

Forth, John; Theodoropoulos, Nikolaos

Working Paper

Earnings Discrimination in the Workplace

IZA Discussion Papers, No. 15357

Provided in Cooperation with:

IZA – Institute of Labor Economics

Suggested Citation: Forth, John; Theodoropoulos, Nikolaos (2022) : Earnings Discrimination in the Workplace, IZA Discussion Papers, No. 15357, Institute of Labor Economics (IZA), Bonn

This Version is available at:

<https://hdl.handle.net/10419/263573>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

DISCUSSION PAPER SERIES

IZA DP No. 15357

Earnings Discrimination in the Workplace

John Forth
Nikolaos Theodoropoulos

JUNE 2022

DISCUSSION PAPER SERIES

IZA DP No. 15357

Earnings Discrimination in the Workplace

John Forth

University of London and IZA

Nikolaos Theodoropoulos

University of Cyprus

JUNE 2022

Any opinions expressed in this paper are those of the author(s) and not those of IZA. Research published in this series may include views on policy, but IZA takes no institutional policy positions. The IZA research network is committed to the IZA Guiding Principles of Research Integrity.

The IZA Institute of Labor Economics is an independent economic research institute that conducts research in labor economics and offers evidence-based policy advice on labor market issues. Supported by the Deutsche Post Foundation, IZA runs the world's largest network of economists, whose research aims to provide answers to the global labor market challenges of our time. Our key objective is to build bridges between academic research, policymakers and society.

IZA Discussion Papers often represent preliminary work and are circulated to encourage discussion. Citation of such a paper should account for its provisional character. A revised version may be available directly from the author.

ISSN: 2365-9793

IZA – Institute of Labor Economics

Schaumburg-Lippe-Straße 5–9
53113 Bonn, Germany

Phone: +49-228-3894-0
Email: publications@iza.org

www.iza.org

ABSTRACT

Earnings Discrimination in the Workplace

This paper provides an overview of theory and empirical evidence on earnings discrimination within the workplace. Earnings discrimination occurs when employees producing work of equal value are differentially remunerated because of their social group. The paper reviews theories of why employers may discriminate in this way. The paper then goes on to review research evidence on earnings discrimination as one source of earnings inequality within the workplace. The ability of empirical studies to identify discrimination is discussed, and evidence on the mechanisms through which discrimination may affect earnings is reviewed, covering observational and experimental studies. The research evidence is most plentiful in respect of discrimination by gender. Accordingly, much of the discussion focuses on the role of discrimination in driving a wedge between the wages of men and women. However, the paper also reviews evidence on earnings discrimination by race or ethnic group. It concludes with a discussion of policy responses.

JEL Classification: J71, J31

Keywords: earnings discrimination, gender wage gap, ethnic wage gap

Corresponding author:

Nikolaos Theodoropoulos
Department of Economics
University of Cyprus
Nicosia
Cyprus

E-mail: n.theodoropoulos@ucy.ac.cy

1. Introduction

Most countries in the world have legislation to outlaw wage discrimination. Historically, the focus has been on discrimination by gender, seeking to enforce the principle of equal pay for men and women undertaking work of equal value – as outlined in the 1951 International Labor Office (ILO) Convention on Equal Remuneration. However, many countries have extended laws on earnings discrimination to cover other social identities, including ethnicity. These laws seek to prevent an employer from using social group membership – either purposefully or inadvertently – as criteria when setting the wage of an individual or group of employees. Despite the existence of these legal restrictions, a minority of employees continue to report, in surveys, that they have been unfairly treated because of their gender or ethnic group, and case law confirms that wage discrimination does indeed persist in some workplaces (e.g. Moss, 2019; Bachmann, 2022).

Economists then face the challenge of explaining why social group membership might feature unfairly in employers' wage setting behavior. They also face the challenge of gauging its economic significance and identifying the arenas of wage setting where it appears to be most prevalent. Meeting these challenges enables economists to identify appropriate policy solutions.

The challenge of explaining why employers might discriminate began to be addressed half a century ago, when Becker (1971), Arrow (1973) and others began to develop theories of discrimination. Whilst Becker focused on employers' tastes or preferences for different types of workers, other theories focused on the role of beliefs or stereotypes in shaping employer behavior.

In the empirical domain, economists have long documented the existence of wage differentials between social groups. The major development in recent decades has been to show that differentials do not merely exist between workers in different firms, but also between observationally-equivalent co-workers within the same firm and, in some instances, within the same job. It is conventional in empirical studies of earnings inequality to label this "earnings discrimination" although, in observational studies there is inevitably some ambiguity as to whether this is true discrimination or whether it might be explained by unobserved differences in employee characteristics. Experimental evidence helps to address this concern, but the difficulty of creating experimental conditions within firms, and subsequent reliance on laboratory experiments, inevitably leads to questions about generalisability.

Empirical studies have also sought to identify the mechanisms through which earnings discrimination is manifest, focusing variously on employers' wage-setting behavior for new hires, the renegotiation of wages on promotion, and processes of performance evaluation and task allocation which may interact with wage determination. Within organizations, much of the focus has been on providing a level playing field in these areas of the employment relationship: seeking to design hiring, performance appraisal and promotion practices so that opportunities for managers to discriminate are squeezed out. However, economists and organizations have each become increasingly aware of the subtle factors – ability to work long hours, differential access to developmental opportunities and so on – that may prevent employees from entering the playing field on equal terms.

This paper provides an overview of theory and empirical evidence on earnings discrimination in the workplace. The paper begins by reviewing theories of discrimination (Section 2). It then goes on to examine evidence on the existence of earnings discrimination by reviewing studies that seek to identify the existence of wage differentials between co-workers (Section 3). The capability of such studies to identify discrimination is discussed, and other explanations for observed wage differences are explored (Section 4). Studies which seek to identify the mechanisms for earnings discrimination are then reviewed (Section 5). As one might expect, the research evidence on wage inequality and wage discrimination is most plentiful in respect of gender. Accordingly, much of the discussion focuses on earnings discrimination as a contributor to the gender wage

gap. However, the paper also reviews evidence on earnings discrimination by ethnic group. The paper concludes with a discussion of policy responses (Section 6).

The paper focuses on discrimination by employers in wage setting for their employees. Discrimination in payments made by customers to self-employed workers are therefore out of scope. However, the discussion recognises, more broadly, that customer pressure may be one of the factors that incentivizes an employer to discriminate.

One limitation of the review is that much of the evidence that is covered concerns North America and Western Europe. The issue of geographical generalisability must be approached with caution due to variation in social norms, legislative context and institutions. Generalising from the empirical literature on earnings discrimination by ethnicity is acknowledged to be particularly hazardous, since the identity of the minority ethnic groups typically differs between countries.

2. Theories of earnings discrimination

Economists have developed two main theories to explain why employers might treat co-workers differently. One set of ideas focuses on employers' tastes or preferences for different types of workers, whilst the other focuses on the role of beliefs or stereotypes.

Becker's (1971) classic model of discrimination argues that economic agents may have tastes or preferences that generate a desire to minimize economic transactions with members of a specific social group (so-called 'taste-based discrimination'). Prejudiced employers will only hire an employee from this non-preferred social group if its members will accept a discounted wage that compensates the employer for the disutility of admitting them into the workforce. The prejudicial tastes of incumbent employees or customers may also influence employer decision-making: employees may demand a wage premium to work alongside employees from the non-preferred group; customers or clients may also be reluctant to interact with employees from the non-preferred group, thus depressing their value to the employer.

The alternative models of Arrow (1973) and Phelps (1972) focus instead on information-deficits. These models propose that employers may use expectations about the average productivity of a social group (correct or otherwise) to make probabilistic judgements about the productivity of individual workers. If the employer perceives that one social group is less productive than another then, in the absence of information to counter the perception, reliance on the stereotype will drive a wedge between the wage offers made to members of each group (so-called 'statistical discrimination'). Such probabilistic judgements are most likely to be made when the employer has limited information on employee productivity, as may be the case when hiring. Wage differentials may then disappear with tenure, as the employer receives updated information; this may in turn reduce the chances of statistical discrimination in hiring wages for the next cohort. Cognitive biases, such as confirmation bias, mean that stereotypical perceptions are difficult to change (Kahnemann, 2011), but they may be more malleable than tastes.

Evidence that supports one or other model of discrimination is discussed later. However, both types of discrimination are expected to be limited by competitive forces. In perfectly competitive product markets, the incidence of wage discrimination is likely to disappear in the long run via competitive pressure from other firms who are making full use of available talent. Wage discrimination is then only expected to persist where product market imperfections allow employers to continue to bear the cost of their discriminatory behavior. Similarly, in perfectly competitive labor markets, job applicants will go in search of unprejudiced working environments, leading to the segregation of social groups into different firms (and no wage differences). In the presence of imperfections such as constraints on worker mobility, on the other hand, members of the social group may accept the discounted wage on offer, thus leading to discriminatory wage differentials within firms.

Policy makers can then affect the incidence of earnings discrimination in a number of ways. They may affect it indirectly by either promoting product market competition, appealing to consumers to reduce their consumption of the goods and services of discriminatory firms, or reducing constraints facing job seekers; they can also affect it directly by imposing costs on such firms via fines.

Employers seeking to eliminate discrimination by managers can seek to limit the space for tastes or stereotypes to affect personnel decisions: for example, by ensuring that wage-setting decisions are taken by committee rather than by individuals. They can also reduce the influence of stereotypes by improving the information set that informs personnel decisions: for example, by ensuring the wage renegotiations are informed by rigorous evaluations of performance. Such actions seek to create a level playing field, in which the opportunities for social group membership to affect decisions are minimized. Many organizations have adopted such approaches, recognising not only that some managers may be explicitly prejudiced, but also that well-meaning managers may sometimes make decisions based on unconscious biases.

Even when the playing field is apparently level, however, earnings differences may arise through what may be referred to as “subtle barriers” (Blau and Kahn, 2017: 828), whereby institutional constraints make it harder for the members of a particular social group to meet the common standards that are required to obtain a particular wage. In case law, this is referred to as “indirect discrimination”. It may occur, for instance, when productivity is unaffected by flexible working, but promotion is preferentially awarded to those working standard hours, who are more likely to be male; alternatively, when wage progression is influenced by networking within the firm, and such networks are less welcoming to women or ethnic minorities. In either example, wage differences may be induced between social groups unrelated to their abilities. Appropriate policy responses are more likely to focus on challenging norms and using personnel data to identify instances in which personnel processes unwittingly disadvantage certain groups.

3. Evidence of earnings inequality between social groups

One important strand of the voluminous body of research on earnings inequality between social groups has tended to focus on the supply-side of the labor market, and in particular differences in individuals’ endowments of productivity-enhancing characteristics, such as qualifications or experience (see Blau and Kahn, 2017, and Longhi, 2020, for two reviews). This body of research, often based on data from surveys of individuals, has had less to say about the demand side of the labor market. Such studies have typically examined the relevance of segregation by occupation, industry sector or region in explaining wage differentials between groups (*ibid.*), but they have tended to lack any more comprehensive understanding of how the firm contributes to wage inequality.

The availability of linked employer-employee data (LEED) has enabled researchers to gain important additional insights, however. The key feature of LEED is that it typically enables the researcher to observe a sample of firms and to observe the wages and (some) characteristics of multiple employees in each of those firms. One is then able to decompose the aggregate wage inequality at a given point in time into two components. The first is that part of aggregate wage inequality that can be explained by the differential segregation of workers across firms. Such segregation will contribute to aggregate wage inequality if there is a firm-specific component to wage levels – arising perhaps from differences in firm-level productivity – which cause average wages in one firm to be higher than those in another. The second component is that part of aggregate wage inequality that can be explained by wage differences between co-workers within the same firm. Criscuolo et al. (2020: 10-11) provide cross-country evidence on the relative contributions of “between-firm” and “within-firm” wage differences to aggregate wage inequality within this framework, showing that the dispersion of average wages between firms, accounts for about half of aggregate wage inequality.

Extensions to this approach go on to estimate human capital earnings functions with firm fixed-effects, thereby obtaining estimates of the dispersion in residual wages after accounting for differences in the observable component of employee productivity and differences in average wages across firms (e.g. Barth et

al., 2016). One limitation of these models is that the component of wage inequality attributed to differences between firms may be biased if unobserved differences between employees cause high-wage workers to sort into high-wage firms. Hence, further extensions use longitudinal LEED to estimate two-way or higher order fixed-effects models that can account for the unobserved characteristics of both workers and firms (Abowd et al., 1999), match-specific effects between workers and firms (Woodcock, 2008) and firm origin and destination effects (Di Addario et al. 2022). One can then obtain a fuller picture of the contributions of worker, firm, and match-specific heterogeneity to residual wage inequality.

Employer discrimination against particular social groups can contribute to either the “between-firm” or “within-firm” components of wage inequality. Specifically, discrimination in hiring and/or firing could depress the relative wages of women if it causes them to be sorted disproportionately into low-wage firms (a “between-firm” component). Discrimination in wage setting could also depress relative wages if it causes women to receive lower pay than otherwise-equivalent men at the same firm (a “within-firm” component). Identical propositions can be made in respect of wage differences between employees from a majority ethnic group and those from a minority group.

The main concern of this paper is earnings discrimination within the workplace, and hence the role of the “within-firm component” in explaining wage differentials between social groups. The discussion begins by examining its contribution to the gender wage gap.

Within-firm gender wage gaps

Most of the early studies of the gender wage gap used LEED data at the cross-sectional level and, hence, took the approach of estimating human capital earnings functions with firm (or workplace) fixed-effects, seeking to isolate that part of the gender wage gap that exists between observationally-equivalent male and female co-workers after removing the contribution of average wage differentials between firms.

Bayard et al. (2003) is one such study for the United States. This study used LEED data on all sectors of the US economy formed by linking the 1990 population census to establishment records. The analysis showed that the raw gender pay gap reduced substantially in size after accounting for the segregation of women into lower-paying occupations, industries, establishments, and occupations within establishments. However, a residual gender earnings differential of 13 percentage points persisted even after controlling for differences in human capital and these various aspects of segregation. Many other studies have gone on to use the same broad approach with cross-sectional data, estimating residual gender earnings gaps after controlling for workplace or firm fixed-effects. Studies in this vein typically find that women are sorted into low-wage firms, but that substantial wage gaps exist after accounting for segregation. These studies have reported residual, within-firm gender wage gaps of 10-11 percentage points for Britain (Drolet and Mumford, 2012; Heywood and Theodoropoulos, 2020; Theodoropoulos et al., 2022), 11-16 percentage points for Canada (Drolet and Mumford, 2012; Javdani, 2015); 12-15 percentage points for Australia (Meng, 2004) and 13 percentage points at the median, rather than the mean, for Japan (Hara, 2018). These estimates are remarkably consistent.

More recently, researchers have begun to use high-dimensional LEED to account for previously unobserved sources of heterogeneity. Studies such as Card et al. (2016) have estimated the two-way fixed-effects models mentioned earlier to arrive at a comprehensive determination of the share of the raw gender wage gap that is accounted for by between and within-firm variance. These studies, which include derivatives by Sorkin (2017), Coudin et al. (2018), Bruns (2019), Gallen et al. (2019), and Casarico and Lattanzio (2019), and which cover a variety of European countries and the United States, are consistent in showing that the majority of the raw gender wage gap is attributable to variance between workers, rather than variance between firms (sorting). In a variant of this approach, Cardoso et al. (2016) and Jewell et al. (2020) seek to estimate wage equations that include worker, occupation and firm fixed effects, to account for gender-occupation sorting within firms. Using data for Portugal, Cardoso et al. (2016) find that 14 log points of the 23 log point conditional wage gap exists within jobs and firms for workers of the same age and firm seniority (Cardoso et al., 2016: Table 1). Using data for Britain, Jewell et al. (2020: Table 4) find that 15 log points of the 17 log point conditional wage gap exists within occupations and firms.

Within-firm ethnic wage gaps

Similar studies of the ethnic wage gap are much less prevalent, no doubt partly because of data limitations. Few administrative sources of LEED include data on ethnicity (indeed, in some countries, administrative recording of employees' ethnicity is prohibited), whilst the numbers of ethnic minority employees observed in LEED derived from surveys is often relatively small.

One of the first studies to look at the role of the firm in the ethnic wage gap was Carrington and Troske's (1998) study of the manufacturing sector in the United States covering the late 1980s and early 1990s. They regress wages on a set of workplace fixed effects after controlling for workers' personal characteristics, finding that none of the black/white gap in residual wages is accounted for by sorting into lower-paying workplaces. In fact, black workers appear to be disproportionately employed in higher-paying workplaces. After accounting for workplace fixed-effects, the within-workplace ethnic wage gap is around six percentage points for men and around three percentage points for women.

Hellerstein and Neumark (2008) draw similar conclusions about the nature of black/white wage differentials in the US from their analysis of matched employer-employee data for 1990. Their analysis, which extends beyond manufacturing to include services, but which has access to fewer employee controls, also finds that black workers are disproportionately employed in higher-paying workplaces. Within workplaces, the black wage penalty relative to whites who are similarly educated is 16 log points. Hellerstein and Neumark (*ibid.*) provide contrary evidence for Hispanics, finding that this group tend to be over-represented in lower-paying workplaces. However, this accounts for less than one-tenth of the overall wage gap. Again, a substantial within-workplace wage penalty (22 log points, conditioning on language skills), is the main contributor to aggregate wage differentials.

Similar studies have been undertaken for Canada by Pendakur and Woodcock (2010) and for Britain by Forth et al. (2021). In common with the US evidence for blacks, Forth et al. (2021) find that ethnic minorities in Britain tend to be over-represented in higher-paying workplaces – a feature that is most evident for women. Non-white male employees then earn, on average, around nine per cent less than observationally equivalent white employees after accounting for wage differences across workplaces. Among female employees, the within-workplace wage penalty for non-whites is around seven per cent on average. Pendakur and Woodcock (2010) focus on wage outcomes for Canadian-born visible minorities. They find that men from minority groups are slightly over-represented in higher-wage firms, but face a substantial within-firm mean wage gap of around five percentage points. Women from minority groups are over-represented in lower-paying firms, but again face a within-firm mean wage gap, in this case amounting to around six percentage points.

4. Is within-workplace earnings inequality due to discrimination?

The previous section summarised the results from a number of empirical studies which indicate that wage gaps exist between observationally-equivalent men and women, and between observationally-equivalent whites and ethnic minorities, after accounting for differences in the firms (and sometimes the jobs) to which they are allocated in the labor market. A standard interpretation of these “within-firm” gaps in residual wages is that they indicate differential treatment of workers in wage-setting on the basis of their social group, that is: discrimination. However, in empirical studies, where it is not possible to control fully for all aspects of unobserved heterogeneity between workers, other explanations are also possible.

One feasible explanation is that observed wage differences reflect differences in worker productivity. Most empirical studies do not have direct measures of worker productivity and so rely on proxies – typically measures of educational attainment, labor market experience and job tenure – which may or may not capture the full extent of any productivity differences. Hellerstein et al. (1999) addressed this issue by using linked employer-employee data for the United States to estimate relative marginal products for male and female employees, which could then be compared with relative wages. They identify the productivity and wage differentials associated with the gender of workers by estimating how output and wages covary with the

demographic composition of the workforce in firm-level production functions and wage equations. They estimated that women were, on average, 16% less productive than men, but earning 45% less than men, all other things equal. These estimates are statistically significant from one another, implying that, on average, women were being paid at a rate below their marginal productivity relative to men, consistent with a scenario in which women are discriminated against in wage-setting.

Bartolucci (2013) later extended Hellerstein et al.'s approach, using data for West Germany, to further divide the gender wage gap into differences in productivity, segregation across firms, disparities in worker mobility and differences in "rent-splitting". In this framework, one is able to provide a more complete portrait of the gender wage gap, in which differences in the "rent-splitting" parameter provide a more specific estimate of the extent of wage discrimination than that previously provided by Hellerstein et al. (1999). Bartolucci (2013) estimates that female workers in West Germany are around 18 per cent less productive than male workers. However, female workers receive wages that are nine percent lower than those of male workers with *equivalent* productivity, after accounting for differences in segregation and mobility. Again, the evidence is consistent with discrimination against women in wage-setting.

Attempts to explain why wage outcomes may differ for equally productive male and female workers have then looked to possible gender differences in attributes which may affect the wage bargaining process. These include attitudes to competition which may mean that women are more likely to choose time-based pay rather than performance-related pay (Niederle and Vesterlund, 2007); this may have the implication of reducing the earnings power of high-productivity women relative to high-productivity men. Another possible factor is gender differences in negotiation skills. Psychological studies suggest that women are less successful at negotiating than men (see Stuhlmacher and Walters, 1999); this gender difference is particularly apparent when information is scarce, as may be the case when negotiating a wage on first entry to a firm. A third candidate is differences in preferences for non-wage amenities. Mas and Pallais (2017) show that women (particularly those with young children) place a higher value on working from home and avoiding irregular work schedules than do men, although the differences found in their study do not appear large enough to explain any substantial share of the gender wage gap. Le Barbanchon et al. (2021) go on to show that women will forego a larger share of offered wages than men in order to secure a short commute: a difference which accounts for around one seventh of the conditional wage gap in their study of job seekers in France. Arriving at overall estimates of the combined contribution of these various differences to gender wage gaps among equally productive male and female employees is necessarily difficult, however.

Turning to ethnic wage gaps, Hellerstein et al.'s (1999) study also examined the relative productivities and relative wages of black and nonblack workers, but found no differences. Nevertheless, studies of professional sport, where good measures of player performance are often available, have shown evidence of wage discrimination against ethnic minorities (Holmes, 2011; Johnson and Minuci, 2020). The general absence of mixed-gender sports teams prohibits similar studies of gender wage gaps.

To summarise the preceding discussion, a wide array of empirical studies show that wage differences exist between male and female co-workers, and between white employees and co-workers from ethnic minority groups. These wage differences persist after accounting for the sorting of workers across firms, and jobs (in other words, they can be taken to exist within-firm and within-job), and they persist after accounting for estimated differences in employee productivity. Other, unobserved attributes, such as differences in negotiation skills or preferences, may explain some part of the remaining wage gaps. However, the balance of evidence seems to indicate that wage discrimination also plays a substantive role. To explore the issue further, the paper now turns to examine some of the mechanisms through which such discrimination might manifest itself, surveying evidence of unfair treatment in firms' personnel practices.

5. How might earnings discrimination manifest itself?

If earnings discrimination within firms is a feature of firm behavior, it is important to understand the specific mechanisms through which it manifests itself. This section reviews evidence on wage discrimination arising in three areas of personnel practice: wage offers made at the time of appointment; task allocation; and performance review.

The identification of discriminatory treatment in these settings is a difficult task for observational studies where, as noted above, it is hard to discount a potential role for unobserved differences in ability or preferences. However, there are several strands of experimental evidence in which researchers seek to use simulated circumstances to study the extent to which there may be a gender or ethnicity-related dimension to wage offers and performance evaluation. The experimental studies may have limited external validity, and may themselves exhibit bias in the absence of the cost-related incentives which may dissuade real-world employers from discriminating in a profit-maximising environment. So, the next two sub-sections also report on observational studies which show that personnel decisions or wage outcomes within firms vary according to the gender or ethnicity of the decision-maker; these are useful in providing indirect evidence of whether discrimination is at work.

Discrimination in wage offers on hiring

There is plentiful evidence from correspondence studies (aka. audit studies) that employers are less willing to hire women or ethnic minorities, on average (see Rich, 2014, for one review and Kline et al., 2022, for recent evidence). However, as noted earlier, an unwillingness to hire women or ethnic minorities on the part of some employers does not necessarily translate into differences in the wages that are offered to acceptable candidates; it may simply lead to segregation between firms. A number of studies do provide evidence of differences in wage offers, however.

Moss-Racusin et al. (2012) administered a randomized double-blind study in which around 130 science faculty from research-intensive universities in the United States were asked to rate the application materials of a student applying for the role of laboratory manager. The gender of the applicant was randomly assigned. The study found that participants rated the male applicant as significantly more competent and hireable than the (identical) female applicant. Participants then also proposed a higher starting salary for the male applicant: around \$30,000 versus around \$26,500 for the female applicant.

Some of this gender wage discrimination at hiring may reflect expectations about parenting. Correll et al.'s (2007) paid undergraduate volunteers to rate a pair of equally qualified, same-gender (either male or female) fictitious job applicants, presented as real, who differed on parental status. Applications were for a senior marketing position in a communications company. Regressions were used to compare ratings between men and women with the same parental status. Among non-parents, male and female applicants received wage offers that were not statistically significant from one another (despite female applicants being rated slightly more competent than male applicants). Among parents, however, mothers were rated as less competent and committed than equally qualified fathers, and received substantially lower wage offers. Differences in wage offers by employers may thus be moderated by parental status, with mothers bearing the greatest disadvantage.

A further study by Benard and Correll (2010) suggests that employer discrimination persists even when there is indisputable evidence of competence. In a variation of the earlier study, a “low ambiguity” condition is included, whereby a highly-favourable performance review is included in each candidate’s application package. In this circumstance, employers continued to penalise mothers in terms of starting salary. A key explanatory factor was that evaluators saw successful mothers as less personable than other candidates. The authors explain this pattern as “normative discrimination”, whereby employers may discriminate against

mothers through a belief (perhaps unconscious) that their success signals that they are assertive or dominant, rather than warm and nurturing: in essence penalising them for violating cultural norms.

These studies focus on initial wage offers and so give little insight into how employers might then behave in subsequent bargaining to arrive at the final hired wage. However, Dittrich et al. (2014) use an experimental setting involving university students to study the process of offer and counter-offer, allowing for differences in which party (employer or employee) initiates the negotiation. They find that male employers make higher initial wage offers to male job candidates than to female job candidates (there is no gender difference in the initial offers of female employers). When the employee makes the initial offer, on the other hand, both male and female employers respond by issuing lower counter-offers to female employees than to male employees. Consequently, female job applicants are disadvantaged by the behavior of the employer at each stage of the negotiation process.

Similar evidence that employers vary their wage offers by ethnicity comes from an observational study by Fryer et al. (2013). They studied longitudinal data on a sample of around 5,000 unemployed workers in New Jersey who completed weekly interviews for up to 12 weeks, examining the job search that each worker undertook, and the job offers that they received. After controlling for differences in age, qualifications, and wages in the previous job, they found that black applicants received hourly wage offers that were around 16% lower than those of white applicants. This study is unusual among observational studies in being able also to account for a wide variety of factors on the applicant side which may affect the size of the wage offer made by the employer. Specifically, the authors are able to control for differences in applicants' search intensity, search strategies, bargaining behavior, discount rates and geographic location. Controlling for these factors brought about almost no change in the estimated racial difference in wage offers. The authors estimated that this substantial difference in wage offers could account for at least one third of the black-white wage gap in the United States.

It was noted earlier that one caveat to the literature on ethnicity is that features may differ considerably by ethnic group. In this vein, Bohren et al.'s (2019a) study is notable. Their experimental design involved recruiting two sets of participants from the United States and India via an online job-task platform. The first set of participants (classified as the "employee") completed a work task to provide realised data on productivity; the second set (classified as "employers") were asked to select among these potential employees and make a wage offer. In this study, the Indian employees were offered higher wages than the American employees, even though their observed performance was worse. The study indicates that statistical discrimination can work in favour of some minority groups with high average levels of attainment.

Discrimination in performance evaluation

Experimental evidence suggests that the process of ongoing performance evaluation post-hiring may – in some cases – reduce any wage gap that may have existed on entry to the firm. Theoretically, if wage discrimination on hiring is based on incorrect beliefs about worker productivity (i.e. statistical discrimination), then the resulting wage gap should decline as the employer observes the employee's performance on the job and so acquires better information about their productivity. If the wage gap persists, however, this suggests that the discrimination is taste-based. Bohren et al. (2019b) provide experimental evidence in support of a temporal reduction in gender discrimination arising from an updating of beliefs. Posting questions and answers on an online Q&A forum used by students and researchers in scientific disciplines, they find that significant gender discrimination exists initially in user evaluations of answers given by men and women with little acquired reputation on the platform; however, the discrimination against women reduces and even reverses as they build experience so that high-reputation women eventually receive better evaluations than high-reputation men.

Despite this evidence, there are plentiful indications that women's careers do not always proceed as well as men's. For instance, many studies have provided evidence of gender differences in promotion to higher paid ranks within an occupation, even after controlling for good measures of workers ability and performance (e.g. Pekkarinen and Vartiainen, 2006; Blau and De Varo, 2007; De Paola et al., 2018). Ransom and Oaxaca (2005) use eleven years of administrative employment records to examine job mobility within a regional chain of grocery stores in the United States. They analyse the probability that a food clerk will be promoted to a managerial position and find that, after controlling for age and seniority, a male food clerk is over six times as likely to be promoted as is a female clerk.

Experimental evidence that managers may apply different promotion standards comes from an early study by Rosen and Jerdee (1974). They asked 95 male bank supervisors to assume the role of the personnel director of a multibranch bank and to review documentary evidence on whether an employee should be promoted. In this study, male candidates were around 1.5 times more likely to be given the promotion than observationally-equivalent female candidates.

This evidence accords with the argument that there exist 'glass ceilings' in place within organizations, whereby women are held to different promotion standards than men, and as a consequence experience slower wage growth within the firm. Lazear and Rosen's (1990) model of job ladders focuses on women's lower probability of promotion as a primary cause of wage discrimination within firms. They argue that the differential in promotion opportunities arises because women have a wider set of outside opportunities than men and so employers prefer to promote men due to their lower likelihood of quitting. Empirical evidence on gender differences in quit rates is mixed (e.g. Ransom and Oaxaca (2005) find no differences in their study) although, of course, women may be more likely to quit if they are presented with fewer opportunities to remain in post.

Differential investments in firm-specific human capital may feed into differential outcomes in performance evaluation. In particular, it may be that women find it more difficult than male co-workers to access the employer-provided investments that support positive evaluations. Rosen and Jerdee (1974) found that the bank supervisors in their study were more willing to give a development opportunity (attending a conference) to a male candidate than a female candidate. Similarly, in Moss-Racusin et al. (2012)'s experimental study, faculty members were more likely to consider offering mentoring opportunities to male applicants than to female applicants. Women may also be disadvantaged through allocation to less rewarding assignments, as evidenced in Madden's (2012) study of the financial industry, where female stockbrokers were allocated to lower-quality accounts than male co-workers.

A further aspect of the evaluation process is the treatment of mistakes and here, also, there is some evidence that employers' decisions may have a gendered dimension. Egan et al. (2022) utilise panel data on all financial advisers registered in the United States from 2005 to 2015, taking advantage of the regulatory requirement for financial advisers to disclose misconduct-related career events. Their data contain detailed information on the nature and cost of misconduct, and allow them to compare male and female advisers working at the same firm, in the same location, at same point in time, and in the same job role, and to control for differences in advisor productivity. They find that, after an incident of misconduct, female advisers are 20% more likely to lose their jobs, despite engaging in misconduct that is 20% less costly and having a substantially lower propensity towards repeat offenses. Such differences in the punishment of errors may not only slow down women's career progression within the firm through discriminatory dismissal, but may also affect future opportunities in the labor market. In Egan et al.'s study, female advisers were 30% less likely to find new jobs relative to male advisers after a misconduct-related dismissal.

Studies of differences in performance evaluation across ethnic groups are less plentiful than gender-focused studies. However, Egan et al.'s (2022) study of financial advisers also examines patterns of punishment

according to the ethnicity of the advisor, finding that male advisors with names that indicate either African or Hispanic origin face a punishment gap that is similar to female advisors. Similarly, Madden's research into stockbroking has shown that biases exist against African American brokers (Madden and Vekker, 2017). Other studies in this vein include those by Elvira and Town (2001) and Castilla (2008).

Elvira and Town (2001) study the annual performance evaluations of around 300 salespersons in a major US corporation. They compare the subjective evaluation given by managers with objective data on the employee's sales performance, finding that black employees receive lower subjective ratings than white co-workers after controlling for actual sales performance. Asian and Hispanic employees do not experience bias, however. A black employee has a 50% lower probability of receiving the highest subjective rating (outstanding) than a white employee with the same attributes. Elvira and Town also find that salary differences between minority and white employees within the firm disappear after controlling for job title and subjective performance ratings, indicating that the bias in performance evaluations is an important contributor to wage differentials.

Wage differences may not only arise through differences in the performance evaluation itself, however; they may also arise if there are differences in the rewards given for good performance. Castilla (2008) examines personnel data from a large service sector organization in the US, finding no evidence of bias in the rating of performance. However, there are gender and racial differences in the salary increases that are granted for observationally equivalent employees who receive the same performance evaluation scores. Castilla attributes this discrepancy to differences in levels of accountability and transparency: the process of performance evaluation in the organization is heavily formalised with high degrees of managerial accountability, whereas the process of awarding salary increases is not. These findings thus provide important context for observational studies on performance-related pay which have shown that white employees receive a wage premium from performance-pay jobs that is greater than that received by black workers (Heywood and Parent, 2012).

Discrimination in task allocation

The discussion in the preceding sub-sections typically assumes that wage discrimination is taking place within-job, that is between equally able and productive workers hired to the same role within the firm (so-called "within-assignment discrimination"). However, employers may also discriminate by assigning equally able and productive workers to more or less lucrative positions within the firm ("cross-assignment discrimination"). In doing so, they may be motivated by their own beliefs about the suitability of different types of workers for certain tasks, by incumbent employees who prefer (or are assumed to prefer) to work alongside a particular type of co-worker, or by similar preferences from customers. Addressing such cross-assignment discrimination in wages is the motivation for legislation which mandates that employers should provide equal pay for work of equal value (aka. comparable worth).

Ransom and Oaxaca's (2005) study of gender wage differentials in a grocery store chain, cited above, not only found evidence of gender differences in promotion, but also differences in initial task allocation on hiring. Women were more likely to be employed as food clerks (stacking shelves and operating cash registers) whereas men were more likely to be employed as meat cutters or wrappers. Variation in job titles accounted for around 95% of variation in wages within the store. One challenge to interpretation here is that Ransom and Oaxaca had little information that would enable them to control for differences in productivity between these different job assignments. However, Bodvarsson et al. (2014) are able to take advantage of the public nature of professional sport to control for various measures of player productivity in a study of job assignment in Major League baseball, using a two-stage estimation process to control for assignment-specific productivity. In this context, white employees are more likely to be employed as pitchers and black or Hispanic employees as hitters. Bodvarsson et al. (2014) find evidence of a white wage premium across these

job assignments, even after controlling for a wide array of demographic variables and position-specific productivity.

Subtle barriers

Taken together, the various studies discussed above provide support to the hypothesis that women and ethnic minorities are victims of direct discrimination in important personnel decisions that contribute to wage outcomes. However, case law also provides for indirect discrimination, whereby personnel practices apply neutrally across social groups, but subtle barriers may affect the ability of the members of a particular social group from meeting the required standards.

The argument is made most forcefully, perhaps, in respect of gender, where it is argued that barriers arise because of how firms reward individuals who differ in their desire for flexibility. Goldin (2014) argues that, in many higher-paying jobs, wage penalties arise when employees take advantage of family-related amenities in the workplace (e.g. parental leave, part-time work and flexibility during the workday). Continuity of service and continuous time on the job (particularly, the working of long hours) are argued to be critical to employee's wage progression in such occupations. Hence, workers in these occupations who desire more flexibility (for example at the onset of motherhood) have to take a compensating wage differential in order to obtain this amenity, leading to the emergence of wage differentials.

Goldin (2014) provides extensive evidence that jobs in which career interruptions are particularly penalised, and work environments that require more interactions or have more time pressure, are those with larger gender earning gaps (see also Goldin and Katz, 2016). Changes in attitudes which break down the norm of "long hours working", and changes in information technology which increase the substitutability of employees during the workday (e.g. through the systematic recording of customer information), then offer the prospect of reducing these wage penalties.

Another growing strand of literature has focused on the role of social capital within the workplace. The typical focus is on the role of referral networks in supporting job search for displaced workers (e.g. Saygin et al., 2021). However, informal networks within the firm can also benefit in-group members in terms of career progression and wage growth. A common focus of debate is the role of the "old boys network" in supporting wage progression for men relative to women. Robust evidence is surprisingly limited. However, in a recent study, Cullen and Perez-Truglia (2019) use administrative and survey data from a large commercial bank in Asia. Using quasi-random variation induced by the rotation of managers, along with the smoking status of managers and employees, they show that employees' social interactions with their managers are advantageous for their careers. These social interactions are found to be particularly advantageous for male employees, and are found to explain one third of the gender gap in promotions at this firm.

The identity of the managerial decision-maker

Many of the experimental studies cited above naturally invite a discriminatory interpretation, by virtue of having controlled away other possible sources of variation in outcomes beyond the social group identity of the employee receiving the wage offer or performance evaluation. However, such interpretations become even more plausible when studies show that outcomes vary also with the social group identity of managerial decision makers, such that any bias against female employees (say) is greater when a higher share of the decision makers are male. A number of the aforementioned studies show such patterns, including Dittrich et al. (2014), Elvira and Town (2001) and Egan et al. (2022). Moss-Racusin et al. (2012) is one exception, where male and female faculty members were equally likely to exhibit bias against female students.

A number of observational studies add to this body of evidence. For example, Tate and Yang (2015) find that US firms with more women in leadership roles have a smaller gender wage gap, and that women in these roles offer equal pay to newly-hired employees at junior levels. Similarly, for Portugal, Cardoso and Winter-Ebmer (2010) show that women's wages rise relative to men's when a workplace switches to being female-

led. Theodoropoulos et al. (2022) provide equivalent evidence for the UK. In respect of promotion prospects, De Paola and Scoppa (2015) find that female candidates are significantly less likely to get promoted when the committee is entirely composed of males, whilst Kunze and Miller's (2017) study of corporate leadership in Norway finds that greater female representation among the higher occupational ranks in the workplace narrows the gender gap in promotion rates at lower ranks. Lucifora and Vigani (2022) directly examine workers' perceptions of whether they have experienced gender discrimination within their firm, and find that having a female manager is associated with lower levels of perceived discrimination. Similar evidence by ethnicity is provided by Giuliano et al. (2011). They analyse quits, dismissals and promotions at a large US retail firm, exploiting changes in management at store-level to identify levels of own-race bias. They find that black, Hispanic and Asian employees have higher relative rates of promotion when their manager is the same race. However, they also find that white employees are more likely to be promoted by non-white managers. They argue that the latter result, which runs counter to expectations about own-race bias, reflect non-white managers' implicit desire to maintain norms that associate white employees with high status.

6. Policy responses

Some of the evidence on sources of wage inequality reviewed in Section 3 indicates that the sorting of workers between firms contributes in some way to gender and ethnic wage gaps. Accordingly, policy makers and employers should continue to focus on reducing hiring discrimination and helping job seekers from disadvantaged groups to identify and access firms with more generous wage policies. However, addressing discrimination in wage-setting within the workplace requires other solutions.

One response, utilised by governments in most developed countries, is to impose financial penalties on firms that are found to have used social group unfairly as criteria in wage determination. As noted earlier, this is one important strand of discrimination law, whereby employers who do not adhere to the principle of equal pay for work of equal value may be subject to fines – often unlimited in value.

Makepeace et al. (1999) show that there was a substantial reduction in the gender wage gap in the UK after the introduction of Equal Pay legislation in the 1970s. Similarly, Beller (1982) presents evidence of improvements in women's relative earnings in the US in the late 1960s and early 1970s after the introduction of Title VII of the Civil Rights Act of 1964, which outlaws discrimination in wage setting and other aspects of employment. Such laws raise the cost of discrimination to employers, although necessarily their efficacy relies on the level of enforcement. In particular, Blau and Kahn (2017: 848) have argued that further progress in the US in the 1980s was limited by the winding down of official efforts to enforce the legislation, whilst Makepeace et al. (1999) suggest that the efficacy of equal pay legislation in the UK may have declined in the 1980s and 1990s as a result of the weakening of trade unions – the support of whom is known to be important for employees wishing to pursue claims.

The limits of regulation imply that there is an important role for voluntary action by firms to limit the space for decision-makers' tastes or stereotypes to affect wage-setting decisions. Diversity training, which seeks to make decision-makers aware of their potential biases, so that they may be suppressed, is a common feature of many organizations' diversity programs. However, a meta-analysis of the effects of such training by Bezrukova et al. (2016) indicates that the effects on decision-makers' behavior is typically small. Such training is more likely to be effective when complemented by other diversity initiatives, targeted to both awareness and skills development, and conducted over a significant period of time.

In addition to the provision of training, firms can help to ensure that decisions affecting wage outcomes are based on detailed performance evaluations. The greater provision of productivity-related information helps to

ensure that the decisions of well-meaning managers are less susceptible to stereotyping, although as Castilla (2008) shows, norms may still play a role.

Managers can be further incentivised to arrive at fair outcomes through an expectation that they may be called upon to justify these decisions to their peers. Experimental evidence shows that bias is less likely to occur when decision makers know they will be held accountable for their decisions (e.g. Tetlock 1983). Accountability incentivises managers to acquire relevant information on which to base a decision, and also prompts them to analyse this information more carefully. Indeed, Kalev et al. (2006) have found accountability to be one of the most effective initiatives underpinning advances in diversity in corporate settings. Accountability is not a universal panacea, however, as there are circumstances in which it can also promote bias towards a socially-desirable outcome (Lerner and Tetlock, 1999).

Firms can also take actions to reduce the subtle barriers which prevent disadvantaged groups from competing on a level playing field. The presence of wage penalties associated with flexible working were cited earlier in the paper. The greater expansion of parental leave and other “family-friendly” policies has increased female labor-force participation, but has also been associated with reductions in women’s earning power in work in cases where the take up of leave and flexible working is primarily among women (Blau and Kahn, 2017). This suggests that policies to equalize working patterns and, perhaps more importantly, the division of labor within the family are needed to drive reductions in gender wage gaps within the workplace. Firms can also address the role of social connections in perpetuating disadvantages within the workplace by making efforts to limit the role of informal influence in wage-setting and/or by promoting co-integration in co-worker networks.

Further, policies have tried to incentivise firms to look for biases via increased public transparency. A requirement for all firms with 250 or more employees in the UK to report publicly on their gender wage gap has been evaluated as a success, reducing the gender wage penalty by 15-20 per cent (Blundell, 2021). Such policies may have a variety of possible effects. Making the situation within the firm apparent to outsiders may affect the firm’s ability to attract talent; it may also affect demand for their goods and services by consumers. Furthermore, it may also reveal previously-hidden problems to decision-makers within the firm. Each of these may act as incentives for the firm to seek out and address the origins of wage inequality among its employees. Since much of the ethnic wage gap exists within workplaces rather than between workplaces (Forth et al., 2021), extending the policy to ethnicity may bring rewards. Although ethnic pay gap reporting has been discussed in UK, it faces challenges because anonymity is more difficult.

Summary

Antidiscrimination legislation is an important public policy and many countries around the world have laws to combat discrimination. However, devising antidiscrimination legislation is extremely difficult, as such legislation might damage labor market efficiency and has limited scope to affect the full range of decisions and interpersonal dynamics that take place within a firm. The relevant literature suggests that the introduction of antidiscrimination legislation in the 1970s significantly reduced pay differentials between genders and ethnicities. Despite this, gender and ethnicity wage gaps still exist.

The early empirical literature on wage differentials used regression analysis from surveys of individuals to estimate differences in earnings that appear attributable to the gender or race characteristics of the individual. However, more recently the availability of data sources that generate information simultaneously on the employee and the employer has been utilised. Thus, the major development in recent decades has been to

show that pay differentials do not only exist between observationally equivalent workers, but also between observationally equivalent co-workers in the same firm and even in the same job.

Economists have tried to identify the mechanisms through which pay discrimination persists by focussing on employers' wage setting behavior. Within organizations much of the focus has been providing a level playing field for women and ethnic minorities. However, economists and organizations have become increasingly aware of the subtle factors that limit the ability of vulnerable groups to enter the playing field in equal terms.

References

- Abowd, J.M., Kramarz, F. and Margolis, D.N. 1999. High wage workers and high wage firms. *Econometrica* 67 (2), pp. 251-333.
- Arrow, K. 1973. The Theory of Discrimination. In *Discrimination in Labor Markets*, edited by Orley Ashenfelter and Albert Rees, 3–33. Princeton and Oxford: Princeton University Press.
- Bachmann E. 2022. \$70 million verdict against Texas company in employment discrimination case. *Forbes*, 21st March 2022. Available online at: <https://www.forbes.com/sites/ericbachman/2022/03/21/70-million-verdict-against-texas-company-in-employment-discrimination-case/?sh=62775ca32ea8> [Retrieved 4th May 2022]
- Barth, E., Bryson, A., Davis, J.C. and Freeman, R. 2016. It's where you work: Increases in the dispersion of earnings across establishments and individuals in the United States. *Journal of Labor Economics* 34 (S2), pp. S67-S97.
- Bartolucci, C. 2013. Gender wage gaps reconsidered: A structural approach using matched employer-employee data. *Journal of Human Resources* 48 (4), pp. 998-1034.
- Bayard, K., Hellerstein, J., Neumark, D. and Troske, K. 2003. New evidence on sex segregation and sex differences in wages from matched employer-employee data. *Journal of Labor Economics* 21 (4), pp. 887-922.
- Becker, G. 1971. *The Economics of Discrimination*, Second edition. Chicago and London: University of Chicago Press.
- Beller, A.H. 1982. The impact of equal opportunity policy on sex differentials in earnings and occupations. *The American Economic Review* 72 (2), pp.171-175.
- Benard, S. and Correll, S.J. 2010. Normative discrimination and the motherhood penalty. *Gender & Society* 24(5), pp. 616-646.
- Bezrukova, K., Spell, C.S., Perry, J.L. and Jehn, K.A. 2016. A meta-analytical integration of over 40 years of research on diversity training evaluation. *Psychological Bulletin* 142 (11), pp. 1227-1274.
- Blau, F.D. and Kahn, L. 2017. The gender wage gap: Extent, trends, and explanations. *Journal of Economic Literature* 55 (3), pp. 789-865.
- Blau, F.D. and DeVaro, J. 2007. New evidence on gender differences in promotion rates: An empirical analysis of a sample of new hires. *Industrial Relations: A Journal of Economy and Society* 46 (3), pp. 511-550.
- Blundell, J. 2021. Wage responses to gender pay gap reporting requirements. CEP Discussion Paper No. 1750, London: Centre for Economic Performance.
- Bodvarsson, Ö.B., Papps, K.L. and Sessions, J.G. 2014. Cross-assignment discrimination in pay: A test case of major league baseball. *Labour Economics* 28, pp. 84-95.
- Bohren, J.A., Haggag, K., Imas, A. and Pope, D.G. 2019a. Inaccurate statistical discrimination: An identification problem. Working Paper No. 25935. National Bureau of Economic Research.

- Bohren, J.A., Imas, A. and Rosenberg, M. 2019b. The dynamics of discrimination: Theory and evidence. *American Economic Review* 109 (10), pp. 3395-3436.
- Bruns, B. 2019. Changes in workplace heterogeneity and how they widen the gender wage gap. *American Economic Journal: Applied Economics* 11 (2), pp. 74 – 113.
- Card, D., Cardoso, A.R. and Kline, P. 2016. Bargaining, sorting, and the gender wage gap: Quantifying the impact of firms on the relative pay of women. *Quarterly Journal of Economics* 131 (2), pp. 633-686.
- Cardoso, A.R., Guimaraes, P. and Portugal, P. 2016. What drives the gender wage gap? A Look at the role of firm and job title heterogeneity. *Oxford Economic Papers* 68 (2), pp. 506-524.
- Cardoso, A.R. and Winter-Ebmer, R. 2010. Female-led firms and gender wage policies. *ILL Review* 64 (1), pp. 143-163.
- Carrington, W. and Troske, K.R. 1998. Interfirm segregation and the black/white wage gap. *Journal of Labor Economics* 16 (2), pp. 231–260.
- Casarico, A. and Lattanzio, S. 2019. What firms do: Gender inequality in linked employer-employee data. Cambridge-INET Working Paper Series No: 2019/15, WorkINPS Papers n. 24.
- Castilla, E.J. 2008. Gender, race and meritocracy in organizational careers. *American Journal of Sociology* 113 (6), pp. 1479-1526.
- Correll, S.J., Benard, S., & Paik, I. 2007. Getting a job: Is there a motherhood penalty? *American Journal of Sociology* 112 (5), pp. 1297-1338.
- Coudin, E., Maillard, S. and To, M. 2018. Family, firms and the gender wage gap in France. IFS Working Paper W18/01.
- Criscuolo, C., Hijzen, A., Schwellnus, C., Barth, E., Chen, W.H., Fabling, R., Fialho, P., Grabska, K., Kambayashi, R., Leidecker, T., Skans, O. N., Riom, C., Roth, D., Stadler, B., Upward, R. and Zwysen, W. 2020. Workforce Composition, Productivity and Pay: The Role of Firms in Wage Inequality. IZA Working Paper No. 13212.
- Cullen, Z.B. and Perez-Truglia, R., 2019. The old boys' club: Schmoozing and the gender gap. Working Paper No. 26530. National Bureau of Economic Research.
- De Paola, M., Ponzio, M. and Scoppa, V. 2018. Are men given priority for top jobs? Investigating the glass ceiling in Italian academia. *Journal of Human Capital* 12 (3), pp. 475-503.
- De Paola, M. and Scoppa, V. 2015. Gender discrimination and evaluators' gender: Evidence from Italian academia. *Economica* 82, pp. 162-188.
- Di Addario, S.L., Kline, P.M., Saggio, R. and Sølvsten, M. (2022). It ain't where you're from, it's where you're at: Hiring origins, firm heterogeneity, and wages. *Journal of Econometrics*
<https://doi.org/10.1016/j.jeconom.2021.12.017>
- Dittrich, M., Knabe, A. and Leipold, K. 2014. Gender differences in experimental wage negotiations. *Economic Inquiry* 52 (2), pp. 862-873.
- Drolet, M. and Mumford, K. 2012. The gender pay gap for private-sector employees in Canada and Britain. *British Journal of Industrial Relations* 50 (3), pp. 529-553.
- Egan, M., Matvos, G. and Seru, A. 2022. When Harry fired Sally: The double standard in punishing misconduct. *Journal of Political Economy* 130 (5), forthcoming. <https://doi.org/10.1086/718964>
- Elvira, M. and Town, R., 2001. The effects of race and worker productivity on performance evaluations. *Industrial Relations: A Journal of Economy and Society* 40 (4), pp. 571-590.

- Forth, J., Theodoropoulos, N. and Bryson, A. 2021. The role of the workplace in ethnic wage differentials. IZA Discussion Paper 14697.
- Fryer, R.G., Pager, D. and Spenkuch, J.L. 2013. Racial disparities in job finding and offered wages. *The Journal of Law and Economics* 56 (3), pp.633-689.
- Gallen, Y., Lesner, R.V. and Vejlin, R. 2019. The labor market gender gap in Denmark: Sorting out the past 30 years. *Labour Economics* 56, pp. 58-67.
- Giuliano, L., Leonard, J. and Levine, D.I. 2011. Racial bias in the manager-employee relationship: An analysis of quits, dismissals and promotions at a large retail firm. *Journal of Human Resources* 46 (1), pp. 26-52.
- Goldin, C. 2014. A grand gender convergence: Its last chapter. *American Economic Review* 104 (4), pp. 1091-1119.
- Goldin, C. and Katz, L. 2016. The most egalitarian of all professions: Pharmacy and the evolution of family-friendly occupations. *Journal of Labor Economics* 34 (3), pp. 705-716.
- Hara, H. 2018. The gender wage gap across the wage distribution in Japan: Within and between-establishment effects. *Labour Economics* 53, pp. 213-229.
- Hellerstein, J. and Neumark, D. 2008. Workplace segregation in the United States: Race, ethnicity, and skill. *Review of Economics and Statistics* 90 (3), pp. 459-477.
- Hellerstein, J.K., Neumark, D. and Troske, K.R. 1999. Wages, productivity, and worker characteristics: Evidence from plant-level production functions and wage equations. *Journal of Labor Economics* 17 (3), pp. 409-446.
- Heywood, J. and Parent, D. 2012. Performance pay and the black-white wage gap. *Journal of Labor Economics* 30 (2), pp. 249-290.
- Heywood, J. and Theodoropoulos, N. 2020. The declining influence of workplace differences on the gender wage gap. *Economics Bulletin* 42 (2), pp. 1194-1200.
- Holmes, P. 2011. New evidence of salary discrimination in major league baseball. *Labour Economics* 18(3), pp. 320-331.
- Javdani, M. 2015. Glass ceiling or glass doors? The role of firms in male-female wage disparities. *Canadian Journal of Economics* 48 (2), pp. 529-560.
- Jewell, S.L., Razzu, G. and Singleton, C. 2020. Who works for whom and the UK gender pay gap. *British Journal of Industrial Relations* 58 (1), pp. 50-81.
- Johnson, C. and Minuci, E. 2020. Wage discrimination in the NBA: Evidence using free agent signings. *Southern Economic Journal* 87 (2), pp. 517-539.
- Kahnemann, D. 2011. *Thinking Fast and Slow*, Macmillan.
- Kalev, A., Dobbin, F. and Kelly, E. 2006. Best practices or best guesses? Assessing the efficacy of corporate affirmative action and diversity policies. *American Sociological Review* 71 (4), pp. 589-617.
- Kline, P.M., Rose, E.K. and Walters, C.R. (2022). Systemic discrimination among large U.S. employers. *Quarterly Journal of Economics* forthcoming.
- Kunze, A. and Miller, A.R. 2017 Women helping women? Evidence from private sector data on workplace hierarchies. *Review of Economics and Statistics* 99 (5), pp. 769-775.
- Lazear, E.P. and Rosen, S. 1990. Male-female wage differentials in job ladders. *Journal of Labor Economics* 8 (1, Part 2), pp. S106-S123.
- Le Barbanchon, T., Rachelot, R. and Roulet, A. 2021. Gender differences in job search: Trading off commute against wage. *Quarterly Journal of Economics* 136 (1), pp. 381-426.

- Lerner, J.S. and Tetlock, P.E. 1999. Accounting for the effects of accountability. *Psychological Bulletin* 125 (2), pp. 255-275.
- Longhi, S. 2020. Racial wage differentials in developed countries. *IZA World of Labor*.
- Lucifora, C. and Vigani, D. 2022. What if your boss is a woman? Evidence on gender discrimination at the workplace. *Review of the Economics of the Household* 20, pp. 389-417.
- Madden, J.F. 2012. Performance-support bias and the gender pay gap among stockbrokers. *Gender & Society* 26 (3), pp. 488-518.
- Madden, J.F. and Vekker, A. 2017. Output-based performance pay, performance-support bias, and the racial pay gap within a large retail stock brokerage. *Industrial Relations: A Journal of Economy and Society* 56 (4), pp. 662-687.
- Makepeace, G., Paci, P., Joshi, H. and Dolton, P. 1999. How unequally has equal pay progressed since the 1970s? A study of two British cohorts. *Journal of Human Resources* 34 (3), pp. 534-556.
- Mas, A. and Pallais, A. 2017. Valuing alternative work arrangements. *American Economic Review* 107 (12), pp. 3722-3759.
- Meng, X. 2004. Gender earnings gap: The role of firm specific effects. *Labour Economics* 11, pp. 555-573.
- Moss, R. 2019. City Banker Wins Sexism Case for Unequal Pay and Victimisation. *Personnel Today*, 12th September. Available online at: <https://www.personneltoday.com/hr/city-banker-wins-sexism-case-for-unequal-pay-and-victimisation/> [Retrieved 4th May 2022]
- Moss-Racusin, C. A., Dovidio, J. F., Brescoll, V. L., Graham, M. J., & Handelsman, J. 2012. Science faculty's subtle gender biases favor male students. *Proceedings of the National Academy of Sciences* 109 (41), pp. 16474-16479.
- Niederle, M. and Vesterlund, L. 2007. Do women shy away from competition? Do men compete too much? *Quarterly Journal of Economics* 122 (3), pp. 1067-1101.
- Pekkarinen, T. and Vartiainen, J. 2006. Gender differences in promotion on a job ladder: Evidence from Finnish metalworkers. *ILR Review* 59 (2), pp. 285-301.
- Pendakur, K. and Woodcock, S.D. 2010. Glass ceilings or glass doors? Wage disparity within and between firms. *Journal of Business and Economics Statistics* 28 (1), pp. 181-189.
- Phelps, E. 1972. The statistical theory of racism and sexism. *American Economic Review* 62 (4), pp. 659-661.
- Ransom, M. and Oaxaca, R.L. 2005. Intrafirm mobility and sex differences in pay. *ILR Review* 58 (2), pp. 219-237.
- Rich, J. 2014. What do field experiments of discrimination in markets tell us? A meta-analysis of studies conducted since 2000. IZA Discussion Paper No. 8584.
- Rosen, B. and Jerdee, T.H. 1974. Influence of sex role stereotypes on personnel decisions. *Journal of Applied Psychology* 59 (1), pp. 9-14.
- Saygin, P.O., Weber, A. and Weynandt, M.A. 2021. Coworkers, networks, and job-search outcomes among displaced workers. *ILR Review* 74 (1), pp. 95-130.
- Sorkin, I. 2017. The role of firms in gender earnings inequality: Evidence from the United States. *American Economic Review (Papers and Proceedings)* 107 (5), pp. 384-387.
- Stuhlmacher, A.F. and Walters, A.E. 1999. Gender differences in negotiation outcome: A meta-analysis. *Personnel Psychology* 52 (3), pp. 653-677.
- Tetlock, P.E. 1983. Accountability and the perseverance of first impressions. *Social Psychology Quarterly* 46 (4), pp. 285-292.

Theodoropoulos, N., Forth, J. and Bryson, A. 2022. Are women doing it for themselves? Gender segregation and the gender wage gap. *Oxford Bulletin of Economics and Statistics* forthcoming.

Woodcock, S.D. 2008. Wage differentials in the presence of unobserved worker, firm, and match heterogeneity. *Labour Economics* 15, pp. 772-794.