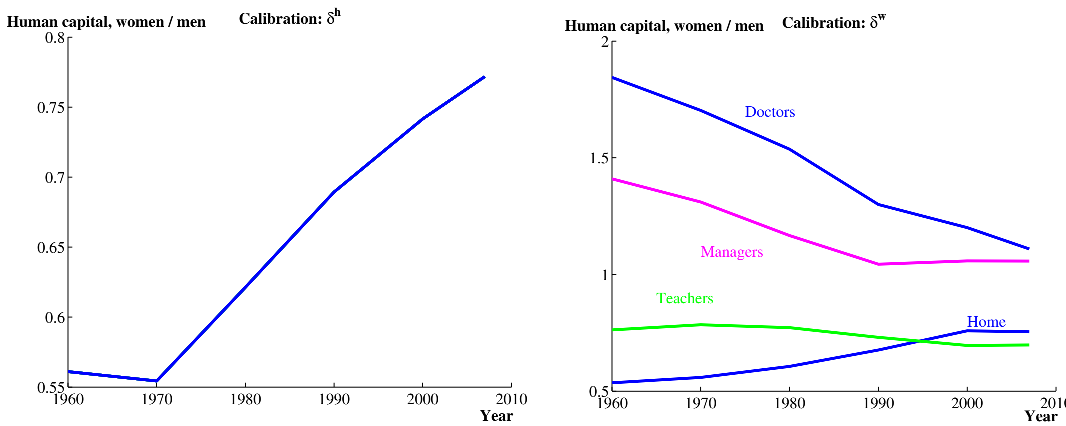

Note: Relative quality (human capital and innate ability) in various occupations for white women versus white men, in the δ^h and δ^w cases. Computed using equation (23).

Figure 8 shows the relative amounts of human capital between white women and white men for these same occupations, as in equation (23). In the δ^h case, as mentioned above, relative qualities are equated for all occupations. The graph shows that the relative quality of women to men in each occupation rose substantially between 1960 and 2008, from 0.56 to 0.77.

The δ^w case presents a very different view of the data. Relative qualities are not the same across occupations, as shown in the right panel. In 1960, human capital per worker was substantially higher for women relative to men for doctors and managers. Only the most talented women overcame frictions to become doctors and managers in 1960, and some lesser talented white men entered these professions instead. According to this case, the difference in quality has faded substantially over time due to declining frictions, but remains present even in 2008.

Of course, the real world likely reflect forces from both the δ^h and the δ^w cases. To this end, independent information on quality trends for occupation-groups could be quite helpful in discriminating between the models empirically (or quantifying their relative contribution).

6. Conclusion

Under construction

A Derivations and Proofs

The propositions in the paper summarize the key results from the model. This appendix shows how to derive the results.

Proof of Proposition 1. Individual Consumption and Schooling

Proposition 1 comes directly from the first order conditions for the individual's optimization problem.

Proof of Proposition 2. Occupational Choice

As given in Proposition 1, the individual's utility from choosing a particular occupation is $U(\delta_i, w_i, \epsilon_i) = \bar{\eta}^{\beta}(\tilde{\delta}_{ig}\epsilon_i)^{\frac{\beta}{1-\eta}}$, where $\tilde{\delta}_{ig} \equiv \delta_{ig} w_i s_i^{\phi_i}(1-s_i)^{\frac{1-\eta}{\beta}}$. The solution to the individual's problem, then, involves picking the occupation with the largest value of $\tilde{\delta}_{ig}\epsilon_i$. To keep the notation simple, we will suppress the g subscript in what follows.

Let p_i denote the probability that the individual chooses occupation i. Then

$$\begin{aligned} p_i &= \Pr\left[\tilde{\delta}_i\epsilon_i > \tilde{\delta}_s\epsilon_s\right] \ \forall i \neq s \\ &= \Pr\left[\epsilon_s < \tilde{\delta}_i\epsilon_i/\tilde{\delta}_s\right] \ \forall s \neq i \\ &= \Pi_{s\neq i} F_s(\tilde{\delta}_i\epsilon_i/\tilde{\delta}_s) \end{aligned} \tag{24}$$

if ϵ_i is known for certain. Since it is not, we must also integrate over the probability distribution for ϵ_i:

$$p_i = \int \Pi_{s\neq i} F_s(\tilde{\delta}_i\epsilon_i/\tilde{\delta}_s) f_i(\epsilon_i) d\epsilon_i, \tag{25}$$

where $f_i(\epsilon) = \theta T_i \epsilon^{-(1+\theta)} \exp\{-T_i\epsilon^{-\theta}\}$ is the pdf of the Fréchet distribution. Substituting in for the distribution and pdf, additional algebra leads to

$$
\begin{aligned}
p_i &= \int \theta T_i \left(\Pi_{s\neq i} \exp\{-T_s(\tilde{\delta}_i\epsilon_i/\tilde{\delta}_s)^{-\theta}\}\right) \epsilon_i^{-(1+\theta)} \exp\{-T_i\epsilon^{-\theta}\} d\epsilon_i \\
&= \int \theta T_i \epsilon_i^{-(1+\theta)} \exp\left\{-\sum_{s=1}^{N} T_s \left(\frac{\tilde{\delta}_i}{\tilde{\delta}_s}\right)^{-\theta} \epsilon_i^{-\theta}\right\} d\epsilon_i. \qquad (26)
\end{aligned}
$$

Now, define $\bar{T}_i \equiv -\sum_{s=1}^{N} T_s \left(\frac{\tilde{\delta}_i}{\tilde{\delta}_s}\right)^{-\theta}$. Then the probability simplifies considerably:

$$
\begin{aligned}
p_i &= \frac{T_i}{\bar{T}_i} \int \theta \bar{T}_i \epsilon_i^{-(1+\theta)} \exp\{-\bar{T}_i\epsilon_i^{-\theta}\} d\epsilon_i \\
&= \frac{T_i}{\bar{T}_i} \int d\bar{F}_i(\epsilon_i) \\
&= \frac{T_i}{\bar{T}_i} \\
&= \frac{T_i\tilde{\delta}_i^{\theta}}{\sum_s T_s\tilde{\delta}_s^{\theta}} \qquad (27)
\end{aligned}
$$

where $\bar{F}_i(\epsilon)$ is the cdf of a Fréchet distribution with parameters $\bar{T}_i$ and θ. The first main result in Proposition 2 then comes from our normalization that $T_i = 1$ for all i.

Total efficiency units of labor supplied to occupation i by group g are

$$H_{ig} = q_g p_{ig} \cdot \mathbb{E}\left[h_i\epsilon_i \mid \text{Person chooses } i\right].$$

Recall that $h(e,s) = s^{\phi_i}e^{\eta}$. Using the results from Proposition 1, it is straightforward to show that

$$h_i\epsilon_i = \tilde{h}_i(\delta_i w_i)^{\frac{\eta}{1-\eta}} \epsilon_i^{\frac{1}{1-\eta}},$$

where $\tilde{h}_i \equiv \eta^{\eta/(1-\eta)} s_i^{\frac{\phi_i}{1-\eta}}$. Therefore,

$$H_{ig} = q_g p_{ig} \tilde{h}_i(\delta_i w_i)^{\frac{\eta}{1-\eta}} \cdot \mathbb{E}\left[\epsilon_i^{\frac{1}{1-\eta}} \mid \text{Person chooses } i\right]. \qquad (28)$$

To calculate this last conditional expectation, we use the extreme value magic of the Fréchet distribution. Let $y_i \equiv \tilde{\delta}_i\epsilon_i$ denote the key occupational choice term.

Then

$$y^* \equiv \max_i \{y_i\} = \max_i \{\delta_i \epsilon_i\} = \delta^* \epsilon^*.$$

Since y_i is the thing we are maximizing, it inherits the extreme value distribution:

$$\begin{aligned}
\Pr[y^* < z] &= \Pi_{i=1}^N \Pr[y_i < z] \\
&= \Pi_{i=1}^N \Pr[\tilde{\delta}_i \epsilon_i < z] \\
&= \Pi_{i=1}^N \Pr[\epsilon_i < z/\tilde{\delta}_i] \\
&= \Pi_{i=1}^N \exp\left\{-T_i \left(\frac{z}{\tilde{\delta}_i}\right)^{-\theta}\right\} \\
&= \exp\{-\sum_{i=1}^N T_i \tilde{\delta}_i^\theta \cdot z^{-\theta}\} \\
&= \exp\{-\bar{T} z^{-\theta}\}. \qquad (29)
\end{aligned}$$

That is, the extreme value also has a Fréchet distribution, with a mean-shift parameter given by $\bar{T} \equiv \sum_s T_i \tilde{\delta}_i^\theta$.

Straightforward algebra then reveals that the distribution of ϵ^*, the ability of people in their chosen occupation, is also Fréchet:

$$G(x) \equiv \Pr[\epsilon^* < x] = \exp\{-T^* x^{-\theta}\} \qquad (30)$$

where $T^* \equiv \sum_{i=1}^N T_i \left(\tilde{\delta}_i/\tilde{\delta}^*\right)^\theta$.

Finally, one can then calculate the statistic we needed above back in equation (28): the expected value of the chosen occupation's ability raised to some power. In particular, let i denote the occupation that the individual chooses, and let α be some positive exponent. Then,

$$\begin{aligned}
\mathbb{E}[\epsilon_i^\lambda] &= \int_0^\infty \epsilon^\lambda dG(\epsilon) \\
&= \int_0^\infty \theta T^* \epsilon^{-(1+\theta)+\lambda} e^{-T^* \epsilon^{-\theta}} d\epsilon \qquad (31)
\end{aligned}$$

Recall that the "Gamma function" is $\Gamma(\alpha) \equiv \int_0^\infty x^{\alpha-1} e^{-x} dx$. Using the change-of-variable $x = T^* \epsilon^{-\theta}$, one can show that

$$\begin{aligned}\mathbb{E}[\epsilon_i^\lambda] &= T^{*\lambda/\theta} \int_0^\infty x^{-\lambda/\theta} e^{-x} dx \\ &= T^{*\lambda/\theta} \Gamma(1 - \lambda/\theta). \end{aligned} \tag{32}$$

Applying this result to our model, we have

$$\begin{aligned}\mathbb{E}\left[\epsilon_i^{\frac{1}{1-\eta}} \mid \text{Person chooses } i\right] &= T^{*\frac{1}{\theta}\cdot\frac{1}{1-\eta}} \, \Gamma\left(1 - \frac{1}{\theta}\cdot\frac{1}{1-\eta}\right) \\ &= p_{ig}^{-\frac{1}{\theta}\cdot\frac{1}{1-\eta}} \, \Gamma\left(1 - \frac{1}{\theta}\cdot\frac{1}{1-\eta}\right). \end{aligned} \tag{33}$$

Substituting this expression into (28) and rearranging leads to the last result of the proposition.

Proof of Proposition 4. Occupational Wage Gaps

The proof of this proposition is straightforward given the results of Proposition 2.

References

Becker, Gary, *The Economics of Discrimination*, Chicago, IL: The University of Chicago Press, 1957.

Biskupic, Joan, *Sandra Day O'Connor: How the First Woman on the Supreme Court Became Its Most Influential Justice*, New York, NY: HarperCollins, 2006.

Blau, Francine D., "Trends in the Well-Being of American Women, 1970-1995," *Journal of Economic Literature*, March 1998, *36* (1), 112–165.

Card, David and Alan B. Krueger, "School Quality and Black-White Relative Earnings: A Direct Assessment," *The Quarterly Journal of Economics*, February 1992, *107* (1), 151–200.

Charles, Kerwin Kofi and Jonathan Guryan, "Prejudice and Wages: An Empirical Assessment of Becker's The Economics of Discrimination," *Journal of Political Economy*, October 2008, *116* (5), 773–809.

Chay, Kenneth Y., Jonathan Guryan, and Bhashkar Mazumder, "Birth Cohort and the Black-White Achievement Gap: The Roles of Access and Health Soon After Birth," NBER Working Papers 15078, National Bureau of Economic Research, Inc June 2009.

Donohue, John J. and James Heckman, "Continuous versus Episodic Change: The Impact of Civil Rights Policy on the Economic Status of Blacks," *Journal of Economic Literature*, December 1991, *29* (4), 1603–43.

Eaton, Jonathan and Samuel S. Kortum, "Technology, Geography, and Trade," *Econometrica*, 2002, *70* (5), 1741–1779.

Erosa, Andres, Tatyana Koreshkova, and Diego Restuccia, "How Important Is Human Capital? A Quantitative Theory Assessment of World Income Inequality," *Review of Economic Studies*, October 2010, *77* (4), 1421–1449.

Fernandez, Raquel, "Culture as Learning: The Evolution of Female Labor Force Participation over a Century," NBER Working Papers 13373, National Bureau of Economic Research, Inc September 2007.

_ , Alessandra Fogli, and Claudia Olivetti, "Preference Formation and the Rise of Women's Labor Force Participation: Evidence from WWII," NBER Working Papers 10589, National Bureau of Economic Research, Inc June 2004.

Goldin, Claudia, *Understanding the Gender Gap: An Economic History of American Women*, New York, NY: Oxford University Press, 1990.

Karabel, Jerome, *The Chosen: the Secret History of Admission and Exclusion at Harvard, Yale, and Princeton*, New York, NY: Houghton Mifflin Harcourt, 2005.

Manuelli, Rodolfo and Ananth Seshadri, "Human Capital and the Wealth of Nations," March 2005. University of Wisconsin working paper.

McFadden, Daniel, "Conditional Logit Analysis of Qualitative Choice Behavior," in P. Zarembka, ed., *Frontiers of Econometrics*, New York, NY: Academic Press, 1974, pp. 105–142.

Rendall, Michelle, "Brain versus brawn: the realization of women's comparative advantage," IEW - Working Papers iewwp491, Institute for Empirical Research in Economics - University of Zurich September 2010.

Roy, A. D., "Some Thoughts on the Distribution of Earnings," *Oxford Economic Papers*, June 1951, *3* (2), 135–146.

Smith, James P. and Finis R. Welch, “Black Economic Progress after Myrdal,” *Journal of Economic Literature*, June 1989, *27* (2), 519–64.

Table 16: Sample Statistics By Census Year

	1960	1970	1980	1990	2000	2006-8
Sample Size	641,686	694,419	4,057,685	4,711,405	5,216,431	3,147,547
Share of White Men in Sample	0.432	0.433	0.435	0.435	0.431	0.431
Share of White Women in Sample	0.475	0.468	0.459	0.447	0.437	0.431
Share of Black Men in Sample	0.042	0.044	0.047	0.054	0.061	0.065
Share of Black Women in Sample	0.052	0.055	0.059	0.065	0.071	0.074
Relative Wage Gap: White Women	-0.578	-0.590	-0.476	-0.372	-0.299	-0.259
Relative Wage Gap: Black Men	-0.379	-0.289	-0.215	-0.158	-0.142	-0.150
Relative Wage Gap: Black Women	-0.875	-0.705	-0.479	-0.363	-0.317	-0.313

Note:

B Data Appendix

Table 17: Occupation Categories for our Base Occupational Specification

Home Sector	Police
Executives, Administrative, and Managerial	Guards
Management Related	Food Preparation and Service
Architects	Health Service
Engineers	Cleaning and Building Service
Math and Computer Science	Personal Service
Natural Science	Farm Managers
Health Diagnosing	Farm Non-Managers
Health Assessment	Related Agriculture
Therapists	Forest, Logging, Fishers, and Hunters
Teachers, Postsecondary	Vehicle Mechanic
Teachers, Non-Postsecondary	Electronic Repairer
Librarians and Curators	Misc. Repairer
Social Scientists and Urban Planners	Construction Trade
Social, Recreation, and Religious Workers	Extractive
Lawyers and Judges	Precision Production, Supervisor
Arts and Athletes	Precision Metal
Health Technicians	Precision Wood
Engineering Technicians	Precision Textile
Science Technicians	Precision Other
Technicians, Other	Precision Food
Sales, All	Plant and System Operator
Secretaries	Metal and Plastic Machine Operator
Information Clerks	Metal and Plastic Processing Operator
Records Processing, Non-Financial	Woodworking Machine Operator
Records Processing, Financial	Textile Machine Operator
Office Machine Operator	Printing Machine Operator
Computer and Communication Equipment Operator	Machine Operator, Other
Mail Distribution	Fabricators
Scheduling and Distributing Clerks	Production Inspectors
Adjusters and Investigators	Motor Vehicle Operator
Misc. Administrative Support	Non Motor Vehicle Operator
Private Household Occupations	Freight, Stock, and Material Handlers
Firefighting	

Note:

Table 18: Examples of Occupations within Our Base Occupational Categories

Management Related Occupations

Accountants and Auditors
Underwriters
Other Financial Officers
Management Analysts
Personnel, Training, and Labor Relations Specialists
Purchasing Agents and Buyers
Construction Inspectors
Management Related Occupations, N.E.C.

Health Diagnosing Occupations

Physicians
Dentists
Veterinarians
Optometrists
Podiatrists
Health Diagnosing Practitioners, N.E.C.

Personal Service Occupations

Supervisors, personal service occupations
Barbers
Hairdressers and Cosmetologists
Attendants, amusement and recreation facilities
Guides
Ushers
Public Transportation Attendants
Baggage Porters
Welfare Service Aides
Family Child Care Providers
Early Childhood Teacher Assistants
Child Care Workers, N.E.C.

Note:

Table 19: Occupation Categories for our Broad Occupation Classification

Home Sector	Sales, All
Executives, Administrative, and Managerial	Administrative Support, Clerks, and Record Keepers
Management Related	Fire, Police, and Guards
Architects, Engineers, Math, and Computer Science	Private Household and Food, Cleaning, and Personal Services
Natural and Social Scientists, Recreation, Religious, Arts, and Athletes	Farm, Related Agriculture, Logging, Forest, Fishing, Hunters and Extraction
Doctors and Lawyers	Mechanics and Construction
Nurses, Therapists, and Other Health Services	Precision Manufacturing
Teachers, Postsecondary	Manufacturing Operators
Teachers, Non-Postsecondary and Librarians	Fabricators, Inspectors, and Material Handlers
Health and Science Technicians	Vehicle Operators

Note: