WORKING PAPER · NO. 2021-40

Task Based Discrimination

Erik Hurst, Yona Rubinstein, and Kazuatsu Shimizu

JUNE 2021

Becker Friedman Institute | Macroeconomic Research Initiative
Becker Friedman Institute | Ronzetti Initiative for the Study of Labor Markets

5757 S. University Ave.
Chicago, IL 60637
Main: 773.702.5599
bfi.uchicago.edu

Task-Based Discrimination*

Erik Hurst   Yona Rubinstein   Kazuatsu Shimizu

June 23, 2021

Abstract

Why did the Black-White wage gap converge from 1960 to 1980 and why has it stagnated since? To answer this question, we introduce a unified model that integrates notions of both taste-based and statistical discrimination into a task-based model of occupational sorting. At the heart of our framework is the idea that discrimination varies by the task requirement of each job. We use this framework to identify and quantify the role of trends in race-specific factors and changing task prices in explaining the evolution of the Black-White wage gap since 1960. In doing so, we highlight a new task measure - *Contact* tasks – which measures the extent to which individuals interact with others as part of their job. We provide evidence that changes in the racial gap in *Contact* tasks serves as a good proxy for changes in taste-based discrimination over time. We find that taste-based discrimination has fallen and racial skill gaps have narrowed over the last sixty years in the United States. However, since the 1980s, the effect of declining racial skill gaps and discrimination on the Black-White wage gap were offset by the increasing returns to *Abstract* tasks which, on average, favored White workers relative to Black workers.

*We thank Jack Hurst for his help in constructing some of the data files used within the paper. We also thank David Autor, Joydeep Bhattacharya, Sandy Black, Kerwin Charles, David Deming, Chinhui Juhn, Patrick Kehoe, Pat Kline, David de Meza, Derek Neal, Matt Notowidigdo, and Elena Pastorino for helpful comments as well as seminar participants at Berkeley, Chicago, Edinburgh, Houston, LSE, the 2021 SOLE meetings, and Virginia. Author contact: erik.hurst@chicagobooth.edu, y.rubinstein@lse.ac.uk, and kazuatsu@uchicago.edu.

1 Introduction

Why did the Black-White wage gap decline so much during the 1960s and the 1970s, and why has it stagnated since? After the passage of the Civil Rights Act, the unconditional Black-White wage gap narrowed substantially from about 50 percent in the early 1960s to about 30 percent by 1980. Some attributed the rapid growth in Black relative wages to declining discrimination (Freeman (1973), Donohue and Heckman (1991)). Others pointed to improvements in Blacks' school quality and market skills (Smith and Welch (1989), Card and Krueger (1992)). However, since 1980, the Black-White wage gap has remained constant. The relative stagnation in labor market progress of Black men during the last forty years has been seen as a puzzle given the documentation of notable declines since 1980 in White's reported discriminatory attitudes (Krysan and Moberg (2016), Lang and Lehmann (2012)) and a racial convergence in characteristics and skills that are rewarded in the labor market (Altonji et al. (2012), Bayer and Charles (2018), Dickens and Flynn (2006), Murray (2007)).[1]

In this paper, we introduce a framework that integrates racial skill gaps and various notions of discrimination into a task-based model of occupational sorting. At the heart of our framework is the idea that discrimination against Black workers varies by the task requirements of each job. For example, one can imagine that taste-based discrimination operates more through tasks that require interactions with others, whereas statistical discrimination is more likely to be present in tasks where there exist large racial differences in underlying required skills.[2]

Merging notions of discrimination and racial skill gaps into a task-based model of occupational sorting has two benefits. First, the framework allows us to explicitly model and quantify how well-documented changes in task returns over time influence Black-White wage gaps recognizing that Black and White workers sort into occupations with different task requirements due to either labor market discrimination among equally skilled workers or differences in endogenously determined labor market skills. Second, by measuring the tasks content of jobs, we can better distinguish among different types of discrimination that contribute to racial wage differentials. Disciplining this framework with detailed micro data allows us to separately quantify the role of changing returns to tasks, declining racial prej-

[1] While some measures of racial skill gaps have narrowed post-1980, like gaps in standardized test scores as measured by the National Longitudinal Survey of Youths (NLSY), Neal (2006) documents that other measures of racial skill gaps have not narrowed between 1980 and the early 2000s.

[2] We wish to stress that taste-based and statistical-based motives for discrimination are neither mutually exclusive nor unrelated. Specifically, the racial gaps in labor market skills reflect the intergenerational transmission of discrimination via skills' formation in early ages (Heckman et al. (2006)) or the influence of schooling and job training later in life (Coate and Loury (1993)). Nevertheless, separating discrimination from racial skill gaps is particularly useful – especially in modern economies – when the labor market returns to certain skills are large and rising (Autor and Dorn (2013), Deming (2017)).

udice, and narrowing racial skills gaps in explaining the decline of the male racial wage gap between 1960 and 1980 and its stagnation thereafter.

The paper has three main sets of findings. First, we document a new set of facts about how the propensities of Black and White men to sort into an occupation differ by types of labor market tasks required in the occupation, and how those differences have evolved over time. These facts are used as ingredients to estimate the evolution of race-specific barriers faced by Black men in each task over the last half century in our model of occupational sorting. Second, using our structurally estimated model and reduced-form estimates, we offer a task-based explanation as to why racial wage gaps converged from 1960 to 1980 and then stopped converging thereafter despite a narrowing of racial skill gaps and declining measures of discrimination during this time period. We show that increasing returns to tasks that require complex analytical activities (*Abstract* tasks) post-1980 relatively disadvantaged Black workers and masked wage gains resulting from improvements in race-specific factors. Finally, bringing in additional micro data which includes measures of pre-labor market skills, we separate changes in taste-based discrimination over time from changes in other race-specific factors. We show that changes in the racial gap with respect to working in occupations that require individuals to interact with customers and co-workers (*Contact* tasks) is a good approximation for changes in taste-based discrimination over time.

We now provide more details on each of these findings. In the first part of the paper we present new facts about racial differences in occupational sorting. Drawing on the work of Autor et al. (2003), Dorn (2009), Autor and Dorn (2013) and Deming (2017), we characterize the task content of occupations along four key labor demand factors: "*Abstract*", "*Routine*", "*Manual*", and "*Contact*". The first three task measures come directly from Dorn (2009) and Autor and Dorn (2013), while the last measure is new and guided by Becker (1957)'s work on taste-based discrimination. Specifically, "*Contact*" measures the extent to which an occupation requires interaction and communication with others within the organization (co-workers) or outside the organization (customers/clients). This task provides a measure of labor market activities where taste-based discrimination is likely to be the most salient because the task requires interacting with others who may have discriminatory preferences. Using data from the US Censuses and American Community Surveys (ACS), we document that there was a large racial gap in the extent to which workers sort into jobs that require *Abstract* tasks in 1960 and that gap has remained essentially constant through 2018. This finding holds regardless of whether or not we control for trends in racial gaps in accumulated levels of schooling. Conversely, we show that there has been a large racial convergence in the *Contact* task content of jobs between 1960 and 2018. The large racial gap in the extent to which workers sort into jobs that require *Contact* tasks that existed in 1960 has almost disappeared by 2018. These differential trends in the racial gaps in occupational sorting associated with various tasks becomes the launching point for the rest of the paper.

In the second part of the paper, we develop a model of task-based discrimination based on the occupational sorting framework of Autor and Handel (2013). In our model, individuals are endowed with task-specific skills that are drawn from a known distribution. There are many potential tasks and, in turn, many different types of skills. Occupations are combinations of tasks with different weights and individuals have different mixtures of skills. We generalize the Autor and Handel (2013) model by allowing for: (1) individuals of differing races to differ in the mean skill levels they have, (2) taste-based discrimination to differ by task, (3) skills to be noisily observed by employers so as to have a meaningful notion of statistical discrimination, and (4) a non-employment option so to match differential trends in employment rates by race. The differences in pre-labor market skills and discrimination facing Black workers give rise to differential sorting into tasks between Black and White individuals in the spirit of Roy (1951). These sorting differences need to be accounted for in order to parse out the effects of race-neutral driving forces (such as time trends in task specific returns) and race-specific driving forces (such as a narrowing of racial skill gaps and/or declining discrimination) when explaining changes in racial wage gaps over time. Using the model structure and empirical moments on the differential occupational sorting of Black and White men, the changing returns to various tasks, and the evolution of the aggregate racial wage gap, we estimate the key driving forces of the model.

Using our estimated model, we find that the stagnation in the racial wage gap post-1980 is a product of two offsetting forces. On the one hand, a narrowing of racial skill gaps and declining discrimination between 1980 and 2018 caused the racial wage gap to narrow by 8 percentage points during this period, all else equal. On the other hand, the changing returns to tasks since 1980 – particularly the increasing return to *Abstract* tasks – widened the racial wage gap by about 7 percentage points during the same period. A rise in the return to *Abstract* tasks disadvantages Black workers because they are underrepresented in these tasks due to a combination of racial skill gaps and discrimination. In sum, the estimated model highlights that race-specific barriers have continued to decline in the U.S. economy post-1980 but the rising relative return to *Abstract* tasks has favored Whites since 1980. As a result, the Black progress stemming from narrowing racial skill gaps and/or declining discrimination did not translate into Black-White wage convergence during this period. On the other hand, we show that the relative wage gains of Black men relative to White men during the 1960 to 1980 period stemmed solely from declining discrimination and a narrowing of racial skill gaps; changing task prices did not undermine any of these gains during this earlier period.

Our structural model provides a road map to empirical researchers looking to uncover changing race specific factors in micro data. Specifically, the model suggests that researchers must not only control for racial differences is skills but also for changes in the returns to different skills when analyzing racial wage gaps over time. Using data from the NLSY, we show that controlling for time-varying returns to skills uncovers a strong convergence in racial wage gaps during the last four decades in the United States. The magnitude of the convergence in the racial wage gap is similar to the effect of declining race specific factors predicted by our structural model. With this discussion we also highlight why our task-based model yields quantitatively different conclusions about the extent to which race-specific forces have changed in the U.S. economy during the last forty years relative to methodologies that rely on purely statistical decomposition procedures (e.g., Juhn et al. (1991)) which ignore task-based sorting forces.

In the third part of the paper, we go one step further and decompose the change in race-specific forces over time into the part that is due to changes in racial skill gaps, the part due to changes in taste-base discrimination, and the part due to changes in statistical discrimination. The taste-based discrimination in our model is something akin to pure racial prejudice in the spirit of Becker (1957) while statistical discrimination stems from employers using easily observable characteristics such as race to forecast the expected productivity of their workers in the spirit of Phelps (1972) and Arrow (1973). To perform this decomposition, we bring in additional data from the 1979 and 1997 National Longitudinal Survey of Youths (NLSY's). Building on the work of Heckman et al. (2006), Altonji et al. (2012), and Deming (2017), we exploit three pre-labor market skill measures in the NLSY: cognitive, non-cognitive and social skills. The former is determined by respondent scores on AFQT tests while the latter two are based on responses to survey questions designed to measure personality traits like self-motivation, self-determination, and extroversion. We use these skill measures to impute racial gaps in model-generated task-specific skills.

We start by documenting the mapping between worker pre-labor market skills and the occupations to which they sort. In particular, we show that cognitive skills are most predictive of entry into occupations that require *Abstract* tasks while social skills are most predictive of entry into occupations that require *Contact* tasks. Additionally, we document large but declining racial gaps in cognitive skills over time but find no racial gaps in social skills in any of the time periods we explore.

We then develop a procedure to translate racial gaps in NLSY skill measures into racial gaps in model-generated task-specific skill gaps. The procedure consists of two steps. First, we load our structurally-imputed average task-specific skills by occupation onto the NLSY measures of average cognitive, non-cognitive, and social pre-labor market skills by occupation, as measured among White workers. Second, we use these loadings and the racial gap in NLSY skills to create a model-based estimate of racial skill gaps associated with each task in each period consistent with the racial skill gaps in the NLSY.

Based on this procedure, our model estimates that the convergence in the sorting into occupations that require *Contact* tasks between 1960 and 2018 is driven almost entirely by declining taste-based discrimination. This result stems from the fact that there are almost no racial gaps in social skills, which implies that the racial barriers we estimate for *Contact* tasks can be mainly attributed to taste-based discrimination. This finding confirms our ex-ante conjecture that the evolution of the racial gap in *Contact* tasks is a good predictor of the change in taste-based discrimination. To further provide evidence for this conclusion, we use data from Charles and Guryan (2008), which provide survey-based measures of taste-based discrimination for each U.S. state. Using cross-state variation, we show that racial gaps in *Contact* tasks are strongly correlated with the Charles-Guryan state-level measures of taste-based discrimination. We find a much weaker correlation with state-level measures of racial gaps in *Abstract* tasks.

Finally, we use the estimated model to quantify how much the changes in each of the driving forces over time contributed to the evolution of the racial wage gap over the last half century. We estimate that at least half of the decline in the overall racial wage gap between 1960 and 2018 can be attributed to declining taste-based discrimination. On the other hand, a racial gap in skills required for *Abstract* tasks explains much of the remaining racial wage gap post-1980.

Before proceeding, we note that our work is closely related to the recent paper by Bayer and Charles (2018). Bayer and Charles (2018) importantly attribute the lack of positional improvement for median Black men since 1940 despite the narrowing of the racial education gaps to differential trends in the returns to high school versus post-secondary schooling. The rising return of college education relative to high school education disadvantaged Blacks as they still disproportionately possess lower levels of school credentials. We use their result as a launching point for our approach and study the trends in Black-White gaps *conditional on schooling*. In particular, focusing on the fact that Black-White labor market progress has stalled even conditional on education, we extend their insights to a task-based model of occupational sorting with multiple tasks and show that higher returns to *Abstract* tasks have disadvantaged Black men relative to White men even conditional on education.

Our paper is also related to the recent paper by Hsieh et al. (2019) which proposes and estimates a multi-sector Roy model of occupational sorting with workers of different races and gender who face differential frictions in both human capital and labor markets.[3] The

[3]There is an extensive literature exploring racial differences in labor market outcomes. Smith and Welch (1989), Altonji and Blank (1999), and Lang and Lehmann (2012) provide excellent surveys of this literature.

goal of Hsieh et al. (2019) is to provide a framework with economically meaningful sorting to assess the role of changes in racial and gender barriers during the last half century to economic growth. Our work complements this paper by extending the occupational sorting decision to a multi-dimensional task frame work in the spirit of Autor et al. (2003), Acemoglu and Autor (2011), and Autor and Handel (2013). Instead of trying to explain U.S. economic growth, our goal is to use our model of sorting to decompose racial wage gaps into taste-based discrimination, racial skill gaps, and statistical discrimination. Additionally, by embedding our model of racial differences into a task model of occupational sorting, we can address the extent to which changes in task returns can help to reconcile the puzzle as to why the Black-White wage gap stagnated since 1980.

2 Data and Measurement

In this section, we provide an overview of the data and measures used throughout the paper. The online appendix provides more detail on both our data and sample selection.

2.1 Task Measures

To assess whether Black and White workers sort into different jobs, preform different tasks and consequently earn different amounts, we measure the skill demands in each occupation using the following data: (i) the U.S. Department of Labor's Dictionary of Occupational Titles (DOT) and (ii) the Occupational Information Network (O*NET) sponsored by the U.S. Department of Labor/Employment and Training Administration (USDOL/ETA). The DOT was constructed in 1939 to help employment offices match job seekers with job openings. It provides information on the skills demanded in over 12,000 occupations. The DOT was updated in 1949, 1964, 1977, and 1991, and replaced by the O*NET in 1998.

We focus on four occupational task measures that are relevant for our study: *Abstract*, *Routine*, *Manual* and *Contact*. The first three measures are taken exactly from Autor and Dorn (2013) and Deming (2017) using the DOT data. Below, we provide a brief summary of these measures. The last task measure is new and was created specifically for this paper to help get at the concept of taste-based discrimination. Building on the work in Deming (2017), *Contact* measures the extent to which an occupation requires interaction and communication with others within the organization (co-workers) or outside the organization (customers/clients). The intensity of this task hence provides a measure of labor market

Surveying this literature is beyond the scope of our paper. But, it should be noted that our paper builds upon the seminal papers modeling both taste-based (Becker (1957)) and statistical discrimination (Phelps (1972), Arrow (1973), Aigner and Cain (1977), and Coate and Loury (1993)).

activities where the intensity of taste-based discrimination is likely to be the most salient. Our conjecture is therefore that the trend in the Black-White difference in the propensity to sort into occupations that intensively require *Contact* tasks, conditional on other task requirements, proxies the trend in the intensity of taste-based discrimination facing Black workers. One of the main objectives of the paper is to provide evidence for this conjecture.

We now summarize the four task measures:

Abstract: indicates the degree to which the occupation (i) demands analytical flexibility, creativity, reasoning, and generalized problem-solving and (ii) requires complex interpersonal communications such as persuading, selling, and managing others. Occupations with high measures of *Abstract* tasks include accountants, software developers, high school teachers, college professors, judges, various medical professionals, engineers, and managers.

Routine: measures the degree to which the task requires the precise attainment of set standards and/or repetitive manual tasks. Occupations with high measures of *Routine* tasks include secretaries, dental hygienists, bank tellers, machinists, textile sewing machine operators, dressmakers, x-ray technology specialists, meter readers, pilots, drafters, auto mechanics, and various manufacturing occupations.

Manual: measures the degree to which the task demands eye, hand, and foot coordination. Occupations with high measures of *Manual* tasks include athletes, police and fire fighters, drivers (taxi, bus, truck), skilled construction (e.g, electricians, painters, carpenters) and landscapers/groundskeepers.

Contact: measures the extent that the job requires the worker to interact and communicate with others (i) within the organization or (ii) with external customers/clients or potential customers/clients. To create our measure of *Contact* tasks we use two 1998 O*NET work activity variables taken from Deming (2017). Specifically, we use the variables *Job-Required Social Interaction (Interact)* and *Deal With External Customers (Customer)*.[4] *Interact* measures how much workers are required to be in contact with others in order to perform the job. *Customer* measures how much workers have to deal with either external customers (e.g., retail sales) or the public in general (e.g., police work). To make our measure of the *Contact* task content of an occupation, we take the simple average of *Interact* and *Customer* for each occupation. Occupations with high measures of *Contact* tasks include various health care

[4] Deming (2017)'s focus is creating a measure of occupational tasks that require social skills and document how the returns to social skills have increased over time. His measure of social skills include measures of whether the job requires the worker to have social perceptiveness and the ability to coordinate, persuade and negotiate with others. His measure of social skills do not include measures for whether the task requires interactions with other co-workers or customers. He uses the measures of customer (*Customer*) and broader social interactions (*Interact*) as controls in some of his specifications. These questions are much more suited to our purpose of trying to measure taste-based discrimination. We explore the relationship between Deming's *Social Skills* task measure and our *Contact* task measure in the online appendix.

workers, waiter/waitress, sales clerks, lawyers, various teachers, and various managers.

Our goal is to stay as close as possible to the definitions of task measures developed by others so as to provide new evidence on the racial differences in these measures. However, in the online appendix, we show that the racial differences in the task content of occupations that we highlight are very similar using alternative task definitions. We directly download all of our DOT and O*NET measures from the replication kit associated with Deming (2017). The occupational task measures are available at the 3-digit occupational code level. We use Deming (2017)'s crosswalk to merge these measures to our samples from the other data sets we use. A full discussion can be found in the online appendix.

Finally, we convert the task measures into z-score space by taking unweighted differences across occupations. This transforms the units of our task measures into standard deviation differences in the task content of a given occupation relative to all other occupations; an *Abstract* task measure of 2.0 in a given occupation means that occupation has an *Abstract* task requirement that is two standard deviations higher than the average occupation. Some occupations require all tasks in relatively high intensities. For example, civil engineers have *Abstract*, *Routine*, *Manual*, and *Contact* task intensities of 2.3, 1.2, 0.6, and 0.1, respectively. Some other occupations require all tasks in relatively low intensities. For example, mail carriers have *Abstract*, *Routine*, *Manual*, and *Contact* task intensities of -0.8, -1.5, -0.7, and 0.0, respectively. Other occupations are mixed in their task demands, and the differences in task demands differentiate between occupations. For example, both physicians and retail sales clerks are high in *Contact* intensities, but physicians are also high in *Abstract* task intensities while retail sales clerks are low in *Abstract* task intensities.

2.2 Census and American Community Survey

To measure long-run trends and cross regional differences in the task content of occupations and wages, we use data from the decennial U.S. Censuses from 1960 through 2000 and the annual American Community Surveys (ACS) thereafter. We pool together the micro data from the annual ACS's between 2010 and 2012 and again between 2016 and 2018. We refer to the former as the 2012 ACS and the latter as the 2018 ACS. Given this, we have seven separate waves of harmonized data for the years 1960, 1970, 1980, 1990, 2000, 2012 and 2018. Within each wave, we restrict our sample to Black and White native born men between the ages of 25 and 54 who do not live in group quarters. We also exclude workers who are self-employed. Finally, we always weight the data using the survey weights provided by the Censuses and the ACS's, respectively.

As discussed above, we use the DOT and O*Net data to define the task content of

Figure 1: Trends in Black-White Wage Gaps Since 1960, Census/ACS Data

Notes: Figure shows the trend in the demographically adjusted Black-White gap in log wages using Census/ACS sample. Wage gaps are conditional on individual age and education dummies.

occupations. We hold the task content of an occupation fixed over time. We measure wages as self-reported annual earnings during the prior year divided by self-reported annual hours worked during the prior year. We only measure wages for individuals who are currently employed working at least 30 hours per week and who reported working at least 48 weeks during the prior year. We treat individuals who are not working as being in the home sector occupation. In most specifications, we control for the worker's age and accumulated years of schooling. All values in the paper are in 2010 dollars.

Figure 1 shows the difference in log wages between Black and White workers conditional on age and education using our sample of Census/ACS individuals. In particular, for each year, we regress an individual's log hourly wage on a race dummy and controls of age (five-year age dummies) and series of dummies indicating the individual's accumulated level of education. Consistent with other findings in the literature, the demographically adjusted racial wage gap narrowed substantially between 1960 and 1980 but has been constant at a gap of roughly 20 log points since 1980.[5] Our goal is to explain both the wage convergence in Figure 1 between 1960 and 1980 as well as its stagnation post-1980.

[5] Chandra (2000), Heckman et al. (2000) and Bayer and Charles (2018) caution the literature about focusing on mean racial wage gaps over time given differential trends in labor force participation between Black and White men. For this reason, we explicitly include a margin of labor force participation in the model we develop below and highlight that our calibrated model matches moments of the wage distribution inclusive of those who do not work.

2.3 National Longitudinal Survey of Youths

We augment our analysis with data from 1979 and 1997 waves of the National Longitudinal Survey of the Youth (NLSY). The NLSY data allows us to link a worker's subsequent occupational choice with pre-labor market measures of cognitive, non-cognitive, and social skills. The NLSY waves are representative surveys of 12,686 and 8,984 individuals, respectively, who were 15 to 22 years old in 1979 or 13-17 years old in 1997 when they were first surveyed. The surveys were conducted either annually or bi-annually every year since for each cohort.

When using the NLSY data, we restrict the main sample to Black and White non-self-employed men 25 years of age and older.[6] As with the Census/ACS data, we hold the task content of occupations constant across both NLSY cohorts. Finally, in specifications where we measure the evolution of the racial gap in the task content of occupations over time, we restrict our NLSY samples to those aged 25-37. The reason for this is the cohort nature of the NLSY data. In the 1980s/early 1990s, respondents from the NLSY-79 were in their mid-20s to mid-30. Likewise in the 2010's, respondents from the NLSY-97 were also aged 25-37. This restriction ensures that we are comparing the occupational sorting and monetary rewards for young adults of similar ages when comparing across cohorts on the NLSY.

The key reason we use the NLSY data is to have measures of racial differences in pre-labor market skills. We use measures of performance on cognitive test and psychometric assessments to generate a set of unified proxies for cognitive, non-cognitive and social traits across the two NLSY waves. These skill measures were primarily collected before the individuals entered the labor market. Again, our goal is to take these skill measures directly from the existing literature. We summarize these measures briefly here and include a more detailed discussion in the online appendix.

Cognitive Skills (COG): We follow the literature and use the respondent's standardized scores on the Armed Forces Qualifying Test (AFQT) as our measure of cognitive skills. The AFQT is a standardized test which is designed to measure an individual's math, verbal and analytical aptitude. The test score was collected from all respondents in their initial year of the survey and was measured in both the 1979 and 1997 waves.[7]

Non-cognitive Skills (NCOG): We use the measures of non-cognitive skills created by Deming (2017). Deming (2017) uses questions pertaining to the Rotter Locus of Control

[6] As in with the Census/ACS data, we measure wages as annual earnings divided by annual hours worked. Following Altonji et al. (2012) and Deming (2017), we trim values of deflated hourly wage that are below \$2/hour and above \$500/hour.

[7] The AFQT score has been used by many in the literature to measure respondent's cognitive skills including Neal and Johnson (1996), Heckman et al. (2006), Neal (2006), Altonji et al. (2012) and more recently Levine and Rubinstein (2017) and Deming (2017). We follow Altonji et al. (2012) and Deming (2017) to generate age-adjusted AFQT scores.

Scale and the Rosenberg Self-Esteem Scale for the NLSY79 cohort to make a measure of non-cognitive skills.[8] Likewise, for the NLSY97 cohort Deming (2017) uses respondent answers (provided prior to entering the labor market) to the question "How much do you feel that conscientious describes you as a person?" to approximate respondents' non-cognitive skill. Deming (2017)'s non-cognitive skill measures are expressed in z-score units.

Social Skills (SOC): We again follow Deming (2017) to generate a unified measure of social skills using a standardized composite of two variables that measure extroversion in both waves. Specifically, for the NLSY79, we use self-reported measures of sociability in childhood and sociability in adulthood. Individuals were asked to assess their current sociability (extremely shy, somewhat shy, somewhat outgoing, or extremely outgoing) and to retrospectively report their sociability when they were age 6. For the NLSY97, we proxy for social skills using the two questions that were asked to capture the extroversion factor from the commonly-used Big 5 personality inventory. For each wave, we normalize the two questions so they have the same scale and then average them together. We then convert the measures into z-score units. Deming (2017) shows that these measures of social skills positively predict individual wages when they are adults even conditional on controlling for individual measures of cognitive skills (AFQT).

When using these skill measures, it is important to keep in mind that racial differences in skills can be the results of current or past discrimination. A large body of research documents the impact of interventions on different skills at different ages, suggesting that skills can be fostered (Knudsen and Shonkoff (2006), Heckman (2008), Almond and Currie (2011), Chetty et al. (2011) and Heckman and Kautz (2012)). Any racial differences in measures of standardized tests or in personality assessments measured among teenagers therefore almost certainly reflect differences in family, neighborhood or school environments stemming from current or past discrimination.

3 Racial Differences in Occupational Tasks

In this section, we document a set of key facts on the differential evolution of the task content of jobs between White and Black men during the last 60 years. We begin by showing aggregate patterns from the Census/ACS data. We then use data from the NLSY to highlight the relationship between pre-labor market skills and task-based occupational sorting.

[8] The Rotter scale measures the degree of control individuals feel they possess over the life. The Rosenberg scale measures perceptions of self-worth. Higher values of both are interpreted as high levels of non-cognitive skills. For example, Heckman and Kautz (2012) documents notable associations between educational attainment, health and labor market performance and these non-cognitive measures using NLSY data.

3.1 Trends in Racial Gaps in Tasks: Census/ACS

While much attention is focused on measuring racial wage gaps, less is known on racial differences in occupational sorting by task requirements. In this subsection, we fill this gap by documenting trends in racial task gaps from the early 1960s to the last 2010s. To measure the racial gaps in task content of occupations, we estimate the following linear probability model separately in each year restricting the sample using our sample of prime age men:

$$Black_{ijt} = \alpha_t + \sum_k \beta_{kt}\tau^k_{ijt} + \Gamma_{kt}X_{it} + \epsilon_{ijt}. \tag{1}$$

where $Black_{ijt}$ is a dummy variable equal to 1 if individual i working in occupation j during year t is Black; τ^k_{ijt} is the task content of task k for individual i working in occupation j in period t; and X_{it} is a vector of individual 5-year age dummies and five dummies measuring educational attainment (less than high school, high school, some college, a bachelor's degree, or more than a bachelor's degree).[9] Our coefficients of interest are the β_{kt}'s, which inform the change in the proportion of Black workers associated with a one standard deviation increase in task k requirements in year t, holding all other task requirements fixed. Each yearly regression yields four β_{kt}'s – one for each of our four task measures.

The results of these regressions are summarized in Figure 2. Panels A and B, respectively, show the patterns excluding and including the vector X_{it} of demographic controls. The main take-away from the figure is that both the level difference in racial task gaps in 1960 and the subsequent time series trend differ markedly by task. The differences are especially pronounced when we compare the racial gaps in *Abstract* and *Contact* tasks. In the early 1960s, Black workers were systematically underrepresented both in occupations that required a high intensity of *Abstract* tasks and in occupations that required a high intensity of *Contact* tasks. In terms of magnitudes, in 1960 a one-standard deviation increase in the *Abstract* task contents of an occupation reduced the probability that an individual working in that occupation was Black by about 3 percentage points, and a one-standard deviation increase in the *Contact* task contents reduced the probability that the individual was Black by about 4 percentage points, both conditional on education. Over the last half a century, however, Black men have made significant progress relative to White men with respect to sorting into occupations that require *Contact* tasks, while they made no progress relative to White men in the extent to which they sort into occupations that require *Abstract* tasks. Whereas

[9] In the online appendix, we show the raw trends in the τ^k_{ijk}'s by year for Black and White men separately. The raw patterns for *Abstract*, *Routine*, and *Manual* tasks for White men are similar to the findings in Autor and Dorn (2013). In the appendix, we also show a different specification where we regress τ^k_{ijt} on a race dummy and controls separately for each task in each year. The time series patterns for the coefficients on the race dummies from this specification matches nearly identically to the patterns shown in Figure 2.

Figure 2: Race Gaps in Task Trends

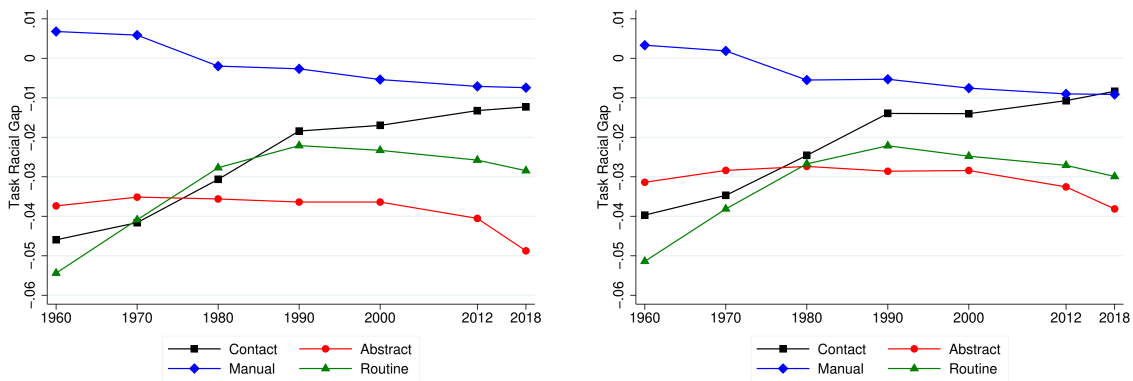

Panel A: No Controls Panel B: With Controls

Notes: Figure shows the estimated β_{kt}'s from the regression specified in equation (1). The coefficients measure the racial gap in the task content of occupations. Sample restricted to native born individuals between the ages of 25 and 54 within the Censuses and ACS years who are not self-employed and who are working more than 30 hours per week. Panel A excludes controls for age and education while Panel B includes those controls. Standard errors on the coefficients (omitted from the figure) had a value of less than 0.001 for all asks in all years.

the racial gap in *Abstract* tasks remained essentially constant through 2000 and widened slightly after 2000, the large racial gap in *Contact* tasks that existed in 1960 has all but disappeared by 2018. This finding persists whether or not we control for individual age and education (Panel A vs. Panel B), although the level of the *Abstract* task gap narrows once we control for them. If discrimination took the form of co-workers and customers not wanting to interact with Black workers in 1960, the patterns in Figure 2 are consistent with that form of discrimination abating over time.[10] Although we focus on such differences to a lesser extent, there were only small racial differences in the propensity for Black and White men to work in occupations that require *Manual* tasks. Likewise, there was some convergence in the propensity to work in occupations that require *Routine* tasks between 1960 and 1980, but large racial gaps still remain post-1980.

There is a large body of research documenting that taste-based discrimination was initially larger in the South region of the U.S. in the 1960s and 1970s (relative to other regions) and subsequently declined more in the South after 1980 (Charles and Guryan (2008), Bobo et al. (2012)). If the racial gap in sorting into occupations that require *Contact* tasks reflects taste-based discrimination, we should expect larger declines in the racial gap of this task

[10]In the online appendix Figure A5, we show that both the sub-components of *Contact* tasks – interacting with co-workers and interacting with customers – had large racial gaps in 1960 with those gaps narrowing sharply through 2018.

Figure 3: Census/ACS Task Content of Occupations: South Region vs Other Regions

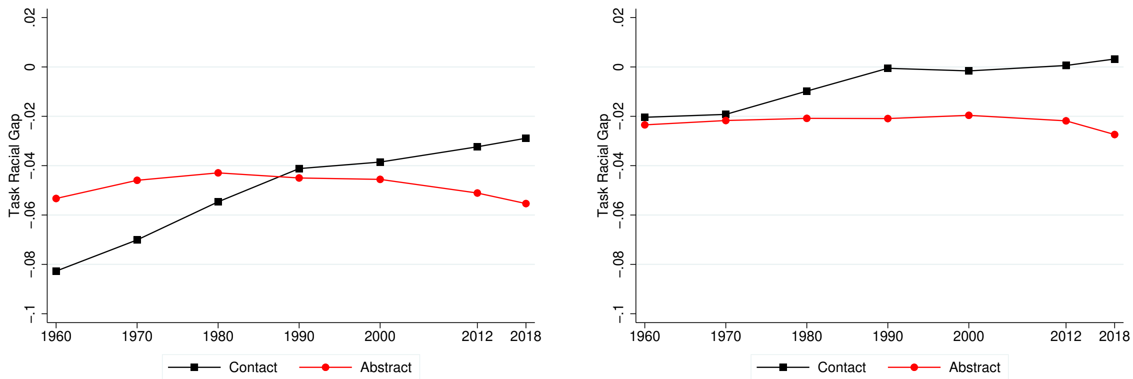

Panel A: South Region

Panel B: All Other Regions

Notes: Figure replicates the analysis in Panel B of Figure 2 separately for individuals residing in the South region (Panel A) and individuals residing in all other regions (Panel B).

measure in the South relative to other regions. Given that, we now explore cross-region variation in racial task gaps. Figure 3 replicates the analysis in Panel B of Figure 2 separately for the individuals in the Census/ACS data living in the South region (Panel A) and all other regions (Panel B). Consistent with our conjecture that the racial gap in *Contact* tasks proxies for taste-based discrimination, the racial gap in *Contact* tasks was much larger in the South relative to all other regions in 1960, and the subsequent convergence in *Contact* tasks over the last half century was also greater in the South relative to the other regions. Note, in both the South and the other regions, there was no racial convergence in *Abstract* tasks over time despite the racial gaps in *Abstract* tasks being larger in the South.

The contrasting trends in racial task gaps – namely the persistence of the large racial gap in *Abstract* tasks and the substantial narrowing of the racial gap in *Contact* tasks – highlight the need for a task-based approach to analyzing racial discrimination in the US labor market. The underlying factors that disadvantaged Blacks in the labor market – whether it may be discrimination, racial skill gaps, or a combination of both – impacted Blacks differently across tasks over time and induced differential trends in occupational sorting between Black and White workers.

The value-added from using a task-based approach to understand trends in racial wage gaps further arises where there are differential trends in task prices over time. To measure how the price of each task has evolved over time, we run the following regressions separately by year for each race group g using the the Census/ACS data:

$$w_{ijt} = \alpha_t^g + \sum_k \beta_{kt}^g \tau_{ijt}^k + \Gamma_{kt}^g X_{it} + \epsilon_{ijt}. \tag{2}$$

where w_{ijt} is the log wage of individual i working in occupation j during year t. Our coefficients of interest are again the β_{kt}^g's, the Mincerian wage premium of task k in year t for group g. For this regression, we use are sample of full-time workers.

Figure 4 reports estimates of the raw wage premium by task requirement for White men (Panel A) and the demographically adjusted Black-White gaps in the wage premium by task requirement (Panel B). Three main findings emerge from this figure. First, the average wage premium of *Abstract* tasks for White men was between 10 and 15 percent higher than the return to the other tasks in 1960. Moreover, the relative return of *Abstract* tasks has been increasing since 1980. This increase in the return to *Abstract* tasks has received lots of attention in the literature (Autor and Dorn (2013), Deming (2017)). Second, in contrast, the wage premium associated with *Contact* tasks was notably lower in the early 1960s and has not changed much since then. Finally, the racial gaps in the wage premiums to tasks are relatively small and roughly constant over time. Most of the racial gaps in the average task returns are slightly negative suggesting that the task return to Black men are systematically smaller than White men of similar age and education.

Figure 4: Minncerian Task Premiums, White Men and Racial Gap

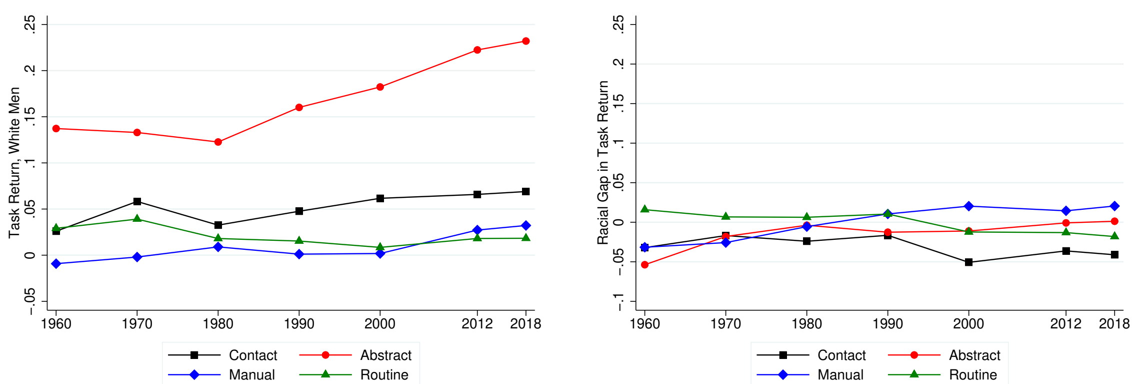

PANEL A: WHITE MEN, UNCONDITIONAL

PANEL B: RACIAL GAP, CONDITIONAL

Notes: Figure shows the average labor market return to occupational task content for White men in Panel A using our primary Census/ACS samples with the additional restriction that individuals report working at least 48 weeks during the prior year. This panel shows coefficients from a regression of log wages on the four task measures, separately by year. Panel B shows the racial gap in average task returns of Black men relative to White Men conditional on education and age.

Before concluding this subsection, we briefly mention the variety of alternate specifications we explored to examine the robustness of our results. All of the details of the robustness exercises are discussed in more detail in the online appendix. One concern is that the task intensities of occupations proxy the demand for general human capital rather than the demand for specific tasks. To explore this concern, we re-estimated all our main analysis separately segmenting our sample by those with less than a bachelor's degree and those with a bachelor's degree or more. Within both samples, we find that there was a racial convergence in the *Contact* tasks and no racial convergence in *Abstract* tasks; although, the convergence in the *Contact* tasks was much stronger among individuals with less than a bachelor's degree. These results highlight that our main findings about the time series patterns in racial task gaps are not being driven by the educational requirement of the occupations associated with the task. We also explored racial gaps in the task content of occupations for different birth cohorts (as opposed to pooling different cohorts together and exploring time series patterns). We find that the same patterns emerge across cohorts. In particular, whereas older cohorts and younger cohorts have the same racial gap in *Abstract* tasks, the racial gap in *Contact* tasks is large for older cohorts and almost zero for younger cohorts.[11]

3.2 Occupational Sorting and Racial Skill Differences

The ACS/Census data highlight racial differences in occupational sorting and how those differences have evolved over time. We now use data from the NLSY to assess the extent to which racial gaps in skill supplies can potentially explain racial gaps in sorting patterns. We begin by reporting whether the mixture of tasks demanded in each occupation predicts the mixture of skills supplied by workers who have sorted into the occupation. Specifically, we focus on the matching between relative task demands and pre-labor market skills for White men between the ages of 25 and 54 pooling together respondents from both the NLSY79 and NLSY97 samples. The results are shown in Table 1. Each column comes from a separate regression projecting an individual's cognitive, non-cognitive or social skills on the relative task content of the occupation in which they work. We define the relative task content of occupation j in which individual i works as $\tau_{ij}^{k} - \bar{\tau}_{ij}$ where $\bar{\tau}_{ij}$ is the simple average of the *Abstract*, *Routine*, *Manual*, and *Contact* task measures for occupation j.[12] Specifically, the

[11] This paper focuses on labor market differences between Black and White men. The appendix, however, also documents differences in task measures between White men and White women, as well as differences between White women and Black women. Like their male counterparts, the gap in *Abstract* tasks between Black and White women remained essentially constant since 1960. Further, the gap in *Contact* tasks between Black and White women narrowed substantively between 1960 and 2018. We choose to focus on Black and White men so as to abstract from the large trends in female labor supply that have also occurred during this time period.

[12] For example, suppose individual i works in occupation $j =$ *Civil Engineer*. As noted above, the *Abstract*, *Routine*, *Manual*, and *Contact* task content for the Civil Engineering occupation are, respectively, 2.3, 1.2,

Table 1: The Matching Between Individual Skills and Relative Job Tasks

	Cognitive Skills	Non-Cognitive Skills	Social Skills
(1) *Abstract* Tasks	0.179	0.043	0.030
	(0.015)	(0.021)	(0.020)
(2) *Routine* Tasks	0.077	0.010	0.004
	(0.019)	(0.025)	(0.025)
(3) *Contact* Tasks	0.117	0.067	0.082
	(0.019)	(0.024)	(0.023)
Difference (1) - (3)	0.062	-0.024	-0.052
	(0.021)	(0.029)	(0.029)
Demographic Controls	Yes	Yes	Yes

Notes: Table shows the relationship between the individual skills and the relative task content of the individual's occupation. Each column is a separate regression. The last row shows the difference between the coefficient on relative *Abstract* tasks and relative *Contact* tasks. Robust standard errors clustered at the individual level show in parenthesis. Data uses the pooled sample of the NLSY 1979 and 1997 waves. Sample restricted to White men between the ages of 25 and 54. Individual skills and occupational task contents measured in z-score units.

regression coefficients in the first column of Table 1 come the following specification:

$$S_{ij,cog}^{NLSY} = \alpha + \sum_k \omega_k(\tau_{ij}^k - \bar{\tau}_{ij}) + \Gamma X_i + \epsilon_{ij} \tag{3}$$

where $S_{ij,cog}^{NLSY}$ is the cognitive skill measure of individual i working in occupation j and X_i is a vector of individual age and education controls. Our coefficients of interest are the ω_k's. Given collinearity, we omit the relative task measure for *Manual* tasks from the regression implying that the ω_k's should be interpreted as the effect of working in an occupation that requires more of task k relative to *Manual* tasks on the cognitive skills of workers. In columns 2 and 3, we replace the dependent variable in equation (3) with the individual's non-cognitive skills ($S_{ij,ncog}^{NLSY}$) and social skills ($S_{ij,soc}^{NLSY}$), respectively.

Table 1 highlights that individual skill supplies respond differentially to relative task demands. Workers with higher cognitive skills are more likely to match with jobs that require higher *Abstract* tasks and workers with higher social skills are more likely to match with jobs

0.6, and 0.1 (in z-score units). For individuals working in Civil Engineering, $\bar{\tau}_{ij}$ would equal 0.9 and the relative task demand for *Abstract* tasks in this occupation would be 1.4 (2.3 - 0.9).

that require higher *Contact* tasks. For example, occupations where their relative *Abstract* task requirement is one-standard deviation higher attract workers who score approximately 0.2 standard deviations higher on cognitive tests. Occupations where their *Contact* tasks requirement is one-standard deviation higher attract workers whose social score is 0.08 standard deviation higher. The results in this table highlight that the NLSY skill measures are informative about the types of jobs into which individuals sort.

Having shown that the occupational task measures are associated with particular NLSY skill measures within the sample of White men, we now explore differences in pre-labor market skills between Black and White men within the NLSY cohorts.[13] Table 2 reports the racial gap in cognitive, non-cognitive, and social skills with various controls for the two separate NLSY samples. The first column for each sample includes all NLSY respondents in the sample without conditioning on employment; each of these samples has only one NLSY respondent per regression. The remaining columns pool over all years and only include individuals that were working. The second column within each sample adds no further controls, while the third column controls for the individual's maximum level of education and the last column controls also for their occupation.

The main takeaway from this table is that the racial gap in cognitive skills is large and narrows over time (especially for working men), whereas the gaps in non-cognitive and social skills are relatively small and constant over time. As seen from Table 2, Black men from the NLSY79 have AFQT scores that are 1.20 standard deviations lower than White men from the NLSY79 cohort. The gap declines to 0.93 when we additionally control for racial education and occupation differences. The race gap in AFQT scores narrowed somewhat between the NLSY79 and NLSY97 cohorts but was still large at -0.58 standard deviations conditional on education and occupation in the later period. On the other hand, the racial skill gaps in non-cognitive and social skills conditional on education and occupation were close to zero for both NLSY cohorts.

4 A Theory of Task Based Discrimination and Occupational Sorting

Autor and Handel (2013) propose a Roy model where workers with differential skill en-

[13]It is worth stressing again that racial differences in measures of pre-labor market skills are likely the result of current or past discrimination. Such a caveat should always be kept in mind when interpreting our results. For our purposes, these skill measures do predict occupational sorting patterns for both Black and White men as highlighted above. But, to the extent that such differences exist, the differences likely stem from factors such as differences in parental background, neighborhood choice, or school quality that resulted from current or past discrimination.

Table 2: Racial Gaps in Skill Measures (Z-Score Differences), NLSY Data

	1979 Cohort				1997 Cohort			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
(A) Cognitive Skills	-1.17	-1.18	-1.01	-0.93	-0.96	-0.80	-0.62	-0.58
	(0.03)	(0.04)	(0.03)	(0.03)	(0.05)	(0.06)	(0.05)	(0.05)
(B) Non-Cog. Skills	-0.20	-0.19	-0.10	-0.05	-0.12	0.06	0.16	0.17
	(0.04)	(0.04)	(0.04)	(0.04)	(0.05)	(0.07)	(0.07)	(0.07)
(C) Social Skills	-0.09	-0.11	-0.09	-0.06	-0.17	-0.15	-0.14	-0.12
	(0.04)	(0.04)	(0.04)	(0.04)	(0.05)	(0.06)	(0.06)	(0.06)
Employed Only Sample	No	Yes	Yes	Yes	No	Yes	Yes	Yes
Education Controls	No	No	Yes	Yes	No	No	Yes	Yes
Occupation Controls	No	No	No	Yes	No	No	No	Yes
Sample Size Clusters	4,226	3,702	3,702	3,702	2,354	1,870	1,870	1,870
Sample Size Observations	4,226	22,479	22,479	22,479	2,354	7,923	7,923	7,923

Note: Table shows the racial gap in various NLSY skill measures for various samples and with various controls. We show results separately for the 1979 cohort (columns (1)-(4)) and the 1997 cohort (columns (5)-(8)). Cognitive skills are measured as normalized AFQT scores. All racial gaps are measured in z-score differences between Black and White men. Columns (1) and (5) shows results for all individuals regardless of employment status; in these specifications each individual is only in the sample once. In the remaining columns we condition on the individual being employed in a given year. In these specifications, individuals can be in the sample multiple times. Robust standard errors are in parentheses.

dowments self-select into occupations according to their task requirements. We utilize the framework to build a task-based model of discrimination where racial differences in underlying skills and the existence of labor market discrimination creates differential sorting patterns between Black and White workers. Discrimination may take a form of taste-based or statistical discrimination, and the extent to which statistical discrimination impacts the returns will differ by tasks based on underlying gaps in average skill levels across racial groups.

As in Autor and Handel (2013), occupations are represented by bundles of K tasks, where the relative importance of tasks differs across occupations. We denote the task content of occupation j with a vector $\Lambda_j = (\lambda_{j1}, ..., \lambda_{jK}) \in \mathcal{R}_+^K$. Workers, in turn, perform tasks by allocating a unit of labor to the occupation of their choice, but each worker has differential efficiencies at performing each type of tasks. We denote the skill-endowment of worker i belonging to race group g with a vector $\Phi_{gi} = \{\phi_{i1}, ..., \phi_{iK}\} \in \mathcal{R}_+^K$. In particular, we suppose that each ϕ_{ik} is drawn from a Frechet distribution with shape parameter θ_k and scale parameter 1, both of which are common across race groups.[14]

[14]The assumption that the scale parameter of the Frechet distribution equals one is innocuous as we let

In absence of discrimination or racial differences in human capital, the potential log output of worker i belonging to race group g in occupation j is given by:

$$y_{gij} = \alpha_j + \sum_K \lambda_{jk}\phi_{ik},$$

where α_j is an occupation-specific constant that represents the potential log output of a worker with no skills in occupation j. If employers possess perfect information on individual worker's skills ϕ_{ik}'s – an assumption we will relax below – workers are paid their marginal revenue product, so that their potential log earnings in occupation j is given by:

$$w_{gij} \equiv p_j + y_{gij} = A_j + \sum_K \lambda_{jk}\phi_{ik},$$

where p_j is the log price of the output of occupation j and $A_j \equiv p_j + \alpha_j$. Conditional on working, each worker self-selects into the occupation j that maximizes her utility, which is the sum of log earnings and her non-pecuniary idiosyncratic preference for occupations $\log \nu_{ij}$:

$$u_{gij} = w_{gij} + \log \nu_{ij}.$$

We suppose that ν_{ij} is drawn from a Frechet distribution with shape parameter ψ and a normalized scale parameter of 1, both of which are common to all race groups.

Workers sort based on their comparative advantage. The optimal occupational choice of worker i in group g is given by

$$j^*_{gi} = \arg\max\nolimits_{j=1,\dots,J,H} \{u_{gij}\}. \tag{4}$$

Everything else equal, occupations that require a large amount of one type of tasks tend to attract workers who are good at performing that type of tasks. So an occupation that requires many *Abstract* tasks will tend to attract workers with higher *Abstract* skills than an occupation that requires few *Abstract* tasks, assuming that the two occupations are identical in other task requirements. This is the basic Autor and Handel framework.[15]

New to our model is differential returns to performing each type of tasks across different

the λ's to be scaled freely when calibrating the model. In addition, below we will introduce a parameter that controls differentials in overall skill-endowment levels across racial groups which is equivalent to allowing for a differential scale parameter between Black and White individuals.

[15] The original Autor and Handel framework does not have the idiosyncratic preference for occupations. It is added here to give the model an extra level of flexibility to fit differentials in employment across otherwise similar occupations. Below, we further expand on the Autor and Handel model by including racial groups that differ by average skill levels, modeling both taste-based and statistical discrimination, and adding a meaningful home sector.

race groups. We allow for three forces that can cause the returns to tasks to differ across races. First, the average skill level may differ across race groups stemming from current or past discrimination.[16] We capture this by writing the skill endowment vector as $\Phi_{gi} = \{\eta_{gk} + \phi_{i1}, ..., \eta_{gK} + \phi_{iK}\}$, where η_{gk} represents the skill (or human capital) level of group g associated with task k. Second, workers of a particular race group may face taste-based discrimination in performing a task. Taste-based discrimination may exist if for example customers do not like interacting with Black employees or White workers do not like interacting with their Black co-workers. In the presence of taste-based discrimination, we assume that employers perceive the efficiency of the discriminated workers to be $\delta_{gk}^{taste} + \eta_{gk} + \phi_{ik}$ rather than $\eta_{gk} + \phi_{ik}$, where δ_{gk}^{taste} is the race-specific taste-based discrimination coefficient in task k. Thus, the potential log wages of a worker belonging to race g becomes:

$$w_{gij} = A_j + \sum_{K} \lambda_{jk}(\delta_{gk}^{taste} + \eta_{gk} + \phi_{ik}).$$

Lastly, workers may face statistical (rather than taste-based) discrimination if their employers do not perfectly observe individual workers' skills. If skills are observed with noise, employers form expectations about a worker's marginal product by using information about the individual's group. The statistical discrimination coefficient, δ_{gk}^{stat}, will therefore differ by task based on underlying gaps in group mean of skills (η_{gkt}'s) used to perform the task.

We formally incorporate the notion of statistical discrimination into the model by introducing noise to skill measurement. Suppose employers cannot observe worker's true efficiency, $\eta_{gk} + \phi_{ik}$, and instead only observe a noisy skill measure given by

$$s_{gik} = (\eta_{gk} + \phi_{ik}) + \epsilon_{ik},$$

where the noise ϵ_{ik} is drawn from a normal distribution with mean zero and variance σ^2 (common to all race groups). Employers, however, observe a worker's group affiliation and know the underlying distributions of $\eta_{gk} + \phi_{ik}$ and ϵ_{ik}. In this environment, employers set the wage of each worker at the worker's expected marginal revenue product conditional on

[16] This difference can proxy for racial differences in parental background, neighborhood choice or school quality that result in differential task specific skill acquisition across racial groups.

observed skills $(\hat{s}_{i1}, ..., \hat{s}_{ik})$ and the worker's group affiliation:[17]

$$w_{gij} = w_{gj}^{cond}(\hat{s}_{i1}, ..., \hat{s}_{ik}; \sigma^2).$$

Normalizing $\eta_{gk} = \delta_{gk}^{taste} = 0$ for the base group $g = g_0$ (which will be White men in our application), the expected marginal revenue product perceived by employers (who may discriminate for taste-based motives) is given by:

$$w_{gj}^{cond}(\hat{s}_{i1}, ..., \hat{s}_{ik}; \sigma^2) = A_j + \sum_K \lambda_{jk} \left(\phi_{g_0 k}^{e}(\lambda_{jk}, \hat{s}_{ik}; \sigma^2) + \delta_{gk}^{taste} + \delta_{gk}^{stat}(\eta_{jk}, \lambda_{jk}, \hat{s}_{ik}; \sigma^2) \right), \quad (5)$$

where

$$\phi_{g_0 k}^{e}(\lambda_{jk}, \hat{s}_{ik}; \sigma^2) = \log E\left[e^{\lambda_{jk}\phi} \mid s_{g_0 ik} = \hat{s}_{ik}\right]^{1/\lambda_{jk}}$$

is the expected efficiency of a base-group worker (e.g., White workers) in task k conditional on observing $\hat{s}_{ik}$, and

$$\delta_{gk}^{stat}(\eta_{jk}, \lambda_{jk}, \hat{s}_{ik}; \sigma^2) = \log E\left[e^{\lambda_{jk}(\phi+\eta_{gk})} \mid s_{gik} = \hat{s}_{ik}\right]^{1/\lambda_{jk}} - \log E\left[e^{\lambda_{jk}\phi} \mid s_{g_0 ik} = \hat{s}_{ik}\right]^{1/\lambda_{jk}} \quad (6)$$

is the statistical discrimination coefficient (measured relative to the base group). In words, the statistical discrimination coefficient equals the gap in the conditional expected efficiency relative to the base group and will be non-zero if η_{gk} is non-zero and σ^2 is positive. Overall, racial wage gaps conditional on identical observed credentials will be a combination of taste-based and statistical discrimination:

$$w_{gj}^{cond}(\hat{s}_{i1}, ..., \hat{s}_{ik}; \sigma^2) - w_{g_0 j}^{cond}(\hat{s}_{i1}, ..., \hat{s}_{ik}; \sigma^2) = \sum_k \lambda_{jk} \left(\delta_{gk}^{taste} + \delta_{gk}^{stat}(\eta_{jk}, \lambda_{jk}, \hat{s}_{ik}; \sigma^2) \right). \quad (7)$$

Conceptually, it would be useful to see the statistical discrimination term δ_{gk}^{stat} as a product of a Bayesian updating process. Before they observe a signal (i.e., the observed skill s_{gik}), the employers' prior on the true efficiency of a worker coincides with the true efficiency distribution for the group to which the worker belongs. They thus expect the true skill of a randomly-chosen worker to differ by η_{gk} across groups. However, upon observing the signal s_{gik}, they update their prior to reflect this new piece of information. The extent of the updating depends on the reliability of the signal, namely the amount of noise with

[17] Strictly speaking, the expected marginal revenue product should be conditional on the worker choosing occupation j. However, note that workers choose occupations based on observable skills $(\hat{s}_{i1}, ..., \hat{s}_{ik})$ and not based on true efficiencies $(\eta_{g1} + \phi_{i1}, ..., \eta_{gK} + \phi_{iK})$, as the wages depend only on the former and not on the latter. Thus, conditional on observed skills, the distribution of $\phi' s$ among workers choosing occupation j is the same as the distribution of $\phi' s$ among all workers. Therefore, we can omit the conditioning on occupational choice.

which employers observe worker skills (σ^2). If the signal is perfect (σ^2=0), employers set the wages solely based on the signal and workers are paid exactly their true marginal product:

$$w_{gij} = A_j + \sum_K \lambda_{jk}(\delta_{gk}^{taste} + \eta_{gk} + \phi_{ik}). \tag{8}$$

In this case, there will be no statistical discrimination and the racial wage gap conditional on observed credentials will only stem from taste-based discrimination:

$$\lim_{\sigma^2 \to 0} \delta_{gk}^{stat}(\hat{s}_{i1}, ..., \hat{s}_{ik}; \sigma^2) = 0, \quad \forall \hat{s}_{i1}, ..., \hat{s}_{ik}.$$

Conversely, if the signal is completely uninformative ($\sigma^2 \to \infty$), no updating takes place and employers pay workers solely based on their initial priors. In this case, the statistical discrimination term for workers of group g (δ_{gk}^{stat}) will equal the mean racial skill gap between the group and the base group (η_{gk}) regardless of the observed credentials:

$$\lim_{\sigma^2 \to \infty} \delta_{gk}^{stat}(\hat{s}_{i1}, ..., \hat{s}_{ik}; \sigma^2) = \eta_{gk}, \quad \forall \hat{s}_{i1}, ..., \hat{s}_{ik}.$$

More generally, when signals are imperfect but not totally uninformative, the expected marginal product conditional on observed skills is something akin to a weighted average of the signal and the prior, where the relative weight on the latter increases with the variance of noise σ^2. Hence, employers will tend to pay more based on the group mean and less based on observed skills of individual workers in a noisier environment.

Another notable implication of equation (6) is the following:

Proposition 1. *The statistical discrimination term,* $\delta_{gk}^{stat}(\eta_{gk}, \lambda_{jk}, \hat{s}_{ik})$*, tends to zero as* $\eta_{gk} \to 0$.

Proof. See online appendix. □

This proposition says that there cannot be any statistical discrimination in tasks where there is no mean gap in skills between Black and White men. When skills are noisily observed by employers, employers put weight on their prior expected difference in skills between workers from different groups when setting individual wages. As racial skill gaps associated with a task tend to zero, statistical discrimination in that task will therefore also tend to zero. We use this proposition later in the paper when decomposing racial task gaps into the portion due to taste-based discrimination, the portion due to statistical discrimination, and the portion due to racial skill gaps.

We complete the model by allowing for differential sorting into employment, within and between groups, by introducing a "home sector", denoted as $j = H$. Specifically, we treat the home sector as another potential occupation (with task requirements $\lambda_{H1}, ..., \lambda_{HK}$ and occupational return A_{gH}) where the returns are non-pecuniary. Hence, the reservation utility u_{giH} of a worker with given observable credentials equals the log of the worker's expected marginal revenue product in an occupation with task requirements $(\lambda_{H1}, ..., \lambda_{HK})$ plus the log of the idiosyncratic preference for home sector, ν_{iH}.[18] Like the other occupational preferences ν_{ij}, the preference for home sector ν_{iH} follows a Frechet distribution with shape ψ and scale 1. We however allow the occupational return to home sector, A_{gH}, to differ by group g unlike with other A_j's. The differences in A_{gH}'s across groups thus capture any forces other than differential task returns that may create labor supply differences between racial groups.[19]

In sum, our task-based model of discrimination predicts differential sorting patterns across race groups as group-specific forces such as discrimination and racial skill differences make the task returns differ by group. As seen from equation (7), the composite discrimination term $(\delta_{gk}^{taste} + \delta_{gk}^{stat})$, which we shall denote by δ_{gk}^{total} hereafter, gives rise to differential returns to tasks across groups conditional on skills. Racial differences in skills, η_{gk}, might further lower the task return for one group relative to the return for another. The differential returns to tasks then induce differential sorting patterns across groups. The differential sorting across groups is a key feature of our model of task-based discrimination. Given the model structure, one can infer the size of racial barriers faced by Black men from the extent of differential sorting between Black and White men. We turn to such an analysis next.

5 Identifying Race-Neutral vs. Race-Specific Driving Forces

In this section, we discuss how we use micro data from the Census/ACS to identify and quantify a combination of discrimination and skill differences between Black and White workers (η_{bkt}'s+δ_{bkt}^{total}'s).

5.1 Calibration and Estimation of Base Model

To estimate and calibrate the baseline model, we proceed in two steps. First, we use micro data from O*Net and DOT combined with the occupational sorting and occupational earnings of White men to discipline the key race-neutral parameters governing occupational and

[18] The online appendix provides more detail on our exact specification of the home sector.

[19] The relative trend over time in the racial gap in the A_{gH}'s is small in our estimated model. This implies that other model forces such as the changing δ_{kt}'s, η_{kt}'s and λ's explain the observed differential trend in employment rate between Black and White men. Given the small quantitative importance in the change in the relative A_{gH}'s for our results, we relegate most of our discussion of these trends to the online appendix.

Table 3: Model Parameters and Data Targets

Panel A: Common Across Race Groups		
Parameter	Variation	Data Target
$\tilde{\tau}_{jk}$'s	Occupational task demands	ONET/DOT Data
β_{kt}'s	Task scaling factors	Mincerian task returns, White men Aggregate task content, White men
A_{jt}'s	Occupational marginal revenue product	Occupational shares, White men Occupational earnings, White men
ψ	Shape parameter occupational tastes	Labor supply elasticity
θ	Shape parameter task skills	Variance of log earnings, White men *Abstract* task returns, White men
σ	Variance of noise of worker skills	Exogenously set
Panel B: Varies Across Race		
$\eta_{gkt} + \delta_{gkt}^{taste}$	Racial differences in task skills plus task-based discrimination	Aggregate racial wage gap Aggregate racial gap in task contents Race gap in empirical task returns
$A_{g,H,t}$	Racial home sector preference	Share full time employed by race

Notes: Table lists key model parameters and data moments used to discipline the parameters.

task sorting (the A_{jt}'s and λ_{jkt}'s). We estimate the race-neutral driving forces separately for each year of our Census/ACS samples. Second, we use racial differences in occupational sorting, task returns and aggregate wages to pin down a composite term comprised of (η_{gk} $+\delta_{gk}^{taste}$) for Black men in each task. We again estimate these combined racial barrier terms separately for each year. Table 3 lists the key parameters of the model and data moments used to help discipline the parameters. We now provide more details.

Specifically, we use the O*NET and DOT data to discipline the task content $\Lambda_j = (\lambda_{j1}, ..., \lambda_{jK}) \in \mathcal{R}_+^K$ of occupations. As in our empirical work above, we will have four types of tasks ($K = 4$): *Abstract*, *Contact*, *Routine*, and *Manual*. However, we cannot directly use the z-scores of task content we defined earlier (the τ's) since $\lambda_{j1}, ..., \lambda_{jK}$ have to be positive. We construct $\lambda_{j1}, ..., \lambda_{jK}$ from the z-scores as follows. First, we linearly project the z-scores of task content to the unit interval [0, 1]. Denote the task content measure thus obtained with $\tilde{\tau}_{jk}$. Then, we define $\lambda_{jk} = \beta_{kt}\tilde{\tau}_{jk}$, where β_{kt} is a scaling factor which can be interpreted as the task's price, which we estimate. The scaling factor allows the task prices to vary over time given that we are holding the $\tilde{\tau}_{jk}$'s constant across years.

As noted above, we calibrate the model for White men using the Census and ACS data year-by-year. Specifically, the model for $g = White$ is given by:

$$\begin{aligned} u_{gij} &= w_{gij} + \log \nu_{gij}, \\ w_{gij} &= \log E\left[e^{p_j + y_{gij}} \mid s_{gi1}, ..., s_{gik}\right], \\ p_j + y_{gij} &= A_j + \sum_K \beta_k \tilde{\tau}_{jk}\phi_{ik}. \end{aligned}$$

The skill endowment ϕ_{ik} follows a Frechet distribution with shape θ, while the occupational preference ν_{ij} follows a Frechet distribution with shape ψ. Furthermore, the conditional expectation depends on a choice of σ, the noise with which employers measure worker skills which, as noted above, which we set to zero for this discussion. We outline how we set θ and ψ below. However, taking these parameters as given, the remaining parameters to be estimated each year for White men are: A_j's for $j = 1, ..., J$; A_{gH} for $g = White$; and the β_k's for $k = 1, ..., 4$. We estimate A_j, A_{gH} and β_k by targeting (i) the average log income of White men in each occupation in each year; (ii) employment share of White men in each occupation in each year; (iii) employment share of White men in the home sector in each year; (iv) the empirical price of each task for White men in each year (shown in Figure 4); (v) the aggregate content of each task for White men in each year.[20] These last two moments help pin down the β_{kt}'s while the first two moments help pin down the A_j's for the market occupations. We compute the mean of squared deviations in each of (i) and (ii), as well as the sum of squared deviation in (iii)-(v), and search for the set of parameter values that minimizes the sum of these numbers.

The Frechet shape parameters θ and ψ are estimated from the average within-occupation variation in log income and the labor supply elasticity for White men, respectively. Intuitively, a smaller θ translates to a higher degree of heterogeneity in skill endowments ϕ_{ik}'s among workers in the same occupation (for given employment shares) and therefore a higher variance in log earnings within each occupation; a smaller ψ translates to stronger occupational preferences (which means workers are less responsive to a change in wages) and hence a lower elasticity of labor supply. We discuss in detail the mapping of the model to these

[20] For the task content of the home sector, we use data from the Census/ACS measuring the individual's last occupation before entering the home sector. We take the average over the years in the sample. However, this normalization plays little role in our main quantitative results given that we allow the A_{gH}'s to match the actual shares in the home sector for White and Black men separately by year.

data moments in the online appendix. Chetty et al. (2013) suggests the extensive margin elasticity of labor supply of about 0.25; the average of the within-occupation variance in log earnings for White men (weighted by employment shares) is about 0.27 in the 1990 Census. Given that some measurement error in the reported earnings is likely to have inflated the variance, we also use the time series of the Mincerian return for *Abstract* tasks to pin down θ. We estimate these shape parameters using the 1990 data and apply the estimates to all years. We choose a value of $\theta = 6$ and a value of $\psi = 4.5$ to roughly match these targets.[21]

In the second step, after we estimate the A_j's, the β_k's, ψ, and θ, we estimate the composite race specific term $\delta_{bkt}^{taste} + \eta_{bkt}$ – the sum of taste-based discrimination and the racial skill gap (η_{bkt}) – for each k. We do so by targeting (i) the conditional racial gaps in aggregate task contents, (ii) the conditional racial gaps in task premiums (for each task k), and (iii) the conditional aggregate wage gap. Specifically, we target (i) the coefficients obtained from regressing $Black_{ijt}$ on task contents τ_{ijt}^k with individual controls for age, education and Census division (Panel B of Figure 2), (ii) the Black-White difference in the Mincerian wage premiums on tasks (Panel B of Figure 4), and (iii) the conditional aggregate wage gap presented in Figure 1. We minimize the weighted sum of squared deviations, where we weight the aggregate task content gaps more heavily than the task price and wage gaps in order to match the sorting pattern closely.

One exception with our estimation strategy is with the *Manual* tasks. Because the empirical wage premium on *Manual* tasks for White men is close to zero, the first step of our model calibration estimates that $\beta_{Manual,t} = 0 \ \forall \ t$. Consequently, the racial barriers $\delta_{bkt}^{taste} + \eta_{bkt}$ for *Manual* tasks contribute neither to overall racial wage gaps nor to sorting given the model structure. Hence, we focus on estimating the η_{bkt}'s and δ_{bkt}'s for *Abstract*, *Contact*, and *Routine* tasks only. We thus exclude the racial gaps in aggregate *Manual* task contents and *Manual* wage premiums from the set of moments we target.

We start by assuming employers observe worker skills without error by setting $\sigma = 0$. We relax this assumption in subsequent sections. However, we note that our results in this section are not overly sensitive to our assumptions about σ.

Appendix Figure A11 compares the key model moments (solid lines) against the corresponding data targets (dashed lines). As seen from the various panels of the figure, our model generally fits the data on racial gaps in tasks and wages very well.

[21] We discuss the robustness of our key results to alternate values of θ and ψ in the online appendix.

Figure 5: Model Estimates of $\eta_{bkt} + \delta_{bkt}^{taste}$ and the Racial Gap in ϕ_{ikt}

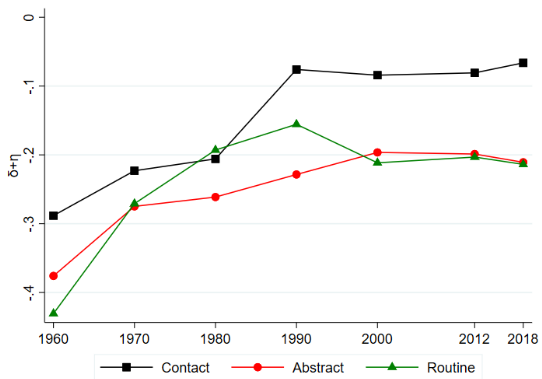

Panel A: Trends in $\eta_{bkt} + \delta_{bkt}^{taste}$

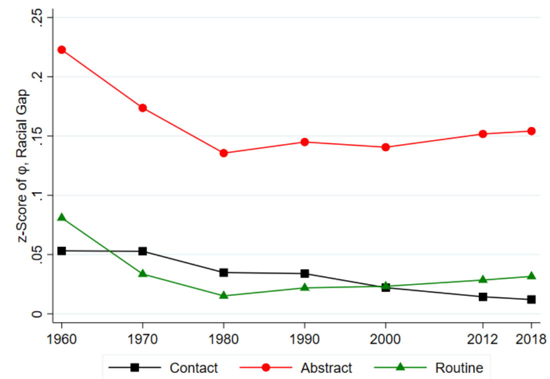

Panel B: Trend in Racial Gap in ϕ_{ikt}

Notes: Panel A of figure shows model generated estimated $\eta_{bkt} + \delta_{bkt}^{taste}$ for *Abstract*, *Contact*, and *Routine* tasks across years. Panel B shows the racial gaps in ϕ_{ikt} defined as the average ϕ for Blacks in task k relative to the average ϕ for Whites in task k during each year t. Positive values imply Blacks have higher ϕ's than Whites. ϕ's measured in z-score units.

5.2 Estimates of Composite Racial Task Differences: $\eta_{bkt} + \delta_{bkt}^{taste}$

Panel A of Figure 5 shows the trend in model implied sum of $\eta_{bkt} + \delta_{bkt}^{taste}$ for Black men in the the *Abstract*, *Contact*, and *Routine* tasks.[22] Given the model, these are the implied racial differences in a combination of mean human capital levels (the η_{bkt}'s) and taste-based discrimination measures (the δ_{bkt}^{taste}'s) for each task. The model says that the combined $\eta_{bkt} + \delta_{bkt}^{taste}$ term explains both racial differences in occupational sorting and racial differences in the returns to task k in year t between Black and White men. As seen in the figure, there was a reduction in the composite term $\eta_{bkt} + \delta_{bkt}^{taste}$ for all three tasks between the 1960s and 2018. In 1960, Black men had a combination of a deficit in mean *Abstract* skills and discrimination in *Abstract* tasks that resulted in 38 percent lower return to *Abstract* tasks than an otherwise comparable White men.[23] That gap fell to 26 percent by 1980 and then only experienced a modest additional decline thereafter. Likewise, the $\eta_{bkt} + \delta_{bkt}^{taste}$ for *Contact* tasks fell from a 28 percent gap in 1960 to a roughly 7 percent gap in 2018.[24]

[22] As a reminder, we calibrate our model to racial gaps in the task content of jobs and racial wage gaps *conditional* on education and age. So the effect of education differences across groups are already controlled for when estimating the η_{bkt}'s and δ_{bkt}^{taste}'s.

[23] The return to tasks discussed here refers to the partial derivative $\partial w_{gij}/\partial \lambda_{jk}$, which equals $(\delta_{gk} + \eta_{gk} + \phi_{ik})$ when employers can measure skills perfectly. Noting that units are in logs, $\eta_{bk} + \delta_{bk}$ represents percentage point gap in the return to tasks between Black and White workers.

[24] Our estimated β_{kt} for *Contact* tasks dips slightly in 1980 as seen in Figure 6 below. This causes the estimated $\eta + \delta^{taste}$ for *Contact* tasks to dip down in 1980. This explains the non-monotonic change in $\eta + \delta^{taste}$ for *Contact* tasks between 1970 and 1990.

Panel B of Figure 5 highlights the evolution of selection over time by plotting model estimates of the reduction in the within-occupation racial skill gaps due to sorting, separately for *Abstract*, *Contact*, and *Routine* tasks in each year. Specifically, the vertical axis measures in standardized units the Black-White gap in average ϕ_{ikt} (individual race-neutral task-specific skill draws) for task k in period t, conditional on task content of occupations. Conceptually, the gap captures how much smaller the racial gap in effective skills $\phi_{ik} + \eta_{bkt}$ becomes when we condition on task requirements. When Black workers face racial barriers (i.e., $\eta_{bkt} + \delta_{bkt}^{taste} > 0$) in task k, they will systematically sort away from occupations requiring this task. However, those Black workers that do sort into jobs that require task k will be positively selected on the race-neutral ϕ_{ikt}'s relative to White workers who select into similar jobs. The differential selection on ϕ_{ik}'s implies that the racial gap in effective skills $\phi_{ik} + \eta_{bkt}$ within occupations with given task requirements (the conditional racial skill gap) will be smaller than the population-wide racial skill gap (the unconditional gap) given by η_{bkt}; the vertical axis in Panel B measures the extent of the reduction (in standardized units).

Three things are of note from the figure. First, for all tasks, Black men are positively selected on their latent ϕ_{ik}'s. Second, the extent of differential selection between Black and White men is greater in *Abstract* tasks relative to both *Contact* and *Routine* tasks. In 1960, the conditional racial skill gap in *Abstract* was about 0.2 standard deviations smaller than the unconditional racial skill gap. The $\eta_{bkt} + \delta_{bkt}^{taste}$ gap in a task has a larger effect on sorting when the Mincerian return to that task is relatively high. A combination of high $\eta_{bkt} + \delta_{bkt}^{taste}$ in *Abstract* tasks and a large return in *Abstract* tasks explains why the selection effects are so much larger in *Abstract* tasks. Finally, the racial gap in ϕ_{ik}'s for *Contact* tasks was large in 1960 but that gap essentially diminished to zero by 2018. The reason that differential selection was close to zero in 2018 for *Contact* tasks is because the composite $\eta_{bkt} + \delta_{bkt}^{taste}$ for *Contact* tasks has fallen close to zero by 2018.

The counterfactuals we explore in subsequent sections rely on the functional form assumptions we make for the various distributions for which individuals draw task specific skills or preferences. To provide confidence in the distributional assumptions underlying our calibrated model, we explore how our model matches the time series trends for many additional non-targeted moments. We leave the full details of these validation exercises to the Appendix Figure A12 and only very briefly summarize them here. First, we show that despite only targeting mean racial wage gaps over time, our calibrated model matches well the racial wage rank gaps empirically documented in Bayer and Charles (2018). Using a sample of all workers regardless of employment status, Bayer and Charles (2018) compute the percentile of the median and 90th percentile of Black men (in the Black men wage distribution) in the wage distribution of White men. Not only does our model match mean racial wage gaps over time, but the model also matches the relative position of Black men in the White men's wage distribution over time at various percentile points. These findings suggest that, even though we are targeting only mean racial wage gaps, our functional form assumptions are sufficient to match the level and trends in relative wages throughout the wage distribution. Second, we show that our model also matches well the empirical time series patterns of racial wage gaps conditional on the task-content of a worker's job. In particular, both in the model and in the data, controlling for the task content of a worker's occupation does not have much of an effect on estimated racial wage gaps in any year. The reason in the model for this result stems from workers with differing task-specific skills and occupational preferences differentially sorting into occupations with larger task-specific barriers. The fact that our model matches various additional non-targeted moments gives us confidence in the counterfactuals we highlight next.

6 The Stagnation of the Racial Wage Gap Post 1980

The prior section discusses our base model estimation and shows that the quantitative model matches many non-targeted moments. We now perform various counterfactuals using the estimated model to shed light on the stagnation of the racial wage gap post-1980.

6.1 Model Counterfactuals

Figure 6 shows the model implied β_{kt}'s for the various task measures. Panel A shows the level of the β's for the different tasks over time while Panel B shows the time series trends in the tasks prices relative to the task price of *Contact* tasks during a given time period, $\beta_{kt}/\beta_{Contact,t}$. As seen in Panel B, we estimate that the task prices for *Abstract* tasks have been rising relative to the task prices of other tasks, particularly after 1980.

In theory, given relatively high and persistent barriers to *Abstract* tasks faced by Black men, the relative increase in the return to *Abstract* tasks will disadvantage Black workers and hence widen the racial wage gap because Black workers are underrepresented in jobs that require these tasks. Panel A of Figure 7 quantifies the extent of this force. Specifically, the black line (with squares) shows what would happen to the racial wage gap relative to the baseline scenario if we hold fixed the task price for *Abstract* tasks ($\beta_{Abstract}$) at 1980 levels.[25] The counterfactual exercise shows that the racial wage gap in 2018 would have fallen by roughly 10 log points relative to the baseline model (from 24 to 14 log points) if the return

[25] We show the results in Figure 7 assuming employers observe worker skills without error ($\sigma = 0$). As we highlight below, our results in time series changes in the combined ($\eta + \delta^{taste}$) for the various tasks are relatively similar across the different signal-to-noise assumptions, especially post-1980.

Figure 6: Task Premium Trends 1980 - 2018

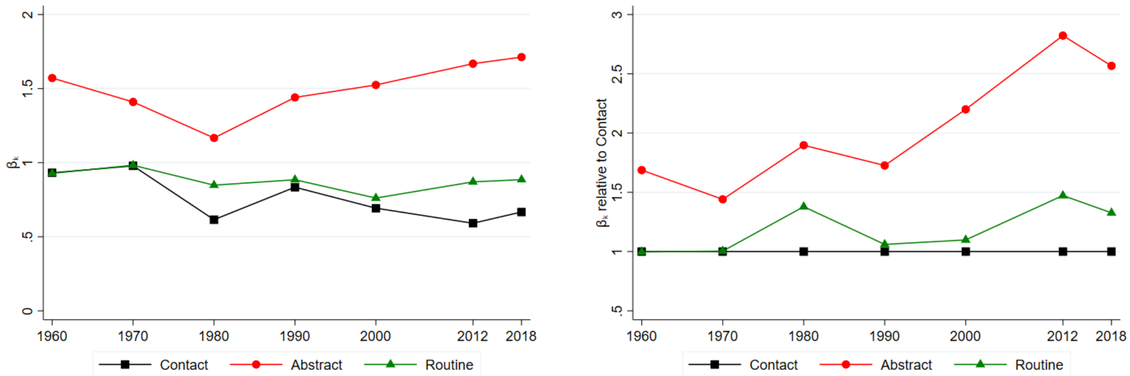

Panel A: β_k's

Panel B: β_k's relative to Contact

Notes: Figure shows trends in model task prices, β_k's; Panel A presents the trends in β_k's as they are, while Panel B normalizes β_k for *Contact* to one and shows the trends in relative values of β_k's.

Figure 7: Counterfactual Racial Wage Gaps Relative to Baseline, 1980 - 2018

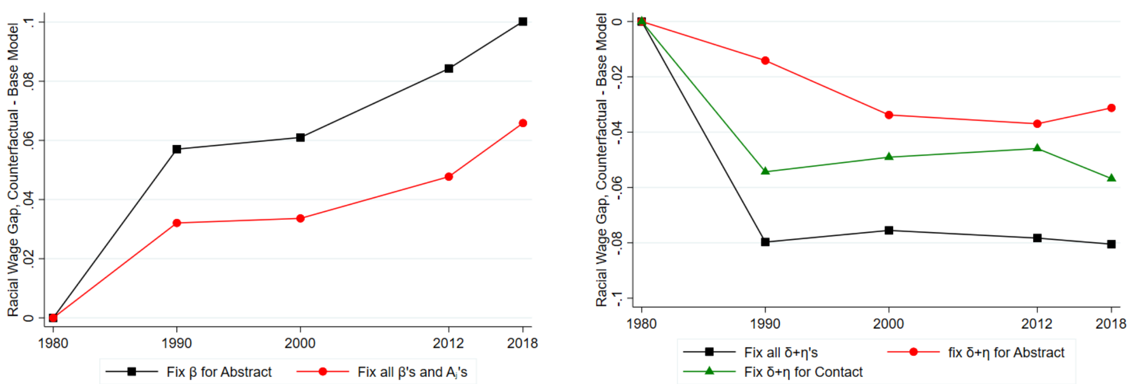

Panel A: Holding Task Returns (β's) Fixed

Panel B: Holding δ's and η's Fixed

Notes: Figure shows counterfactual racial wage gaps relative to the racial wage gaps in the baseline scenario assuming various β_{kt}'s and A_{jt}'s are held fixed at 1980 levels (Panel A) and various $\delta_{kt}^{taste} + \eta_{kt}$'s are held fixed at 1980 levels (Panel B). Both figures show the log differences between the counterfactual and baseline wage gaps in 1990, 2000, 2012, and 2018 under the various counterfactuals.

to *Abstract* tasks were held at 1980 levels. In the red line (with circles) in this panel, we hold the β_{kt}'s for all the tasks and the A_{jt}'s for all occupations at their 1980 levels, rather than just fixing the β_{kt} for *Abstract* tasks. In this case, we find that that the Black-White wage gap in 2018 would have narrowed by 7 percentage points relative to the baseline model.

Panel B of the figure shows the flip side of our analysis. In this set of counterfactual exercises, we hold the composite $\delta_{bkt}^{taste} + \eta_{bkt}$ fixed at 1980 levels for various tasks. As seen from the figure, if $\delta_{bkt}^{taste} + \eta_{bkt}$ for *Contact* tasks remained at 1980 levels, the racial wage gap in 2018 would have risen by 6 log points (from 24 log points to 30 log points). In other words, the decline in $\delta_{bkt}^{taste} + \eta_{bkt}$ for *Contact* tasks reduced the racial wage gap by 6 percentage points since 1980 with almost all of that decline occurring between 1980 and 1990. Likewise, holding the $\delta_{bkt}^{taste} + \eta_{bkt}$ for all tasks at 1980 levels would have increased the racial wage gap in 2018 by 8 percentage points (from 24 to 32 log points).

To summarize, the racial wage gap has remained relatively constant since 1980 because of two offsetting effects. On the one hand, a combination of declining discrimination and a narrowing of racial skill gaps reduced the racial wage gap between 1980 and 2018 by about 8 percentage points. On the other hand, changes in race neutral forces such as the increasing return to *Abstract* tasks (perhaps due to forces such as skill biased technological change) widened the gap by about 10 percentage points during the same period. Because Black workers face discrimination in *Abstract* tasks and have a gap in skills associated with *Abstract* tasks, Black workers were not able to capture as much of the gains from the increasing returns in these activities. These two sets of forces have combined to keep the racial wage gap relatively unchanged between 1980 and 2018.[26]

6.2 Additional Model Validation: Cross Region Wage Gaps

In this subsection, we show additional empirical evidence exploiting cross region variation that supports the model counterfactuals. Our model suggests that during the 1980-2018 period, racial wage gaps should have risen in regions where there were small declines in discrimination. In those regions, the primary effect on racial wage gaps should have been the increasing returns to *Abstract* tasks which relatively favored White workers. As discussed in Section 3, the literature has found that taste-based measures of discrimination were higher in 1980 and declined more between 1980 and 2018 in the South region relative to other regions. Likewise, using our NLSY data, we find that racial gaps in AFQT scores were also larger in the South region in the NLSY79 wave and converged more in the South region than in the non-South region between the NLSY79 and NLSY97 waves. These results in combination suggest that race-specific forces in the South region were large in 1980 and declined more post-1980 than they did in other regions.

Figure 8 shows the male racial wage gap in the South region (black line with squares) and

[26]We have explored the robustness of Figure 7 to alternate values of θ (the shape parameter from the distribution from which task skills are drawn) or ψ (the shape parameter from the distribution from which occupational preferences are drawn). To conserve space, we relegate the details of these robustness exercises to the online appendix. The main takeaways are always qualitatively similar and often quantitatively similar across the different values of θ and ψ we explored.

Figure 8: Racial Wage Gaps by Region: Census/ACS Data

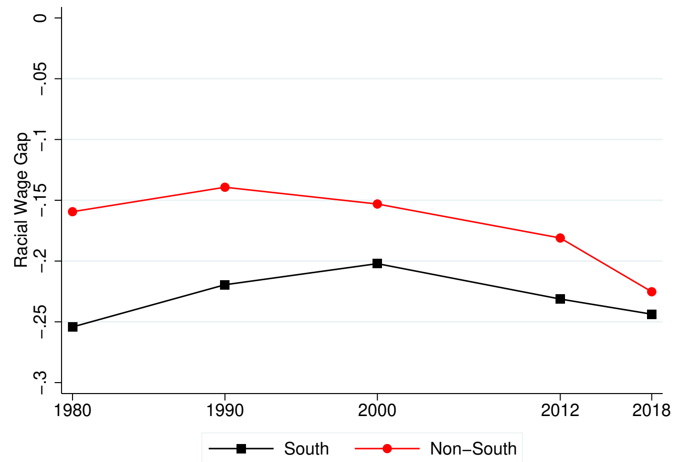

Notes: Figure shows racial wage gaps for men aged 25-54 over time (conditional on age and education) using individual level data from the Census/ACS separately for individuals in the South and non-South regions.

non-South region (red line with circles) using our main Census/ACS samples between 1980 and 2018. The regressions underlying the results in this panel are similar to those shown Figure 1 and are conditioned on education and age controls. The results in in the figure confirm our model's prediction. First, racial wage gaps in the non-South region *increased* by about 7 log points between 1980 and 2018 (from -0.16 to -0.23). This is consistent with the counterfactual findings in Figure 7 that increasing task returns post-1980 increased the racial wage gap in the US as a whole by about 7 log points. In regions where the changes in race-specific factors were small, the only effect on racial wage gaps would be the negative effect stemming from rising task returns. Second, racial wage gaps in the South converged toward the gaps in the non-South region. In 1980, the racial wage gap was about 10 log points larger in the South region relative to the non-south region. By 2018, the racial wage gap in the South was essentially the same as the racial wage gap in the non-South region. Moreover, unlike the racial wage gap in the non-South regions, the racial wage gap in the South region was relatively constant between 1980 and 2018.

Collectively, these cross-region patterns are consistent with the model counterfactual results suggesting that Black workers in the aggregate gained from declining discrimination and narrowing racial skill gaps post-1980, but these gains were masked in aggregate data by rising returns to *Abstract* tasks. If one focuses on regions where the change in race specific factors were small, our model says that racial wage gaps should have increased post-1980; this is exactly what we find in the non-South regions of the United States.

6.3 Model Implications for Reduced Form Wage Regressions

Additionally, our structural model provides a road map to empirical researchers looking to uncover race-specific factors in micro data. In particular, the model suggests that one must control not only for racial differences in skills but also for changes in the return to different skills when analyzing the evolution of Black-White wage differences over time. In this subsection, we exploit the panel structure of the NLSY data to illustrate this point. Specifically, we use our pooled NLSY sample comprising multiple cohorts to run the following regression:

$$w_{it} = \omega^0 + \omega_t^1 D_t Black_i + \sum_k \omega_{kt}^2 D_t S_{ki}^{NLSY} + \Gamma X_{it} D_t + \mu_i + \epsilon_{it} \tag{9}$$

where w_{isjt} is the log wage of individual i from the NLSY in period t and S_{ki}^{NLSY}'s ($k = cog, ncog, soc$) are the pre-labor market measures of cognitive, non-cognitive, and social skills discussed above. Our parameters of interest are the estimated ω_t^1's, which measure the Black-White wage gap relative to the benchmark time-period controlling for differences in pre-labor market skills. Guided by our structural estimation, however, we estimate relative Black progress in log wages after controlling for changing skill returns that can mask this progress. Specifically, when we control for the individual skills, we allow the labor market returns to the various NLSY skills – the regression coefficients on the S_{ki}^{NLSY}'s – to evolve over time; note that the individual skills are interacted with time dummies. This conditions out the effects of changing skill prices over time on the racial wage differences.

In addition to controlling for changing skill returns, we control for omitted time-invariant factors – such as unmeasured skills that are constant within an individual over time – by including individual fixed effects (μ_i). We hence identify the year-specific race dummies (the ω_t^1's) by exploiting within-individual changes over time. We also include demographic controls (X_{isjt}) consisting of age and education dummies again interacted with time dummies. In terms of estimation, we segment the NLSY into four year periods: 1980-1989, 1990-1999, 2000-2009, and 2010-2018. We set the 1980-1989 period to be the benchmark year group so all other differences in the racial wage gap over time are relative to the 1980-1989 period.

The results from the regressions are shown in Table 4. To illuminate the effects of including various controls, we start with specifications inclusive of just the basic controls. Specifically, column 1 shows the evolution of racial wage gaps in the NLSY controlling only for our standard demographics. As with the patterns in the Census/ACS data, the racial wage gap in the NLSY has been roughly constant between the early 1980s and the late 2010s. In column 2, we include individual fixed effects; we still find no trend in racial wage gaps

Table 4: The Evolution of Racial Wage Gaps Over Time in the NLSY: The Importance of Controlling for Time-Varying Skill Returns

	(1)	(2)	(3)	(4)
Racial Wage Gap: 1990s	0.002	0.020	0.040	0.042
	(0.022)	(0.019)	(0.022)	(0.022)
Racial Wage Gap: 2000s	0.051	0.034	0.093	0.094
	(0.028)	(0.032)	(0.035)	(0.035)
Racial Wage Gap: 2010s	0.029	0.026	0.100	0.099
	(0.038)	(0.039)	(0.042)	(0.043)
Demographic Controls * Year Dummies	Yes	Yes	Yes	Yes
Individual Fixed Effects	No	Yes	Yes	Yes
Cognitive Skills * Year Dummies	No	No	Yes	Yes
Non-Cognitive and Social Skills * Year Dummies	No	No	No	Yes

Notes: Table shows the evolution of the racial log wage gap over time in the NLSY data with various sets of controls. Data uses the pooled sample of the NLSY 1979 and 1997 waves. Sample restricted to Black and White men between the ages of 25 and 54. Robust standard errors clustered at the individual level shown in parentheses.

between 1980 and 2018. Omitted time-invariant factors thus cannot explain the stagnation in Blacks' relative wages over the last 40 years. Once we control for the rising return to cognitive skills over time, however, we find a strong convergence in racial wage gaps post-1980. Specifically, in column 3, we control for time-varying return to just cognitive skills. In this column, we find a narrowing of the racial wage gap relative to the 1980s of about 4 log points in the 1990s, about 9 log points in the 2000s, and about 10 log points in the 2010s. The results are nearly identical when we additionally control for time-varying returns to non-cognitive and social skills (column 4).[27] As suggested by our model, conditioning out the effects of time-varying skill returns – the rising return to cognitive skills in particular – unveils the convergence in the racial wage gap due to changing race-specific factors. Strikingly, the magnitude of the convergence we estimate in the NLSY between 1980 and 2018 once properly controlling for the changing returns to skills (column 4 of Table 4) is close to the magnitude we estimate from our structural model (Panel B of Figure 7).

The above findings highlight why our estimated model yields quantitatively different conclusions regarding the extent to which race-specific factors have improved in the United

[27] Although not shown, we estimate that the return to cognitive skills increased by 8 log points between the 1980s and the 2010s while the return to non-cognitive and social skills remained roughly constant.

States during the last forty years relative to a popular statistical decomposition method developed by Juhn et al. (1991) (henceforth known as JMP). In the online appendix, we perform the JMP decomposition on our data from the Census/ACS and show that the decomposition dramatically understates the importance of both skill price changes in widening the racial wage gap and declining race specific factors in narrowing the racial wage gap over the 1980-2018 period relative to our model. This is because the JMP procedure assumes that White workers with a given wage have a similar skill bundle and furthermore perform a similar mixture of tasks as Black workers with the same wage. In our multi-task model with selection, that assumption does not hold; a White worker with a given wage is more likely to have a skill bundle tilted towards *Abstract* (cognitive) skills and also more likely to have sorted into occupations with high *Abstract* task requirement than a Black worker with the same wage. These appendix results highlight the quantitative importance of accounting for selection on multiple skills when decomposing the effect of changing skill prices, declining discrimination or narrowing racial skill gaps on racial wage gaps.

7 Mapping Empirical Measures of Individual Skills to Model Task Specific Skills

At the heart of our model is the idea that the productivity of individual skills varies across tasks; cognitive, non-cognitive and social skills are not equally rewarded in different jobs. While much research has focused on accounting for individual skills in explaining racial wage gaps (e.g., Neal and Johnson (1996)), our framework emphasizes workers' task-specific skills.

We now lay out our procedure mapping individual measures of skills from the NLSY into model-based measures of task-specific skills. Our procedure has two steps. First, restricting ourselves to the sample of White men, we map NLSY measures of cognitive, non-cognitive, and social skills into task-specific skills in the model (up to a scalar) using the following regression:

$$\overline{s}_{wjkt} = a_{kt} + b_{cog,kt}\overline{S}^{NLSY}_{cog,wjt} + b_{ncog,kt}\overline{S}^{NLSY}_{ncog,wjt} + b_{soc,kt}\overline{S}^{NLSY}_{soc,wjt} + \epsilon_{jkt}, \tag{10}$$

where the dependent variable $\overline{s}_{wjkt}$ is the occupational-average of observed task-specific skills s_{wjkt} for White men generated by the model, and the regressors are the empirical measures of average cognitive ($\overline{S}^{NLSY}_{cog,wjt}$), non-cognitive ($\overline{S}^{NLSY}_{ncog,wjt}$) and social skills ($\overline{S}^{NLSY}_{soc,wjt}$) for White men in the corresponding occupations from our sample of NLSY respondents. Intuitively, the first stage produces a weighting (the b's) of NLSY individual skill measures for each task-specific skill by exploiting cross-occupation variation for White men.

In the second stage, we use the estimated weights for White men and the Black-White gap in measured individual skills from the NLSY to impute the racial gaps in task-specific skills in each occupation. Formally, using the coefficients from the first stage regression ($\hat{b}_{cog,kt}$, $\hat{b}_{ncog,kt}$, and $\hat{b}_{soc,kt}$), we predict racial gaps in task-specific skills $\overline{s}^{gap}_{jkt}$ – whose predicted values we denote with $\widehat{s}^{gap}_{jkt}$ – in each occupation based on the racial gaps in the NLSY skills:

$$\widehat{s}^{gap}_{jkt} = \hat{b}_{cog,kt}\overline{S}^{gap}_{cog,jt} + \hat{b}_{ncog,kt}\overline{S}^{gap}_{ncog,jt} + \hat{b}_{soc,kt}\overline{S}^{gap}_{soc,jt}, \tag{11}$$

where $\overline{S}^{gap}_{cog,jt}$, $\overline{S}^{gap}_{ncog,jt}$, and $\overline{S}^{gap}_{scog,jt}$ respectively are the racial gaps in NLSY measures of cognitive, non-cognitive, and social skills in each occupation estimated *conditional on education* (as in columns 3 and 7 of Table 2). Recall that we estimate all race-specific factors ($\delta^{taste}_{bkt}+\eta_{bkt}$'s) controlling for education; we thus estimate the racial gaps in task-specific skills analogously. Once we obtain the NLSY-based predictions, we infer the η_{bkt}'s that make the model-generated $\overline{s}^{gap}_{jkt}$'s consistent with the NLSY-based predicted $\widehat{s}^{gap}_{jkt}$'s. The procedure just ensures the model estimate of racial skill gaps matches the weighted average of the racial gaps in NLSY skills separately for each task.

Once the η_{kt}'s are estimated, the model structure can be used to infer the discrimination parameters. As a reminder, in Section 5 we developed a procedure to identify the combined term ($\delta^{taste}_{bkt} + \eta_{bkt}$) using data on racial gaps in task contents, task prices, and aggregate wages. Given our estimates of η_{bkt}'s using the above procedure and an assumed level of noise in worker skills (σ), we can therefore infer the amount of taste-based discrimination (δ^{taste}_{bkt}'s) faced by each worker in each occupation. In other words, we attribute the residual task-specific barriers facing Black men to taste-based discrimination (δ^{taste}_{bkt}'s) after accounting for racial skill differences (η_{bkt}'s) and possible statistical discrimination (δ^{stat}_{bkt}).[28]

In terms of estimating our first stage regression, we pool together data from multiple years to estimate the $b_{cog,kt}$'s, $b_{ncog,kt}$'s, and $b_{soc,kt}$'s, assuming each of the b_{kt}'s to be constant over time; we do, however, allow the a_{kt}'s to differ across t's.[29] With respect to implementation, we map the model estimates from 1990 to the data for the NLSY-79 cohort; given our age restrictions, 1990 is about the average year of data for the NLSY-79 cohort. Likewise, we map the model estimates from 2012 to the data from the NLSY-97 cohort. When estimating (10) for our first stage regression, we use cross occupational variation aggregating the data to 66 unique broader occupations within each year.[30]

[28]In the actual implementation of our decomposition procedure, the δ^{taste}_{bkt}'s and η_{bkt}'s are estimated jointly given that the second stage of the projection procedure depends on sorting of Black workers across occupations, which in turn is affected by δ^{taste}_{bkt}'s.

[29]In the online appendix, we show a robustness exercise computing all of our key decomposition results when we allow the b's to differ by year.

[30]We use the same broad occupation classification as in Hsieh et al. (2019), which uses the occupational

Given the NLSY data with skill measures do not extend back to 1960, we need to make assumptions about the projection in 1960 if we want to discuss long run trends in δ^{taste}. To this end, we use the fact that the racial task gaps in the South Census region of the U.S. in 1990 were similar to the racial task gaps in the entire U.S. in 1960. Specifically, the demographically adjusted racial gap in *Contact*, *Abstract*, and *Routine* task content of occupations for the U.S. as a whole in 1960 were, respectively, -0.040, -0.031, and -0.051 (see Panel B of Figure 2). The corresponding values for individuals living in the South region in 1990 Census/ACS data were -0.041, -0.045, and -0.044 (see Panel A of Figure 3). Relative to the observed time series trends over the 1960-2018 period, these values are relatively close to the 1960 national levels. Given this, for our 1960 decomposition, we load the average occupational efficiency units in 1960 on the average occupational skill levels of White men in the South in 1990. We then use racial differences in skill levels in the South in 1990 as a proxy for racial skill differences nationally in 1960. Given that the estimated b's are relatively constant over time when we estimate equation (10) separately by year, the first part of our assumption for the 1960 projection is not overly restrictive. The stronger assumption is that the observed racial gap in skills in the NLSY in the South for the 1979 cohort is a good proxy for the racial gap in skills for the country as a whole in 1960. There is some existing empirical support for this assumption. Chay et al. (2009) using data from National Assessment of Educational Progress finds a Black-White gap in standardized cognitive test scores for a nationally representative sample of individuals born between 1953 and 1961 of about -1.25 standard deviations. For male NLSY79 respondents in the South, we find an unconditional AFQT racial gap of about -1.2 standard deviations. The fact that the Black-White gaps in both cognitive test scores and occupational sorting for men in the NSLY79 cohort are roughly similar to the Black-White gaps in cognitive test scores and occupational sorting for the U.S. as a whole in 1960 gives us some confidence in using our imputation procedure to infer 1960 relationships.

The first stage regressions are shown in Table 5. The table reports the first stage mapping for *Abstract* (column 1), *Contact* (column 2) and *Routine* tasks (column 3). Each column reflects the estimates of $b_{cog,kt}$'s, $b_{ncog,kt}$'s, and $b_{soc,kt}$'s from separate regressions of equation (11) for the various tasks. A few things are of note from Table 5. First, cognitive skills are most predictive of the skills required for *Abstract* tasks. Occupations where NLSY workers have high cognitive skills on average are also the same occupations that the model predicts that workers have higher levels of *Abstract* task-specific skills. Second, social skills are positively predictive only of the skills required for *Contract* tasks. Social skills, conditional on cognitive and non-cognitive skills, are unrelated to the skills required for *Abstract* tasks

sub-heading groups provided in the 1990 Census.

Table 5: First Stage Regression of Average Model Task Skills on Average NLSY Individual Skills, Cross-Occupation Variation

	Abstract	*Contact*	*Routine*
Cognitive	0.16	0.04	-0.02
	(0.03)	(0.01)	(0.02)
Non-Cognitive	0.05	0.02	0.01
	(0.03)	(0.02)	(0.03)
Social	-0.02	0.12	-0.10
	(0.05)	(0.03)	(0.03)
Year Fixed Effects	Yes	Yes	Yes
Adj. R-Squared	0.41	0.37	0.07
F-Stat	20.9	10.1	4.6

Notes: Table shows estimate coefficients from first stage regression equation (10) for White men. Each column is a separate regression exploiting cross-occupation variation. We use 66 broad occupation categories. For these regressions, we pool together observations from 1960, 1990, and 2012 so that each regression will have 198 observations (3*66). See the text for additional details.

and are negatively related to the skills required for *Routine* tasks. Third, our first stage procedure has large F-stats for both *Abstract* and *Contact* tasks. However, we have little first stage power predicting *Routine* tasks. As a result, we have less confidence in the ability to use our first stage procedure to decompose the η_{kt} from the δ_{kt}^{taste} for *Routine* tasks. Given this, we focus our main decomposition results on *Abstract* and *Contact* tasks. The patterns in this table are consistent with the reduced form estimates shown in Table 1 which related the NLSY skill measures to the actual task content of occupations for NLSY respondents. Despite these skill measures coming from relatively narrow survey questions in the NLSY, the skill measures are quite predictive of task specific occupational sorting for *Abstract* and *Contact* tasks when viewed through the lens of the model. This predictive power gives us confidence with respecting to performing the decomposition exercises for these tasks below.

We end this section by discussing how any misspecification in decomposition equations (11) and (10) can bias our estimates of the change in η_{bkt} over time. In particular, if there is an omitted trait not measured in the NLSY that predicts an individual's task-based skills, and if that omitted variable changes differentially between Black and White men over time, our estimates of $\Delta\eta_{bkt}$ between two periods will be biased. Both within the main paper and in the appendix, we perform various exercises to assess whether such omitted skills could be an issue. We highlight two such exercises here. First, in Table 4 above, we exploit the panel structure of the NLSY and show that controlling for unmeasured traits by including individual fixed effects hardly affects the estimated changes in the racial wage gap over time (compare columns 1 and 2 of Table 4). This suggests that omitted skills play little role in the evolution of the racial wage gap over the last forty years. Second, in the appendix, we examine whether the labor market returns to skills differ between Black and White men in the NLSY. We find that the labor market returns to social skills are similar between Black and White men. This finding is consistent with there being no differential bias between Black and White men with respect to predicting *Contact* task efficiency from measured traits. On the other hand, consistent with the findings in Neal (2006), the wage return to cognitive skills is higher for Black men than for White men with the same occupation and education. This is suggestive of the possibility that missing traits associated with *Abstract* tasks differ systematically between Black and White men.

Overall, these results give us some confidence that changing racial gaps in omitted skills are not biasing our estimates of the $\Delta\eta_{bkt}$ and $\Delta\delta^{taste}_{bkt}$ for *Contact* tasks. This is crucial because most of our key model findings in the next section hinge on our estimates of the $\Delta\eta_{bkt}$ and $\Delta\delta^{taste}_{bkt}$ for *Contact* tasks being unbiased. It being a pivotal concern for our paper, we will later revisit the possibility of a bias in predicting *Contact* task efficiency and provide some additional reassurance by exploiting state-level survey-based measures of taste-based discrimination. On the other hand, we are less certain about our ability to estimate the $\Delta\eta_{bkt}$ for *Abstract* tasks without a bias. We will acknowledge this concern when we decompose the change in $\delta^{total}_{bkt} + \eta_{bkt}$ for *Abstract* tasks.

8 Estimates of Task-Based Discrimination

Given the mapping procedure described above, we can use the observed racial gaps in skills from the NLSY to inform our decomposition of racial gaps in $\eta_{bkt} + \delta^{total}_{bkt}$ into its various components. The results of our decomposition are shown in Table 6. As the decomposition depends on the value of σ – the variance of noise in skills observed by employers – we perform the decomposition at three levels of strength of signal ratios: 1.0, 0.9, and 0.75. These ratios correspond to values of σ equal to 0.0, 0.17, and 0.30, respectively. We measure the strength of the signal as the variance of the truth divided by the variance of the truth plus the variance of the noise; thus a strength of signal ratio of 1 implies that the variance of the noise is zero. In first set of rows, we explore the model predictions for δ^{total}_{kt} assuming the various levels of σ. Columns (1)-(3) show the estimates for *Contact* tasks in 1960, 1990, and 2012 while columns (5)-(7) show the estimates for *Abstract* tasks. Columns (4) and (8)

Table 6: Task Decomposition of Racial Skill Gap and Task-Based Discrimination

	Contact				*Abstract*			
	1960	1990	2012	$\Delta(12-60)$	1960	1990	2012	$\Delta(12-60)$
δ^{total}, $\sigma = 0$	-0.20	-0.01	-0.02	0.18	-0.10	0.00	-0.03	0.07
δ^{total}, $\sigma = 0.17$	-0.22	-0.02	-0.03	0.18	-0.18	-0.06	-0.07	0.11
δ^{total}, $\sigma = 0.30$	-0.22	-0.03	-0.04	0.18	-0.20	-0.08	-0.09	0.12
δ^{taste}, $\sigma = 0$	-0.20	-0.01	-0.02	0.18	-0.10	0.00	-0.03	0.07
δ^{taste}, $\sigma = 0.17$	-0.19	0.00	-0.02	0.18	-0.08	0.02	-0.02	0.06
δ^{taste}, $\sigma = 0.30$	-0.18	0.00	-0.01	0.17	-0.05	0.04	0.00	0.05
δ^{stat}, $\sigma = 0$	—	—	—	—	—	—	—	—
δ^{stat}, $\sigma = 0.17$	-0.02	-0.02	-0.02	0.01	-0.10	-0.08	-0.05	0.05
δ^{stat}, $\sigma = 0.30$	-0.04	-0.03	-0.03	0.01	-0.15	-0.12	-0.08	0.07
η, $\sigma = 0$	-0.09	-0.07	-0.06	0.03	-0.28	-0.23	-0.17	0.11
η, $\sigma = 0.17$	-0.09	-0.07	-0.06	0.03	-0.29	-0.24	-0.17	0.11
η, $\sigma = 0.30$	-0.10	-0.08	-0.07	0.03	-0.30	-0.25	-0.18	0.12

Notes: Table shows model decomposition of racial differences in δ^{total}_{bk}, δ^{taste}_{bk}, δ^{stat}_{bk}, and η_{bk} for *Abstract* and *Contact* task in 1960, 1990, and 2012. We do the decomposition across various assumptions with respect to which employers can accurately observe the skill levels of their employees. σ = 0, 0.17, and 0.30 corresponds to implied strength of signal ratios of worker skills by employers of 1, 0.9, and 0.75, respectively.

show the changes in the respective δ^{total}_{kt}'s between 1960 and 2012. The next two sets of rows of Table 6 show the results separately for δ^{taste}_{kt} and δ^{stat}_{kt}. The final set of rows shows the results for η_{kt}. Notice, the sum $(\delta^{taste} + \eta)$ is stable across different values of σ assumed for each task. This is because $(\delta^{taste} + \eta)$'s determine for each task the size of task-specific racial barriers unconditional on sorting, which we discipline with the targeted moments.

8.1 Decomposing Racial Gaps in *Contact* Tasks

A few key results are notable for the decomposition for *Contact* tasks. First, for *Contact* tasks, almost all of the δ^{total} can be attributed to δ^{taste} in each of the years regardless of the amount of noise we assume. This result stems directly from Proposition 1, which states that when racial gaps in skills associated with a task are close to zero, statistical discrimination in that task must be close to zero. Specifically, we find only small racial gaps in social skills in the NLSY data, where *Contact* task skills in the model are predominantly associated with social skills (Table 5). The fact that there are only very small racial gaps in the skills associated with *Contact* tasks implies that statistical discrimination cannot be that important for *Contact* tasks, as seen in the δ^{stat} rows of Table 6. Second, to the extent that there are small changes in our level estimates of δ^{taste}, δ^{stat}, and η for *Contact* tasks in any one year as we change the level of σ, the estimated *change* in the discrimination and racial skill gap parameters for *Contact* tasks are essentially invariant to our assumption about how well employers observe a worker's skills. When we decompose the time series trends in the racial gaps in *Contact* tasks, we explain essentially all of the trends with declining δ^{taste}. Hence, the large inferred decline in taste-based discrimination associated with *Contact* tasks between 1960 and 2012 – with essentially all of these gains occurring prior to 1990 – closely mirrors the trend in racial gaps in *Contact* tasks. We thus conclude that changes in the racial gaps in *Contact* tasks over time are a good proxy for changes in taste-based discrimination measures over time. Again, this stems from the fact that racial gaps in social skills in the NLSY are small and social skills predict *Contact* task efficiency.

To further validate our conclusion that racial gaps in *Contact* tasks is a good proxy for taste-based discrimination, we compare racial gaps in *Contact* tasks to survey based measures of taste-based discrimination created by Charles and Guryan (2008), exploiting cross-state variation in these measures. Charles and Guryan (2008) (henceforth CG) use confidential location data from the General Social Survey (GSS) conducted during the 1970s through the early 1990s. The GSS asked a nationally representative sample dozens of questions eliciting potential prejudice against Blacks.[31] Focusing on a sample of White individuals, CG create measures of state level prejudice against Blacks.[32] Their measure is standardized with higher values indicating larger levels of taste-based discrimination among Whites within the state.

Panel A of Figure 9 correlates measures of racial gaps in the *Contact* tasks for each state with the CG state-level taste-based discrimination measures. Specifically, for each state we measure the conditional race gap in *Contact* tasks using the specification in equation (1). Given the GSS was conducted in the mid-1970s through the early 1990s, we map the CG measures to our 1980 data. As seen from the figure, there is a strong correlation between the state-level racial gaps in the *Contact* task content of jobs in 1980 and the CG measure of state-level taste-based discrimination; a simple regression line through the scatter plot yields a slope coefficient of -0.11 (standard error = 0.02) and a R-squared of 0.52. That is, states with high survey-based measures of taste-based discrimination are systematically the

[31]For example, respondents were asked how they would feel if a close relative was planning to marry someone who was Black, whether they would ever vote for a Black president, or whether they were in favor of laws restricting interracial marriage.

[32]Charles and Guryan (2008) produce measures of the average level of discrimination in the state as well as the discriminatory preferences of the marginal individual. We use their average measure in our work below, but the results are very similar using their marginal measure.

Figure 9: Racial Gaps in *Contact* and *Abstract* Tasks vs Survey Measures of Taste-Based Discrimination, State Level Variation

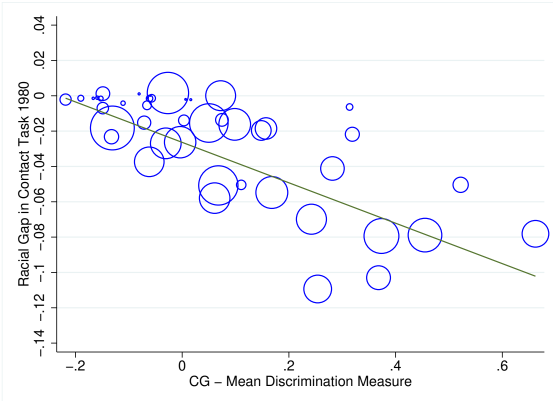

Panel A: Contact Tasks

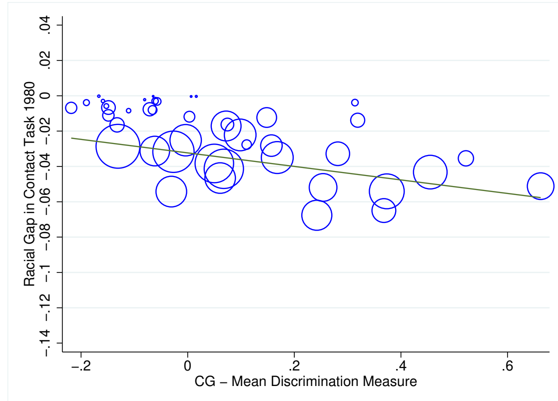

Panel B: Abstract Tasks

Notes: Figure shows state-level conditional racial gaps in the *Contact* task content of jobs (Panel A) and the *Abstract* task content of jobs (Panel B) against the Charles-Guryan mean measures of state level prejudice. Racial gaps in the task content of jobs measured using the 1980 U.S. Census. Gaps are conditioned on age and education as in equation (1). Each observation is a U.S. state with the size of circle measuring the number of Black individuals in the state in the 1980 Census.

states with a larger racial gap in *Contact* task content of jobs.

Panel B, on the other hand, illustrates the relationship between the CG measures of taste-based discrimination and state-level gaps in *Abstract* tasks. As seen from this figure, the relationship between survey-based measures of taste-based discrimination and the racial gap in *Abstract* tasks is much weaker than the relationship with the racial gap in *Contact* tasks. In particular, the simple regression line has a slope coefficient of -0.04 (standard error = 0.01) and a R-squared of 0.25. Consistent with our model findings, racial gaps in *Contact* tasks are much more predictive of taste-based measures of discrimination than are *Abstract* tasks. Collectively, these results provide some support for our finding that changes in the racial gaps in *Contact* tasks are informative measures of changing taste-based discrimination.

8.2 Decomposing Racial Gaps in *Abstract* Tasks

We now turn to discussing the decomposition of the racial barriers to *Abstract* tasks shown in Table 6. First, relative to *Contact* tasks, racial skill gaps are much more important in explaining the differential sorting of Black men into occupations requiring *Abstract* tasks. For all measures of noise and in all years, η_{kt} for *Abstract* tasks is at least three times as large as δ_{kt}^{taste}. Second, consistent with the results from the NLSY discussed above, the racial gap in skills associated with *Abstract* tasks narrowed from 1990 to 2012. However, racial

Figure 10: Counterfactual Racial Gap in *Abstract* Tasks 1980 - 2018

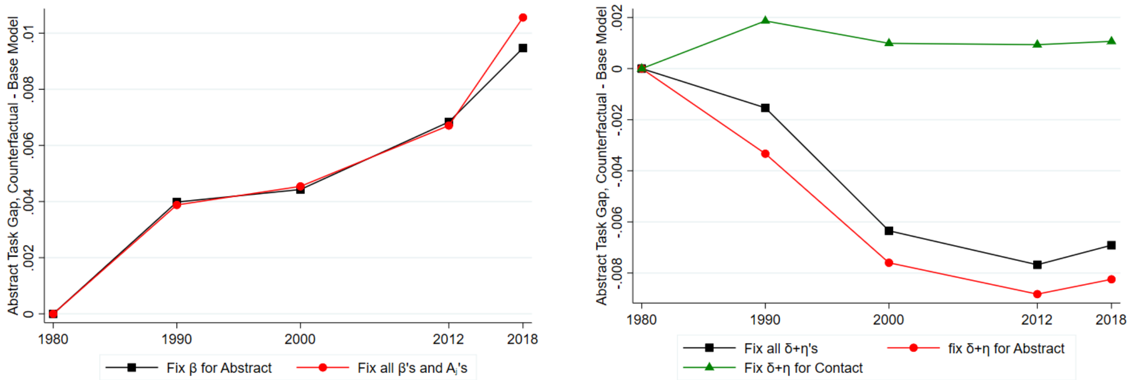

Panel A: Holding Task Returns (β's) Fixed

Panel B: Holding δ's and η's Fixed

Notes: Figure shows counterfactual racial task gaps assuming various β_{kt}'s and A_{jt}'s are held fixed at 1980 levels (Panel A) and various $\delta_{kt}^{taste} + \eta_{kt}$'s are held fixed at 1980 levels (Panel B). Both figures show the percentage point change in the racial task gaps in 1990, 2000, 2012, and 2018 – relative to 1980 – from the various counterfactuals. See the text for additional details.

skill gaps still remain large as of 2012. These results are also robust to the level of noise assumed. Finally, as a result of the large racial skill gaps, even a little bit of noise on the part of employers measuring skills results in a sizeable amount of statistical discrimination. For example, when we assume the strength of the signal is only 0.9 ($\sigma = 0.17$), the majority of δ^{total} for *Abstract* tasks in all years can be attributed to δ^{stat}. Given that skills are almost certainly measured with some error by employers, these results suggest that statistical discrimination against Black men in *Abstract* tasks is still a prominent feature of the U.S. labor market.

Figure 10 uses the model to isolate how the racial gap in *Abstract* tasks would have evolved if task returns were held fixed (Panel A) and if various η's and δ's were held fixed (Panel B). This figure assumes employers observe worker skills without error. Like the wage gaps, the relatively constant racial gap in *Abstract* tasks over time is the result of two offsetting effects. First, the increase *Abstract* task returns since 1980 favored Whites who had a comparative advantage in *Abstract* tasks. By holding task returns fixed, the racial gap in *Contact* tasks would have narrowed by about 1 percentage points between 1980 and 2018. This is due to the fact that, as discussed above, racial skill gaps narrowed during this period. Conversely, if taste-based discrimination and racial skill gaps were held fixed, the racial gap in *Abstract* tasks would have increased. As the relative return to *Abstract* skills increased, it drew relatively more White men into occupations requiring *Abstract* tasks given that White men have comparative advantage in these tasks. This latter effect resulted in the overall racial gap in *Abstract* tasks being relatively constant over the last forty years despite a narrowing of the racial gap in *Abstract* skills ($\eta_{Abstract,t}$) during this period.

8.3 Further Decomposing Trends in Racial Wage Gaps

In Table 7, we use the model to assess how much of the *change* in the racial wage gap between two periods can be attributed to changes in various model driving forces. Panel A performs our counterfactuals assuming no noise in the extent to which skills are measured by employers ($\sigma = 0$). Panel B performs our counterfactuals assuming a strength of signal of 0.75 ($\sigma = 0.30$). The counterfactuals are performed as follows. Consider the first entry in row 1 of Panel A. For this entry, we first compute the counterfactual wage gap trend over the 1960-1990 period fixing δ^{taste} for *Abstract* and *Contact* tasks at the mean of the 1960 and 1990 levels. We then take the difference between the counterfactual and actual change in the racial wage gap over the period. Taking the difference between the counterfactual and actual changes in the racial wage gap gives us a measure of the change in the racial wage gap that can be attributed to those factor that were held fixed. For the second entry in the same row, we compute the counterfactual wage gap change similarly for the 1990-2012 period, sum the estimated counterfactual changes over the 1960-1990 and 1990-2012 periods, and subtract the actual change over the 1960-2012 periods. Other rows are calculated analogously while fixing different sets of parameters.

According to our fully estimated model, nearly 60 percent of the decline in the racial wage gap between 1960 and 1990 can be attributed to declining taste-based discrimination associated with *Abstract* and *Contact* tasks. The fraction attributed to declining taste-based discrimination for these two tasks during the 1960-1990 period is relatively unchanged with higher levels of σ. Over the longer 1960-2012 period, declining taste-based discrimination for *Contact* and *Abstract* tasks contribute to about 40 percent of the decline in the racial wage gap. Declining taste-based discrimination contributed to a smaller extent during the longer 1960-2012 period because a narrowing of racial skill gaps in *Contact* and *Abstract* tasks explained more of the decline in the racial wage gap during the 1990-2012 period. As seen from Table 6, all of the decline in η occurred within *Abstract* tasks. The narrowing of racial gaps in skills associated with *Abstract* tasks can explain nearly 45% of the declining racial wage gap over the 1960-2012 period. The narrowing of racial skills help to offset the widening of the racial wage gap post-1980 resulting from rising returns to *Abstract* tasks.

The table also highlights counterfactuals inclusive of the changing δ's and η's for *Routine* tasks. As seen from Table 5, our first stage regression needed for the decomposition of *Routine* tasks lacks power. That is a reason we do not focus on the *Routine* decomposition in Table 6.

Table 7: Contribution of Various Forces to Changing Racial Wage Gaps Over Time

	1960-1990	1960-2012
Panel A: $\sigma = 0$		
δ_{bkt}^{taste} for *Abstract* and *Contact* tasks	57%	48%
η_{bkt} for *Abstract* and *Contact* tasks	11%	33%
δ_{bkt}^{taste} for all tasks	98%	85%
η_{bkt} for all tasks	13%	36%
A_{jt}'s and β_{kt}	-2%	-12%
Panel B: $\sigma = 0.30$		
δ_{bkt}^{taste} for *Abstract* and *Contact* tasks	56%	43%
η_{bkt} for *Abstract* and *Contact* tasks	13%	38%
δ_{bkt}^{taste} for all tasks	98%	79%
η_{bkt} for all tasks	15%	41%
A_{jt}'s and β_{kt}	-2%	-11%

Note: Table shows how much of the change in the racial wage gap between 1960 and 1990 (column 1) and between 1960 and 2012 (column 2) can be can be attributed to changes in in taste-based discrimination (δ_{bkt}^{taste}), the narrowing of racial skill gaps (η_{bkt}), or a combination of changing race-neutral occupation-specific and task-specific returns (the A_{jt}'s and β_{kt}'s) during the period. To perform these counterfactuals, we set the designated variables fixed at their average level over the period and allow all other variables to change over the time period. We do this calculation separately for the 1960 to 1990 period and again during the 1990 to 2012 period. We then compare the actual change in the wage gap during the period to the counterfactual model predicted change in the wage gap during the period to estimate how much of the change in the wage gap can be attributed to each factor. Column 2 sums together the results from the 1960 to 1990 period and the separate results from the 1990 to 2012 period. Panel A performs the counterfactual assuming skills are observed without error ($\sigma = 0$) while Panel B assumes $\sigma = 0.30$.

The decomposition from that first stage regression implies that essentially all of the decline in the estimated gap in $\delta^{taste} + \eta$ for *Routine* tasks can be attributed to declining δ^{taste} for *Routine* tasks. While there is almost certainly noise in this decomposition, attributing all of the declining task gap for *Routine* tasks to declining δ^{taste} allows us to provide an upper-bound on the importance of declining overall taste-based discrimination for all tasks in contributing to the declining racial wage gaps during the period.[33]

[33] Note, the sum of rows 3-5 in each panel sum to over 100%. This is because there are non-trivial co-variances between the various model driving forces.

9 Conclusion

In this paper, we present new facts about differences in the extent to which Black and White men sort into occupations that require different tasks and how those differences have evolved over time. We then develop a unified framework of occupational sorting that jointly incorporates notions of taste-based and statistical discrimination alongside group differences in skills and changing returns to labor market tasks in order to explain the evolution of wage gaps across groups. Using detailed micro data on racial differences in occupational sorting, we calibrate and estimate the model so as to explain changes in racial wage gaps between prime age Black and White men in the United States since 1960.

There are two important results highlighted in the paper. First, our paper provides an explanation for both why the Black-White wage gap narrowed between the 1960s and 1970s and why the racial wage gap has remained constant thereafter. We find that the constant wage gap between Black and White men post-1980 is due to two offsetting effects. Both the racial skill gap narrowed and taste-based discrimination fell post-1980 resulting in the wages of Black men converging to those of White men, all else equal. However, during the same period, the return to *Abstract* skills rose disadvantaging Blacks relative to Whites. This latter effect resulted in increasing racial wage gaps during the 1980-2018 period. The magnitude of these two effects were roughly similar resulting in a relatively constant racial wage gap post-1980. On the other hand, we show that the relative wage gains of Black men relative to White men during the 1960 to 1980 period stemmed solely from declining discrimination and a narrowing of racial skill gaps; changing task prices did not undermine any of these gains during this earlier period. We also provide a road map to empirical researchers looking to uncover changing race specific factors in micro data.

Second, our paper establishes that the declining racial gap in *Contact* tasks between 1960 and 2018 is a good proxy for declining taste-based discrimination during this period. We motivated the introduction of this novel task measure by conjecturing *ex-ante* that occupations which require many interactions with others are more likely to be susceptible to taste-based discrimination; our model and data work confirm this conjecture *ex-post*. Specifically, the fact that there are very small racial gaps in social skills – combined with the fact that measures of pre-labor market social skills in the NLSY are highly predictive of subsequent entry into occupations that require *Contact* tasks – implies that racial gaps in *Contact* tasks must stem almost entirely from taste-based discrimination and very little from racial skill differences or statistical discrimination. Our model thus implies that the changes in racial gaps in *Contact* tasks over time is a good proxy for taste-based discrimination. To further provide evidence for this conclusion, we document that state-level racial gaps in *Contact* tasks correlate strongly with state-level survey measures of taste-based discrimination, while state-level racial gaps in *Abstract* tasks correlate with them only weakly.

Related, a growing body of influential research uses an individual's name to proxy for their group affiliation when estimating labor market discrimination in experimental and observational data (Bertrand and Mullainathan (2004); Rubinstein and Brenner (2014)). Going forward, this literature can leverage the insights in our paper to further separate between taste-based and statistical discrimination through examining heterogeneous effects by the *Contact* and *Abstract* task intensity of the positions.

As pointed out by Spriggs (2020), for too long, economists attribute too much of the racial disparities in labor market outcomes to racial differences in skills and to the potential statistical discrimination that can result when skills are noisily observed by employers. While racial differences in parental background, school quality, and neighborhood sorting have had a notable impact on racial gaps in people's ability to earn, Spriggs (2020) argues that any racial gap in earnings stemming from such differences are the result of man-made discrimination. Our paper uses a task-based sorting framework to try to directly identify the existence of taste-based discrimination. In doing so, we highlight the importance of taste-based discrimination in explaining why Black workers have earned less than White workers in the United States during the last sixty years. Our estimated model finds that modelling taste-based discrimination is important also for understanding the evolution of the racial wage gap over the last half century. Quantitatively, we find that declining taste-based discrimination explains at least half of the declining racial wage gap between 1960 and 2018.

While there was a narrowing in racial skill gaps over time, we estimate that large racial skill gaps remain. We want to stress that these racial gaps in skills are themselves endogenous and subject to discrimination. Current or past levels of taste-based discrimination are almost certainly responsible for Black-White differences in measures of cognitive test scores. Such caveats should be kept in mind when trying to segment current racial wage gaps into parts due to taste-based discrimination and parts due to differences in market skills. To the extent that we identify taste-based discrimination as being an important barrier to labor market equality between Black and White workers, these estimates should be viewed as a lower bound given that the racial skill gaps themselves stem from past racial prejudice. However, we also wish to stress that regardless of the reason for the racial skill gaps associated with a given task, the existence of such gaps imply that changes in task returns can have meaningful effects on the evolution of racial wage gaps. Our paper highlights that it is becoming even more important today to equalize opportunities in early childhood to close the racial *Abstract* skill gap given that the return to *Abstract* skills has been rising over time.

References

Acemoglu, Daron and David Autor, "Skills, Tasks and Technologies: Implications for Employment and Earnings," in "Handbook of Labor Economics, Volume 4. Amsterdam: Elsevier-North" 2011, pp. 1043–1171.

Aigner, Dennis J. and Glen G. Cain, "Statistical Theories of Discrimination in Labor Markets," *ILR Review*, 1977, *30* (2), 175–187.

Almond, Douglas and Janet Currie, *Human Capital Development Before Age Five*, Vol. 4B of *Handbook of Labor Economics*, Elsevier, 2011.

Altonji, Joseph G. and Rebecca M. Blank, "Chapter 48 Race and gender in the labor market," in "in," Vol. 3 of *Handbook of Labor Economics*, Elsevier, 1999, pp. 3143 – 3259.

_ , Prashant Bharadwaj, and Fabian Lange, "Changes in the Characteristics of American Youth: Implications for Adult Outcomes," *Journal of Labor Economics*, 2012, *30* (4), 783–828.

Arrow, Kenneth, *The Theory of Discrimination*, Princeton University Press, 1973.

Autor, David and David Dorn, "The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market," *American Economic Review*, August 2013, *103* (5), 1553–97.

_ and Michael Handel, "Putting Tasks to the Test: Human Capital, Job Tasks, and Wages," *Journal of Labor Economics*, 2013, *31* (S1), S59 – S96.

_ , Frank Levy, and Richard Murnane, "The Skill Content of Recent Technological Change: An Empirical Exploration," *The Quarterly Journal of Economics*, 2003, *118*, 1279–1333.

Bayer, Patrick and Kerwin Kofi Charles, "Divergent Paths: A New Perspective on Earnings Differences Between Black and White Men Since 1940," *The Quarterly Journal of Economics*, 01 2018, *133* (3), 1459–1501.

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

Bertrand, Marianne and Sendhil Mullainathan, "Are Emily and Greg More Employable Than Lakisha and Jamal? A Field Experiment on Labor Market Discrimination," *American Economic Review*, September 2004, *94* (4), 991–1013.

Bobo, Lawrence D., Camille Z. Charles, Maria Krysan, and Alicia D. Simmons, "The Real Record on Racial Attitudes," in Peter V. Marsden, ed., *Peter V. Marsden, ed.*, Vol. 1 of *Social Trends in American Life: Findings from the General Social Survey Since 1972*, Princeton University Press, 2012, pp. 38–83.

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

Chandra, Amitabh, "Labor-Market Dropouts and the Racial Wage Gap: 1940-1990," *American Economic Review*, May 2000, *90* (2), 333–338.

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

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

Chetty, Raj, Adam Guren, Dayanand Manoli, and Andrea Weber, "Does Indivisible Labor Explain the Difference between Micro and Macro Elasticities? A Meta-Analysis of Extensive Margin Elasticities," *NBER Macroeconomics Annual*, 2013, *27* (1), 1 – 56.

_ **, John N. Friedman, Nathaniel Hilger, Emmanuel Saez, Diane Whitmore Schanzenbach, and Danny Yagan**, " How Does Your Kindergarten Classroom Affect Your Earnings? Evidence from Project Star *," *The Quarterly Journal of Economics*, 11 2011, *126* (4), 1593–1660.

Coate, Stephen and Glenn C. Loury, "Will Affirmative-Action Policies Eliminate Negative Stereotypes?," *The American Economic Review*, 1993, *83* (5), 1220–1240.

Deming, David J., "The Growing Importance of Social Skills in the Labor Market," *Quarterly Journal of Economics*, 2017, *132* (4), 1593–1640.

Dickens, William and James Flynn, "Black Americans Reduce the Racial IQ Gap Evidence From Standardization Samples," *Psychological science*, 11 2006, *17*, 913–20.

Donohue, John J. and James Heckman, "Continuous Versus Episodic Change: The Impact of Civil Rights Policy on the Economic Status of Blacks," *Journal of Economic Literature*, 1991, *29* (4), 1603–1643.

Dorn, David, "Essays on Inequality, Spaital Interaction, and the Demand for Skill," Dissertation, University of St. Gallen Dissertation 2009.

E., J. Heckman J. Cameron Knudsen and J. Shonkoff, "Economic, Neurobiological, and Behavioral Perspectives on Building America's Future Workforce," *Proceedings of the National Academy of Sciences*, 2006, *103*, 10155–10162.

Freeman, Richard B., "Changes in the Labor Market for Black Americans, 1948-72," *Brookings Papers on Economic Activity*, 1973, *20*, 67–120.

Heckman, James J., "Schools, Skills and Synapses," *Economic Inquiry*, 2008, *46*, 289–324.

_ **and Tim Kautz**, "Hard Evidence on Soft Skills," *Labour Economics*, 2012, *19*, 451–464.

_ **, Jora Stixrud, and Sergio Urzua**, "The effects of cognitive and noncognitive abilities on labor market outcomes and social behavior," *Journal of Labor economics*, 2006, *24* (3), 411–482.

_ , **Thomas M. Lyons, and Petra E. Todd**, "Understanding Black-White Wage Differentials, 1960-1990," *American Economic Review*, May 2000, *90* (2), 344–349.

Hsieh, Chang-Tai, Erik Hurst, Charles Jones, and Peter Klenow, "The Allocation of Talent and U.S. Economic Growth," *Econometrica*, 2019, *87* (5), 1439–1474.

Juhn, Chinhui, K. Murphy, and B. Pierce, "Accounting for the Slowdown in Black-White Wage Convergence," in Marvin Kosters, ed., *Workers and Their Wages: Changing Patterns in the United States*, American Enterprise Institute, 1991.

Krysan, Maria and Sarah Moberg, "Trends in Racial Attitudes," Working Paper, University of Illinois Institute of Government and Public Affairs August 2016.

Lang, Kevin and Jee-Yeon K. Lehmann, "Racial Discrimination in the Labor Market: Theory and Empirics," *Journal of Economic Literature*, December 2012, *50* (4), 959–1006.

Levine, Ross and Yona Rubinstein, "Smart and illicit: who becomes an entrepreneur and do they earn more?," *The Quarterly Journal of Economics*, 2017, *132* (2), 963–1018.

Murray, Charles, "The magnitude and components of change in the Black–White IQ difference from 1920 to 1991: a birth cohort analysis of the Woodcock–Johnson standardizations," *Intelligence*, 07 2007, *35*, 305–318.

Neal, Derek, "Chapter 9 Why Has Black–White Skill Convergence Stopped?," in E. Hanushek and F. Welch, eds., *E. Hanushek and F. Welch, eds.*, Vol. 1 of *Handbook of the Economics of Education*, Elsevier, 2006, pp. 511–576.

_ **and William Johnson**, "The Role of Premarket Factors in Black-White Wage Differences," *Journal of Political Economy*, 1996, *104* (5), 869–895.

Phelps, Edmund S., "The Statistical Theory of Racism and Sexism," *The American Economic Review*, 1972, *62* (4), 659–661.

Roy, A. D., "Some Thoughts on the Distribution of Earnings," *Oxford Economic Papers*, 06 1951, *3* (2), 135–146.

Rubinstein, Yona and Dror Brenner, "Pride and Prejudice: Using Ethnic-Sounding Names and Inter-Ethnic Marriages to Identify Labour Market Discrimination," *The Review of Economic Studies*, 1 2014, *81* (1), 389–425.

Ruggles, Steve, Sarah Flood, Ronald Goeken, Josiah Grover, Erik Meyer, Jose Pacas, and Matthew Sobek, "IPUMS USA: Version 10.0 [dataset]," 2021. https://doi.org/10.18128/D010.V10.0.

Smith, James P. and Finis R. Welch, "Black Economic Progress After Myrdal," *Journal of Economic Literature*, 1989, *27* (2), 519–564.

Spriggs, William, "Is Now a Teachable Moment for Economists?," *An Open Letter to the Economics Profession*, 2020.

Appendix A Data Description

In our empirical work, we primarily use data from three sources: cross-sectional labor market data from the Census/ACS, occupational task measures from DOT and O*Net, and panel micro data from the NLSY79 and NLSY97 that contain measures of worker pre-labor market skills.

Appendix A.1 Census/ACS Sample

To access the Census/ACS data, we download the micro data directly from the IPUMS USA website (Ruggles et al. (2021)). We use data from the 1960, 1970, 1980, 1990, and 2000 US Censuses. Additionally, we pool together data from the 2010-2012 and the 2016-2018 American Community Surveys. We refer to the former as the 2012 ACS sample and the latter as the the 2018 ACS. We restrict our Census and ACS samples to those between the ages of 25 and 54 (inclusive), those who report their race as "White" (race = 1) or "Black" (race = 2), and those born within the United States. We exclude from our sample anyone who is living in group quarters (keep gq = 1) and those who are self employed (classwkr = 2). Finally, we exclude any employed worker whose occupation has missing task values. This last restriction reduces the overall sample by less than one percent.

Appendix A.2 NLSY Data

Our primary data sources for individual skills are the 1979 and the 1997 National Longitudinal Survey of Youth, NLSY79 and NLSY97, respectively. The NLSY79 is a representative survey of 12,686 individuals who were 15-22 years old when they were first surveyed in 1979. Individuals were interviewed annually through 1994 and biennially since then. The NLSY97, which follows a nearly identical structure to the NLSY79, is a nationally representative panel survey of 8,984 individuals who were 12-16 years old when they were first surveyed in 1997. Individuals were interviewed annually through 2011 and biennially since then.

The NLSY79 and the NLSY97 waves provide detailed demographic information, such as age, gender, race, and educational attainment. The files also contain measures of cognitive ability, personality traits, and sociability. We follow a large body of research, including Neal and Johnson (1996), Heckman et al. (2006), Altonji et al. (2012) and Deming (2017), and aggregate these measures into three categories (i) cognitive, (ii) non-cognitive, and (iii) social skills. These measures are taken directly from (Deming, 2017). Specifically, we downloaded these variables from Deming's replication files at `https://dataverse.harvard.edu/dataset.xhtml?persistentId=doi:10.7910/DVN/CYPKZH`.

Cognitive skills are proxied using the Armed Forces Qualifications Test (AFQT). This measure is available for both the NLSY79 and the NLSY97 waves. Altonji et al. (2012) developed a mapping of the AFQT score across the NLSY79 and NLSY97 waves that accounts for differences in age-at-test and test format. Deming (2017) normalized these to have mean zero and standard deviation one. We use his measures for all of our analysis.

While both the NLSY79 and the NLSY97 include AFQT scores, these waves contain different measures of non-cognitive and social traits. Deming (2017) provides a set of unified measures of non-cognitive and social skills which we adopt. Specifically, the Deming definition for non-cognitive skills uses (i) the normalized average of the Rotter Locus of Control and the Rosenberg Self-Esteem Scale (see Heckman et al. (2006)) in the NLSY79 wave and (ii) the Big-5 factor conscientiousness, normalized and standardized, in the NLSY97 wave. The Deming definition for social skills uses (i) an average of four self-reported normalized and standardized measures, including sociability at age 6, sociability in 1981, number of clubs each respondent participated in high school, and participation in high school sports in the NLSY79 wave and (ii) an average of two items, normalized and standardized, that capture the extroversion factor from the Big-5 personality test in the NLSY97 wave.

We restrict our primary sample to Black and White men only. We exclude observations with missing demographics or missing measures of cognitive, non-cognitive, or social skills. Our wage and employment sample focuses on prime-aged male who are full-time and full-year workers. We exclude observations that report less than 1,750 annual worked hours or hourly wages lower than 2 or higher than 500 in 2010 CPI prices. We further exclude observations with missing occupation codes. When comparing over time and across cohorts of birth, we restrict the NLSY79 sample to individuals aged 25-37 for comparability to the NLSY97 wave.

Appendix A.3 Task Measures Creation

To assess the extent to which Black and White workers sort into different occupations, perform different tasks and consequently earn different amounts, we use data from the following to measure the skills demanded in each occupation: (i) the U.S. Department of Labor's Dictionary of Occupational Titles (DOT) and (ii) the Occupational Information Network (O*NET) sponsored by the U.S. Department of Labor/Employment and Training Administration (USDOL/ETA). The DOT was constructed in 1939 to help employment offices match job seekers with job openings. It provides information on the skills demanded of over 12,000 occupations. The DOT was updated in 1949, 1964, 1977, and 1991, and replaced by the O*NET in 1998.

The DOT and the O*NET measure task requirements associated with many detailed occupations. For example, one O*Net question asks whether the occupation requires dealing with external customers - survey respondents provide responses on an ordinal scale of 0 to 5 where the higher values signify that the job requires more of that task. Different questions have answers that range on different ordinal scales (e.g., 0-5, 1-7, 0-10, etc.). We again downloaded the tasks measures directly from the replication package for Deming (2017). For all questions we use from both surveys, we follow Deming (2017) and re-scale the answers so they range from zero to ten to ensure consistency in units when we combine questions. We convert the answers into z-score units after combining them into different tasks.

We focus on four occupational task measures that are relevant for our study: (i) *Abstract*; (ii) *Routine*; (iii) *Manual* and (iv) *Contact*. The first three measures were created following the definitions in Autor and Dorn (2013) using the DOT data while the last measure builds on Deming (2017) using the O*Net data. Our goal is to stay as close to possible to the definitions of task measures developed by others to focus our analysis on the racial differences in these measures. Throughout the main paper, we define the key task measures as follows:

Abstract: indicates the degree to which the occupation demands (i) analytical flexibility, creativity, reasoning, and generalized problem-solving, and (ii) complex interpersonal communications such as persuading, selling, and managing others. Following Dorn (2009) and Autor and Dorn (2013), we measure *Abstract* tasks in practice by using the 1977 DOT data using the average scores from questions measuring *General Educational Development in Math (GED-Math)* and *Direction, Control, and Planning of Activities (DCP)*. Higher levels of *GED-Math* are associated with higher quantitative abstract tasks. Occupations with high measures of *GED Math* include various medical professionals, various engineers, accountants, and software developers. Higher levels of *DCP* are associated with higher levels of abstract thinking associated with management, organizational, and teaching tasks. Occupations with high measures of *DCP* include various managers, high school teachers, college professors and judges. To create our measure of the *Abstract* task content of an occupation, we follow Autor and Dorn (2013) and Deming (2017) and take the simple average of *GED-Math* and *DCP* for each occupation.

Routine: measures the degree to which the task requires the precise attainment of set standards and/or repetitive manual tasks. Following Dorn (2009) and Autor and Dorn (2013), we measure *Routine* task using the 1977 DOT data taking the average scores from questions measuring *Finger Dexterity (FINGDEX)* and *Set Limits, Tolerances, or Standards (STS)*. *FINGDEX* measures the ability to move fingers and manipulate small objects with fingers and serves as a proxy for repetitive routine manual tasks. Occupations with high measures of *FINGDEX* include secretaries, dental hygienists, bank tellers, machinists, textile sewing machine operators, dressmakers, and x-ray technology specialists. *STS* measures the adaptability to work situations requiring setting of limits and measurements and serves as a proxy for routine cognitive tasks. Occupations with high measures of *STS* include meter readers, pilots, drafters, auto mechanics, and various manufacturing occupations. To create our measure of the *Routine* task content of an occupation, we follow Autor and Dorn (2013) and Deming (2017) and take the simple average of *FINGDEX* and *STS* for each occupation.

Manual: measures the degree to which the task demands eye, hand, and foot coordination. Following Dorn (2009), Autor and Dorn (2013) and and Deming (2017), we measure *Manual* using the 1977 DOT data using the question *EYEHAND* which measures the ability to coordinately move hand and foot in accordance with visual stimuli. Occupations with high measures of *EYEHAND* include athletes, police and fire fighters, drivers (taxi, bus, truck), skilled construction (e.g, electricians, painters, carpenters) and landscapers/groundskeepers. To create our measure of the *Manual* task content of an occupation, we just use the *EYEHAND* measure for that occupation.

Contact: measures the extent that the job requires the worker to interact and communicate with others whether (i) within the organization or (ii) with external customers/clients or potential customers/clients. For this measure of *Contact* tasks we use two 1998 O*NET work activity variables taken from Deming (2017). Specifically, we use the variables *Job-Required Social Interaction (Interact)* and *Deal With External Customers (Customer)*.[34] *Interact* measures how much workers are required to be in contact with others in order to perform the job. *Customer* measures how much workers have to deal with either external customers (e.g., retail sales) or the public in general (e.g., police work). To make our measure of the *Contact* task content of an occupation, we take the simple average of *Interact* and *Customer* for each occupation. Occupations with high measures of *Contact* tasks include various health care workers, waiter/waitress, sales clerks, lawyers, various teachers, and various managers.

The data we use from Deming (2017) are available at the 3-digit occupational code level. We use Deming (2017)'s crosswalk to merge these measures to (i) the Census and the American Community Surveys (ACS) and (ii) the National Longitudinal Survey of the

[34] Deming (2017)'s focus is creating a measure of occupational tasks that require social skills and document how the returns to social skills have increased over time. His measure of social skills include measures of whether the job requires the worker to have social perceptiveness and the ability to coordinate, persuade and negotiate with others. For example, his measure of social skills includes O*NET questions assessing whether the job requires "adjusting actions in relation to others' action", "being aware of others' reactions and understanding why they react the way they do", and "persuading others to approach things differently". His measure of social skills do not include measures for whether the task requires interactions with other co-workers or customers. He uses the measures of customer (*Customer*) and broader social interactions (*Interact*) as controls in some of his specifications. These questions are much more suited to our purpose of trying to measure taste-based discrimination. We explore the relationship between Deming's *Social Skills* task measure and our *Contact* task measure in the online appendix.

Youth (NLSY 1979 and 1997 waves) which we use for our analysis. Again, we download these data directly from Deming's replication file at https://dataverse.harvard.edu/dataset.xhtml?persistentId=doi:10.7910/DVN/CYPKZH.

Appendix B Robustness of Racial Task Gaps: Alternate Specifications

In this section of the appendix, we show the robustness of our results pertaining to racial task gaps. We start by showing the raw task trends separately for Black and White men (in the main text, we only show the racial gaps). We then show the robustness of the racial task gaps documented in the main text to an alternate specifications. Finally, we show the racial task gaps separately for different education groups.

Appendix B.1 Occupational Task Sorting, Separately by Race

Appendix Figure A1 plots the raw trends in occupational tasks separately for White (Panel A) and Black (Panel B) men since 1960 using the Census/ACS data. As in the main text, we restrict our sample to native born men between the ages of 25 and 54 who are not self employed and who report currently working full time (e.g., at least 30 hours per week). Specifically, Appendix Figure A1 reports the coefficients on the year dummies (β_t) from the following regressions using our individual Census/ACS data:

$$\tau_{ijgt}^{k} = \sum_{t} \beta_t^g D_t + \epsilon_{ijgt} \tag{A1}$$

where τ_{ijgt}^{k} is the task content of task k for individual i from group g working in occupation j in period t. Task contents are expressed in z-score units. We run this regression separately for two groups g: White men and Black men. As a result, all coefficients have g superscripts. We explore the four tasks k highlighted in the main text. D_t is a vector of dummies that take the value of 1 if the year is, respectively, 1960, 1970, 1980, 1990, 2000, 2012, or 2018. The coefficient on the year dummies from these regressions, β_t^g are plotted in the figure.

A few things are noticeable from Figure A1. First, focusing on Panel A, there were large shifts in the task content of occupations to which White men sorted into over the last six decades. During this time period, White men have sorted more into *Abstract* occupations and sorted away from *Routine* occupations. Likewise, White men have sorted into occupations which require meaningful *Contact* tasks during the last half century. These patterns mimic the patterns for *Abstract*, *Routine* and *Manual* tasks highlighted in Dorn (2009), Autor

Figure A1: Raw Task Trends: White and Black Men

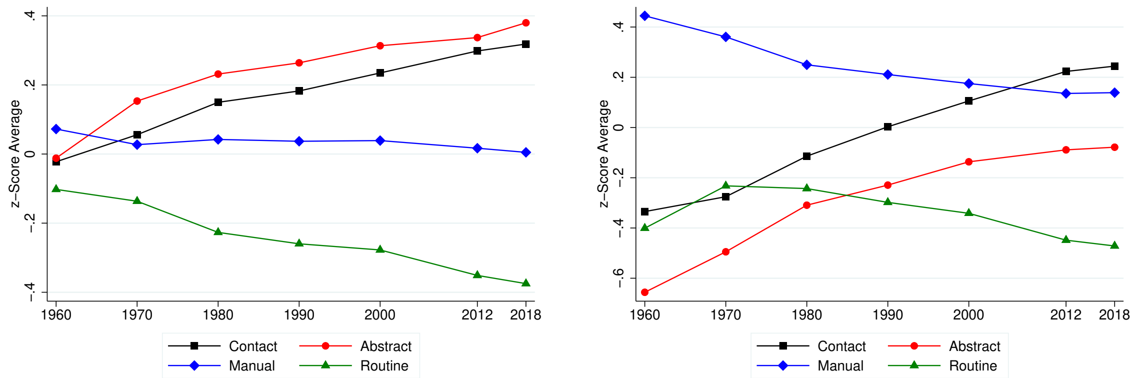

Panel A: White Men

Panel B: Black Men

Notes: Figure shows the raw trend in the task content of jobs for White and Black men using Census and ACS data. Sample restricted to native born individuals between the ages of 25 and 54 who are not self-employed but who are working full time. Tasks are expressed as z-scores across occupations. Regressions are run separately for White men (Panel A) and Black men (Panel B) and were weighted using Census/ACS individual sampling weights.

and Dorn (2013), and Deming (2017). This is not surprising given we are using their task measures in this analysis. The main text highlights the racial gaps in these measures by comparing the patterns between Panels A and B with and without additional demographic controls.

Appendix B.2 Alternate Specification of Racial Task Gaps

We next provide an alternate specification of racial gaps in the task content of occupations. This is an alternate way to display the data shown in Figure 2 of the main text. Particularly, we estimate the following:

$$\tau^k_{ijt} = \alpha^k_t + \beta^k_t Black_{it} + \sum_{s \neq k} \omega^k_{st} \tau^s_{ijt} + \Gamma^k X_{it} + \epsilon^k_{ijt} \tag{A2}$$

where τ^k_{ijt}, $Black_{it}$ and X_{it} are defined as in the main text. In this specification, we isolate the partial effect of the racial gap in the task content of occupations by controlling directly for the other tasks in the regression ($\sum_{s \neq k} \tau^s_{ijt}$). We run this regression separately for each year and for each task yielding 28 estimates of β^k_t. These coefficients are plotted in Appendix Figure A2. Panel A shows the results excluding the X vector of demographic controls while Panel B shows the results including the additional controls. The racial gaps are expressed in z-score units.

Figure A2: Race Gap in Tasks Without and With Demographic Controls, Alternate Specification

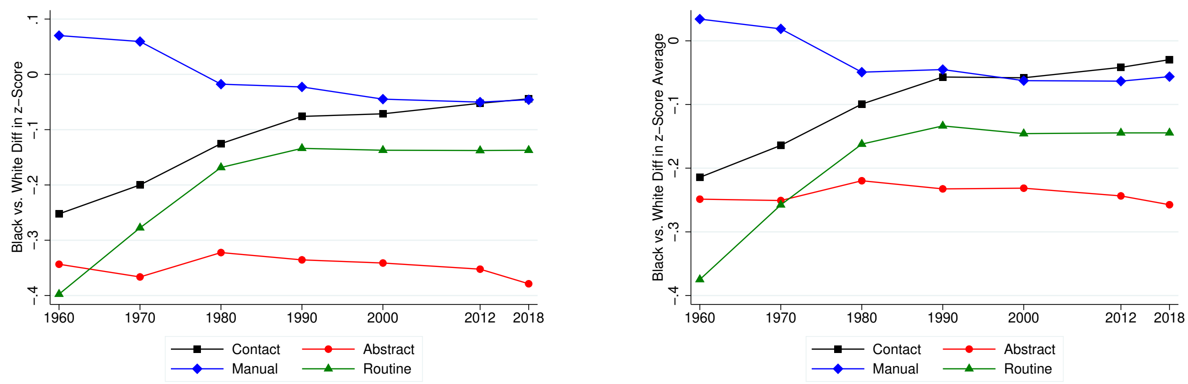

Panel A: No Demographic Controls Panel B: With Demographic Controls

Notes: Figure shows the racial gap in task trends using our alternate specification in the appendix without (Panel A) and with (Panel B) demographic controls. Sample restricted to native born individuals between the ages of 25 and 54 who are not self-employed but who are working full time. Tasks are expressed as z-scores across occupations.

This figure is the analog to Figure 2 in the main text. The time series patterns from this specification are nearly identical to the patterns in the main text. The only difference is the units in which the racial task gaps are expressed. The racial gap in *Abstract* tasks was essentially constant between 1960 and 2018 at a level a 0.25 standard deviation gap (with controls). However, the racial gap in *Contact* tasks converged from roughly a 0.25 standard deviation gap in 1960 to a close to zero gap in 2018.

Appendix B.3 Racial Task Gaps, by Education Levels

Finally, we show robustness of the time series patterns in racial task gaps within different education groups using our main specification described in the text. Panel A of Appendix Figure A3 redoes the main results of Figure 2 of the main text (with demographic controls) but segmenting the sample to only those individuals with education less than a bachelor's degree. Panel B shows the same specification but restricting the sample to those individuals with a bachelors degree or more. These figures shed light on whether our time series patterns of the changing racial task gaps that we highlight in the main paper are found in both higher and lower education samples. For both education groups, there was a convergence in *Contact* tasks and a slight divergence in *Abstract* tasks. The magnitude of the *Contact* convergence is

Figure A3: Race Gap in Tasks: By Educated Groups

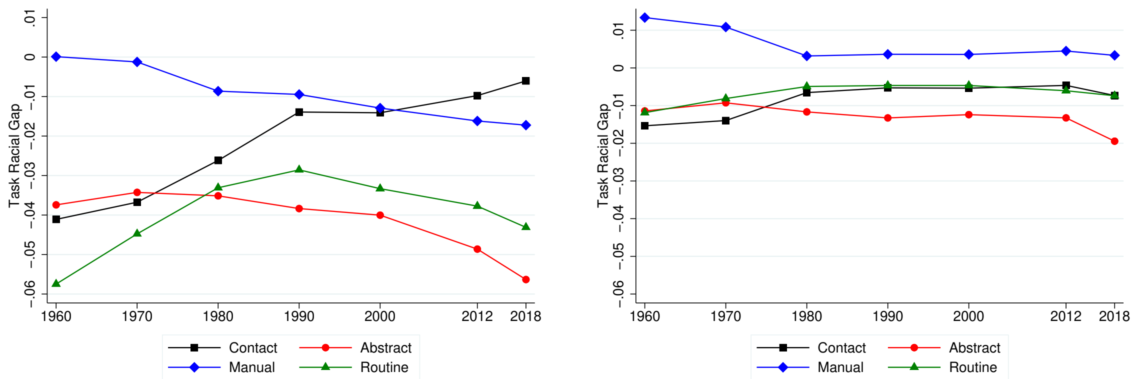

Panel A: Less than a Bachelors Degree

Panel B: Bachelors Degree of More

Notes: Figure re-estimates Panel B of Figure 2 of the main text separately by those with less than a bachelors degree (Panel A) and those with a bachelors degree or more (Panel B).

much larger for less educated individuals, but given selection (Panel A represents between 70 and 75 percent of the sample depending on the year), it is not surprising that the convergence in *Contact* tasks is smaller for higher educated individuals.

Appendix C Robustness of Racial Task Gaps: Alternate Task Definitions

In this section, we explore the robustness of our results to alternate task definitions. We begin by disaggregating our current task measures into their separate task components. We then explore the racial gaps in alternate definitions of four main task categories. Finally, we compare our *Contact* tasks measure to Deming's *Social* task measure. As seen in this section, our results are quite robust to alternate task definitions.

Appendix C.1 Decomposing Task Measures into Sub-Components

We used three task measures emphasized in the recent literature using DOT data: *Abstract*, *Routine* and *Manual* tasks. As discussed above, these three measures of tasks were created using five separate questions from the DOT data. *Abstract* task is a combination of $GED-Math$ and DCP. *Routine* task is a combination of $FINGDEX$ and STS. In this subsection of the appendix, we move from using four tasks measures (*Abstract*, *Routine*, *Manual*, and *Contact*) to six tasks measures (*GED-Math*, *DCP*, *FINGDEX*, *STS*, *Manual* and *Contact*).

Figure A4: Race Gap in Tasks: Disaggregated Task Measures

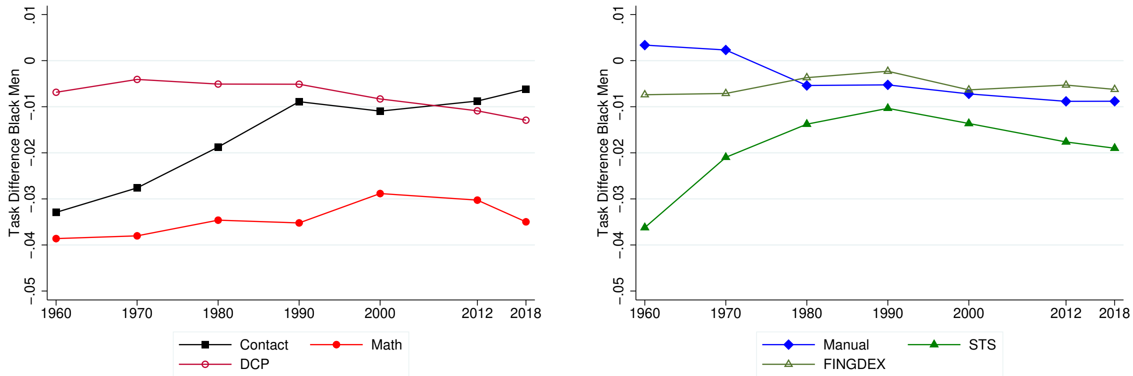

PANEL A PANEL B

Notes: Figure re-estimates Panel B of Figure 2 of the main text with six task components instead of four. In particular, we disaggregate *Abstract* tasks into its (1) *Math* and (2) DCP sub-components. Likewise, we disaggregate *Routine* tasks into its (1) STS and (2) Finger subcomponents.

In particular, we re-estimate the results in Panel B of Figure 2 using six task measures instead of four. The sample used is the same as in Panel B of Figure 2 of the main text. The coefficients on the task measures from these yearly regressions are plotted in Appendix Figure A4. We plot the coefficients in two panels instead of one for readability.

The figure shows that the main take-aways highlighted in the text are unaltered when using the six task measures. Specifically, there have been no relative gains by Blacks with respect to either component of *Abstract* tasks; Blacks were underrepresented in both *GED Math* and *DCP* in 1960 and the race gap was constant through 2018. However, Blacks made large gains in *Contact* tasks over this time period.

Appendix Figure A5 shows the results from the regression but with seven tasks measures. We still include *GED-Math*, *DCP*, *FINGDEX*, *STS* and *Manual*. But, we now disaggregate *Contact* into its two sub-components: *Interact* and *Customer*. The former measures the extent to which the job requires social interactions with others while the latter measures whether the job requires individuals to deal with external customers. Instead of showing all seven coefficients, we only show the coefficients on *Interact* tasks and *Customer* tasks.[35] There was racial convergence in both tasks requiring contact within the firm (*Interact*) and tasks requiring contact with external customers (*Customer*). These results highlight that Blacks were moving into occupations (relatively) that require both forms of contact with others.

[35]The coefficients on the other five tasks were essentially unchanged relative to Appendix Figure A4.

Figure A5: Race Gap in Disaggregated *Contact* Task Measures

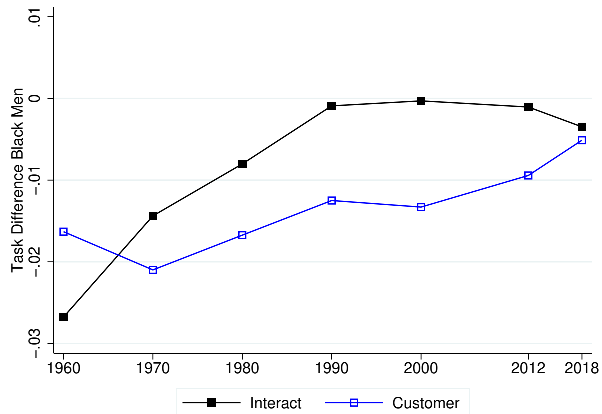

Notes: Figure re-estimates Panel B of Figure 2 of the main text with seven task components instead of four. In particular, we disaggregate *Abstract* tasks into its (1) *Math* and (2) *DCP* sub-components. Likewise, we disaggregate *Routine* tasks into its (1) STS and (2) *FINGDEX* subcomponents. Finally, we disaggregate *Contact* tasks into (1) *Interact* and (2) *Customer* sub-components. Only the coefficients on the *Interact* and *Customer* task measures from these yearly regressions are plotted in the figure.

Appendix C.2 Robustness to O*Net Measures of *Math* and *Routine* Tasks

Deming (2017) used data from 1998 O*Net survey to make two alternate measures of *Math* and *Routine* occupations. For his alternate *Math* task measure, he combines O*Net questions measuring (i) the extent to which an occupation requires mathematical reasoning, (ii) whether the occupation requires using mathematics to solve problems, and (iii) whether the occupation requires knowledge of mathematics. The measure of the *GED-Math* task content of an occupation created using DOT data is highly correlated with Deming's *Math* task content of an occupation created using the O*Net data; the correlation between the two series (weighted by 1990 population in each occupation) is 0.81.

For his alternate *Routine* task measure, Deming again uses the 1998 O*Net and combines the questions measuring (i) how automated is the job and (ii) how important is repeating the same physical activity (e.g. key entry) or mental activities (e.g., checking entries in a ledger over and over, without stopping to perform the job). This measure is highly correlated with the *STS* portion of *Routine* tasks within the DOT data. However, conditional on controlling for the *STS* content of a job, the Deming *Routine* task measure using the O*Net data is uncorrelated with the occupations *FINGDEX* task content.[36] Given this, we treat Deming's

[36]Regressing the Deming *Routine* task content of an occupation on the occupation's *STS* and *FINGDEX* task content (weighted by 1990 population counts in each occupation) yields a coefficient on STS of 0.50 (standard error = 0.05) and a coefficient on *FINGDEX* of -0.06 (standard error = 0.06).

Figure A6: Race Gap in Tasks: Alternate Measures of *Routine* and *Math* Task Measures

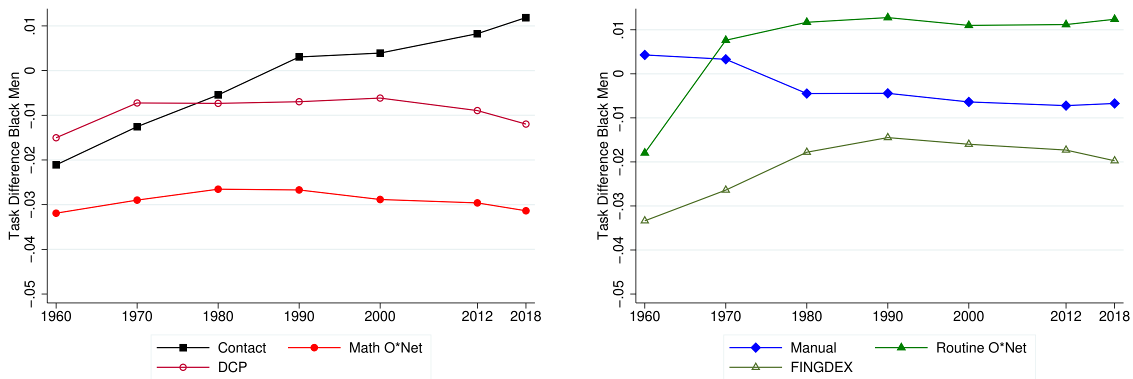

Panel A   Panel B

Notes: Figure re-estimates Panel B of Figure 2 of the main text with six task components instead of four. In particular, we disaggregate *Abstract* tasks into its (1) *Math* and (2) *DCP* sub-components. For this figure, we use Deming's measure of occupational *Math* task measures using the O*Net data. Likewise, we disaggregate the DOT *Routine* tasks into its (1) *STS* and (2) *Finger* subcomponents. However, we replace the DOT *STS* measure with Deming's *Routine* task measure using O*Net data.

Routine task measure created using the 1998 O*Net data as being an alternative for the *STS* task measure within the DOT data.

With this in mind, we explore the sensitivity of our results to using Deming's *Math* and *Routine* measure using the O*Net data as alternative task measures for the *GED-Math* and *STS* measures using the DOT data. We re-estimate the patterns in Appendix Figure A4 with the six task measures but we use the alternate Deming measures for *Math* and *STS*. The results of this regression are shown in Appendix Figure A6. Again, we display the results over two panels for readability. Our main results are unchanged with these two alternative task measures. Primarily, there has still been no racial progress in the *Math* task content of an occupation over the last 60 years. However, there have been a large convergence in the racial gap in occupational *Contact* tasks.

Appendix C.3 Alternate Measures of *Contact* Tasks

The key finding in our paper is the racial convergence in *Contact* tasks relative to *Abstract* tasks in the U.S. over the last half century. In this sub-section, we explore the sensitivity of our results to using other measures of *Contact* tasks. Deming's *Social Skills* task measure is highly correlated with our *Contact* task measure. This is not surprising given that Deming's measure of *Social Skills* tasks measures whether the occupation requires skills associated with the ability to coordinate, negotiate, and persuade. The ability to coordinate, negotiate, and

Figure A7: Correlation Between *Contact* Task and Social Task, Cross-Occupation Variation

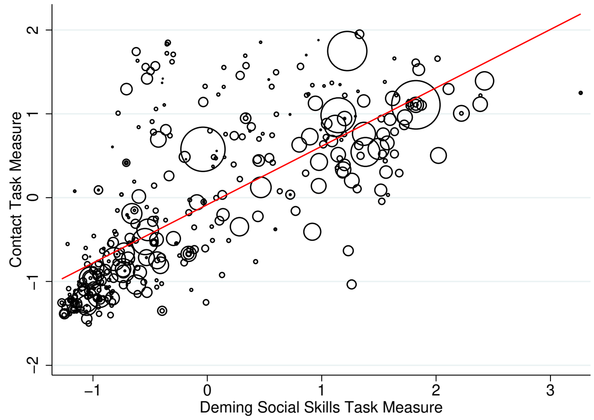

Notes: Figure shows a scatter plot of the correlation between the *Contact* task content of an occupation and Deming's *Social Skills* task content of an occupation. Each observation in the figure is an occupation. *Contact* and *Social Skills* tasks are measured in z-score space. The size of the circle represents the number of prime age men working in that occupation in 1990. Figure also includes the weighted simple regression line through the scatter plot. The coefficient on the z-score for *Social* tasks is 0.70 (standard error = 0.03) and an adjusted R-squared of 0.65.

persuade is needed when the job requires workers to come into contact with other co-workers, clients and customers. The simple correlation between Deming's *Social Skills* task measure and our *Contact* task measure is about 0.7 (weighted by 1990 population counts within in each occupation). We show the simple scatter plot by occupation of the two measures in Appendix Figure A7.

Appendix Figure A8 is the analog to Appendix Figure A6 except we replace our *Contact* task measure with Deming's *Social Skills* task measure. As highlighted in Deming (2017), the *Social Skills* task content of an occupation is highly correlated with the *Math* and the *DCP* task content of an occupation. As a result, the racial gap in *Abstract* Skills is smaller and the racial gap in *Social Skills* is larger in 1960. Despite that, our key patterns remain. There was a substantial narrowing of the racial gap in the *Social Skills* task content of an occupation since 1960. When we use this measure, there is also a slight divergence in the task content of the two components of *Abstract* skills. Black men are gaining relative to White men not because of a convergence in *Abstract* tasks but a convergence in tasks that require interactions with others.

As one final robustness exercise, we created a combined *Contact*/Social task measure by taking the simple average of our *Interact* task measure, our *Customer* task measure, and Deming's *Social Skills* task measure for each occupation. Appendix Figure A9 shows the analog of Panel B of Figure 2 of the main text but replacing our main *Contact* task measure with the combined *Contact*/Social Skills task measure. As seen from the figure, our key

Figure A8: Race Gap in Tasks: Social Skills Tasks

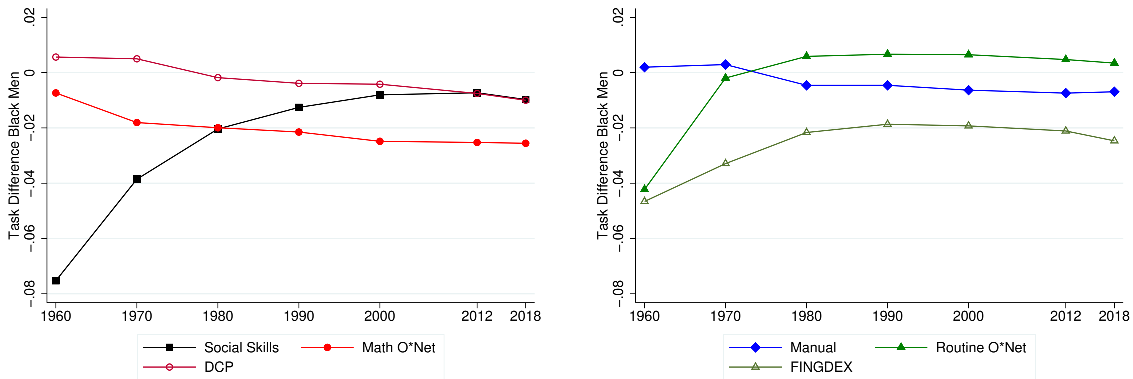

Panel A Panel B

Notes: Figure re-estimates Appendix Figure 2 with six task measures further replacing Deming's *Social Skills* task measure for our *Contact* task measure. As with Appendix Figure A6, we use Deming's measure of occupational *Math* and *Routine* task measures along with our measures of DCP, FINGDEX, and *Manual* tasks.

findings remain with this broader measure of the social interactions needed in occupations.

Appendix D Task Gaps by Gender

Appendix Figure A10 shows the occupational task differences between White men and White women (panel A) and between White women and Black women (panel B) using data from the Census/ACS. This figure uses the same specification as Panel B of Figure 2 in the main text. Panel A of this appendix figure restricts the sample to native born White men and White women between the ages of 25 and 54. Panel B restricts the sample to native born White women and Black women between the ages of 25 and 54. Both panels also restrict the sample to those individuals working full time and excludes the self-employed. As with the figures in the main text, we condition on education and age when we measure the gaps in the task content of jobs.

As seen from Panel A, women are much more likely to be in *Contact* and *Routine* tasks and are much less likely to be in *Manual* and *Abstract* tasks. Unlike the gaps between Black and White men, the gaps between White men and White women were fairly stable over the last 60 years. One exception is the gap in *Abstract* tasks. In the 1960, White women were 16 percentage points less likely to work in occupations that require 1 standard deviation higher *Abstract* tasks relative to White men (conditional on age and education). By 2018, that gap fell to 7 percentage points.

Figure A9: Predicted Race Gap in Task Content, with Combined *Contact*/Social Skills Task Measure

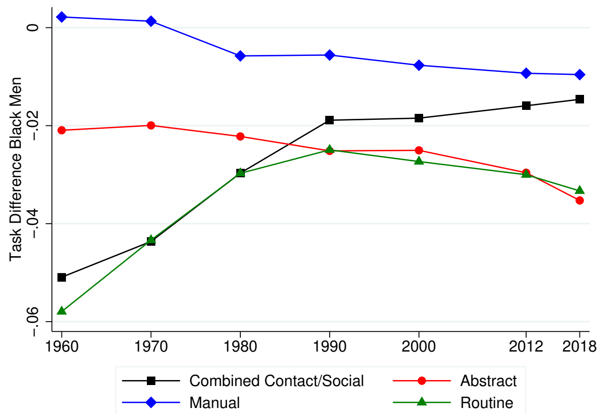

Notes: Figure is analogous to Panel B of Figure 2 in the main text. The only differences is that in this figure we combine Deming's *Social Skills* task measure and our *Contact* task measure into one combined task measure.

Figure A10: Task Differentials between White Men and White Women and between White Women and Black Women

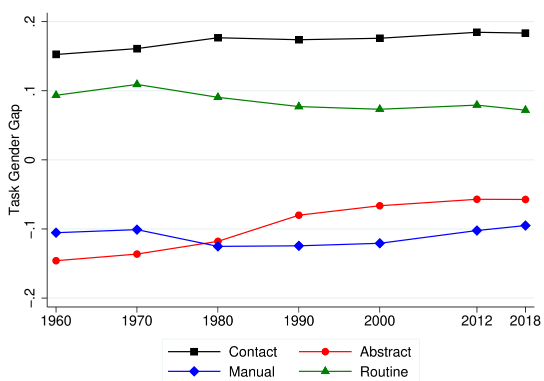

Panel A: White Men vs White Women

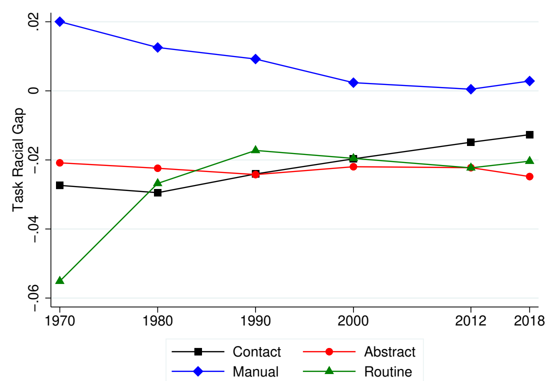

Panel B: White Women vs Black Women

Notes: Figure shows the extent to which the task content of an occupation can predict whether an individual employed in that occupation is a White woman (Panel A) or a Black woman (Panel B). For the regressions in Panel A, we use the Census/ACS sample pooling together prime-age White men and women. Panel shows the coefficients from a regression of a dummy variable equal to one if the individual is a White woman on the four task measures and controls for individual education, age and Census division, separately by year. For the regressions in Panel B, we use the Census/ACS sample pooling together prime-age White women and Black women. Panel shows the coefficients from a regression of a dummy variable equal to one if the individual is a Black woman on the four task measures and controls for individual education and age, separately by year. All samples for both regressions are also restricted to full time workers who are not self employed and who are native born.

The time series patterns in Panel B between White women and Black women mirror the patterns in Panel B of Figure 2 of the main text showing differences between White men and Black men although the level gaps are smaller. The gap in the *Abstract* task content of jobs between White and Black women was roughly constant between 1960 and 2018. However, Black women converged to White women in the *Contact* task content of jobs over this period.

Appendix E The Relationship Between Log Wages and Skills, NLSY Data

In this section of the appendix, we assess the extent to which there are racial differences in the responsiveness of wages to individual skill measures. The coefficients in Panel A of Appendix TableA1 comes from a regression of log individual wages of NLSY respondents on NLSY cognitive, non-cognitive and social skill measures and those skill measures interacted with a race dummy. The regression also includes age, education and occupation fixed effects and those fixed effects interacted with a race dummy. Finally, the regression also includes year and NLSY97 sample fixed effects and those fixed effects interacted with a race dummy. For this regression, we pool together both the NLSY79 and NLSY97 samples. As with the rest of the paper, we only include in our sample Black and White men between the ages of 25 and 54.

Panel A of the table highlights the labor market returns to cognitive, non-cognitive and social skills for White men (first three rows). As seen from the table, having more of any of the three skill measures raises labor market earnings for White men (even conditional on individual education and occupation). Furthermore, we find no differential labor market returns for Black men for non-cognitive and social skills. However, similar to the findings in Neal (2006), the coefficient on AFQT in a regression of log wages on AFQT scores is larger for Blacks than for Whites. This is consistent with the conjecture that Black men who receive the same AFQT test score relative to White men (conditional on education and occupation) may be positively selected in traits not measured in the NLSY that are rewarded in the labor market.

Panel B of the regression runs a similar regress as Panel A but replaces the three NLSY skill measures with individual measures of *Abstract* and *Contact* skills using the prediction equation from equation (10) of the main text. Otherwise, the sample and demographic controls are the same as in Panel A. Specifically, to the get the model imputed task specific skill measures, we simply multiply the coefficients shown in Table 5 by the individual's

Table A1: Racial Differences in the Relationship between Log Wages and Skill Measures, Model vs. Data

Panel A: NLSY Skill Measures	
Cognitive	0.070
	(0.009)
Non-Cognitive	0.034
	(0.007)
Social	0.014
	(0.007)
$Black * Cognitive$	0.037
	(0.015)
$Black *$ *Non-Cognitive*	-0.004
	(0.011)
$Black * Social$	0.001
	(0.011)
Panel B: Model Imputed Skill Measures	
Abstract	0.990
	(0.121)
Contact	0.459
	(0.100)
$Black * Abstract$	0.474
	(0.217)
$Black * Contact$	-0.122
	(0.164)

Note: Panel A of table shows key coefficients from a regression of log wages on a cognitive, non-cognitive, and social skills and those skill measures interacted with a Black dummy using the NLSY micro data. Panel B of table shows key coefficients from a regression of log wage on model predicted Abstract and Social skills and those skill measures interacted with a Black dummy using NLSY micro data. All regressions include controls for individual age, education, and occupation and those controls interacted with a race dummy. Additionally, the regression also includes year and sample fixed effects plus those fixed effects interacted with a race dummy. The sample includes all Black and White men in both waves of the NLSY data between the ages of 25 and 54. Robust standard errors clustered at the individual level are shown in parentheses.

reported measures of cognitive, non-cognitive and social skills.[37] We then multiply the imputed task specific individual skills by the respective task requirement of the individual's occupation. This ensures that that the skills are in consistent model units.[38] Three sets of results can be seen from Panel B. First, like in our calibrated model, the return to *Abstract* skills for White men is roughly twice the return to *Contact* skills for White men. Second, the return to model implied *Abstract* skills is estimated to be larger for Black men relative to White men. Finally, the estimated return to model implied *Contact* skills is not statistically different between Black and White men.

Appendix F Additional Results on Estimated Model Fit and Model Validation

In this section of the appendix, we show additional results on how well our calibrated model matches both additional targeted and non-targeted moments.

Appendix F.1 Model Fit

Figure A11 compares the key model moments (solid lines) against the corresponding data targets (dashed lines). As seen from the various panels of the figure, our model generally fits the data quite well. The model fit for the racial gap in the *Manual* task content of jobs – the moment we do not target – is naturally less tight (not shown), but nonetheless the model is able to match the fact that the racial gap in *Manual* tasks is close to zero. This makes us confident that our estimate of $\beta_{Manual,t}$ being equal to zero (which means that racial barriers in *Manual* tasks have no effect on sorting or wages) has little impact on our key paper results.

Appendix F.2 Model Validation

The counterfactuals we explore in the paper rely on the functional form assumptions we made for the various distributions from which individuals draw task specific skills or preferences. In this subsection of the appendix, we explore whether such distributional assumptions are grossly at odds with the data by assessing the extent to which our estimated model matches other non-targeted moments.

[37]Given the low first stage power for *Routine* skills, we do not include a prediction of model implied *Routine* skills in this regression.

[38]Specifically, according to equation 8, the model structure implies that an individual's log wages are linear in $\tilde{\tau}_{jk}\phi_{ik}$. Note the δ_k's and η_k's are absorbed in the occupation-race fixed effects.

Figure A11: Model versus Data Moments

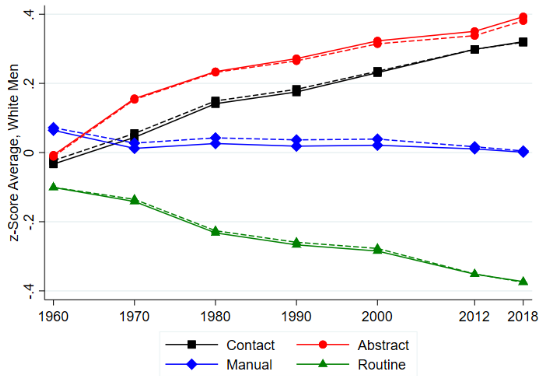

Panel A: Task Contents, White Men

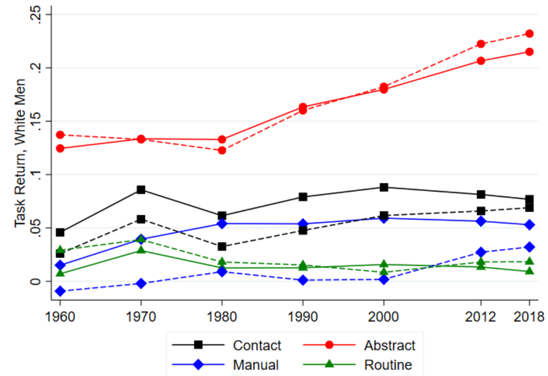

Panel B: Task Prices, White Men

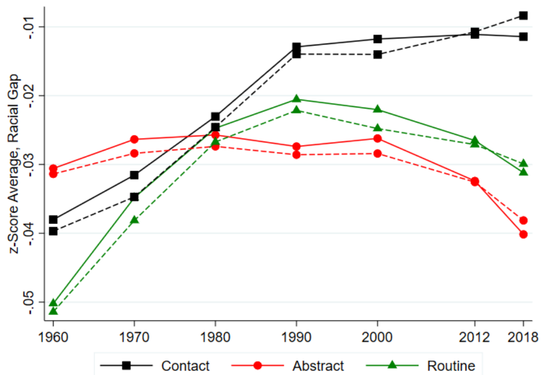

Panel C: Task Contents, Gap

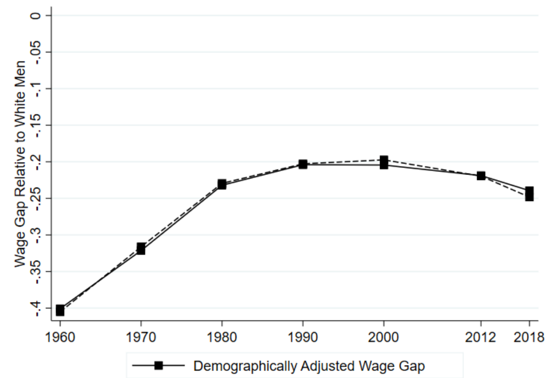

Panel D: Aggregate Wage Gap

Notes: Figure shows how selected model moments (solid lines) compare to their corresponding data moments (dashed lines). The data moments are the ones used to discipline the model. Panels A and B are data for White Men and are unconditional on education. Panels C and D are the racial gaps in wages and task content of occupations conditional on age and education as highlighted in Figures 1 and 2 to account for these demographic differences between Black and White men.

When calibrating our model, we targeted the mean wage gap between Black and White men as one of our key moments. We now explore how our model performs in matching the trends in racial wage rank gaps for different percentiles as documented by Bayer and Charles (2018). Specifically, we compute (separately by year) the median and 90th percentile of the Black wage distribution, and find out the positions of these Black wages in the White wage distribution. The differences in positions of these Black wages in Black and White distributions constitute the "wage rank gaps" at the median and 90th percentile, respectively. For example, a relative wage rank gap of -30 for the median series implies that the median wage of Black men is at the 20th percentile of the White men wage distribution or 30 percentage points lower than the median. Likewise, a relative rank gap of -30 for the 90th percentile series implies that the 90th percentile in the Black man wage distribution is at

Figure A12: Model Performance Against Non-Targeted Empirical Moments

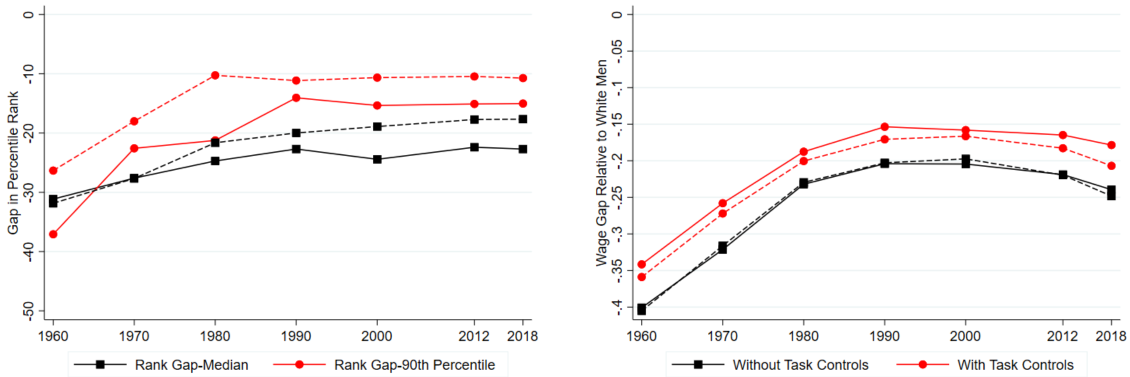

Panel A: Racial Gap in Percentile Rank of Wages

Panel B: Racial Wage Gaps Conditional on Tasks

Notes: Panel A shows the model implied racial rank gaps for different percentiles against their empirical analogs. In particular, the solid black line (with squares) shows the relative rank gap. Panel B shows model based estimates (solid lines) and data estimates from the Census/ACS (dashed lines) of demographically adjusted racial wage gaps with and without controlling for the task content of occupations. For these figures, we show model results assuming $\sigma^2 = 0$. Model results in Panels A and B are robust to alternate values of σ^2.

the 60th percentile of the White man wage distribution. For this analysis, we follow Bayer and Charles (2018) and include both working and non-working individuals in our analysis with the wages of non-working individuals set to zero.

Panel A of Appendix Figure A12 shows our results. The dashed black line (with squares) represents the relative racial rank gap for the median series while the dashed red line (with circles) represents the relative rank gap for the 90th percentile, both using our Census/ACS data. The black and red solid lines, respectively, show the analogs from the model. It should be noted that the empirical findings from the Census/ACS data in Panel A are similar to those documented in Bayer and Charles (2018). The median Black man in 1960 had a wage that was equal to the 20th percentile of the White wage distribution. Between 1960 and 2018, the relative rank gap of the median Black made little progress. Both in 1980 and 2018, the median Black man had wages that was equal to about the 25th percentile of the White wage distribution. Conversely much more relative progress was made for Blacks at the top of the wage distribution. In 1960, the 90th percentile of the Black wage distribution was at about the 60th percentile of the White wage distribution. By 2018, the 90th percentile of the Black wage distribution had a value that was equal to roughly the 80th percentile of White distribution. However, even for the 90th percentile, little progress was made in the racial rank gap since 1980. Notice, our model (in solid lines) roughly matches these patterns even though they were not targeted. This suggests that model driving forces and racial sorting that we estimate can explain relative racial wage patterns throughout the wage distribution.

Panel B of Appendix Figure A12 shows the demographically-adjusted racial wage gap (Black lines with squares) and the racial wage gap conditional on task controls (red lines with circles), where the solid lines are model-implied and the dashed lines are their data analogs using the Census/ACS samples. Specifically, to get the red lines we regress the log wages on a race dummy and the τ_{jk}'s for each of the four tasks, separately for each year, first with the model-generated data and then with the Census/ACS data. As the comparison of the black and red solid lines reveals, the model predicts that controlling for occupational tasks only has a small effect on the estimated racial wage gap. This model finding closely matches what we find in the data. Again, these results were not targeted when calibrating the model. The similarity stems from the fact that the sorting on skills in the model is close to the sorting on skills in the data. Collectively, the fact that our estimated model matches a variety of non-target moments gives us confidence in the counterfactuals we highlight next.

Appendix G Model Estimates of Home Sector Preferences

In this section, we report the model estimates of the racial gap between preferences for the home sector in each year: $A_{bHt} - A_{wHt}$. The racial gap in the A_{gH}'s ensures that the model matches labor force participation of Black and White men in each year. The results are shown in Appendix Figure A13. For the most part, Panel A shows that the racial gap in the A_{gH}'s are relatively constant over time. However, it should be noted that the model does generate a slight increase in the preference for the home sector between Black and White men between 1980 and 2000 and a slight decline in the relative home sector preference thereafter. The relatively preferences are essentially unchanged in 1970, 1980, 2012 and 2018 relative to 1960. Panel B shows that the racial gap in non-employment rates between Black and White men both in the model (solid line) and the data (dashed line). It is not surprising that our model is matching the empirical racial gap in employment rates because we are targeting the moment. Our model estimates of $A_{bHt} - A_{wHt}$ basically just tracks the racial gap in non-employment rates between Black and White men over time.

Figure A13: Racial Differences in Home Sector Preferences

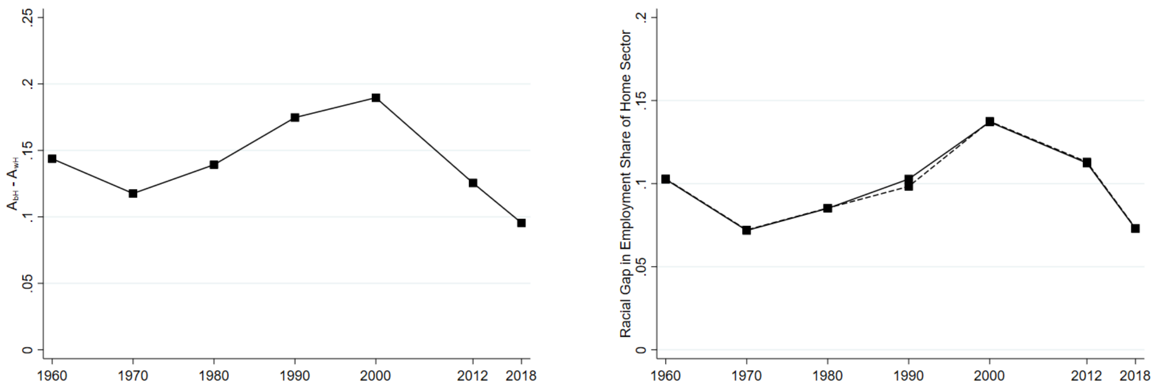

Panel A: Difference in A_{gH} Panel B: Gap in Home Sector Shares

Notes: Panel A of Figure shows the estimated differences in race-specific home sector preference parameters, A_{bH} and A_{wH}. Panel B shows the racial difference in non-employment rates in the model (solid line) and Census/ACS data (dashed line) for prime age Black and White men.

Appendix H Additional Results on Model Selection

Our task-based model captures the extent to which Black and White workers sort into different tasks in response to racial barriers. A simple statistical model with multiple skills but without a sorting mechanism will miss this force. In this section, we highlight the quantitative importance of allowing for differential sorting into tasks when we consider the effect of changing task returns on the racial wage gap.

To better understand the importance of differential selection into tasks, it would be useful to break down wages into components by task. Specifically, we define the task wage on k to be the part of the wage earned from performing task k, namely $\beta_{kt}\tilde{\tau}_{jk}(\phi_{ikt}+\eta_{gkt}+\delta^{taste}_{gkt})$; the log wage of worker i in occupation j in period t (w_{ijt}) is the sum of the wage on each task plus the occupational return A_{jt}. The Black-White gap in the *wage associated with a given task* can thus reflect three forces: (i) task racial barriers $\eta_{bkt}+\delta^{taste}_{bkt}$, (ii) differential selection on ϕ_{ikt} between Black and White men into the task, and (iii) the size of the task return β_{kt}. Importantly, the Black-White gap in task wages approximately equals the elasticity of the racial wage gap to the task price, β_{kt}. That is, a one-percent increase in β_{kt} for *Abstract* will widen the racial wage gap roughly by the size of the *Abstract* task wage gap.[39] We can therefore compare how large the racial task wage gaps in the model would be with and

[39]If workers did not re-optimize their occupational sorting in response to task return changes, the racial gap in task wages exactly equals the elasticity of the racial wage gap to β_k. But the effect of re-optimization is quantitatively small because of a logic similar to the envelope theorem. Hence, the racial gap in task wages roughly equals the elasticity.

Figure A14: *Abstract* Task Wage Gaps, with and without Sorting

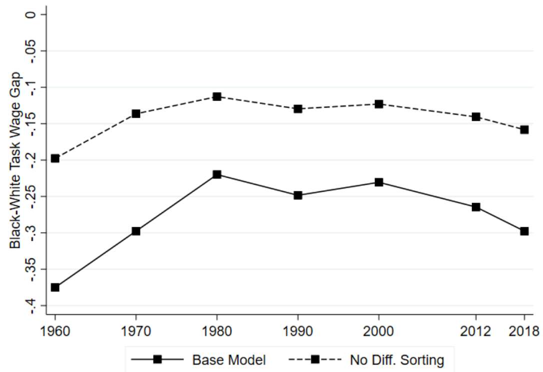

Notes: Figure shows the Black-White gap in wage associated with *Abstract* tasks, namely $\beta_{kt}\tilde{\tau}_{jk}(\phi_{ikt} + \eta_{gkt} + \delta^{taste}_{gkt})$ for $k = Abstract$, in the baseline scenario (solid line) and in the counterfactual scenario where we assume Black workers sort exactly as White workers do (dotted line).

without the differential sorting into tasks and infer the quantitative difference our task-based model makes relative to a multi-skill model without differential sorting.

Figure A14 shows the Black-White gap in *Abstract* task wages over time with and without differential sorting into tasks. We focus on *Abstract* task wages because the rising *Abstract* task return post-1980 is the driving force of our model result on the stagnation of racial wage convergence. Specifically, the solid black line plots the racial gap in task wages in the base model (which captures differential sorting), whereas the dotted black line plots the counterfactual task wage gap when Black men sort exactly as White men do. Generally, even in absence of differential sorting, changes in aggregate tasks returns can disproportionately raise the wages of one group relative to the other when one group suffers from skill gaps or discrimination. For example, if Black workers sorted exactly as Whites do, a one-percent rise in $\beta_{Abstract}$ would widen the racial wage gap by about 0.1 log points in 1980 (dotted black line). Differential sorting into tasks, however, further amplifies this effect. Notably, when Black and White workers sort differentially, the effect of a one-percent rise in $\beta_{Abstract}$ on the racial wage gap in 1980 is doubled to about 0.2 log points (solid black line). Since Black workers face high racial barriers in *Abstract* tasks, they tend to sort away from occupations with high *Abstract* task requirements. But this means that Black workers will miss out relative to Whites when the return to *Abstract* task rises. Consequently, when the *Abstract* task return rises, the racial wage gap in our task-based model widens over and beyond what the composite racial barrier $\eta_{bkt} + \delta^{taste}_{bkt}$ alone would imply in absence of differential sorting.

In sum, differential sorting into tasks amplifies the impact of changing task returns on the racial wage gap. We will miss this quantitatively significant amplification mechanism even in a model with multiple skills so long as it ignores sorting. To appreciate the true impact of changing task returns, we need a structural model with a meaningful sorting mechanism. The next section of the appendix compares the results of our structural model to other statistical decomposition methods.

Appendix I Comparison of Model Based Decomposition Method to Juhn-Murphy-Pierce Statistical Decomposition Method

Our estimated model yields quantitatively different conclusions about the extent to which race-specific factors (like a narrowing of racial skill gaps or a decline in discrimination) have improved in the United States during the last forty years relative to what would have been concluded using a popular statistical decomposition method developed by Juhn et al. (1991). Juhn et al. (1991) attribute the slowdown of convergence in the racial wage gap to rising skill prices. Central to their analysis is the racial wage *rank* gap, i.e., the position (percentile rank) of Black workers in the White earnings distribution. Specifically, they decompose trends in racial wage gaps into Black workers changing their position in the White distribution ("positional" convergence) and a change in the variance of the (White) earnings distribution ("distributional" convergence).[40] Their key insight is that changes in the level of inequality within the White earnings distribution can impact the racial wage gap even if Blacks maintained the same position, simply because Blacks and Whites occupy different initial positions in the earnings distribution. In their attempt to distinguish race-specific forces from general forces such as skill price changes, they perform the statistical decomposition of the racial wage gap trends into distributional and positional convergence, and then interpret the former as stemming from trends in skill prices and the latter as proxies for trends in race-specific forces. Such an interpretation is valid in a univariate skill model. In such a model, two workers with the same earnings will have the same underlying levels of aggregate skills, so changes in aggregate skill returns will affect them in the same way. Said differently, White men of a given wage is a good control group to proxy for the unobserved skills of Black men in a model with one aggregate skill price. This means trends in skill prices cause distributional convergence but not positional convergence; hence, when there is only one aggregate skill price, it is correct to attribute positional convergence to trends in race-specific forces.

However, in a multivariate skill model like ours, the distributional convergence fails to

[40]The terms positional convergence and distributional convergence were introduced in Bayer and Charles (2018).

capture the full effects of relative skill return changes. This is because White workers with identical initial wages are not a good control group for Black workers. A change in one skill price (relative to other skill prices) can affect two workers with the same initial earnings differently depending on the exact mix of skills they possess ($\eta_{gk}+\phi_{ik}$'s), the level of discrimination they face (δ_{gk}^{total}'s), as well as the task requirements ($\tilde{\tau}_{jk}$'s) in the occupations they have sorted into. Hence, changes in relative skill (or task) returns can shift the percentile ranks of Black workers in the White earnings distribution therefore also causing positional convergence (or divergence). As we have documented throughout the paper, Black workers have lower *Abstract* skills, face higher discrimination in *Abstract* tasks, and as a result are less likely to be in occupations with high *Abstract* task content. Given that, a rising *Abstract* task return will on average benefit Black workers less than White workers with the same initial earnings and will therefor shift down their relative positions in the earnings distribution. To the extent that this force is ignored, measured distributional convergence understates the impact of the rising *Abstract* task return on the racial wage gap. By the same token, the shifting down of Black percentile ranks in the earnings distribution (due to the rising *Abstract* task return) will dampen any estimated gains Blacks have made in reducing racial wage rank gaps, so the positional convergence will also understate the effects of declining discrimination and narrowing racial skill gaps.

Table A2: Model Decomposition vs Juhn-Murphy-Pierce Decomposition

	Change in Wage Gap	Task Model Decomposition		JMP Decomposition	
Time Period	Total	β's/A's	$(\delta+\eta)$'s	Distributional Convergence	Positional Convergence
$1960-1980$	0.169	0.056	0.120	0.035	0.134
$1980-2018$	-0.007	-0.066	0.080	-0.034	0.027

Notes: Table shows counterfactual wage gaps using our task-based model and then separately the Juhn-Murphy-Pierce (JMP) decomposition. The first column shows the actual change in the Black-White wage gap during the given time period. The next two columns decomposes how much of the change in the wage gap is due to the changing β_{kt}'s and A_{jt}'s and how much is due to the changing $\delta_{bkt}+\eta_{bkt}$. The final two columns show the JMP decomposition where the distributional convergence refers to how much of the racial wage gap is due to the changing aggregate price of skill throughout the wage distribution. Positional convergence refers to the sifts in the relative positions of Blacks and Whites within the earnings distribution.

To illustrate the quantitative difference between our model and a model with one aggregate skill price, we compare the estimated effects of changing β's and A_j's and changing

δ's and η's presented in Figure 7 to what we would find if we did a Juhn-Murphy-Pierce (JMP) decomposition on the same model-generated data. We perform this comparison during two time periods: 1960-1980 and 1980-2018.[41] The results of this comparison are shown in Table A2. During the early period, our model based decomposition and the JMP decomposition yield very similar results.[42] This is not surprising given the results in Panel B of Figure 6 showing that there was no differential trend in task prices during the 1960-1980 period. When relative task prices evolve similarly, the implications of a one-skill model and a multi-task model are similar. However, during the post-1980 period, the JMP decomposition dramatically understates the importance of skill price changes in widening the racial wage gap relative to our model. In particular, we find that the changing task prices caused the racial wage gape to increase by 6.5 log points during this period while the JMP decomposition concludes that changing skill prices increased the racial wage gap by half that amount. Because the distributional convergence is understated relative to our model, the JMP model also substantially understates the importance of declining discrimination and the narrowing of racial skill gaps in improving relative Black wages during the last forty years. Collectively, the results in Table A2 highlight that analyzing racial wage gaps in a multi-skill task model can lead to quantitatively different conclusions relative to a standard JMP decomposition, particularly in periods when relative task prices are changing.

Appendix J Additional Model Wage Decompositions

In Table 6, we show estimates of racial skill differences η_{bkt} and taste-based discrimination intensities δ_{bkt}^{taste}. However, it is hard to see just from the table how much each parameter contributes to the overall racial wage gap. Appendix Table A3 uses the model to decompose the racial wage gap in 1960, 1990, and 2012 into various additional components. To do this, we re-solve the model setting the δ_{bkt}^{taste} for both *Contact* and *Abstract* tasks to zero (row 1 of all panels) or setting η_{bkt} for both *Contact* and *Abstract* tasks to zero (row 2 of all panels). In

[41] On the model-side, we fix β_{kt}'s (and A_{jt}'s) or $\delta_{bkt}^{total}+\eta_{bkt}$'s at the levels at the beginning of the period, and report the differences between the counterfactual racial wage gap thus computed and the actual racial wage gap at the end of the period as the estimated effects of changing β_{kt}'s (second column) and $\delta_{bkt}^{total}+\eta_{bkt}$ (third column), respectively. As for the JMP decomposition, we use the model-generated earnings distributions to compute the changes in the percentile rank of each Black worker in the White earnings distribution over each period, and estimate what their wages would have been at the end of the period if the White earnings distribution were fixed at the beginning of the period; the difference between the counterfactual racial wage gap thus computed and the racial wage gap at the beginning of the period gives an estimate for positional convergence (fifth column), while the difference between the actual racial wage gap at the end of the period and the counterfactual wage gap gives an estimate for distributional convergence (fourth column).

[42] Our decomposition does not exactly sum to the empirical change in the wage gap because of an unreported covariance term.

all cases, we hold all other model parameters fixed during these respective counterfactuals. We do note, however, that when we set $\eta_{bkt} = 0$, δ_{bkt}^{stat} will also be set to zero; by definition, there can be no statistical discrimination when the average racial skill gap is set to zero. We again estimate the model under three different assumptions about the extent to which employers noisily observe worker skills. The percent explained does not sum to 100 percent because we are not doing counterfactuals on the $\eta_{bkt} + \delta_{bkt}^{total}$ for *Routine* tasks nor are we reporting the covariance terms between the various η_{bkt}'s and δ_{bkt}'s.

A few results can be seen from Appendix Table A3. First, in a world where skills are measured without error ($\sigma = 0$), tasted-based discrimination in *Contact* and *Abstract* tasks explain 28 percent of the racial wage gap in 1960 and roughly 0 to 10 percent of the racial wage gap in 1990 and 2012. Conversely, racial differences in skills associated with these tasks explain 30 percent of the racial wage gap in 1960 and 70 to 85 percent of the racial wage gap in 1990 and 2012. Second, as skills are measured with more error, taste-based discrimination explains a smaller share of the racial wage gap in all years and differences in skills explain a larger fraction in all years. Third, regardless of our assumptions about the extent to which skills are measured with error, our model finds that a substantial amount of the current racial wage gap is due to racial differences in skills. Even when skills are perfectly measured, our model concludes that half of the current racial wage gap can be explained by a remaining racial skill gap in primarily *Abstract* tasks.

Appendix K Counterfactual Robustness to Alternate θ's and ψ's

We chose our base estimate of θ in part to match the racial wage gaps at different parts of the distribution. In this section of the appendix, we show the fit of different θ's in matching Figure A12 of the main text. We then show the robustness of many of our key counterfactual findings to alternate values of θ and ψ.

In Appendix Figure A15, we show how our model matches the racial gap in percentile ranks of wages with $\theta = 4$ and $\theta = 8$. As seen from the figure, lower levels of θ fit the distributional racial wage gaps less well. However, choosing a high θ (like $\theta = 8$) does not improve fit on this dimension in any meaningful way. This is partly why we chose $\theta = 6$ as our baseline parameterization.

Appendix Table A4 highlights that many of our key findings are quite robust to our choice of θ and ψ. The table shows the robustness of the results in Table 7 (row 1) and Panels A and B of 7 (rows 2 - 6). In column 1 of Appendix Table A4, we re-report our baseline results. In

Table A3: Decomposition of Racial Wage Gaps

	Counterfactual Wage Gap			Percent Explained		
	1960	1990	2012	1960	1990	2012
	(1)	(2)	(3)	(4)	(5)	(6)
Base Model	-0.40	-0.20	-0.22	–	–	–
Panel A: Signal To Noise = 1.0 ($\sigma = 0$)						
Setting $\delta_{bkt}^{taste} = 0$ ($k = Abstract, Contact$)	-0.29	-0.20	-0.19	28%	0%	12%
Setting $\eta_{bkt} = 0$ ($k = Abstract, Contact$)	-0.28	-0.03	-0.07	30%	84%	68%
Panel B: Signal To Noise = 0.9 ($\sigma = 0.17$)						
Setting $\delta_{bkt}^{taste} = 0$ ($k = Abstract, Contact$)	-0.30	-0.21	-0.20	26%	-5%	9%
Setting $\eta_{bkt} = 0$ ($k = Abstract, Contact$)	-0.27	-0.02	-0.06	32%	92%	72%
Panel C: Signal To Noise = 0.75 ($\sigma = 0.30$)						
Setting $\delta_{bkt}^{taste} = 0$ ($k = Abstract, Contact$)	-0.31	-0.22	-0.21	23%	-11%	3%
Setting $\eta_{bkt} = 0$ ($k = Abstract, Contact$)	-0.26	0.00	-0.04	36%	102%	80%

Note: Table shows the racial wage gap in our base model (row 1) as well the racial wage gap in various counterfactuals were we separately set the η_{bkt}'s and δ_{bkt}^{taste}'s for *Abstract* and *Contact* tasks to 0 for Black men. When setting the η's to zero, δ_{bkt}^{stat} will also equal zero by definition. The three panels show the counterfactuals under different assumptions about the which skills are noisily observed. Columns (1)-(3) show the level of the racial wage gap setting the various parameters to 0 in 1960, 1990, and 2012, respectively. Columns (4)-(6) show the share of the wage gap explained in counterfactual. We compute the share explained by computing the counterfactual wage gap relative to the base model wage gap and dividing that difference by the base model wage gap.

Figure A15: The Role of Theta in Matching the Racial Gap in Percentile Ranks of Wages

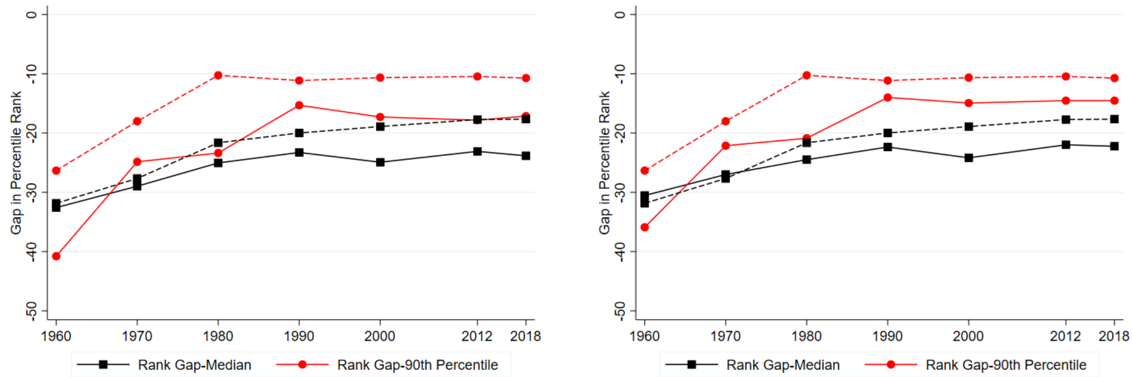

PANEL A: $\theta = 4$ PANEL B: $\theta = 8$

Notes: Figure reproduces the results from Figure A12 of the main text with different values of θ.

columns 2 and 3, we show the robustness of results when we set $\theta = 4$ and $\theta = 8$, respectively. In columns 4 and 5, we show the results for $\psi = 3.5$ and $\psi = 5.5$. Importantly, across all the robustness exercises, declining taste-based discrimination in *Abstract* and *Contact* tasks explains between 30 and 50 percent of the decline in the racial wage gap between 1960 and 2018. Likewise, across all the robustness specifications, the increasing returns to tasks post-1980 exacerbated the racial wage gap while declining δ_{bkt}'s and η_{bkt}'s narrowed the gap. In all cases, the two effects roughly offset each other such that racial wage gaps were constant during the 1980 to 2018 period.

Appendix L Counterfactual Robustness to Alternate First Stage Decomposition Projection

For our main decomposition of the composite gap in $\delta_{kt}^{total} + \eta_{kt}$ into the individual components of δ_{kt}^{taste}, η_{kt} and δ_{kt}^{stat}, we used the first stage projections in from Table 5 from the main text. Specifically, by estimating equation 10 for each task-specific skill, we produced a weighting (the b's) of each NLSY individual skill measure for each of the model task-specific skills. Our baseline projection, however, restricted the estimated b's from the first stage regression to be the same across all years. The purpose of this restriction was to increase the power in estimation by pooling the data for all three years. Yet, it is possible that the importance of each individual skill in performing certain tasks – the b's – changed over time, for example due to changing technology. As a robustness exercise, we re-estimate the first stage regressions separately using the 1979 NLSY data (which we map to model year 1990) and using the

Table A4: Key Result Robustness to Alternate Values of θ and ψ

	Base	θ		ψ	
		4	8	3.5	5.5
Share of 1960-2012 Wage Gap Explained					
Holding δ_{bkt}^{taste} fixed ($k = Abstract, Contact$)	48%	49%	47%	33%	48%
Counterfactual: Wage Change 1980-2018					
Holding $\beta_{Abstract}$ at 1980 Levels	0.100	0.073	0.119	0.106	0.096
Holding all β_k's and A_i's at 1980 Levels	0.066	0.049	0.073	0.058	0.062
Holding $\delta^{total} + \eta$ for *Contact* at 1980 Levels	-0.057	-0.055	-0.058	-0.027	-0.058
Holding $\delta^{total} + \eta$ for *Abstract* at 1980 Levels	-0.031	-0.028	-0.032	-0.034	-0.027
Holding $\delta^{total} + \eta$ for all tasks at 1980 Levels	-0.080	-0.067	-0.086	-0.047	-0.082

Note: Table shows the robustness of many of our key results to alternate values of θ and ψ. In our baseline model, we use $\theta = 6$ and $\psi = 5$. We replicate the results from our baseline model in column 1. We show the robustness of our results to values of $\theta = 4$ and $\theta = 6$ (columns 2 and 3) and values of $\psi = 4$ and $\psi = 6$ (columns 4 and 5). The first row replicates results results from row 1 of Table 7 showing the percent of the actual decline in the racial wage gap between 1960 and 2012 explained by the falling δ^{taste} for *Contact* and *Abstract* tasks. The remaining rows show the counterfactual changes in the racial wage gap between 1980 and 2018 from the various counterfactual exercises done in Figure 7.

1997 NLSY data (which we map to model year 2012).

Appendix Table A5 shows the coefficients in the alternate first stage regressions where we allow the b's to vary by year.[43] Three things are of note in the table. First, the power of the alternate first-stage regressions is lower than that of the baseline regressions. For example, the F-statistics in the baseline specification are 21 for *Abstract* and 10 for *Contact*; the respective F-stats in the alternate specification are roughly around 12 and 9 in 1990 and 2012, and they are even lower in 1960. Second, however, the estimated b's are roughly stable over time. In all years, cognitive skills strongly predict *Abstract* task efficiency and social skills strongly predict *Contact* task efficiency. The stability of coefficients over time justifies our restriction in the baseline specification that the b's are constant across years. Last, that said, the alternate first stage regressions seem to suggest that the weighting on cognitive skills for *Abstract* task efficiency increased between 1990 and 2012 (from 0.13 to 0.19). Of course, given the lower power of the alternate regressions, it is hard to tell whether the estimated trends reflect actual changes in the structure of the economy or they solely capture noise. Nevertheless, we can proceed with the decomposition with the alternative

[43] We omit the results for *Routine* task skills because the power is too low for the results to be credible.

Table A5: First Stage Regression of Average Model Task Skills on Average NLSY Individual Skills, Cross-Occupation Variation

	Pane A: 1960		Pane B: 1990		Panel C: 2012	
	Abstract	*Contact*	*Abstract*	*Contact*	*Abstract*	*Contact*
Cognitive	0.14	0.03	0.13	0.02	0.19	0.05
	(0.09)	(0.03)	(0.04)	(0.03)	(0.03)	(0.01)
Non-Cognitive	0.04	0.05	0.09	0.05	0.11	-0.02
	(0.07)	(0.03)	(0.07)	(0.04)	(0.04)	(0.02)
Social	-0.01	0.11	0.02	0.18	-0.04	0.11
	(0.11)	(0.04)	(0.09)	(0.06)	(0.08)	(0.03)
Constant	1.10	1.11	1.03	1.09	1.04	1.10
	(0.02)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
Adj. R-Squared	0.30	0.34	0.46	0.45	0.54	0.43
F-Stat	4.95	3.00	10.5	8.2	15.0	9.9

Notes: Table shows estimate coefficients from first stage regression equation (10) for White men. Each column is a separate regression exploiting cross-occupation variation. We use 66 broad occupation categories. For these regressions, we estimate these regressions separately for 1960 (using data in the South region from the 1979 NLSY), 1990 (using data from the 1979 NLSY) and for 2012 (using data from the 1997 NLSY). See the text for additional details.

weightings and assess their implications in case that these trends are real.

Appendix Table A6 shows estimates of δ_{kt}^{taste} and η_{kt} in different years calculated using our alternate first-stage specification. For illustrative purposes, we only show the results assuming employers observe worker skills without noise ($\sigma = 0$). Importantly, the estimated decline in η for *Abstract* tasks between 1990 and 2012 is much smaller under the alternate first-stage specification. In the baseline specification, η for *Abstract* tasks declines from -0.26 to -0.18 over the period, reflecting the narrowing cognitive skill gap observed in the NLSY data; in the alternate specification, however, it declines from -0.24 to -0.22 only. This is because an increase in the importance of cognitive skills in performing *Abstract* tasks – which the alternative first-stage regressions imply over the 1990-2012 period – disadvantages Blacks relative to Whites with respect to efficiency in performing *Abstract* tasks given that Black men have a relative deficit in cognitive skills. Hence, despite the significant decline in the cognitive skill gap between 1990 and 2012, the racial gap in *Abstract* task-specific skills declined little over the period. Put differently, the increase in the importance of cognitive skills for *Abstract* tasks might be another force that offset the decline in race-specific barriers and caused the stagnation of the racial wage gap over the last couple of decades, along with the rising return to *Abstract* tasks we discuss in Section 6. If this is the case, then declining taste-based discrimination accounts for almost all the decline in $\delta + \eta$ for *Abstract* tasks over the 1990-2012 period.

Table A6: Task Decomposition of Racial Skill Gap and Task-Based Discrimination, Alternate First Stage Regression ($\sigma = 0$)

	Contact				*Abstract*			
	1960	1990	2012	$\Delta(12-60)$	1960	1990	2012	$\Delta(12-60)$
δ^{taste}, $\sigma = 0$	-0.18	-0.00	-0.05	0.13	-0.14	0.02	0.03	0.16
η, $\sigma = 0$	-0.11	-0.08	-0.03	0.08	-0.24	-0.24	-0.22	0.02

Notes: Table shows model decomposition of racial differences in δ_{bk}^{taste}, and η_{bk} for *Abstract* and *Contact* task in 1960, 1990, and 2012 using our alternate first stage regression where the b's are allowed to differ over time. We do the decomposition assuming employers can accurately observe the skill levels of their employees ($\sigma = 0$).

Appendix M Additional Model Validation Using Cross-Region Analysis

To further highlight the importance of declining taste-based discrimination in explaining the racial wage convergence, we show one additional set of results in this appendix section exploiting regional variation in the micro data from the Census/ACS. Appendix Figure A16 compares the cross-regional difference in racial wage gaps to the cross-regional differences in various racial task gaps. Specifically, the solid black line (with squares) shows the convergence of racial wage gaps (conditional on education) in the South relative to the non-South regions since 1960 using the Census/ACS data (left axis). This line is just the difference between the two lines in Figure 8 discussed above. The other two lines in Figure A16 show the difference in the racial gaps in *Contact* and *Abstract* tasks between the South and non-South regions in each year (right axis). This data is just the difference between the patterns shown in Panels A and B of Figure 3.

Three facts emerge from the figure. First, the racial wage gap in the South fell sharply relative to the racial wage gap in the non-South region consistently during the 1960 to 2018 period. Specifically, the racial wage gap in the South was roughly 20 log points higher than that in the non-South in 1960, but by 2018, the racial wage gap in the South essentially converged to that in the non-South region. Second, the racial gap in *Contact* tasks narrowed

Figure A16: Trends in Racial Wage and Task Gaps by Region: Census/ACS Data

Notes: Figure shows *differential* racial gaps in wages and tasks between the South region and non-South region over time. The black line (with squares) is the difference between the two racial wage gap series in Figure 8. The dashed red line (with squares) shows the difference in the racial gap in *Contact* tasks over time while the dashed blue line (with triangles) shows the difference in the racial gap in *Abstract* tasks over time. The the data from the latter two series come from differencing the patterns in Panels A and B of Figure 3.

sharply in the South relative to the non-South region over the same period, from a 6 percentage point gap in 1960 to a 2 percentage point gap in 2012, closely tracking the convergence of racial wage gaps across the two regions. Last, on the other hand, there was almost no change in the racial gap in *Abstract* tasks between the South and non-South regions during this time period. Given the model results in Section 8.1 suggesting that the trends in *Contact* task gaps are a good proxy for trends in taste-based discrimination, the patterns imply that the faster decline in taste-based discrimination in the South region primarily drove the convergence of the racial wage gaps across the two regions. Overall, the regional result reinforces the model finding that the declining taste-based discrimination was the primary force behind the racial wage gap trends post-1960.

Appendix N Additional Model Results

This section of the appendix provides details on additional model results.

Appendix N.1 Conditional Wages and Statistical discrimination

In this section, we derive expressions for conditional wages and statistical discrimination when skills are observed with noise by employers. Employers set the wage of each worker at the worker's expected marginal revenue product given observed skills $(\hat{s}_1, ..., \hat{s}_K)$. Since p_j

and y_{gij} are in logs, the conditional expectation of marginal revenue product is given by

$$E\left[e^{p_j+y_{gij}} \mid (s_{g1},...,s_{gK}) = (\hat{s}_1,...,\hat{s}_K)\right] = E\left[\exp\left(A_j + \sum_K \lambda_{jk}(\phi+\delta_{gk}+\eta_{gk})\right) \mid (s_{g1},...,s_{gK}) = (\hat{s}_1,...,\hat{s}_K)\right].$$

Taking the logs and noting the independence of ϕ_{ik}'s and ϵ_{ik}'s from one another, we obtain that the conditional wage is given by

$$w_{gij}^{cond}(\hat{s}_1,...,\hat{s}_K) = A_j + \sum_K \log E\left[e^{\lambda_{jk}(\phi+\eta_{gk})} \mid s_{gk} = \hat{s}_k\right] + \sum_K \lambda_{jk}\delta_{gk}^{taste},$$

which yields the expression in the text with some rearrangements.

Next, we derive the expression for $E\left[e^{\lambda_{jk}(\phi+\eta_{gk})} \mid s_{gk} = \hat{s}_k\right]$. Define f_{ϕ_K} and f_{ϵ_k} to be the probability density functions for ϕ_{ik} and ϵ_{ik}, respectively. The joint density function for s_k and ϕ_k is then given by

$$f_k(s,\phi) = f_{\phi_k}(\phi)f_{\epsilon_k}(s-\eta_k-\phi).$$

Thus, we have

$$E\left[e^{\lambda_{jk}(\phi+\eta_{gk})} \mid s_{gk} = \hat{s}_k\right] = \frac{\int_0^\infty e^{\lambda_{jk}(\phi+\eta_{gk})}\, f_{\phi_k}(\phi)f_{\epsilon_k}(\hat{s}_k-\phi-\eta_{gk})d\phi}{\int_0^\infty f_{\phi_k}(\phi)f_{\epsilon_k}(\hat{s}_k-\phi-\eta_{gk})d\phi}.$$

We evaluate the integrals numerically.

Lastly, we highlight one interesting fact about the statistical discrimination term. Note that the statistical discrimination term can be written as

$$\delta_{gk}^{stat}(\eta_{gk},\lambda_{jk},\hat{s}_k) = \eta_{gk} + \left[\phi_{g_0ik}^e(\lambda_{jk},\hat{s}_k-\eta_{bk}) - \phi_{g_0ik}^e(\lambda_{jk},\hat{s}_k)\right], \tag{A3}$$

where we defined $\phi_{g_0ik}^e(\lambda_{jk},\hat{s}_k) = \log E\left[e^{\lambda_{jk}\phi} \mid s_{g_0k} = \hat{s}_k\right]^{1/\lambda_{jk}}$ to be the log expected efficiency of a base-group worker conditional on observing $\hat{s}_k$. The expression shows that the statistical discrimination term depends on (i) the difference in group means of effective skills (η_{gj}) and (ii) the expected within-group position (ϕ_{ik}) of each worker in the skill distribution conditional on observed skills (the term in brackets). To see this, note that the position of $\hat{s}_k$ in the observable skill distribution for group g is equivalent to the position of $\hat{s}_k - \eta_{bk}$ in the observable skill distribution of the base group. Thus, supposing $\eta_{gk} < 0$, the expected ϕ_{ik} is higher for workers of group g than for workers of the base group with the same observed skills (assuming a finite σ). The term in brackets captures how much higher the ϕ_{ik}'s of workers of group g is expected to be relative to those of base group workers with the same observed skills. For example, skills are observed perfectly by employers, then the expected ϕ_{ik} for workers of group g exceeds that of the base group workers with the same observed skills by $-\eta_{gk}$ and completely offsets the gap in mean skill levels η_{gk}. If, on the other hand, observed skills are so noisy and give no information about true levels of worker skills, then no difference in ϕ_{ik}'s can be expected between groups even conditional on observed skills (i.e., the term in brackets equals zero) and workers are paid entirely based on group means. Said differently, the term in brackets reflects how much the information from observed skills can compensate for the known gaps in the mean human capital levels across groups.

Appendix N.1.1 Proof of Proposition 1

We prove Proposition 1 in two steps. First, it is clear from equation A3 that $\delta^{stat}_{gk}(\eta_{gk}, \lambda_{jk}, \hat{s}_k) = 0$ if $\eta_{gk} = 0$. Second, the continuity of $\phi^e_{g_0 ik}(\lambda_{jk}, \hat{s}_k)$ in $\hat{s}_k$ implies that $\delta^{stat}_{gk}(\eta_{gk}, \lambda_{jk}, \hat{s}_k)$ is continuous in η_{gk}. Thus, $\delta^{stat}_{gk}(\eta_{gk}, \lambda_{jk}, \hat{s}_k)$ tends to zero as $\eta_{gk} \to 0$, as desired. $\square$

Appendix N.2 Home Sector and Reservation Utility

The model allows for the possibility that workers may choose to work in the home sector, denoted as $j = H$. Specifically, we treat the home sector as another potential occupation (with task requirements $\lambda_{H1}, ..., \lambda_{HK}$ and occupational return A_{gH}) where the returns are non-pecuniary. That is, the reservation utility u_{giH} of a worker with given observable credentials equals the log of the worker's expected marginal revenue product in an occupation with task requirements $(\lambda_{H1}, ..., \lambda_{HK})$ plus the log of the idiosyncratic preference for home sector, ν_{iH}:

$$u_{giH} \equiv A_{gH} + \sum_K \lambda_{Hk} \left(\phi^e_{g_0 k}(\lambda_{Hk}, \hat{s}_{ik}) + \delta^{taste}_{gk} + \delta^{stat}_{gk}(\eta_{jk}, \lambda_{Hk}, \hat{s}_{ik}) \right) + \log \nu_{giH},$$

where ν_{giH} is the idiosyncratic preference for home sector. Like the other occupational preferences ν_{gij}, the preference for home sector ν_{giH} follows a Frechet distribution with shape ψ and scale 1. Note that the term A_{gH} plays the role of a scale parameter of the Frechet distribution. We allow A_{gH} to differ by race, unlike with A_j's. The differences in A_{gH}'s across groups capture any forces other than differential task returns that may create labor supply differences between racial groups.

Appendix N.3 Employment Share of Occupations

In this section, we derive an expression for the employment share of each occupation. Recall that each worker i of race group g chooses occupation j by maximizing the utility given by

$$u_{gij} = w^{cond}_{gj}(s_{gi1}, ..., s_{giK}) + \log \nu_{ij}.$$

where the occupational preference ν_{ij} follows a Frechet distribution with scale 1 and shape ψ. Let f and F be the pdf and cdf of the Frechet distribution, respectively. Among workers with the same observed skills $(s_{gi1}, ..., s_{giK}) = (\hat{s}_1, ..., \hat{s}_K)$ and of the racial group g, the proportion of those who chooses occupation j is given by

$$\begin{aligned}\rho_{gj}(\hat{s}_1, ..., \hat{s}_K) &= \Pr\left[\exp\left(w_{gj}^{cond}(\hat{s}_1, ..., \hat{s}_K)\right)\nu_j > \exp\left(w_{gj'}^{cond}(\hat{s}_1, ..., \hat{s}_K)\right)\nu_{j'},\ \forall j' \neq j\right] \\ &= \int f(\nu)\cdot \Pi_{j'\neq j} F\left(\exp\left(w_{gj}^{cond}(\hat{s}_1, ..., \hat{s}_K) - w_{gj'}^{cond}(\hat{s}_1, ..., \hat{s}_K)\right)\nu\right) d\nu.\end{aligned}$$

Defining

$$\omega_{gj} = \sum_{j'=1,...,J,H} \exp\left(\psi w_{gj'}^{cond}(\hat{s}_1, ..., \hat{s}_K) - \psi w_{gj}^{cond}(\hat{s}_1, ..., \hat{s}_K)\right),$$

we obtain

$$\begin{aligned}\rho_{gj}(\hat{s}_1, ..., \hat{s}_K) &= \int \psi\nu^{-\psi-1}\cdot\exp\left[-\omega_g \nu^{-\psi}\right] d\nu \\ &= \frac{1}{\omega_{gj}}.\end{aligned}$$

Appendix N.4 Elasticity of Labor Supply

Lastly, we derive an expression for labor supply elasticity and show how targeting the moment can help us pin down ψ, the shape parameter for the Frechet distribution from which we draw idiosyncratic preferences for occupations.

Given the expression for the employment share in each occupation, the aggregate labor supply for race group g is given by

$$L_g = \int_{S_g} \left(1 - \rho_{gH}(\hat{s}_1, ..., \hat{s}_K)\right),$$

where the integral is taken over the support of $(\hat{s}_1, ..., \hat{s}_K)$, denoted by S_g. We want to find out the labor supply changes in response to one percent increase in wages for all occupations (excluding, of course, the home sector). To this goal, write

$$\exp\left(w_{gj}^{cond}(\hat{s}_1, ..., \hat{s}_K)\right) \equiv W\cdot\exp\left(w'^{cond}_{gj}(\hat{s}_1, ..., \hat{s}_K)\right)$$

for all $j' \neq H$ for some constant W. The labor supply elasticity for group g is then given by

$$\varepsilon_g = -\left[\int_{S_g} \frac{\partial \rho_{gH}(\hat{s}_1, ..., \hat{s}_K)}{\partial W}\right]\cdot\frac{W}{L_g}.$$

Using the results from the previous section, we obtain

$$\frac{\partial \rho_{gH}(\hat{s}_1, ..., \hat{s}_K)}{\partial W} = W^{-1}\psi\rho_{gH}(\hat{s}_1, ..., \hat{s}_K)\left(1 - \rho_{gH}(\hat{s}_1, ..., \hat{s}_K)\right),$$

Hence, the labor supply elasticity is given by

$$\begin{aligned}
\varepsilon_g &= \psi \frac{\int_{S_g} \rho_{gH}(\hat{s}_1, ..., \hat{s}_K)(1 - \rho_{gH}(\hat{s}_1, ..., \hat{s}_K))}{\int_{S_g}(1 - \rho_{gH}(\hat{s}_1, ..., \hat{s}_K))} \\
&= \psi\left[1 - \frac{\int_{S_g}(1 - \rho_{gH}(\hat{s}_1, ..., \hat{s}_K))^2}{\int_{S_g}(1 - \rho_{gH}(\hat{s}_1, ..., \hat{s}_K))}\right].
\end{aligned}$$

Note that the expression for ε_g involves the parameter ψ, the shape parameter for the Frechet distribution from which we draw idiosyncratic preferences for occupations. Hence, targeting the elasticity of labor supply helps us discipline the parameter ψ.