

THE ALLOCATION OF TALENT AND U.S. ECONOMIC GROWTH

CHANG-TAI HSIEH

Booth School of Business, University of Chicago and NBER

ERIK HURST

Booth School of Business, University of Chicago and NBER

CHARLES I. JONES

Graduate School of Business, Stanford University and NBER

PETER J. KLENOW

Department of Economics, Stanford University and NBER

In 1960, 94 percent of doctors and lawyers were white men. By 2010, the fraction was just 62 percent. Similar changes in other highly-skilled occupations have occurred throughout the U.S. economy during the last 50 years. Given that the innate talent for these professions is unlikely to have changed differently across groups, the change in the occupational distribution since 1960 suggests that a substantial pool of innately talented women and black men in 1960 were not pursuing their comparative advantage. We examine the effect on aggregate productivity of the convergence in the occupational distribution between 1960 and 2010 through the prism of a Roy model. Across our various specifications, between 20% and 40% of growth in aggregate market output per person can be explained by the improved allocation of talent.

KEYWORDS: Economic growth, discrimination, misallocation, Roy model.

1. INTRODUCTION

THE LAST 50 YEARS HAVE SEEN A REMARKABLE CONVERGENCE in the occupational distribution between white men, women, and black men. For example, 94 percent of doctors and lawyers in 1960 were white men. By 2010, the fraction was just over 60 percent. Similar changes occurred throughout the economy, particularly in highly-skilled occupations.¹ Yet no formal study has assessed the effect of these changes on aggregate economic performance. Since the innate talent for a profession among members of a group is unlikely to change over time, the change in the occupational distribution since 1960 suggests that a substantial pool of innately talented women and black men in 1960 were not pursuing their comparative advantage. The resulting (mis)allocation of talent could potentially have important aggregate consequences.

Chang-Tai Hsieh: Chang-Tai.Hsieh@chicagobooth.edu

Erik Hurst: Erik.Hurst@chicagobooth.edu

Charles I. Jones: chad.jones@stanford.edu

Peter J. Klenow: klenow@stanford.edu

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

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

This paper measures the aggregate effects of the changing allocation of talent from 1960 to 2010. We examine labor market outcomes for race and gender groups through the prism of a Roy (1951) model of occupational choice. Within the model, every person is born with a range of talents or preferences across occupations. Each individual chooses the occupation where she obtains the highest utility given her talents and preferences.

We introduce three forces that will cause individuals to choose occupations where they do not have a comparative advantage. First, we allow for discrimination in the labor market. Consider the world that Supreme Court Justice Sandra Day O'Connor faced when she graduated from Stanford Law School in 1952. Despite being ranked third in her class, the only job she could get in 1952 was as a legal secretary (Biskupic (2006)). We model labor market discrimination as an occupation-specific wedge between wages and marginal products. This "tax" is a proxy for many common formulations of discrimination in the literature.²

Second, the misallocation of talent can be due to barriers to forming human capital. We model these barriers as increased monetary costs associated with accumulating occupation-specific human capital. These costs are a proxy for parental and teacher discrimination in favor of boys in the development of certain skills, historical restrictions on the admission of women to colleges or training programs, differences in school quality between black and white neighborhoods, and differences in parental wealth and schooling that alter the cost of investing in their children's human capital.³

Finally, we allow for differences in preferences or social norms to drive occupation differences across groups. For example, there might have been strong social norms against women and black men working in high-skilled occupations in the 1960s. This possibility has been highlighted in the work of, among others, Johnson and Stafford (1998), Altonji and Blank (1999), and Bertrand (2011). We treat the home sector as additional occupation. As a result, we also allow for differences across groups in the extent to which they want to work in the home sector. This factor can capture changes in social norms related to women working at home. However, we can interpret the change in the preference for the home sector over time broadly so that it also includes changes in the preference for children or the ability to control the timing of fertility.⁴

To measure these three forces, we make a key assumption that the distribution of innate talent of women and black men—relative to white men—is constant over time. With this assumption, we back out the *changes* in labor market frictions, human capital frictions, and occupational preferences from synthetic panel data on the occupations and wages of women and black men relative to white men from 1960 to 2010. We infer that preferences changed and/or labor and human capital frictions declined from 1960 to 2010 to jointly explain the convergence in occupations and wages of women and black men relative to

²See Becker (1957), Phelps (1972), and Arrow (1973), and a summary in Altonji and Blank (1999).

³Karabel (2005) documented that Harvard, Princeton, and Yale systematically discriminated against blacks, women, and Jews in admissions until the late 1960s. Card and Krueger (1992) showed that public schools for blacks in the U.S. South in the 1950s were underfunded relative to schools for white children. Goldin and Katz (2002), Bailey (2006), and Bailey, Hershbein, and Miller (2012) documented that innovations related to contraception had important consequences for female labor market outcomes and educational attainment. Neal and Johnson (1996) documented differences in AFQT scores across race and how controlling for AFQT explains a portion of black-white gaps. Akcigit, Grigsby, and Nicholas (2017) highlighted how parental liquidity constraints can affect investments in their children's education.

⁴See Fernández, Fogli, and Olivetti (2004) and Fernández (2013) on the role of cultural forces, Greenwood, Seshadri, and Yorukoglu (2005) on the role of home durables, and Goldin and Katz (2002) on the role of birth control in explaining changes in female labor supply over time. Surveys of this extensive literature can be found in Costa (2000) and Blau, Ferber, and Winkler (2013).

white men. When we view these facts through the lens of our general equilibrium model, we find that these shifts account for roughly two fifths of growth in U.S. market GDP per person between 1960 and 2010. They also account for most of the rise in labor force participation over the last five decades.

We use the model to decompose the contribution of each force. In our base specification, individuals draw a vector of idiosyncratic productivities across occupations. With this assumption, wage differences across groups within an occupation discipline our estimates of changing group preferences. If women did not like being lawyers in 1960, the model says women must have been paid more to compensate for this disamenity. Second, we use the life-cycle structure of the model to distinguish between barriers to human capital investment and labor market discrimination. In our setup, human capital barriers affect an individual's choice of human capital prior to entering the labor market. The effect of these barriers remains with a cohort throughout their life-cycle. In contrast, labor market discrimination affects all cohorts within a given time period. We then use the evolution of life-cycle wages across groups to distinguish occupation-specific human capital barriers (akin to "cohort" effects) from occupation-specific labor market discrimination (akin to "time" effects).

We find that declining obstacles to accumulating human capital were more important than declining labor market discrimination: the former explains 36 percent of growth in U.S. GDP per person between 1960 and 2010, while the latter explains 8 percent of growth. Changing group-specific occupational preferences explain little of U.S. growth during this time period.

Our main findings are robust to having workers draw a vector of occupation-specific preferences instead of productivities. Even if individuals sort only on preferences, we find that one-fifth of growth in market GDP per person over the last five decades can be traced to declining occupational barriers. A key reason is that women and black men are moving into high-skilled occupations over time. When individuals have occupation-specific abilities, this reallocation represents a better allocation of talent. When workers have the same ability in all occupations and choose occupations based on idiosyncratic preferences, the movement of women and black men into high-skilled occupations increases the average market return to their ability.

To recap, this paper makes a conceptual point and an empirical point. Conceptually, we show that quantities (occupational shares) are more robustly related to group-specific occupational frictions than are wage gaps. Empirically, we demonstrate that there could be substantial gains in GDP as a result of declining occupational barriers facing women and black men. Both our empirical and conceptual points hold as long as individuals sort at least partially on ability.

The rest of the paper proceeds as follows. Section 2 presents the model. Section 3 discusses data and inference for our baseline in which individuals differ in occupational productivities. Section 4 presents the main results for this setting. Section 5 explores robustness when individuals sort based on preferences or on both preferences and productivities. Section 6 discusses other robustness checks. Section 7 concludes.

2. MODEL

The economy consists of a continuum of workers, each in one of M discrete sectors, one of which is the home sector. Workers are indexed by occupation i , group g (such as race and gender), and cohort c . A worker possesses heterogeneous abilities ϵ_i or preferences μ_i over occupations. Some people are better teachers while others derive more utility from working as a teacher.

2.1. *Workers*

In a standard Roy (1951) model, workers are endowed with idiosyncratic talent ϵ in each occupation. We add two additional forces to this setup. First, we assume workers are heterogeneous in either their talent or their preferences over occupations, but not both; heterogeneity on both dimensions hinders tractability. Second, we allow for forces that distort the allocation of workers across occupations. We have in mind forces such as discrimination in the labor market, barriers to human capital accumulation, and group-specific social norms.

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

Lifetime utility of a worker from group g and cohort c who chooses occupation i is a function of lifetime consumption, time spent accumulating human capital, and occupational preferences:

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

$C(c, t)$ is consumption of cohort c in year t , s denotes time allocated to human capital acquisition in the pre-period, z_{ig} is the common utility benefit of all members of group g from working in occupation i , μ is the idiosyncratic utility benefit of the *individual* from the occupation, and β parameterizes the trade-off between lifetime consumption and time spent accumulating human capital.⁵ We normalize the time endowment in the pre-period to 1, so $1 - s$ is leisure time in the pre-period. Changes in social norms for women working in the market sector or changing preferences for fertility can be thought of as changes in z in the home sector for women. The idiosyncratic preference of a specific woman in an occupation is represented by μ .

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

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

\bar{h}_{ig} captures permanent differences in human capital endowments and γ parameterizes the return to experience. We assume γ is only a function of age $= t - c$ and \bar{h}_{ig} is fixed for a given group-occupation. \bar{h}_{ig} reflects any differences in talent common to a group in a given occupation. ϕ_i is the occupation-specific return to time investments in human capital, while η is the elasticity of human capital with respect to human capital expenditures.

Consumption equals “after-tax” earnings net of expenditures on education:

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

Net earnings are the product of $1 - \tau_{ig}^w$ and a person’s efficiency units of labor, which in turn is the product of the price per efficiency unit w_i , the worker’s idiosyncratic talent in their chosen occupation ϵ , and their human capital h . Individuals borrow $e(c)(1 + \tau_{ig}^h(c))$

⁵In the first period of cohort c , $t = c$. We omit subscripts on other individual-specific variables for ease of notation, but z_{ig} has subscripts to emphasize that it varies across groups and occupations.

in the first period to purchase $e(c)$ units of human capital, a loan they repay over their lifetime subject to the lifetime budget constraint $e(c) = \sum_{t=c}^{c+2} e(c, t)$.

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

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

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

$$\begin{aligned}
 s_i^* &= \frac{1}{1 + \frac{1 - \eta}{3\beta\phi_i}}, \\
 e_{ig}^* &= \left(\frac{\eta(1 - \tau_{ig}^w)w_i\bar{\gamma}h_{ig}\bar{s}_i^{\phi_i}\epsilon}{1 + \tau_{ig}^h} \right)^{\frac{1}{1-\eta}},
 \end{aligned}
 \tag{3}$$

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

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

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

where

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

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

and

$$\tilde{z}_{ig} \equiv z_{ig}^{\frac{1-\eta}{3\beta}}.$$

The effect of labor market discrimination and human capital barriers is summarized by the “composite” τ_{ig} . More human capital barriers or labor market discrimination increase τ_{ig} , which lowers indirect utility for an individual from group g when choosing occupation i . Group-specific disutility from working in occupation i is represented as a low value of \tilde{z}_{ig} . We represent group-specific preferences by \tilde{z} instead of z to make the units of group preferences comparable to those of τ . Higher innate talent ϵ or preferences μ also increases the rewards for choosing an occupation.

Finally, turning to the distribution of the idiosyncratic talent ϵ and preferences μ , we borrow from [McFadden \(1974\)](#) and [Eaton and Kortum \(2002\)](#). Each person gets either a skill draw ϵ_i or a preference draw μ_i in each of the M occupations. To be clear, if a worker gets a skill draw, we assume that $\mu = 1$ for the worker. If the person gets a skill draw, talent in each occupation is drawn from a multivariate Fréchet distribution:

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

The parameter θ governs the dispersion of skills, with a higher value of θ corresponding to *smaller* dispersion. We normalize the mean parameter of the skill distribution to 1 in all occupations for all groups, but this mean parameter is isomorphic to \bar{h}_{ig} .

If the individual instead gets a preference draw, these preferences are also drawn from a multivariate Fréchet distribution, where the shape parameter for the Fréchet distribution of preferences is equal to $\frac{\theta(1-\eta)}{3\beta}$. This assumption makes the elasticity of labor supply to an occupation of individuals with heterogeneous preferences the same as that of workers with ability heterogeneity. We assume the ability of workers who sort on preferences is the same in all occupations and given by $\epsilon_i = \Gamma^{1-\eta}$ where $\Gamma \equiv \Gamma(1 - \frac{1}{\theta(1-\eta)})$ is the Gamma function. This assumption makes average ability the same for the two groups of workers.

2.2. Occupational Choice

Given the above assumptions, the occupational choice problem thus reduces to picking the occupation that delivers the highest value of U_{ig}^* . Because heterogeneity 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 people⁶:

⁶Proofs are in the Supplemental Material (Hsieh, Hurst, Jones, and Klenow (2019)).

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

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

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

Occupational sorting depends on \tilde{w}_{ig} , which is the overall reward that someone from group g with the mean talent obtains by working in occupation i , relative to the power mean of \tilde{w} for the group over all occupations. The occupational distribution is driven by relative returns and not absolute returns: forces that change \tilde{w} for all occupations have no effect on the occupational distribution. Occupations where the wage per efficiency unit w_i is high will attract more workers of *all* groups. In contrast, differences between groups in occupational choice are driven by differences in \bar{z} , \bar{h}_{ig} , τ^w , and τ^h . The fraction of group g who choose occupation i is low when the group dislikes the occupation (\bar{z}_{ig} is low), has low ability in the occupation (\bar{h}_{ig} is low), is discriminated against (τ_{ig}^w is high), or faces a barrier in accumulating human capital (τ_{ig}^h is high).

We view home production as simply another occupation, so the share of a group in the home sector is also given by equation (4). The labor force participation rate therefore depends on the return in the home sector relative to the market. For example, the decline in the labor force participation rate of white men since the 1960s can be driven by higher returns in the home sector (such as better video games), a decline in labor market opportunities (such as the decline of blue-collar jobs), or changing preferences for the market sector relative to the home sector. The increase in female labor force participation rates from 1960 to 2010 can be due to less labor market and human capital discrimination in market occupations.

2.3. Worker Quality

For individuals with heterogeneous abilities, sorting affects the average quality of workers in an occupation. For individuals with heterogeneous preferences, sorting has no effect on the average quality of workers in an occupation. But sorting on productivities or preferences has different effects on occupational wages. Average worker quality in an occupation is therefore a weighted average of the quality of workers who sort on ability and those who sort on preferences:

PROPOSITION 2—Average Quality of Workers: *For a given cohort c of group g at time t , the geometric average of worker quality in each occupation, including both human capital and talent, is*

$$e^{\bar{\epsilon} \log[h_{ig}(c,t)\epsilon_{ig}(c)]} = \bar{s}_i(c)^{\phi_i(t)} \gamma(t - c) \left(\frac{\eta s_i(c)^{\phi_i(c)} \bar{\gamma} \bar{h}_{ig} w_i(c) [1 - \tau_{ig}^w(c)]}{1 + \tau_{ig}^h(c)} \right)^{\frac{\eta}{1-\eta}} \left(\frac{1}{p_{ig}(c)} \right)^{\frac{1-\delta}{\theta(1-\eta)}}. \tag{5}$$

The parameter δ denotes the share of the population with idiosyncratic preferences (so $1 - \delta$ is the share of workers with idiosyncratic ability) and $\bar{\Gamma}$ is a constant.⁷ By varying δ , we can explore the robustness of our results to sorting that occurs completely on talent ($\delta = 0$), sorting that occurs completely on preferences ($\delta = 1$), or sorting that occurs on both margins. When all individuals possess heterogeneous abilities ($\delta = 0$), average quality is inversely related to the share of the group working in the occupation $p_{ig}(c)$. This captures the selection effect. For example, the model predicts that if the labor market discriminated against female lawyers in 1960, only the most talented female lawyers would have chosen to work in this occupation. And if the barriers faced by female lawyers declined after 1960, less talented female lawyers would move into the legal profession and thus lower the average quality of female lawyers. Conversely, in 1960, the average quality of white male lawyers would have been lower in the presence of labor market discrimination against women and black men. At the other extreme, when $\delta = 1$ (all workers sort on preferences), this selection effect is absent.

2.4. Occupational Wages

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

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

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

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

For individuals in the young cohort, $t = c$, which implies $\frac{s_i(c)^{\phi_i(t)}}{s_i(c)^{\phi_i(c)}} = 1$ and $\frac{(1-\tau_{ig}^w(t))w_i(t)}{(1-\tau_{ig}^w(c))w_i(c)} = 1$. When all individuals sort on ability ($\delta = 0$), average earnings for a given group among the young differ across occupations only because of differences in s_i and \tilde{z}_{ig} . Occupations in which schooling is especially productive (a high ϕ_i and therefore a high s_i) will have higher average earnings. Occupations where individuals have a strong common disutility from being in the profession (\tilde{z}_{ig} is small) have higher wages as compensation for the lower utility. These are the only two forces that generate differences in wages across occupations for the young when individuals sort completely on talent ($\delta = 0$). Average earnings are no higher in occupations where a group faces less discrimination in the labor market, lower frictions in human capital attainment, a higher wage per efficiency unit, or where the group has more talent in the sector. The reason is that each of these factors leads lower quality workers to enter those jobs. This composition effect exactly offsets the direct effect on earnings when the distribution of talent is Fréchet.

The composition effect would not be present if selection was driven by forces other than occupational ability. When all workers select based on idiosyncratic preferences ($\delta = 1$),

⁷ $\bar{\Gamma}$ is defined in equation (A7) in the Supplemental Material (Hsieh et al. (2019)).

selection affects the average *utility* of workers in an occupation, but has no effect on average ability. In this case, there is no quality offset due to selection and the average wage in an occupation varies with p_{ig} —so the average wage and the occupational share will both be higher in occupations where a group faces less discrimination or where the wage per efficiency unit is higher.

The general point is that the wage gap is not a robust measure of the frictions faced by a group in a given occupation. The elasticity of the wage gap with respect to occupational frictions depends on the share of individuals who sort on preferences instead of ability. When individuals sort entirely on ability ($\delta = 0$), the wage gap is uncorrelated with these frictions because of the offsetting effect of selection.

Equation (6) for the average wage also identifies the forces behind wage changes over a cohort’s life-cycle. For a given cohort-group in an occupation, s_i , \tilde{z}_{ig} , and p_{ig} are fixed. Therefore, the average wage increases over time when the price of skills in the occupation w_i increases, labor market discrimination τ^w falls, return to experience is positive, or the return to schooling increases.

2.5. Relative Propensities

Putting together the equations for the occupational shares and wages in each occupation, and assuming the experience profiles are the same across groups, we get the relative propensity of a group to work in an occupation:

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

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

where the subscript “wm” denotes white men.

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

2.6. Relative Labor Force Participation

The labor force participation rate of a group relative to white men is given by equation (7). We normalize $\tilde{z} = 1$, $\tau^w = 0$, and $\tau^h = 0$ for the home sector. With these normalizations, the labor force participation rate relative to white men is given by the following:

PROPOSITION 5—Relative Labor Force Participation: *Let $LFP_g \equiv 1 - p_{\text{home},g}$ denote the share of group g in the market occupations. The share of group g in the home sector relative*

to white men is then

$$\frac{1 - \text{LFP}_g(c)}{1 - \text{LFP}_{\text{wm}}(c)} = \frac{m_{\text{wm}}(c)}{m_g(c)} = \left(\frac{\overline{\text{wage}}_{ig}(c, c)}{\overline{\text{wage}}_{i,\text{wm}}(c, c)} \right)^{-\theta(1-\eta)} \left(\frac{\tilde{z}_{ig}(c)}{\tilde{z}_{i,\text{wm}}(c)} \right)^{-\theta} \left(\frac{p_{ig}(c)}{p_{i,\text{wm}}(c)} \right)^\delta \quad \forall i \in \text{market}, \quad (8)$$

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

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

The second equation in (8) says that the relative return to market work is given by a power function of the gap in market wages in *any* market sector, the relative occupational preference term in that sector, and the relative occupational propensity in the sector with an elasticity that depends on the share of people that sort on preferences δ . We will use this insight to back out \tilde{z} in the market sectors from data on labor force participation of the group and wage gaps, both relative to white men.

2.7. Firms

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

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

where $H_i \equiv \sum_g \sum_c q_g(c) p_{ig}(c) \mathbb{E}[h_{ig}(c) \epsilon_{ig}(c)]$ denotes total efficiency units of labor in occupation i , A_i is the exogenously-given productivity of occupation i , and σ is the elasticity of substitution across occupations in aggregate production.

2.8. Equilibrium

Sections B and C of the Supplemental Material (Hsieh et al. (2019)) contain the remaining details of the model. Section B endogenizes τ^w and τ^h as a function of the discriminatory preferences of firm owners. Section C defines the general equilibrium of the model and contains a proposition describing how the equilibrium allocation and prices can be solved for.

2.9. Intuition

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

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

In the case with no selection on ability ($\delta = 0$), aggregate output is given by

$$Y = \underbrace{q_m \cdot \left(\sum_{i=1}^M A_i^\theta \right)^{\frac{1}{\theta} \cdot \frac{1}{1-\eta}}}_{\text{output from men}} + \underbrace{\frac{q_w}{1 - \bar{\tau}^w} \cdot \left(\sum_{i=1}^M \left(\frac{A_i(1 - \tau_i^w) \tilde{z}_i}{(1 + \tau_i^h)^\eta} \right)^\theta \right)^{\frac{1}{\theta} \cdot \frac{1}{1-\eta}}}_{\text{output from women}}, \tag{10}$$

where q_w and q_m denote the number of women and men, and $\bar{\tau}^w$ denotes the earnings-weighted average of the labor market friction facing women.⁸ The first term in (10) is aggregate output produced by men and is not affected by the occupational distortions facing women because occupations are perfect substitutes here. The second term is aggregate output produced by female labor. The effect of τ^w , τ^h , and \tilde{z} on aggregate output shows up in the second term; it is increasing in the number of people in the discriminated group q_w . Also note that the effect of \tilde{z} on aggregate output is isomorphic to the effect of τ^w and τ^h . Societal preferences shift the allocation of talent in exactly the same way as labor market and human capital distortions.

We illustrate how this setup can be used to gain intuition by focusing on τ^w ; the effects of τ^h or \tilde{z} can be analyzed in a similar fashion. Assuming $\tau^h = 0$, $\tilde{z} = 1$, and that τ^w and A are jointly log-normally distributed, aggregate output produced by women Y_w (the second term in equation (10)) is given by

$$\ln Y_w = \ln q_w + \ln \left(\sum_{i=1}^M A_i^\theta \right)^{\frac{1}{\theta} \cdot \frac{1}{1-\eta}} + \frac{\eta}{1 - \eta} \cdot \ln(1 - \bar{\tau}^w) - \frac{1}{2} \cdot \frac{\theta - 1}{1 - \eta} \cdot \text{Var} \ln(1 - \tau_i^w). \tag{11}$$

τ^w affects output via the last two terms in equation (11). The mean of τ^w changes the return to investment in human capital. This effect is captured by the third term in equation (11) and its magnitude depends on elasticity of output with respect to human capital η . The dispersion of τ^w across occupations affects aggregate output via a different channel. Here, dispersion of τ^w affects the allocation of female labor across occupations. A decline in the dispersion of τ^w improves the allocation, which increases aggregate output. This effect is captured by the fourth term in equation (11).

Equation (11) suggests that the effect of unequal barriers on aggregate output is increasing in θ . While this is true for a given variance of labor distortions, our inference about the magnitude of that variance from observed data also depends on θ . Using the equation for relative propensities, the variance in the labor distortion is given by⁹

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

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

⁸ $\bar{\tau}^w \equiv \sum_{i=1}^M \omega_i \tau_i^w$, where $\omega_i \equiv \frac{p_{iw} \overline{\text{wage}}_w}{1 - \tau_i^w} / \sum_{j=1}^M \frac{p_{jw} \overline{\text{wage}}_w}{1 - \tau_j^w}$.

⁹ We maintain the assumption that τ^w is the only source of variation.

output from female labor is

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

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

Finally, now consider the case in which individuals only sort on preferences ($\delta = 1$). The effect of dispersion in τ on average utility is exactly the same as in the case when individuals only sort on ability, but the effect on aggregate output is different.¹⁰ Specifically, maintaining the same assumptions behind equation (11), the log of aggregate output of women is given by¹¹

$$\ln Y_w = \ln q_w + \ln \frac{\sum_{i=1}^M A_i^{\theta+1}}{\sum_{j=1}^M A_j^\theta} + \theta \text{Cov}[\ln A_i, \ln(1 - \tau_i^w)]. \quad (12)$$

Now aggregate output increases in the *covariance* of A_i and $1 - \tau_i^w$. Aggregate output falls when a group is under-represented in high A_i occupations. Similarly, a change in labor market frictions can raise aggregate output if τ^w declines in high A_i occupations.

3. INFERENCE WITH SELECTION ONLY ON ABILITY ($\delta = 0$)

We now explain how we identify the driving forces of our model given data on wages and occupational shares for different groups of workers. We begin by considering the case where all individuals only draw occupational talent ($\delta = 0$). In Section 5, we discuss inference when selection occurs on both ability and preferences ($\delta > 0$). In that section, we also provide an estimate of δ . The results with our estimated δ are not far from our results with $\delta = 0$, so we use $\delta = 0$ as our benchmark specification.

The inference exercise—for any value of δ —is based on two key assumptions. First, we assume the relative mean latent occupational talent of a group relative to white men $\bar{h}_{ig}/\bar{h}_{i,wm}$ is constant over time. This is a key assumption and we cannot proceed without it. It implies that the *change* in the occupational distribution of women and black men relative to white men since 1960 must be driven by changes in labor market or human capital frictions or by changes in common occupational preferences. Second, we assume that idiosyncratic occupational abilities or preferences are distributed i.i.d. Fréchet. This assumption is not as crucial, but it buys us enormous tractability because it leads to simple

¹⁰ Average utility of women is given by $(\sum_{i=1}^M (\frac{A_i(1-\tau_i^w)}{(1+\tau_i^h)^\eta})^\theta)^{\frac{1}{\theta}} \cdot \frac{1}{1-\eta}$ in the two cases.

¹¹ For simplicity, equation (12) also assumes $\eta = 0$.

expressions for occupational shares and wages as a function of the occupational frictions. Relaxing this assumption is a valuable direction for future research but would not change the fact that the τ 's and \tilde{z} 's must have changed since the 1960s to explain the observed changes in the distribution of occupations of white women and blacks relative to white men over the last fifty years.¹²

This section proceeds as follows. First, we describe the data. Second, we explain what in the data allows us to measure the composite friction τ and the group-specific occupational preferences \tilde{z} . Third, we show how to decompose the composite τ into labor market frictions (τ^w) and human capital barriers (τ^h). Finally, we explain how we infer productivity and the return to skill in each occupation.

3.1. Data

We use data from the 1960, 1970, . . . 2000 decennial Censuses and the 2010–2012 American Community Surveys (ACS).¹³ We restrict the sample to four groups: white men, white women, black men, and black women. We include individuals between the ages of 25 and 54. This restriction focuses the analysis on individuals after they finish schooling and prior to retirement. Finally, we exclude individuals who report being unemployed (not working but searching for work) or on active military duty.¹⁴

We create pseudo-panels by following synthetic cohorts from 1960 to 2010. We define three ages within a cohort's life-cycle: young (age 25–34), middle (age 35–44), and old (age 45–54). For example, a synthetic cohort would be the young in 1960, the middle aged in 1970, and the old in 1980. We have information on eight cohorts for the timespan we study. We observe information at all three life-cycle points for four cohorts (the young in 1960, 1970, 1980, and 1990) and one or two life-cycle points for the remaining cohorts.

We define a person as either in the home sector or in the market sector based on their number of hours worked. We classify a person as being in the home sector if she is not currently employed or works less than ten hours per week. Individuals working more than thirty hours per week are classified as employed in one of 66 market occupations.¹⁵ Those who are employed but usually work between ten and thirty hours per week are classified as part-time workers. We split the sampling weight of part-time workers equally between the home sector and the reported market occupation.

We measure earnings as the sum of labor, business, and farm income in the previous year. When calculating earnings, we restrict the sample to individuals who are currently working, who worked at least 48 weeks during the prior year, and who earned at least 1000 dollars (in 2010 dollars) in the previous year. We convert all earnings data from the Census to constant dollars. Our measure of wage gaps across occupations and groups is the difference in the log of the geometric average of earnings.¹⁶

¹²Lagakos and Waugh (2013) and Adão (2016) estimated selection models with arbitrary correlation but with only two or three sectors. We do not know how to do something similar for the 67 occupations we have.

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

¹⁴The Supplemental Material (Hsieh et al. (2019)) reports summary statistics. For all analysis, we apply survey sample weights.

¹⁵The 66 market occupations (shown in Supplemental Material (Hsieh et al. (2019)) Table FII) are based on occupation codes in the 1990 Census. We chose the 1990 codes as they are available in Census and ACS years since 1960.

¹⁶Our results are robust to adjusting for hours worked across groups. This is not surprising given we already condition on full-time work status. When computing average earnings by occupation, we include top-coded and imputed data; excluding such data had little effect on our estimated τ 's.




FIGURE 1.—Standard deviation of relative occupational shares. Note: The figure shows the standard deviation of $\ln\left(\frac{p_{ig}}{p_{wm}}\right)$ across occupations for each group weighting each occupation by the share of earnings in that occupation. Specifically, we show the data for young white women (middle line), young black men (bottom line), and young black women (top line) relative to young white men.

3.2. Composite Frictions versus Occupational Preferences

Equation (4) says that differences in occupational choice between women and black men relative to white men are driven by differences in the ratio of occupational preferences to occupational frictions \tilde{z}/τ . Figure 1 plots the standard deviation of the shares of women and black men relative to white men across market occupations for the young cohort in each decade. The sorting of women and blacks has converged toward that of white men over time. Viewed through the lens of (4), this fact indicates that the τ and/or \tilde{z} of women and black men must have converged toward that of men. This is one key fact behind our finding that the allocation of talent has improved over the last five decades.

Equation (6) says that wage gaps across occupations for the young are proportional to gaps in $\tilde{z}^{\frac{1}{1-\eta}}$. For example, if white women are poorly compensated (relative to white men) as lawyers compared to secretaries, it must be the case that women receive higher utility from working as secretaries compared to lawyers. When $\delta = 0$, for a given estimate of η we can infer relative \tilde{z} across groups by fitting the occupational wage gaps across groups and occupations for the young.

Equation (7) then says that, conditional on having an estimate of the parameters θ and η , the composite friction τ can be recovered from data on relative occupational shares after controlling for the average wage gap. Intuitively, the wage gap controls for the effect of preferences on occupational choice. The “residual” occupational choice is therefore only driven by the effect of τ . Our base results normalize $\bar{h}_{ig}/\bar{h}_{i,wm} = 1$ and assume the occupational choice of white men is undistorted (i.e., $\tau_{i,wm} = 1$).¹⁷ So when the share of some group in an occupation is low relative to white men after we control for the wage gap, we infer that the group faces a high τ^w or a high τ^h in the occupation.

We need estimates of θ and η to infer τ 's and \tilde{z} 's from the data. To estimate θ , we use the fact that distributional assumptions imply that wages within an occupation for a given group follow a Fréchet distribution with the shape parameter $\theta(1 - \eta)$. This reflects both comparative advantage (governed by $1/\theta$) and amplification from endogenous human capital accumulation (governed by $1/(1 - \eta)$). Using micro data from the U.S. Population

¹⁷We will later show robustness to these two normalizations.

Census/ACS, we estimate $\theta(1 - \eta)$ to fit the distribution of the residuals from a cross-sectional regression of log hourly wages on $66 \times 4 \times 3$ occupation-group-age dummies in each year.¹⁸ The resulting estimates for $\theta(1 - \eta)$ range from a low of 1.24 in 1980 to a high of 1.42 in 2000, and average 1.36 across years.¹⁹

The parameter η denotes the elasticity of human capital with respect to education spending and is equal to the fraction of output spent on human capital accumulation. Spending on education (public plus private) as a share of GDP in the United States averaged 6.6 percent over the years 1995, 2000, 2005, and 2010.²⁰ Since the labor share in the United States in the same four years was 0.641, this implies an η of 0.103.²¹ With our base estimate of $\theta(1 - \eta) = 1.36$, $\eta = 0.103$ gives us $\theta = 1.52$.

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

With these values for θ and η , we can now infer τ and \bar{z} from data on occupational propensities and wage gaps. Figure 2 shows the mean of τ of each group across the 67 occupations. For white women, the mean of τ fell from about 7 in 1960 to around 3 in 2010, with most of the decline occurring prior to 1990. Average τ facing black women declined from around 8 to about 3 from 1960 through 2010. Black men experienced a decline in mean τ from around 3 to 1.5 during the five decades. For both black women and black men, most of the decline occurred between 1960 and 1980.




FIGURE 2.—Mean of composite occupational frictions. Note: Figure shows earnings-weighted mean of τ for each group.

¹⁸We use MLE, taking into account the number of observations which are top-coded in each year.

¹⁹Sampling error is minimal because there are 300–400k observations per year for 1960 and 1970 and 2–3 million observations per year from 1980 onward. We did not use 2010 data because top-coded wage thresholds differed by state in that year.

²⁰See <http://www.oecd.org/education/eag2013.htm>.

²¹Labor share data are from <https://research.stlouisfed.org/fred2/series/LABSHPUA156NRUG>. The young's share of earnings is from the U.S. Population Census/ACS.




FIGURE 3.—Variance of composite frictions and occupational preferences. Note: Figure shows the earnings-weighted variance of $\ln \tau$ (left panel) and $\ln \tilde{z}$ (right panel).

Figure 3 shows the dispersion of $\ln \tau$ (left panel) and $\ln \tilde{z}$ (right panel) across all 67 occupations. For all three groups, the variance of $\ln \tau$ fell by about 0.4 log points between 1960 and 2010. The right panel shows that the decline in the dispersion of $\ln \tilde{z}$ is much smaller than the decline in the dispersion of τ . For black men and white women, there is essentially no change in the dispersion of occupational preferences relative to white men so almost all of the occupational convergence is due to τ . So for black men and white women, almost all of the convergence in occupational propensities is due to the convergence in τ . For black women, the variance of relative $\ln \tilde{z}$ fell, but the magnitude of the decline is only 17% of the decline in the dispersion of $\ln \tau$. So even for black women, most of the occupational convergence is due to τ convergence.

Figure 4 displays τ for white women for a subset of occupations. The composite friction was high for women in 1960 working in construction, as lawyers, and as doctors relative to working as teachers and secretaries. For white women lawyers and doctors, τ in 1960 was around 10. If τ reflected labor market discrimination only, the implication would be that women lawyers in 1960 were paid only one-tenth of their marginal product relative to their male counterparts. The model infers large τ 's for white women in these occupations in 1960 because there were few white women doctors and lawyers in 1960, even after controlling for the gap in wages. Conversely, a white woman in 1960 was 24 times more




FIGURE 4.—Occupational barriers (τ_{ig}) for white women. Note: Author's calculations based on equation (7) using Census data and imposing $\theta = 2$ and $\eta = 0.103$.

likely to work as a secretary than was a white man. The model explains this huge gap by assigning a τ below 1 for white women secretaries.

Over time, τ of white women in the lawyer and doctor professions fell. By 2010, white women faced composite frictions below 2 in the lawyer, doctor, and teacher occupations. The barrier facing white women in the construction sector remained large. This fact could be the result of women having a comparative disadvantage (relative to men) as construction workers, a possibility we consider later in our robustness checks.

3.3. Labor Market versus Human Capital Discrimination

The occupational frictions shown in Figures 3 and 4 are a composite of labor market discrimination (τ^w) and human capital barriers (τ^h). We now show how we distinguish between these two forces by exploiting life-cycle variation. The key assumption is that individuals make an active choice to obtain human capital prior to entering the labor market. This assumption implies that human capital discrimination is akin to a cohort effect, whereas labor market discrimination affects all cohorts in the labor market at the same point in time and thus is like a time effect.

The wage gap of cohort c and group g (relative to white men) in occupation i at time t relative to the wage gap at time c (when cohort c was young) is

$$\frac{\text{gap}_{ig}(c, t)}{\text{gap}_{ig}(c, c)} \propto \frac{1 - \tau_{ig}^w(t)}{1 - \tau_{ig}^w(c)}. \tag{13}$$

The change in the wage gap over the life-cycle depends on the change in τ^w over time. If labor market discrimination diminishes over time, this raises the average wage (relative to white men) in occupations where the group previously faced discrimination. We therefore use the change in the wage gap over a cohort-group’s life-cycle to infer the change in τ^w over time. We then use $\tau_{ig} \equiv (1 + \tau_{ig}^h)^\eta / (1 - \tau_{ig}^w)$ to infer the change in τ^h from the change in τ after controlling for the change in τ^w . Intuitively, the change in τ^h is calculated as the difference in the wage gap of the young between successive cohorts after controlling for the slope of the life-cycle wage gap for a given cohort.

Figure 5 shows the wage gap of white women (left panel) and black men (right panel) vis à vis white men data for different cohorts over their life-cycle. A decline in τ^w in a given year steepens a given life-cycle profile. As τ^w falls, the wage of a given group relative to white men converges during an individual’s life-cycle. On the other hand, a decline in either τ^w or τ^h will shift up the *intercept* of the life-cycle wage gap profiles. Figure 5 shows clearly the large increases in the intercept of the wage gap of the young white women across successive cohorts. However, there are only small changes in the slopes of the cohort profiles over time, which suggests that most of the shift of the intercept is due to a decline in human capital barriers. For black men, however, we see both shifts in the intercepts and steepening slopes particularly during the 1960s to 1980s, suggesting a role for both declining human capital and labor market frictions.²²

²²We weight equally wage growth from young to middle age and from middle age to old to infer the change in τ^w . We also need to normalize the initial split of the composite τ between τ^w and τ^h . For our baseline, we assume a split of 50/50 in 1960. In subsequent years, we let the data speak to the importance of τ^h relative to τ^w . Finally, we place an additional constraint on the τ breakdown to keep aggregate “revenue” from changing by more than 10 percent of GDP over our sample period. This requires τ^h to be no lower than 0.8, to keep subsidies for women secretaries from getting too large.

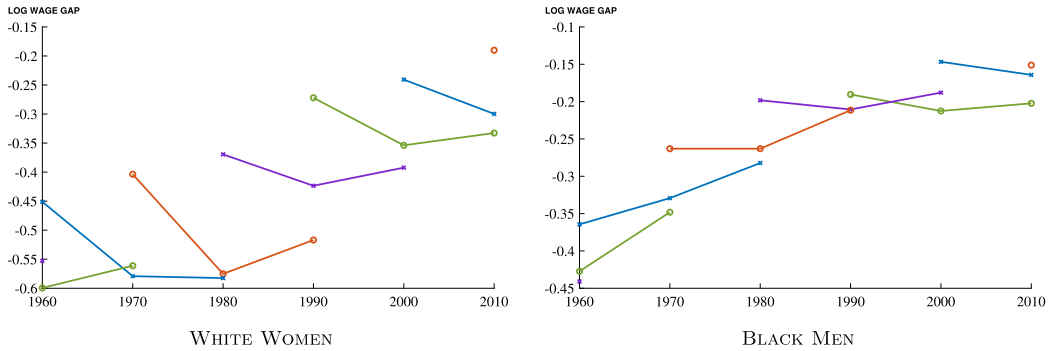


FIGURE 5.—Wage gaps relative to white men by time and cohort. Note: Log wage gaps are shown for the life-cycle of each cohort by connected line segments for young, middle-aged, and old periods.

3.4. Home Sector, Technology, and Return to Skill Parameters

We now show how we pin down the parameters that determine the labor force participation rate. Remember that the home sector is simply another sector, so the labor force participation rate is simply determined by the returns in the market sector relative to the returns in the home sector. We assume the home sector is undistorted in that τ^h and τ^w in the home sector are zero for all groups. This implies that distortions in the market sector lower the labor force participation rate. We also normalize the common home sector preference term (\tilde{z}_{home}) to 1 for all groups. This implies that the \tilde{z}_i 's we estimate for the market occupations are relative to the home sector. So the labor force participation rate is decreasing in τ^w , τ^h , \tilde{z} and in the ratio of w_i and ϕ_i in the home sector relative to the market sectors. And the *gap* in labor force participation rate relative to white men is decreasing in τ and \tilde{z} (w_i and ϕ_i have the same effect on labor force participation for all groups).

In Figure 3, we showed the dispersion of \tilde{z} imputed from the dispersion of the wage gap across occupations, but not the level of \tilde{z} . We now use equation (5) to infer the level of \tilde{z} (relative to white men) from the gap in the labor force participation rate after conditioning on the wage gap. Intuitively, the wage gap captures the effect of τ on labor force participation, so the residual has to be driven by preferences for market work relative to the home sector. Figure 6 shows that mean \tilde{z} for white women was about 0.5 in 1960. Intuitively, the gap in labor force participation rates between white women and men in 1960 was larger than can be explained by the wage gap, so we infer that white women did not like working in the market sector relative to the home sector in 1960. Over time, to match the fact that female labor force participation rates rose relative to those of white men, the model infers that white women's preference for working in the market relative to the home sector must have increased. For black men in 1960, the gap in labor force participation rates relative to white can be entirely explained by the wage gap, so average \tilde{z} is about 1. Over time, the labor force participation rate of black men fell relative to white men. The model “explains” this fact as a result of a decline in the mean \tilde{z} of black men in the market sectors relative to white men.

The last thing are the parameters that determine the *level* of the labor force participation rates of white men, which are w_i and ϕ_i in all sectors (including the home sector) and \tilde{z} in the market sectors for white men. We pick ϕ_i in each year to match data on school-




FIGURE 6.—Mean of occupational preferences (relative to white men). Note: Figure shows earnings-weighted mean of \tilde{z} for each group relative to white men.

ing differences for young white men across occupations in each year.²³ Then, conditional on estimates of ϕ_i , we pick \tilde{z}_i in each occupation to fit the average wage by occupation. Then, given \tilde{z}_i and ϕ_i , we pick w_i to exactly fit occupational shares for young white men. Occupations with a large share of young white men in a given year are ones where the price of skills w_i is high. With estimates of w_i , we then back out the technology parameter A_i .²⁴

3.5. Recap and Model Fit

Table I summarizes the identifying assumptions and normalizations for our base parameterization of the model wherein individuals only draw occupational talent ($\delta = 0$).

TABLE I
IDENTIFYING ASSUMPTIONS AND NORMALIZATIONS

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

²³The first-order conditions for schooling (equation (3)) says that s_i is a function of ϕ_i and the parameters η and β . We assume the pre-market period is 25 years long so that $s_i = \frac{\text{Years of Education}}{25}$. We already have an estimate of η , so all we need is β . The average wage of group g in occupation i is proportional to $(1 - s_i)^{\frac{1}{\beta}}$, so the Mincerian return $\psi \pm 1$ year around mean schooling \bar{s} satisfies $e^{2\psi} = (\frac{1-\bar{s}+0.04}{1-\bar{s}-0.04})^{\frac{1}{\beta}}$ so $\beta = \ln(\frac{1-\bar{s}+0.04}{1-\bar{s}-0.04}) / (6\psi)$. Since the average Mincerian return from a cross-sectional regression of log income on years of schooling (with group dummies) averages 12.7% in our data, this gives us $\beta = 0.231$.

²⁴We need the elasticity of substitution among occupations (in aggregating to final output) σ to infer A_i from w_i . We choose $\sigma = 3$ as our baseline value, but we have no information on this parameter. Given this, we explore the robustness of our results to alternate values of σ .

TABLE II
BASELINE PARAMETER VALUES

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

To reiterate, we assume τ^w and τ^h are zero for white men in all occupations. We cannot identify the level of τ^w and τ^h , only their levels relative to a given group. We assume that relative innate talent levels are the same across all groups and are normalized to 1 ($\bar{h}_{ig} = 1$). We also assume τ^w and τ^h in the home sector are zero for all groups. We normalize preferences in the home sector to be 1 for all groups. Again, we cannot identify the level of preferences, only their level relative to the home sector.

Table II summarizes the key parameters and Table III the endogenous variables and the target data for their indirect inference. Some forcing variables depend on cohorts and some on time, but never both. Variables changing by cohort include the human capital barriers (τ^h), common group-specific occupational preferences (z), and the elasticity of human capital with respect to time investment (ϕ). Labor market barriers (τ^w) and technology parameters (A) vary over time. Human capital barriers, labor market discrimination, and occupational preferences vary across occupation-groups.

Finally, Table IV compares the data and the model's predictions for aggregate earnings per worker and labor force participation by year. Remember that the model only targets the occupational shares and labor force participation rates of the young. Despite this, predicted per-capita earnings and labor force participation rates in the model are not very far from the data. For example, in 2010, predicted earnings in the model are within 2 percent of the actual earnings in the data. In the model, labor force participation rate increases by 15 percentage points between 1960 and 2010. The actual increase between 1960 and 2010 is 16 percentage points.

4. RESULTS WITH SELECTION ONLY ON ABILITY ($\delta = 0$)

Given the discussion of inference above, we can now answer the key question of the paper: how much of the overall growth from 1960 to 2010 can be explained by the changing

TABLE III
FORCING VARIABLES AND EMPIRICAL TARGETS WHEN $\delta = 0^a$

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

^aThe variable values are chosen jointly to match the empirical targets. $\delta = 0$ refers to the polar case where individuals draw idiosyncratic ability for each occupation (but not idiosyncratic tastes).

TABLE IV
MODEL VERSUS DATA: EARNINGS AND LABOR FORCE PARTICIPATION^a

Year	Earnings Data	Earnings Model	LFP Data	LFP Model
1960	18,383	18,615	0.599	0.599
1970	24,645	25,000	0.636	0.614
1980	27,088	27,900	0.702	0.653
1990	33,953	34,265	0.764	0.720
2000	39,419	41,134	0.747	0.743
2010	41,541	42,717	0.759	0.748

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

labor market outcomes of women and black men during this time period? In this section, we explore this question assuming that all individuals only draw occupation-specific talent. In the next section, we explore how the results change if individuals also draw occupational preferences.

Real earnings per person in our census sample grew by 1.8 percent per year between 1960 and 2010. According to our model, this observed earnings growth can come from five sources. First is general occupational productivity growth (changing A 's). Second is growth in the returns to schooling, which results in more human capital attainment (changing ϕ 's). Third, changing preferences can reallocate labor across occupations and generate earnings growth (changing \bar{z} 's). Fourth, growth in the relative share of each group in the working age population can also mechanically change earnings per capita (changing q 's). Finally, as described in Section 2.9, changing gender- and race-specific barriers to occupational choice can result in economic growth (changing τ 's).

To assess how much the changing τ 's contributed to economic growth, we hold the τ 's fixed while allowing the A 's, ϕ 's, \bar{z} 's, and q 's to evolve. We then use the difference between the actual path in the data and the counterfactual “no change in τ 's” path to measure the contribution of changing τ 's.

4.1. Income and Productivity Gains

The results of our baseline counterfactual are shown in the first column of Table V. The changes in τ 's account for 41.5% of growth from 1960 to 2010 in market GDP per person (row 1) and 38.4% in market earnings per person (row 2). Market earnings and market GDP differ due to changing “revenue” from labor market discrimination over time.²⁵ The decline in market revenue from the declining τ 's results in the market earnings growth of workers being slightly larger than market GDP growth over the sample.

A portion of the growth in both market GDP per person and market earnings per person reflects rising labor force participation of women in response to falling frictions. Aggregate labor force participation rates rose steadily in the data, from 60% in 1960 to 76% in 2010, primarily due to increased female labor supply. Changes in the τ 's account for 90% of this increase, as seen in row 3 of Table V. Declining τ 's also contributed to the growth in market GDP per person by raising the average wages of those who work in the

²⁵See the Supplemental Material (Hsieh et al. (2019)) for the time series trends in the model's implied revenue from τ^w and τ^h .

TABLE V
SHARE OF GROWTH DUE TO CHANGING FRICTIONS (ALL AGES)^a

	Share of growth accounted for by			
	τ^h and τ^w	τ^h, τ^w, \bar{z}	τ_h only	τ_w only
Market GDP per person	41.5%	40.8%	36.0%	7.7%
Market earnings per person	38.4%	37.5%	18.9%	26.0%
Labor force participation	90.4%	112.7%	24.9%	56.2%
Market GDP per worker	24.0%	15.0%	40.0%	-9.8%
Home + market GDP per person	32.7%	32.1%	30.6%	4.4%

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

market. As seen in row 4 of the table, declining τ 's account for 24% of the increase in market GDP *per worker*. Declining labor market frictions allowed women and black men to better exploit their comparative advantage reducing misallocation in the economy. Given that the occupations with the highest τ 's in 1960 were more likely to be high-skilled occupations, the declining τ 's resulted in women and black men accumulating more human capital which also contributed to aggregate growth in market output per worker.

The final row of column 1 of Table V shows that the declining τ 's explain about one-third of the growth in total GDP (inclusive of home sector output) during the last fifty years. The reduction in labor market discrimination and barriers to human capital growth drew more women into the market sector, which had a direct effect of *raising* market GDP per person and simultaneously *lowering* home sector GDP per person. On net, however, declining labor market frictions for women and black men substantially increased the sum of market and home output per person.

Figure 7 shows the time series decomposition of growth in market GDP per person coming from the changing τ 's. The top line in the figure shows growth in market GDP per person implied by the model. The bottom line is the counterfactual growth in market GDP if the τ 's were held fixed. Not surprisingly, the productivity effect of the τ 's has grown over time. Additionally, our results suggest that productivity growth would have




FIGURE 7.—GDP per person, data and model counterfactual. Note: The graph shows the cumulative growth in GDP per person (market), in the data (overall), and in the model with no changes in τ 's as in Table V.

been close to zero during the 1970s had it not been for the reduction in labor market barriers to women and black men during that time period.

Column 2 of Table V assesses how much of growth can be explained by declining labor market and human capital frictions (the τ 's) and changing common occupational preferences (the \bar{z} 's). Comparing the first and second columns, it can be seen that changing occupational preferences has only a modest effect on productivity growth. If anything, changing occupational preferences actually had a slightly negative effect on market GDP per person. The combined effect of the τ 's and \bar{z} 's explained 40.8% of the growth in market GDP per person since 1960, whereas the τ 's alone explained 41.5%. Changing preferences do explain a small amount of the change in aggregate labor force participation, with most of the effect being driven by women. The results in column 2 also imply that the majority of the growth in market GDP per person over the last half-century was due to changes in the A_i 's and the ϕ_i 's. These forces are not group-specific and explained 59% of growth between 1960 and 2010.

Why can't changing preferences for market work explain the growth in market GDP per person? If women simply did not like some occupations in 1960, the model with only sorting on talent says they would have been paid more in occupations in which they were under-represented. The data show no such patterns. The gender (wage) gap was no lower in skilled occupations, and it did not fall faster in skilled occupations as the share of women rose. So while preference changes did result in the reallocation of women and black men across occupations and did explain some of the rise in their labor force participation, it did not generate substantial economic growth.

The last two columns of Table V report growth contributions separately for falling barriers to human capital accumulation (τ^h) and falling labor market discrimination (τ^w). Falling human capital barriers alone would have accounted for 36% of growth in market GDP per person, versus 8% from falling labor market discrimination. Ebbing labor market discrimination looms larger for growth in market earnings (26% of growth). The reason is that declining discrimination in the labor market contributes directly to earnings growth *relative* to output growth. When we look at growth in home plus market GDP, declining barriers to human capital are again more important (31% of growth) than is diminishing labor market discrimination (4%). These results are consistent with Figure 5, where most of the changes in wage gaps for white women show up as a cohort effect rather than a time effect, suggesting a more important role for τ^h .

Table V suggests falling labor market discrimination drove much (over 56%) of the rise in labor force participation. Falling barriers to human capital accumulation played a lesser role since human capital is also useful in the home sector, albeit less so than in some market occupations. The breakdown into contributions from human capital versus labor market barriers is also revealing for why the contribution to growth in market GDP per worker is smaller than the contribution to growth in market GDP per person (24% vs. 42%). Falling human capital barriers, on their own, would have explained 40% of the growth in market GDP per worker. But falling labor market discrimination actually lowered growth in market GDP per worker (-10%) by enticing workers with marginal talent to move out of the home sector and into market occupations.²⁶

Table VI shows how the changing τ 's (column 1) and combined τ 's and \bar{z} 's (column 2) affect earnings and wage gaps across groups. The third column shows the explanatory

²⁶In four of the five rows in Table V, the combined effect of changing τ^h and τ^w is smaller than the sum of the effects from eliminating them individually. The explanation for this is that misallocation is convex in barriers. Reducing one of the barriers individually yields the largest gains to be had by moving highly misallocated workers to the right occupation.

TABLE VI
WAGE GAPS AND EARNINGS BY GROUP AND CHANGING FRICTIONS^a

	Share of growth accounted for by		Full Model
	τ^h and τ^w	τ^h, τ^w, \bar{z}	
Earnings, WM	-12.2%	-17.4%	91.4%
Earnings, WW	76.9%	86.3%	107.0%
Earnings, BM	28.7%	20.4%	100.5%
Earnings, BW	50.6%	52.5%	101.7%
Wage gap, WW	148.1%	98.3%	176.7%
Wage gap, BM	98.0%	115.4%	143.4%
Wage gap, BW	84.6%	71.4%	131.4%
LF Participation	90.4%	112.7%	93.6%

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

effect of our full model (τ 's, \bar{z} 's, A 's, and ϕ 's) on group-specific earnings and wage gaps. A few things are of note from Table VI. First, our model collectively does very well in explaining the average earnings growth of all groups between 1960 and 2010. The model explains about 100% of earnings growth for black men and black women while only slightly overestimating earnings growth for white women and underestimating earnings growth for white men. Second, falling labor market frictions account for 77% of earnings growth for white women, 29% for black men, and 51% for black women. The declining τ 's, particularly for women, were the primary source of earnings growth over the last half-century. Third, Table VI highlights that the changing τ 's actually lowered wage growth of white men. This is because falling barriers to women and black men in high-skilled occupations caused white men to shift to lower wage occupations. For men (both black and white), wage growth was driven primarily by changes in technology and skill requirements (A 's and ϕ 's). Finally, the changing \bar{z} 's again had only modest positive effects on the earnings growth of women and modest negative effects on the earnings growth on men. As women entered the labor force due to changing preferences, this increased their market earnings and reduced the earnings of men.

Our model concludes that most of the change in wage gaps between groups and white men can be explained by falling τ 's. Our model actually over-predicts the changing wage gaps for all groups. This is because the model slightly under-predicts the earnings growth of white men. With that in mind, declining τ 's explain 148% of the declining wage gap of white women, while the model in total explains 177%. Other than this, our model does fairly well in predicting the changing wage gaps over time. Our model, collectively, over-predicts slightly the rising wages of women and black men relative to white men during the 1960–2010 period. For women, the changing τ 's more than explain the shrinking gender gap in wages observed in the data. Declining barriers to human capital attainment and declining labor market discrimination were primarily responsible for the declining gender and racial wage gaps during the last fifty years.

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

TABLE VII
SHARE OF GROWTH IN MARKET GDP PER PERSON DUE TO DIFFERENT GROUPS^a

1960–2010	τ^h and τ^w	τ^h only	τ^w only
All groups	41.5%	36.0%	7.7%
White women	33.8%	29.8%	6.1%
Black men	1.2%	0.7%	0.6%
Black women	3.7%	3.2%	0.9%

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

contributed much more to growth than did declining labor market barriers. However, as can be seen in Figure 5, declining human capital barriers and labor market discrimination were roughly equally as important for black men, while for white women it was the decline in human capital barriers that was primarily important.

Finally, we can ask: how much additional growth could be achieved by reducing frictions (τ 's) all the way to zero? If the remaining frictions in 2010 were removed entirely, we calculate that GDP today would be 9.9% higher. These remaining gains result from the fact that, even in 2010, there are still some differences in occupational choice and average wages across groups. However, through the lens of our model, there are only modest potential gains in GDP from reducing the τ 's for women and black men fully to zero. Most of the large productivity gains from the occupational convergence across groups occurred between 1970 and 2000. This is one reason to be less optimistic about U.S. economic growth after 2010 compared to growth in the last half-century.

4.2. Model Gains versus Back-of-the-Envelope Gains

Our baseline estimate in Table V suggests that τ^w and τ^h account for 42% of the gains in market GDP per person and 33% of the gains in total GDP per person. Is this number large or small relative to what one might have expected? We have two ways of thinking about this question. First, in the log-normal approximation to the model with only τ^w variation that we presented back in Section 2.9, the elasticity of GDP to 1 minus the mean of τ^w is $q_w \cdot \frac{\eta}{1-\eta}$. If we assume that the share of women in the population q_w is 1/2 and $\eta = 0.1$, then this elasticity is $\frac{1}{2} \cdot \frac{1}{9}$. Figure 2 showed that the mean of the composite τ of women fell from about 10 in 1960 to 3 in 2010. This decline in $\bar{\tau}$ can thus account for a 7% increase in total GDP per person.²⁷ Figure 3 shows that $\text{Var} \ln \bar{\tau}$ fell from about 0.9 to 0.6 from 1960 to 2010. In the log-normal approximation to the model, the semi-elasticity of GDP to $\text{Var} \ln \tau$ is $q_w \cdot \frac{1}{2} \cdot \frac{\theta-1}{1-\eta} \approx 0.3$.²⁸ A 0.3 decrease in the variance of $\ln \tau$ thus could explain an 8% increase in total GDP per person. Thus, according to this back-of-the-envelope calculation, changing τ 's for women boosted GDP about 15%. A similar calculation for black men suggests that changing τ for black men boosted GDP by about 2%.²⁹ The overall increase of GDP per person in our setup was about 138%, so the changing τ explains $0.15/1.38 \approx 12.3\%$ of growth in total GDP. This is significantly lower

²⁷ $\frac{1}{2} \cdot \frac{1}{9} \cdot \ln(10/3) \approx 0.07$.

²⁸ $\frac{1}{2} \cdot \frac{1}{2} \cdot \frac{2-1}{1-0.1} \approx 0.28$.

²⁹ Black men 8% of the population, mean τ going from 3 to 1.5, and the variance of $\ln \tau$ falling by 0.3.

than the 33% contribution we estimate without imposing any parametric assumption on the distribution of the barriers. Clearly, log-normal is not a reasonable assumption. Also, the back-of-the-envelope calculation assumes no general equilibrium effects on wages of white men from changes in the labor supply of women and black men across occupations.

A second way to answer the question is to compare our model-based contribution in market earnings per person growth to what one would infer from the falling gaps in earnings per person for women and black men relative to white men. The narrowing gaps in earnings per person—including both declining wage gaps and rising labor force participation—mechanically account for 37% of growth in market per person.³⁰ Coincidentally, the model-based estimate of market earnings per person growth is nearly the same as this naive back-of-the-envelope calculation. The back-of-the-envelope calculation assumes no general equilibrium effect of falling frictions on the earnings growth of white men. Yet we reported above that white men's wages fell 11% relative to what they would have done without the changing barriers facing women and black men (see Table VI). Moreover, this back-of-the-envelope calculation assumes that earnings gaps would not have changed in the absence of falling frictions. That is, this calculation implicitly attributes the entire decline in earnings gaps to changing frictions.

As we also show in Table VI, changes in occupational productivity and returns to schooling also contribute to wage gaps between groups over time. It just so happens that these effects roughly cancel out such that our model predictions and the back-of-the-envelope predictions are approximately equal. Still, our back-of-the-envelope calculations suggest that our model-based estimates of the growth consequences from falling labor market barriers to women and black men are not implausibly large.

4.3. Robustness to Alternate Values of θ , η , and σ

Table VIII explores the robustness of our productivity gains to different parameter values. In each case, we recalculate the τ , \tilde{z} , A , and ϕ values so that the model continues to fit the occupation shares, wage gaps, etc. by group and year. The first row of Table VIII replicates the gains from changing the τ 's in explaining the growth in market GDP per

TABLE VIII
ROBUSTNESS TO PARAMETER VALUES^a

	Growth in GDP per person accounted for by		
	τ^h and τ^w	τ^h alone	τ^w alone
Benchmark	41.5%	36.0%	7.7%
$\theta = 1.5$	45.3%	41.2%	6.8%
$\theta = 4$	32.2%	25.2%	7.0%
$\eta = 0.05$	39.8%	33.8%	8.2%
$\eta = 0.20$	44.4%	39.4%	7.2%
$\sigma = 1.05$	22.2%	19.8%	3.6%
$\sigma = 10$	42.9%	36.5%	8.5%

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

³⁰For this calculation, we held fixed earnings per person relative to white men at 1960 levels, and found only 63% as much growth in earnings per person as seen in the data.

person under baseline parameter values for comparison. The next row considers a lower value for the Fréchet shape parameter θ , which is inversely related to the dispersion of comparative advantage across occupations. With $\theta = 1.5$ rather than the baseline $\theta = 2$, changing barriers explain modestly more of growth in market GDP per person (45.3%) than in the baseline (41.5%).

Recall that our baseline θ was estimated from wage dispersion within occupation-groups. This might overstate the degree of comparative advantage because some of the wage variation is due to *absolute* advantage. We thus entertain a much higher value ($\theta = 4$) than in our baseline ($\theta = 2$). With this higher θ , the share of growth from changing τ 's falls to 32.2% (vs. 41.5% in the baseline). Less discrimination is needed to explain occupational choices when comparative advantage is weak. Even with this higher value of θ , however, declining τ 's explain about one-third of growth in market GDP per person over the last half-century.

Table VIII also varies η , the elasticity of human capital with respect to goods invested in human capital. The gains rise slightly from 40% with $\eta = 0.05$, to 42% with our baseline $\eta = 0.103$, to 44% with $\eta = 0.20$.

The last rows of Table VIII show the sensitivity of the results to the elasticity of substitution σ between occupations in production. When the elasticity of substitution across occupations is higher, the declining τ 's explain a higher portion of the growth in market GDP per person. The gains in market GDP per person to changing τ 's when $\sigma = 1.05$ (close to Cobb–Douglas) are 22.2%. While it may appear our results are quite sensitive to changes in σ , it should be noted that changing σ simply reallocates how much of the growth occurs in the market sector versus home sector. When $\sigma = 1.05$, the declining τ 's explain 33.5% of the growth in total GDP per person (inclusive of the home sector). This is nearly identical to our base results in Table V. However, when $\sigma = 1.05$, the declining τ 's explain only 70% of the rise in labor force participation. With a lower elasticity of substitution, fewer workers migrate from the home sector to the market sector when labor market frictions fall.

The moderate sensitivity of our results to θ , η , and σ may seem puzzling. But remember that, as we entertain different parameter values, we simultaneously change the A 's and τ 's to fit observed wages and employment shares of the young in each occupation and group in each year. However, our results are more sensitive if we vary the key parameters holding all other parameter values and forcing variables fixed (the A 's, ϕ 's, τ 's, etc.). That is, the sensitivity increases if we do not recalibrate. Consistent with the intuition provided in Section 2.9, the gains from changing τ 's rise dramatically as we raise θ holding the other forcing variables fixed. Specifically, when ability is less dispersed (θ is higher), comparative advantage is weaker and the allocation of talent is *more* sensitive to changing τ 's. The higher is θ , the more occupational decisions are distorted by given barriers, and hence the bigger the gains from removing them. For example, if we hold the A 's, \tilde{z} 's, τ 's, and ϕ 's fixed at baseline estimates, the declining τ 's explain 177% of the growth in market GDP per person when $\theta = 3$.³¹

5. INFERENCE AND RESULTS WITH IDIOSYNCRATIC PREFERENCES ($\delta > 0$)

5.1. Inference With $\delta = 1$

When selection is entirely on idiosyncratic preferences rather than idiosyncratic ability, the way we identify the τ 's and \tilde{z} 's changes. Table IX summarizes how, italicizing the

³¹In the Supplemental Material (Hsieh et al. (2019)), we highlight the robustness of our results to other empirical choices.

TABLE IX
FORCING VARIABLES AND EMPIRICAL TARGETS WITH $\delta = 1^a$

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

^aVariables are chosen to match the empirical targets. $\delta = 1$ is the polar case where individuals draw idiosyncratic tastes for each occupation, with no idiosyncratic heterogeneity in ability across occupations.

entries that switch in this polar case. With selection only on preferences, average preferences for an occupation do not affect average wages in that occupation. Using equation (7) and imposing $\delta = 1$, the average wage gap in an occupation between a group and white men can be expressed as

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

With $\delta = 1$ and $\bar{h}_{ig}/\bar{h}_{i,wm} = 1$, the wage gap between groups within an occupation is pinned down by the relative τ 's between groups in that occupation. Conditional on the τ 's implied by the wage gaps, the relative \tilde{z} 's can be inferred from the relative share of a group in an occupation. In the polar case ($\delta = 0$) considered earlier, the inference was precisely the opposite: the \tilde{z} 's were inferred from the occupational wage gaps and the τ 's were inferred from the differences in occupational sorting (conditional on the \tilde{z} 's).

5.2. Calibrating δ

If selection is based at all on tastes for an occupation ($\delta > 0$) instead of only on ability ($\delta = 0$), then average wages in an occupation should be positively correlated with occupational shares. Figure 8 shows there is no systematic relationship between the share of young white women relative to young white men in an occupation in 1980 and the occupational wage gaps between young women and men in the same year.³² For example, young white women were 64 times more likely to work as secretaries and one-fourth as likely to work as lawyers in 1980. Yet the wage gap between young white women and young white men among secretaries was nearly identical to the gender wage gap among lawyers. Figure 9 shows that there is similarly no systematic correlation between the *change* in relative gender occupational shares and the *change* in the gender occupational wage gap over that time, both from 1960 to 2010.

Of course, it is possible to reconcile any pattern of wage gaps—and how they change over time—with just the right levels and changes in preferences (\tilde{z} 's) across occupation-groups and years. To offset selection on idiosyncratic preferences, group-specific occupational preferences (the \tilde{z} 's) would need to be negatively correlated with occupational barriers. Women would need to dislike working in occupations in which they are discriminated against. As occupational barriers fall, their preferences would need to move in favor

³²A weighted regression of the scatter plot yields a slope coefficient of 0.01 with a standard error of 0.01.

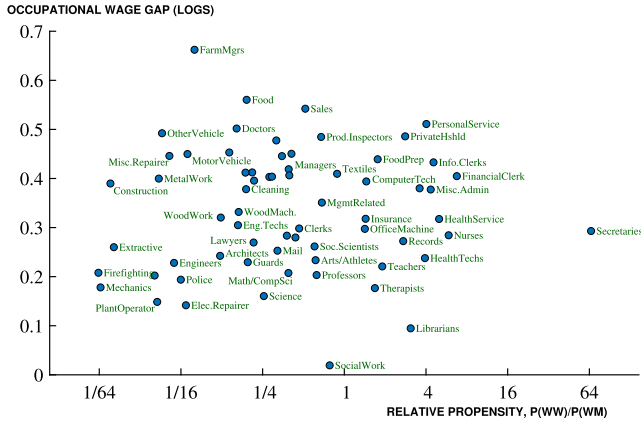


FIGURE 8.—Wage gaps versus propensities across occupations for white women in 1980. Note: The figure shows the relationship between the (log) occupational earnings gap for white women compared to white men (both in the young cohort) and the relative propensity to work in the occupation for the two groups, $p_{i,ww}/p_{i,wm}$ in 1980. A simple regression line through the scatter plot in the top panel yields a slope coefficient of 0.01 with a standard error of 0.01.

of those occupations. In this sense, the τ 's and \tilde{z} 's could both be barriers to occupational choice and human capital formation for women and black men.

The gains from changing τ 's and \tilde{z} 's, however, may depend on the value of δ . We therefore wish to show how the growth contributions of changing τ 's and \tilde{z} 's vary with δ . Just as important, we would like to estimate a plausible value for δ . Rearranging (6) for young white men in year t , we can express schooling-adjusted wages as

$$\frac{\overline{\text{wage}}_{i,wm}(t)}{[1 - s_i(t)]^{-\frac{1}{3\beta}}} \propto p_{i,wm}(t)^{\frac{\delta}{\theta(1-\eta)}} \tilde{z}_{i,wm}(t)^{-\frac{1}{1-\eta}}.$$

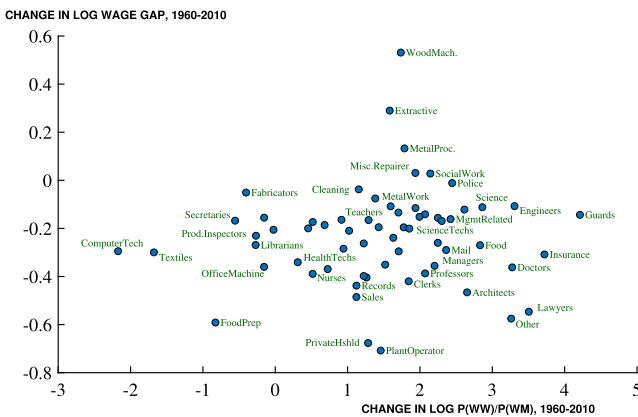


FIGURE 9.—Changes in wage gaps versus propensity gaps, young white women 1960–2010. Note: The figure shows the relationship between the change in (log) occupational earnings gaps for white women compared to white men (both in the young cohort) and the change in the relative propensity to work in the occupation for the two groups, $p_{i,ww}/p_{i,wm}$, between 1960 and 2010. A simple regression line through the scatter plot in the top panel yields a slope coefficient of 0.06 with a standard error of 0.05.

Letting Δ denote first differences, we obtain

$$\Delta \ln \left(\frac{\overline{\text{wage}}_{i,\text{wm}}(t)}{[1 - s_i(t)]^{-\frac{1}{3\beta}}} \right) = \text{constant} + \frac{\delta}{\theta(1 - \eta)} \Delta \ln p_{i,\text{wm}}(t) - \frac{1}{1 - \eta} \Delta \ln \tilde{z}_{i,\text{wm}}(t). \quad (14)$$

As (14) demonstrates, $\delta > 0$ implies that wages (conditional on years of schooling s_i) should be rising in occupations with rising propensities, controlling for preferences. There will be no such relationship between wage changes and propensity changes when $\delta = 0$, that is, when there is selection only on ability. When there is selection on preferences, marginal workers with little preference for an occupation must be attracted by higher wages to enter growing occupations.

If we run OLS on (14) for young white men from 1960 to 2010, treating preferences as the residual, the implied δ is 0.077 with a standard error of 0.055.³³ Figure 10 plots the empirical counterpart of such an OLS regression. There is a simultaneity problem with OLS, however, because the residual preference shifters are likely correlated with changing occupational shares. We thus propose to instrument for changing p 's using the model-implied changes in occupation technology parameters (A 's). This is a valid instrument assuming the orthogonality condition:

$$\Delta \ln A_{i,\text{wm}} \perp \Delta \ln \tilde{z}_{i,\text{wm}}.$$

That is, we assume that changes in technology are uncorrelated with changes in preferences across occupations for young white men. We iterate on δ until we obtain a value




FIGURE 10.—Changes in wages versus propensities, young white men 1960–2010. Note: The figure shows the relationship between the log change in (log) occupational earnings for young white men (adjusted for schooling) and the change in the log of their propensity to work in the occupation, $p_{i,\text{wm}}$, between 1960 and 2010.

³³ Average earnings and schooling by occupation are from the Census. We convert years of schooling into s by dividing by 25 years, our assumed pre-work time endowment. We use the benchmark values of $\theta = 2$ and $\eta = 0.103$ to calculate the implied δ from the OLS coefficient. Running this regression across occupations in each year, the coefficient averages 0.123 across years, with standard errors of around 0.04.

such that the implied growth rates of technology and preferences are uncorrelated with each other. This yields an estimate of $\delta = 0.22$. This is higher than the OLS estimate because it adjusts for the covariance between changing wages and changing preferences in (14). The fact that changes in occupational propensities and wages (adjusted for years of schooling) are weakly correlated for young white men suggests that individuals primarily sort based on talent instead of preference heterogeneity.

5.3. Results With $\delta > 0$

Table X examines the contribution of falling barriers to growth when we entertain $\delta = 0.22$, and also $\delta = 1/2$ and $\delta = 1$. The first column repeats our results with $\delta = 0$ for comparison. With our estimated value of $\delta = 0.22$, the share of growth attributed to changing τ 's falls modestly. The differences are even smaller if we consider the share of growth coming from changing τ 's and changing preferences (\bar{z} 's).

The growth contributions from changing τ 's fall more markedly when we move to $\delta = 0.5$, but still remain economically important. For example, changing τ 's account for 28% of growth in market GDP per person with $\delta = 0.5$ versus 42% when $\delta = 0$. When $\delta = 0.5$, the share of growth coming from the τ 's and \bar{z} 's combined falls less, from 41% to 34%. As δ increases, the growth effects attributed to changing τ 's fall while the growth effects from changing \bar{z} 's rise, leaving the total effect on growth from the changing τ 's and \bar{z} 's less affected. In terms of economic growth, it matters little whether women in 1960 faced labor market or human capital frictions in becoming doctors and lawyers or whether social norms kept them from those professions.

Only when we go all the way to pure selection on idiosyncratic preferences, $\delta = 1$, do growth contributions from changing τ 's plummet. The contributions of changing τ 's and \bar{z} 's combined fall less dramatically. With idiosyncratic preferences only, changes in τ and

TABLE X
ALLOWING FOR SELECTION ON IDIOSYNCRATIC PREFERENCES^a

	Share of growth accounted for by τ^h and τ^w			
	$\delta = 0$	$\delta = 0.22$	$\delta = 0.5$	$\delta = 1$
Market GDP per person	41.5%	37.4%	28.3%	5.6%
Market earnings per person	38.4%	37.0%	31.2%	11.5%
Labor force participation	90.4%	88.1%	74.9%	21.2%
Market GDP per worker	24.0%	20.6%	14.2%	0.1%
Home + market GDP per person	32.7%	31.3%	25.7%	1.6%
	Share of growth accounted for by τ^h , τ^w , and \bar{z}			
	$\delta = 0$	$\delta = 0.22$	$\delta = 0.5$	$\delta = 1$
Market GDP per person	40.8%	38.9%	33.7%	19.8%
Market earnings per person	37.5%	36.8%	33.4%	20.8%
Labor force participation	112.7%	117.6%	117.1%	71.2%
Market GDP per worker	15.0%	12.7%	8.6%	1.8%
Home + market GDP per person	32.1%	32.3%	29.8%	7.5%

^aEntries in the table show the share of growth in the model attributable to changing frictions τ^h (human capital frictions), τ^w (labor market frictions), and \bar{z} (occupational preferences), for alternate values of δ , the fraction of workers who draw occupation-specific preferences rather than ability. $\delta = 0$ is our baseline with selection based only on talent. $\delta = 1$ is the opposite pole, with selection based only on preferences.

\tilde{z} contribute to growth only if they result in more representation of the group in high A_i occupations. Changes in dispersion of the τ 's and \tilde{z} 's by themselves do not affect aggregate output, at least in the case where the distortions and A_i are log-normally distributed. The fact that about 20% of growth was due to the change in τ and \tilde{z} in the model where individuals have the same ability in all occupations indicates that τ/\tilde{z} of women and black men fell by more in high A_i occupations over the last 50 years.

6. ROBUSTNESS TO ALTERNATE IDENTIFYING ASSUMPTIONS

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

Table XI shows how our results change with alternative assumptions about the evolution of \bar{h} across groups within different occupations over time under our base scenario of $\delta = 0$.³⁴ The first row of the table redisplayes our baseline estimates for market GDP per person growth from Table V. The second row relaxes the assumption that men and women draw from the same distribution of talent in all occupations. In particular, we consider the possibility that some occupations rely more on physical strength than others, and that this reliance might have changed over time because of technological progress. To see the potential importance of this story, we go to the extreme of assuming no frictions at all faced by women in any of the occupations where physical strength is arguably important (i.e., $\tau_{ig}^h = \tau_{ig}^w = 0$ for women in these occupations). These occupations include construction,

TABLE XI
ROBUSTNESS TO ALTERNATIVE ASSUMPTIONS ABOUT GROUP DIFFERENCES
IN TALENT^a

	Market GDP per person growth accounted for by τ^h and τ^w
Benchmark	41.5%
No frictions in "brawny" occupations	38.7%
No frictions in 2010	35.8%

^aEntries are the share of market GDP per person growth in the model attributable to changing frictions τ^h (human capital) and τ^w (labor market). A key identifying assumption is that any talent differences across groups, to the extent they exist, are constant over time. In our baseline specification, we assume that $\bar{h}_{ig} = 1$ for all occupations and all groups; in other words, that there are no innate talent differences between groups. The first row of the table recounts are baseline estimates. In the second row, we allow men and women to have different \bar{h} 's in "brawny" occupations. Specifically, we assume no gender-specific τ 's in these occupations. Instead, we allow the \bar{h} 's to evolve to exactly fit the quantity data for these occupations. "No frictions in 2010" (the third row) assumes that there are no frictions in 2010 for any group, so that differences in \bar{h}_{ig} explain all group differences in that year; we then calculate τ 's for earlier years assuming the mean value of the distribution of market skills in 2010 apply to earlier years. For the results in this table, we assume $\delta = 0$.

³⁴Given that the growth results from $\delta = 0$ are so similar to the growth results with our estimates of $\delta = 0.22$, we focus our robustness results on the $\delta = 0$ scenario.

firefighters, police officers, and most of manufacturing.³⁵ We estimate differences in \bar{h}_{ig} for young women to fully explain their allocation to these occupations in 1960, 1970, . . . , 2010. As shown in Table XI, the fraction of market GDP growth per person explained by changing frictions falls only slightly from 41.5% to 38.7% with this alternative identifying assumption. Our results are not sensitive to this alternative because most of the gains we attribute to changing τ 's come from the rising propensity of women becoming lawyers, doctors, scientists, professors, and managers—occupations where physical strength is not important.

The last row in Table XI makes a more extreme assumption. In this alternative, we allow all groups to have different levels of innate talent in all occupations. We assume, however, that these innate talent differences are constant over time. Specifically, we assume all group differences among the young in 2010 reflect talent rather than distortions. Specifically, we set the 2010 τ 's to zero for all groups and all occupations and assume differences in \bar{h}_{ig} fully account for group differences in occupational choice among the young in 2010. We keep talent in prior years at the 2010 values for each group, but back out distortions in earlier years. In essence, this specification allows for arbitrary talent differences between men and women to fit the 2010 data. Under this more flexible alternative, eliminating the τ 's in the earlier years still generates 36% of growth in market GDP per person. Thus, our gains are not an artifact of assuming the allocation of talent was far from optimal in 2010.

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

Another assumption that facilitates our identification is that white men face no labor market or human capital frictions. An alternative assumption might be that there was no discrimination in 1960 at all, but growing discrimination *against* men and *in favor* of women since then. If we assume women and men have the same mean talent, as we do in our baseline specification of $\delta = 0$, this would imply identical average wages and occupational distributions for women and men in 1960. This is something we do not observe in the data. Assuming relative talent stays constant over time, this alternative would also require women to earn increasingly more than men and be increasingly over-represented in high-skilled occupations after 1960. All of these predictions are at odds with the patterns documented above. If men and women have the same level of innate talent, the data strongly reject the hypothesis that men have been increasingly discriminated against over time.

Another alternative would be to assume discrimination *in favor* of men and no discrimination against women in 1960, with the discrimination in favor of men abating over time. This would fit the facts on wages and sorting over time, and would imply falling misallocation. But it is not isomorphic to our baseline assumption. First, it would imply falling education spending by men over the decades. Second, it would entail huge subsidies for men that diminish over time. When we calibrate our baseline model with $\delta = 0$, we find that earnings of men must exceed their marginal product by orders of magnitude. The implied total subsidy to men would be multiples of 1960 GDP. Men must be paid massive

³⁵Rendall (2017) classified occupations based on the importance of physical strength, and we define brawny occupations for our analysis as those occupations in the top half of her brawny distribution.

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

Yet another alternative would be to assume—contrary to our presumption—that women are somehow innately less talented than men, supposedly explaining women's lower wages and under-representation in skilled occupations in 1960. Rising discrimination in *favor* of women since 1960 might then account for the closing gaps between men and women. This hypothesis would entail *rising* misallocation and a drag on aggregate growth. Data on individual test scores suggest women are not less talented than men. The Armed Forces Qualifying Test (AFQT) was administered in both the NLSY 1979 and the NLSY 1997. The NLSY tracks a sample of individuals who were 12–16 years old when the surveys started. The AFQT scores in the NLSY are very similar for men and women in both 1979 and 1997.³⁶ According to these scores, women seem no less talented than men in their early teens. If we condition on working, women likewise have similar scores to men in both 1979 and 1997.³⁷ If one believes the story of rising discrimination in favor of women, one would have expected the relative test scores of working women to fall along with their rising participation rates. AFQT scores do not support the hypothesis that women are innately less talented than men.

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

7. CONCLUSION

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

We apply this framework to measure the changes in barriers to occupational choice facing women and black men in the United States from 1960 to 2010. We find large reductions in these barriers, concentrated in high-skilled occupations. We then use our general

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

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

equilibrium setup to isolate the aggregate effects of the reduction in occupational barriers facing these groups. Our baseline calculations suggest that falling barriers explain roughly 40% of aggregate growth in market GDP per person.

In our baseline model, we assume that individuals are heterogeneous only in their draws of occupational talent. We show that our baseline results are robust to instead assuming that individuals draw idiosyncratic occupational preferences. Even under this polar assumption, we find that one-fifth of U.S. market GDP growth can be explained by falling labor market barriers, falling human capital barriers, and shifting occupational preferences. Much of the productivity gains come from drawing women and black men into high-skilled occupations. Whether women increased their propensity to become lawyers and doctors because of declining labor market frictions or because of changing social norms, the growth implications are similar. That said, we estimate that occupational sorting based on talent draws better fits the data than occupational sorting based on preference draws.

Our general equilibrium Roy model assumed no correlation between an individual's absolute advantage and their comparative advantage, and that comparative advantage or preferences followed a Fréchet distribution. These assumptions aided tractability, but we hope they can be relaxed in future work.

REFERENCES

- ADÃO, R. (2016): "Worker Heterogeneity, Wage Inequality, and International Trade: Theory and Evidence From Brazil," Unpublished Paper, University of Chicago. [1451]
- AKCIGIT, U., J. GRIGSBY, AND T. NICHOLAS (2017): "The Rise of American Ingenuity: Innovation and Inventors of the Golden Age," NBER Working Paper No. 4096. [1440]
- ALTONJI, J. G., AND R. M. BLANK (1999): "Race and Gender in the Labor Market," *Handbook of Labor Economics*, 3, 3143–3259. [1439,1440]
- ARROW, K. (1973): "The Theory of Discrimination," in *Discrimination in Labor Markets*, 3–33. [1440]
- BAILEY, M. J. (2006): "More Power to the Poll: The Impact of Contraceptive Freedom on Women's Life Cycle Labor Supply," *Quarterly Journal of Economics*, 121 (1), 289–320. [1440]
- BAILEY, M. J., B. HERSHBEIN, AND A. R. MILLER (2012): "The opt-in Revolution? Contraception and the Gender Gap in Wages," *American Economic Journal: Applied Economics*, 4 (3), 225–254. [1440]
- BECKER, G. (1957): *The Economics of Discrimination*. University of Chicago Press. [1440]
- BERTRAND, M. (2011): "New Perspectives on Gender," in *Handbook of Labor Economics*, Vol. 4, 1543–1590. [1439,1440]
- BISKUPIC, J. (2006): *Sandra Day O'Connor: How the First Woman on the Supreme Court Became Its Most Influential Justice*. HarperCollins. [1440]
- BLAU, F. D. (1998): "Trends in the Well-Being of American Women, 1970–1995," *Journal of Economic Literature*, 36 (1), 112–165. [1439]
- BLAU, F. D., P. BRUMMUND, AND A. Y.-H. LIU (2013): "Trends in Occupational Segregation by Gender 1970–2009: Adjusting for the Impact of Changes in the Occupational Coding System," *Demography*, 50 (2), 471–492. [1439]
- BLAU, F. D., M. A. FERBER, AND A. E. WINKLER (2013): *The Economics of Women, Men, and Work*. Prentice-Hall. [1439,1440]
- CARD, D., AND A. B. KRUEGER (1992): "School Quality and Black-White Relative Earnings: A Direct Assessment," *Quarterly Journal of Economics*, 107 (1), 151–200. [1440]
- CHETTY, R., A. GUREN, D. MANOLI, AND A. WEBER (2012): "Does Indivisible Labor Explain the Difference Between Micro and Macro Elasticities? A Meta-Analysis of Extensive Margin Elasticities," in *NBER Macroeconomics Annual*, 1–56. [1453]
- COSTA, D. L. (2000): "From Mill Town to Board Room: The Rise of Women's Paid Labor," *Journal of Economic Perspectives*, 14 (4), 101–122. [1440]
- EATON, J., AND S. S. KORTUM (2002): "Technology, Geography, and Trade," *Econometrica*, 70 (5), 1741–1779. [1444]
- FERNÁNDEZ, R. (2013): "Cultural Change as Learning: The Evolution of Female Labor Force Participation Over a Century," *American Economic Review*, 103 (1), 472–500. [1440]

- FERNÁNDEZ, R., A. FOGLI, AND C. OLIVETTI (2004): “Mothers and Sons: Preference Formation and Female Labor Force Dynamics,” *Quarterly Journal of Economics*, 1249–1299. [1440]
- GOLDIN, C. (1990): *Understanding the Gender Gap: An Economic History of American Women*. Oxford University Press. [1439]
- GOLDIN, C., AND L. F. KATZ (2002): “The Power of the Pill: Oral Contraceptives and Women’s Career and Marriage Decisions,” *Journal of Political Economy*, 110 (4), 730–770. [1440]
- (2012): “The Most Egalitarian of All Professions: Pharmacy and the Evolution of a Family-Friendly Occupation,” NBER Working Paper 18410. [1439]
- GREENWOOD, J., A. SESHADRI, AND M. YORUKOGLU (2005): “Engines of Liberation,” *Review of Economic Studies*, 72 (1), 109–133. [1440]
- HSIEH, C.-T., E. HURST C. I. JONES, AND P. J. KLENOW (2019): “Supplement to ‘The Allocation of Talent and U.S. Economic Growth,’” *Econometrica Supplemental Material*, 87, <https://doi.org/10.3982/ECTA11427>. [1444,1446,1448,1451,1459,1465]
- JOHNSON, G. E., AND F. P. STAFFORD (1998): “Alternative Approaches to Occupational Exclusion,” *Routledge Research in Gender and Society*, 2, 72–90. [1440]
- KARABEL, J. (2005): *The Chosen: The Secret History of Admission and Exclusion at Harvard, Yale, and Princeton*. Houghton Mifflin Harcourt. [1440]
- LAGAKOS, D., AND M. E. WAUGH (2013): “Selection, Agriculture, and Cross-Country Productivity Differences,” *American Economic Review*, 103 (2), 948–980. [1451]
- McFADDEN, D. (1974): “Conditional Logit Analysis of Qualitative Choice Behavior,” in *Frontiers of Econometrics*, ed. by P. Zarembka. Academic Press, 105–142. [1444]
- NEAL, D. A., AND W. R. JOHNSON (1996): “The Role of Pre-market Factors in Black-White Wage Differences,” *Journal of Political Economy*, 104 (5), 869–895. [1440]
- PAN, J. (2015): “Gender Segregation in Occupations: The Role of Tipping and Social Interactions,” *Journal of Labor Economics*, 33 (2), 365–408. [1439]
- PHELPS, E. S. (1972): “The Statistical Theory of Racism and Sexism,” *The American Economic Review*, 62 (4), 659–661. [1440]
- RENDALL, M. (2017): “Brain versus Brawn: The Realization of Women’s Comparative Advantage,” Unpublished Paper, Monash University. [1471]
- ROY, A. D. (1951): “Some Thoughts on the Distribution of Earnings,” *Oxford Economic Papers*, 3 (2), 135–146. [1440,1442]
- SMITH, J. P., AND F. R. WELCH (1989): “Black Economic Progress After Myrdal,” *Journal of Economic Literature*, 27 (2), 519–564. [1439]

Co-editor Daron Acemoglu handled this manuscript.

Manuscript received 22 February, 2013; final version accepted 13 May, 2019; available online 13 May, 2019.