

The Allocation of Talent and U.S. Economic Growth

Chang-Tai Hsieh

Chicago Booth and NBER

Erik Hurst

Chicago Booth and NBER

Charles I. Jones

Stanford GSB and NBER

Peter J. Klenow *

Stanford University and NBER

February 22, 2013 – Version 3.0

Abstract

In 1960, 94 percent of doctors and lawyers were white men. By 2008, 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. This paper measures the macroeconomic consequences of the remarkable convergence in the occupational distribution between 1960 and 2008 through the prism of a Roy model. We find that 15 to 20 percent of growth in aggregate output per worker over this period may be explained by the improved allocation of talent.

*We are grateful to Raquel Fernandez, Kevin Murphy, and seminar participants at Boston College, Boston University, Brown, Chicago, Columbia, CREI, Duke, Harvard, LSE, Michigan, MIT, an NBER EFG meeting, Penn, Princeton, Stanford, Toronto, UBC, UCLA, USC and Yale for helpful comments, and to Jihee Kim, Huiyu Li and Gabriel Ulyssea for 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 2008, the fraction was just 62 percent. Similar changes occurred throughout the economy during the last fifty years, particularly among highly-skilled occupations.¹

A large literature attempts to explain why white men differ in their occupational distribution relative to women and blacks and why those differences have been changing over time. 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 misallocation of talent could potentially have important effects on aggregate productivity.²

This paper measures the aggregate productivity effects of the potential misallocation of talent among women and blacks from 1960 to 2008. To do this, we examine the differences in labor market outcomes between race and gender groups through the prism of a Roy (1951) model of occupational choice. While it is not our goal to causally identify any specific friction that explains differences in occupational sorting, our model is broad enough to encompass many of the common explanations highlighted in the literature. Specifically, we begin by assuming that every person is born with a range of talents across all possible occupations and chooses the occupation with the highest return. Differences in the occupational distribution between men and women can be driven by differences in the distribution of talent between groups. Rendall (2010), for example, shows that brawn-intensive occupations (such as construction) in the U.S. are dominated by men, and that changes in the returns to brawn- vs. brain-intensive occupations can explain changes in the occupational distribution of women vs. men since 1960; see also Pitt, Rosenzweig and Hassan (2012). Related,

¹These statistics are based on the 1960 Census and the 2006-2008 American Community Surveys. We discuss the sample in more detail below. A large literature provides more extensive documentation of these facts. See Blau (1998), Blau, Brummund and Liu (2012), Goldin (1990), Goldin and Katz (2012), and Smith and Welch (1989) for assessments of this evidence.

²Hsieh and Klenow (2009) suggest that differences in barriers that affect the allocation of inputs can explain some of the differences in productivity across countries.

Goldin and Katz (2002) and Bertrand, Goldin and Katz (2010) provide evidence that innovations in contraception and increased labor market flexibility for women had important effects on the occupational choices of women.

However, we also allow other forces to play a role in explaining differences in occupational sorting. Consider the world that Supreme Court Justice Sandra Day O'Connor faced when she graduated from Stanford Law School in 1952. Despite being ranked third in her class, the only private sector job she could get immediately after graduating was as a legal secretary (Biskupic, 2006). Such barriers might explain why white men dominated the legal profession at that time. And the fact that private law firms are now more open to hiring talented female lawyers might explain why the share of women in the legal profession has increased dramatically over the last fifty years. Similarly, the Civil Rights movement of the 1960s was surely important for the changing occupational distribution of blacks.³

To capture these forces, we make several changes to the canonical Roy framework. First, we allow for the possibility that each group faces different occupational frictions in the labor market. We model these frictions as a group/occupation-specific "tax" on earnings that drives a wedge between a group's marginal product in an occupation and their take home pay. One interpretation of these "taxes" is that they represent preference-based discrimination as in Becker (1957). For example, one reason why private law firms would not hire Justice O'Connor is that the law firms' partners (or their customers) viewed the otherwise identical legal services provided by female lawyers as somehow less valuable.⁴

Second, we allow for frictions in the acquisition of human capital. We model these frictions as a group-specific tax for each occupation on the inputs into human capital production. These human capital frictions could represent the fact that some groups were restricted from elite higher education institutions, that black public schools are underfunded relative to white public schools, that there are differences in prenatal or early life health investments across groups, or that social forces steered certain groups towards certain occupations.⁵

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

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

⁵Here is an incomplete list of the enormous literature on these forces. Karabel (2005) documents how

Finally, we allow for changes in the returns to skill across occupations. If these changes are common to all groups, then they will not affect the occupations of men and women differently. However, some technological changes may be group-specific. The innovations related to contraception mentioned earlier are a prime example. Skill biased technical change may have changed the occupational distribution of women relative to men if women are relatively more endowed with skill.

In our augmented Roy model, all three forces — barriers to occupational choice, relative ability across occupations, and relative returns to occupational skills — will affect the occupational distribution. To make progress analytically, we follow McFadden (1974) and Eaton and Kortum (2002) and assume that talent obeys an extreme value distribution. This assumption gives us two key results. First, we get a closed-form expression relating the share of a group in an occupation to the frictions faced by the group in the occupation relative to frictions faced by the group in *all* occupations. Talent is misallocated only when the frictions differ across occupations. Second, the misallocation of talent due to the dispersion of occupational frictions lowers the average wage of the group in all occupations. Larger barriers in an occupation lead to a selection effect in which only the most talented choose that occupation, and these two forces exactly net out in the model. As a result, frictions specific to an occupation show up in quantities rather than in wages in that occupation. These two results allow us to back out the occupation-specific frictions for each group from data on occupational shares and average wages. Using data from the decadal U.S. Censuses and the American Community Surveys, we find that the dispersion of the occupational frictions faced by women and blacks, and thus the misallocation of talent, decreased substantially over the last fifty years.

We close the model by introducing the demand for skills in each occupation, where the price of skills is determined by the supply and demand for skills in each occupation. In our general equilibrium Roy model, the observed changes in occupational choice,

Harvard, Princeton, and Yale systematically discriminated against blacks, women, and Jews in admissions until the late 1960s. Card and Krueger (1992) documents that public schools for blacks in the U.S. South in the 1950s were underfunded relative to schools for white children. See Chay, Guryan and Mazumder (2009) for evidence on the importance of improved access to health care for blacks. See Fernandez (2012) and Fernandez, Fogli and Olivetti (2004) on the role of social forces in women's occupational choice. Goldin and Katz (2002), Bailey (2006), Bertrand, Goldin and Katz (2010), and Bailey, Hershbein and Milleri (2012) document that innovations related to contraception had important consequences for female labor market outcomes and educational attainment. Fernandez and Wong (2011) stress rising divorce rates as a force behind women's rising labor force participation and educational attainment.

relative wages, and aggregate productivity can be explained by a combination of our three forces (occupational barriers, talent distribution, occupation-specific technical change). We then use this framework to isolate the effect of improved allocation of talent among women and blacks from 1960 to 2008.

We freely admit this calculation makes no allowance for model misspecification and thus should be viewed as only an illustration of the potential magnitude of the effect of declining occupational barriers. In addition, without further information, we cannot disentangle the effect of labor market discrimination from that of frictions in the human capital market. However, while only illustrative, this calculation captures forces that a simple back-of-the-envelope calculation (based on changing wage gaps alone) does not. First, our calculation isolates the potential effect of changes in occupational barriers, whereas the observed changes in the wage gap are also driven by innovations in sector-specific productivities, skill requirements, and demographics. Second, our calculation hones in on the effect of talent misallocation which, in the model, is only driven by the dispersion in occupational barriers. Third, our calculation incorporates the impact of changes in occupational barriers on white men, whereas the wage gap calculation does not.

Our results imply that changes in occupational barriers facing blacks and women can potentially explain 15 to 20 percent of aggregate growth in output per worker between 1960 and 2008. These estimates are 40 percent larger than what we find with a simple back-of-the-envelope calculation. Furthermore, essentially all of the gain is driven by the movement of women into high-skilled occupations. We infer that changes in occupational barriers may have raised real wages by roughly 40% for white women, 60% for black women, and 45% for black men, but *lowered* them by about 5% for white men. Again, that wages of white men may have suffered is one important reason why a simple back-of-the-envelope calculation could underestimate the importance of declining labor market frictions in explaining aggregate productivity growth.

The paper proceeds as follows. Section 2 lays out the basic model of occupational choice. In Section 3, we provide micro evidence for one of the key predictions of our sorting model. We then use our framework to measure the frictions in occupational choice between blacks and women versus white men in Section 4. In Section 5, we explore the macroeconomic consequences of the changes in occupational frictions

across groups. We offer some closing thoughts in the final section.⁶

2. Occupational Sorting and Aggregate Productivity

We start with the occupational choice decision. The economy consists of a continuum of people working in N possible occupations, one of which is the home sector. Each person possesses heterogeneous abilities — some people are good economists while others are good nurses; some people may even be better than others at many jobs. The basic allocation to be determined in this economy is how to match workers with occupations.

2.1. People

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

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

where s represents time spent on human capital accumulation, and β parameterizes the tradeoff between consumption and time spent accumulating human capital.

Each person works one unit of time in an occupation indexed by i . Another unit of time — think “when young” — is divided between leisure and schooling. A person’s human capital is produced by combining time s and goods e . The production function for human capital in occupation i is

$$h(e, s) = \bar{h}_{ig} s^{\phi_i} e^\eta. \tag{2}$$

Note that we will omit subscripts on individual-specific variables (such as s and e in this case) to keep the notation clean. The elasticity of human capital with respect to

⁶Several recent papers are worth noting for related contributions. Ellison and Swanson (2010) show that high-achieving girls in elite mathematical competitions are more geographically concentrated than high-achieving boys, suggesting many girls with the ability to reach these elite levels are not doing so. Cavalcanti and Tavares (2007) use differences in wage gaps across countries in a macro model to measure the overall costs of gender discrimination and find that it is large. Albanesi and Olivetti (2009b) study the gender earnings gap in a model of home production, while Dupuy (2012) studies the evolution of gender gaps in world record performances in sport. Beaudry and Lewis (2012), looking across cities, suggest that much of the change in the gender wage gap can be explained by a change in the relative price of skills.

time, ϕ_i , however, does have a subscript to emphasize that this elasticity varies across occupations. As we discuss below, this plays a key role in generating wage differences across occupations.

The parameter \bar{h}_{ig} allows for potentially different efficiency in human capital accumulation across groups. We have two interpretations in mind for \bar{h}_{ig} . One is that family background (e.g., nutrition and health care) could differ across groups, thereby affecting the human capital payoff to investments in schooling quantity and quality. Cunha, Heckman and Schennach (2010) provide evidence of such complementarity between early and later human capital investments. A second interpretation is that women's childbearing may disrupt human capital investment, a force that could change over time with fertility and technology. For example, Goldin and Katz (2002) and Bertrand, Goldin and Katz (2010) provide evidence that innovations in contraception had important effects on the timing of childbearing and, in turn, the education and occupational choices of women.

In addition, we allow for two other frictions. The first affects human capital choices. We model this friction as a "tax" τ_{ig}^h that is applied to the goods e invested in human capital and that varies across both occupations and groups. We think of this tax as representing forces that affect the cost of acquiring human capital for different groups in different occupations. For example, τ_{ig}^h might represent discrimination against blacks or women in admission to universities, or differential allocation of resources to public schools attended by black vs. white children, or parental liquidity constraints that affect children's education. Additionally, it can represent the differential investments made toward building up math and science skills in boys relative to girls.

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

Consumption equals labor income less expenditures on education:

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

Note that pre-distortion labor income is the product of the wage received per efficiency unit of labor, the idiosyncratic talent draw ϵ in the worker's chosen occupation, and the

individual's acquired human capital h .

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

$$U(\tau^w, \tau^h, \bar{h}, w, \epsilon) = \max_{c, e, s} (1-s)c^\beta \quad \text{s.t.} \quad c = (1 - \tau_{ig}^w)w\epsilon h(e, s) - e(1 + \tau_{ig}^h). \quad (4)$$

This yields the following for the amount of time and goods spent on human capital:

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

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

Time spent accumulating human capital is increasing in ϕ_i . Individuals in high ϕ_i occupations acquire more schooling and have higher wages to compensate them for the 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 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, we get the following expression for indirect utility in occupation i :

$$U(\tau_{ig}, w_i, \epsilon_i) = \left(\frac{w_i s_i^{\phi_i} (1 - s_i)^{\frac{1-\eta}{\beta}} \epsilon_i \cdot \eta^\eta (1 - \eta)^{1-\eta}}{\tau_{ig}} \right)^{\frac{\beta}{1-\eta}} \quad (5)$$

Here, we define τ_{ig} as a “gross” tax rate that summarizes the frictions:

$$\tau_{ig} \equiv \frac{(1 + \tau_{ig}^h)^\eta}{1 - \tau_{ig}^w} \cdot \frac{1}{\bar{h}_{ig}}. \quad (6)$$

It turns out that \bar{h}_{ig} and $(1 + \tau_{ig}^h)^\eta$ are observationally equivalent in our setup (in particular if e is unobserved). This occurs despite the fact that they may have very dif-

ferent implications for efficiency. Given that it does not matter at all for our quantitative results, we set $\bar{h}_{ig} = 1$ recognizing that τ_{ig}^h could have this alternative interpretation.

2.2. Occupational Skills

Turning to talent, we borrow from McFadden (1974) and Eaton and Kortum (2002). Each person gets a skill draw ϵ_i in each occupation. These skills are allowed to be correlated across occupations: some people may be better than others at many jobs. Talents are drawn from a multivariate Fréchet distribution:

$$F_g(\epsilon_1, \dots, \epsilon_N) = \exp \left\{ - \left[\sum_{i=1}^N (\tilde{T}_{ig} \epsilon_i^{-\tilde{\theta}})^{\frac{1}{1-\rho}} \right]^{1-\rho} \right\}. \quad (7)$$

The parameter $\tilde{\theta}$ governs the dispersion of skills, with a higher value of $\tilde{\theta}$ corresponding to *smaller* dispersion. The parameter ρ determines the correlation of an individual's skills. If $\rho = 0$, a person's talents are uncorrelated across occupations while if $\rho = 1$, talents are perfectly correlated for that person.⁷

It turns out to be convenient for notation to transform the parameters slightly. In particular, define $\theta \equiv \tilde{\theta}/(1 - \rho)$ and $T_{ig} \equiv \tilde{T}_{ig}^{\frac{1}{1-\rho}}$. The joint distribution can then be written as

$$F_g(\epsilon_1, \dots, \epsilon_N) = \exp \left\{ - \left[\sum_{i=1}^N T_{ig} \epsilon_i^{-\theta} \right]^{1-\rho} \right\}. \quad (8)$$

The T_{ig} parameters differ across occupations and, potentially, groups. Across occupations, differences in T 's are easy to understand. For example, talent is easy to come by in some occupations and scarce in others. The way we formulate the model, the differences in T 's across occupations (for all groups) will be isomorphic to the sector-specific productivities that we introduce below. When we observe very few individuals in a given occupation, it could be that talent for this occupation is scarce or that this occupation is less productive than other occupations.

More important for our purposes is the potential that the T 's differ across groups within a given occupation. We allow for this possibility between men and women but not between blacks and whites. Specifically, in some occupations, brawn may be a

⁷This generalization of the more common i.i.d. case was suggested in footnote 14 of Eaton and Kortum (2002) and explored by Ramondo and Rodriguez-Clare (2013) in the trade context.

desirable attribute. If men are physically stronger than women on average, then one would expect to observe more men in occupations requiring more physical strength, such as firefighting or construction. To account for this, T_{ig} may be higher in these occupations for white and black men relative to white and black women.

2.3. Occupational choice

The occupational choice problem reduces to picking the occupation that delivers the highest value of U_{ig} . Because talent is drawn from an extreme value distribution, the highest utility can also be characterized by an extreme value distribution, a result reminiscent of McFadden (1974). The overall occupational share can then be obtained by aggregating the optimal choice across people, as we show in the next proposition. (Proofs of the propositions are given in the [online appendix](#).)

Proposition 1 (Occupational Choice): *Let p_{ig} denote the fraction of people in group g that work in occupation i . Aggregating across people, the solution to the individual's occupational choice problem leads to*

$$p_{ig} = \frac{\tilde{w}_{ig}^\theta}{\sum_{s=1}^N \tilde{w}_{sg}^\theta} \quad \text{where} \quad \tilde{w}_{ig} \equiv \frac{T_{ig}^{1/\theta} w_i s_i^{\phi_i} (1 - s_i)^{\frac{1-\eta}{\beta}}}{\tau_{ig}}. \quad (9)$$

Equation (9) says that 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.⁸ The occupational distribution is driven by *relative* returns and not absolute returns: forces that only change \tilde{w} for *all* occupations have no effect on the occupational distribution. This reward depends on the mean talent parameter in the occupation T_{ig} , the post-friction wage per efficiency unit in the occupation $\frac{w_{ig}}{\tau_{ig}}$, and the amount of time spent accumulating human capital by a person in that occupation s_i .⁹ Technological change affects occupational choice through the price per unit of skill, w_i . For example, technological innovations in the home sector emphasized by Greenwood, Seshadri and Yorukoglu (2005) can be viewed as a decline in w_i in the home sector.

⁸See Luttmer (2008) for a similar result.

⁹Notice that human capital enters twice, once as a direct effect on efficiency units and once indirectly, capturing the fact that a person who gets a lot of education has lower leisure.

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

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

$$\mathbb{E}[h_i \epsilon_i] = \gamma \left[\eta^\eta s_i^{\phi_i} \left(\frac{w_i(1 - \tau_i^w)}{1 + \tau_i^h} \right)^\eta \left(\frac{T_i}{p_{ig}} \right)^{\frac{1}{\theta}} \right]^{\frac{1}{1-\eta}} \quad (10)$$

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

Notice that average quality is inversely related to the share of the group in the occupation p_{ig} . This captures the selection effect. For example, the model predicts that only the most talented female lawyers (such as Sandra Day O'Connor) would have chosen to be lawyers in 1960. 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.

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

Proposition 3 (Occupational Wage Gaps): *Let \overline{wage}_{ig} denote the average earnings in occupation i by group g . Its value satisfies*

$$\overline{wage}_{ig} \equiv (1 - \tau_{ig}^w) w_i \mathbb{E}[h_i \epsilon_i] = (1 - s_i)^{-1/\beta} \gamma \bar{\eta} \left(\sum_{s=1}^N \tilde{w}_{sg}^\theta \right)^{\frac{1}{\theta} \cdot \frac{1}{1-\eta}}. \quad (11)$$

In turn, the occupational wage gap between any two groups is the same across all occupations. Specifically,

$$\frac{\overline{wage}_{ig}}{\overline{wage}_{i,wm}} = \left(\frac{\sum_s \tilde{w}_{sg}^\theta}{\sum_s \tilde{w}_{s,wm}^\theta} \right)^{\frac{1}{\theta} \cdot \frac{1}{1-\eta}}. \quad (12)$$

Equation (11) states that average earnings for a given group differs across occupations only because of the first term, $(1 - s_i)^{-1/\beta}$. Occupations in which schooling is especially productive (a high ϕ_i and therefore a high s_i) will have higher average earnings, and that is the only reason for differences across occupations in the model. Average earnings are no higher in occupations where a group faces less discrimination 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. Because of the selection effect, the wage gap between two groups is the *same* for all occupations.

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.

Putting together the equations for the occupational share and the wage gap, we get the propensity of a group to work in an occupation:

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

Equation (13) states that the propensity of a group to work in an occupation (relative to white men) depends on three terms: relative mean talent in the occupation (arguably equal to one for many occupations), the relative occupational friction, and the average 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 of relative mean talent and occupational frictions. This will be the key equation we take to the data.

2.4. Aggregate Productivity

To close the model, assume a representative firm produces aggregate output Y from labor in N occupations:

$$Y = \left(\sum_{i=1}^N (A_i H_i)^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}} \quad (14)$$

where H_i denotes the total efficiency units of labor and A_i is the exogenously-given productivity in occupation i . In turn, H_i is defined as

$$H_i = \sum_{g=1}^G q_g p_{ig} \cdot \mathbb{E} [h_{ig} \epsilon_{ig} \mid \text{Person chooses } i]. \quad (15)$$

where q_g denotes the total number of people in group g . H_i is the product of average human capital and the number of people in the occupation, summed over groups.

2.5. Equilibrium

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

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

$$U(\tau^w, \tau^h, w, \epsilon) = \max_{c, e, s} (1-s)c^\beta \text{ s.t. } c = (1 - \tau_{ig}^w)w\epsilon h(e, s) - e(1 + \tau_{ig}^h). \quad (16)$$

2. Each individual chooses the occupation that maximizes his or her utility: $i^* = \arg \max_i U(\tau_{ig}^w, \tau_{ig}^h, w_i, \epsilon_i)$, taking $\{\tau_{ig}^w, \tau_{ig}^h, w_i, \epsilon_i\}$ as given.
3. A representative firm hires H_i in each occupation to maximize profits:

$$\max_{\{H_i\}} \left(\sum_{i=1}^N (A_i H_i)^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}} - \sum_{i=1}^N w_i H_i \quad (17)$$

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

$$H_i = \sum_{g=1}^G q_g p_{ig} \cdot \mathbb{E} [h_{ig} \epsilon_{ig} \mid \text{Person chooses } i]. \quad (18)$$

5. Total output is given by the production function in equation (14).

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

Proposition 4 (Solving the General Equilibrium): *The general equilibrium of the model is $\{p_{ig}, H_i^{\text{supply}}, H_i^{\text{demand}}, w_i\}$ and Y such that*

1. p_{ig} satisfies equation (9).

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

$$H_i^{supply} = \gamma \bar{\eta} w_i^{\theta-1} (1 - s_i)^{(\theta(1-\eta)-1)/\beta} s_i^{\theta\phi_i} \sum_g q_g T_{ig} \frac{(1 - \tau_{ig}^w)^{\theta-1}}{(1 + \tau_{ig}^h)^{\eta\theta}} \left(\sum_{s=1}^N \tilde{w}_{sg}^\theta \right)^{\frac{1}{\theta} \cdot \frac{1}{1-\eta} - 1} \quad (19)$$

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

$$H_i^{demand} = \left(\frac{A_i^{\frac{\sigma-1}{\sigma}}}{w_i} \right)^\sigma Y \quad (20)$$

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

5. Total output is given by the production function in equation (14).

Note that we have modeled the barriers as taxes driving a wedge between the wages paid by firms and those received by women and blacks (or between education spending by women and blacks and the education output sold by firms). For this reason firms earn zero profits. One could equivalently think of the taxes as rents earned by discriminating firms. Alternatively, we could have assumed firms both discriminate against women and blacks and discriminate in favor of white men such that these rents are zero in each occupation even in the absence of taxes. Our results are robust to this alternative.

2.6. Intuition

To develop intuition, consider the following simplified version of the model. First, assume only two groups, men and women, and assume 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, τ_{ig} affects the average wage and occupational choices of group g but has no effect on other groups. Third, assume $\phi_i = 0$ (no schooling time) and $T_{ig} = 1$ (mean occupational talent is the same for every group). Aggregate output is then equal to the sum of wages paid to men and wages paid to women

(gross of the labor market friction):

$$Y = q_m \cdot \overline{wage}_m + q_w \cdot \frac{\overline{wage}_w}{1 - \bar{\tau}^w} \quad (21)$$

where $\bar{\tau}^w$ denotes the earnings-weighted average of the labor market friction facing women.¹⁰

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

$$\overline{wage}_m = \left(\sum_{i=1}^N A_i^\theta \right)^{\frac{1}{\theta} \cdot \frac{1}{1-\eta}} \quad (22)$$

$$\overline{wage}_w = \left(\sum_{i=1}^N \left(\frac{A_i (1 - \tau_i^w)}{(1 + \tau_i^h)^\eta} \right)^\theta \right)^{\frac{1}{\theta} \cdot \frac{1}{1-\eta}} \quad (23)$$

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 productivities and distortions.

To see the effect of the distortions, assume $\tau^h = 0$ and that $1 - \tau_i^w$ and A_i are jointly log-normally distributed. The average female wage is then equal to:

$$\ln \overline{wage}_w = \ln \left(\sum_{i=1}^N A_i^\theta \right)^{\frac{1}{\theta} \cdot \frac{1}{1-\eta}} + \frac{1}{1-\eta} \cdot \ln(1 - \bar{\tau}^w) - \frac{1}{2} \cdot \frac{\theta - 1}{1-\eta} \cdot \text{Var}(\ln(1 - \tau_i^w)). \quad (24)$$

The first term says that the average female wage is increasing in the power mean of occupational productivities. The second term states 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 third term says that the average female wage is decreasing in the dispersion of $1 - \tau^w$, and more so the greater the importance of goods for human capital.

This simple example shows the fundamental effects of labor market frictions in our model. Both the mean and dispersion of τ_w reduce productivity and lower the average female wage. The productivity losses come from two sources: underinvestment in

¹⁰That is, $\bar{\tau}^w \equiv \sum_{i=1}^N \omega_i \tau_i^w$ where $\omega_i \equiv \frac{\frac{p_{i,w} \overline{wage}_w}{1 - \tau_i^w}}{\sum_{j=1}^N \frac{p_{j,w} \overline{wage}_w}{1 - \tau_j^w}}$.

human capital (tied to weighted mean levels of τ_w) and misallocation of female talent across occupations (tied to the dispersion of τ_w across occupations).

Though we will not impose these simplifying assumptions, with this motivation in mind we will later isolate the potential productivity loss due to the misallocation of talent versus underinvestment in human capital for all groups facing discrimination.

3. Empirically Evaluating the Occupational Sorting Model

3.1. Data

We use data from the 1960, 1970, 1980, 1990, and 2000 Decennial Censuses and the 2006-2008 American Community Surveys (ACS).¹¹ We make four restrictions to the data. 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 55. 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. [Appendix Table B.1](#) reports the summary statistics of our sample.¹²

A key to our analysis is a consistent definition of occupations over time. First, we treat the home sector as a separate occupation. We define a person who is not currently employed or who works less than ten hours per week as working in the home sector. Those who are employed but usually work between ten and thirty hours per week are classified as being part-time workers. We split the sampling weight of part-time workers equally between the home sector and the occupation in which they are working. Individuals working more than thirty hours per week are considered to be working full-time in an occupation outside of the home sector. Second, we define the non-home occupations using the roughly 70 occupational sub-headings from the 1990 Census occupational classification system.¹³ [Appendix Table B.2](#) reports the 67 occupations we

¹¹When using the 2006-2008 ACS data, we pool all three years together and treat them as one cross section. Henceforth, we refer to the pooled 2006-2008 sample as the 2008 sample.

¹²For all analysis in the paper, we apply the sample weights available in the different surveys.

¹³<http://usa.ipums.org/usa/volii/99occup.shtml>. We use the 1990 occupation codes as our basis because the 1990 codes are available in all Census and ACS years since 1960. We start our analysis in 1960, as

analyze. Some samples of the occupational categories are “Executives, Administrators, and Managers”, “Engineers”, “Natural Scientists”, “Health Diagnostics”, “Health Assessment”, and “Lawyers and Judges”.¹⁴ To assess robustness, we also use a more detailed classification of occupations into 340 occupations defined consistently since 1980 as well as 20 broad occupational groups defined consistently since 1960.

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, with more than 1000 dollars of earnings (in 2007 dollars) in the previous year, and who worked on average more than 30 hours per week. We define the hourly wage as total annual earnings divided by total hours worked in the previous year.¹⁵

3.2. Occupational Sorting and Wage Gaps By Group

We begin by documenting the degree of convergence in the occupational distribution between white men and the other groups over the last fifty years. We measure similarity in the occupational distribution with the following index:

$$\Psi_g \equiv 1 - \frac{1}{2} \sum_{i=1}^N |p_{i,wm} - p_{ig}| \quad (25)$$

Ψ_g is a function of the sum across occupations of the absolute value of the difference in the propensity of the group relative to white men in the occupation. We normalize the index so that $\Psi_g = 0$ implies no overlap in the occupational distribution and $\Psi_g = 1$ implies an identical occupational distribution between the group and white men.

Panel A of Table 1 presents the occupational similarity index for white women, black men, and black women for all members of each group and separately for less-educated (high school degree or less) and highly-educated individuals (more than high school).¹⁶

A few things of note from this table. First, each group saw substantial occupa-

this is the earliest year for which the 1990 occupational crosswalk is available.

¹⁴Appendix Table B.3 gives a more detailed description of some of these occupational categories.

¹⁵We impute the average wage in the home sector from the group composition and average schooling of individuals in the home sector, assuming that the relationship between income and these characteristics are the same in the home sector as in the market sector.

¹⁶We exclude the home sector in the estimates shown in Table 1, but broad patterns are very similar – particularly the index for white women — when the home sector is included.

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

Panel A: Occupational Similarity Index, Relative to White Men					
Race-Sex Group	1960	1980	2008	1980-1960 Difference	2008-1980 Difference
White Women: All	0.42	0.49	0.55	0.07	0.06
White Women: High Educated	0.38	0.49	0.59	0.10	0.12
White Women: Low Educated	0.40	0.43	0.46	0.02	0.04
Black Men: All	0.56	0.72	0.76	0.16	0.04
Black Men: High Educated	0.51	0.74	0.77	0.23	0.03
Black Men: Low Educated	0.59	0.75	0.75	0.16	0.00
Black Women: All	0.28	0.42	0.50	0.14	0.08
Black Women: High Educated	0.31	0.44	0.53	0.13	0.11
Black Women: Low Educated	0.27	0.41	0.44	0.14	0.03
Panel B: Conditional Log Difference in Wages, Relative to White Men					
Race-Sex Group	1960	1980	2008	1980-1960 Difference	2008-1980 Difference
White Women: All	-0.57	-0.47	-0.26	0.10	0.21
White Women: High Educated	-0.50	-0.40	-0.24	0.10	0.16
White Women: Low Educated	-0.56	-0.47	-0.27	0.09	0.20
Black Men: All	-0.38	-0.22	-0.15	0.16	0.07
Black Men: High Educated	-0.29	-0.16	-0.18	0.13	-0.02
Black Men: Low Educated	-0.38	-0.23	-0.12	0.15	0.11
Black Women: All	-0.86	-0.48	-0.31	0.38	0.17
Black Women: High Educated	-0.62	-0.39	-0.31	0.23	0.08
Black Women: Low Educated	-0.88	-0.48	-0.28	0.40	0.20

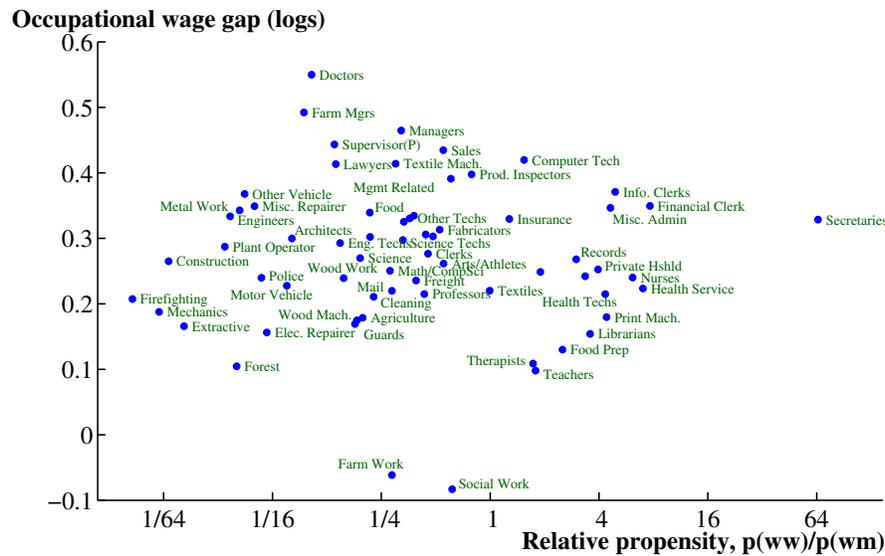
Note: Panel A of the table reports our occupational similarity index for white women, black men, and black women relative to white men in 1960, 1980, and 2008. The occupational similarity index runs from zero (no overlap with the occupational distribution for white men) and one (identical occupational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

tional convergence relative to white men between 1960 and 2008. Second, the timing of the convergence differed between women and black men. The occupations of women (both black and white) converged towards those of white men over the entire time period, while those of black men converged largely prior to 1980. Third, convergence in occupations is largely driven by highly-educated individuals in each group, with very little convergence for the less-educated. This pattern is particularly striking among white women. In 1960, there were substantial occupational differences both between high educated white women and high educated white men, and between low educated white women and low educated white men. Low educated white men worked primarily in construction and manufacturing, while low educated white women worked primarily as secretaries or in low skilled services like food service. High educated white men in 1960 were spread out across many high skilled occupations, while high educated white women worked primarily as teachers and nurses. Between 1960 and 2008, the occupations of higher educated white men and women converged dramatically, while occupational similarity between lower educated white men and women remained largely unchanged. Today, low skilled women still primarily work in services and office support occupations, while low skilled men still primarily work in construction and manufacturing.¹⁷

A strong prediction of our model is that the convergence in the occupational distribution documented in Table 1 will narrow the wage gap in *all* occupations. Moreover, it should not have a larger effect on the wage gap in an occupation where relative propensities converged versus one where relative propensities remain unchanged. We begin to examine these predictions in Panel B of Table 1. As is well known, wages of women have converged with those of white men over the last fifty years. What is perhaps less well known is that the wage convergence is exactly the same for less-educated women as for more educated women. From 1960 to 2008, the wage gap between highly-educated women and comparable white men narrowed by 29 log points. Over the same time period, the wage gap between less-educated women and white men narrowed by 26 log points, despite the fact that the gap in the occupational distribution between less-educated women and white men remained essentially unchanged. This is exactly what our model predicts: changes in the τ_{ig} 's for white women in high ϕ occupations results

¹⁷Blau, Brummund and Liu (2012) documents similar results pertaining to the occupational convergence of women relative to men by education.

Figure 1: Occupational Wage Gaps for White Women in 1980



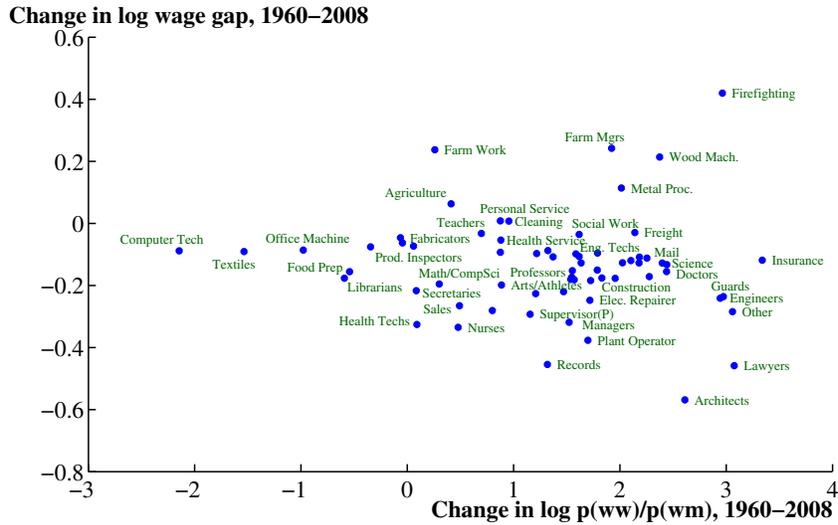
Note: The figure shows the relationship between the (log) occupational wage gap for white women compared to white men and the relative propensity to work in the occupation between white women and white men, $p_{i,ww}/p_{i,wm}$.

in entry and lower average quality of women in these occupations and exit and higher average quality in low ϕ occupations. The overall wage for women increases due to the improved allocation, but by the same magnitude in the high ϕ as in the low ϕ occupations.¹⁸

Figure 1 provides additional evidence that the wage gap in an occupation is uncorrelated with the frictions in that occupation. Specifically, the figure plots the (log) wage gap of white women relative to white men in an occupation against the relative propensity to work in that occupation $p_{i,ww}/p_{i,wm}$ in 1980. A white woman was 65 times more likely than a white man to work as a secretary, but only 1/7 as likely to work as a lawyer. Given this enormous variation, the difference in the wage gaps between these two occupations is remarkably small. White women secretaries earned about 33 percent less

¹⁸For black men, the wage gap evolved similarly for less vs. more educated individuals from 1960 to 1980. After 1980, however, there was little change in occupational similarity for either high or low skilled black men, and there was no change in the wage gap for high skilled black men. The wage gap for low skilled black men, however, continued to narrow after 1980. This may be due to the rapid decline in labor market participation of low skilled black men during the last thirty years, if it was not random. As currently formulated, our model would not predict such a trend. However, as we discuss in Section 5, the change in labor market outcomes for black men between 1980 and 2008 do not materially affect our estimates of aggregate productivity gains.

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



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

than white men secretaries in 1980, while the gap was 41 percent for lawyers. Looking across the 67 occupations, there is no systematic relationship between the wage gap and the relative propensity in the occupation.¹⁹ The patterns in other years and for other groups were quite similar.²⁰

Figure 2 provides similar evidence over time, again for white women relative to white men. This figure shows that the change in the wage gap from 1960 to 2008 is also uncorrelated with the change in the relative propensities to work in the occupation. For example, the relative fraction of white women who are doctors increased by 144 percent between 1960 and 2008. For nurses, in contrast, the relative fraction who are white women decreased by 52 percent. Yet the relative wage gap between white

¹⁹The coefficient from a regression of the occupational wage gap on $\log p_{i,ww}/p_{i,wm}$ was 0.002 with a standard error of 0.008 and an adjusted R-squared of essentially zero. For interpretation, the standard deviation of the independent variable was 1.96 and the mean of the dependent variable was -0.31. The regression weighted occupations by the share of all workers (across all groups) in the occupation. Albanesi and Olivetti (2009b) find similar evidence using the PSID.

²⁰There is, however, a relationship between the occupational wage gaps and the average earnings of individuals in the occupations. On average, high income occupations tended to have larger wage gaps. Nonetheless, the magnitude of this correlation was almost always small. For example, in 2006-2008, white working women had about a 3 percentage point larger wage gap relative to white men in response to a one-standard deviation increase in occupational log income. As seen from Table 1, the average wage gap was 26 percentage points.

men and white women narrowed by 20 to 30 log points in both occupations. From the perspective of the model, the weak relationship between wages gaps and propensities is not surprising. Within the model, the relative propensity, not the wage gap, reveals frictions facing a group in an occupation.

4. Estimating the Frictions

Motivated by our model, we now use data on occupational propensities and average wage gaps to infer the frictions. Specifically, given equations (13) and (12), we can define the composite friction measure for each group (relative to white men) in each occupation as:

$$\hat{\tau}_{ig} \equiv \frac{\tau_{ig}}{\tau_{i,wm}} \left(\frac{T_{i,wm}}{T_{i,g}} \right)^{\frac{1}{\theta}} = \left(\frac{p_{ig}}{p_{i,wm}} \right)^{-\frac{1}{\theta}} \left(\frac{\overline{\text{wage}}_g}{\overline{\text{wage}}_{wm}} \right)^{-(1-\eta)}. \quad (26)$$

This equation has the following interpretation. If a group is either underrepresented in an occupation or if it faces a large average wage gap, the right-hand side will be high. The model can explain this in one of two ways (on the left side): either the group faces a large composite barrier, or it has a relatively low mean talent in that occupation (e.g. women in occupations where brute strength is important). We observe the right-hand side of this equation in the data and therefore use it to back out the average relative distortion (or talent) between groups, $\hat{\tau}_{ig}$.

To implement this calculation, we require estimates of θ and η . The parameter θ governs the dispersion of comparative advantage. Given the occupational choice model above, the dispersion of wages across people within an occupation-group obeys a Fréchet distribution with the shape parameter $\theta(1-\rho)(1-\eta)$: the lower is this shape parameter, the *more* wage dispersion there is within an occupation. Wage dispersion therefore depends on the dispersion of comparative advantage (governed by $1/\theta$), the correlation of ability across occupations (ρ), and amplification from human capital (governed by $1/(1-\eta)$). In particular, the coefficient of variation of wages within an occupation-group in our model satisfies:

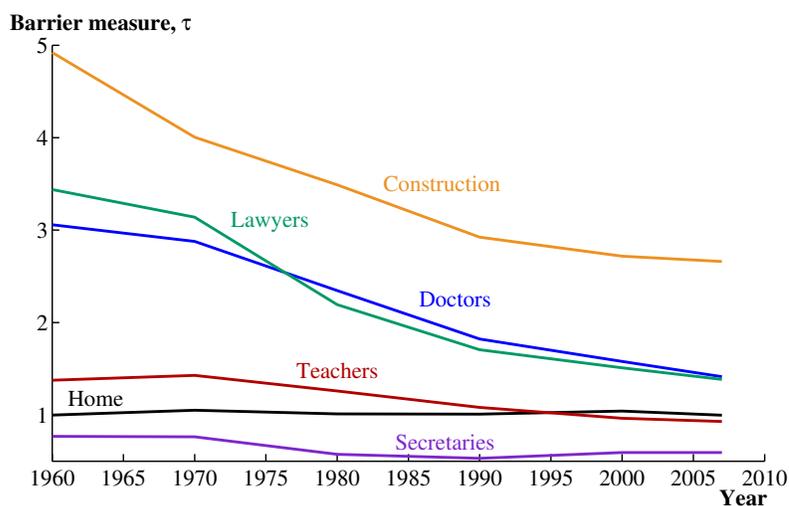
$$\frac{\text{Variance}}{\text{Mean}^2} = \frac{\Gamma\left(1 - \frac{2}{\theta(1-\rho)(1-\eta)}\right)}{\left(\Gamma\left(1 - \frac{1}{\theta(1-\rho)(1-\eta)}\right)\right)^2} - 1. \quad (27)$$

To estimate $\theta(1 - \rho)(1 - \eta)$, therefore, we look at wage dispersion within occupation-groups. We take residuals from a cross-sectional regression of log worker wages on 66x4 occupation-group dummies. These span the 66 market occupations and the four groups of white men, white women, black men, and black women. The wage is the hourly wage, and the sample includes both full-time and part-time workers. The dummies should capture the impact of schooling (ϕ_i levels) on average wages in an occupation, as well as the wage gaps created by frictions (the average τ_{ig} across occupations). We calculate the mean and variance across workers of the exponent of these wage residuals. We then numerically solve equation (27) for $\theta(1 - \rho)(1 - \eta)$. Sampling error is minimal because there are 300-400k observations per year for 1960 and 1970 and 2-3 million per year for 1980 onward. The point estimates for $\theta(1 - \rho)(1 - \eta)$ average 3.12. They drift down over time, from 3.3 in 1960 to 2.9 in 2008, consistent with rising residual wage inequality.

To estimate the distortions above, we need to net out the contribution of absolute advantage (governed by ρ) and temporary factors to wage dispersion. We thus make several adjustments, all of which serve to reduce residual wage dispersion. First, we compress the variance 4% for the share of wage variance explained by AFQT scores in Rodgers and Spriggs (1996). AFQT scores arguably pick up absolute ability for many jobs. Second, we control for individual education, hours worked, and potential experience in the Census data. A worker's education could proxy for absolute advantage across many occupations, there could be compensating differentials associated with the workweek, and experience is not the same as lifelong comparative advantage. Finally, we compress the variance of the residuals by another 14% to reflect an estimate of transitory wage movements from Guvenen and Kuruscu (2009). Temporary wage differences across workers are not a source of enduring comparative advantage.

These adjustments cumulatively explain 25% of wage variation within occupation-groups. Attributing the remaining 75% of wage dispersion to comparative advantage, we arrive at a baseline value of $\theta(1 - \eta) = 3.44$. We also explore the sensitivity of our results to alternate values of ρ : when computing our counterfactuals in Section 5, we show results where we set $\theta(1 - \eta)$ such that only 50%, 25% or 10% of wage dispersion within occupation-groups is due to comparative advantage.

The parameter η denotes the elasticity of human capital with respect to educa-

Figure 3: Estimated Barriers ($\hat{\tau}_{ig}$) for White Women

Note: Author's calculations based on equation (26) using Census data and imposing $\theta = 3.44$ and $\eta = 1/4$.

tion spending. Related parameters have been discussed in the literature, for example by Manuelli and Seshadri (2005) and Erosa, Koreshkova and Restuccia (2010). In our model, η will equal the fraction of output spent on human capital accumulation. Absent any solid evidence on this parameter, we set $\eta = 1/4$ in our baseline and explore robustness to $\eta = 0$ and $\eta = 1/2$. This parameter affects the *level* of the τ_{ig} parameters, but little else in our results.

Figure 3 presents our estimates of $\hat{\tau}_{ig}$ for white women for a select subset of our occupations. Consider the results for white women in the “home” occupation in 1960. Despite white women being 7 times as likely to work in the home sector as white men, we estimate $\hat{\tau}_{ig}$ for white women in that sector to be just below 1 (0.99) — that is, we estimate zero frictions in the home sector for these women. This implies that white women in 1960 did not have an absolute advantage over white men in the home sector. In essence, we find that white women were choosing the home sector because they were facing disadvantages in *other* occupations.

Figure 3 shows that $\hat{\tau}_{ig}$ is close to 1 for white women in the home sector in all years of our analysis. This suggests women did not move out of the home sector because they lost any absolute advantage in the home sector. Instead, our results suggest that women moved into market occupations due to declining barriers in the market. Below

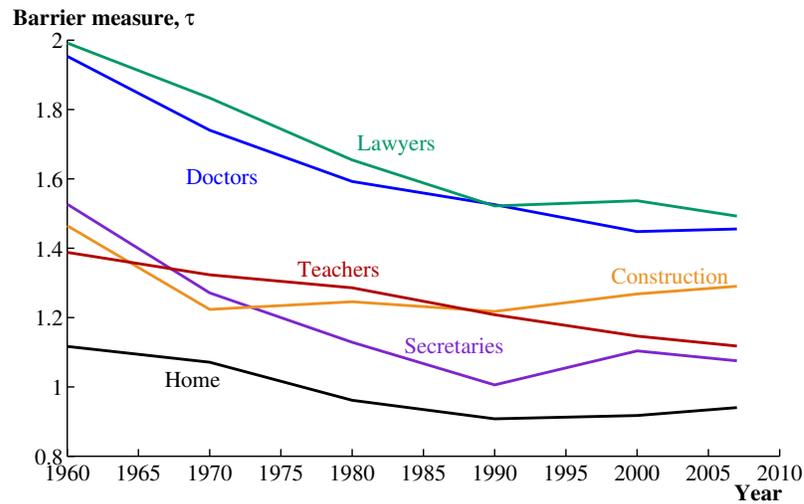
we will show that changes in the productivity of the home sector relative to the market sector for all groups also contributed to women exiting the home sector. To preview our results, we find that changing productivity at home versus in the market explains roughly 25 percent of the movement of white women out of the home sector. The remaining 75 percent is due to changes in the $\hat{\tau}_{ig}$ in the market sector.

The remainder of the results from Figure 3 highlight that the $\hat{\tau}_{ig}$'s for white women changed dramatically in certain occupations. For example, the $\hat{\tau}_{ig}$ for white women lawyers and doctors in 1960 ranged from 3.0 to 3.5. Interpreting these as labor market frictions, it's as if women lawyers in 1960 received only 1/3 of their marginal products. The low participation of white women in these occupations in 1960 causes us to infer high values of $\hat{\tau}_{ig}$.

Interestingly, the $\hat{\tau}_{ig}$ for white women teachers is also greater than one in 1960. While white women were 1.7 times more likely than white men to work as teachers, this propensity is more than offset by the overall wage gap in 1960, where women earned about 0.57 times what men earned. If white women were not facing some friction or lower absolute advantage in the teacher occupation, our model predicts there should have been an even higher fraction ending up as teachers in 1960.

Contrast this with secretaries. A white woman in 1960 was 24 times more likely to work as a secretary than was a white man. The model explains this enormous discrepancy by assigning a $\hat{\tau}_{ig}$ of 0.76 for white women secretaries. A τ below 1 is like a subsidy, so the model says either white women had an absolute advantage as secretaries, or there was discrimination that encouraged women into this profession. White women also had very high $\hat{\tau}_{ig}$ values in the construction, firefighting and vehicle mechanic professions, which could reflect frictions or low T_{ig} for women if physical strength was important in these occupations.

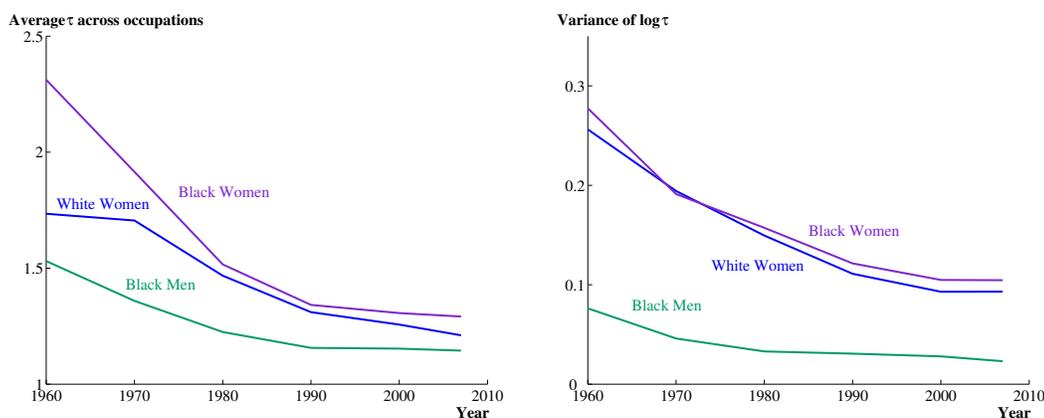
For lawyers, and doctors, the $\hat{\tau}_{ig}$'s fell from around 3.0 to 1.4 between 1960 and 2008. School teachers also saw a substantial fall in their average $\hat{\tau}_{ig}$ from 1.4 to around 1. While barriers facing white women fell in many skilled professional occupations, their $\hat{\tau}_{ig}$ values did not change much for low skilled occupations. This is particularly true after 1980. For example, the estimated $\hat{\tau}_{ig}$ for white women barely changed (or rose) for secretaries or construction workers between 1980 and 2008. Yet the $\hat{\tau}_{ig}$'s for doctors, lawyers, and teachers continued to fall during this period. These results are consistent

Figure 4: Estimated Barriers ($\hat{\tau}_{ig}$) for Black Men

with our earlier evidence that occupational convergence from 1980–2000 was primarily among high skilled individuals.

The $\hat{\tau}_{ig}$'s for black men — for these same select occupations — are shown in Figure 4. A similar overall pattern emerges, with the $\hat{\tau}_{ig}$'s being substantially above 1 in general in 1960, but falling through 2008. Still, they remained above 1 by 2008, especially for the high-skilled occupations, suggesting barriers remain. Unlike for white women, almost the entire change in the $\hat{\tau}_{ig}$ for black men occurred prior to 1980. The plots for black women look like a combination of those for white women and black men.

Figure 5 presents the mean and variance of $\hat{\tau}_{ig}$ across occupations for each group over time. The left panel shows the average $\hat{\tau}_{ig}$ falling over time for each of the groups. For women, the decline in average $\hat{\tau}_{ig}$ occurred throughout the period. For black men, the decline was concentrated prior to 1980. The right panel of Figure 5 shows that, in 1960, the $\hat{\tau}_{ig}$'s were also dispersed across occupations for blacks and (especially) women. This dispersion leads to misallocation of talent across occupations (at least to the extent that it reflects frictions instead of changing T 's or \bar{h} 's). The decline in the mean $\hat{\tau}_{ig}$ for the groups can also explain some of U.S. productivity growth over the last half century. The extent to which the changing $\hat{\tau}_{ig}$'s contribute to productivity growth depends on the root causes. We turn to these issues next.

Figure 5: Means and Variances of $\hat{\tau}_{ig}$ Over Time

Note: The left panel shows the average level of $\hat{\tau}$ by group, weighted by total earnings in each occupation in 2008. The right panel shows the variance of $\log \hat{\tau}$, weighted in the same way.

5. Estimating Productivity Gains

5.1. Parameter Values and Exogenous Variables

The key parameters of the model — assumed to be constant over time — are η , θ , σ , and β . We discussed the estimation and assumptions for η (the elasticity of human capital with respect to goods invested) and θ (the parameter governing the dispersion of talent) above. The parameter 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 an elasticity of substitution $\sigma = 3$ for our baseline value. We explore robustness to a wide range of values for σ .

The parameter β is the geometric weight on consumption relative to time in an individual's utility function (1). As schooling trades off time for consumption, wages must increase more steeply with schooling in equilibrium when people value time more (i.e. when β is lower). We choose $\beta = 0.693$ to match the Mincerian return to schooling across occupations, which averages 12.7% across the six decades.²¹ Our results are es-

²¹Workers must be compensated for sacrificing time to schooling the more they care about time relative to consumption. The average wage of group g in occupation i is proportional to $(1 - s_i)^{\frac{-1}{\beta}}$. If we take a log linear approximation around average schooling \bar{s} , then β is inversely related to the Mincerian return to schooling across occupations (call this return ψ): $\beta = (\psi(1 - \bar{s}))^{-1}$. We calculate s as 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 $\beta = 0.693$, the simple average of the implied β values

entially invariant to this parameter, as documented later.

As our model is static, we infer exogenous variables separately for each decade. In each year, we have $6N$ variables to be determined. For each of the $i = 1, \dots, N$ occupations these are A_i , ϕ_i , and τ_{ig} , where g stands for white men, white women, black women, or black men. We also allow population shares of each group q_i to vary by year to match the data. To begin, we normalize average ability to be the same in each occupation-group ($T_{ig} = 1$). Differences in ability across occupations (T_i) are isomorphic to differences in the production technology A_i . Across groups within an occupation, we think the natural starting point is *no* differences in mean ability in any occupation; this assumption will be relaxed in our robustness checks. Additionally, we set $\tau_{i,wm} = 1$ in all occupations. This restriction implies that white men face no occupational barriers. Again, we think this is a natural benchmark to consider.²² Finally, we normalize the \bar{h}_{ig} 's to be one for all groups in all occupations. As discussed above, the \bar{h}_{ig} 's are isomorphic to the τ_{ig}^h 's. Even though the magnitude of the productivity gains are the same, however, the underlying economic mechanisms differ under in the two cases. We discuss the \bar{h}_{ig} interpretation in greater depth below.

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

- ($4N - 4$) The fraction of people from each group working in each occupation, p_{ig} . (Fewer than $4N$ moments because the p_{ig} sum to one for each group.)
- (N) The average wage in each occupation.
- (N) The assumption that $\tau_{i,wm} = 1$ in each occupation.
- (3) Average wage gaps across all occupations between white men and each of our 3 other groups.
- (1) Average years of schooling in one occupation.

As shown above, the $3N$ $\hat{\tau}_{ig}$ variables are easy to identify from the data given our setup. Assuming that $\tau_{i,wm} = 1$, $T_{ig} = 1$, and $\bar{h}_{ig} = 1$, the τ_{ig} 's for the other groups across years. This method allows the model to approximate the Mincerian return to schooling across occupations.

²²Our results are robust to instead setting the $\tau_{i,wm}$ such that wages equal output in each occupation.

are easy to infer using equation (26). But recall that $\tau_{ig} \equiv \frac{(1+\tau_{ig}^h)^\eta}{1-\tau_{ig}^w}$. From the data we currently have, we cannot separately identify the τ^h and τ^w components of τ . That is, we cannot distinguish between human capital barriers and labor market barriers. We proceed by considering two polar cases. At one extreme, we assume all of the τ_{ig}^w 's are zero, so that τ_{ig} solely reflects τ_{ig}^h . At the other extreme, we set all of the $\tau_{ig}^h = 0$ and assume the τ_{ig}^w 's are responsible for the τ 's. In short, we either assume only human capital barriers (the τ^h case) or only labor market barriers (the τ^w case).

The A_i levels and the relative ϕ_i 's across occupations involve the general equilibrium solution of the model, but the intuition for what pins down their values is clear. We already noted that A_i is observationally equivalent to the mean talent parameter in each occupation T_i . The level of A_i helps determine the overall fraction of the population that works in each occupation. We also noted that ϕ_i is the key determinant of average wage differences across occupations. Thus, the data on employment shares and wages by occupation pin down the values of A_i and ϕ_i .²³

5.2. Productivity Gains

Given our model, parameter values, and the forcing variables we infer from the data, we can now answer one of the key questions of the paper: how much of overall growth from 1960 to 2008 can be explained by the changing labor market outcomes of blacks and women during this time period?

In answering this question, the first thing to note is that output growth in our model is a weighted average of earnings growth in the market sector and in the home sector. Earnings growth in the market sector can be measured as real earnings growth in the census data. Deflating by the NIPA Personal Consumption Deflator, real earnings in the census data grew by 1.32 percent per year between 1960 and 2008.²⁴ We impute wages in the home sector using the relationship between average earnings and average education across market occupations and from wage gaps by group in market occupations.

²³From wages in each occupation, we can infer the *relative* values of ϕ_i across occupations. But we cannot pin down the ϕ_i levels, as absolute wage levels are also affected by the A_i productivity parameters. Thus we use a final moment – average years of schooling in one occupation – to determine the ϕ_i levels. We choose to match schooling in the lowest wage occupation, which is Farm Non-Managers in most years. Calling this the “min” occupation, we set ϕ_{min} in a given year to match the observed average schooling among Farm Non-Managers in the same year: $\phi_{min} = \frac{1-\eta}{\beta} \frac{s_{min}}{1-s_{min}}$.

²⁴This might be lower than standard output growth measures because it is calculated solely from wages; for example, it omits employee benefits.

Table 2: Productivity Gains: Share of Growth due to Changing Frictions

	τ^h case	τ^w case
Frictions in all occupations	20.4%	15.9%
Counterfactual: wage gaps halved	12.5%	13.7%
Counterfactual: zero wage gaps	2.9%	11.8%
No frictions in “brawny” occupations	18.9%	14.1%
No frictions in 2008	20.4%	12.3%
Market sector only	26.9%	23.5%

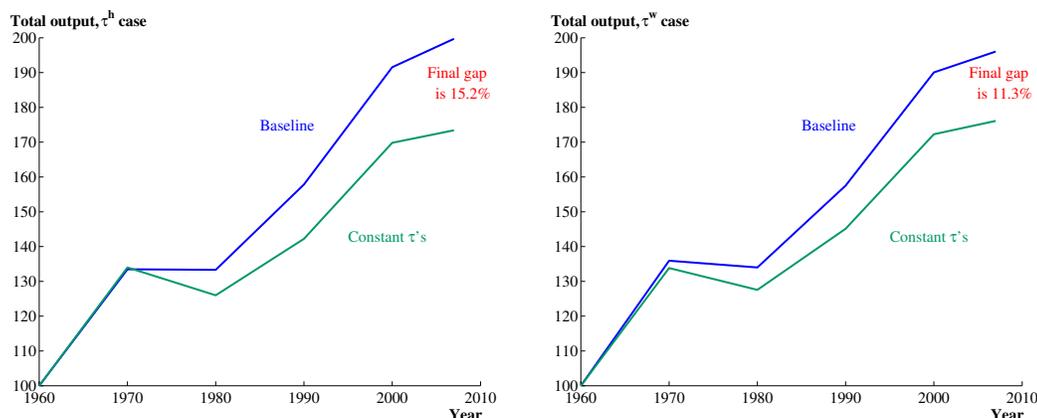
Note: Average annual wage growth between 1960 and 2008 was 1.47%. Entries in the table show the share of labor productivity growth attributable to changing frictions according to our model under various assumptions. In the penultimate line, we assume that there are no frictions for white women in occupations where physical strength is important. Instead, we allow $T_{i,ww}$ to change over time to match the occupational allocation for white women. For blacks in this case, we do allow for frictions, but also assume $T_{i,bw} = T_{i,ww}$. In the last row we assume that there are no frictions in 2008, so that $T_{i,g}$ differences explain all group differences in that year; we then calculate τ 's for earlier years assuming the T 's for 2008 apply to earlier years.

(See the discussion in section 3.1 for additional details.) Taking a weighted average of the imputed wage in the home sector and the wage in the census data, we estimate that output (as defined by our model) grew by 1.47 percent per year between 1960 and 2008.

How much of this growth is due to changing τ 's, according to our model? We answer this question by holding the A 's (productivity parameters by occupation), ϕ 's (schooling parameters by occupation), and q 's (group shares of the working population) constant over time and letting the τ 's change.²⁵

The results of this calculation are shown in Table 2. When the frictions are interpreted as occurring in human capital accumulation (the τ^h case), the change in occupational frictions explain 20.4 percent of overall growth in output per worker over the last half century. If we instead interpret the frictions as occurring in the labor market

²⁵More specifically, we follow the standard approach of chaining. That is, we compute growth between 1960 and 1970 allowing the τ 's to change but holding the other parameters at their 1960 values. Then we compute growth between 1960 and 1970 from changing τ 's holding the other parameters 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 1960–2008.

Figure 6: Counterfactuals: Output Growth due to A , ϕ versus τ 

Note: These graphs show the counterfactual path of output in the model if the τ 's were kept constant over time (in a chained calculation). That is, how much of cumulative growth is due to changing A 's and ϕ 's versus changing τ 's. The left panel is for the τ^h case, and the right is for the τ^w case.

(the τ^w case), the changing τ^w 's account for 15.9 percent of the cumulative growth from 1960 to 2008. The gains are smaller in the τ^w case because some of the wage gaps are accounted for directly by labor market discrimination in this case, with no direct implications for productivity. There are still indirect effects operating through human capital accumulation and occupational choice, of course.

A related calculation is to hold the τ 's constant and calculate growth due to changes in the A 's, ϕ 's, and q 's. Figure 6 does this. The left panel considers the τ^h case. The vast majority of growth is due to increases in A_i and ϕ_i over time, but an important part is attributable to reduced frictions. Allowing the τ^h 's to change as they did historically raises output by 15.2 percent in the τ^h case. The right panel of Figure 6 presents the τ^w case. Here, reduced frictions raised overall output by 11.3% between 1960 and 2008.

Could the productivity gains we estimate be inferred from a simple back-of-the-envelope calculation involving the wage gaps alone? In particular, suppose one takes white male wage growth as fixed, and calculates how much of overall wage growth comes from the faster growth of wages for the other groups. The answer is that faster wage growth for blacks and white women accounts for 13 percent of overall wage growth from 1960 to 2008. This is compared to our estimate of the productivity gains from changing τ 's of 20 percent in the τ^h case and 16 percent in the τ^w case.

Our counterfactuals differ from the back-of-the-envelope calculation in two funda-

mental ways. First, we are isolating the contribution of changing τ 's to labor productivity growth. The back-of-the-envelope calculation includes the effect of changes in A 's, ϕ 's, and q 's on the wage growth of women and blacks relative to white men. Second, our counterfactuals take into account the impact of changing τ 's on white men. In our counterfactuals, we will show shortly, the wage gains to women and blacks come partly at the expense of white males. As women and blacks move into high-skill occupations, this crowds out white men by lowering the wage per unit of human capital in those occupations so long as $\sigma < \infty$, i.e. occupations are not perfectly substitutable. As a result of the above forces, a simple back-of-the-envelope calculation underestimates the baseline models' productivity gains by between 25 and 75 percent depending on the τ^h or τ^w interpretation.

The rest of Table 2 explores some additional counterfactuals which shed light on the robustness of our results. The start, the next two rows illustrate how average wage gaps between groups relate to the productivity gains that we estimate. This counterfactual illustrates how our results are not simply linked one-for-one with the wage gaps we feed into the model. In the second row, we show that if we cut the wage gaps in half in all years, we reduce the share of growth explained from 20.4% to 12.5% in the τ^h case. Setting the average wage gaps in the data to zero leaves only 2.9% of growth explained by changes in the human capital frictions. This is not surprising given the theoretical results shown in Section 2. The wage gap is affected by frictions faced by a group. As the wage gap goes to zero, so do the average frictions. In the τ^w case, in contrast, the gains are relatively insensitive to the wage gaps, falling from 15.9% to 13.7% to 11.8% as wage gaps are halved and then eliminated. In the τ^w case, 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. In the τ^h case, the distortions operate by affecting average human capital investments that do show up in the wage gaps.²⁶ As these contrasting cases show, model productivity gains cannot be gleaned from the wage gaps alone.

The next row in Table 2 considers the robustness of our productivity gains to relaxing the assumption that men and women draw from the same distribution of talent

²⁶The τ^h results are identical to the case where we assume human capital efficiency parameters \bar{h}_{ig} are driving the labor market outcomes of groups. As a result, the wage gap is closely linked to the productivity gains in this case.

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.²⁷ That is, we estimate values for T_{ig} for white women that fully explain their observed allocation to these occupations for 1960, 1970, . . . , 2008. Our hypothesis going into this check was that most of the productivity gains were coming from the rising propensity of women to become lawyers, doctors, scientists, professors, and managers, occupations where physical strength is not important. The results in Table 2 support our hypothesis. The amount of growth explained by changing frictions falls only slightly — for example, from 20.4% to 18.9% in the τ^h case — if we assume that all the movement of women into manufacturing, construction, police, firefighting and other brawny occupations was due to changes in relative comparative advantages in these occupations as opposed to changing τ 's in these occupations.

The penultimate row in Table 2 assumes that all group differences that we currently observe in 2008 are due to differences in talents T_{ig} in that year. These talent differences are then imposed in all prior years. Given these values for T_{ig} , we then estimate the τ_{ig} 's in previous years to fully account for group differences in those years. In the τ_h case, the gains from changing distortions are identical to the baseline case (20.4% of growth). It doesn't matter at all for 1960–2008 gains whether the remaining group differences in 2008 are due to talent or human capital distortions. In the τ_w case, however, the gains are smaller (12.3% of growth vs. 15.9% in the baseline). The labor market distortions are smaller in each year when there are large differences in talent, so removing them from 1960–2008 yields smaller productivity gains.

The final row in Table 2 focuses on the market sector. In the prior examples, we focused on total output inclusive of the home sector. Given that the home sector wages are imputed, one might wonder how much of the growth in market output that we are explaining. Here we calculate market output growth in the model using a Divisia

²⁷These occupations are assigned based on Rendall (2010), who classifies occupations based on the importance of physical strength. We define brawny occupations as those below the median in Rendall (2010).

Table 3: Potential Remaining Output Gains from Zero Barriers

	τ^h case	τ^w case
<i>Frictions in all occupations</i>		
Cumulative gain, 1960–2008	15.2%	11.3%
Remaining gain from zero barriers	14.3%	10.0%
<i>No frictions in “brawny” occupations</i>		
Cumulative gain, 1960–2008	14.0%	10.0%
Remaining gain from zero barriers	11.7%	9.1%

index with Tornqvist shares. The τ changes explain an even higher fraction of growth in *market* output: 26.9% in the τ^h case and 23.5% in the τ^w case (vs. 20.4% and 15.9% when combining the home sector with market occupations). Our results do not hinge on the inclusion of the nonmarket sector in our output measures.

How much additional growth could be had from reducing the frictions all the way to zero? The answer is in Table 3. Consider first the τ^h case. Between 1960 and 2008, changing frictions raised output by 15.2% with our baseline parameters. If the remaining frictions in 2008 were removed entirely, output would be higher by an additional 14.3%. For the τ^w case the gain from eliminating all frictions in 2008 would be 10.0%. Thus there remain substantial gains from removing frictions. We also calculated the remaining gains assuming face different innate ability in the brawny occupations.

As we have stressed, productivity gains can come from reducing misallocation across occupations and from boosting average human capital investments. These, in turn, are tied to the variance and mean of distortions across occupations for a given group. To illustrate their relative importance, consider the following thought experiment: in a given year, first (a) eliminate the variance of barriers across occupations for each group; then (b) eliminate the mean barrier across occupations for each group. Table 4 does this for 1960 and 2008. In the τ^h case, over 80% of the gains from eliminating all barriers arise from reducing misallocation in 1960. In 2008 more than 100% of the gains come from reducing misallocation. In the τ^w case, around 80% of the gains come from

Table 4: Decomposing the Output Gains: Variance vs. Mean Barriers

		τ^h case	τ^w case
1960	Eliminating Dispersion	22.2%	14.9%
	Eliminating Mean and Variance	26.9%	18.6%
2008	Eliminating Dispersion	16.6%	7.8%
	Eliminating Mean and Variance	14.3%	10.0%

Note: Entries in the table represent the percentage increase in labor productivity from either (a) eliminating the dispersion of distortions across occupations for each group; or (b) eliminating all distortions (setting their mean and variance to zero for each group).

reducing misallocation in both years. These results suggest that the declining variance — and therefore misallocation — plays a central role in our productivity gains. This insight is not something that can easily be discerned from a simple back-of-the-envelope calculation.

5.3. Robustness

How sensitive is the growth contribution of changing τ 's to our parameter choices? Many parameter estimates are needed to calculate our productivity gains. Tables 5 and 6 explore robustness to different parameter values. For each set of parameter values considered, we recalculate the τ_{ig} , A_i , and ϕ_i values so that the model continues to fit the occupation shares and wage gaps.

The first row checks sensitivity to the elasticity of substitution (σ) between occupations in production. In the τ^h case, the share of growth explained ranges from 19.7 percent when the occupations are almost Leontief ($\sigma = .01$) to 21.0 percent when they are almost perfect substitutes ($\sigma = 20$). This compares to 20.4 percent with our baseline value of $\sigma = 3$. Outcomes are more sensitive in the τ^w case, with the share of growth explained by changing τ^w 's going from 12.3 to 18.4 percent (vs. 15.9 percent baseline). The gains are increasing in substitutability. Our intuition is that distortions to the total amount of human capital in one occupation versus another are greater with

Table 5: Robustness Results: Percent of Growth Explained in the τ^h case

	Baseline				
	$\sigma = 3$	$\sigma = .01$	$\sigma = 1/2$	$\sigma = 1.5$	$\sigma = 20$
Changing σ	20.4%	19.7%	19.9%	20.2%	21.0%
	3.44	4.16	5.61	8.41	
Changing θ	20.4%	20.7%	21.0%	21.3%	
	$\eta = 1/4$	$\eta = .01$	$\eta = .05$	$\eta = .1$	$\eta = .5$
Changing η	20.4%	20.5%	20.5%	20.5%	20.3%

Note: Entries in the table represent the share of labor productivity growth that is explained by the changing τ^h 's using the chaining approach. Each entry changes one of the parameter values relative to our baseline case. In the "Changing θ " row, the parameter values reported are $\theta(1 - \eta)$.

higher substitutability across occupations. We not only have too few women doctors, for example, but too little total human capital of doctors when women face barriers to the medical profession. This is particularly true when the allocation of talent is being directly distorted as in the τ^w case.

The second row indicates that the gains from changing τ 's rise modestly as $\theta(1 - \eta)$ rises above our baseline value (holding η fixed at 0.25). Recall that our baseline $\theta(1 - \eta)$ of 3.44 was estimated from wage dispersion within occupation-groups controlling for hours worked, potential experience, and education — and making adjustments for AFQT scores and transitory wage movements. This baseline value attributes 75% of wage dispersion within occupation-groups to comparative advantage. Our baseline θ may overstate the degree of comparative advantage, as it imperfectly controls for absolute advantage. We thus entertain higher values of θ that attribute 50%, 75% and 90% of wage dispersion within occupation-groups to absolute advantage. These higher $\theta(1 - \eta)$ values of 4.16, 5.61, and 8.41 attribute the remainder — 50%, 25% and 10% of wage dispersion — to comparative advantage. As shown in the second row of the robustness table, counterfactual gains rise from 20.4 percent to 21.3 percent of growth as θ rises in the τ^h case. In the τ^w case, the percent of growth explained falls from 15.9 percent to 11.2 percent across the θ range.

The insensitivity of our results to θ (and other parameters such as σ) may seem puz-

Table 6: Robustness Results: Percent of Growth Explained in the τ^w case

	Baseline				
	$\sigma = 3$	$\sigma = .01$	$\sigma = 1/2$	$\sigma = 1.5$	$\sigma = 20$
Changing σ	15.9%	12.3%	13.3%	14.7%	18.4%
	3.44	4.16	5.61	8.41	
Changing θ	15.9%	14.6%	12.9%	11.2%	
	$\eta = 1/4$	$\eta = 0$	$\eta = .05$	$\eta = .1$	$\eta = .5$
Changing η	15.9%	13.9%	14.4%	14.8%	17.5%

Note: Entries in the table represent the share of labor productivity growth that is explained by the changing τ^w 's using the chaining approach. Each entry changes one of the parameter values relative to our baseline case. In the "Changing θ " row, the parameter values reported are $\theta(1 - \eta)$.

zling. But note that, as we entertain different values of θ , we simultaneously change the baseline A 's and τ 's to fit observed wages and employment shares for each occupation and group in each year. When the A 's and τ 's are not distributed lognormally, changing their covariance can push the gains up or down. In the τ^h case, the net result is that the productivity gains are actually increasing in θ . In the τ^w case, productivity gains are decreasing in θ .

Table 7 isolates the impact of changing θ holding fixed the baseline paths for all variables. This means we no longer match the observed wages and occupations. But it usefully illustrates that higher θ alone decisively shrinks the gains from changing τ 's. This is true for both the τ^h and τ^w cases. As people are more similar in ability (θ increases), the gains are smaller from reducing occupational frictions holding all else equal. The reason that our results are so robust to changes in θ in Tables 5 and 6 is that when we reestimate the model, all else is not equal.

Returning to Tables 5 and 6, the third row considers different values of the elasticity of human capital with respect to goods invested in human capital (η). In the τ^h case, the gains hardly change as η rises from 0.01 to 0.25.²⁸ The τ^h values adjust as we change η to explain the observed wage gaps, and the productivity gains are tied to the wage

²⁸We must have $\eta > 0$ in the τ^h case as the only source of wage and occupation differences across groups is different human capital investments in this case.

Table 7: Gains When Changing Only the Dispersion of Ability

	Changing θ			
	3.44	4.16	5.61	8.41
τ^h case	20.4%	18.6%	9.5%	8.4%
τ^w case	15.9%	15.1%	8.0%	3.9%

Note: Entries in the table represent the share of labor productivity growth that is explained by the changing τ 's using the chaining approach. Each entry alters only the value of θ relative to our baseline case. The baseline paths of the τ 's and A 's are not recalculated. The parameter values listed are for $\theta(1 - \eta)$. The baseline case features $\theta(1 - \eta) = 3.44$

gaps in the τ^h case. Gains are much more sensitive to η in the τ^w case, rising from 13.9 percent to 17.5 percent as η goes from 0 to 0.5.

Although not shown in the robustness tables, the gains are not at all sensitive to β , the weight placed on time vs. goods in utility. The gains do not change to one decimal point as we move its value from 0.5 to 0.8 around the baseline value of $\beta = 0.693$. The gains are also not sensitive to how we impute wages in the home sector.²⁹

Another robustness check we carry out is to look only at workers between 25 and 35 years old (inclusive). The reason for this counterfactual is that our model is static and the data are inherently dynamic. By focusing on the young, we will include each cohort only once in our analysis, when they first enter the labor market. Using this younger sample, the changing τ 's account for an even higher fraction of growth, 28.7% in the τ^h case and 23.6% in the τ^w case.³⁰ Related, the gains from dismantling barriers that remain in 2008 are smaller when looking only at the young: 8.0% and 7.6% for the τ^h and τ^w cases, vs. 14.3% and 10.3% when looking at all ages. These results suggest that it is the young that respond most to changing frictions over time and that the current young face much lower frictions relative to the current old.

Finally, we test the sensitivity of our findings to the number of occupations consid-

²⁹Specifically, we raise or lower home sector wages in every year by one standard error from the regression of market wages on education across occupation-groups controlling for group dummies — between 15% and 23% depending on the year. We find the gains fluctuate modestly around the baseline values. For example, in the τ^h case the fraction of growth explained rises to 21.9% with low home sector wages and falls to 18.6% with higher home sector wages, versus 20.4% in the baseline.

³⁰Wage growth was notably slower for the young at 1.04% per year, so there was less to explain.

ered. The gains are surprisingly robust. When we look at a broader set of 20 occupations (vs. 67 in our baseline), the gains are 20.1% vs. 20.4% baseline in the τ^h case, and 14.4% vs. 15.4% baseline in the τ^w case. For 1980 onward we can construct a consistent set of 331 more detailed occupations. Looking at the 1980–2008 subperiod, the gains are similar with more finely divided occupations: 21.1% vs. 20.9% baseline in the τ^h case, and 16.8% vs. 15.2% in the τ^w case.

5.4. Further Results

In this subsection, we describe a number of additional insights from the model.

In the Census data, the share of women working in the market rose from 32.9 percent in 1960 to 69.2 percent in 2008. One explanation is that women’s market opportunities rose, say due to declining discrimination or better information. See Jones, Manuelli, and McGrattan (2003), Albanesi and Olivetti (2009a), and Fogli and Veldkamp (2011) for empirical analysis of these hypotheses. As illustrated in Figure 3, the τ ’s fell in market occupations relative to the home sector for women. How much of the rising female labor-force participation rate can be traced to changing τ ’s? Table 8 provides the answer. Of the 36.4 percentage point increase, the changing τ ’s contributed 23.5 or 26.2 percentage points, or around 75 percent of the total increase. According to our model, the remaining 25 percent can be attributed to changes in technology such as changes in the A ’s or ϕ_i ’s. The latter is in the spirit of the work by Greenwood, Seshadri and Yorukoglu (2005) on “engines of liberation”. It is also consistent with studies attributing women’s rising work to changes in the wage structure, such as Jones, Manuelli, and McGrattan (2003) and Fernandez and Wong (2011).

As we report in Table 9, gaps in average years of schooling narrowed from 1960 to 2008 for all three groups vs. white males: by 0.4 years for white women, 1.8 years for black men, and 1.55 years for black women. If the τ ’s for blacks and women fell faster in high schooling occupations, then the changing τ ’s contributed to some of this educational convergence. The table indicates how much. For white women, the changing τ ’s account for the trend and then some (0.6 years, vs. 0.4 in the data). For black men, falling frictions might have narrowed the schooling gap by 0.65 years, about one-third of the convergence in the data. For black women, declining distortions might explain three-quarters (1.17/1.55) of their catch-up in schooling.

Table 8: Female Participation Rates

	τ^h case	τ^w case
<i>Women's LF participation</i>	1960 = 0.329	2008 = 0.692
<i>Change, 1960 – 2008</i>		0.364
Due to changing τ 's	0.235	0.262
(Percent of total)	(72.3%)	(78.7%)

Note: Results are for white women and black women combined. Participation is defined as working in market occupations. Italicized entries in the table are data; non-italicized entries are results from the model.

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

	Actual 1960	Actual 2008	Actual Change	Change vs. WM	Due to τ 's
White men	11.11	13.47	2.35
White women	10.98	13.75	2.77	0.41	0.63
Black men	8.56	12.73	4.17	1.81	0.65
Black women	9.24	13.15	3.90	1.55	1.17

Table 10: Contribution of Each Group to Total Earnings Growth

	1960–1980	1980–2008	1960–2008
<i>τ^h case:</i>			
All groups	19.7%	20.9%	20.4%
White women	11.3%	18.2%	15.3%
Black men	3.3%	0.9%	1.9%
Black women	5.1%	1.9%	3.2%
<i>τ^w case:</i>			
All groups	21.1%	19.2%	20.0%
White women	8.7%	15.4%	12.6%
Black men	3.4%	0.8%	1.9%
Black women	4.7%	1.4%	2.8%

How much of the productivity gains reflect changes in the occupational frictions facing women vs. those facing blacks? Tables 10 provides the answer. The columns presents the productivity gain from setting the τ 's to their levels at the end of each period (1960–1980, 1980–2008, and 1960–2008). The first row does this for all groups, and the next rows do this for white women, black men, and black women, respectively. Take the τ^h case. Three-quarters (15.3 out of 20.4) of the total gains from reduced occupational frictions over the last fifty years can be explained by the changes facing white women. Falling frictions faced by blacks accounted for the remaining one-quarter of the gains. The primary reason for this is that women are a much larger fraction of the population compared to blacks.

The share of gains associated with falling frictions for white women vs. blacks differs across the time periods. Again, consider the τ^h case. Blacks accounted for a larger share of the gains in the 1960s and 1970s than in later decades. From 1960 to 1980, reduced frictions for blacks account for over 40% of the overall gains. From 1980 to 2008, reduced frictions for blacks account for less than 15% of the overall gains. This timing might link the gains for blacks to the Civil Rights movement of the 1960s.

What was the consequence of shifting occupational frictions for the wage growth of different groups? Table 11 tries to answer this question. The first column presents

Table 11: Group Changes in Wages

	Actual Growth	Due to τ^h 's	Due to τ^w 's
White men	77.0 percent	-5.8%	-7.1%
White women	126.3 percent	41.9%	43.0%
Black men	143.0 percent	44.6%	44.3%
Black women	198.1 percent	58.8%	59.5%

the actual growth of real wages for the different groups from 1960 to 2008. Real wages increased by 77 percent for white men, 126 percent for white women, 143 percent for black men, and 198 percent for black women. For brevity, consider the τ^h case. In the absence of the change in occupational frictions, the model says real wages for white men would have been almost 6 percent higher. Put differently, real income of white men declined due to the changing opportunities for blacks and women. At the aggregate level, this loss was swamped by the wage gains for blacks and women. Over 40 percent of the wage growth for these groups was due to the change in occupational frictions. The model explains the remainder of growth as resulting from changes in technology (A 's) and skill requirements (ϕ 's).

Tables 12 looks at the regional dimension of the decline in frictions confronting blacks and women. Here, we assume that workers are immobile across regions. With this assumption, a decline in occupational frictions in the South will increase average wages in the South relative to the North. From 1960 to 2008, wages in the South increased by 10 percent relative to wages in the Northeast. In the τ^h case, about 7 percentage points of this convergence was due to reduced occupational frictions facing blacks and women in the South — with the bulk of the effect due to falling τ 's for blacks.

From 1980 to 2008, we see a reversal of the North-South convergence, perhaps driven by the reverse migration of blacks to the U.S. South. Reverse migration is what one would expect to see if workers are responding to improved labor market outcomes in that region. Persistent wage gaps might then reflect skill differences between regions.

Table 12: Contributions to Northeast - South Convergence

	1960–1980	1980–2008	1960–2008
τ^h case:			
Actual wage convergence	20.7%	-16.5%	10.0%
Due to all τ 's changing	4.9%	1.5%	6.9%
Due to black τ 's changing	3.6%	1.9%	5.6%
τ^w case:			
Actual wage convergence	20.7%	-16.5%	10.0%
Due to all τ 's changing	2.2%	0.4%	2.8%
Due to black τ 's changing	2.3%	1.2%	3.6%

5.5. Average Quality of Workers by Occupation

Using equations (10) and (11), the average quality of workers — including both innate ability and human capital — for group g in occupation i is given by

$$\frac{H_{ig}}{q_g p_{ig}} = \gamma \bar{\eta} \cdot \frac{1}{(1 - \tau_{ig}^w) w_i} \cdot (1 - s_i)^{-1/\beta} \cdot \left(\sum_{s=1}^N \tilde{w}_{sg}^\theta \right)^{\frac{1}{\theta} \cdot \frac{1}{1-\eta}} \quad (28)$$

Average quality for a group *relative* to white men is therefore

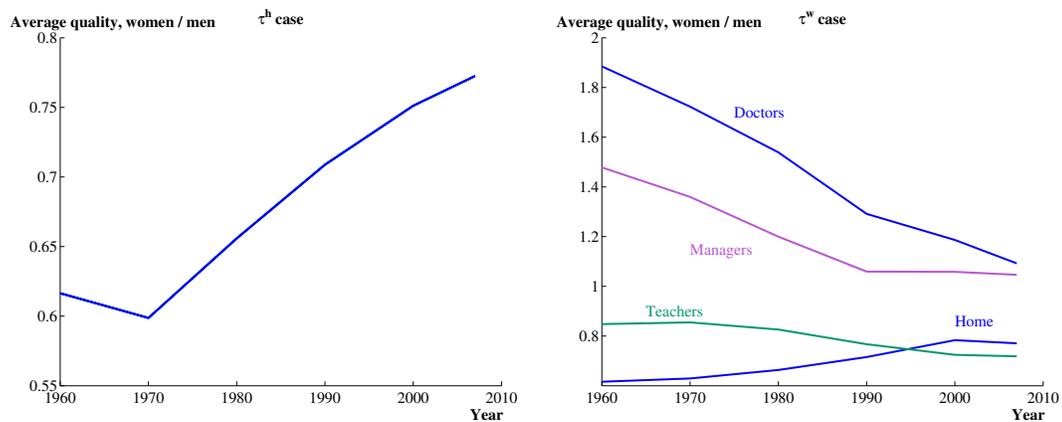
$$\frac{H_{ig}/q_g p_{ig}}{H_{i,wm}/q_{wm} p_{i,wm}} = \frac{1 - \tau_{i,wm}^w}{1 - \tau_{ig}^w} \cdot \frac{\overline{\text{wage}}_g}{\overline{\text{wage}}_{wm}}. \quad (29)$$

Relative quality in an occupation is simply the wage gap adjusted by the τ^w frictions.

In the τ^h case (where the τ^w variables are set to zero), equation (29) implies that average quality for a group relative to white men is *the same across all occupations*. In particular, relative quality is precisely equal to the wage gap. When the labor market frictions are introduced, this changes. In this case, wages are not equal to marginal products, so that average quality differs from wages. More specifically, wages are less than marginal products, so average quality is higher when the frictions are larger.

One way to think about these quality differences is to consider the following question: if you were to see a doctor chosen at random in 1960 for a fixed fee, would you

Figure 7: Relative Average Quality, White Women vs. White Men



Note: The panels show relative average quality (human capital and innate ability) in various occupations for white women versus white men, in the τ^h and τ^w cases. Computed using equation (29).

rather see a male doctor or a female doctor? Figure 7 shows the ratio of average quality for white women vs. white men for several occupations, as in equation (29). In the τ^h case, relative qualities are equated in all occupations. Because the wage gap reflects quality, the average female doctor has less human capital than the average male doctor. Over time, as the wage gaps have declined, the relative quality of women in each occupation rose substantially between 1960 and 2008, from 0.56 to 0.77.

The τ^w case presents a very different view of the data, as shown in the right panel. The relative quality of women is higher than their wages suggest because they are paid less than their marginal products. In 1960, average quality was substantially higher for women vs. men doctors and managers. Only the most talented women overcame frictions to become doctors and managers in 1960, and some lesser talented white men entered these professions instead. In this case, the difference in quality has faded substantially over time due to declining frictions, but remains present even in 2008.

The real world presumably has elements of both the τ^h and the τ^w cases. To this end, independent information on quality trends for occupation-groups could help us separate and identify human capital and labor market frictions.

6. Conclusion

How does discrimination in the labor market and in the acquisition of human capital affect occupational choice? And what are the consequences of the resulting 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 and discrimination in the acquisition of human capital. Second, we impose the assumption that the distribution of an individual's ability over all possible occupations 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 2008. 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 15 to 20 percent of aggregate wage growth, 75 percent of the rise in women's labor force participation, and essentially all of the wage convergence between women and blacks and white men.

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 could be a poor approximation of the U.S. labor market. It would also be useful to quantify the extent to which the barriers are due to labor market discrimination versus barriers to the acquisition of human capital. Independent data on quality trends would be useful to distinguish between these two forces. Finally, we have focused on the gains from reducing barriers facing women and blacks over the last fifty years. But we suspect that barriers facing children from less affluent families and regions have worsened in the last few decades. If so, this could explain both the adverse trends in aggregate productivity and the fortunes of less-skilled Americans in recent decades. We hope to tackle some of these questions in future work.³¹

³¹One could also investigate groups in other countries. Hnatkovska, Lahiri and Paul (2011) look at castes in India and find narrowing differences in education, occupations, and wages in recent decades.

References

- Albanesi, Stefania and Claudia Olivetti, "Gender Roles and Medical Progress," NBER Working Papers 14873, National Bureau of Economic Research, Inc April 2009.
- and —, "Home Production, Market Production and the Gender Wage Gap: Incentives and Expectations," *Review of Economic Dynamics*, 2009, 12 (1), 80–107.
- 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.
- Beaudry, Paul and Ethan Lewis, "Do Male-Female Wage Differentials Reflect Differences in the Return to Skill? Cross-City Evidence From 1980-2000," Technical Report, UBC unpublished paper 2012.
- Becker, Gary, *The Economics of Discrimination*, Chicago, IL: The University of Chicago Press, 1957.
- Bertrand, Marianne, Claudia Goldin, and Lawrence F. Katz, "Dynamics of the Gender Gap for Young Professionals in the Financial and Corporate Sectors," *American Economic Journal: Applied Economics*, July 2010, 2 (3), 228–55.
- 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.
- , 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," Working Paper 17993, National Bureau of Economic Research April 2012.
- 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.
- Cavalcanti, Tiago and Jos Tavares, "The Output Cost of Gender Discrimination: A Model-Based Macroeconomic Estimate," CEPR Discussion Papers 6477, C.E.P.R. Discussion Papers September 2007.

- Charles, Kerwin Kofi and Jonathan Guryan, "Prejudice and Wages: An Empirical Assessment of Becker's The Economics of Discrimination," *Journal of Political Economy*, October 2008, 116 (5), 773–809.
- Chay, Kenneth Y., Jonathan Guryan, and Bhashkar Mazumder, "Birth Cohort and the Black-White Achievement Gap: The Roles of Access and Health Soon After Birth," NBER Working Papers 15078, National Bureau of Economic Research, Inc June 2009.
- Cunha, Flavio, James J. Heckman, and Susanne M. Schennach, "Estimating the Technology of Cognitive and Noncognitive Skill Formation," *Econometrica*, 2010, 78 (3), 883–931.
- Donohue, John J. and James Heckman, "Continuous versus Episodic Change: The Impact of Civil Rights Policy on the Economic Status of Blacks," *Journal of Economic Literature*, December 1991, 29 (4), 1603–43.
- Dupuy, Arnaud, "An Economic Model of the Evolution of the Gender Performance Ratio in Individual Sports," *International Journal of Performance Analysis in Sport*, 2012, 12 (1), 224–245.
- Eaton, Jonathan and Samuel S. Kortum, "Technology, Geography, and Trade," *Econometrica*, 2002, 70 (5), 1741–1779.
- Ellison, Glenn and Ashley Swanson, "The Gender Gap in Secondary School Mathematics at High Achievement Levels: Evidence from the American Mathematics Competitions," *Journal of Economic Perspectives*, 2010, 24 (2), 109–128.
- Erosa, Andres, Tatyana Koreshkova, and Diego Restuccia, "How Important Is Human Capital? A Quantitative Theory Assessment of World Income Inequality," *Review of Economic Studies*, October 2010, 77 (4), 1421–1449.
- Fernandez, Raquel, "Cultural Change as Learning: The Evolution of Female Labor Force Participation over a Century," *American Economic Review*, forthcoming 2012.
- , 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.
- and Joyce Chen Wong, "The Disappearing Gender Gap: The Impact of Divorce, Wages, and Preferences on Education Choices and Women's Work," 2011. unpublished paper.
- Fogli, Alessandra and Laura Veldkamp, "Nature or Nurture? Learning and the Geography of Female Labor Force Participation," *Econometrica*, July 2011, 79 (4), 1103–1138.

- 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,” *Review of Economic Studies*, January 2005, 72 (1), 109–133.
- Güvenen, Fatih and Burhanettin Kuruscu, “A Quantitative Analysis of the Evolution of the U.S. Wage Distribution, 1970–2000,” in Daron Acemoglu, Kenneth Rogoff, and Michael Woodford, eds., *NBER Macroeconomics Annual*, Cambridge, MA: MIT Press, 2009, pp. 227–276.
- Hnatkovska, Viktoria, Amartya Lahiri, and Sourabh B. Paul, “Breaking the Caste Barrier: Inter-generational Mobility in India,” 2011. U.B.C. manuscript.
- Hsieh, Chang-Tai and Peter J. Klenow, “Misallocation and Manufacturing TFP in China and India,” *Quarterly Journal of Economics*, 2009, 124 (4), 1403–1448.
- Jones, Larry E., Rodolfo E. Manuelli, , and Ellen R. McGrattan, “Why Are Married Women Working So Much?,” Staff Report 317, Federal Reserve Bank of Minneapolis June 2003.
- Karabel, Jerome, *The Chosen: the Secret History of Admission and Exclusion at Harvard, Yale, and Princeton*, New York, NY: Houghton Mifflin Harcourt, 2005.
- Luttmer, Erzo G.J., “On the Mechanics of Firm Growth,” Working Papers 657, Federal Reserve Bank of Minneapolis 2008.
- Manuelli, Rodolfo and Ananth Seshadri, “Human Capital and the Wealth of Nations,” March 2005. University of Wisconsin working paper.
- McFadden, Daniel, “Conditional Logit Analysis of Qualitative Choice Behavior,” in P. Zarembka, ed., *Frontiers of Econometrics*, New York, NY: Academic Press, 1974, pp. 105–142.
- Pitt, Mark M., Mark R. Rosenzweig, and Nazmul Hassan, “Human Capital Investment and the Gender Division of Labor in a Brawn-Based Economy,” *American Economic Review*, December 2012, 102 (7), 3531–3560.
- Ramondo, Natalia and Andres Rodriguez-Clare, “Trade, Multinational Production, and the Gains from Openness,” *Journal of Political Economy*, forthcoming 2013.

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.

Rodgers, William and William Spriggs, "What Does the AFQT Really Measure: Race, Wages, Schooling and the AFQT Score," *Review of Black Political Economy*, Spring 1996, pp. 13-46.

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.