

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

August 26, 2016 – Version 4.0

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 that innate talent for these professions is unlikely to differ across groups, the occupational distribution in 1960 suggests that a substantial pool of innately talented black men, black women, and white women 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 seminar participants at Boston College, Boston University, Brown, Chicago, Columbia, Cornell, CREI, Duke, Essex, Harvard, LSE, Michigan, MIT, an NBER EFG meeting, Oxford, Penn, Princeton, Stanford, Toronto, UBC, UCL, UCLA, USC and Yale for helpful comments. Jihee Kim, Munseob Lee, Huiyu Li, Zihong 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.

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, in 1960, 94 percent of doctors and lawyers were white men. By 2010, the fraction was just over 60 percent.¹ Similar changes occurred throughout the economy during the last fifty years, particularly among highly-skilled occupations. A large literature attempts to explain these facts.² Yet no formal study has assessed the effect of these changes on aggregate productivity. Given that innate talent for many professions is unlikely to differ across groups, the occupational distribution in 1960 suggests that a substantial pool of innately talented blacks and women 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 for women and blacks 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 and chooses the occupation with the highest expected utility. A key assumption we make is that the distribution of innate talent is similar across sex/race groups. This implies that the likelihood that a woman or black man receives a high talent draw in a particular occupation (e.g., doctors, lawyers, nurses) is the same as that for a white man. Put differently, we assume that occupational differences between groups are not driven by differences in innate talent.³

Instead, we consider four forces that potentially drive differences in occupational choices across groups. First, the model incorporates 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

¹These statistics are based on the 1960 Census and the 2010-2012 American Community Surveys. We discuss the sample in more detail below.

²See, 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).

³We relax this assumption in our robustness section when we allow men to have an absolute advantage in "brawny" occupations.

class, the only private sector job she could get immediately after graduating was as a legal secretary (Biskupic, 2006). We model labor market discrimination as an occupation-specific wedge between the wages of white women, black men, and black women and their marginal products. This “tax” is a proxy for many common formulations of taste-based and statistical discrimination found in the literature.⁴

Second, our model allows for barriers to human capital investment across groups and occupations. 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, and social norms that steer groups away from certain occupations.⁵

Third, we allow preferences for the various occupations to differ across groups. Such differences have been highlighted in the work of, among others, Johnson and Stafford (1998), Altonji and Blank (1999), and Bertrand (2011).

Finally, the model allows for group-specific differences in productivity at home. This factor, which could also have been modeled as a preference for the home sector, can capture changes in social norms related to women working at home and changes in fertility over time.⁶

⁴See, for example, Becker (1957), Phelps (1972) and Arrow (1973). A summary of such theories can be found in Altonji and Blank (1999).

⁵Here 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. See Fernandez (2012) and Fernandez, Fogli and Olivetti (2004b) on the role of social forces in women’s occupational choice. 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.

⁶The literature on changes in female labor supply due to changes in productivity, preferences, and social norms is extensive. See Fernández, Fogli and Olivetti (2004a) 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 relative changes in labor force participation rates for black men. Neal and Rick (2014), for example, highlights the importance of differential incarceration rates in explaining changes in labor market outcomes between black and white men.

We then show how observable data on wages, occupational choices, and labor force participation can be used to tease apart the relative importance of these four forces. Wage differences across groups within an occupation 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 disamenity.

Next, the life cycle structure of our model allows us 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 life-cycle patterns of wages and employment to distinguish occupation-specific human capital barriers (which are akin to "cohort" effects) from occupation-specific labor market discrimination (akin to "time" effects). Finally, we use differences in labor force participation across groups to discipline home sector productivity/preferences.

Economic growth in our model comes from two sources. First are forces common to all groups, including occupation-specific productivity growth and changes in occupational returns to schooling. We discipline the importance of these factors by matching changes in the number and wages of white men working in each occupation over time. Second, growth occurs because of forces that affect other groups relative to white men. These factors include changes in labor market discrimination, changes in occupational barriers to human capital attainment, changes in occupational preferences, and changes in preferences/productivity for the home sector. Declining barriers affect growth both by increasing human capital investment and by improving the allocation of talent.

We find that roughly one-quarter of growth in U.S. GDP per person between 1960 and 2010 can be explained by declining labor market discrimination and barriers to human capital attainment facing white women, black men and black women. These factors also contribute to large gains in average consumption and consumption-equivalent utility during the same time period. They also account for about one-half of rising labor force participation over the last five decades. Falling barriers facing white women were

particularly important given their large share of the population.

Declining obstacles to accumulating human capital were more important than declining labor market discrimination in our base specification. For example, declining barriers to human capital attainment explain 18 percent of growth in U.S. GDP per person between 1960 and 2010, while declining labor market discrimination explains 8 percent of growth; however, these numbers are somewhat sensitive to parameter choices. Meanwhile, changing occupation-specific preferences across groups explain little of U.S. growth during this time period.

In the final part of the paper, we use external data to assess the plausibility of our model. First, we compute a U.S. state-level version of our model's summary statistic for labor market discrimination and barriers to human capital attainment for black men. We then correlate this composite friction with survey-based state measures of racial discrimination as calculated by Charles and Guryan (2008) for a similar time period. Our model-based estimates match up closely with Charles and Guryan's survey-based measures. In addition, we compare our model's female labor supply elasticities for different cohorts to empirical estimates by Blau and Kahn (2007). Our model matches well the decline in female labor supply elasticities that they document. The fact that we can replicate measures of discrimination and labor supply elasticities from other authors — even though such moments were not used to discipline our model — provides some validation that our model can capture salient features of the U.S. labor market during this time period.

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 Section 6 concludes.

2. Model

The economy consists of a continuum of workers, each in one of M market occupations or 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 market occupation. We add two things to this standard framework. First, we allow for forces that alter the allocation of talent across market occupations. These forces can take the form of discrimination in the labor market (which we denote τ^w), barriers to human capital accumulation (τ^h), and group-specific preferences for an occupation (z). Second, we assume that individuals are also endowed with an idiosyncratic talent in the home sector (e^{home}).

Individuals invest in their human capital and choose a market occupation in an initial “pre-period”, after which they work for three periods (“young”, “middle”, and “old”). We assume that human capital investments and the choice of market occupation are fixed after the pre-period. Individuals then choose between the home sector and their chosen market occupation each 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, 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 in year t , $1 - s$ is leisure time during the pre-period when human capital investments are made, and z_{ig} denotes the utility benefit of working in occupation i among members of group g .⁷ Note that we assume no discounting of consumption for simplicity. β parameterizes the tradeoff between lifetime consumption and time spent accumulating human capital.⁸

Individuals acquire human capital h in the initial period, and this human capital remains fixed over their lifetime. We 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. Individuals use time s and goods e to produce h according to $h = s^{\phi_i} e^{\eta}$.

⁷We define the cohort index c as the time when the cohort is young so time $t = c$ denotes the year when cohort c works for the first time.

⁸We omit subscripts on other individual-specific variables (such as s and C here) to keep the notation clean. However, z_{ig} does have subscripts to emphasize that it varies across groups and occupations.

In terms of timing, we temporally separate the decision to accumulate human capital from the decision to work in the market versus the home sector. Individuals assume that they will work in their chosen market sector when making their human capital decisions. Once they have accumulated their occupation-specific human capital, individuals then decide whether to allocate their labor supply to working in the market sector (in the chosen occupation) or in the home sector. While this assumption may seem stark, it may not be grossly at odds with the data. For example, in the 1960s and 1970s many women went to college and accumulated human capital only to later decide whether to work in the home sector.

Household consumption can be generated either by allocating time to the market during working years and purchasing consumption or by allocating time to home production. If a worker chooses to work in the market sector, her consumption in each period is her net market income minus a portion of her expenditures on education:

$$C(c, t) = (1 - \tau^w(t)) w(t) \epsilon T h - e(c, t)(1 + \tau^h(c)). \quad (2)$$

Net market income is the product of $1 - \tau^w$ and total efficiency units of labor, which is the product of the price per efficiency unit of skill w , the idiosyncratic talent in the worker's chosen occupation ϵ , the return to experience T , and acquired human capital h . For our baseline case, we assume that the return to experience is the same for all cohorts, occupations, and groups so T is simply a function of age $t - c$.⁹ Individuals borrow $e(c)(1 + \tau^h(c))$ 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^w , works as a "tax" on individual earnings. Given our assumption that the firm owner discriminates against all workers of a given group, τ^w affects all the cohorts of group g equally at a given point in time. Barriers to human capital attainment, τ^h , affect consumption directly by increasing the cost of e in (2) as well as indirectly by lowering human capital h . Because the human capital decision is only made once and fixed thereafter, τ^h for a given occupation varies across cohorts and groups, but is fixed for a given cohort-group over time.

⁹Because labor force participation by age differs across groups and cohorts, the *amount* of experience differs even when the return to experience does not. We allow for this as a robustness check later. Separately, we experimented with allowing the returns to experience to differ by occupation. Doing so did not alter our quantitative conclusions in any meaningful way.

Under these assumptions, 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. In turn, lifetime income depends on the expected values of w_i , τ_{ig}^w , and T in the three periods. For simplicity, we assume that individuals anticipate that the return to experience grows with 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.

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

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

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

where $\bar{T} = 1 + T(1) + T(2)$ is the sum of the experience terms over the three periods. 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 , τ_{ig}^h , and τ_{ig}^w do not affect s because they have the same effect on the return and on the opportunity cost of time. In contrast, these forces change the returns of investment in *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 utility for an individual from group g working in occupation i is $U_{ig}^* = (\tilde{w}_{ig}\epsilon_i)^{\frac{3\beta}{1-\eta}}$ where $\tilde{w}_{ig} \equiv \frac{(\bar{T}/3)w_i s_i^{\phi_i} [(1-s_i z_{ig})^{\frac{1-\eta}{3\beta}}]}{\tau_{ig}}$ and $\tau_{ig} \equiv \frac{(1+\tau_{ig}^h)^\eta}{1-\tau_{ig}^w}$ 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 .

After the individual chooses a market occupation, she then decides whether to work in the chosen market sector or in the home sector. If the individual chooses the home sector, consumption is income from home production minus payments for past investments in human capital:

$$C(c, t) = \Omega_g^{home}(c) \epsilon^{home} h - e(c, t)(1 + \tau_{ig}^h(c)). \quad (3)$$

Income in the home sector is the product of mean home talent $\Omega_g^{home}(c)$, the idiosyncratic talent draw ϵ^{home} in the home sector, and accumulated human capital h in the pre-period. We assume that talent for the home sector $\Omega_g^{home}(c)$ varies across groups g and cohorts c but is fixed for a given group and cohort. For example, technological innovations in the home sector emphasized by Greenwood, Seshadri and Yorukoglu (2005) can be viewed as a decline in $\Omega_g^{home}(c)$ that affect men and women equally. Innovations in contraception emphasized by Goldin and Katz (2002), 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 $\Omega_g^{home}(c)$ of women (but not of men). Notice that acquired human capital h and payments for human capital $e(c, t)(1 + \tau_{ig}^h(c))$ are the same regardless of whether the individual chooses to work in the home or in the market sector.

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 occupation. Talent in the market occupations are 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.¹⁰ Similarly, we assume talent in the home sector ϵ^{home} is drawn

¹⁰We normalize the mean parameter of the Fréchet distribution distribution to one for all groups. A common mean parameter across groups turns out to be isomorphic to the occupation-specific productivity parameters (A_i) we later introduce in the production function. Another consideration is whether all groups have the same mean talent. While this is a reasonable assumption for different racial groups, it may not be when comparing men and woman in certain occupations. For example, physical strength is important in some occupations, and men may be more endowed with this attribute on average. In our robustness checks, we consider this possibility.

from a Fréchet distribution with the same dispersion parameter θ .

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}^* in the first period. 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, as we show in our first proposition.¹¹

Proposition 1 (Occupational Choice): *Let $\tilde{p}_{ig}(c)$ 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*

$$\tilde{p}_{ig}(c) = \frac{\tilde{w}_{ig}(c)^\theta}{\sum_{s=1}^M \tilde{w}_{sg}(c)^\theta} \quad \text{where} \quad \tilde{w}_{igc} \equiv \frac{(\bar{T}/3)w_i(c)s_i(c)^{\phi_i(c)}[(1-s_i(c))z_{ig}(c)]^{\frac{1-\eta}{3\beta}}}{\tau_{ig}(c,c)}. \quad (4)$$

Recall that $\tau_{ig}(c,c) \equiv \frac{(1+\tau_{ig}^h(c))^\eta}{1-\tau_{ig}^w(c)}$ 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 return to skill w_i is high relative to other occupations are ones where more people choose to work. However, heterogeneity in w_i will have the same effect on the occupational choice for every group. In contrast, z , τ^w , and τ^h can explain differences across groups in occupational choice. The fraction of members of group g that choose occupation i is low when a group does not have a preference for that occupation (z_{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).

¹¹Proofs of the propositions are given in <http://www.stanford.edu/~chadj/HHJKAppendix.pdf>.

2.3. Labor Force Participation

After choosing an occupation, each individual decides whether to work in the market or stay at home. Consumption when staying home is given by equation (3). Because the talent in the home sector ϵ^{home} is drawn from an extreme value distribution, indirect utility in the home sector can also be characterized as an extreme value distribution. The probability of choosing the market sector is thus the probability that draws from one extreme value distribution (utility when working in the chosen market sector) is greater than draws from another extreme value distribution (utility when staying at home). We aggregate the individual choices from this decision to obtain the share of individuals that work in the market.

Proposition 2 (Labor Force Participation): *Let $LFP_{ig}(c, t)$ denote the fraction of people in occupation i , cohort c and group g at time t who decide to work rather than stay at home. This fraction is given by*

$$LFP_{ig}(c, t) = \frac{1}{1 + \tilde{p}_{ig}(c) \cdot \left[\frac{\Omega_g^{home}(c)}{\bar{T}(t-c) \cdot (1 - \tau_{ig}^w(t)) \cdot w_i(t)} \right]^\theta}. \quad (5)$$

We do not observe \tilde{p} or LFP in the data, but their product is the fraction of people of a cohort-group actually working in an occupation, p_{ig} , which is the model's counterpart of the occupational shares we do observe in the data:

$$p_{ig}(c, t) = \tilde{p}_{ig}(c) \cdot LFP_{ig}(c, t). \quad (6)$$

Remember $\tilde{p}_{ig}(c)$ is fixed for a given cohort because individuals choose their occupation when they are young. The model thus interprets changes in the occupational share of a given cohort over time as being driven entirely by changes in the labor force participation rate LFP . For a given cohort, LFP increases when the wage per unit of human capital w_i rises, labor market discrimination τ_w declines, and when the return to experience is positive. For example, a decline in labor market discrimination facing female lawyers will increase the fraction of women of a given cohort working as lawyers. A decline in w_i in low-skilled occupations will lower the share of individuals (of a given cohort) working in those occupations.

2.4. Worker Quality

The sorting model then generates the average quality of workers in an occupation for each group. We show this in the following proposition:

Proposition 3 (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*

$$\mathbb{E}[h_{ig}(c, t) \cdot \epsilon_{ig}(c, t)] = \gamma s_i(c)^{\phi_i(t)} \left[\left(\frac{\eta \cdot s_i(c)^{\phi_i(c)} \cdot \bar{T} \cdot w_i(c) \cdot (1 - \tau_{ig}^w(c))}{1 + \tau_{ig}^h(c)} \right)^\eta \left(\frac{1}{p_{ig}(c, t)} \right)^{\frac{1}{\theta}} \right]^{\frac{1}{1-\eta}} \quad (7)$$

where $\gamma \equiv \Gamma(1 - \frac{1}{\theta} \cdot \frac{1}{1-\eta})$ 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 $p_{ig}(c, t)$. This captures the selection effect. For example, the model predicts that if the labor market discriminated against lawyers in 1960, only the most talented female lawyers would have chosen to work in this occupation. And as the barriers faced by female lawyers declined after 1960, less talented female lawyers moved into the legal profession and thus lowered the average quality of female lawyers. Conversely, in 1960, the average quality of white male lawyers would have been lower than it would have been if there were no discrimination against women and blacks.

2.5. 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 4 (Occupational Wages): *Let $\overline{wage}_{ig}(c, t)$ denote the average earnings in occupation i by cohort c of age a of group g . Its value satisfies*

$$\begin{aligned} \overline{wage}_{ig}(c, t) &\equiv (1 - \tau_{ig}^w(t)) \cdot w_i(t) \cdot T(t - c) \cdot \mathbb{E}[h_{ig}(c, t) \cdot \epsilon_{ig}(c, t)] \\ &= \gamma \bar{\eta} \left(\frac{m_g(c, t)}{LFP_{ig}(c, t)} \right)^{\frac{1}{\theta} \cdot \frac{1}{1-\eta}} \cdot [(1 - s_i(c)) z_{ig}(c)]^{-\frac{1}{3\beta}} \cdot \frac{(1 - \tau_{ig}^w(t)) w_i(t)}{(1 - \tau_{ig}^w(c)) w_i(c)} \cdot \frac{T(t - c)}{T} \cdot \frac{s_i(c)^{\phi_i(t)}}{s_i(c)^{\phi_i(c)}} \end{aligned} \quad (8)$$

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

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 $\frac{(1-\tau_{ig}^w(t))w_i(t)}{(1-\tau_{ig}^w(c))w_i(c)} = 1$. Thus, after controlling for the labor force participation rate, 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 disutility 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 (controlling for labor force participation). Average earnings are no higher in occupations where a group faces less discrimination or lower frictions in human capital attainment or a better talent pool or a higher wage per efficiency unit. 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. In the empirical section, we will examine the extent to which changes in the occupational distortions account for the narrowing of wage gaps.

The equation 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 $T(t - c) > T(0)$, the return to schooling increases ($\phi(t) > \phi(c)$), or the share of the cohort-group in the occupation falls. 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 (after controlling for the effect of changes in the share of the group in the occupation). We will use this insight to estimate the change in τ^w in the empirical section.

Putting together the equations for the occupational shares, labor force participation, and wages in each occupation, we get the relative propensity of a group to work in an occupation for the young cohort in each year:

Proposition 5 (Relative Propensities): *Focusing on the young cohort in each year, the fraction of a group working in an occupation — relative to white men — is given by*

$$\frac{p_{ig}(c, c)}{p_{i,wm}(c, c)} = \left(\frac{\tau_{ig}(c, c)}{\tau_{i,wm}(c, c)} \right)^{-\theta} \left(\frac{\overline{wage}_{ig}(c, c)}{\overline{wage}_{i,wm}(c, c)} \right)^{-\theta(1-\eta)} \quad (9)$$

This expression for the share of a group working in an occupation relative to white men when young is a key theoretical relationship we will take to the data. The proposition states that the propensity of a group to work in an occupation when young (relative to white men) depends on two terms: the relative composite occupational frictions and the average occupational wage gap between the groups. From Proposition 2, the wage gap itself is a function of the distortions faced by 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 and barriers to human capital attainment. 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. We will use these insights to recover the composite frictions and the preference parameters.

2.6. Firms and the Rents from Discrimination

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

$$Y = \left[\sum_{i=1}^M \left(A_i \cdot \sum_g H_{ig} \right)^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma}{\sigma-1}} \quad (10)$$

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 (1957), we assume the owner of the firm in the final goods sector discriminates against workers of certain groups. We model the “taste” for discrimina-

tion as lower utility of the owner when she employs workers from groups she dislikes. Her utility is given by

$$U_{owner} = Y - \underbrace{\sum_i \sum_g w_i H_{ig}}_{\text{Base profit}} - \underbrace{\sum_i \sum_g d_{ig} H_{ig}}_{\text{Utility loss via discrimination}} + \underbrace{\sum_i \sum_g \tau_{ig}^w w_i H_{ig}}_{\text{Rents from discrimination}} . \quad (11)$$

The first two terms denote the firm's base profit. The third 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 . The last term reflects the fact that the revenues from discriminating are captured by the firm. Intuitively, when the owner hires a worker from a group she dislikes, the owner's utility loss is compensated by the lower wage for these workers. Because all employers are assumed to have these racist and sexist preferences, in equilibrium $\tau_{ig}^w = d_{ig}/w_i$ and the last two terms sum to zero. Thus the frictions are ultimately pinned down by the discriminatory tastes of (homogeneous) owners.

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

$$U_{school} = \underbrace{\sum_i \sum_g (R_{ig} - 1) \cdot e_{ig}}_{\text{Base profit}} - \underbrace{\sum_i \sum_g d_{ig}^h e_{ig}}_{\text{Utility loss via discrimination}} + \underbrace{\sum_i \sum_g \tau_{ig}^h e_{ig}}_{\text{Rents from discrimination}} \quad (12)$$

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 $\tau_{ig}^h = d_{ig}^h$ so that the last two terms sum to zero in equilibrium.

2.7. Equilibrium

A competitive equilibrium in this economy consists of individual choices $\{C, e, s\}$, an occupational choice in the pre-period, a labor force participation decision in each subsequent period, total efficiency units of labor of each group in each occupation H_{ig} , final market 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 from market work 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 from market work: $i^* = \arg \max_i U(\tau_{ig}^w, \tau_{ig}^h, z_{ig}, w_i, \epsilon_i)$, taking as given $\{\tau_{ig}^w, \tau_{ig}^h, z_{ig}, w_i, \epsilon_i\}$.
3. Each individual chooses between the market occupation and the home sector in each period — taking as given the choice of market occupation, human capital $h, w, \tau^w, \tau^h, T, \Omega^{home}, z$, and idiosyncratic ability in both the market ϵ and in the home sector ϵ^{home} .
4. 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 (11).
5. A representative firm in the education sector maximizes profit net of the utility cost of discrimination given by equation (12).
6. Perfect competition in the final goods and education sectors generates $\tau_{ig}^w = d_{ig}/w_i$ and $\tau_{ig}^h = d_{ig}^h$ so that the last two terms in equations (11) and (12) sum to zero.
7. $w_i(t)$ clears each occupational labor market.
8. Total market output is given by the production function in equation (10).

The equations characterizing the general equilibrium are given in the next result.¹²

¹²Additional results are

1. Total output from labor in the home sector is given by

$$\sum_c \sum_g \sum_i q_g(c) (1 - p_{ig}) \Omega^{home}(t) \Omega_{ig}^{home} \mathbb{E} \left[h_{ig}(c) \epsilon^{home} \mid \text{Person chooses } i \text{ \& home} \right]$$

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 such that*

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

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

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 (7).

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 market output is given by the production function in equation (10) and equals aggregate market wages plus total revenues from τ^w .

2.8. 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, $\tau_{i,ww}$ affects the average wage and occupational choices of women but has no effect on white men. Third, assume $\phi_i = 0$

2. Total output (market plus home) is the sum of aggregate consumption, tax revenues from τ^w and τ^h , and payments for prior investments in human capital given by $\sum_c \sum_g \sum_i q_g(c) e_{ig}(c, t) (1 + \tau_{ig}^h)$.
3. Net borrowing from the rest of the world is total investment in human capital net of payments for prior investments human capital:

$$\text{Net Borrowing}(t) = \underbrace{\sum_g \sum_i q_g(t) e_{ig}(t) (1 + \tau_{ig}^h(t))}_{\text{investment in human capital}} - \underbrace{\sum_c \sum_g \sum_i q_g(c) e_{ig}(c, t) (1 + \tau_{ig}^h(c))}_{\text{payment for investments in human capital}}$$

(no schooling time) and $z_{ig} = 0$. Finally, assume that each cohort lives for one period and all workers choose to work in the market.

The average wages of white men and white women, respectively, are then given by:

$$\overline{wage}_m = \left(\sum_{i=1}^M A_i^\theta \right)^{\frac{1}{\theta} \cdot \frac{1}{1-\eta}}$$

$$\overline{wage}_w = \left(\sum_{i=1}^M \left(\frac{A_i (1 - \tau_i^w)}{(1 + \tau_i^h)^\eta} \right)^\theta \right)^{\frac{1}{\theta} \cdot \frac{1}{1-\eta}}$$

The average male wage is a power mean of the occupational productivity terms and is not affected by the occupational distortions facing women (this is driven by the assumption that occupations are perfect substitutes). The average wage of women is a power mean of the occupational productivity and distortions.

Aggregate consumption and output can then be expressed as a function of the average wage of men and women. Aggregate consumption of workers is proportional to the aggregate wage. Aggregate output is the sum of aggregate wages and rents from τ^w :

$$Y = \underbrace{q_m \cdot \overline{wage}_m + q_w \cdot \overline{wage}_w}_{\text{aggregate wage}} + \underbrace{q_w \cdot \overline{wage}_w \cdot \left(\frac{\bar{\tau}^w}{1 - \bar{\tau}^w} \right)}_{\text{rents from } \tau^w}$$

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.¹³

We illustrate how this setup can be used to gain intuition for the results by focusing on τ^w ; the effects of τ^h can be analyzed in a similar fashion.¹⁴ Assuming $\tau^h = 0$ and that τ^w and A are jointly log-normally distributed, the average female wage is given by:

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

¹³ $\bar{\tau}^w \equiv \sum_{i=1}^M \omega_i \tau_i^w$ where $\omega_i \equiv \frac{p_{iw} \overline{wage}_w}{1 - \tau_i^w} / \sum_{j=1}^M \frac{p_{jw} \overline{wage}_w}{1 - \tau_j^w}$

¹⁴ Assuming $\tau^w = 0$ and that τ^h and A are jointly log-normally distributed, the average female wage is:

$$\ln \overline{wage}_w = \ln \left(\sum_{i=1}^M A_i^\theta \right)^{\frac{1}{\theta} \cdot \frac{1}{1-\eta}} - \frac{\eta}{1-\eta} \cdot \ln(1 + \bar{\tau}^h) - \frac{\eta^2}{2} \cdot \frac{\theta-1}{1-\eta} \cdot \text{Var} \ln(1 + \tau_i^h).$$

The first term says that the average female wage is increasing in the power mean of occupational productivity. The second and third terms state that the average female wage is decreasing in the weighted average of the labor market frictions, and more so the higher is η (the greater the importance of goods for human capital); the reason for not combining these terms is explained below. The fourth term says that the average female wage is decreasing in the dispersion of $1 - \tau^w$.

Both the mean and dispersion of τ^w lower the average female wage and aggregate consumption. However, the effect of the mean of τ^w on aggregate output is different from the effect of its dispersion. The mean has two effects on the average wage. First, $\bar{\tau}^w$ has a direct effect on the average wage, where the elasticity of the average wage to $1 - \bar{\tau}^w$ equals one. This effect is given by the second term in the equation for the female wage. This effect is simply a redistribution from firms: workers get a larger share and firm owners a smaller share, with no change in aggregate output.

The mean of τ^w also has an indirect effect on the average wage by changing the return to investment in human capital. The magnitude of this effect depends on η , and is captured by the third term in the equation. The net effect of a change in $\bar{\tau}^w$ on the aggregate wage is larger than on aggregate output, with the difference depending on η .

The dispersion of τ^w across occupations affects the wage via a different channel. Here, dispersion of τ^w affects the allocation of labor across occupations. A decline in the dispersion of τ^w improves the allocation, which increases the average wage. This effect is captured by the fourth term in the equation for the female wage.

Finally, the equation for the female wage suggests that the effect of misallocation on the wage is increasing in θ . While this is true for a given amount of misallocation, what we need to remember is 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:¹⁵

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

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, the

¹⁵We maintain the assumption that τ^w is the only source of variation.

average female wage is:

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

The elasticity of the wage 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 the wage. 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 shall see later, the effect of changes in occupational shares on wage 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).¹⁶ We make four restrictions to the data when performing our analysis. First, we restrict the sample to white men (wm), white 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.¹⁷

We do not have panel data. Instead, we follow synthetic cohorts over time. We define three age periods within a cohort's lifecycle: the young (those aged 25-34), the middle aged (those aged 35-44) and the old (those aged 45-54). A given synthetic cohort, for example, would be the young in 1960, the middle aged in 1970, or the old in 1980. We have information on 8 cohorts for the time periods we study. For 4 cohorts

¹⁶When using the 2010-2012 ACS data, we pool all three years together and treat them as one cross section. Henceforth, we refer to the pooled 2010-2012 sample as the 2010 sample.

¹⁷Appendix Table C1 reports summary statistics from our sample. For all analysis in the paper, we apply the sample weights in each survey.

(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.

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.¹⁸

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 year, and who reported working more than 30 hours per week. We define the hourly wage as total annual earnings divided by total hours worked in the previous year. We convert all earnings data from the Census to constant dollars.¹⁹

4. Data Inference

As discussed, τ^h , τ^w , and z are isomorphic in terms of their effect on occupational choice. We exploit two features of the model to discriminate between them. First, while τ^h and τ^w do not affect wages in the corresponding occupation (they lower overall wages for the discriminated group), preferences for an occupation do affect a group's average wage in that occupation. This allows us to distinguish the effect of a composite of τ^w and τ^h from the effect of preferences z . Second, we assume that labor market discrimination affects all working cohorts of a discriminated group equally in a given

¹⁸See <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.

¹⁹When 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.

year. In contrast, human capital barriers in a given year only affect the young cohort in that year. This distinction allows us to separately identify τ^w and τ^h .

4.1. Composite Frictions vs. Occupational Preferences

Suppose we normalize τ^h and τ^w of white men to zero.²⁰ Then we can rearrange equation (9) to solve for the composite $\tau_{ig} \equiv \frac{(1+\tau_{ig}^h)^\eta}{1-\tau_{ig}^w}$

$$\tau_{ig}(c, c) = \left(\frac{p_{ig}(c, c)}{p_{i,wm}(c, c)} \right)^{-1/\theta} \left(\frac{\overline{\text{wage}}_{ig}(c, c)}{\overline{\text{wage}}_{i,wm}(c, c)} \right)^{-(1-\eta)} \quad (13)$$

Equation (13) says that, conditional on estimates of θ and η , 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 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.

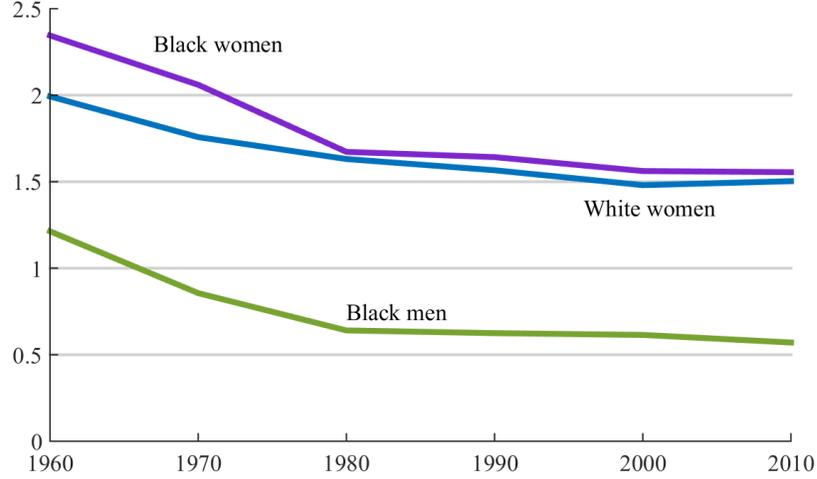
We now present the ingredients needed to measure τ_{ig} from equation (13). Figure 1 plots the standard deviation of $\ln(p_{ig}/p_{i,wm})$ across occupations for the young cohort in each decade.²¹ As shown in Figure 1, the occupations of white men and white women have converged over time. In particular, 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 (13), 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 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 should be similar in all occupations. Specifically, the wage in occupation i for group g relative to

²⁰This will be our base normalization throughout the paper. We will also explore other normalizations, including the assumption that the τ 's for white women are zero and that men have an absolute advantage in brawny occupations relative to women.

²¹We 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.

Figure 1: 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.

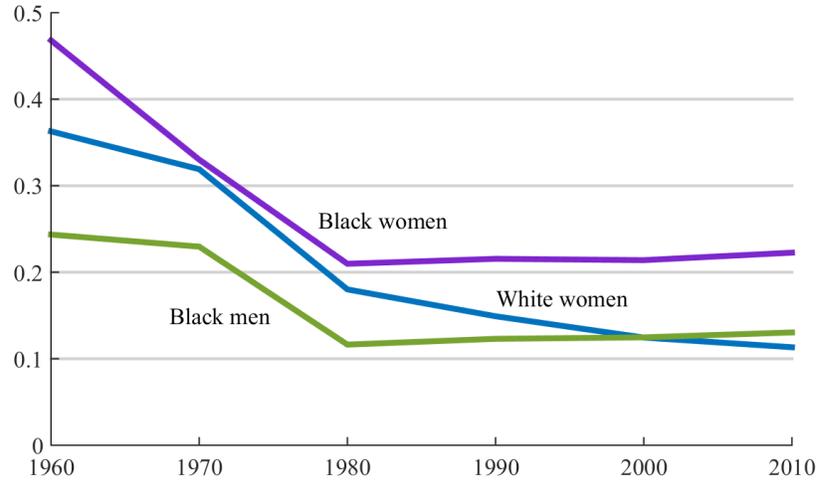
the wages for white men is:

$$\frac{\overline{\text{wage}}_{ig}(c, c)}{\overline{\text{wage}}_{i,wm}(c, c)} = \left(\frac{m_g(c, c)}{m_{wm}(c, c)} \cdot \frac{LFP_{i,wm}(c, c)}{LFP_{ig}(c, c)} \right)^{\frac{1}{\theta \cdot (1-\eta)}} \cdot z_{ig}(c)^{-1/3\beta} \quad (14)$$

where we normalize z to one for white men. The wage gap in an occupation thus isolates the effect of z_{ig} on the occupational choice patterns in Figure 1 (conditional on labor force participation). In addition, this implies that when we condition the occupational gaps on the wage gap in the occupation in (13), we isolate the effect of the labor market and human capital distortions on occupational choice.

Figure 2 plots the weighted standard deviation of the wage gap across market occupations for young (25-34 year old) white women, black men, and black women relative to white men in each year. 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 narrowing of the gap in the occupational distribution over this time period.

Figure 2: 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.

4.2. Estimating θ and η

Wages within an occupation for a given group should 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 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.²²

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

²²Sampling 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.

percent over the years 1995, 2000, 2005, and 2010.²³ Since the labor share in the U.S. in the same four years was 0.641, this implies an η of 0.103.²⁴ With our base estimate of $\theta(1 - \eta) = 1.36$, $\eta = 0.103$ gives us $\theta = 2.12$.

We can also estimate θ from information on the elasticity of labor supply. In our model, the extensive margin elasticity of labor supply with respect to a wage change is equal to $\theta \cdot \frac{1-LFP_g(c,t)}{LFP_g(c,t)}$. With data on labor force participation rates for different cohort-groups, we can back out θ from estimates of the labor supply elasticity. The meta analysis in Chetty et al. (2012) suggests an extensive margin labor supply elasticity of about 0.26. 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 2.31, which is a little higher than the estimate of θ we get from the wage dispersion. We find it comforting that both methods provide estimates of θ that are broadly similar. We will use $\theta = 2.12$ as our base case, but will also provide estimates for θ of 4, which is on the upper end of the estimates for labor supply elasticities.

4.3. Composite Frictions vs. Occupational Preferences: Results

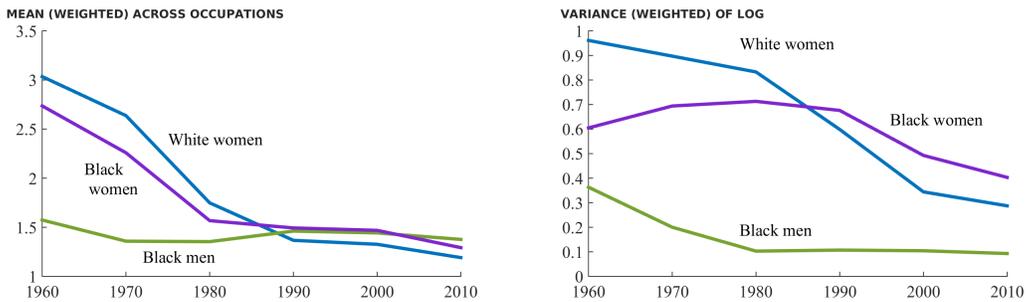
Now that we have estimates of θ and η , we can decompose the dispersion in the occupational propensities into the contribution of τ_{ig} and z_{ig} using equations (13) and (14). Figure 3a summarizes the mean (left panel) and the dispersion (right panel) of τ_{ig} across all 67 occupations for each of the three groups.²⁵ For white women, the mean τ_{ig} fell from about 3 in 1960 to 1.25 in 2010. The decline was steady through 1990 and slowed thereafter. The mean τ_{ig} facing black women declined from around 2.7 to about 1.7 from 1960–1980, then leveled off. Black men, in contrast, experienced no decline in mean τ_{ig} . The dispersion of τ_{ig} fell for all three groups, although the timing differed. For white women the variance of $\ln \tau$ fell continuously from 1 to about one-third, though most sharply between 1980 and 2000. For black women it rose before falling, and for black men it fell from around 0.37 to around 0.1 from 1960 to 1980 and stayed flat thereafter.

²³<http://www.oecd.org/education/eag2013.htm>

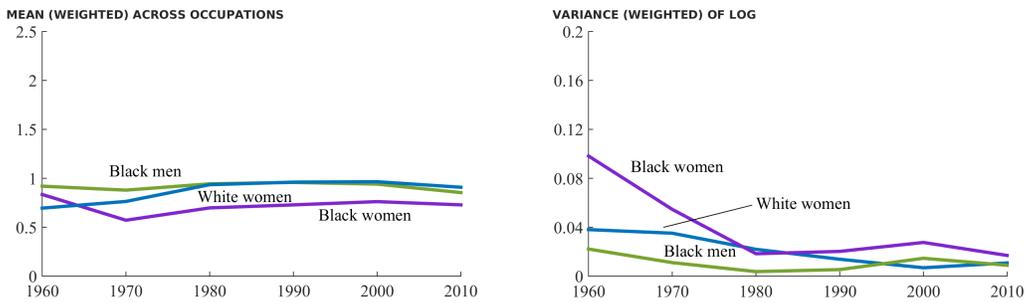
²⁴Labor share data are from <https://research.stlouisfed.org/fred2/series/LABSHPUA156NRUG>. The young's share of earnings is from the U.S. Population Census/ACS.

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

Figure 3: Mean and Variance of τ_{ig} and z_{ig} by Group



(a) The Composite Barriers, τ_{ig}



(b) The Occupational Preferences, z_{ig}

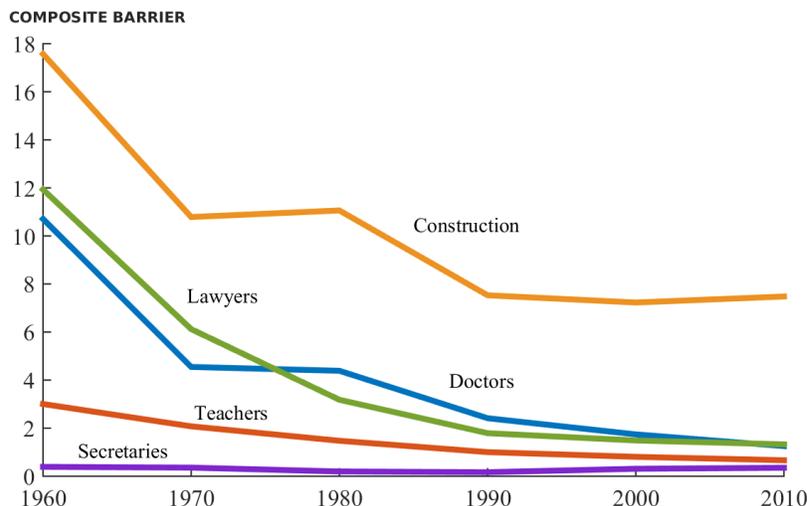
Note: The left panel of each pair shows the average level of the frictions, weighted by total earnings in each occupation in each year. The right panel shows the variance of the log frictions, weighted in the same way. The axis scales are similar but not the same so that it is possible to read the line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

Figure 3b 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 discuss later, 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.

Figure 4 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 4: Estimated Barriers (τ_{ig}) for White Women

Note: Author's calculations based on equation (13) using Census data and imposing $\theta = 2.12$ 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 the changes in τ_{ig}^w versus τ_{ig}^h by exploiting life-cycle variation. The key assumption that allows us to do this is that τ_{ig}^w affects all cohorts in the labor market at the same point in time, whereas τ_{ig}^h only affects the young. This insight can be applied in two different ways.

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

$$\frac{gap_{ig}(c,t)}{gap_{ig}(c,c)} \propto \left(\frac{p_{i,wm}(c,t)/p_{i,wm}(c,c)}{p_{ig}(c,t)/p_{ig}(c,c)} \right)^{\frac{1}{\theta} \cdot \frac{1}{1-\eta}} \cdot \frac{1-\tau_{ig}^w(t)}{1-\tau_{ig}^w(c)} \quad (15)$$

For a given cohort, differences across occupations in the change in the wage gap are a function of the change in occupational shares relative to white men and in τ_{ig}^w over time. Therefore, we can back out the change in τ_{ig}^w from data on the change in the wage gap for a given cohort in a given occupation, after controlling for the effect of changes in the occupational shares. Intuitively, if labor market discrimination diminishes over time, this raises the average wage in occupations where the group previously faced

discrimination.

An alternative to using data on wages is to use information on the change in occupational shares of a given cohort over time. Specifically, equation (5) can be expressed as the ratio of the share of a cohort-group in an occupation at time t relative to time c :

$$\frac{p_{ig}(t)}{p_{ig}(c)} = \frac{1 - p_{ig}(t) \cdot \left[\frac{\Omega_g^{home}(c)}{\Gamma(t-c) \cdot (1 - \tau_{ig}^w(t)) \cdot w_i(t)} \right]^\theta}{1 - p_{ig}(c) \cdot \left[\frac{\Omega_g^{home}(c)}{(1 - \tau_{ig}^w(c)) \cdot w_i(c)} \right]^\theta} \quad (16)$$

Equation (16) indicates that, *ceteris paribus*, a decline in τ_{ig}^w increases the share of individuals of a given cohort and group that chose to work in the market occupation i . In our model, the change in the share of individuals of a given cohort that work in each market sector can be measured by the change in p_{ig} . We can thus use this information to back out the implied change in τ_{ig}^w .

In practice, we pick the value of the change in τ^w that provides the best fit to a weighted average of the change in the wage gap and occupational shares (but fits neither of the two data moments perfectly). Specifically, we assume the return to experience is given by the return to experience for white men, where we make an adjustment for the different amounts of work experience of each group. This adjustment allows women and men to have different observed lifecycle wage profiles in part because older women may have had less labor market experience than older men.²⁶ We then pick the change in τ_{ig}^w to fit a weighted average of the change in the wage gap and occupational shares in each occupation.²⁷ We put more weight (75%) on fitting wages than on fitting occupational choice (25% weight) because sampling error is smaller for wages than for occupations, especially for occupations with low shares.²⁸ For each cohort group, we have $2M$ moments (M wage changes and M changes in occupation shares for each of the M market occupations) to estimate M parameters (M estimates of the change in τ_{ig}^w).

The last step is to back out the change in τ_{ig}^h as the residual of the change in τ_{ig}

²⁶Appendix B provides the full details of our estimation and inference. In that appendix we also discuss in greater depth our procedure to separate the contributions of τ^w and τ^h .

²⁷The estimation takes as given the values of w_i and Ω_g^{home} . We discuss later how we infer their values.

²⁸Below we explore the robustness of our results to specifications where we estimate τ_{ig}^w by putting all the weight on changing wage gaps and then by putting equal weight on changing wage gaps and changing occupational shares.

after controlling for the change in τ_{ig}^w . Specifically, we use $\frac{1+\tau_{ig}^h(t)}{1+\tau_{ig}^h(c)} = \left(\frac{\tau_{ig}(t)}{\tau_{ig}(c)} \cdot \frac{1-\tau_{ig}^w(t)}{1-\tau_{ig}^w(c)} \right)^{\frac{1}{\eta}}$. Again, 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.

This 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. We do this by making an initial guess (a 50/50 split in each occupation) and then computing the implied split for each occupation in subsequent decades. We use the average of this implied split as a new guess for the split in 1960 and iterate until we reach a fixed point. In our baseline case, this gives a (geometric) weight of 0.53 for calculating $1/(1 - \tau^w)$ in 1960.

4.5. Other Parameter Values and Exogenous Variables

The parameters we take as constant over time are η , θ , σ , and β . We discussed our estimates of η and θ earlier. σ governs the elasticity of substitution among our 67 occupations in aggregating up to final output. We have little information on this parameter and choose $\sigma = 3$ as our baseline value. β is the geometric weight on consumption relative to time in an individual's utility function (1). 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 $\beta = \frac{1}{3} \cdot 0.693$ to match the Mincerian return to schooling across occupations, which averages 12.7% across the six decades.²⁹

The remaining variables we need are A_i , ϕ_i , w_i , and Ω^{home} . The first three vary by occupation over time. The last varies by cohort-group over time. The variable ϕ_i governs the occupation-specific return to time invested in human capital. In the model, higher ϕ_i raises time spent in (say) schooling and 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.

²⁹The average wage of group g in occupation i is proportional to $(1 - s_i)^{\frac{-1}{3\beta}}$. If we take a log linear approximation around average schooling \bar{s} , then β is inversely related to the Mincerian return to schooling across occupations (call this return ψ): $3\beta = (\psi(1 - \bar{s}))^{-1}$. We calculate s as average years of schooling divided by a pre-work time endowment of 25 years, and 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. We then set $\beta = \frac{1}{3} \cdot 0.693$, the simple average of the implied β values across years. This method allows the model to approximate the Mincerian return to schooling across occupations.

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.

To estimate w_i and Ω^{home} (for white men), we use the fact that w_i , ϕ_i , and Ω^{home} collectively determine the observed share of young white men in each occupation and the average wage of young white men across all occupations.³⁰ Using the estimates of ϕ_i obtained from the data on wage gaps, we pick w_i and Ω^{home} for white men 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. Conditional on w_i , Ω^{home} of white men pins down the the aggregate labor force participation rate of young white men. With estimates of w_i 's and the production function (10), we then back out the A_i 's.

The last thing we need are estimates of Ω^{home} for non-white men. We pick this parameter such that the aggregate labor force participation rate of the young cohort of each group is exactly equal to the aggregate labor force participation rate observed in the data (conditional on the values of w_i , τ^w , and \tilde{p}_{ig}).

Table 1 summarizes the parameters and other normalizations. The key normalizations are $\tau^h = \tau^w = 0$ and $z = 1$ for white men (in all occupations). This implies that white men face no barriers in either their labor market or human capital decisions and that z_{ig} is the preference of group g for occupation i relative to white men.

Table 2 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. Human capital barriers, labor market discrimination, and occupational preferences vary across occupation-groups. The return to talent in the home sector (Ω^{home}) varies across cohorts for each group.

³⁰This can be seen by combining equations (4), (5), and (6).

Table 1: Baseline Parameter Values and Variable Normalizations

Parameter	Definition	Determination	Value
θ	Fréchet shape	Wage dispersion	2.12
η	Goods elasticity of human capital	Education spending	0.103
σ	EoS across occupations	Arbitrary	3
β	Consumption weight in utility	Wage return to education	$\frac{1}{3} \cdot 0.693$
$z_{i,wm}$	Occupational preferences (white men)	Normalized	1
$\tau_{i,wm}^h$	Human capital barriers (white men)	Normalized	0
$\tau_{i,wm}^w$	Labor market barriers (white men)	Normalized	0

Table 2: Endogenous Variables and Empirical Targets

Parameter	Definition	Empirical Target
$A_i(t)$	Technology by occupation	Occupations of young white men
$\phi_i(c)$	Time elasticity of human capital	Average wages by occupation, white men
$\tau_{i,g}^h(c)$	Human capital barriers	Occupations of young by group
$\tau_{i,g}^w(t)$	Labor market barriers	Life-cycle wage changes by group
$z_{ig}(c)$	Occupational preferences	Occupation wage gaps of young by group
$\Omega_g^{home}(c)$	Home sector talent/taste	Labor force participation

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

4.6. Estimating Human Capital vs. Labor Market Frictions

We now present the split of τ_{ig} into τ_{ig}^w and τ_{ig}^h in Figure 5. Remember this split is estimated from the lifecycle variation in wages and employment of a given cohort-group. The left panel shows the weighted average and the right panel the weighted dispersion τ_{ig}^w and τ_{ig}^h . The decline in average τ_{ig} shown in Figure 5a for white women is driven both by the decline in τ_{ig}^h and τ_{ig}^w . Falling dispersion in τ_{ig} for white women was also due both to declines in τ_{ig}^h and τ_{ig}^w . The decline in the mean occurred throughout, but the decline in the variance occurred after 1980.

Figure 5b presents the same variables for black men. Only the variance of τ_{ig}^w exhibits a notable decline, with the action occurring between 1960 and 1980. Figure 5c presents the path of frictions for black women. Black women, for the most part, exhibit a combination of the results for black men and white women.

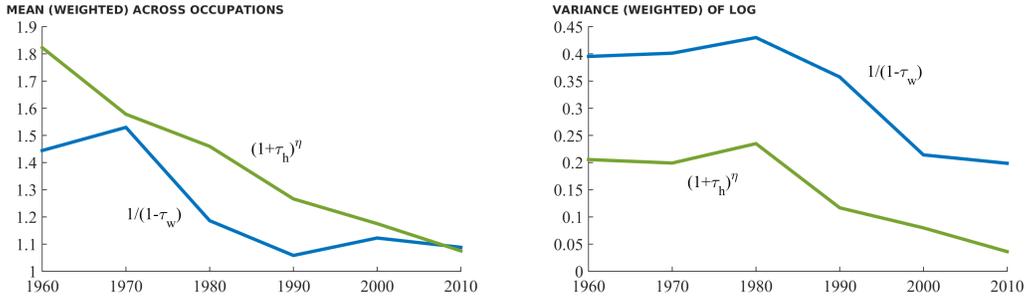
4.7. Model Fit

We remind the reader that we pick the forcing variables and parameters to perfectly match earnings in each occupation and the aggregate labor force participation rates of the *young* in each cohort and group. Furthermore, we pick the changes in τ^w and τ^h to match a weighted average of the changes in average wages and occupational shares of each cohort-group, but we do not match either the wage or quantities of the middle-aged perfectly. Furthermore, we never use data on wages or occupational shares of the old to estimate the model, nor do we use labor force participation rates of middle-aged white men. The model, however, does generate predictions for wages and labor force participation rates for the middle-aged and the old.

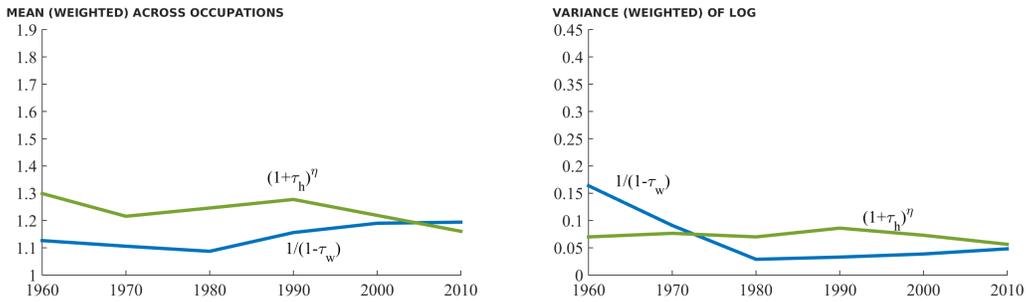
Table 3 illustrates the fit of the model by comparing data and model predictions for aggregate earnings per worker and labor force participation by year. Despite the fact that the model does not target wages for the old and only imperfectly targets wages for the middle-aged, predicted earnings in the model are not very far from earnings in the data. The discrepancy is driven by the middle-aged and the old. For example, in 2010 predicted earnings for the old in the model is 85%, 96% and 85% of the average earnings in the data among white women, black men, and black women, respectively.

The model does less well at matching labor force participation rates of the middle-aged and the old. This gap is largely because the model systematically underpredicts

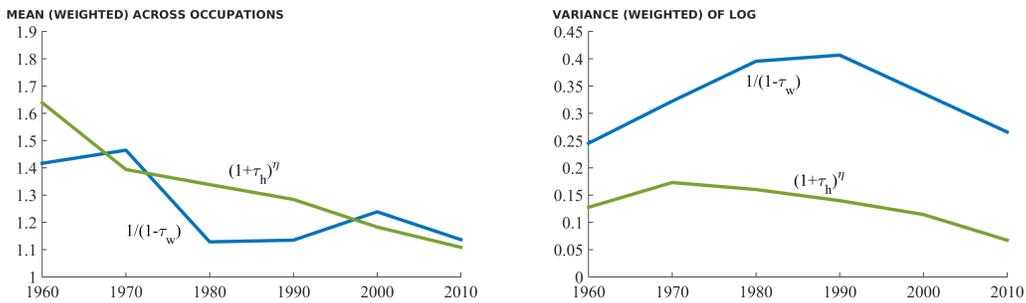
Figure 5: Mean and Variance of τ^h and τ^w



(a) White Women



(b) Black Men



(c) Black Women

Note: The left panel 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 the same across the three groups.

Table 3: Model versus Data: Earnings and Labor Force Participation

Year	Earnings Data	Earnings Model	LFP Data	LFP Model
1960	26,191	26,199	0.599	0.599
1970	35,593	36,142	0.636	0.597
1980	32,925	33,703	0.702	0.643
1990	38,026	39,357	0.764	0.708
2000	47,772	50,195	0.747	0.689
2010	50,981	53,898	0.759	0.723

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

labor force participation for middle and older ages for women who entered the labor market in 1960, 1970, and 1980. For example, middle-age women in 1980 and 2010 have labor force participation rates of 52% and 68%, respectively. Our model predicts labor force participation rates of 36% and 68%. Older women in 1980 and 2010 have labor force participation rates of 51% and 69%. In those years, our model predicts rates of 20% and 57%. This may not be overly surprising given that we abstract from forces such as the pill that allow women to alter their lifecycle profile of family planning.

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?

5.1. Productivity Gains

Real earnings per person in our census sample grew by 1.85 percent per year between 1960 and 2010.³¹ How much of this growth is due to changing τ 's, according to our

³¹These earnings omit employee benefits.

Table 4: 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
Earnings per person	28.7%	29.2%
GDP per person	26.6%	27.3%
Labor force participation	55.1%	41.9%
GDP per worker	19.1%	23.5%

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). We follow the standard approach of chaining – see footnote 32.

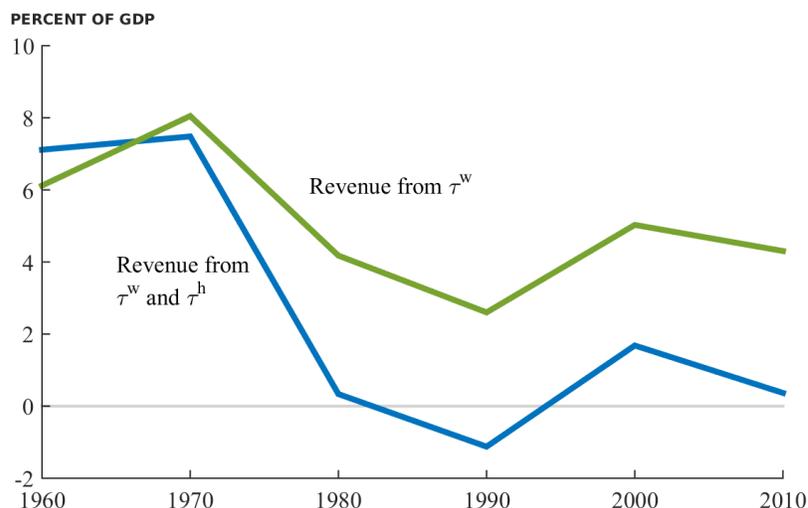
model? We answer this question by changing the τ 's while holding the following fixed: A 's (productivity by occupation), ϕ 's (schooling intensity by occupation), z 's (preferences for each occupation), q 's (group shares of the working population), and Ω^{home} 's (home sector talent/tastes).³²

The results of this calculation are shown in the first column of Table 4. The changes in τ 's account for 28.7% of growth from 1960 to 2010 in earnings per person. The Table also reveals that the changing frictions account for a modestly smaller share of growth in market GDP per person, 26.6%, because some of the earnings growth reflects rising wages relative to marginal products, and in turn declining rents for employers. Figure 6 shows that employer rents were around 7% of GDP in 1960 and 1970, but fell to roughly 4% to 5% of GDP thereafter.

Part of the 26.6% contribution of changing τ 's to growth in GDP per person stems from individuals shifting to the market sector from home production. Aggregate labor force participation rates rose steadily in the data from 59.9% in 1960 to 75.9% in 2010, primarily due to increased female labor supply. Changes in the τ 's explain 55% of this

³² We follow the standard approach of chaining. For example, we compute growth between 1960 and 1970 allowing the τ 's to change but holding the other forcing variables at their 1960 values. Then we compute growth between 1960 and 1970 from changing τ 's holding the other forcing variables at their 1970 values. We take the geometric average of these two estimates of growth from changing τ 's. We do the same for other decadal comparisons (1970 to 1980 and so on) and cumulate the growth to arrive at an estimate for our entire sample from 1960–2010.

Figure 6: Rents as share of GDP in the Model



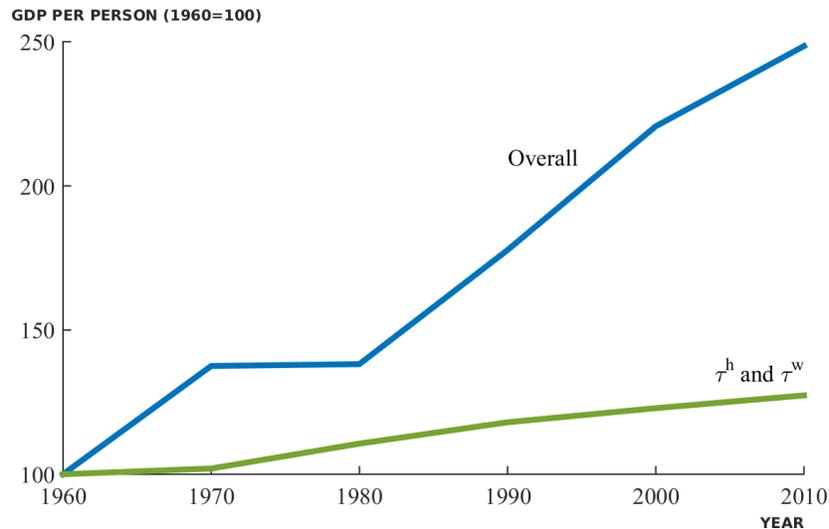
Note: The graph shows the employer rents from discrimination in the labor and human capital markets as a percent of GDP.

increase. Still, the changing τ 's account for 19% of the 1960–2010 growth in GDP *per worker*. These gains come from better allocation of talent and investment in human capital in response to the falling variance of the τ 's, as well as rising human capital investments in response to the falling mean of the τ 's.

Figure 7 shows the time series decomposition of growth. The vast majority of growth in GDP per person is due to increases in A_i and ϕ_i over time, but an important part is attributable to reduced frictions.

If one compares the first and second columns of Table 4, one can see modest effects from changing preferences for each market occupation (z_{ig}). 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. However, as we discuss more in Table 7, productivity/preferences for the home sector (Ω_g^{home}) also explains a large amount of the change in female labor supply. Table 5 also shows that the changing τ 's account for

Figure 7: GDP per person, Data and Model Counterfactual



Note: The graph shows the cumulative growth in GDP per person in the data (overall) and in the model (the contribution of changing τ 's as in Table 4).

a slightly larger share of growth in market earnings per person than GDP per person for the young (41% vs 39%). Again, one reason for this is falling mean barriers and hence falling rents for discriminating owners.

Of perhaps greater interest is the impact of changing τ 's on welfare. To this end, Table 5 focuses on the young to characterize how welfare changes across cohorts. We focus on the young because we can follow them through all six waves of the Census/ACS data. The table shows the effects on GDP per person, earnings per person, market consumption per person, home plus market consumption per person, and consumption-equivalent utility per person. Changing τ 's explain almost 39% of growth in market output for the young, a higher fraction than for all workers because the older cohorts cannot respond in terms of their schooling and human capital decisions in 1970 (and the same for the oldest cohort in 1980).

Table 5 also shows that the changing τ 's account for a smaller share of market consumption growth (32%) than earnings growth (42%) for the young. The share of expenditures on schooling as a share of lifetime earnings is constant in the model, which implies that lifetime consumption is a constant share of lifetime income. Since individuals in the model attempt to smooth consumption over their lifetime income, the

Table 5: Share of Growth due to Changing Frictions (young only)

	Share of growth accounted for by τ^h and τ^w
GDP per person (young)	38.8%
Earnings per person (young)	41.6%
Consumption per person (market, young)	31.8%
Consumption per person (home+market, young)	34.7%
Utility per person (consumption equivalent, young)	56.5%

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).

fact that τ explains a smaller share of the growth of market consumption among the young reflects increases in the return to experience over time. The changing τ 's explain a slightly higher share of total consumption growth (35%) – i.e., including consumption from home production – relative to just market consumption growth (32%). More strikingly, the last row indicates the calibrated model accounts for over 56% of growth in consumption-equivalent utility. Much of growth in GDP 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, the efficiency gains from a better allocation of talent entail no such cost. The falling employer rents also translate into utility gains for workers relative to employers.

Table 6 splits the growth into contributions from falling barriers to human capital accumulation (τ^h 's) vs. falling labor market barriers (τ^w 's). Of the 26.6% contribution to growth in GDP per person, falling human capital barriers account for about 18 percentage points and falling labor market barriers around 8 percentage points. The larger role played by human capital barriers also holds for growth across cohorts (GDP, consumption, and consumption-equivalent welfare). As we shall see below, however, the relative contribution of human capital and labor market barriers is sensitive to various parameter choices. While the magnitudes differ across our robustness exercises, the modal finding is that declining human capital barriers contributed more to U.S. growth

Table 6: Share of Growth due to Changing Labor- vs. Product-Market Frictions

	Share of growth accounted for by		
	τ^h and τ^w	τ^h only	τ^w only
GDP per person	26.6%	18.3%	8.4%
GDP per person (young)	38.8%	26.9%	12.3%
Earnings per person (young)	41.6%	21.0%	20.5%
Consumption per person (market, young)	31.8%	16.3%	15.5%
Consumption per person (home+market, young)	34.7%	21.8%	13.0%
Utility per person (consumption equivalent, young)	56.5%	37.4%	15.7%

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).

than declining labor market barriers.

Table 7 shows how the changing τ 's affect wage gaps and earnings across groups. For women, the changing τ 's more than explain the shrinking gender gap in wages. The model says that, in the absence of changing τ 's, the rising labor force participation rate of women (due to changing Ω^{home} 's) 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).³³

Table 7 also says that the changing τ 's account for almost none of the wage growth for white men, not surprisingly. But they do account for 67% of earnings growth for

³³As we explain in the appendix, the productivity term in the home sector Ω^{home} has a component that is fixed over time for a given group but varies across occupations as well as a component that is the same across occupations but varies across time. We then estimate the component of Ω^{home} that is fixed over time to force the labor force participation rate in each occupation among the young to be roughly equal to the aggregate labor force participation rate. When we do this, we get that women who entered higher skilled occupations had a higher taste/productivity for the home sector. As market productivity growth was higher for higher skilled occupations, all individuals migrated towards those occupations. As women moved into those occupations, however, their higher home sector taste/productivity resulted in some of those women staying at home. This put downward pressure on female labor supply, all else equal. This force explains why Table 7 suggests that changes in occupational productivity and returns to occupation-specific schooling (A and ϕ in the model) *reduced* female labor supply (compare columns 3 and 4). The model then matches observed labor supply patterns by forcing the component of Ω^{home} that varies over time to fall sharply for women in all occupations for new cohorts relative to older ones.

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	$\tau^h, \tau^w, z, \Omega_g^{home}$	
Wage gap, WW	158.0%	171.5%	88.3%	104.9%
Wage gap, BM	85.4%	93.4%	81.0%	104.0%
Wage gap, BW	110.2%	124.6%	81.8%	98.0%
Earnings, WM	0.2%	0.0%	1.0%	104.6%
Earnings, WW	67.6%	68.2%	86.8%	100.2%
Earnings, BM	20.7%	20.4%	22.5%	96.0%
Earnings, BW	48.0%	49.5%	61.5%	96.9%
LF Participation	55.1%	41.9%	185.4%	79.4%

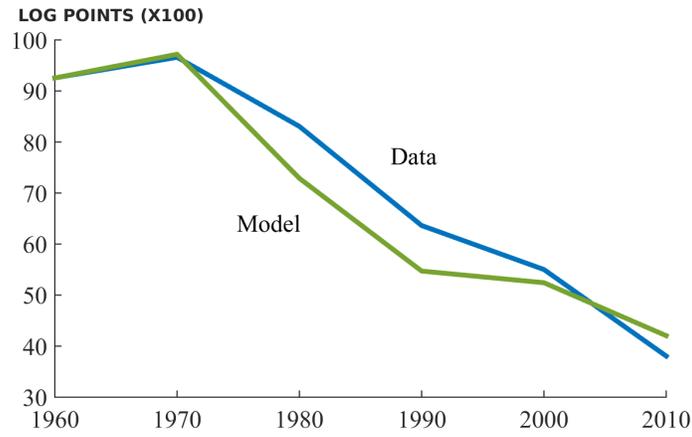
Note: Entries in the table show the share of growth in the model attributable to changing frictions and other variables. The frictions are τ^h (human capital) and τ^w (labor market), and the other variables are z (occupational preference) and Ω_g^{home} (labor force participation). The last column reports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

white women, 21% for black men, and 48% 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). This highlights the value in our structural approach, which does not assume all changes in relative earnings stem from changing τ 's, compared to a back-of-the-envelope approach which does.

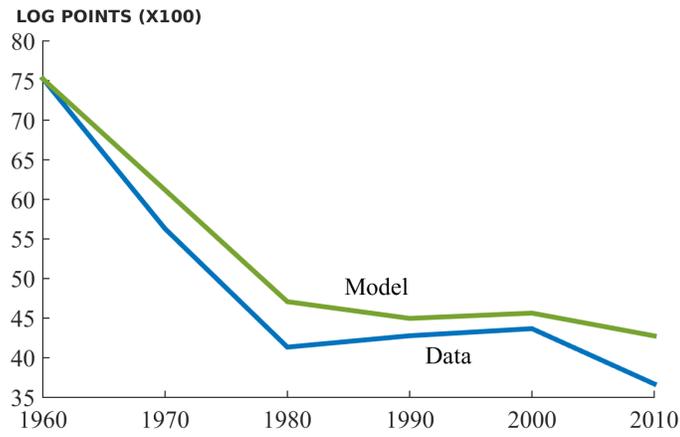
Figure 8 shows that the model also does a good job of matching the *timing* of falling wage gaps. The wage gap for white women fell continuously after 1970 in both model and data. The gap closed for black men vs. white men from 1960 to 1980 but remained stagnant afterward in both model and data.

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 output per worker during the 1960-2010 period. This is primarily due to the fact that white women are a much larger share of the population.

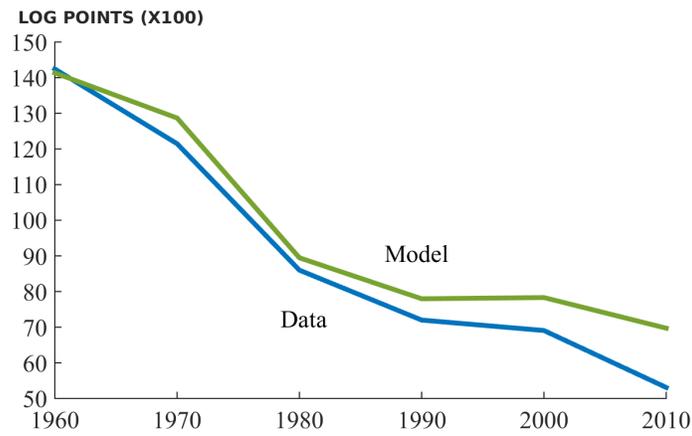
Figure 8: Wage Gaps: Model vs. Data



(a) White Women



(b) Black Men



(c) Black Women

Note: Model lines show the impact of changing τ^h , τ^w , z , and Ω^{home} .

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	26.6%	18.3%	8.4%
White women	22.3%	15.2%	7.3%
Black men	1.4%	1.1%	0.3%
Black women	2.9%	2.1%	0.8%
1960–1980			
All groups	31.2%	12.6%	19.0%
White women	24.9%	9.2%	16.1%
Black men	2.8%	1.5%	1.3%
Black women	3.5%	1.8%	1.8%
1980–2010			
All groups	24.0%	21.5%	2.6%
White women	20.8%	18.5%	2.5%
Black men	0.6%	0.8%	-0.2%
Black women	2.6%	2.2%	0.3%

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

The impact of changing τ 's for black men was most pronounced from 1960-1980. There was little change in the τ 's of black men after 1980, so changes in the τ 's of black men explain essentially none of the post-1980 growth in GDP per person. Table 8 also shows that, while pre-labor market barriers to human capital accumulation contribute about two-thirds of the gains overall from 1960–2010, a majority of the gains from 1960–1980 came from declining labor market barriers. In contrast, whereas declining labor market discrimination was important prior to 1980, changes in labor market discrimination after 1980 had essentially no effect on earnings per capita. After 1980, growth occurred, in part, due to the continued decline in barriers to human capital attainment for women.

Finally, we can ask how much additional growth could be had from reducing the frictions all the way to zero? If the remaining frictions in 2010 were removed entirely, we calculate that GDP per person would be higher by an additional 21.4%. GDP for the young would rise 16.4%. These remaining gains result from the fact that even in 2010 occupational barriers exist across groups.

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

Our baseline estimate in Table 4 suggests that τ^w and τ^h account for 26.6% of the gains in 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 2.8, the elasticity of GDP to $1 - \bar{\tau}^w$ 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 3a shows that the mean of the composite τ of women fell from 3 in 1960 to 1.25 in 2010. This decline in $\bar{\tau}$ can thus account for a .05 log point increase in GDP per person.³⁴ Figure 3a also shows that $\text{Var} \ln \bar{\tau}$ fell from 1 to 0.3 from 1960 to 2010. In the log-normal approximation to the model, the semi-elasticity of GDP to $\text{Var} \ln \tau$ is $q_w \cdot \frac{1}{2} \cdot \frac{\theta-1}{1-\eta} \approx 0.3$.³⁵ A 0.7 log point decrease in the variance of $\ln \tau$ thus can explain a 0.21 log point increase in GDP per person. This back-of-the-envelope calculation thus suggests that changes in τ account for a 0.26 log point increase in GDP per person, of which 80% comes from a more efficient allocation of talent across occupations (as measured by the decline in

³⁴ $\frac{1}{2} \cdot \frac{1}{9} \cdot \ln(3/1.25) \approx .05$

³⁵ $\frac{1}{2} \cdot \frac{1}{2} \cdot \frac{2.12-1}{1-.1} \approx 0.3$

τ dispersion.) The overall increase of GDP per person in our setup was 0.91 log points, so the changing τ explains $0.26/0.91 \approx 28.6\%$ of growth. This is quite similar to the estimated gains we get in the full model.

A second way to answer the question is to compare our 26.6% 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 37.3% of growth in earnings per person.³⁶ Why is our model-based estimate lower than this back-of-the-envelope calculation? The back-of-the-envelope calculation assumes no general equilibrium effect of falling frictions on the earnings growth of white men. We show below that this GE effect is quite small in our base-line calibration. Much more important, the back-of-the-envelope calculation assumes 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 been important.

5.3. Robustness

Table 9 explores the robustness of our results to alternative counterfactuals. 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 23.3% or 21.5% of growth in GDP per person, vs. 26.6% 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 wage gap for white women would have widened due to the changing Ω^{home} 's in the absence of the changing τ 's. The main take away from this analysis is that productivity gains from changing labor market discrimination and barriers to human capital accumulation cannot be gleaned from the wage gaps alone.

³⁶For this calculation, we held fixed earnings per person relative to white men at 1960 levels, and found only 62.7% as much growth in earnings per person as seen in the data.

Table 9: Robustness to Alternative Counterfactuals

	GDP per person growth accounted for by τ^h and τ^w
Benchmark	26.6%
Wage gaps halved	23.3%
Zero wage gaps	21.5%
No frictions in “brawny” occupations	22.9%
No frictions in 2010	26.4%

Note: Entries in the table show the share of growth in the model attributable to changing frictions τ^h (human capital) and τ^w (labor market). In the fourth line, we assume that there are no frictions for white women in occupations where physical strength is important. Instead, we allow the mean value of the Fréchet distribution of market skills in these occupations to change over time to match the occupational allocation for white women. For blacks in this case, we do allow for frictions, but also assume the mean value of the Fréchet distribution of market skills are the same as that for white men. “No frictions in 2010” assumes that there are no frictions in 2010, so that differences in the mean value of the Fréchet distribution of market skills 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.

The next row in Table 9 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 because of technological progress. To see the potential importance of this story, we go to the extreme of assuming no frictions faced by white women in any of the occupations where physical strength is arguably important. These occupations include construction, firefighters, police officers, and most of manufacturing.³⁷ We estimate mean skill for young white women to fully explain their allocation to these occupations in 1960, 1970, . . . , 2010. As shown in Table 9, the fraction of growth explained by changing frictions falls slightly from 26.5% to 22.9%. Thus most of the productivity 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.

For the last row in Table 9, we assume all group differences among the young in 2010 reflect talent rather than distortions. I.e., we set 2010 τ 's to zero and assume absolute advantage fully accounts for group differences in occupational choice among the young in 2010. We keep talent in previous years at the 2010 values for each group, but back out distortions in earlier years. Eliminating the τ 's in the earlier years still generates 26.4% of observed growth in GDP per person. Thus our productivity gains are not an artefact of assuming the allocation of talent was far from optimal in 2010.

Table 10 explores robustness of our productivity gains to different parameter values. For each set of parameter values considered, we recalculate the τ , z , Ω^{home} , A , and ϕ values so that the model continues to fit the occupation shares, wage gaps, etc. The first row of Table 10 replicates the gains under baseline parameter values for comparison. The next row considers a higher value of the Fréchet shape parameter θ , which is inversely related to the dispersion of comparative advantage across occupations. Recall that our baseline θ was estimated from wage dispersion within occupation-groups. This might overstate the degree of comparative advantage because some wage variation is due to *absolute* advantage. We thus entertain a higher value ($\theta = 4$) than in our baseline ($\theta = 2.12$). As shown in Table 10, the productivity gains from changing τ 's rise only slightly to 27% with a higher θ .

³⁷Rendall (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.

Table 10: Robustness to Parameter Values

	GDP per person growth accounted for by		
	τ^h and τ^w	τ^h alone	τ^w alone
Benchmark	26.6%	18.3%	8.4%
$\theta = 4$	27.0%	15.2%	12.5%
$\eta = 0.05$	24.7%	6.4%	18.4%
$\eta = 0.20$	28.2%	25.0%	3.1%
$\sigma = 1.05$	27.0%	18.7%	8.4%
$\sigma = 10$	26.3%	18.1%	8.5%

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 $\theta = 2.12$, $\eta = 0.103$, and $\sigma = 3$.

Table 10 also varies η , the elasticity of human capital with respect to goods invested in human capital. As one might expect, the higher the value of η , the bigger the gains from reducing human capital and labor market frictions: the gains rise from 24.7% with $\eta = 0.05$ to 26.6% with our baseline $\eta = 0.103$ to 28.2% with $\eta = 0.20$. Recall that η should be the share of output invested in human capital, and its baseline value was set to match education spending as a share of labor earnings. Intuitively, the gains from falling human capital barriers are greater the higher is η .

The last rows of Table 10 show the (in)sensitivity of the results to the elasticity of substitution (σ) between occupations in production. As shown, the gains to changing τ 's are very similar with $\sigma = 1.05$ (close to Cobb-Douglas aggregation) and with $\sigma = 10$ as with the baseline of $\sigma = 3$. Although not shown in Table 10, 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. For Table 11 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.

Table 11: Changing Only the Dispersion of Ability

Value of θ	GDP per person growth accounted for by τ^h and τ^w
1.9	13.0%
2.12 (baseline)	26.6%
3	67.1%
4	99.8%
5	128.4%

Note: Here we hold fixed all other parameters and forcing variables at their baseline values. Unlike the other robustness checks, we do not re-calibrate.

Consistent with the intuition provided in Section 2.8, 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.

Table 12 considers additional robustness checks related to the details of our empirical estimation. This highlights one of the key places where our results are not robust: the relative importance of τ^h versus τ^w depends on the weight we place on cohorts when middle-aged in matching quantities (the p_{ig}) versus the wage growth they experience between youth and middle age. Our benchmark case puts a weight of 1/4 on the p_{ig} and therefore a weight of 3/4 on wage growth. The first part of Table 12 shows how the results change when we vary this weight. The overall share of growth in GDP per person accounted for by the τ^h and τ^w together is very robust to this choice. However, the relative importance of τ^w falls considerably if we put all our weight on matching p_{ig} , reaching as low as 2.0%. Conversely, if we put all our weight on wage growth, the importance of τ^w rises to 19.1% while the contribution from τ^h falls to only 8.1%. Putting equal weights on the two moments turns out to be very similar to putting all the weight on p_{ig} , while our benchmark weight of 1/4 is more intermediate. We therefore chose this value as most representative of our results.

The other robustness checks in Table 12 highlight that other parameter choices we

Table 12: More Robustness

	— GDP growth accounted for by —		
	τ^h and τ^w	τ^h only	τ^w only
Benchmark	26.6%	18.3%	8.4%
Weight on $p_{ig} = 1$	23.8%	21.9%	2.0%
Weight on $p_{ig} = 1/2$	25.2%	22.7%	2.4%
Weight on $p_{ig} = 0$	27.2%	8.1%	19.1%
50/50 split of $\hat{\tau}_{i,g}$ in 1960	26.6%	19.1%	7.7%
50/50 split of $\hat{\tau}_{i,g}$ in all years	28.8%	19.8%	9.3%
LFP minimum factor = 1/3	26.5%	18.6%	8.2%
LFP minimum factor = 2/3	26.4%	17.9%	8.8%
No constraint on τ^h	26.4%	21.8%	4.6%

Note: See notes to Table 4. The baseline parameter values include an LFP minimum factor of 0.5, and a Weight = 0.25 on fitting p_{ig} . Weight=1 chooses the split of τ into τ^h vs. τ^w to match p_{ig} for the middle-age group, while Weight=0 puts all the weight on matching wage growth for the middle-age growth; the benchmark case has Weight=0.25. Alternatives to this include splitting τ evenly into τ^h and τ^w in 1960, or in all years.

make for technical reasons make little difference to our overall findings.³⁸

6. Further Model Implications

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

6.1. Trends in Female Labor Supply Elasticities

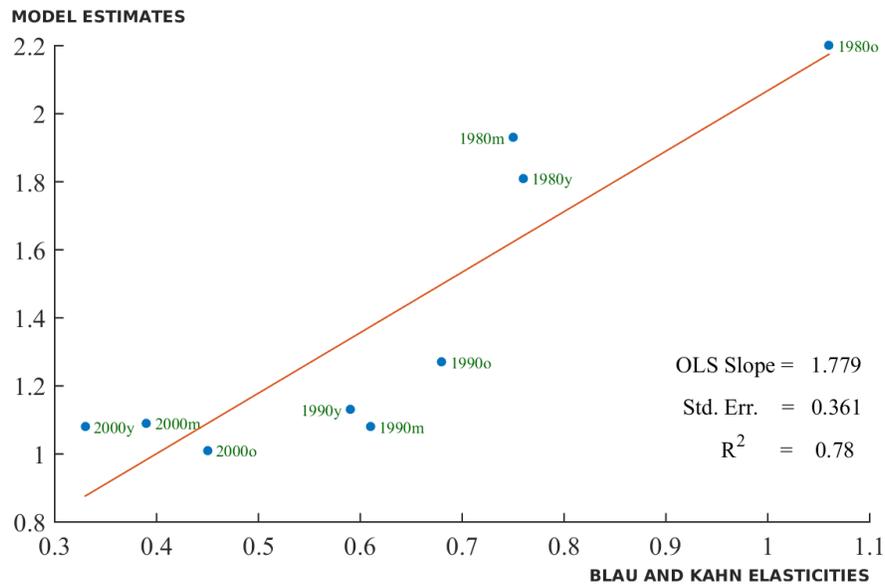
Using data from the Current Population Survey, Blau and Kahn (2007) estimate that there was a dramatic decline in female labor supply elasticity during the 1980–2000 period. Helpful for comparing with the predictions of our model, they report female labor supply elasticities for 25-34 year olds, 35-44 year olds, and 45-54 year olds in 1980, 1990, and 2000. They find that labor supply elasticities of women (with respect to changes in their own wage) fell by roughly 50 percent during the 1980–2010 period.

Figure 9 plots the estimates of female labor supply elasticities from Blau and Kahn for various age-year groups against the model’s implied labor supply elasticities for white women for the same age-year groups. In total, there are nine points, corresponding to the young, middle age, and old groups in the years 1980, 1990, and 2000.

As seen in Figure 9, the implied labor supply elasticities for white women from our model track closely the estimates by Blau and Khan. Consider married women aged 45-54 in 1980 and 2000. Blau and Kahn estimate labor supply elasticities in these two years of 1.1 and 0.5, respectively. Our model implies labor supply elasticities for these two years of 2.2 and 1.0, respectively. Our estimates are systematically higher than the Blau and Khan estimates — potentially due to the fact that we are focusing on all white women as opposed to married women. But, the key finding we want to highlight is that the proportional decline in labor supply elasticities from our model is nearly identical to the estimated decline in labor supply elasticities documented by Blau and Khan.

³⁸The 50/50 split cases involve alternative ways of decomposing the composite $\hat{\tau}_{ig}$ into τ^w and τ^h , as opposed to our benchmark approach of finding the decomposition in 1960 that endogenously matches the average decomposition in subsequent decades. The “LFP minimum factor” is the maximum extent to which we allow labor force participation in any occupation to fall below the aggregate for a cohort/group. Finally, our benchmark case constrains the values of τ^h to be no smaller than -0.8. Allowing even larger subsidies (e.g. for white women secretaries) can lead to large rents in some cases.

Figure 9: Labor Supply Elasticities for White Women



Note: Figure compares the estimates of female labor supply elasticities for young, middle age, and older women in 1980, 1990, and 2000 as reported in Blau and Kahn (2007) (x-axis) against the model implied labor supply elasticities for the same age-year groups (y-axis). For the Blau and Khan estimates, we use the results from their Model 1 reported in their Table 2.

For example, in our model, women aged 45-55 experienced a decline in labor supply elasticities between 1980 and 2000 of 55 percent. Blau and Khan also report a decline in labor supply elasticities for this group of 55 percent during the same time period. The results are not just concentrated among the older group. Between 1980 and 2000, our model matches nearly identically the percentage decline in labor supply elasticities for the young and middle age groups.

Nothing in our model is calibrated to match 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 female labor supply as well as the trend for different cohorts shows that our model is consistent with empirical moments outside the ones we used to calibrate the model.

6.2. Cross State Measures of Discrimination

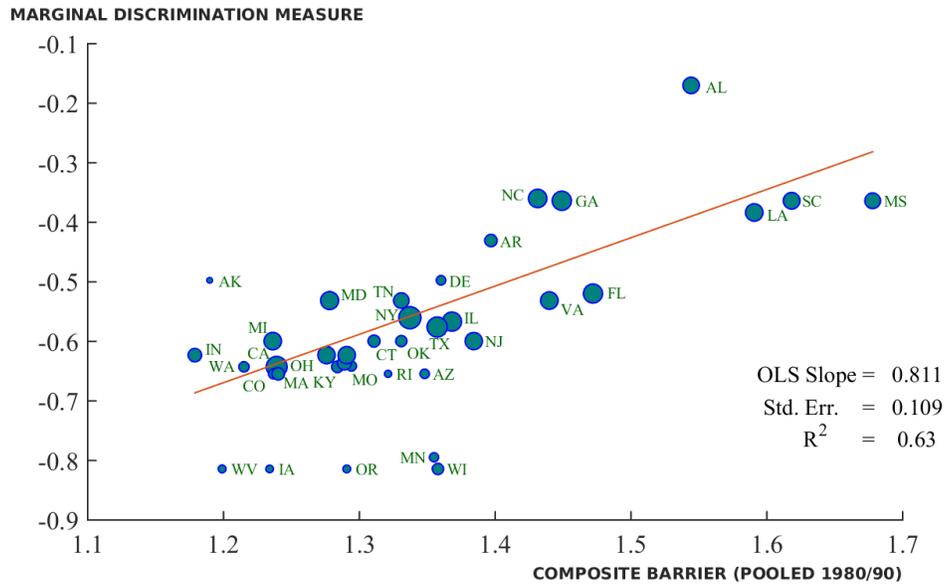
There are very few micro-based measures of discrimination to which we can compare our estimated τ 's. One such exception is the recent work by Charles and Guryan (2008). Charles and Guryan (CG) used data from the General Social Survey (GSS) to construct a measure of the taste for discrimination against 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 values of the CG discrimination measure imply more discrimination. They compute their measure for 44 states.

Figure 10 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 0.81 with a standard error of 0.11. Places we identify as having a high τ_{bm} are the same places Charles and Guryan find as being highly discrim-

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 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 10: 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.

inatory based on survey data from the GSS. The findings in Figure 10 provide additional external validity that our procedure is measuring salient features of the U.S. economy over the last five decades.

7. 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 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 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. In addition, 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

- 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.

- 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.
- Fernandez, Raquel, "Cultural Change as Learning: The Evolution of Female Labor Force Participation over a Century," *American Economic Review*, forthcoming 2012.

- 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.
- Fernandez, Raquel, Alessandra Fogli, and Claudia Olivetti, “Preference Formation and the Rise of Women’s Labor Force Participation: Evidence from WWII,” NBER Working Papers 10589, National Bureau of Economic Research, Inc June 2004.
- Goldin, Claudia, *Understanding the Gender Gap: An Economic History of American Women*, New York, NY: Oxford University Press, 1990.
- 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.
- 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

August 26, 2016

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, $U(\tau_i, w_i, \epsilon_i)$, is proportional to $(\tilde{w}_{ig}\epsilon_i)^{\frac{3\beta}{1-\eta}}$, where $\tilde{w}_{ig} \equiv \frac{(\bar{T}/3)w_i s_i^{\phi_i} [(1-s_i)z_{ig}]^{\frac{1-\eta}{3\beta}}}{\tau_{ig}}$ and $\frac{\bar{T}}{3} \equiv \frac{1}{3}(1 + T(1) + T(3))$ is the average of the lifetime 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

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

Proof of Proposition 2. Labor Force Participation

In each period, conditional on choosing market sector i , the individual compares her consumption from market work (equation (2)) with consumption from work in the home sector (equation (3)). She works in the market when

$$\Omega_g^{home} \epsilon^{home} < (1 - \tau_{ig}^w) w_i T \epsilon^*$$

where ϵ^* is the ability of people in their chosen occupation. The probability of working in the market is thus

$$\begin{aligned}
 LFP &= \Pr[\Omega_g^{home} \epsilon^{home} < (1 - \tau_{ig}^w) w_i T \epsilon^*] \\
 &= \Pr[\epsilon^{home} < \mu \epsilon^*]
 \end{aligned}$$

where $\mu \equiv \frac{(1 - \tau_{ig}^w) w_i T}{\Omega_g^{home}}$.

To calculate this probability, we need to know the distribution of ϵ^* . We use the extreme value magic of the Fréchet distribution. 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:

$$\begin{aligned} \Pr [y^* < z] &= \Pr [y_i < z] \quad \forall i \\ &= \Pr [\epsilon_i < z/\tilde{w}_i] \quad \forall i \\ &= F \left(\frac{z}{\tilde{w}_1}, \dots, \frac{z}{\tilde{w}_M} \right) \\ &= \exp \left[- \sum_s \tilde{w}_s^\theta z^{-\theta} \right] \\ &= \exp \{-mz^{-\theta}\}. \end{aligned} \tag{19}$$

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:

$$G(x) \equiv \Pr [\epsilon^* < x] \equiv \exp \left[-m^* x^{-\theta} \right] \tag{20}$$

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

With the distribution of ϵ^* in hand, we can now evaluate the probability of working in the market sector:

$$\begin{aligned} LFP &= \Pr [\epsilon^{home} < \mu\epsilon] \\ &= \int F_H(\mu\epsilon^*) dG(\epsilon^*) d\epsilon^* \end{aligned}$$

where $F_H(\cdot)$ is the CDF of ability in the home sector. Since $F_H(\mu\epsilon^*) = \exp \left[-\mu^{-\theta} \epsilon^{*-\theta} \right]$ and $dG(x) = \theta(1/\tilde{p}^*)x^{-\theta-1}G(x)$, after substituting and solving the integral, the probability of working in the market is given by:

$$LFP = \frac{1}{1 + \tilde{p}^* \mu^{-\theta}}$$

which simplifies to equation (5) after we replace μ with $\frac{(1-\tau_{ig}^w)w_i T}{\Omega_g^{home}}$.

Proof of Proposition 3. Average Quality of Workers

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

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

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

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

Therefore,

$$H_{ig} = q_g p_{ig} (s_i^{\phi_i})^{\frac{1}{1-\eta}} \left(\frac{\eta w_i (1 - \tau_i^w) \bar{T}}{1 + \tau_i^h} \right)^{\frac{\eta}{1-\eta}} \cdot \mathbb{E} \left[\epsilon_i^{\frac{1}{1-\eta}} \mid \text{Person chooses } i \& \text{ work} \right]. \quad (21)$$

We need to know the distribution of ϵ^* conditional on choosing to work raised to some power. The distribution of ϵ^* conditional on choosing to work is:

$$\begin{aligned} H(x) &\equiv \Pr[\epsilon^* < x \mid \epsilon^{home} < \mu \epsilon^*] \\ &= \frac{\Pr[\frac{1}{\mu} \epsilon^{home} < \epsilon^* < x]}{LFP} \end{aligned}$$

We already have an expression for LFP so all we need is to solve for the numerator in this expression.

$$\begin{aligned} \Pr\left[\frac{1}{\mu} \epsilon^{home} < \epsilon^* < x\right] &= \int_0^x \int_0^{\mu \epsilon^*} dF(\epsilon^{home}) dG(\epsilon^*) \\ &= \int_0^x F(\mu \epsilon^*) dG(\epsilon^*) \\ &= \frac{1}{1 + \frac{\mu}{\tilde{p}^*}} \cdot \exp\left[-\left(\frac{1}{\tilde{p}^*} + \mu^{-\theta}\right) x^{-\theta}\right] \end{aligned}$$

After combining this result with equation (5), the distribution of ϵ^* conditional on

working in the market is:

$$\begin{aligned} H(x) &\equiv \Pr [\epsilon^* < x | \epsilon^{home} < \mu \epsilon^*] \\ &= \exp \left[- \left(\frac{1}{\tilde{p}^*} + \mu^{-\theta} \right) x^{-\theta} \right] \\ &= \exp \left[- \left(\frac{1}{p_{ig}} \right) x^{-\theta} \right] \end{aligned}$$

where $p_{ig} \equiv LFP_{ig} \cdot \tilde{p}_{ig}^*$ is the share of a group observed working in the occupation. The distribution of talent conditional on working is Fréchet with dispersion parameter θ and mean parameter $\frac{1}{p_{ig}}$.

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

$$\begin{aligned} \mathbb{E}[\epsilon_i^\lambda] &= \int_0^\infty \epsilon^\lambda dG(\epsilon) \\ &= \int_0^\infty \theta \left(\frac{1}{p_{ig}} \right) \epsilon^{-\theta-1+\lambda} e^{-\left(\frac{1}{p_{ig}}\right)\epsilon^{-\theta}} d\epsilon \end{aligned} \quad (22)$$

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

$$\begin{aligned} \mathbb{E}[\epsilon_i^\lambda] &= \left(\frac{1}{p_{ig}} \right)^{\lambda/\theta} \int_0^\infty x^{-\frac{\lambda}{\theta}} e^{-x} dx \\ &= \left(\frac{1}{p_{ig}} \right)^{\lambda/\theta} \Gamma \left(1 - \frac{\lambda}{\theta} \right). \end{aligned} \quad (23)$$

Applying this result to our model, we have

$$\mathbb{E} \left[\epsilon_i^{\frac{1}{1-\eta}} \mid \text{Person chooses } i \text{ \& works} \right] = \left(\frac{1}{p_{ig}} \right)^{\frac{1}{\theta} \cdot \frac{1}{1-\eta}} \Gamma \left(1 - \frac{1}{\theta} \cdot \frac{1}{1-\eta} \right). \quad (24)$$

Substituting this expression into (21) and rearranging leads to equation (7)

Proof of Proposition 4. Occupational Wage Gaps

The proof of this proposition is straightforward given the results of Proposition 3. Note that $\bar{\eta} \equiv \eta^{\eta/(1-\eta)}$.

Proof of Proposition 5. Relative Propensities

The fraction of a group working in an occupation relative to white men is given by:

$$\frac{p_{ig}}{p_{i,wm}} \equiv \left(\frac{LFP_{ig}}{LFP_{i,wm}} \right) \cdot \left(\frac{\tilde{p}_{ig}}{\tilde{p}_{i,wm}} \right)$$

After substituting the results from propositions 1, 2, and 4 into the expression for relative propensities above and rearranging, we get equation (6).

B Identification and Estimation

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

B1. Key Equations

To estimate the model, we add two additional features to the model. First, in our base case, we assume the return to experience is the same for all occupations, groups, and cohorts, but we allow for differences across groups in the average amount of market experience over the life-cycle. For example, the market experience of the average middle-aged female may be lower than that of a comparable male (of the same cohort). Furthermore, in our robustness checks, we allow the returns to experience to also differ by occupations. We thus index T (and the sum of the experience terms \bar{T}) by group g and occupation i in the equations that follow.

Second, we also introduce a home talent term that is fixed over time but varies across occupations and group. Specifically, we assume the home talent term is given by $\Omega_{ig}^{home} \Omega_g^{home}(c) \epsilon^{home}$. Ω_{ig}^{home} varies across occupation-groups but is fixed across cohorts. $\Omega_g^{home}(c)$ varies across cohorts for a given group but is the same across occupations. After we account for these two changes, the key equations we use for estimation are listed below.

$$\begin{aligned} \text{Occupational Choice : } \tilde{p}_i &= \frac{\tilde{w}_i^\theta}{\sum_s \tilde{w}_s^\theta} \\ \text{where } \tilde{w}_{ig} &\equiv \frac{w_i \bar{T}_{ig} s_i^{\phi_i} [(1-s_i)z_{ig}]^{\frac{1-\eta}{3\beta}}}{\tau_{ig}} \\ \text{and } \tau_{ig} &\equiv \frac{(1+\tau_{ig}^h)^\eta}{1-\tau_{ig}^w} \end{aligned}$$

$$\text{Labor Force Participation Rate : } LFP_{ig}(c, t) = \frac{1}{1 + \tilde{p}_{ig}(c) \cdot \left[\frac{\Omega^{\text{home}(c)} \cdot \Omega_{ig}^{\text{home}}}{(1-\tau_{ig}^w(t))w_i(t) \cdot T_{ig}(t-c)} \right]^\theta}$$

$$\text{Occupational Share : } p_{ig}(c, t) = \tilde{p}_{ig}(c) \cdot LFP_{ig}(c, t)$$

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

$$\begin{aligned} \text{Average Wage : } \overline{\text{wage}}_{ig}(c, t) &\equiv (1-\tau_{ig}^w(t))w_i(t) \cdot T_{ig} \cdot \mathbb{E}[h_{ig}(c)\epsilon_{ig}(c, t)] \\ &= \gamma \bar{\eta} \left[\frac{m_g(c, t)}{LFP_{ig}(c, t)} \right]^{\frac{1}{\theta} \cdot \frac{1}{1-\eta}} [(1-s_i(c))z_{ig}(c)]^{-\frac{1}{3\beta}} \cdot \frac{(1-\tau_{ig}^w(t))w_i(t)T_{ig}s_i(c)^{\phi(t)}}{(1-\tau_{ig}^w(c))w_i(c)\bar{T}_{ig}s_i(c)^{\phi(c)}} \\ \text{where } m_g(c, t) &= \sum_{i=1}^M \tilde{w}_{ig}(c, t)^\theta \end{aligned}$$

$$\text{Relative Propensity : } \frac{p_{ig}(c, c)}{p_{i,wm}(c, c)} = \left(\frac{T_{ig}}{T_{i,wm}} \right)^{\theta(1-\eta)} \left(\frac{\tau_{ig}(c, c)}{\tau_{i,wm}(c, c)} \right)^{-\theta} \left(\frac{\overline{\text{wage}}_{ig}(c, c)}{\overline{\text{wage}}_{i,wm}(c, c)} \right)^{-\theta(1-\eta)} \left(\frac{\bar{T}_{ig}}{\bar{T}_{i,wm}} \right)^{\theta\eta}$$

B2. Estimate Wages & Schooling from data of young white men

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

First, we assume a value of ϕ in the farming occupation. Using this normalization, we back out s in farming 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:

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

where $i = \text{farm}$. We assume that LFP in the farm sector is equal to the aggregate labor force participation rate of young white men and the average wage in the farm sector $\overline{wage}_{\text{farm},wm}$ is data. Furthermore, we need to make an initial guess about the returns to experience terms T and \bar{T} . (We describe later how we estimate the return to experience terms).

Third, now that we have an estimate of m_{wm} , we recover w_{wm}^{home} and $\Omega_{i,wm}^{\text{home}}$ such that $LFP_{i,wm}$ in every occupation is exactly equal to the observed aggregate labor force participation rate for young white men. Specifically, after substituting the expression for \tilde{p}_{ig} into the expression for LFP_{ig} and simplifying, we get:

$$LFP_{i,wm} = \left[1 + \left[\Omega_{wm}^{\text{home}} \Omega_{i,wm}^{\text{home}} \bar{T}_{i,wm} s_i^{\phi_i} (1-s_i)^{\frac{1-\eta}{3\beta}} \right]^{\theta} \cdot \frac{1}{m_{wm}} \right]^{-1}$$

where we assume $LFP_{i,wm}$ is the same in all occupations and equal to the aggregate labor force participation rate observed in the data for young white men. Remember that $\tau_{i,wm} = 1$ (we normalize τ^w and τ^h to zero for white men) and $z_{i,wm} = 1$ for white men which is why these terms do not show up in the above equation. We have M observations and $M + 1$ parameters (M parameters for $\Omega_{i,wm}^{\text{home}}$ and one parameter for Ω^{home} so we normalize the average of $\Omega_{i,wm}^{\text{home}}$ to one).

Fourth, 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.

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

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

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

Sixth, 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:

$$\frac{\overline{\text{wage}}_{i,wm}(c, t)}{\overline{\text{wage}}_{i,wm}(c, c)} = \left[\frac{LFP_{wm}(c, c)}{LFP_{wm}(t, c)} \right]^{\frac{1}{\theta} \cdot \frac{1}{1-\eta}} \cdot \frac{w_i(t) T_{ig}(c, t) s_i^{\phi(t)}}{w_i(c) T_{ig}(c, c) s_i^{\phi(c)}}$$

We normalize $T_{i,wm}(c, c) = 1$ and estimate $T_{i,wm}(c, t)$ from the change in the average wage in an occupation, after controlling from the effect of changes in the labor force participation rate (which is simply the change in the observed occupational shares) and the change in w_i . We then adjust T_{ig} for the other groups using the observed changes in labor force participation rates for a given cohort. Specifically, we assume T_{ig} is the product of the experience terms of white men of the same cohort and $\frac{p_{ig}(c, c)}{p_{ig}(c, t)}$ in the cases when this ratio falls with age. When the average occupational share in the group declines, we assume T_{ig} is equal to that of white men. In our base case, we assume T_{ig} is the same across all occupations and cohorts so simply take the average across all occupations and cohorts. In our robustness checks, we allow T_{ig} to vary across occupations.

Finally, we assume $\Omega_{i,wm}^{home}$ for all cohorts is the minimum of the $\Omega_{i,wm}^{home}$ observed across all cohorts. We then iterate over steps (2)-(5) with this fixed value of Ω^{home} to obtain final values of w_i and ϕ_i .

B3. Estimating $\hat{\tau}$

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 $\tau_{i,wm}^w = \tau_{i,wm}^h = 0$. This normalization implies that we can express the relative propensity expression as:

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

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

B4. 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 programs `estimatetauz`, `eval_young`, and `eval_middle`. 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 .

Our first step is to estimate Ω_{ig}^{home} and Ω_g^{home} for the other groups (other than white men). We start with an initial guess of α and allow Ω_{ig}^{home} to vary across cohorts such that the predicted labor force participation rates in every occupation is exactly equal to the aggregate labor force participation rate of the young cohort of the group (we assume $\Omega_g^{home} = 1$ in this first step). Specifically, the labor force participation equation can be expressed as a function of Ω_{ig}^{home} :

$$\Omega_{ig}^{home} = \left[\frac{1 - LFP_{ig}}{p_{ig}} \right]^{\frac{1}{\theta}} \cdot \frac{(1 - \tau_{ig}^w) w_i T_{ig}}{\Omega_g^{home}}$$

To be clear, we assume LFP_{ig} is equal for all occupations and given by the aggregate

LFP observed in the data (for the young of group g of the specific cohort).

Second, we set Ω_{ig}^{home} equal to its minimum across all cohorts. Without further changes, the aggregate labor force participation rate in the model will not be equal to that in the data. To correct this, we do two things. For the young in 1960, we pick the value of α (the split of τ into τ^w and τ^h) such that LFP in every occupation is equal to the aggregate LFP observed in the data. This pins down the split of τ into τ^w and τ^h in 1960.

For the young in later years (after 1960), we use the same equation for Ω_{ig}^{home} above and allow Ω_g^{home} to vary across cohorts such that predicted aggregate labor force participation rate in the model is equal to the aggregate labor force participation rate observed in the data. To be clear, we do not impose the condition that LFP in every occupation is equal to the aggregate LFP in the data. Instead, LFP varies across occupations but we choose Ω^{home} such that the weighted average of LFP in the model is equal to the aggregate LFP observed in the data for young cohorts after 1960. We also impose the constraint that the lowest possible value of LFP in an occupation is 50 percent of the aggregate LFP of each cohort-group observed in the data.

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

$$m_g(c) = LFP_{farm,g}(c, c) \left[\frac{\overline{wage}_{farm,g}(c, c)(1 - s_{farm}(c))^{\frac{1}{3\beta}}}{\gamma\bar{\eta}} \cdot \frac{\bar{T}_{farm,g}}{T_{farm,g}} \right]^{\theta(1-\eta)}$$

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

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

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. We have two ways to do this. First, we can use the labor force participation equation for the young and the middle-aged of

the same cohort. After some manipulation, this can be expressed as:

$$\tau_{ig}^w(c, t) = \left[\frac{LFP_{ig}(c, t)}{1 - LFP_{ig}(c, t)} \cdot \tilde{p}_{ig}(c) \right]^{\frac{1}{\theta}} \cdot \left[\frac{\Omega^{home}(c) \cdot \Omega_{ig}^{home}}{w_i(t) \cdot T_{ig}(t - c)} \right] - 1$$

A second option is to use the change in the average wage for a given occupation and cohort. Specifically, after some manipulation, τ^w for the years after 1960 can be expressed as a function of the change in the average wage of a given cohort and τ^w in the previous year:

$$1 - \tau_{ig}^w(t) = \frac{\overline{\text{wage}}_{ig}(c, t)}{\overline{\text{wage}}_{ig}(c, c)} \cdot \left[\frac{LFP_{ig}(c, t)}{LFP_{ig}(c, c)} \right]^{\frac{1}{\theta} \cdot \frac{1}{1-\eta}} \cdot \frac{w_i(c) s_i^{\phi(c)}}{w_i(t) T_{ig}(t - c) s_i^{\phi(t)}} \cdot (1 - \tau_{ig}^w(c))$$

We estimate the last two equations sequentially, starting with the young in 1960, then the young in 1970, and so forth, subject to the constraints that α (the split of τ between τ^h and τ^w) is between 0 and 1 and the minimum value of τ^h is -0.8. We put more weight (75%) on fitting the change in the wage than on fitting the change in occupational shares (25%) between the young and middle-aged of each group.

C Appendix Tables

Table C1: 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.

Table C2: 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.