

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

April 6, 2018 – Version 5.0

Abstract

In 1960, 94 percent of doctors and lawyers were white men. By 2010, the fraction was just 62 percent. Similar changes in other highly-skilled occupations have occurred throughout the U.S. economy during the last fifty years. Given the innate talent for these professions has unlikely changed differentially over time across groups, the change in the occupational distribution since 1960 suggests that a substantial pool of innately talented blacks and women in 1960 were not pursuing their comparative advantage. We examine the effect on aggregate productivity of the remarkable convergence in the occupational distribution between 1960 and 2010 through the prism of a Roy model. About one-quarter of growth in aggregate output per person over this period can be explained by the improved allocation of talent.

*We are grateful to Daron Acemoglu, Raquel Fernandez, Kevin Murphy, Rob Shimer, Ivan Werning, three anonymous referees, and numerous seminar participants for helpful comments. Jihee Kim, Munseob Lee, Huiyu Li, Ziho Park, Cian Ruane, and Gabriel Ulyssea provided excellent research assistance. Hsieh and Hurst acknowledge support from the University of Chicago's Booth School of Business, and Klenow from the Stanford Institute for Economic Policy Research (SIEPR).

1. Introduction

Over the last 50 years, there has been a remarkable convergence in the occupational distribution between white men, women, and blacks. For example, 94 percent of doctors and lawyers in 1960 were white men. By 2010, the fraction was just over 60 percent. Similar changes occurred throughout the economy during the last fifty years, particularly among highly-skilled occupations. A large literature attempts to explain these facts.¹ Yet no formal study has assessed the effect of these changes on aggregate productivity. Since the innate talent for a profession among members of group is unlikely to change over time, the change in the occupational distribution since 1960 suggests that a substantial pool of innately talented blacks and women in 1960 were not pursuing their comparative advantage. The resulting (mis)allocation of talent could potentially have important aggregate consequences.

This paper measures the aggregate productivity effects of the changing allocation of talent from 1960 to 2010. To do so, we examine the differences in labor market outcomes between race and gender groups through the prism of a Roy (1951) model of occupational choice. Within the model, every person is born with a range of talents across all possible occupations. In an efficient allocation, each individual chooses the occupation where she obtains the highest return for her talent.

We introduce three forces that will cause individuals to choose occupations where they do not have a comparative advantage. First, we allow for discrimination in the labor market. 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 after graduating was as a legal secretary (Biskupic, 2006). We model labor market discrimination as an occupation-specific wedge between wages and marginal products. This "tax" is a proxy for many common formulations of taste-based and statistical discrimination found in the literature.²

Second, the misallocation of talent can also be due to barriers to human capital in-

¹See, for example, Blau (1998), Blau, Brummund and Liu (2013b), Goldin (1990), Goldin and Katz (2012), Smith and Welch (1989) and Pan (2015). Detailed surveys of this literature can be found in Altonji and Blank (1999), Bertrand (2011), and Blau, Ferber and Winkler (2013a).

²See, for example, Becker (1957a), Phelps (1972) and Arrow (1973). A summary of such theories can be found in Altonji and Blank (1999).

vestment. We model these barriers as increased monetary costs associated with accumulating occupation-specific human capital. These costs are a proxy for many different race- and gender-specific factors. Examples include parental and teacher discrimination in favor of boys in the development of certain skills, historical restrictions on the admission of women to colleges or training programs, differences in school quality between black and white neighborhoods, differences in parental wealth and schooling levels across groups which alter the cost of investing in their children’s human capital.³

Finally, we allow for differences in “preferences” or social norms to drive occupational choice differences across groups. For example, there might have been strong social norms against women and blacks in high skilled occupations in the 1960s. The potential for preference or social norm differences across groups has been highlighted in the work of, among others, Johnson and Stafford (1998), Altonji and Blank (1999), and Bertrand (2011). We treat the home sector as additional occupation. As a result, we also allow for differences across groups in the extent to which they want to work in the home sector. This factor can capture changes in social norms related to women working at home. However, without the loss of generality, we can interpret the change in the preference for the home sector over time broadly so that it also includes changes in the preference for children or the ability to control the timing of fertility.⁴

To measure these three forces from the data, we make a key assumption that the distribution of innate talent of blacks and women — relative to white men — is constant over time. With this assumption, we back out the *change* in labor market frictions,

³Here 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) document 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. Goldin and Katz (2002), Bailey (2006), and Bailey, Hershbein and Milleri (2012) document that innovations related to contraception had important consequences for female labor market outcomes and educational attainment. Neal and Johnson (1996) document differences in AFQT scores across race and how controlling for AFQT explains a portion of Black-White gaps. Akcigit, Grigsby and Nicholas (2017) highlight the potential importance of parental liquidity constraints affecting investments into their children’s education. Therefore, educational attainment differences across groups can be determined by differences in parental background and parental wealth.

⁴The literature on changes in female labor supply due to changes in productivity, preferences, and social norms is extensive. See Fernández, Fogli and Olivetti (2004) and Fernández (2013) on the role of cultural forces, Greenwood, Seshadri and Yorukoglu (2005) on the role of home durables, and Goldin and Katz (2002) on for the role of birth control. Surveys on this literature can be found in Costa (2000) and Blau, Ferber and Winkler (2013a). There is also a literature on the changes in labor force participation rates for black men. Neal and Rick (2014), for example, highlights the importance of incarceration rates in explaining changes in labor market outcomes for black versus white men.

human capital frictions, and occupational preferences from synthetic panel data on the occupational distribution and wages of women and blacks relative to white men from 1960 to 2010. Specifically, we infer that occupational preferences, labor market frictions, and human capital frictions must have declined from 1960 to 2010 to jointly explain the observed convergence in the occupational distribution and wages between blacks and women relative to white men. When we filter these facts through the lenses of our general equilibrium model of occupational choice, we find that that changes in these frictions account for roughly one-fourth of growth in US GDP per person between 1960 and 2010. They also account for the entire rise in labor force participation over the last five decades.

We also use the structure of the model to decompose the contribution of each force. First, we use wage differences across groups within an occupation to discipline the role of preferences in explaining cross-group differences in occupational choice. If women did not like being lawyers in 1960, the model implies that women should have been paid more to compensate for this dis-amenity. Second, we use the life cycle structure of the model to distinguish between barriers to human capital attainment and labor market discrimination. In our setup, human capital barriers affect an individual's choice of human capital prior to entering the labor market. The effect of these barriers remains with a cohort throughout their life cycle. In contrast, labor market discrimination affects all cohorts within a given time period. Under these assumptions, we can use the changing differential life-cycle patterns of wages between groups to distinguish changing occupation-specific human capital barriers (which are akin to "cohort" effects) from changing occupation-specific labor market discrimination (which are akin to "time" effects).

We find that declining obstacles to accumulating human capital were much more important than declining labor market discrimination. Declining barriers to human capital attainment explain 24 percent of growth in U.S. GDP per person between 1960 and 2010, while declining labor market discrimination explains 6 percent of growth. Meanwhile, changing occupation-specific preferences across groups explain little of U.S. growth during this time period.

Our paper adds to the large literature explaining differences in occupational sorting and wage gaps between race and gender groups in two important ways. First, we extend

a Roy model of selection to explain differences in both occupational outcomes and wages across different race-sex groups. The model nests many of the stories highlighted in the literature to explain differential sorting patterns like labor market discrimination, barriers to human capital accumulation, and preference differences. Second, we use this model to assess how the changing occupational choice of women and blacks have contributed to US economic growth over the last fifty years.

The rest of the paper proceeds as follows. Section 2 presents the model. Sections 3 and 4 discuss data and identification. Section 5 presents the main results, and many robustness checks. Section 6 explores the implications of our model and its calibration for female labor supply elasticities, black-white discrimination, and education trends. Section 7 concludes.

2. Model

The economy consists of a continuum of workers, each in one of M discrete sectors, one of which is the home sector. Workers are indexed by occupation i , group g (such as race and gender), and cohort c . Each worker possesses heterogeneous abilities — some people are good teachers while others are good lawyers. The basic allocation to be determined in this economy is how to match workers with occupations.

2.1. Workers

As in a standard Roy (1951) model of occupational choice, workers are endowed with idiosyncratic talent ϵ in each sector (including the home sector). We add to this standard framework forces that alter the allocation of talent across occupations. These forces can take the form of discrimination in the labor market (τ^w), barriers to human capital accumulation (τ^h), and group-specific preferences for an occupation (z).

Individuals invest in their human capital and choose an occupation in an initial “pre-period”, after which they work in the chosen market occupation or home sector for three working life cycle periods (“young”, “middle”, and “old”). We assume that human capital investments and the choice of occupation are fixed after the pre-period.

Lifetime utility of a worker from group g and cohort c that chooses occupation i is a function of lifetime consumption, time spent on human capital accumulation, and the

preferences associated with choosing the occupation:

$$\log U = \left[\beta \sum_{t=c}^{c+2} \log C(c, t) \right] + \log [1 - s(c)] + \log z_{ig}(c) \quad (1)$$

Here $C(c, t)$ is consumption of cohort c in year t , s denotes time allocated to human capital acquisition in the pre-period, and z_{ig} is the utility benefit of group g from working in occupation i , and β parameterizes the trade-off between lifetime consumption and time spent accumulating human capital.⁵ We normalize the time endowment in the pre-period to 1 so $1 - s$ is leisure time in the pre-period. Forces such as changes in social norms for women working in the market sector or changing preferences in fertility and marriage patterns can be thought of as changes in z of the home sector of women (but not of men). Note that we assume no discounting of consumption for simplicity.

Individuals acquire human capital in the initial period, and this human capital remains fixed over their lifetime.⁶ Individuals use time s and goods e to produce h according to:

$$h_{ig}(c, t) = \bar{h}_{ig} \gamma(t - c) s_i(c)^{\phi_i} e_{ig}(c)^\eta.$$

\bar{h}_{ig} captures *permanent* differences by group-occupation pairs in human capital endowments and γ captures human capital due to experience. We assume γ is only a function of $age = t - c$ and \bar{h}_{ig} is fixed for a given group-occupation. \bar{h}_{ig} could include early investments in nutrition, health, cognitive development resulting from differing socioeconomic backgrounds, or simply natural differences in talent common to a group in a given occupation. ϕ_i is the occupation specific return to time investments in human capital accumulation while η is the elasticity of human capital with respect to human capital expenditures.

Consumption in each period is net income minus a portion of expenditures that are spent on education:

$$C(c, t) = [1 - \tau_{ig}^w(t)] w_i(t) \epsilon h_{ig}(c, t) - e_{ig}(c, t) [1 + \tau_{ig}^h(c)]. \quad (2)$$

⁵We define the cohort index c as the time when the cohort is young so time $t = c$ is the first period of cohort c . We omit subscripts on other individual-specific variables to keep the notation clean. However, z_{ig} does have subscripts to emphasize that it varies across groups and occupations.

⁶We do not allow workers to return to school after the pre-period. Given that in our empirical implementation our pre-period extends to age 25, this assumption is not too restrictive.

Net income is the product of $1 - \tau_{ig}^w$ and the total efficiency units of labor, which is the product of the price per efficiency unit of skill w_i , the idiosyncratic talent in the worker's chosen occupation ϵ and human capital h . Individuals borrow $e(c)(1 + \tau_{ig}^h(c))$ in the first period to purchase $e(c)$ units of human capital, a loan they repay over their lifetime subject to the lifetime budget constraint $e(c) = \sum_{t=c}^{c+2} e(c, t)$.

Labor market discrimination τ_{ig}^w works as a "tax" on individual earnings. Given our assumption that the firm owner discriminates against all workers of a given group, τ_{ig}^w affects all the cohorts of group g within occupation i equally at a given point in time. Barriers to human capital attainment τ_{ig}^h affect consumption directly by increasing the cost of e for a given group-occupation pair in (2) as well as indirectly by lowering acquired human capital e . We interpret τ_{ig}^h broadly to incorporate even early differences in childhood environments across groups, as long as these differences affect accumulation of human capital. That is, τ_{ig}^h reflects more than just discrimination in access to quality schooling. Because the human capital decision is only made once and fixed thereafter, τ_{ig}^h for a given occupation varies across cohorts and groups, but is fixed for a given cohort-group over time.

Given an occupational choice, the occupational wage w_i , and idiosyncratic ability ϵ in the occupation, the individual chooses consumption in each period and e and s in the initial pre-period to maximize lifetime utility given by (1) subject to the constraints given by (2) and $e(c) = \sum_{t=c}^{c+2} e(c, t)$. Individuals will choose the time path of $e(c, t)$ such that expected consumption is constant and equals one third of expected lifetime income. Lifetime income depends on τ_{ig}^h in the first period (when the individual is young) and the expected values of w_i , τ_{ig}^w , and γ in middle and old age. For simplicity, we assume that individuals anticipate that the return to experience varies by age but that the labor tax τ_{ig}^w and returns to market skill w_i they observe when young will remain constant over time. Because individuals expect the same conditions in future periods as in the first period (except for the accumulation of experience), expected lifetime income is proportional to income in the first period.

The amount of time and goods an individual spends on human capital are then:

$$s_i^* = \frac{1}{1 + \frac{1-\eta}{3\beta\phi_i}}$$

$$e_{ig}^* = \left(\frac{\eta(1 - \tau_{ig}^w)w_i\bar{\gamma}\bar{h}_{ig}s_i^{\phi_i}\epsilon}{1 + \tau_{ig}^h} \right)^{\frac{1}{1-\eta}}$$

where $\bar{\gamma} \equiv 1 + \gamma(1) + \gamma(2)$ is the sum of the experience terms over the life-cycle with $\gamma(0)$ set to 1. Time spent accumulating human capital is increasing in ϕ_i . Individuals in high ϕ_i occupations acquire more schooling and have higher wages as compensation for time spent on schooling. Forces such as w_i , \bar{h}_{ig} , τ_{ig}^h , and τ_{ig}^w do not affect s because they have the same effect on the wage gains from schooling and on the opportunity cost of time. These forces do change the return to investing *goods* in human capital (relative to the cost) with an elasticity that is increasing in η . These expressions hint at why we use both time and goods in the production of human capital. Goods are needed so that distortions to human capital accumulation matter. As we show below, time is needed to explain average wage differences across occupations.

After substituting the expression for human capital into the utility function, indirect expected utility for an individual from group g working in occupation i is

$$U_{ig}^* = [\bar{\gamma}\tilde{w}_{ig}\epsilon_i]^{\frac{3\beta}{1-\eta}}$$

where

$$\tilde{w}_{ig} \equiv \frac{\bar{h}_{ig}w_i s_i^{\phi_i} [(1 - s_i) z_{ig}]^{\frac{1-\eta}{3\beta}}}{\tau_{ig}}$$

and

$$\tau_{ig} \equiv \frac{(1 + \tau_{ig}^h)\eta}{1 - \tau_{ig}^w}. \quad (3)$$

τ_{ig} is a “composite” distortion that summarizes the effect of labor market discrimination and human capital barriers. More human capital barriers or labor market discrimination increase τ_{ig} , which lowers indirect utility for an individual from group g when choosing occupation i . Similarly, disutility from working in occupation i is represented as a low value of z_{ig} , which lowers the indirect utility for an individual from group g when choosing occupation i . Higher innate talent or human capital endowment \bar{h}_{ig} also increases the rewards for choosing an occupation.

Finally, turning to the distribution of the idiosyncratic talent, we borrow from McFadden (1974) and Eaton and Kortum (2002). Each person gets a skill draw ϵ_i in each of the M occupations. Talent in each occupation is drawn from a multivariate Fréchet

distribution:

$$F_g(\epsilon_1, \dots, \epsilon_M) = \exp \left[- \sum_{i=1}^M \epsilon_i^{-\theta} \right].$$

The parameter θ governs the dispersion of skills, with a higher value of θ corresponding to *smaller* dispersion. We normalize the mean parameter of the skill distribution to one in all occupations for all groups, but this mean parameter is isomorphic to \bar{h}_{ig} .

2.2. Occupational choice

Given the above assumptions, the occupational choice problem thus reduces to picking the occupation that delivers the highest value of U_{ig}^* .

Because talent is drawn from an extreme value distribution, the highest utility can also be characterized by an extreme value distribution, a result reminiscent of McFadden (1974). The overall occupational share can then be obtained by aggregating the optimal choice across people:⁷

Proposition 1 (Occupational Choice): *Let $p_{ig}(c)$ denote the fraction of people from cohort c and group g who choose occupation i , a choice made when they are young. Aggregating across people, the solution to the individual's choice problem leads to*

$$p_{ig}(c) = \frac{\tilde{w}_{ig}(c)^\theta}{\sum_{s=1}^M \tilde{w}_{sg}(c)^\theta} \quad (4)$$

where $\tilde{w}_{ig}(c) \equiv \frac{\bar{h}_{ig} w_i(c) s_i(c)^{\phi_i(c)} [(1-s_i(c)) z_{ig}(c)]^{\frac{1-\eta}{3\beta}}}{\tau_{ig}(c)}$.

Recall from (3) that $\tau_{ig}(c)$ is a composite of τ^h and τ^w facing cohort c when young ($t = c$). Occupational sorting depends on \tilde{w}_{ig} , which is the overall reward that someone from group g with the mean talent obtains by working in occupation i , relative to the power mean of \tilde{w} for the group over all occupations. The occupational distribution is driven by relative returns and not absolute returns: forces that change \tilde{w} for all occupations have no effect on the occupational distribution. Occupations where the wage per efficiency unit w_i is high will attract more workers. In contrast, z , \bar{h}_{ig} , τ^w , and τ^h can explain differences across groups in occupational choice. The fraction of members of

⁷Proofs of the propositions are given in the Appendix which can be found at <http://www.stanford.edu/~chadj/HHJKAppendix.pdf>.

group g that choose occupation i is low when a group dislikes like the occupation (z_{ig} is low), have low ability in the occupation (\bar{h}_{ig} is low), employers discriminate against the group in the occupation (τ_{ig}^w is high), or when the group faces a barrier in accumulating human capital associated with that occupation (τ_{ig}^h is high).

We view the home sector as simply another sector so the share of a group in the home sector is also given by equation (4). The labor force participation rate therefore depends on the returns in the home sector relative to the returns in the market sectors. For example, the decline in the labor force participation rate of white men since the 1960s can be driven by higher returns in the home sector (such as better video games) or by a decline in labor market opportunities (such as the decline of blue-collar jobs).

2.3. Worker Quality

Sorting generates an average quality of workers in an occupation for each group. We show this in the following proposition:

Proposition 2 (Average Quality of Workers): *For a given cohort c of group g at time t , the average quality of workers in each occupation, including both human capital and talent, is*

$$\mathbb{E}[h_{ig}(c, t) \epsilon_{ig}(c)] = \Gamma s_i(c)^{\phi_i(t)} \gamma(t-c) \left[\left(\frac{\eta s_i(c)^{\phi_i(c)} \bar{\gamma} \bar{h}_{ig} w_i(c) (1 - \tau_{ig}^w(c))}{1 + \tau_{ig}^h(c)} \right)^\eta \left(\frac{1}{p_{ig}(c)} \right)^{\frac{1}{\theta}} \right]^{\frac{1}{1-\eta}} \quad (5)$$

Here $\Gamma \left(1 - \frac{1}{\theta} \cdot \frac{1}{1-\eta}\right)$ is related to the mean of the Fréchet distribution for abilities. Notice that average quality is inversely related to the share of the group working in the occupation $p_{ig}(c)$. This captures the selection effect. For example, the model predicts that if the labor market discriminated against female lawyers in 1960, only the most talented female lawyers would have chosen to work in this occupation. And if the barriers faced by female lawyers declined after 1960, less talented female lawyers would move into the legal profession and thus lower the average quality of female lawyers. Conversely, in 1960, the average quality of white male lawyers would have been lower in the presence of labor market discrimination against women and blacks.

2.4. Occupational Wages

Next, we compute the average wage for a given group working in a given occupation — the model counterpart to what we observe in the data.

Proposition 3 (Occupational Wages): *Let $\overline{wage}_{ig}(c, t)$ denote the average earnings in occupation i by cohort c at date t of group g . Its value satisfies*

$$\begin{aligned} \overline{wage}_{ig}(c, t) &\equiv (1 - \tau_{ig}^w(t)) w_i(t) \mathbb{E} [h_{ig}(c, t) \epsilon_{ig}] \\ &= \Gamma \bar{\eta} m_g(c)^{\frac{1}{\theta} \frac{1}{1-\eta}} \cdot [(1 - s_i(c)) z_{ig}(c)]^{-\frac{1}{3\beta}} \frac{(1 - \tau_{ig}^w(t)) w_i(t)}{(1 - \tau_{ig}^w(c)) w_i(c)} \frac{\gamma(t-c)}{\bar{\gamma}} \frac{s_i(c)^{\phi_i(t)}}{s_i(c)^{\phi_i(c)}} \end{aligned} \quad (6)$$

where $m_g(c) \equiv \sum_{i=1}^M \tilde{w}_{ig}(c)^\theta$ and $\bar{\eta} \equiv \eta^{\eta/(1-\eta)}$.

For individuals in the young cohort, $t = c$ which implies $\frac{s_i(c)^{\phi_i(t)}}{s_i(c)^{\phi_i(c)}} = 1$ and $\frac{(1 - \tau_{ig}^w(t)) w_i(t)}{(1 - \tau_{ig}^w(c)) w_i(c)} = 1$. Average earnings for a given group among the young differs across occupations only because of the term $[(1 - s_i(c)) z_{ig}(c)]^{-1/3\beta}$. Occupations in which schooling is especially productive (a high ϕ_i and therefore a high s_i) will have higher average earnings. Similarly, occupations where individuals have a strong dis-utility from being in the profession (z_{ig} is small) have higher wages as compensation for the lower utility. And these are the only two forces that generate differences in wages across occupations for the young. Average earnings are no higher in occupations where a group faces less discrimination in the labor market, lower frictions in human capital attainment, a higher wage per efficiency unit, or where the group has more talent in the sector. The reason is that each of these factors leads lower quality workers to enter those jobs. This composition effect exactly offsets the direct effect on earnings when the distribution of talent is Fréchet.

The exact offset due to selection is a feature of the Fréchet distribution, and we would not expect this feature to hold more generally. However, the general point is that when the selection effect is present, the wage gap is a poor measure of the frictions faced by a group in a given occupation. Such frictions lower the wage of the group in *all* occupations, not just in the occupation where the group encounters the friction.

The composition effect would not be present if selection was driven by forces other than occupational ability. For example, if individuals differ by *tastes* for an occupation instead of ability, then selection has no effect on average ability. In this case, the average

wage in an occupation will vary with \tilde{w} so the average wage and the occupational share will both be higher in occupations where a group faces less discrimination or where the wage per efficiency unit is higher.⁸ We will later see that this prediction generated by a model of selection on preferences is overwhelmingly rejected by the data.

Reverting back to the model where ϵ represents occupational skill, equation (6) for the average wage also identifies the forces behind wage changes over a cohort's life-cycle. For a given cohort-group in an occupation, s_i and z_{ig} are fixed. Therefore, the average wage increases over time when the price of skills in the occupation w_i increases, labor market discrimination τ^w falls, return to experience is positive, or the return to schooling increases.

Comparing wage changes across groups, the effect of the returns to schooling, experience, and returns to skill have the same effect on all groups (of a given cohort in the occupation). Thus, differences in the growth rate of wages between groups (say between men and women) can only be due to differences in the change in τ^w between the groups. We will use this insight to estimate the change in τ^w in the empirical section.

2.5. Relative Propensities

Putting together the equations for the occupational shares and wages in each occupation and assuming the experience profiles are the same across groups, we get the relative propensity of a group to work in an occupation. These equations provide us a way to explain differences in occupational choices of groups through the lens of our model that can be mapped directly to observable data moments.

Proposition 4 (Relative Propensities): *The fraction of a group working in an occupation — relative to white men — is given by*

$$\frac{p_{ig}(c)}{p_{i,wm}(c)} = \left(\frac{\tau_{ig}(c)}{\tau_{i,wm}(c)} \right)^{-\theta} \left(\frac{\bar{h}_{ig}}{\bar{h}_{i,wm}} \right)^{\theta} \left(\frac{\overline{wage}_{ig}(c, c)}{\overline{wage}_{i,wm}(c, c)} \right)^{-\theta(1-\eta)} \quad (7)$$

where the subscript “wm” denotes white men.

⁸We sketch the model of selection on preferences in Appendix B.

The propensity of a group to work in an occupation (relative to white men) depends on three terms: the relative composite occupational frictions, relative talent in the sector, and the average wage gap between the groups in the occupation. From Proposition 3, the wage gap itself is a function of the distortions faced by the group, the talent of the group, and the price of skills in *all* occupations. With data on occupational shares and wages, we can measure a composite term that measures the combined effect of labor market discrimination, barriers to human capital attainment, and talent in the sector. The preference parameters z_{ig} do not enter this equation once we have controlled for the wage gap; instead, they influence the wage gaps themselves.

2.6. Relative Labor Force Participation

Remember we treat the home sector as another sector so the labor force participation rate of a group relative to white men is also given by equation (7). As we discuss below in Section 4, we normalize $z = 1$, $\tau^w = 0$ and $\tau^h = 0$ for the home sector and $z_{i,wm} = 1$. With these assumptions, the relative labor force participation rate is given by the following proposition:

Proposition 5 (Relative Labor Force Participation): *Let $LFP_g \equiv 1 - p_{home,g}$ denote the share of group g in the market sectors. The share of group g in the home sector relative to white men is then*

$$\begin{aligned} \frac{1 - LFP_g(c)}{1 - LFP_{wm}(c)} &= \frac{m_{wm}(c)}{m_g(c)} \\ &= z_{ig}(c)^{-\frac{\theta(1-\eta)}{3\beta}} \left(\frac{\overline{wage}_{ig}(c, c)}{\overline{wage}_{i,wm}(c, c)} \right)^{-\theta(1-\eta)} \quad \forall i \in market \end{aligned} \quad (8)$$

where $\frac{m_{wm}(c)}{m_g(c)} \equiv \frac{\sum_{i=1}^M \tilde{w}_{i,wm}(c)^\theta}{\sum_{i=1}^M \tilde{w}_{ig}(c)^\theta}$.

Since the return to the home sector sector is the same for all groups (this is an implication of the normalization that the home sector is undistorted), $\frac{m_{wm}(c)}{m_g(c)}$ is the return to *market* work of white men relative to group g . For example, if women are discriminated against in the labor market or in accumulating human capital for the market sector, this will drive down female labor force participation rates. Or if social norms discourage women from the market sector (low z in market sectors), this will also lower female labor force participation rates.

The second term in equation (8) says that the relative return to market work is given by a power function of the gap in market wages in *any* market sector and the occupational preference term in that sector. We will use this insight to back out z_{ig} in the market sectors from data on labor force participation of the group (relative to white men) and wage gaps.

2.7. Firms and Determinants of Labor Market and Human Capital Frictions

A representative firm produces final output Y from workers in M occupations:

$$Y = \left[\sum_{i=1}^M \left(A_i \cdot \sum_g H_{ig} \right)^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma}{\sigma-1}} \quad (9)$$

where H_{ig} denotes the total efficiency units of labor provided by group g in occupation i and A_i is the exogenously-given productivity of occupation i . The parameter σ represents the elasticity of substitution across occupations in aggregate production.

Following Becker (1957a), we assume the owner of the firm in the final goods sector discriminates against workers of certain groups. We model the “taste” for discrimination as lower utility of the owner when she employs workers from groups she dislikes. Her utility is given by

$$U_{owner} = \underbrace{Y - \sum_i \sum_g (1 - \tau_{ig}^w) w_i H_{ig}}_{\text{Profit}} - \underbrace{\sum_i \sum_g d_{ig} H_{ig}}_{\text{Utility loss via discrimination}}. \quad (10)$$

The first term denotes profits and the second term captures the extent to which owners are prejudiced: d_{ig} is the utility loss associated with employing workers from group g in occupation i . Because all employers are assumed to have these racist and sexist preferences, perfect competition implies that $\tau_{ig}^w = d_{ig}/w_i$. Intuitively, when the owner hires a worker from a group she dislikes, she needs to be compensated for her utility loss via a lower wage for these workers. In equilibrium the utility loss is exactly offset by the lower wage. Thus the frictions are ultimately pinned down by the discriminatory tastes of (homogeneous) owners.⁹

⁹What is important is not that all firms discriminate but that the marginal firm discriminates (Becker (1957b)). We abstract from having a continuum of firms within each occupation and instead assume

A second firm (a “school”) sells educational goods e to workers who use it as an input in their human capital. We assume the school’s owner dislikes providing e to certain groups. The utility of the school’s owner is

$$U_{school} = \underbrace{\sum_i \sum_g \left(R_{ig} - (1 - \tau_{ig}^h) \right) \cdot e_{ig}}_{\text{Profit}} - \underbrace{\sum_i \sum_g d_{ig}^h e_{ig}}_{\text{Utility loss via discrimination}} \quad (11)$$

where e_{ig} denotes educational resources provided to workers from group g in market sector i , R_{ig} denotes the price of e_{ig} , and d_{ig}^h represents the owner’s distaste from providing educational resources to workers from group g in sector i . We think of this as a shorthand for complex forces such as discrimination against blacks or women in admission to universities, or differential allocation of resources to public schools attended by black vs. white children, or differential parental investments made toward building up math and science skills in boys relative to girls. Groups that are discriminated against in the provision of human capital pay a higher price for e , and the higher price compensates the school owner for her disutility. Perfect competition ensures that $R_{ig} = 1$, and that $\tau_{ig}^h = d_{ig}^h$.

2.8. Equilibrium

A competitive equilibrium in this economy consists of a sequence of individual choices $\{C, e, s\}$, occupational choices in the pre-periods, total efficiency units of labor of each group in each occupation H_{ig} , 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 expected lifetime utility given by (1) subject to the constraints given by (2) and $e(c) = \sum_{t=c}^{c+2} e(c, t)$.
2. Each individual chooses the occupation that maximizes expected lifetime utility: $i^* = \arg \max_i U(\tau_{ig}^w, \tau_{ig}^h, z_{ig}, w_i, \epsilon_i)$, taking as given $\{\tau_{ig}^w, \tau_{ig}^h, z_{ig}, w_i, \bar{h}_{ig}, \epsilon_i\}$.
3. A representative firm in the final good sector hires H_{ig} in each occupation to maximize profits net of utility cost of discrimination given by equation (10).

all firms within an occupation discriminate. This simplifies the analysis but still allows us to match key features of the data.

4. A representative firm in the education sector maximizes profit net of the utility cost of discrimination given by equation (11).
5. Perfect competition in the final goods and education sectors generates $\tau_{ig}^w = d_{ig}/w_i$ and $\tau_{ig}^h = d_{ig}^h$.
6. $w_i(t)$ clears each occupational labor market.
7. Total output is given by the production function in equation (9).

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

Proposition 6 (Solving the General Equilibrium): *The general equilibrium of the model is H_{ig}^{supply} , H_i^{demand} , w_i , and market output Y at each point in time such that*

1. $H_{ig}^{supply}(t)$ aggregates the individual choices:

$$H_{ig}^{supply}(t) = \sum_c q_g(c) p_{ig}(c) \mathbb{E} [h_{ig}(c) \epsilon_{ig}(c) \mid \text{Person chooses } i]$$

where $q_g(c)$ denotes the number of workers of group g and cohort c and the average quality of workers is given in equation (5).

2. $H_i(t)^{demand}$ satisfies firm profit maximization:

$$H_i^{demand}(t) = \left(\frac{A_i(t)^{\frac{\sigma-1}{\sigma}}}{w_i(t)} \right)^\sigma Y(t)$$

3. $w_i(t)$ clears each occupational labor market: $\sum_g H_{ig}^{supply}(t) = H_i^{demand}(t)$.
4. Total output is given by the production function in equation (9) and equals aggregate wages plus total revenues from τ^w .

2.9. Intuition

To develop intuition, consider the following simplified version of the model. First, assume only two groups, men and women, and assume that men face no distortions. Second, assume occupations are perfect substitutes ($\sigma \rightarrow \infty$) so that $w_i = A_i$. With this assumption, the production technology parameter pins down the wage per unit

of human capital in each occupation. In addition, labor market and human capital frictions affect aggregate output produced by women but have no effect on output produced by men. Third, assume $\phi_i = 0$ (no schooling time), $z_i = 1$, and $\bar{h}_i = 1$. Finally, assume that each cohort lives for one period.

Aggregate output can then be expressed as the sum of aggregate output produced from male labor and female labor:

$$Y = \underbrace{q_m \cdot \left(\sum_{i=1}^M A_i^\theta \right)^{\frac{1}{\theta} \cdot \frac{1}{1-\eta}}}_{\text{output from men}} + \underbrace{\frac{q_w}{1 - \bar{\tau}^w} \cdot \left(\sum_{i=1}^M \left(\frac{A_i (1 - \tau_i^w)}{(1 + \tau_i^h)^\eta} \right)^\theta \right)^{\frac{1}{\theta} \cdot \frac{1}{1-\eta}}}_{\text{output from women}} \quad (12)$$

where q_w and q_m denote the number of women and men and $\bar{\tau}^w$ denotes the earnings-weighted average of the labor market friction facing women.¹⁰ The first term in equation (12) is aggregate output produced by men and is not affected by the occupational distortions facing women (this is driven by the assumption that occupations are perfect substitutes). The second term is aggregate output produced by female labor. So the effect of τ^w and τ^h on aggregate output shows up in the second term, and their effect on aggregate output is increasing in the number of people in the discriminated group q_w .

We illustrate how this setup can be used to gain intuition by focusing on τ^w ; the effects of τ^h can be analyzed in a similar fashion.¹¹ Assuming $\tau^h = 0$ and that τ^w and A are jointly log-normally distributed, aggregate output produced by women Y_w (the second term in equation (12)) is given by

$$\ln Y_w = \ln q_w + \ln \left(\sum_{i=1}^M A_i^\theta \right)^{\frac{1}{\theta} \cdot \frac{1}{1-\eta}} + \frac{\eta}{1-\eta} \cdot \ln(1 - \bar{\tau}^w) - \frac{1}{2} \cdot \frac{\theta - 1}{1-\eta} \cdot \text{Var} \ln(1 - \tau_i^w). \quad (13)$$

τ^w affects output via the last two terms in equation (13). The mean of τ^w changes the return to investment in human capital. This effect is captured by the third term in equation (13) and its magnitude depends on elasticity of output with respect to

¹⁰ $\bar{\tau}^w \equiv \sum_{i=1}^M \omega_i \tau_i^w$ where $\omega_i \equiv \frac{\frac{p_{i,w} \bar{w}^\alpha g^\epsilon \bar{e}_w}{1 - \tau_i^w}}{\sum_{j=1}^M \frac{p_{j,w} \bar{w}^\alpha g^\epsilon \bar{e}_w}{1 - \tau_j^w}}$.

¹¹ Assuming $\tau^w = 0$ and that τ^h and A are jointly log-normally distributed, aggregate output produced by women is $\ln Y_w = \ln \left(\sum_{i=1}^M A_i^\theta \right)^{\frac{1}{\theta} \cdot \frac{1}{1-\eta}} - \frac{\eta}{1-\eta} \cdot \ln(1 + \bar{\tau}^h) - \frac{\eta^2}{2} \cdot \frac{\theta - 1}{1-\eta} \cdot \text{Var} \ln(1 + \tau_i^h)$.

human capital η . The dispersion of τ^w across occupations affects aggregate output via a different channel. Here, dispersion of τ^w affects the allocation of female labor across occupations. A decline in the dispersion of τ^w improves the allocation, which increases aggregate output. This effect is captured by the fourth term in equation (13).

Finally, equation (13) suggests that the effect of misallocation on aggregate output is increasing in θ . While this is true for a given amount of misallocation, remember that the inference about the magnitude of misallocation from observed data also depends on θ . Using the equation for relative propensities, the variance in the labor distortion is given by:¹²

$$\text{Var} \ln(1 - \tau^w) = \frac{1}{\theta^2} \cdot \text{Var} \ln \frac{p_{ig}}{p_{i,wm}}$$

This says that, conditional on data on occupational shares, the implied dispersion of τ^w is *decreasing* in θ . Expressed as a function of data on occupational propensities, aggregate output from female labor is:

$$\ln Y_w = \ln q_w + \ln \left(\sum_{i=1}^M A_i^\theta \right)^{\frac{1}{\theta}} + \frac{\eta}{1-\eta} \cdot \ln(1 - \bar{\tau}^w) - \frac{1}{2} \cdot \frac{\theta - 1}{(1-\eta)\theta^2} \cdot \text{Var} \ln \left(\frac{p_{ig}}{p_{i,wm}} \right)$$

The elasticity of Y_w with respect to the variance in the observed propensities in the data is $\frac{1}{2} \cdot \frac{\theta-1}{(1-\eta)\theta^2}$ while the elasticity with respect to the variance in τ^w is $\frac{1}{2} \cdot \frac{\theta-1}{1-\eta}$. Intuitively, a higher value of θ implies that a given amount of misallocation has a larger effect on aggregate output. On the other hand, given the observed data on occupational shares, a higher θ also implies a smaller amount of misallocation. For this reason, as we document later, the effect of changes in occupational shares on output growth will not be very sensitive to the values we use for θ .

3. Data

We use data from the 1960, 1970, 1980, 1990, and 2000 decennial Censuses and the 2010-2012 American Community Surveys (ACS).¹³ We make four restrictions to the data when performing our analysis. First, we restrict the sample to white men (wm), white

¹²We maintain the assumption that τ^w is the only source of variation.

¹³When using the 2010–2012 ACS data, we pool all three years together for power and treat them as one cross section. Henceforth, we refer to the pooled 2010-2012 sample as the 2010 sample.

women (ww), black men (bm) and black women (bw). These will be the four groups we analyze in the paper. Second, we only include individuals between the ages of 25 and 54. This restriction focuses the analysis on individuals after they finish schooling and prior to retirement. Third, we exclude individuals on active military duty. Finally, we exclude individuals who report being unemployed (not working but searching for work). Our model is not well suited to capture transitory movements into and out of employment.¹⁴

We do not have actual panel data. Instead, we create pseudo-panel data by following synthetic cohorts over time. We define three age periods within a cohort's life cycle: the young (those aged 25-34), the middle aged (those aged 35-44) and the old (those aged 45-54). For example, a synthetic panel for a given cohort would be the young in 1960, the middle aged in 1970, and the old in 1980. We have information on 8 cohorts for the time periods we study. For 4 cohorts (the young in 1960, 1970, 1980, and 1990), we observe information at all three life cycle points. We observe either one or two life cycle points for the remaining cohorts.

We made a few other adjustments to our data. First, we define a person who is not currently employed or who works less than ten hours per week as being in the home sector. Those who are employed but usually work between ten and thirty hours per week are classified as part-time workers. We split the sampling weight of part-time workers equally between the home sector and the occupation in which they are working. Individuals working more than thirty hours per week are considered to be full-time in a market occupation. Second, we define the market occupations using the roughly 67 occupational sub-headings from the 1990 Census occupational classification system.¹⁵

We measure earnings as the sum of labor, business, and farm income in the previous year. For earnings we restrict the sample to individuals who worked at least 48 weeks during the prior year, who earned at least 1000 dollars (in 2007 dollars) in the previous

¹⁴The Appendix reports summary statistics from our sample. For all analysis in the paper, we apply the sample weights in each survey.

¹⁵See <http://www.bls.gov/nls/quex/r1/y97r1cbka1.pdf>. We chose the 1990 occupation codes because they are available in all Census and ACS years since 1960. Appendix Table C2 reports the 67 occupations we analyze. Some samples of the occupational categories are "Executives, Administrators, and Managers", "Engineers", "Natural Scientists", "Health Diagnostics", "Health Assessment", and "Lawyers and Judges." We have also experimented with a more detailed classification of occupations by using 340 three digit occupation groupings that were defined consistently since 1980, as well as aggregating occupations into 20 broad occupational groups defined consistently since 1960. Our results were broadly similar at these different levels of occupation aggregation.

year, and who reported working more than 30 hours per week. We convert all earnings data from the Census to constant dollars. Our measure of wage gaps across groups is the difference in log earnings between groups.¹⁶

4. Inference

We now show how we use the synthetic panel data on wages and occupational shares from 1960 to 2010 to uncover the change in τ^h , τ^w , and z . Our inference exercise is based on a number of key assumptions.

First, we assume that talent of women and blacks and men is the same. This is a key identification assumption for us, and we cannot proceed without it. For our base case, we normalize \bar{h}_{ig} to one. We show through a series of robustness exercises, however, that our main results change very little with alternative approaches as long as the ratio $\bar{h}_{ig}/\bar{h}_{i,wm}$ is constant over time. I.e., we really need to assume that *relative* talent of women and blacks to white men is constant over time. This assumption implies that the *change* in the occupational distribution of women and blacks relative to white men since 1960 must be driven by *changes* in τ^h , τ^w , and z .

Second, we assume that idiosyncratic occupational abilities are distributed iid Fréchet. We need to impose structure on the distribution of occupation skills to infer τ^w and τ^h over time. Relaxing this distributional assumption is a valuable direction for future research and would change the precise magnitudes of our estimates, but would not change the fact that these frictions must have changed since the 1960s (as long as the innate talent of a group relative to white men is constant over time).¹⁷


If selection is based on tastes for an occupation instead of ability, then average wages in an occupation should be positively correlated with occupational shares.¹⁸ Figure 1 shows that there is no relationship between the share of young white women in an occupation relative to young white men in 1980 and the corresponding occupational

¹⁶Our results were not altered when adjusting for hours worked across groups. This is not surprising given that we already condition on full time work status. When computing average earnings by occupation, we include both top-coded and imputed data. We experimented with excluding top-coded and imputed data and it had no effect on our estimated τ 's.

¹⁷Lagakos and Waugh (2013) and Adão (2016) have estimated selection models with arbitrary correlation but only with 2 or 3 sectors. We do not know how to do something similar for the nearly 70 occupations we have in our data.

¹⁸See Appendix B for the model where selection is based on preferences.

Figure 1: Occupational Wage Gaps vs Relative Occupational Propensities: White Women in 1980



Note: The figure shows the relationship between the (log) occupational earnings gap for white women compared to white men (both in the young cohort) and the relative propensity to work in the occupation for the two groups, $p_{i,ww}/p_{i,wm}$ in 1980. A simple regression line through the scatter plot in the top panel yields a slope coefficient of 0.01 with a standard error of 0.01.

wage gaps between young women and men in the same year. For example, young white women were 64 times more likely to work as secretaries as young white men in 1980 and were one-fourth as likely to work as lawyer in 1980. Yet, the wage gap between young white women and young white men among secretaries was nearly identical to the gender wage gap among lawyers.¹⁹ While the absence of a tight relationship between occupational wage gaps and occupational propensities suggests that there is no systematic selection based on preferences, gender wage gaps are not perfectly equated across occupations. In our model, this variation will discipline the z 's (which are common across group-occupation pairs).

If we were only interested in uncovering τ_{ig} , we would not need any further key identifying assumptions. However, to decouple τ_{ig}^h from τ_{ig}^w , we rely on our modeling assumption that agents only make an active choice to obtain human capital prior to entering the labor market. By making this assumption, we can treat the human capital

¹⁹A weighted regression of the scatter plot yields a slope coefficient of 0.01 with a standard error of 0.01. In the Online Appendix, we also show that there is no systematic correlation between the *change* in relative occupational shares between 1960 and 2010 and the *change* in the occupational wage gap over that time.

frictions as something akin to a cohort effect. Labor market discrimination, on the other hand, affects all cohorts in the labor market at the same point in time. This allows us to treat the labor market discrimination as something akin to a time effect. Without this additional model structure, we would not be able to separately identify how human capital frictions and labor market discrimination have affected U.S. productivity growth over the last 50 years. However, we want to stress that this assumption is not essential to inferring the joint effect of changes in these variables on aggregate productivity growth.

We need three further assumptions to complete our inference. The procedure discussed above leverages the changing nature of occupational sorting and wages to pin down changes in occupational frictions and occupational preferences across groups over time. However, without additional assumptions, we cannot identify their level. We therefore need to pick one sector and one group as being undistorted. For our base-case, we pick the home sector and white men as the undistorted sector and group. So, in what follows, z_{ig} , τ_{ig}^w , and τ_{ig}^h should be interpreted as preferences and distortions relative to the home sector and white men. Below, we explore the robustness of our base results to alternate assumptions.

4.1. Composite Frictions vs. Occupational Preferences

We now recover a composite of the frictions and occupational preferences from data on wages and occupational shares. Remember the three normalizations discussed above. First, we normalize $\bar{h}_{ig}/\bar{h}_{i,wm} = 1$. Second, we assume that occupational choice of white men is undistorted. Third, we assume the home sector is undistorted for all groups.


Given these normalizations, we rearrange equation (7) to solve for the composite

$$\tau_{ig} \equiv \frac{(1+\tau_{ig}^h)^\eta}{1-\tau_{ig}^w}$$

$$\tau_{ig}(c) = \left(\frac{p_{ig}(c)}{p_{i,wm}(c)} \right)^{-1/\theta} \left(\frac{\overline{\text{wage}}_{ig}(c, c)}{\overline{\text{wage}}_{i,wm}(c, c)} \right)^{-(1-\eta)} \quad (14)$$

Equation (14) says that, conditional on estimates of θ (the shape parameter of the Fréchet distribution) and η (the elasticity of human capital to education expenditure), we need two pieces of data to recover τ_{ig} . These are the share of the group working in the occupation relative to that of white men and the average wage gap of a group

Figure 2: Standard Deviation of Relative Occupational Shares



Note: Figure shows earnings-weighted standard deviation of the log of occupational propensities for young white women (middle line), young black men (bottom line), and young black women (top line) relative to young white men.

relative to white men in the occupation. Intuitively, when the share of some group in an occupation is low (after we control for the wage gap), we infer that the group faces discrimination in the labor market or barriers to acquiring the human capital necessary for the occupation. Note, that (14) uses data only for the young cohorts to infer the composite τ_{ig} 's.

We now present the ingredients needed to measure τ_{ig} from equation (14). Figure 2 plots the standard deviation of $\ln(p_{ig}/p_{i,wm})$ across market occupations for the young cohort in each decade.²⁰ As shown in Figure 2, the occupations of white men and white women have converged over time. In particular, the standard deviation of $\ln(p_{ig}/p_{i,wm})$ fell sharply from 1960 through 2000. For black men, the standard deviation of $\ln(p_{ig}/p_{i,wm})$ also fell sharply between 1960 and 1980 and has remained relatively constant since. When filtered through equation (14), the decline in the dispersion of $\ln(p_{ig}/p_{i,wm})$ implies that the dispersion of the combination of τ^w , τ^h and z has declined

²⁰We compute the standard deviation of $(\ln(p_{ig}) - \ln(p_{i,wm}))$ across occupations weighting each occupation by the share of earnings in that occupation. We exclude the home sector in this calculation.

over time.

To determine how much of the occupational convergence is driven by convergence in τ_{ig} versus convergence in z_{ig} , we use the fact that convergence in z_{ig} also narrows the dispersion in wages across occupations. Intuitively, wage gaps are driven by differences in utility across occupations. In the absence of differences in z_{ig} , wage gaps for young cohorts should be equal in all occupations. Specifically, for the young cohort ($t = c$), the wage in occupation i for group g relative to the wages for white men is:

$$\frac{\overline{\text{wage}}_{ig}(c, c)}{\overline{\text{wage}}_{i,wm}(c, c)} = \left(\frac{m_g(c)}{m_{wm}(c)} \right)^{\frac{1}{\theta \cdot (1-\eta)}} \cdot z_{ig}(c)^{-1/3\beta} \quad (15)$$


where $m_g(c) \equiv \sum_{i=1}^M \tilde{w}_{ig}(c)$. So up to the term m_g/m_{wm} (which is the same across all occupations for the group), the wage gap in an occupation isolates the effect of heterogeneity in z_{ig} on the occupational choice patterns in Figure 2. Putting together (15) with (14) implies that when we condition the occupational gaps on the wage gap, we isolate the effect of the labor market and human capital distortions on occupational choice.

Figure 3 plots the weighted standard deviation of the wage gap across market occupations for young women and blacks relative to white men. The standard deviation of the wage gap across occupations fell for each group (relative to white men) between 1960 and 1980. The decline in the standard deviation of the wage gap suggests that the dispersion of z_{ig} declined from 1960 to 1980, and this decline is likely to be partially responsible for the narrowing of the gap in the occupational distribution over this time period.

4.2. Estimating θ , η , and β

In addition to wages and occupational choice of the young, we also need estimates of θ and η to infer τ_{ig} . Given our assumptions, wages within an occupation for a given group follow a Fréchet distribution with the shape parameter $\theta(1 - \eta)$. This reflects both comparative advantage (governed by $1/\theta$) and amplification from endogenous human capital accumulation (governed by $1/(1 - \eta)$). Using micro data from the U.S. Population Census/ACS, we estimate $\theta(1 - \eta)$ to fit the distribution of the residuals from a cross-sectional regression of log hourly wages on 66x4x3 occupation-group-age

Figure 3: Standard Deviation of Wage Gaps by Decade



Note: Figure shows the earnings-weighted standard deviation of the log of the average wage of young white women (middle line), young black men (bottom line), and young black women (top line) relative to young white men.

dummies in each year. We use MLE, with the likelihood function taking into account the number of observations which are top-coded in each year. The resulting estimates for $\theta(1 - \eta)$ range from a low of 1.24 in 1980 to a high of 1.42 in 2000, and average 1.36 across years.²¹

The parameter η denotes the elasticity of human capital with respect to education spending and is equal to the fraction of output spent on human capital accumulation. Spending on education (public plus private) as a share of GDP in the U.S. averaged 6.6 percent over the years 1995, 2000, 2005, and 2010.²² Since the labor share in the U.S. in the same four years was 0.641, this implies an η of 0.103.²³ With our base estimate of $\theta(1 - \eta) = 1.36$, $\eta = 0.103$ gives us $\theta = 1.52$.

We can also estimate θ from the elasticity of labor supply. In our model, the exten-

²¹Sampling error is minimal because there are 300-400k observations per year for 1960 and 1970 and 2-3 million observations per year from 1980 onward. We did not use 2010 data because top-coded wage thresholds differed by state in that year.

²²See <http://www.oecd.org/education/eag2013.htm>.

²³Labor share data are from <https://research.stlouisfed.org/fred2/series/LABSHPUA156NRUG>. The young's share of earnings is from the U.S. Population Census/ACS.

sive margin elasticity of labor supply with respect to a wage change is $\theta(1 - LFP_g)$. The meta analysis in Chetty et al. (2012) suggests an extensive margin labor supply elasticity of about 0.26 for men. The underlying data in their meta analysis come from the 1970-2007 period. In 1990, roughly in the middle of their analysis, 89.9 percent of men aged 25-34 were in the labor force. To match a labor supply elasticity of 0.26, our model implies that θ would equal 2.57. This is higher than the estimate of θ we get from wage dispersion. As a compromise between our two estimates, we will use $\theta = 2$ as our base case, but will also provide results with θ of 1.5 and 4.

Finally, β is the geometric weight on consumption relative to time in an individual's utility function (1). This parameter is needed to help distinguish preferences from labor market and human capital frictions. As schooling trades off time for consumption, wages must increase more steeply with schooling when people value time more (i.e. when β is lower). We choose $\beta = 0.231$ to match the Mincerian return to schooling across occupations, which averages 12.7% across the six decades.²⁴


4.3. Composite Frictions vs. Occupational Preferences: Results

Now that we have estimates of θ , η , and β , we can decompose the dispersion in occupational propensities into the contribution of τ_{ig} and z_{ig} using equations (14) and (15). Figure 4a summarizes the mean (left panel) and the dispersion (right panel) of τ_{ig} across all 67 occupations for each of the three groups.²⁵ For white women, the mean τ_{ig} fell from about 10 in 1960 to around 3 in 2010 then leveled off. The decline continued through 1990 and slowed thereafter. The mean τ_{ig} facing black women declined from around 11 to about 4 from 1960-1980, then fell more slowly. Black men experienced a decline in mean τ_{ig} from around 3.5 to 2 from 1960-1980, after which no progress was made. The dispersion of τ_{ig} fell for all three groups. For white and black women the variance of $\ln \tau$ fell continuously from about 0.9 to about 0.6. For black men it fell from


²⁴We find the Mincerian return across occupations ψ from a regression of log average wages on average schooling across occupation-groups, with group dummies as controls. The average wage of group g in occupation i is proportional to $(1 - s_i)^{\frac{1}{3\beta}}$. We let s be years of schooling divided by a pre-work time endowment of 25 years. Thus the Mincerian return ψ +/- 1 year around mean schooling \bar{s} should satisfy $e^{2\psi} = \left(\frac{1-\bar{s}+.04}{1-\bar{s}-.04}\right)^{\frac{1}{3\beta}}$. The implied $\beta = \ln\left(\frac{1-s+.04}{1-s-.04}\right)/(6\psi)$. We set $\beta = 0.231$, the average of the implied β values across years. This method allows the model to approximate the Mincerian return to schooling across occupations.

²⁵The weights are the occupation's share of earnings out of total earnings for each group in each year.

Figure 4: Mean and Variance of τ_{ig} and z_{ig} by Group



(a) The Composite Barriers, τ_{ig}



(b) The Occupational Preferences, z_{ig}

Note: The left panel of each pair shows the average level of the frictions, weighted by total earnings in each occupation in each year. The right panel shows the variance of the log frictions, weighted in the same way. The axis scales are similar but not the same across panels so that it is possible to read the line labels on each panel.


about 0.4 to around 0.1 from 1960 to 1980 and stayed flat thereafter.

Figure 4b summarizes the mean and the dispersion of z_{ig} . Consistent with the evidence that the dispersion of wages across occupations narrowed over time, the dispersion of z_{ig} narrowed over time as well, most noticeably for black women.

To compare the magnitude of the decline in dispersion of z_{ig} vs. the dispersion in τ_{ig} , recall that each occupational propensity is proportional to $z_{ig}^{\frac{1-\eta}{3\beta}} / \tau_{ig}$ in equation (4). As we discussed earlier, we choose $\beta = \frac{0.693}{3}$, which combined with $\eta = 0.103$ gives us $\frac{1-\eta}{3\beta} = 1.29$. Therefore, to gauge the relative contribution of z_{ig} vs. τ_{ig} to the change in occupational shares, we need to multiply the dispersion of $\ln z_{ig}$ by 1.29. With this in mind, note that the decline in the standard deviation of z_{ig} is much smaller than the decline in the standard deviation of τ_{ig} . Intuitively, although wage gaps across occupations narrowed in the 1960s and 1970s, the magnitude of this decline was swamped by the decline in the gaps in occupational propensities. Therefore, although z plays a role, it is not the main force behind the changes in occupational shares. This finding will also be present in our model counterfactuals discussed later in the paper.

Figure 5 displays τ_{ig} for white women for a select subset of occupations. As shown, τ_{ig} was very high for women in 1960 in the construction, lawyer, and doctor occupations relative to the teacher and secretary occupations. τ_{ig} levels for white women lawyers and doctors in 1960 were at 10 or higher. If τ_{ig} reflected labor market discrimination only, the implication would be that women lawyers in 1960 were paid only one-tenth of their marginal product relative to their male counterparts. The model infers large τ_{ig} 's for white women in these occupations in 1960 because there were few white women doctors and lawyers in 1960, even after controlling for the gap in wages. Conversely, a white woman in 1960 was 24 times more likely to work as a secretary than was a white man. The model explains this huge gap by assigning a τ_{ig} below 1 for white women secretaries.

Over time, white women saw large declines in τ_{ig} for lawyers and doctors. As of 2010, white women faced composite frictions below 2 in the lawyer, doctor, and teacher occupations. The barrier facing white women in the construction sector remained large. This fact could be the result of women having a comparative disadvantage (relative to men) as construction workers, a possibility we consider later in our robustness checks.

Figure 5: Estimated Barriers (τ_{ig}) for White Women

Note: Author's calculations based on equation (14) using Census data and imposing $\theta = 2$ and $\eta = 0.103$.

4.4. Labor Market Discrimination vs. Human Capital Barriers

Our estimates of τ_{ig} above are a combination of labor market discrimination and human capital barriers. We distinguish changes in τ_{ig}^w from changes in τ_{ig}^h by exploiting life-cycle variation. As noted above, the key identifying assumption is that labor market discrimination equally affects all cohorts of the discriminated group in the labor market at the same point in time, whereas discrimination in schooling only affect individuals in the human capital accumulation stage of their life-cycle.

We proceed in two steps. First, the wage gap of cohort c and group g (relative to white men) in occupation i at time t relative to the wage gap at time c (when cohort c was young) is

$$\frac{gap_{ig}(c, t)}{gap_{ig}(c, c)} \propto \frac{1 - \tau_{ig}^w(t)}{1 - \tau_{ig}^w(c)} \quad (16)$$

The change in the wage gap depends on the change in τ_{ig}^w over time. Intuitively, if labor market discrimination diminishes over time, this raises the average wage (relative to white men) in occupations where the group previously faced discrimination. We will

therefore use the change in the wage gap to infer the change in τ_{ig}^w over time.


The second step is to back out the change in τ_{ig}^h as the residual of the change in τ_{ig} after controlling for the change in τ_{ig}^w . Specifically, we use $\frac{1+\tau_{ig}^h(t)}{1+\tau_{ig}^h(c)} = \left(\frac{\tau_{ig}(t)}{\tau_{ig}(c)} \cdot \frac{1-\tau_{ig}^w(t)}{1-\tau_{ig}^w(c)} \right)^{\frac{1}{\eta}}$. Figure 6 shows the data we use to disentangle τ_{ig}^h from τ_{ig}^w for white women (top panel) and black men (bottom panel). Specifically, the top panel shows the wage gap of white women relative to white men for different cohorts at differing points of their life cycle. Each line is a different cohort. Our model implies that a decline in τ_{ig}^h will shift up the intercept of the life cycle wage gap profiles as declining barriers to human capital allow individuals to accumulate more human capital that stays with them throughout their working lives. Conversely, a decline in τ_{ig}^w in a given year causes a steepening of a given life cycle profile. As τ_{ig}^w falls, the wage of a given group relative to white men converges during an individual's life cycle.

As seen from the top panel of Figure 6, there are large increases in the intercept of the wage gap profiles for white women suggesting declining τ_{ig}^h . Conversely, there are only slight changes in the slopes of the cohort profiles over time suggesting a potential smaller role for declining τ_{ig}^w . For black men (bottom panel), one sees both shifts in the intercepts and steepening slopes particularly during the 1960s to 1980s suggesting a role for both declining human capital and labor market frictions. It is this underlying variation that is at the heart of our decomposition of τ_{ig}^h from τ_{ig}^w in our model below.


We infer τ_{ig}^w by using differential wage growth across groups within an occupation between young and middle age and between middle age and old. For our base specification, we weight equally the variation in the data coming from young to middle age and from middle age to old. Our procedure gives us the *changes* in τ^w and τ^h over time. To get the initial levels, we need to determine how to split the composite τ in 1960. For our baseline specification, we assume an initial split of 50/50 in 1960. In all subsequent years, we let the data speak to the relative importance of τ_{ig}^h to τ_{ig}^w . Given that we do not have a direct moment in the data to pin down the initial split, we do extensive robustness around our baseline assumption.²⁶

²⁶We place one additional constraint on the τ breakdown to keep aggregate “revenue” from changing more than 10 percent of GDP over our sample period. This requires that we constrain τ^h to be no lower than -0.8 . This keeps subsidies for women secretaries from getting too large. In the Appendix we show that dropping this constraint has only a modest effect on the gains from changing barriers.

Figure 6: Wage Gaps For White Women Relative to White Men by Time and Cohort



(a) White Women



(b) Black Men

Note: Log wage gaps are shown for the life cycle of each cohort by connected line segments for young, middle-aged, and old periods.

4.5. Schooling and Production Function Parameters

We now turn to the schooling and production function parameters. The variable ϕ_i governs the occupation-specific return to time invested in human capital. Higher ϕ_i raises time spent in human capital accumulation (e.g., schooling). The higher time spent acquiring human capital for occupation i necessitates higher average wages in occupation i . We therefore infer ϕ_i from data on average wages in each occupation among young white men in each year. The $M - 1$ wage gaps pin down the *relative* ϕ_i values across occupations.²⁷ We set the *levels* of the ϕ_i 's so that “schooling” levels implied by the model (s_i values multiplied by 25 years, the pre-work ages) match average years of schooling across all occupations for young white men in the data.

Finally, w_i and ϕ_i collectively determine the observed share of young white men in each occupation and the average wage of young white men across all occupations.²⁸ Using the estimates of ϕ_i obtained from the data on wage differences across occupations, we pick w_i to exactly fit the observed occupational shares and the average wage for young white men in each year. The intuition is that, conditional on estimates of ϕ_i , the average wage for young white men pins down a weighted average of w_i . The differences in occupational shares then pin down the heterogeneity in w_i across occupations: occupations with a large share of young white men are ones where the price of skills w_i is high. With estimates of w_i , we then back out the technology parameter A_i .²⁹

4.6. Occupational Preferences

We now turn to the occupational preference parameters. We rewrite equation (8) as

$$z_{ig}(c) = \left(\frac{1 - LFP_{wm}(c)}{1 - LFP_g(c)} \right)^{\frac{3\beta}{\theta(1-\eta)}} \left(\frac{\overline{wage}_{ig}(c, c)}{\overline{wage}_{i,wm}(c, c)} \right)^{-3\beta}.$$

So with data on the labor force participation rate of a group (relative to white men) and the wage gap in one market occupation, we back out occupational preference z_{ig} in

²⁷We impute the average wage in the home sector for young white men. We use average schooling of young white men in the home sector, combined with the average years of schooling in market occupations and the Mincerian return to schooling we estimate for young men across market occupations.

²⁸Equations (4) and (6).

²⁹We need the elasticity of substitution among occupations (in aggregating to final output) σ to infer A_i from w_i . We choose $\sigma = 3$ as our baseline value, but we have no information on this parameter. Given this, we explore the robustness of our results to alternate values of σ . As we highlight below, extreme values of σ (e.g., 1.05 and 10) do not alter substantively our results.

that occupation. And since we know the occupational preference for a market sector relative to another market sector from the wage gaps (from equation (15)), we can then infer the *level* of the occupational preference term in each market sector.³⁰

4.7. Recap and Model Fit

Table 1 summarizes the identifying assumptions we make and the normalizations we choose. The latter are without loss of generality, but the former are crucial. We assume $\tau^h = \tau^w = 0$ for white men in all occupations. This implies that white men face no barriers in either the labor market or in their human capital acquisition. We also assume the mean talent draw is the same for all occupation-groups ($\bar{h}_{i,g} = 1$). We let technologies (the A_i 's) differ across occupations, so equating talent across occupations is without loss of generality. Our assumption that women, blacks and white men have the same talent in a given occupation is what matters. As discussed above and below, what ends up mattering quantitatively is that the relative talent of women, blacks and men does not change over time.

We also normalize τ_w and τ_h in the home sector to zero for all groups. This implies that individuals face no barriers in the home sector. We cannot identify the level of distortions, only their level in the market relative to the home sector. The decision to participate in the home sector is influenced by preferences for the home sector (for groups other than white men), the price of home-sector talent (w_{home}), and opportunities in the market (w_i , τ_w , and τ_h in the market occupations).

We normalize men's preferences to be the same for all market occupations ($z_{ig} = 1$). We use how wage gaps differ across market occupations to identify the preference of group g for market occupation i *relative* to that of white men in each year. We do not identify preference levels. Similarly, because we do not observe home wage gaps, we normalize preferences for the home sector to one for all occupation-groups in all years ($z_{home,g} = 1$).

Table 2 summarizes our key parameters. We choose the value of the Fréchet shape

³⁰To infer the price for home sector skill w_{home} , we use the equation for occupational propensity (4) for white men in the home sector. In turn, labor force participation of white men only depends on w_{home} , w_i in the market sectors, and ϕ_i (in all sectors). We choose ϕ_{home} to fit average years of schooling of young white men in the home sector in each year. We already have estimates of ϕ_i and w_i for the market occupations, so we use them along with data on the labor force participation rate of white men to infer the value of w_{home} .

Table 1: Identifying Assumptions and Normalizations

Parameter	Definition	Determination	Value
$\tau_{i,wm}^h$	Human capital barriers (white men)	Assumption	0
$\tau_{i,wm}^w$	Labor market barriers (white men)	Assumption	0
$\bar{h}_{i,g}$	Talent in each occupation (all groups)	Assumption	1
$\tau_{home,g}^h$	Home human capital barriers (all groups)	Assumption	0
$\tau_{home,g}^w$	Home labor market barriers (all groups)	Assumption	0
$z_{i,wm}$	Occupational preferences (white men)	Normalization	1
$z_{home,g}$	Home occupational preference (all groups)	Normalization	1

Table 2: Baseline Parameter Values

Parameter	Definition	Determination	Value
θ	Fréchet shape	Wage dispersion, Frisch elasticity	2
η	Goods elasticity of human capital	Education spending	0.103
σ	EoS across occupations	Arbitrary	3
β	Consumption weight in utility	Mincerian return to education	0.231

parameter ($\theta = 2$) based on wage dispersion within occupations and the elasticity of labor supply. We set the elasticity of human capital with respect to goods invested in human capital to match spending on education relative to earnings ($\eta = .103$). We set the elasticity of substitution across occupations arbitrarily ($\sigma = 3$), but check robustness to a wide range of alternative values below. Finally, we set weight on consumption vs. leisure in utility to match the observed Mincerian return to education across occupations ($\beta = 0.231$).

Table 3 summarizes the endogenous variables and the target data for their indirect inference. Some forcing variables depend on cohorts and some on time, but never both. Variables changing by cohort include the human capital barriers (τ^h), occupational preferences (z), and the elasticity of human capital with respect to time investment (ϕ). Labor market barriers (τ^w) and technology parameters (A) vary over time.

Table 3: Forcing Variables and Empirical Targets

Parameter	Definition	Empirical Target
$A_i(t)$	Technology by occupation	Occupations of young white men
$\phi_i(t)$	Time elasticity of human capital	Wages by occupation, white men
$\tau_{ig}^h(c)$	Human capital barriers	Occupations of the young, by group
$\tau_{ig}^w(t)$	Labor market barriers	Life-cycle wage growth, by group
$z_{ig}(c)$	Occupational preferences	Wage gaps by occupation for the young
$\gamma(1), \gamma(2)$	Experience terms	Age earnings profile of white men

Note: The variable values are chosen jointly to match the empirical targets.

Human capital barriers, labor market discrimination, and occupational preferences vary across occupation-groups. As discussed above, the structure of our model yields expressions that can be matched directly to empirical moments that allow us to infer the underlying driving forces of occupational choice.

Finally, Table 4 compares the data and the model's predictions for aggregate earnings per worker and labor force participation by year. Remember that the model only targets the occupational shares (and labor force participation rates) of the young. Despite this, predicted per-capita earnings and labor force participation rates in the model are not very far from the data. For example, in 2010 predicted earnings in the model is within 3 percent of the actual earnings in the data. In the model, labor force participation rate increases by 15.1 percentage points between 1960 and 2010. The actual increase between 1960 and 2010 is 16 percentage points.

5. Main Results

We can now answer the key question of the paper: how much of the overall growth from 1960 to 2010 can be explained by the changing labor market outcomes of blacks and women during this time period? Real earnings per person in our census sample grew by 1.8 percent per year between 1960 and 2010. According to our model, this observed earnings growth can come from four sources. First, growth in per capita

Table 4: Model versus Data: Earnings and Labor Force Participation

Year	Earnings Data	Earnings Model	LFP Data	LFP Model
1960	14,718	14,720	0.599	0.599
1970	20,802	20,728	0.636	0.614
1980	22,153	22,348	0.702	0.651
1990	28,281	27,956	0.764	0.719
2000	33,888	34,718	0.747	0.743
2010	37,036	38,165	0.759	0.750

Note: This table shows average market earnings per worker in 2009 dollars and labor force participation in the Census/ACS data alongside the corresponding model values by year.

earnings comes from general occupational productivity growth (changing A 's). Second, earnings growth results from growth in the returns to schooling resulting in more human capital attainment (changing ϕ 's). Third, changing in preferences for each occupation including the home sector can reallocate labor across occupation resulting in earnings growth (changing z 's). Fourth, growth in the relative share of each group in the working age population can also mechanically result in changing earnings per capita (changing q 's). Finally, as described in Section 2.9., changing gender and race specific barriers to occupational choice in both the labor market and human capital market can result in economic growth (changing τ 's).

The goal of our model is to assess how much of economic growth can be attributed to the changing τ 's. We answer this key economic question by holding τ 's fixed while allowing the A 's, ϕ 's, z 's and q 's to evolve. For each variable, we calculate the difference between the actual path in the data and the counterfactual “no change in τ 's” path to gauge the effect of changing τ 's.

5.1. Income and Productivity Gains

The results of our counterfactuals are shown in the first column of Table 5. The changes in τ 's account for 27% of growth from 1960 to 2010 in GDP per person. This includes market output and our estimate of home sector output. If one compares the first and

second columns of Table 5, one can see only a modest effect on productivity growth from changing preferences for each market occupation (z_{ig}). Collectively, these results imply that most of the growth in the economy over the last half century was due to increases in A_i and ϕ_i over time, but an important part (27%) is attributable to reduced frictions.

Why can't changing preferences for market work explain women's rising labor force participation or women's movement into high skilled occupations relative to white men? If women simply did not like some occupations, the model says they would have been paid more in occupations in which they were underrepresented. The data show no such patterns. The gender (wage) gap was no lower in skilled occupations, and it did not fall faster in skilled occupations as the share of women rose. So while preference changes did result in the reallocation of women and blacks across occupations, the reallocation was not systematically in a way to result in substantive economic growth.

Figure 7 shows the time series decomposition of growth. The top line shows growth in GDP per person. The bottom line is growth if the τ 's were held fixed. Not surprisingly, the productivity effect of the τ 's have grown over time. Additionally, our results suggest that productivity growth would have been negative during the 1970s had it not been for the reduction in labor market barriers to blacks and women during that time period.

Table 5 also reveals that changing frictions account for a bigger share of growth in market earnings per person (36%) and market GDP per person (38%). A big part of earnings growth reflects rising labor force participation of women in response to falling barriers. Aggregate labor force participation rates rose steadily in the data, from 60% in 1960 to 76% in 2010, primarily due to increased female labor supply. Changes in the τ 's account for more than 100% of this increase, according to Table 5.

Note that market earnings and market GDP differ due to changing "revenue" from labor market discrimination over time. Figure 8 shows how such revenue evolves from 1960 to 2010. The figure also displays how combined revenue from both barriers combined shrinks from around 5% of GDP in 1960 to -3% of GDP in 2010. This evolution boosts consumption of workers relative to GDP over the sample.

Changing τ 's contributed to faster growth still in market GDP per person, at 38%. The changing τ 's account for only 9% of the 1960–2010 growth in market GDP *per worker*. These gains come from better allocation of talent and more investment in

Table 5: Share of Growth due to Changing Frictions (all ages)

	Share of growth accounted for by			
	τ^h and τ^w	τ^h, τ^w, z	τ_h only	τ_w only
GDP per person (home+market)	26.7%	27.2%	24.5%	5.7%
Market earnings per person	36.1%	29.2%	16.1%	28.5%
Market GDP per person	38.1%	36.6%	30.9%	11.1%
Labor force participation	126.6%	134.0%	27.3%	90.6%
Market GDP per worker	9.2%	7.8%	32.0%	-15.0%


Note: Entries in the table show the share of growth in the model attributable to changing frictions under various assumptions. The variables are τ^h (human capital frictions), τ^w (labor market frictions), and z (occupational preferences).

human capital in response to falling barriers.


The last two columns of Table 5 report growth contributions from falling barriers to human capital accumulation (τ^h only) vs. falling labor market discrimination (τ^w only). Falling human capital barriers alone would have accounted for 24.5% of growth, and falling labor market discrimination around 6% of growth. Falling labor market discrimination loom larger for growth in market earnings (28.5% of growth). The reason is that declining discrimination in the labor market contributes directly to earnings growth *relative* to output growth. When we look at growth in market GDP per person, declining barriers to human capital are again more important (31% of growth) than diminishing labor market discrimination (11%).

Table 5 suggests that falling labor market discrimination drove much (over 90%) of the rise in labor force participation. Falling barriers to human capital accumulation played a lesser role since human capital is also useful in the home sector, albeit less so than in some market occupations. The breakdown into contributions from human capital vs. labor market barriers is also revealing for why the contribution to growth in market GDP per worker is modest (9%). Falling human capital barriers, on their own, would have explained 32% of growth. But falling labor market discrimination actually lowered growth (-15%) by enticing workers with marginal talent to move out of the home sector and into market occupations.

Figure 7: GDP per person, Data and Model Counterfactual



Note: The graph shows the cumulative growth in GDP per person (home+market), in the data (overall) and in the model with no changes in τ 's as in Table 5).

Figure 8: Revenue from τ as share of GDP in the Model

Note: The graph shows the employer revenue from discrimination in the labor and human capital markets as a percent of GDP.

Table 6: Share of Growth due to Changing Frictions (young only)

	Share of growth accounted for by τ^h and τ^w
Market GDP per person (young)	49.8%
Market earnings per person (young)	47.7%
Consumption per person (home+market, young)	35.2%
Utility per person (consumption equivalent, young)	50.7%

Note: Entries in the table show the share of growth in the model attributable to changing frictions. The variables are τ^h (human capital barriers) and τ^w (labor market frictions).

In four of the five rows in Table 5, the combined effect of changing the τ^h 's and τ^w 's is smaller than the sum of the effects from eliminating them individually. The explanation for this is that misallocation is convex in barriers. Reducing one of the barriers individually yields the largest gains to be had by moving highly misallocated workers to the right occupation.

Table 6 focuses on the young to characterize how welfare changes across cohorts. We can follow them through all six waves of the Census/ACS data. Moreover, the young are most responsive to the changes in the combined τ 's because they can optimize their human capital decisions in response to the changes. The table shows the effects on market GDP per person, market earnings per person, market+home consumption per person, and consumption-equivalent utility per person. Changing τ 's explain 50% of growth in market output for the young, a higher fraction than for all workers (38%) because the older cohorts cannot alter their earlier schooling and human capital decisions. Falling barriers explain 48% of growth in market earnings per person for the young. The changing τ 's account for a smaller share of home+market consumption growth (35%) than market earnings growth (48%) for the young, as the falling barriers enticed young women into the market.

Strikingly, the last row of Table 6 says that changing τ 's accounts for over 50% of growth in consumption-equivalent utility for the young. Much of growth in market GDP and earnings per person came from a rising fraction of time spent investing in human capital (in part due to rising ϕ 's), which came at a utility cost. In contrast,

Table 7: Wage Gaps and Earnings by Group and Changing Frictions

	— Share of growth accounted for by —		Full Model
	τ^h and τ^w	τ^h, τ^w, z	
Wage gap, WW	113.5%	103.8%	121.4%
Wage gap, BM	85.6%	102.6%	123.4%
Wage gap, BW	75.9%	82.5%	110.6%
Earnings, WM	-10.7%	-10.8%	100.4%
Earnings, WW	85.0%	89.3%	104.3%
Earnings, BM	28.4%	24.9%	102.6%
Earnings, BW	55.6%	57.4%	100.2%
LF Participation	126.6%	134.0%	94.6%

Note: Entries in the table show the share of growth in the model attributable to changing frictions and other variables. The frictions are τ^h (human capital) and τ^w (labor market), and z are occupational preferences. The last column reports the share of observed growth explained by the full model solution, including the A and ϕ variables.

the efficiency gains from a better allocation of talent entail no such cost. The falling employer revenue from discrimination also translate into utility gains for workers.

Table 7 shows how the changing τ 's affect wage gaps and earnings across groups. The last column shows that our model does fairly well in predicting the changing wage gaps over time. Our model, collectively, over-predicts slightly the rising wages of women and blacks relative to white men during the 1960–2010 period. For women, the changing τ 's more than explain the shrinking gender gap in wages observed in the data. The model also says that, in the absence of changing τ 's, the rising labor force participation rate of women would have *widened* the gender gap by bringing in women with less of a comparative advantage in market occupations (compared to other women in the market, not men). For blacks, the changes in the τ 's explain most of the shrinking wage gaps. Declining barriers to human capital attainment and declining labor market discrimination was primarily responsible for the declining gender and racial wage gaps during the last fifty years.

Table 7 also says that the changing τ 's actually lowered wage growth of white men.

Table 8: Share of growth in GDP per person due to different groups

1960–2010	τ^h and τ^w	τ^h only	τ^w only
All groups	26.7%	24.5%	5.7%
White women	22.1%	20.1%	4.9%
Black men	1.0%	0.7%	0.4%
Black women	2.1%	2.1%	0.5%

Note: Entries are the share of growth in GDP per person (home+market) from changing frictions for various groups over different time periods. The variables are τ^h (human capital barriers), and τ^w (labor market frictions).

This is because falling barriers to women and blacks in high skilled occupations caused white men to shift to lower wage occupations. Falling barriers account for 85% of earnings growth for white women, 28% for black men, and 56% for black women. For men (both black and white), wage growth was driven primarily by changes in technology and skill requirements (A 's and ϕ 's).

Table 8 breaks down the growth from changing τ 's into contributions by each group. Changes in the τ 's of white women were much more important than changes in the τ 's of blacks in explaining growth in home+market output per person during the 1960-2010 period. This is primarily because white women are a much larger share of the population. Table 8 also shows that falling pre-labor market barriers to human capital accumulation contributed much more to growth did declining labor market barriers.

Finally, we can ask: how much additional growth could be achieved by reducing the frictions all the way to zero? If the remaining frictions in 2010 were removed entirely, we calculate that GDP today would be 16.8% higher. These remaining gains result from the fact that, even in 2010, occupational barriers exist across groups. Now, the middle-aged and old in 2010 cannot respond to this hypothetical lifting of all barriers in terms of their occupational choices and human capital accumulation. When we look at output from the young, who can respond in 2010, the potential gains are larger at 19.1%. By comparison, the gains from eliminating all barriers on the young in 1960 would have been 69.9%. Thus the vast majority of gains from eliminating barriers have already been reaped. This is one reason to be less optimistic about growth after 2010 than in

the half-century leading up to it.

5.2. Model Gains vs. Back-of-the-Envelope Gains

Our baseline estimate in Table 5 suggests that τ^w and τ^h account for 27% of the gains in home+market GDP per person. Is this number large or small relative to what one might have expected? We have two ways of thinking about this question. First, in the log-normal approximation to the model with only τ^w variation that we presented back in Section 2.9., the elasticity of GDP to 1 minus the mean of τ^w is $q_w \cdot \frac{\eta}{1-\eta}$. If we assume that the share of women in the population q_w is $1/2$ and $\eta = 0.1$ then this elasticity is $\frac{1}{2} \cdot \frac{1}{9}$. Figure 4a showed that the mean of the composite τ of women fell from about 10 in 1960 to 3 in 2010. This decline in $\bar{\tau}$ can thus account for a 7% increase in GDP per person.³¹ Figure 4a also shows that $\text{Var} \ln \bar{\tau}$ fell from about 0.9 to 0.6 from 1960 to 2010. In the log-normal approximation to the model, the semi-elasticity of GDP to $\text{Var} \ln \tau$ is $q_w \cdot \frac{1}{2} \cdot \frac{\theta-1}{1-\eta} \approx 0.3$.³² A 0.3 decrease in the variance of $\ln \tau$ thus could explain an 8% increase in home+market GDP per person. Thus, according to this back-of-the-envelope calculation, changing τ 's boosted GDP about 15%. The overall increase of GDP per person in our setup was about 138%, so the changing τ explains $0.15/1.38 \approx 11\%$ of growth. This is significantly lower than the 27% contribution we estimate when not assuming a lognormal distribution of barriers.

A second way to answer the question is to compare our 27% model-based growth contribution to what one would infer from the falling gaps in earnings per person for women and blacks relative to white men. The narrowing gaps in earnings per person — including both declining wage gaps and rising labor force participation — mechanically account for 37% of growth in earnings per person.³³ Why is our model-based estimate of 27% lower than this back-of-the-envelope calculation of 37%? The back-of-the-envelope calculation assumes no general equilibrium effect of falling frictions on the earnings growth of white men. Yet we reported above white men's wages fell 11% relative to what they would have done without the changing barriers facing women and blacks (see Table 7). Moreover, this back-of-the-envelope calculation assumes

³¹ $\frac{1}{2} \cdot \frac{1}{9} \cdot \ln(10/3) \approx .07$.

³² $\frac{1}{2} \cdot \frac{1}{2} \cdot \frac{2-1}{1-.1} \approx 0.28$

³³ For this calculation, we held fixed earnings per person relative to white men at 1960 levels, and found only 63% as much growth in earnings per person as seen in the data.

that earnings gaps would not have changed in the absence of falling frictions. That is, this calculation implicitly attributes the entire decline in earnings gaps to changing frictions. As we show below, other forces such as changes in occupational productivity and returns to schooling have also had an effect in explaining changing wage gaps between groups over time.

6. Robustness

In this section, we explore the robustness of our results to alternate parameterizations, identifying assumptions, and data moments used to discipline the model.

6.1. Alternative Parameter Values

Table 9 explores robustness of our productivity gains to different parameter values. For each set of parameter values considered, we recalculate the τ , z , A , and ϕ values so that the model continues to fit the occupation shares, wage gaps, etc. The first row of Table 9 replicates the gains under baseline parameter values for comparison. The next row considers a lower value for the Fréchet shape parameter θ , which is inversely related to the dispersion of comparative advantage across occupations. With $\theta = 1.5$ rather than the baseline $\theta = 2$, changing barriers explain modestly more of growth (31%) than in the baseline (27%).

Recall that our baseline θ was estimated from wage dispersion within occupation-groups. This might overstate the degree of comparative advantage because some of the wage variation is due to *absolute* advantage. We thus entertain a much higher value ($\theta = 4$) than in our baseline ($\theta = 2$). With this higher θ , the share of growth from changing τ 's falls to 16% (vs. 27% in the baseline). Our explanation is that less discrimination is needed to explain occupational choices when comparative advantage is weak. Even with this higher value of θ , however, declining τ 's explain over one-seventh of growth in GDP per person over the last half-century.

Table 9 also varies η , the elasticity of human capital with respect to goods invested in human capital. Intuitively, the gains from falling human capital barriers are greater the higher is η : the gains rise slightly from 26% with $\eta = 0.05$ to 27% with our baseline $\eta = 0.103$ to 28% with $\eta = 0.20$.

Table 9: Robustness to Parameter Values

	GDP per person (home+market) growth accounted for by		
	τ^h and τ^w	τ^h alone	τ^w alone
Benchmark	26.7%	24.5%	5.7%
$\theta = 1.5$	31.0%	28.7%	5.2%
$\theta = 4$	15.8%	22.2%	4.6%
$\eta = 0.05$	25.7%	24.1%	4.8%
$\eta = 0.20$	28.0%	24.5%	7.8%
$\sigma = 1.05$	27.5%	24.8%	5.2%
$\sigma = 10$	25.8%	24.0%	6.0%

Note: Entries show the share of growth in the model attributable to changing frictions τ^h (human capital) and τ^w (labor market). The baseline parameter values are $\theta = 2$, $\eta = 0.103$, and $\sigma = 3$.

The last rows of Table 9 show the (in)sensitivity of the results to the elasticity of substitution σ between occupations in production. The gains to changing τ 's with $\sigma = 1.05$ (close to Cobb-Douglas) and $\sigma = 10$ are within one percent of the gains under the baseline of $\sigma = 3$. Although not shown in Table 9, the gains are not at all sensitive to β , the weight placed on time vs. goods in utility.

The moderate sensitivity of our results to θ , η , σ and β may seem puzzling. But remember that, as we entertain different parameter values, we simultaneously change the A 's and τ 's to fit observed wages and employment shares of the young in each occupation and group in each year. In Table 10 we vary θ while holding all other parameter values and forcing variables fixed (the A 's, ϕ 's, τ 's, etc.). That is, we do not re-calibrate. Consistent with the intuition provided in Section 2.9., the gains from changing τ 's rise dramatically as we raise θ . When ability is less dispersed (θ is higher), comparative advantage is weaker and the allocation of talent is *more* sensitive to changing τ 's. The higher is θ , the more occupational decisions are distorted by given barriers, and hence the bigger the gains from removing them.

Table 10: Changing Only the Dispersion of Ability

Value of θ	GDP per person growth accounted for by τ^h and τ^w
1.5	-108.5%
2 (baseline)	26.7%
3	125.2%
4	162.9%

Note: Here we hold fixed all other parameters and forcing variables at their baseline values. Unlike the other robustness checks in Table 9, we do not re-calibrate. These results confirm that for a given time path of the barriers, a higher value of θ leads declining barriers to account for more of growth.

6.2. Alternative Identifying Assumptions

One key identifying assumption that underlies our estimation is that any innate talent differences between men and women are constant over time. Under this assumption, changes in occupational sorting and wage gaps between groups inform us about changes in the τ 's and z 's. In our base specification, we go even farther and assume there are no innate talent differences between group in any period ($\bar{h}_{ig} = 1$ for all i and g in all time periods). In this section, we explore alternative assumptions while still holding relative talent fixed over time.

Table 11 shows how our results change with alternative assumptions about \bar{h} across groups within different occupations over time. The first row of the table redisplay our baseline estimates. The second row relaxes 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 over time because of technological progress. To see the potential importance of this story, we go to the extreme of assuming no frictions at all faced by women in any of the occupations where physical strength is arguably important (i.e., $\tau_{ig}^h = \tau_{ig}^w = 0$ for women in these occupations). These occupations

Table 11: Robustness to Alternative Assumptions about Group Differences in Talent

	GDP per person (home+market) growth accounted for by τ^h and τ^w
Benchmark	26.7%
No frictions in “brawny” occupations	25.5%
No frictions in 2010	24.2%

Note: Entries are the share of growth in the model attributable to changing frictions τ^h (human capital) and τ^w (labor market). A key identifying assumption is that the any talent differences across groups, to the extent they exist, are constant over time. In our baseline specification, we assume that $\bar{h}_{ig} = 1$ for all occupations and all groups. In other words, that there are no innate talent differences between groups. The first row of the table recounts are baseline estimates. In the second row, we allow men and women to have different \bar{h} 's in “brawny” occupations. Specifically, we assume no gender specific τ 's in these occupations. Instead, we allow the \bar{h} 's to evolve to exactly fit the quantity data for these occupations. “No frictions in 2010” (the third row) assumes that there are no frictions in 2010 for any group, so that differences in \bar{h}_{ig} explain all group differences in that year; we then calculate τ 's for earlier years assuming the mean value of the distribution of market skills in 2010 apply to earlier years.

include construction, firefighters, police officers, and most of manufacturing.³⁴ We estimate differences in \bar{h}_{ig} for young women to fully explain their allocation to these occupations in 1960, 1970, . . . , 2010. As shown in Table 11, the fraction of growth explained by changing frictions falls only slightly from 26.7% to 25.5% with this alternative identifying assumption. Our results are not sensitive to this alternative because most of the gains we attribute to changing τ 's come from the rising propensity of women to become lawyers, doctors, scientists, professors, and managers — occupations where physical strength is not important.

The last row in Table 11 makes a more extreme assumption. In this alternative, we allow all groups to have different levels of innate talent in all occupations. We assume, however, that these innate talent differences are constant over time. Specifically, we assume all group differences among the young in 2010 reflect talent rather than distortions. We set the 2010 τ 's to zero for all groups and all occupations and assume differences in \bar{h}_{ig} fully account for group differences in occupational choice among the

³⁴Rendall (2010) classifies occupations based on the importance of physical strength, and we define brawny occupations for our analysis as those occupations in the top half of her brawny distribution.

young in 2010.³⁵ We keep talent in prior years at the 2010 values for each group, but back out distortions in earlier years. In essence, this specification allows for arbitrary talent differences between men and women to fit the 2010 data. Under this more flexible alternative, eliminating the τ 's in the earlier years still generates 24% of growth in home+market GDP per person. Thus our gains are not an artifact of assuming the allocation of talent was far from optimal in 2010.

These exercises highlight our key identifying assumption. What is important is not that different groups have the same level of innate talent in all occupations. Instead, what is important is that, whatever the talent differences are across groups, those talent differences remain constant over time. This assumption is particularly important for high skilled occupations like doctors and lawyers and less important for “brawny” occupations like construction workers.

Another assumption that facilitates our identification is that white men face no labor market or human capital frictions. An alternative assumption might be that there was no discrimination in 1960 at all, but growing discrimination *against* men and *in favor* of women since then. If we assume women and men have the same mean talent, as we do in our baseline specification, this would imply identical average wages and occupational distributions for women and men in 1960. This is something we do not observe in the data. Assuming relative talent stays constant over time, this alternative would also require women to earn increasingly more than men and be increasingly overrepresented in high skill occupations after 1960. All of these predictions are at odds with the patterns documented above. If men and women have the same level of innate talent, the data strongly reject the hypothesis that men have been increasingly discriminated against over time.

Another alternative would be to assume discrimination in *favor* of men and no discrimination against women in 1960, with the discrimination in favor of men abating over time. This would fit the facts on wages and sorting over time, and would imply falling misallocation. But it is not isomorphic to our baseline assumption. First, it would imply falling education spending by men over the decades. Second, it would entail huge subsidies for men that diminish over time. When we calibrate the model, we find that earnings of men must exceed their marginal product by orders of magnitude.

³⁵We still normalize $\bar{h}_{iwm} = 1$ in 2010.

The implied total subsidy to men would be multiples of 1960 GDP. Men must be paid massive subsidies in order to induce so many of them, relative to women, to choose high skilled occupations. Earnings would need to vastly exceed GDP in 1960, which of course we do not observe. Such an extreme outcome does not arise under our baseline assumption because no revenue is collected from qualified women who are driven out of occupations by discrimination.

Yet another alternative would be to assume — contrary to our presumption — that women are somehow innately less talented than men, supposedly explaining women's lower wages and underrepresentation in skilled occupations in 1960. Rising discrimination in *favor* of women since 1960 might then account for the closing gaps between men and women. This hypothesis would entail *rising* misallocation and a drag on aggregate growth.

Data on individual test scores suggests women are not less talented than men. The Armed Forces Qualifying Test (AFQT) was administered in both the NLSY 1979 and the NLSY 1997. The NLSY tracks a sample of individuals who were 12-16 years old when the surveys started. The AFQT scores in the NLSY are very similar for men and women in both 1979 and 1997.³⁶ According to these scores, women seem no less talented than men in their early teens. If we condition on working, women likewise have similar scores to men in both 1979 and 1997.³⁷ If one believes the story of rising discrimination in favor of women, one would have expected the relative test scores of working women to fall along with their rising participation rates. AFQT scores do not support the hypothesis that women are innately less talented than men.

Collectively, these results suggest that alternate assumptions do not fit aspects of the data as well as our baseline. We therefore prefer our baseline assumption that women and blacks faced human capital and labor market frictions in 1960 relative to white men, and that these frictions fell over time.

Table 12: Additional Robustness

	GDP per person growth accounted for by		
	τ^h and τ^w	τ^h only	τ^w only
Benchmark	26.7%	24.5%	5.7%
Wage gaps halved	23.8%	20.9%	5.9%
Zero wage gaps	21.2%	18.1%	6.1%
Half the return to experience	29.2%	27.6%	4.5%
2/3, 1/3 split of $\tau_{i,g}$ in 1960	21.7%	20.7%	3.6%
1/3, 2/3 split of $\tau_{i,g}$ in 1960	29.5%	25.0%	9.4%
50/50 split of $\tau_{i,g}$ in all years	25.6%	21.7%	7.1%
No constraint on τ^h	29.3%	26.8%	5.7%

Note: See notes to Table 5. GDP includes home+market. The baseline splits τ in 1960 evenly into τ^h and τ^w in 1960, but not in future years. The baseline also constrains τ^h to be at most -0.8 .

6.3. Other Robustness

Table 12 explores an additional set of robustness of exercises. The first row repeats our benchmark results for comparison. The next two rows show that the productivity gains we estimate are *not* proportional to the gender and race wage gaps we fed into the model. We can halve the wage gaps in all years, or even eliminate them in all years, and the implied τ 's still explain 24% or 21% of growth in home+market GDP per person, vs. 27% in the baseline. One reason is that misallocation of talent by race and gender can occur even if average wages are similar. The misallocation of talent is tied to the *dispersion* in the τ 's, whereas the wage gaps are related to both the mean and variance of the τ 's. Another reason is that the wage gap for white women would have widened in the absence of the changing τ 's. A key take away from this exercise is that productivity gains from changing labor market discrimination and barriers to human

³⁶In 1979, the average normalized AFQT score was 54.3 for white men and 53.6 for white women. In 1997, the respective averages were 55.5 and 57.4.

³⁷In 1979, the average normalized AFQT score for working white men was 51.8 and for working white men was 52.3. In 1997, the respective averages were 52.6 and 54.9.

capital accumulation cannot be gleaned from the wage gaps alone.

Another assumption we make in our base specification is that the returns to experience are constant across groups and occupations over time. We want to stress that allowing for general returns to experience is not adding much to our inference. The fourth row in Table 12 illustrates this point. Specifically, in this robustness exercise, we cut productivity growth over the life cycle (old/middle and middle/young) in half for each group. Such a change barely alters our baseline results.

The final four rows of Table 12 considers additional robustness checks. In our benchmark we split the composite τ_{ig} in 1960 evenly into τ^w and τ^h . Our procedure estimates changes in τ 's over time but we need to make an assumption on the initial split between τ^h and τ^w . If we put more weight (2/3) on τ^w we account for 22% of growth in home+market GDP per person, versus 27% in the baseline. If we put less weight (1/3) on τ^w we account for 30% of growth. We also tried splitting τ_{ig} evenly into τ^w and τ^h in all years, which differed modestly from the benchmark. Finally, our benchmark case constrains the values of τ^h to be no smaller than -0.8 . If we put no constraint on how negative τ^h can get (e.g., how large subsidies for white women secretaries can get), we explain 29% of growth vs. 27% in the baseline.³⁸

7. Further Model Implications

While our model is stylized in many respects, it is able to match at least three other important facts that were not targeted in the estimation: trends in female labor supply elasticities, cross-state variation in survey measures of racial discrimination, and changes in educational attainment by group.

7.1. Trends in Female Labor Supply Elasticities

Using data from the Current Population Survey, Blau and Kahn (2007) estimate that there was a dramatic decline in female labor supply elasticity during the 1980–2000

³⁸In Appendix Table D3 we consider a few more robustness checks. Our benchmark divides τ into τ^h and τ^w based on wage growth for women and blacks as they age from both young to middle-aged and middle-aged to old. We can explain more of the gains if we do the division based only on the young to middle-aged (31% vs. 27% in the baseline). The benchmark also adjusts our estimate of wage growth due to accumulated experience for rising participation rates of men over the years of our sample. We explain more (32%) of growth in GDP per person if we make no such adjustment.

period. Helpful for comparing with the predictions of our model, they report female labor supply elasticities specifically for 25-34 year olds. We compare the model's implied labor supply elasticities — equal to $\theta(1 - LPF_g)$ — for young white women to the estimated labor supply elasticities reported in Blau and Kahn (2007). Using our baseline θ , the model matches both the level and the trend female labor supply elasticities well. Blau and Kahn (2007) report labor supply elasticities for women aged 25-34 of 0.75, 0.60 and 0.35, respectively in 1980, 1990, and 2000 — a change of 0.40 over the time period. Our comparable model estimates for young women are 0.90, 0.70, and 0.65 for the three years - a change of 0.25 over the time period. Our estimates are only slightly higher in levels than the Blau and Khan estimates over the three years with a roughly similar trend.

Nothing in our model is calibrated to match either the level or the trend in labor supply elasticities for women. As discussed earlier, we estimated θ to match the labor supply elasticity of men in 1980. With that parameter pinned down, our model implies that women's labor supply elasticity is only a function of female labor force participation. The fact that we can roughly match the level of the labor supply elasticity for young women in three different time periods suggests that our model is consistent with empirical moments outside the ones we used to calibrate the model.

7.2. Cross State Measures of Discrimination

There are very few micro-based measures of discrimination to which we can compare our estimated τ 's. One such exception is the recent work by Charles and Guryan (2008). Charles and Guryan (CG) used data from the General Social Survey (GSS) to construct a measure of the taste for discrimination against blacks for every state. The GSS asks a large nationally representative sample of individuals about their views on a variety of issues. A series of questions have been asked over the years assessing the respondents attitudes towards race. For example, questions were asked about individuals' views on cross-race marriage, school segregation, and the ability for homeowners to discriminate with respect to home sales. Pooling together survey questions from the mid 1970s through the early 1990s and focusing only a sample of white respondents, Charles and Guryan make indices of the extent of racial discrimination in each state.³⁹ Higher

³⁹We focus on their marginal discrimination measure. The concept of the marginal discriminator comes from Beckers theory of discrimination. If there are 10 percent of blacks in the state labor market, it is only

values of the CG discrimination measure imply more discrimination. They compute their measure for 44 states.

Figure 9 shows a simple scatter plot between the CG measure of discrimination and our measure τ_{bm} at the state level.⁴⁰ Each observation in the scatter plot is a U.S. state where the size of the circle represents the number of black men within our Census sample. We also show the weighted OLS regression line on the figure. As seen from the figure, there is a very strong relationship between our measures of τ_{bm} and the CG discrimination index. The adjusted R-squared of the simple scatter plot is 0.6 and the slope of the regression line is 0.45 with a standard error of 0.06. Places we identify as having a high τ_{bm} are the same places Charles and Guryan find as being highly discriminatory based on survey data from the GSS. The findings in Figure 9 provide additional external validity that our procedure is measuring salient features of the U.S. economy over the last five decades.


7.3. Trends in Educational Attainment

As we report in Table 13, our benchmark model does fairly well in capturing trends in educational attainment seen in the U.S. Census from 1960 to 2010. The data exhibit convergence toward the educational attainment of white men: by 0.9 years for white women, 1.6 years for black men, and 1.65 years for black women. The model, meanwhile, features a fraction s of time spent in education during a pre-work period. We multiply this fraction by 25 years to arrive at educational attainment predicted by the model. (Recall that our working years start with 25-34 year olds.) Because the τ 's for blacks and women fell faster in higher schooling occupations, the changing τ 's contributed in a major way to educational convergence. For white women, the

the discrimination preferences of the white person at the 10th percentile of the white distribution that matters for outcomes (with the first percentile being the least discriminatory).

⁴⁰From our earlier estimates, we compute a composite τ measure for black men relative to white men in each U.S. state. To ensure we have enough observations in each state, we make a few simplifying assumptions. First, we assume that there are no cohort effects in our composite measure of τ . This allows us to pool together all cohorts within a year when computing our measure of τ . Next, we collapse our 67 occupations to 20 occupations; see Appendix Table C2. Also, we pool together data from 1980 and 1990; we do this because the CG discrimination measure is based on data pooled from the GSS between 1977 and 1993. We then aggregate $\tau_{i,bm}$ from our 20 different occupations to one measure of τ_{bm} for each state by taking a weighted average of the occupation level τ s where the weights are based on share of the occupations income (for the country as whole) out of total income across all occupations (for the country as a whole). Finally, we exclude states with an insufficient number of black households to compute our measure of τ_{bm} . Given the CG restrictions from the GSS and our restrictions from the Census data, we are left with 37 states.

Figure 9: Model τ 's for Black Men vs. Survey Measures of Discrimination, by U.S. State



Note: Figure plots measures of our model's implied composite τ 's for black men for each state using pooled data from the 1980 and 1990 census (x-axis) against survey-based measures of discrimination against blacks for each state as reported in Charles and Guryan (2008). The Charles and Guryan data are compiled using data from the General Social Survey between 1977 and 1993. We use their marginal discrimination measure for this figure. See text for additional details.

Table 13: Years of Educational Attainment for the Young

	Actual 1960	Actual 2010	Actual Change	Model Change	Due to τ 's
White men	11.78	13.47	1.69	1.90	-0.29
White women	11.46	14.04	2.59	2.07	0.66
Black men	9.49	12.76	3.28	3.13	1.36
Black women	10.08	13.42	3.34	2.18	0.86

Note: Actual educational attainment is from the U.S. Census. We multiply the model's fraction of time spent in schooling s by 25 years (a pre-work time endowment) to arrive at educational attainment predicted by the model.

changing τ 's fully account for the trend and then some (0.95 years narrowing in the model, vs. 0.9 in the data). For black men, falling frictions likewise explain all of the narrowed schooling gap (1.65 years in the model vs. 1.6 years in the data). For black women, declining distortions generate 1.15 years of the 1.65 year catch-up in schooling relative to white men from 1960–2010.

8. Conclusion

How does discrimination in the labor market and barriers to the acquisition of human capital for white women, black men, and black women affect occupational choice? And what are the consequences of the altered allocation of talent for aggregate income and productivity? We develop a framework to tackle these questions empirically. This framework has three building blocks. First, we use a standard Roy model of occupational choice, augmented to allow for labor market discrimination, barriers to the acquisition of human capital and occupation-specific preferences. Second, we impose the assumption that an individual's talent in each occupation follows an extreme value distribution. Third, we embed the Roy model in general equilibrium to account for the effect of occupational choice on the price of skills in each occupation and to allow for the effect of technological change on occupational choice. We use synthetic cohort data measuring changes in relative occupational sorting and wage gaps across time to discipline our model. A key identifying assumption is that the distribution of innate

talent across groups is constant over time.

We apply this framework to measure the changes in barriers to occupational choice facing women and blacks in the U.S. from 1960 to 2010. We find large reductions in these barriers, concentrated in high-skilled occupations. We then use our general equilibrium setup to isolate the aggregate effects of the reduction in occupational barriers facing these groups. Our calculations suggest that falling barriers may explain one-quarter of aggregate growth in GDP per person.

It should be clear that this paper provides only a preliminary answer to these important questions. The general equilibrium Roy model we use is a useful place to start, but it is possible that a different framework can do a better job. We abstract from allowing for correlations between an individual's absolute advantage and their comparative advantage. Additionally, the ease with which our model can be matched to observable moments of the data is facilitated by our assumption that comparative advantage is distributed according to a Fréchet distribution. These assumptions have the benefits of tractability but may abstract from other important features of the data. Some structure is needed in order to assess how the substantive changes in occupational sorting across gender and race affected US economic growth. We provided one such framework as a starting point. Our results suggest that the decline in occupational and human capital barriers to women and blacks was a very important source of growth to the US economy and the leveling out of changes may be one reason why growth has slowed down. However, we realize that our model is only a launching off point to address these important questions and expect some of our assumptions used for tractability to be relaxed as the literature progresses.

Finally, we have focused on the gains from reducing barriers facing women and blacks over the last fifty years. But we suspect that barriers facing children from less affluent families and regions have worsened in the last few decades. If so, this could explain both the adverse trends in aggregate productivity and the fortunes of less-skilled Americans in recent decades. We hope to tackle some of these questions in future work.

References

Adão, Rodrigo, "Worker heterogeneity, wage inequality, and international trade: Theory and evidence from Brazil," *Unpublished paper, University of Chicago*, 2016.

- Akcigit, Ufuk, John Grigsby, and Tom Nicholas, "The Rise of American Ingenuity: Innovation and Inventors of the Golden Age," 2017. NBER Working Paper No. 4096.
- Altonji, Joseph G and Rebecca M Blank, "Race and Gender in the Labor Market," *Handbook of Labor Economics*, 1999, 3, 3143–3259.
- Arrow, Kenneth, "The Theory of Discrimination," in "Discrimination in Labor Markets," Princeton University Press, 1973, pp. 3–33.
- Bailey, Martha J., "More Power to the Poll: The Impact of Contraceptive Freedom on Women's Life Cycle Labor Supply," *Quarterly Journal of Economics*, February 2006, 121 (1), 289–320.
- , Brad Hershbein, and Amalia R. Milleri, "The Opt-In Revolution? Contraception and the Gender Gap in Wages," NBER Working Papers 17922, National Bureau of Economic Research, Inc March 2012.
- Becker, Gary, *The Economics of Discrimination*, Chicago, IL: The University of Chicago Press, 1957.
- , *The Economics of Discrimination*, University of Chicago Press, 1957.
- Bertrand, Marianne, "New Perspectives on Gender," *Handbook of labor economics*, 2011, 4, 1543–1590.
- 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.
- Blau, Francine D and Lawrence M Kahn, "Changes in the Labor Supply Behavior of Married Women: 1980-2000," *Journal of Labor Economics*, 2007, 25 (3), 348–352.
- , Marianne A Ferber, and Anne E Winkler, *The Economics of Women, Men, and Work*, Prentice-Hall Englewood Cliffs, NJ, 2013.
- , Peter Brummund, and Albert Yung-Hsu Liu, "Trends in occupational segregation by gender 1970–2009: Adjusting for the impact of changes in the occupational coding system," *Demography*, 2013, 50 (2), 471–492.
- 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 Beckers The Economics of Discrimination," *Journal of Political Economy*, 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.
- Chetty, Raj, Adam Guren, Day Manoli, and Andrea Weber, "Does indivisible labor explain the difference between micro and macro elasticities? A meta-analysis of extensive margin elasticities," in "NBER Macroeconomics Annual 2012, Volume 27," University of Chicago Press, 2012, pp. 1–56.
- Costa, Dora L, "From Mill Town to Board Room: The Rise of Women's Paid Labor," *Journal of Economic Perspectives*, 2000, 14 (4), 101–122.
- Eaton, Jonathan and Samuel S. Kortum, "Technology, Geography, and Trade," *Econometrica*, 2002, 70 (5), 1741–1779.
- Fernández, Raquel, "Cultural Change as Learning: The Evolution of Female Labor Force Participation over a Century," *American Economic Review*, 2013, 103 (1), 472–500.
- , Alessandra Fogli, and Claudia Olivetti, "Mothers and Sons: Preference Formation and Female Labor Force Dynamics," *The Quarterly Journal of Economics*, 2004, pp. 1249–1299.
- Goldin, Claudia, *Understanding the Gender Gap: An Economic History of American Women*, New York, NY: Oxford University Press, 1990.
- and Lawrence F. Katz, "The Power of the Pill: Oral Contraceptives and Women's Career and Marriage Decisions," *Journal of Political Economy*, 2002, 110 (4), 730–770.
- and —, "The Most Egalitarian of All Professions: Pharmacy and the Evolution of a Family-Friendly Occupation," NBER Working Paper 18410 September 2012.
- Greenwood, Jeremy, Ananth Seshadri, and Mehmet Yorukoglu, "Engines of Liberation," *The Review of Economic Studies*, 2005, 72 (1), 109–133.
- Johnson, George E and Frank P Stafford, "Alternative Approaches to Occupational Exclusion," *Routledge Research in Gender and Society*, 1998, 2, 72–90.
- Karabel, Jerome, *The Chosen: the Secret History of Admission and Exclusion at Harvard, Yale, and Princeton*, New York, NY: Houghton Mifflin Harcourt, 2005.

- Lagakos, David and Michael E Waugh, "Selection, agriculture, and cross-country productivity differences," *American Economic Review*, 2013, 103 (2), 948–80.
- 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.
- Neal, Derek A and Armin Rick, "The Prison Boom and the Lack of Black Progress After Smith and Welch," *NBER Working Paper*, 2014.
- and William R Johnson, "The Role of Premarket Factors in Black-White Wage Differences," *The Journal of Political Economy*, 1996, 104 (5), 869–895.
- Pan, Jessica, "Gender Segregation in Occupations: The Role of Tipping and Social Interactions," *Journal of Labor Economics*, 2015, 33 (2), 365–408.
- Phelps, Edmund S, "The Statistical Theory of Racism and Sexism," *The American Economic Review*, 1972, 62 (4), 659–661.
- 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.

Online Appendix to “The Allocation of Talent and U.S. Economic Growth”

(Not for publication)

Chang-Tai Hsieh, Erik Hurst, Charles I. Jones, Peter J. Klenow

April 6, 2018

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. Occupational Choice

The individual’s utility from choosing a particular occupation, $U(\tau_{ig}, w_i, \epsilon_i)$, is proportional to $(\bar{\gamma} \tilde{w}_{ig} \epsilon_i)^{\frac{3\beta}{1-\eta}}$, where $\tilde{w}_{ig} \equiv \frac{\bar{h}_{ig} w_i s_i^{\phi_i} [(1-s_i) z_{ig}]^{\frac{1-\eta}{3\beta}}}{\tau_{ig}}$ and $\bar{\gamma} \equiv 1 + \gamma(2) + \gamma(3)$ is the sum of the experience terms. The solution to the individual’s problem, then, involves picking the occupation with the largest value of $\tilde{w}_{ig} \epsilon_i$. To keep the notation simple, we will suppress the g subscript in what follows.

Without loss of generality, consider the probability that the individual chooses occupation 1, and denote this by p_1 . Then

$$\begin{aligned} p_1 &= \Pr[\tilde{w}_1 \epsilon_1 > \tilde{w}_s \epsilon_s] \quad \forall s \neq 1 \\ &= \Pr[\epsilon_s < \tilde{w}_1 \epsilon_1 / \tilde{w}_s] \quad \forall s \neq 1 \\ &= \int F_1(\epsilon, \alpha_2 \epsilon, \dots, \alpha_M \epsilon) d\epsilon, \end{aligned} \tag{17}$$

where $F_1(\cdot)$ is the derivative of the cdf with respect to its first argument and $\alpha_i \equiv \tilde{w}_1 / \tilde{w}_i$.

Recall that

$$F(\epsilon_1, \dots, \epsilon_M) = \exp \left[\sum_{s=1}^M \epsilon_s^{-\theta} \right].$$

Taking the derivative with respect to ϵ_1 and evaluating at the appropriate arguments gives

$$F_1(\epsilon, \alpha_2 \epsilon, \dots, \alpha_M \epsilon) = \theta \epsilon^{-\theta-1} \cdot \exp \left[\bar{\alpha} \epsilon^{-\theta} \right] \tag{18}$$

where $\bar{\alpha} \equiv \sum_s \alpha_s^{-\theta}$.

Evaluating the integral in (17) then gives

$$\begin{aligned}
 p_1 &= \int F_1(\epsilon, \alpha_2\epsilon, \dots, \alpha_M\epsilon) d\epsilon \\
 &= \frac{1}{\bar{\alpha}} \int \bar{\alpha}\theta\epsilon^{-\theta-1} \cdot \exp[\bar{\alpha}\epsilon^{-\theta}] d\epsilon \\
 &= \frac{1}{\bar{\alpha}} \cdot \int dF(\epsilon) \\
 &= \frac{1}{\bar{\alpha}} \\
 &= \frac{1}{\sum_s \alpha_s^{-\theta}} \\
 &= \frac{\tilde{w}_1^\theta}{\sum_s \tilde{w}_s^\theta}
 \end{aligned}$$

A similar expression applies for any occupation i , so we have

$$p_i = \frac{\tilde{w}_i^\theta}{\sum_s \tilde{w}_s^\theta}.$$

Proof of Proposition 2. Average Quality of Workers

Efficiency units of labor of an individual of cohort c in occupation i at time t is given by $h_i(c, t) = \bar{h}_i s(c)^{\phi_i(t)} e_i(c)^\eta$. Using the results from the individual's optimization problem, it is straightforward to show that

$$h_i(c, t) \epsilon_i = s_i(c)^{\phi_i(t)} \gamma(t-c) \left(\frac{\eta s_i(c)^{\phi_i(c)} w_i(c) (1 - \tau_i^w(c)) \bar{h}_i \bar{\gamma}}{1 + \tau_i^h(c)} \right)^{\frac{\eta}{1-\eta}} \epsilon_i^{\frac{1}{1-\eta}}.$$

Therefore, *average* efficiency units of labor in an occupation is given by

$$\mathbb{E}[h_i(c, t) \epsilon_i \mid \text{choose } i] = s_i(c)^{\phi_i(t)} \gamma(t-c) \left(\frac{\eta s_i(c)^{\phi_i(c)} w_i(c) (1 - \tau_i^w(c)) \bar{h}_i \bar{\gamma}}{1 + \tau_i^h(c)} \right)^{\frac{\eta}{1-\eta}} \mathbb{E}\left[\epsilon_i^{\frac{1}{1-\eta}} \mid \text{choose } i\right].$$

Let ϵ^* denote ability in the chosen occupation. We need to know the distribution of ϵ^* raised to some power. Let $y_i \equiv \tilde{w}_i \epsilon_i$ denote the key occupational choice term. Then

$$y^* \equiv \max_i \{y_i\} = \max_i \{\tilde{w}_i \epsilon_i\} = \tilde{w}^* \epsilon^*.$$

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

$$\begin{aligned}
\Pr [y^* < z] &= \Pr [y_i < z] \quad \forall i \\
&= \Pr [\epsilon_i < z/\tilde{w}_i] \quad \forall i \\
&= F \left(\frac{z}{\tilde{w}_1}, \dots, \frac{z}{\tilde{w}_M} \right) \\
&= \exp \left[- \sum_s \tilde{w}_s^\theta z^{-\theta} \right] \\
&= \exp \{ -mz^{-\theta} \}.
\end{aligned}$$

That is, the extreme value also has a Fréchet distribution, where $m \equiv \sum_s \tilde{w}_s^\theta$.

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

$$G(x) \equiv \Pr [\epsilon^* < x] \equiv \exp \left[-m^* x^{-\theta} \right]$$

where $m^* \equiv \sum_{s=1}^M (\tilde{w}_s/\tilde{w}^*)^\theta = 1/p^*$.

The last thing we need is an expression for the expected value of the chosen occupation's ability raised to some power. Let λ be some positive exponent. Then,

$$\begin{aligned}
\mathbb{E}[\epsilon^{*\lambda}] &= \int_0^\infty \epsilon^{*\lambda} dG(\epsilon^*) \\
&= \int_0^\infty \theta \left(\frac{1}{p^*} \right) \epsilon^{*(-\theta-1+\lambda)} e^{-\left(\frac{1}{p^*}\right)\epsilon^{*- \theta}} d\epsilon^*
\end{aligned}$$

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

$$\begin{aligned}
\mathbb{E}[\epsilon^{*\lambda}] &= \left(\frac{1}{p^*} \right)^{\lambda/\theta} \int_0^\infty x^{-\frac{\lambda}{\theta}} e^{-x} dx \\
&= \left(\frac{1}{p^*} \right)^{\lambda/\theta} \Gamma \left(1 - \frac{\lambda}{\theta} \right).
\end{aligned}$$

Applying this result to our model, we have

$$\mathbb{E} \left[\epsilon_i^{\frac{1}{1-\eta}} \mid \text{choose } i \right] = \left(\frac{1}{p_{iq}} \right)^{\frac{1}{\theta} \cdot \frac{1}{1-\eta}} \Gamma \left(1 - \frac{1}{\theta} \cdot \frac{1}{1-\eta} \right).$$

Substituting this expression into the equation above for $\mathbb{E}[h_i(c, t)\epsilon_i \mid \text{choose } i]$ leads to equation (5).

Proof of Proposition 3. Occupational Wage Gaps

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

Proof of Proposition 4. Relative Propensities

The proof of this proposition is straightforward after substituting the results from propositions 1 and 3 into the expression for relative propensities $\frac{p_{ig}}{p_{i,wm}}$.

Proof of Proposition 5. Relative Labor Force Participation

This proposition is an application of proposition 4 to the home sector, assuming no distortions for white men in all sectors and no distortions in the home sector for all groups.

B Heterogeneity in Occupational Preferences

This section lays out an alternative model where individuals differ in terms of their *preference* for an occupation instead of their talent in an occupation. For simplicity, suppose individuals live for one period, $\phi = 0$, $\bar{h} = 1$, $z = 1$, $\eta = 0$, and $\tau^h = 0$, so the only occupational distortion is τ^w . Utility is now given by $U = (1 - \tau_{ig}^w)w_i\epsilon$ where ϵ now represents idiosyncratic *preferences* in the chosen occupation. The share of a group in an occupation is still given by equation (4) where $\tilde{w}_{ig} \equiv (1 - \tau_{ig}^w)w_i$. The equation for average worker quality (equation (5)) now refers to the average of the idiosyncratic preference ϵ . Specifically, average idiosyncratic preferences of workers in an occupation is given by:

$$\mathbb{E}[\epsilon_i \mid \text{choose } i] = \Gamma p_{ig}^{-\theta}$$

And the equation for average earnings in an occupation (equation (6)) now refers to average *utility*:

$$\begin{aligned} \mathbb{E} [(1 - \tau_{ig}^w)w_i\epsilon_i | \text{choose } i] &= (1 - \tau_{ig}^w)w_i \mathbb{E} [\epsilon_{ig} | \text{choose } i] \\ &= \left(\sum_{j=1}^M \tilde{w}_{jg} \right)^\theta \end{aligned}$$

So average *utility* is the same across occupations. But average *earnings* differs across occupations. Market earnings in an occupation is given by $(1 - \tau_{ig}^w)w_i$ so a key prediction of this alternative model is that the average earnings in an occupation is positively related to the occupational share. Intuitively, average ϵ is lower and average earnings are higher in an occupation with a large occupational share, and these two effects exactly offset.

C Identification and Estimation

This section explains how we identify and estimate the frictions and other parameters, carried out in the program `EstimateTauZ.m`.

C1. Key Equations

To estimate the model, we add one additional feature to the model. In our base case, we assume the return to experience is the same for all occupations, groups, and cohorts. In our robustness checks, however, these parameters may be allowed to vary. We thus index γ (and the sum of the experience terms $\bar{\gamma}$) by group g and occupation i in the equations that follow.

The key equations underlying our estimation are listed below.

- Occupational Choice

$$p_i = \frac{\tilde{w}_i^\theta}{\sum_s \tilde{w}_s^\theta}$$

where $\tilde{w}_{ig} \equiv \frac{w_i \bar{h}_{ig} \bar{\gamma}_{ig} s_i^{\phi_i} [(1 - s_i) z_{ig}]^{\frac{1-\eta}{3\beta}}}{\tau_{ig}}$

and $\tau_{ig} \equiv \frac{(1 + \tau_{ig}^h)^\eta}{1 - \tau_{ig}^w}$

- Average Quality

$$\mathbb{E}[h_{ig}(c, t) \epsilon_{ig}(c)] = s_i(c)^{\phi(t)} \gamma_{ig}(t-c) \left[\eta \frac{1 - \tau_{ig}^w(c)}{1 + \tau_{ig}^h(c)} w_i(c) \bar{h}_{ig} \bar{\gamma}_{ig} s_i(c)^{\phi(c)} \right]^{\frac{\eta}{1-\eta}} \Gamma \left(\frac{1}{p_{ig}(c)} \right)^{\frac{1}{\theta} \frac{1}{1-\eta}}$$

- Average Wage

$$\begin{aligned} \overline{\text{wage}}_{ig}(c, t) &\equiv (1 - \tau_{ig}^w(t)) w_i(t) \gamma_{ig}(t-c) \mathbb{E}[h_{ig}(c, t) \epsilon_{ig}(c)] \\ &= \Gamma \bar{\eta} [m_g(c)]^{\frac{1}{\theta} \frac{1}{1-\eta}} [(1 - s_i(c)) z_{ig}(c)]^{-\frac{1}{3\beta}} \frac{(1 - \tau_{ig}^w(t)) w_i(t) \gamma_{ig}(t-c) s_i(c)^{\phi(t)}}{(1 - \tau_{ig}^w(c)) w_i(c) \bar{\gamma}_{ig} s_i(c)^{\phi(c)}} \end{aligned}$$

where $m_g(c) = \sum_{i=1}^M \tilde{w}_{ig}(c)^\theta$

- Relative Propensity

$$\frac{p_{ig}(c, c)}{p_{i,wm}(c, c)} = \left(\frac{\bar{h}_{ig}}{\bar{h}_{i,wm}} \right)^\theta \left(\frac{\tau_{ig}(c, c)}{\tau_{i,wm}(c, c)} \right)^{-\theta} \left(\frac{\overline{\text{wage}}_{ig}(c, c)}{\overline{\text{wage}}_{i,wm}(c, c)} \right)^{-\theta(1-\eta)} \left(\frac{\bar{\gamma}_{ig}}{\bar{\gamma}_{i,wm}} \right)^{\theta\eta}$$

C2. Estimate Wages & Schooling from data of young white men

The following refers to the program `solveWMfor_w_phi.m`. This program uses data on wages and occupational shares of young white men to estimate w_i and ϕ_i .

First, we pin down the level of ϕ in the farming occupation and in the home sector such that the average years of schooling in the model is equal to the years observed in the data. Using this normalization, we back out s in farming and the home sector from the following equation:

$$s^{\text{farm}} = \frac{1}{1 + \frac{1-\eta}{3\beta\phi_{\text{farm}}}}$$

Second, we use equation for the average wage and s_{farm} (from the previous step) to back out m_{wm} . After omitting the indices for cohort and time, the specific equation is:

$$m_{wm} = \left[\frac{\overline{wage}_{i,wm} (1 - s_i)^{\frac{1}{3\beta}}}{\Gamma \bar{\eta}} \cdot \frac{\bar{\gamma}_{i,wm}}{\gamma_{i,wm}} \right]^{\theta(1-\eta)}$$

where $i = farm$. Furthermore, we need to make an initial guess about the return to experience term γ (We describe later how we do this).

Third, we estimate s_i for the other occupations (non-farming) from the equation we use above to back out m_{wm} from data on wages in the farm sector. In this case, we use data on the average wage on the occupation, the aggregate labor force participation rate for the group and cohort, and the estimate for m_{wm} we obtained from step 2 to back out the s_i that fits the wage equation. This value of s_i then allows us to back out ϕ_i for the occupation. This approach also provides an estimate of \overline{wage} for white men in the home sector.

Fourth, we estimate w_i from the observed occupational shares. After some algebra, the occupational share equation can be expressed as:

$$w_i = \frac{[p_{i,wm} \cdot m_{wm}]^{\frac{1}{\theta}}}{\bar{\gamma}_{i,wm} \cdot s_i^{\phi_i} (1 - s_i)^{\frac{1-\eta}{3\beta}}}$$

Again, $\tau = 1$ and $z = 1$ for white men so these two terms do not show up.

Fifth, we estimate T and \bar{T} (remember we assumed a value for the experience terms for the previous steps) from the change in the average wage of a given cohort and occupation over time. Specifically, the ratio of the average wage in an occupation at time t to that at time c is:

$$\frac{\overline{wage}_{i,wm}(c, t)}{\overline{wage}_{i,wm}(c, c)} = \frac{w_i(t)(\gamma_{ig}(t - c))s_i^{\phi(t)}}{w_i(c)s_i^{\phi(c)}}$$

We estimate $\gamma_{i,wm}(t - c)$ from the change in the average wage in an occupation, after controlling for the change in w_i and the returns to schooling. In our base case, we assume $\gamma_{i,wm}(t - c)$ is the same across all occupations and cohorts so simply take the average across all occupations and cohorts.

C3. Estimating τ

The next part of the estimation obtains the composite of the distortions $\tau_{ig} \equiv \frac{(1+\tau^h)^\eta}{1-\tau^w}$. Remember we assume $\tau_{i,wm}^w = \tau_{i,wm}^h = 0$ and $\bar{h}_{ig} = \bar{h}_{i,wm}$. These two normalizations imply that we can express relative propensities as:

$$\tau_{ig} = \hat{p}_{ig}^{-\frac{1}{\theta}} \cdot \widehat{\text{wage}}_{ig}^{-(1-\eta)} \cdot \hat{\gamma}_{ig}^\eta$$

where a “hat” denotes the value of the variable relative to white men. In this equation, $\widehat{\text{wage}}_i$ and \hat{p}_{ig} are data and $\hat{\gamma}_{ig}$ and $\widehat{\gamma}_{ig}$ are estimated from the previous step.

C4. Estimating τ^w , τ^h , and z

The next step is to estimate z and the components of τ (i.e. τ^w and τ^h) for the other groups (non-white men). This is done in the program `estimatetauz.m`. We define α as the Cobb-Douglas split of τ that recovers $1 - \tau^w$. Specifically,

$$\tau^\alpha = \frac{1}{1 - \tau^w} \text{ and } \tau^{1-\alpha} = (1 + \tau^h)^\eta$$

This implies the following definitions of τ^w and τ^h as a function of τ and α :

$$\begin{aligned} \tau^w &= 1 - \tau^{-\alpha} \\ \tau^h &= (\tau^{1-\alpha})^{\frac{1}{\eta}} - 1 \end{aligned}$$

Our estimation of τ^w and τ^h is expressed in terms of α .

First, the home sector is assumed to be undistorted, so τ^w and τ^h for that sector are set to zero. We then use the “relative propensity” key equation for \hat{p}_{ig} at the start of this section, together with the wage in the home sector for white men, to recover the wage at home for the other groups.

Second, we normalize $z = 1$ for the home sector and back out m_g for the group based on data on the average wage in the home sector. Specifically, after some manipulation, the average wage equation for the sector can be expressed as:

$$m_g(c) = \left[\frac{\widehat{\text{wage}}_{\text{home},g}(c, c)(1 - s_{\text{home}}(c))^{\frac{1}{3\beta}}}{\Gamma \bar{\eta}} \right]^{\theta(1-\eta)}$$

For the other sectors, we use the same wage equation to back out z . Specifically, the wage equation can be expressed as:

$$z_{ig} = \frac{1}{1 - s_i} \cdot \left[\Gamma \bar{\eta} m_g^{\frac{1}{\theta} \frac{1}{1-\eta}} \frac{1}{\bar{\gamma}_{ig}} \frac{1}{\overline{\text{wage}}_{ig}} \right]^{3\beta}$$

We now have z for all cohorts and τ^w and τ^h for the young cohort in 1960. What is left is to pin down τ^w and τ^h for the years after 1960. From the ‘‘Average Wage’’ equation in our list of key equations, we can express wage growth in a given group-occupation as

$$\frac{\overline{\text{wage}}_{ig}(c, t+1)}{\overline{\text{wage}}_{ig}(c, t)} = \frac{1 - \tau_{ig}^w(t+1)}{1 - \tau_{ig}^w(t)} \cdot \frac{w_i(t+1)}{w_i(t)} \cdot \frac{\gamma_{ig}(t+1-c)}{\gamma_{ig}(t-c)} \cdot \frac{s_i(c)^{\phi_i(t+1)}}{s_i(c)^{\phi_i(t)}} \quad (19)$$

We use solve this equation for $\tau_{ig}^w(c, t+1)$ and this becomes our estimate since everything else in the equation is now observed. Then, $\tau_{ig}^h(t+1)$ is obtained from $\tau_{ig}^w(t+1)$. In other words, τ^w is the time effect in wage growth, while τ^h is the cohort effect.

There are two small modifications we make to this in practice. First, we set the minimum value of τ^h to -0.80 , though we relax this constraint in the robustness checks (without this constraint, the revenue required to subsidize women secretaries with τ^h gets implausibly large).

Second, in our model, occupations are chosen when young, so all groups have the same labor-force participation when middle-aged and old. In the data, this is clearly not the case. Therefore, we strip out from wage growth for a given group-occupation using our model’s estimate of the selection effect from differential participation. Based on the ‘‘Relative Propensity’’ equation in our ‘‘Key Equation’’ list, this effect has an elasticity of $\theta(1 - \eta)$. Absent data on labor-force participation by group, we use a common adjustment across all occupations to obtain the wage growth estimate used in equation (19):

$$\left(\frac{\text{wagegrowth}_{ig}}{\text{wagegrowth}_{i,wm}} \right)^{\text{for estimation}} = \left(\frac{\text{wagegrowth}_{ig}}{\text{wagegrowth}_{i,wm}} \right)^{\text{data}} \left(\frac{\text{LFPgrowth}_{ig}}{\text{LFPgrowth}_{i,wm}} \right)^{\frac{1}{\theta(1-\eta)}} \quad (20)$$

We also report results without making this adjustment in our robustness checks.


D Appendix Tables and Figures

Table D1: Sample Statistics by Census Year

	1960	1970	1980	1990	2000	2010
Sample Size	624,579	674,059	3,943,034	4,607,829	5,084,891	2,889,513
Share of Sample:						
White Men, Age 25-34	0.142	0.148	0.185	0.172	0.132	0.133
White Men, Age 35-44	0.155	0.140	0.131	0.156	0.162	0.136
White Men, Age 45-55	0.136	0.145	0.118	0.108	0.139	0.157
White Women, Age 25-34	0.158	0.160	0.194	0.175	0.130	0.134
White Women, Age 35-44	0.171	0.150	0.138	0.159	0.163	0.137
White Women, Age 45-55	0.145	0.157	0.127	0.113	0.143	0.162
Black Men, Age 25-34	0.015	0.016	0.021	0.023	0.021	0.022
Black Men, Age 35-44	0.015	0.015	0.014	0.019	0.024	0.022
Black Men, Age 45-55	0.013	0.014	0.012	0.012	0.017	0.023
Black Women, Age 25-34	0.019	0.021	0.026	0.027	0.024	0.024
Black Women, Age 35-44	0.018	0.019	0.018	0.023	0.027	0.025
Black Women, Age 45-55	0.015	0.016	0.015	0.015	0.020	0.026

Note: Data comes from the 1960-2000 U.S. Censuses and the pooled 2010 American Community Survey (ACS). Samples restricted to black and white, men and women between the ages of 25 and 54. Those in the military are excluded. Also, excluded are those not working but actively searching for a job. Sample shares are weighted using Census and ACS provided sample weights.

Figure D1: Changes in Occupational Wage Gaps vs. Changes in Relative Propensities: White Women 1960-2010



Note: The figure shows the relationship between the change in (log) occupational earnings gaps for white women compared to white men (both in the young cohort) and the change in the relative propensity to work in the occupation for the two groups, $p_{i,ww}/p_{i,wm}$, between 1960 and 2010. A simple regression line through the scatter plot in the top panel yields a slope coefficient of 0.06 with a standard error of 0.05.

Table D2: Occupation Categories for our Base Occupational Specifications

0.	Home Sector (0)	34.	Police (12)
1.	Executives, Administrative, and Managerial (1)	35.	Guards (12)
2.	Management Related (2)	36.	Food Preparation and Service (13)
3.	Architects (3)	37.	Health Service (6)
4.	Engineers (3)	38.	Cleaning and Building Service (13)
5.	Math and Computer Science (3)	39.	Personal Service (13)
6.	Natural Science (4)	40.	Farm Managers (14)
7.	Health Diagnosing (5)	41.	Farm Non-Managers (14)
8.	Health Assessment (6)	42.	Related Agriculture (14)
9.	Therapists (6)	43.	Forest, Logging, Fishers, & Hunters (14)
10.	Teachers, Postsecondary (7)	44.	Vehicle Mechanic (15)
11.	Teachers, Non-Postsecondary (8)	45.	Electronic Repairer (15)
12.	Librarians and Curators (8)	46.	Misc. Repairer (15)
13.	Social Scientists and Urban Planners (4)	47.	Construction Trade (15)
14.	Social, Recreation, and Religious Workers (4)	48.	Executive (14)
15.	Lawyers and Judges (5)	49.	Precision Production, Supervisor (16)
16.	Arts and Athletes (4)	50.	Precision Metal (16)
17.	Health Technicians (9)	51.	Precision Wood (16)
18.	Engineering Technicians (9)	52.	Precision Textile (16)
19.	Science Technicians (9)	53.	Precision Other (16)
20.	Technicians, Other (9)	54.	Precision Food (16)
21.	Sales, All (10)	55.	Plant and System Operator (17)
22.	Secretaries (11)	56.	Metal and Plastic Machine Operator (17)
23.	Information Clerks (11)	57.	Metal & Plastic Processing Operator (17)
24.	Records Processing, Non-Financial (11)	58.	Woodworking Machine Operator (17)
25.	Records Processing, Financial (11)	59.	Textile Machine Operator (17)
26.	Office Machine Operator (11)	60.	Printing Machine Operator (17)
27.	Computer & Communication Equip. Operator (11)	61.	Machine Operator, Other (19)
28.	Mail Distribution (11)	62.	Fabricators (18)
29.	Scheduling and Distributing Clerks (11)	63.	Production Inspectors (18)
30.	Adjusters and Investigators (11)	64.	Motor Vehicle Operator (19)
31.	Misc. Administrative Support (11)	65.	Non Motor Vehicle Operator (19)
32.	Private Household Occupations (13)	66.	Freight, Stock, & Material Handlers (18)
33.	Firefighting (12)		

Notes: Our 66 market occupations are based on the 1990 Census Occupational Classification System. We use the 66 sub-headings (shown in the table) to form our occupational classification. See <http://www.bls.gov/nls/quex/r1/y97r1cbka1.pdf> for the sub-heading as well as detailed occupations that correspond to each sub-heading. As discussed in the text, we include the home sector as an additional occupation. When computing racial barriers at the state level, we use only twenty broader occupations. The number in parentheses refers to how we group these 67 occupations into the twenty broader occupations for the cross-state analysis. For example, all occupations with a 11 in parentheses refers to the fact that these occupations were combined to make the 11th occupation in our broader occupation classification.

Table D3: Robustness on Life Cycle Wages

	GDP per person growth accounted for by		
	τ^h and τ^w	τ^h only	τ^w only
Benchmark	26.7%	24.5%	5.7%
Wage growth young to middle-aged	31.1%	29.5%	3.6%
Wage growth middle-aged to old	16.8%	14.2%	7.3%
No experience adjustment	32.2%	31.0%	3.4%

Note: See notes to Table 5. GDP includes home+market. The benchmark divides τ into τ^h and τ^w based wage growth for women and blacks as they age from both young to middle-aged and middle-aged to old. The benchmark also adjusts our estimate of wage growth due to accumulated experience for rising participation rates of men over the years of our sample.