

DISCUSSION PAPER SERIES

IZA DP No. 13877

Gender Economics: An Assessment

Almudena Sevilla

NOVEMBER 2020

DISCUSSION PAPER SERIES

IZA DP No. 13877

Gender Economics: An Assessment

Almudena Sevilla

University College London and IZA

NOVEMBER 2020

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

Gender Economics: An Assessment*

Concerns about gender equality have jumped to the forefront of public debate in recent years, and Gender Economics is slowly affirming its place as a major field of study. This assessment examines where we are in terms of gender equality. It reviews the theoretical foundations that can explain existing inequalities, and documents the empirical findings supported by the theories, identifying avenues for future research and providing a fruitful framework to think about the effectiveness of policies and interventions targeting gender inequality. In doing so, I provide the foundations against which the contributions in this issue can be placed.

JEL Classification: D31, D63, F60

Keywords: gender economics, gender inequality, technology, norms, policy

Corresponding author:

Almudena Sevilla
University College London
55-59 Gordon Square
London WC1H 0NT
United Kingdom
E-mail: A.Sevilla@ucl.ac.uk

* I am grateful to Abi Adams, Chris Adam, Cameron Hepburn, Alex Teytelboym, David Vines, and Ken Mayhew for comments on an earlier draft. I thank the European Research Council Consolidator Grant (CoG), SH3, ERC-2017-COG for funding my time through the PARENTime project. This paper is forthcoming in *Oxford Review of Economic Policy*, vol. 36 no. 4

I. Introduction

This issue on gender economics comes 15 years after the first issue on a related topic, ‘Gender and the Life Cycle’, was published by the *Oxford Review of Economic Policy*.¹ It joins a new wave of academic activity around gender in economics, as evidenced by other special issues in leading economic journals, a surge in gender economic courses in many universities around the world, and an awakening by economic societies after revelations about sexism, racism, and harassment in the discipline.² This movement in the economics profession is set against a background in which gender has jumped to the forefront of policy and the public debate, as evidenced by the relatively recent #MeToo movement. In this assessment I review the theoretical and empirical progress made in relation to our understanding of gender inequalities. In doing so, I provide a fruitful framework in which to think about the effectiveness of policies and interventions targeting gender inequality.

Gender economics builds on economic theories and empirical methods to understand the mechanisms behind gender inequalities. The crucial question is: Why should we care? Here I argue that efficiency considerations can take us a long way to understanding why, as Christine Lagarde put it while Managing Director of the IMF, ‘excluding women simply makes no economic sense—and including women can be a tremendous boon to the 21st century global economy’.³

At the macro level, a growing literature directly associates women’s participation in the labour force with economic growth and development, as well as with other outcomes that indirectly increase economic productivity, such as an increase in competitiveness (see the review by Lisa Kolovich *et al.* (2020), the opening paper in this issue).⁴ An immediate channel through which levelling the playing field to include

¹ Here I focus on gender as it encompasses male and female; other gendered identities fall beyond the scope of this issue, and they remain a pending issue for future generations of researchers.

² *The Economic Journal* recently published an issue on gender economics. Coverage of the claims against Roland G. Fryer Jr about the harassment and bullying of women can be found at <https://www.nytimes.com/2020/01/07/business/economy/economics-race-gender.html>

³ ‘The Economic Power of Women’s Empowerment’, keynote speech by Christine Lagarde, Managing Director, International Monetary Fund, available at <https://www.imf.org/external/np/speeches/2014/091214.htm>.

⁴ This issue focuses mainly on developed countries. Developing countries differ in the extent of disadvantage and discrimination, their courses over time, and policy responses.

more women in economic decision-making can improve economic outcomes is through a more efficient allocation of scarce resources (namely, human talent). In the United States, for example, the improved allocation of talent has been associated with increases in economic growth of between 20 and 40 per cent (Hsieh *et al.*, 2019). In the developed world, women are now more likely than men to be enrolled in tertiary education (Goldin, 2015), and low fertility and tighter immigration controls make the economic costs to excluding women in the economy simply not affordable. At the micro level, the ‘business case’ for increasing the talent pool relies on women being more productive (for example, women invest better; see Kolovich *et al.*(2020)), and diversity may lead to better decision-making and increased productivity (see the paper by Ghazala Azmat and Anne Boring (2020, in this issue)). Gender inequality may not only result in an inefficient allocation of resources, but this inefficient allocation may be perpetuated if lower returns for women mean fewer girls choose less than the efficient amount of human capital investment (see the paper by Chiara Cavaglia *et al.* (2020, this issue)).

Beyond the obvious economic loss from not allocating scarce resources efficiently, an important efficiency aspect of gender inequality that is commonly overlooked is that of externalities, which sometimes can be a matter of life and death. Economic models of the household predict that, *ceteris paribus*, the partner with greater opportunity cost (as measured by potential wages) has a higher bargaining power within the household, and thus the ability to tilt household decisions in his or her favour (Becker, 1965; Chiappori, 1997). Lack of opportunities relative to their male counterparts in the labour market puts some women at a bargaining disadvantage *vis-à-vis* their male partners. In separate spheres models, sub-optimal equilibrium happens when divorce is too costly and is not a credible threat point (Lundberg and Pollak, 1993). Here the prospect of domestic violence serving as a threat can act as an alternative non-cooperative fallback that shifts cooperative allocations.

Dan Anderberg and Gloria Moroni (2020, this issue) report that, according to figures from the Office for National Statistics, about 1.3m women in the UK (or 8 per cent of women) experienced domestic abuse in 2018. In line with a growing literature that investigates the effects of domestic violence on children’s later development outcomes (Aizer, 2011; Almond *et al.*, 2018), the authors find that exposure to domestic abuse

brings negative consequences to child development. This negative effect is particularly crucial in key developmental periods such as early childhood, given the fact that at least 50 per cent of the variability of lifetime earnings results from attributes determined early on in a child's life, with irreversible devastating consequences for later-in-life economic outcomes and economic growth more generally (Keane and Wolpin, 1997; Cunha *et al.*, 2005).

Section II examines where we stand in terms of gender equality and documents the trends in gender equality in recent decades. Women shy away from maths-intensive fields with higher economic returns (Chiara Cavaglia *et al.*, 2020, this issue) earn lower wages (Alex Bryson *et al.*, 2020, this issue), and continue to be penalized upon the arrival of children (Monica Costa Dias *et al.*, 2020, this issue). Women continue to be underrepresented in the labour market, particularly at the top and in certain occupations (Gozde Corekcioglu *et al.*, 2020, this issue). It also brings to the forefront the importance of gender in important economic areas that have traditionally remained gender-neutral in their research approach. For example, Guido Matias Cortes *et al.* (2020, this issue) explore how automating technology is likely to affect men and women in different ways.

In section III I bring efficiency considerations to the forefront of the equality debate by providing a brief historical overview of the economic theories and empirical findings that have been brought forward in economics to explain gender inequalities. In doing so, I situate the papers contained in this issue within the broader theoretical frameworks in economics aimed at understanding the driving forces behind gender economic inequalities. I argue that we find ourselves at a critical point at which to move the current research frontier. Insights from psychology have spurred a new line of research in economics looking at unconscious bias as a driver of discriminating behaviour (see the paper by Marina Della Giusta and Steven Bosworth (2020, this issue). Social norms and culture can tell us a lot about the origins of conscious and unconscious stereotypes and how they evolve, as described by Paola Giuliano (2020, this issue). Creative ways of analysing 'big data' are producing new evidence, from representative samples, of what was previously perceived to be sporadic and non-generalizable discriminatory behaviour (as evidenced by Danula Gamage *et al.* (2020, this issue)). This new body

of evidence is spurring new theoretical and empirical research that is moving the field to a new research frontier.

In section IV I review traditional policies stemming from these models, as well as the potential for innovative policy interventions based on new theories and empirical findings from unconscious bias and cultural change. Section V concludes by setting out the remaining gaps in our understanding and lessons for policy.

II. Gender equality: where do we stand?

One of the more prevalent topics in gender economics deals with gender inequalities in the labour market in terms of the gender wage and representation gaps. Bryson *et al.* (2020) offer a brief and informative historical overview of the gender wage gap and female representation in the labour market in the UK by comparing the 1958 and 1970 cohorts of women over the life cycle. Pre-war labour force participation was always high among women who were single, divorced, or widowed, and increased from around 65 per cent in the 1930s to about 79 per cent in 2019. What changed dramatically over this period was the increase in labour force participation by married women from around 10 per cent in the 1930s to about 75 per cent in 2019. Thus, in the aggregate, single and married women work now at similar rates.

Important changes can be hidden within the aggregate. A striking revelation in this study is that downward trends in the overall gender gap shown in other historical studies hide a tremendous life-cycle variation. Whereas in their twenties the raw gender wage gap more than halved for the younger cohort, to about 9 percentage points, it remained at about 30 percentage points at age 42 for both cohorts. This finding links with the issue of career progression, which is analysed by Corekcioglu *et al.* (2020) who document that, even in a country like Norway, with high female labour force participation rates and smaller gender gaps than in the UK, professional women find it harder than men to reach the top. Azmat and Boring (2020) make a case for diversity policies (gender quotas, mentoring and network programmes, diversity training to change firm culture, and family-friendly policies) as a way to promote gender diversity at the top of the corporate hierarchy.

Differences in men's and women's work experience in mid-life account for much but not all of the raw gender wage gap in both cohorts. This point is nicely picked up by Costa Dias *et al.* (2020), who show that the UK has one of the biggest shares of mothers working part-time (33 per cent), together with Austria, Netherlands, and Germany. Women are also penalized for the loss of experience in the labour market as a result of childbirth. Simulations show that if women were to work full-time over the life cycle, as men do, the pay gap would be halved 18 years after childbirth, and if they were to work at the same rate as men, the pay gap would be a third of what is observed 18 years after child birth.

Inequalities in the labour market go back to inequalities in education, posing important supply-side constraints for firms to be able to increase diversity in employment and promote the career progression for women to managerial positions (Azmat and Boring, 2020). As Cavaglia *et al.* (2020) document, over the last 20 years gender gaps in educational achievement in England have evolved to a point at which they favour girls in several areas (including enrolment in tertiary education). There is one substantial exception: boys have continued to be far more likely to engage in maths-intensive courses at university level, with little change in the gender gap over the last 20 years in subjects like computing, engineering, physics, and mathematics at university. The documented gender difference in maths-intensive courses has important implications for future labour market outcomes because of the well-documented fact that science, technology, engineering, and mathematics (STEM) occupations offer higher wages (Brown and Corcoran, 1997; Black *et al.*, 2008; Blau and Kahn, 2017).

Technological progress and the process of routinization have dramatically changed the labour market landscape in the last decades. Because of occupational segregation, technology shocks can affect men and women in different ways. Yet the gender effects have only been recently investigated. Cortes *et al.* (2020) find that women tended to move out of declining routine-intensive occupations at a faster rate than men. However, in Portugal women seem to be underrepresented in jobs where wage growth as a result of technological change was strongest, whereas the opposite is observed in the US. The

authors explain these opposite findings partly because of social norms in the two countries.

III. The role of gender economics in understanding gender inequality

Discrimination, or unexplained gaps between men and women, is observed in labour market and educational settings, as well as in other markets such as access to credit, housing and health services, offers for products and services, politics, and law enforcement (Knowles *et al.*, 2001; Bertrand and Mullainathan, 2004; List, 2004; Nelson, 2009; Rodgers, 2009; Ewens *et al.*, 2014; Alesina *et al.*, 2018).

Economics has generally found creative ways to show that discrimination exists, and that remaining gender differences in pay and representation in labour markets, for example, are not the result of comparing apples to pears. A technique regularly used to empirically uncover discrimination is natural experiments. The logic behind a natural experiment is that it is similar to a laboratory experiment, but takes place outside of the laboratory and involves humans instead of guinea pigs. In gender economics, natural experiments exploit an institutional set-up that randomly manipulates the perceptions of the potential discriminator of the subject's gender. The seminal paper by Claudia Goldin, the first woman to make it to professor in the economics department at Harvard, compared the same woman musician as she played in blind and non-blind auditions. A woman musician playing in a blind audition behind a screen was more likely to be advanced and hired than when she played in front of a hiring committee who could see her. There was no difference for men (Goldin and Rouse, 2000).

Other types of experiment consist in sending the same CV (or application) to potential employers, with the only variation being the name of the applicant (male or female). The seminal paper by Bertrand and Mullainathan (2004) found that CVs that had been randomly given 'black-sounding' names were significantly less likely to receive callbacks than otherwise identical CVs with 'white-sounding' names. Using the same kind of experiment, women are also less likely to be called to interview despite having the same CV (Azmat and Petrongolo, 2014). One criticism of these studies is that, because 'black-sounding names' are less common in the names distribution, it is not clear that

one can interpret the findings of lower call-back rates for blacks with racial discrimination *per se*. It could be that ‘female-sounding’ names carry other types of information that are picked up by the recruitment team and that have nothing to do with gender (see Guryan and Charles (2013) for a review).

Despite some of the shortcomings in quasi-experimental designs, the evidence gathered from them made it increasingly hard to justify gender gaps using alternative explanations not based on some form of discrimination that benefits men over women. These findings, however, remain silent about what the sources of discrimination really are, and it is precisely our understanding of what determines discrimination that crucially matters for the design and effectiveness of policy. Back in the 1970s, two models were developed to answer whether discrimination is motivated by prejudice (‘taste-based’), or whether it arises because of a lack of information that forces people to rely on some statistical information on the group (‘statistical discrimination’). The theory of taste-based discrimination was first brought forward by Gary Becker in his book *The Economics of Discrimination* (Becker, 1957). The model of taste-based discrimination assumes that employers have an overt prejudice or dislike for a particular group, and are willing to pay a premium to avoid having individuals from that group. The same basic result follows when taste-based discrimination arises from the prejudice of customers or co-workers. Under a taste-based discrimination model employers may be willing to sacrifice the higher productivity (and thus higher firm profits) brought about by a particular woman to favour a less productive man, leading to adverse hiring and promotion decisions for women.

In theory, taste-based discrimination should disappear under perfect markets, as openly biased employers are outcompeted by non-discriminatory employers who are able to hire talented women at a cheaper rate. However, statistical discrimination does not. Under statistical discrimination, the inability to overcome asymmetric information failures ensures that gender inequality may exist and persist even when employers are rational and non-prejudiced. Statistical discrimination arises as a result of a market failure, namely asymmetric information problems between employers and employees. Because employers do not know enough about a potential applicant, or even enough about an existing employee asking for promotion, employers use some of the group’s statistical characteristics to infer the applicant’s qualities. *Statistical discrimination* is

thus a product of both: asymmetric information, coupled with a ‘statistical’ characterization of a particular group (Phelps, 1972; Arrow, 1973). Because women are on average more likely to stop working when they have children, an employer may assign a higher probability of quitting a job after having a child to a woman than to a man (regardless of whether that particular woman plans to have children, or whether her labour market attachment will be lower after child bearing than that of a comparable man). Statistical discrimination can explain the fact, for example, that the gender wage gap persists even among those women who have never had children (Manning and Swaffield, 2008).

(i) Crossing the next research frontier

Data limitations have been one of the factors that have made it difficult to empirically distinguish the type of discrimination at play. There are some exceptions to the problems with economists’ attempts at distinguishing between taste-based and statistical discrimination. For example, some interesting findings come from sports economics, where the data are very granular and natural experiments abound. Using experimental methods, List (2004) rules out taste-based discrimination in the baseball card market, assigning all black and white differences to statistical discrimination. However, taking the type of experiments that are run in sports economics to other markets is difficult because of the amount and quality of information needed. As a result, research in economics took a (more politically correct) stand and ignored taste-based discrimination for the most part to concentrate on statistical discrimination. A survey of the literature found that, of 105 papers in top journals testing for discrimination between 1990 and 2018, only 11 per cent discussed the possibility that statistical discrimination might be based on odious beliefs (Bohren *et al.*, 2019). Similarly, Azmat and Boring (2020) discuss the relatively few papers on the study of sexual harassment in the workplace (an extreme case of taste-based discrimination) and its impact on women’s careers.

New ways of collecting, storing, and analysing rich data from images, text, and other forms of media has provided a new set of rich resources and prompted a singular shift in terms of how we economists think about discrimination. These developments have brought to the surface that systematic discriminating behaviour based on animus may

be ubiquitous rather than anecdotal. Gamage *et al.* (2020) provide a telling summary of this new evidence in the economics profession. Processing of image data has revealed that US economics textbooks feature only men (and if there is any woman, it is always the same one—Janet Yellen, the first women president of the Fed). Similarly, systematic analysis of on-line video content reveals that women get asked more questions in economics seminars than men do—and more questions that are deemed to be unfair (Dupas *et al.*, 2020). A text analysis of words used anonymously by economists in an online discussion forum found that the set of words most predictive of the subject of the online discussion being female are ‘hot, hotter, attractive, gorgeous, pregnant, tits, lesbian, bang, horny’, while for men the top words were ‘homo, motivated, keen, slides, textbook, Nordic’ (Wu, 2018). This new evidence has brought taste-based discrimination models back to the surface.

Second, there is a growing recognition that discriminatory behaviour may arise not from animus, but from either explicit or implicit stereotyping, i.e. fixed ideas about what someone is like. Examples of stereotypes in gender economics are that women are worse than men at mathematics, or that women are better at raising children than men. Not all women are necessarily represented by these stereotypes. Discrimination as a result from stereotyping does not easily fall within traditional taste-discrimination models. Stereotypes do not necessarily entail animus or ill feelings directed at people outside one’s own group as in taste-based discrimination models. Instead, an interesting new line of research attempts to model the formation and evolution of stereotypical beliefs into a statistical discrimination framework. For example, Bordalo *et al.* (2016) assume that a decision-maker forms their beliefs based on the representativeness heuristic, i.e. by overweighting its representative types, defined as the types that occur more frequently in that group than in a baseline reference group.

Stereotyping behaviour can be the result of a conscious thought process, and gender economics has looked at culture to understand precisely where stereotypes that relate to gender roles come from.⁵ There is now mounting evidence that cultural beliefs about the role of women in society can explain much of the gender gap in mathematical attainment between boys and girls across countries (see Guiso *et al.* 2008; Fryer and

⁵ Culture is defined as ‘those customary beliefs and values that ethnic, religious, and social groups transmit fairly unchanged from generation to generation’ (Guiso *et al.*, 2006).

Levitt, 2010; Nollenberger *et al.*, 2016). Sport choices by boys and girls are also affected by cultural values and beliefs, and US states with more gender-equal social norms have been found (*ceteris paribus*) to have a higher percentage of boys and girls breaking stereotypes when making sport choices in high school (Marcén *et al.*, 2020). Similarly, prevailing sexism, i.e. negative or stereotypical beliefs concerning the ability or appropriateness of women engaging in market work rather than home production, has been shown to lower a woman's wages, labour force participation, ages of marriage and childbearing, and labour market outcomes (Charles *et al.*, 2018).

Giuliano (2020) provides a comprehensive review of the literature on culture in gender economics, from the early literature establishing the importance of gender economics in explaining gender gaps, to the literature looking at the origins of cultural beliefs. Geography, rather than biology, is an important determinant of gender roles, either through its direct impact on them, or through the formation of language and technology. For example, based on Boserup's (1970) argument that differences in the role of women in society originate in the adoption of either shifting agriculture, which uses hand-held tools, or plough agriculture, which requires significantly more physical power, Alesina *et al.* (2013) show that societies characterized by plough agriculture developed the belief that the natural place for women was in the home. They empirically show that these beliefs persist to this day and affect the participation of women in activities performed outside the home, including market employment, entrepreneurship, and politics. Similarly, Galor *et al.* (2020) show that the complementarity of early agricultural technology with gender roles also fostered the emergence and prevalence of grammatical gender in a language, and that sex-based grammatical gender, and its association with gender bias, has a significant adverse and persistent impact on female educational attainment. As I discuss in the next section, culture research in gender economics is rapidly moving into a highly policy-relevant line of research literature investigating the conditions that make culture evolve.

Borrowing from the social psychology literature, one very recent line of research in gender economics embraces the notion that stereotyping does not need to be the result of a conscious 'slow' thought process (explicit stereotyping), but rather a 'fast' implicit and unconscious process (implicit stereotyping) (Kahneman, 2011). Discrimination as a result of unconscious mental associations may have remained undetected, even by

discriminatory actors themselves who cannot control automatic thought processes (Bertrand *et al.*, 2005). Della Giusta and Bosworth (2020) offer an excellent review of this new line of research, currently pushing the frontier of knowledge in gender economics.

IV. Effective interventions to address gender inequality

Traditional economic theories of discrimination provide a clear justification for policy to target gender inequalities on efficiency grounds. Taste-based discrimination based on animus calls for interventions to address the market imperfections that perpetuate discriminating behaviour in the first place, or to address discriminating behaviour directly. Gender inequalities stemming from statistical discrimination call for policies to overcome the asymmetric information market failures underlying the discrimination process. Here I review traditional policies stemming from these models, as well as the potential for innovative policy interventions based on new theories and empirical findings from unconscious bias and cultural change.

(i) Tackling taste-based discrimination (animus)

The underlying contradiction in taste-based discrimination has long troubled economists and puts policy-makers in a perfect catch-22. On the one hand, these models predict that perfect competition is the main tool for solving taste-based discrimination. On the other hand, profit-maximizing behaviour on the part of employers, a bedrock assumption of perfect competitive models, is assumed away in taste-based discrimination models where individuals running firms are not profit maximizing because of their tastes. The *Cambridge Dictionary* defines prejudice as ‘an unfair and unreasonable opinion or feeling, especially when formed without enough thought or knowledge’. It is precisely the fact that prejudice cannot be rationalized, the fact that prejudice is ‘unreasonable’, that makes taste-based discrimination models particularly unappealing to economists and leaves policy-makers with a tautological choice: to solve taste-based discrimination policy-makers need to encourage perfect competition; to have perfect competition, policy-makers need to get rid of taste-based discrimination. As a result, interventions have attempted both.

One camp of economists favours the notion of addressing the market imperfections that prevent prejudice against women from disappearing over time (Stiglitz, 1973). A variety of models have shown how wage discrimination arising from prejudice can survive when there are search frictions, as it can generate monopsony powers in unprejudiced firms (Black, 1995; Bowlus and Eckstein, 2002; Lang *et al.*, 2005). The results from this literature call for policies addressing employers' monopsony wage-setting powers that allow women to be paid less than their marginal product because women are often less mobile or because their bargaining power is limited due to caring responsibilities (see Bryson *et al.* (2020) and Blau and Kahn (2017) for a review).

It is, however, generally difficult to provide causal evidence of monopsonistic behaviour, which makes it harder for policy-makers to tackle it (Mayhew and Wills, 2019). Instead, policy-makers have opted to regulate overt prejudice upfront. At the time that taste-based discrimination models were developed back in the 1970s, policies were designed to tackle overt and explicit prejudice stemming from animus. Bryson *et al.* (2020) document a reduction in the gender pay gap of between 15 and 18 per cent as a result of the UK Equal Pay Act mandating equal pay for equal jobs. Similarly, following the passage of Title IX in 1972 in the US, requiring schools to provide equal access to all sport activities by 1978, the number of high-school girls participating in sports as a percentage of female high-school enrolment increased ten-fold from close to 3 in 100 girls in 1972 to almost 30 in 100 girls in 1978 (Stevenson, 2007).

The new evidence pointing towards the ubiquitous nature of taste-based discrimination described in the previous section has recently spurred a new wave of measures aimed at tackling it. Whereas Title IX brought about substantial litigation, anti-discrimination legislation in the labour market proved harder to enforce. It remains very costly for an employee to make a discrimination case, which limits the potential for change (Bryson *et al.*, 2020). Yet, there remains a lot of scope for the implementation of this kind of intervention. A case in point is sexual harassment. Azmat and Boring (2020) highlight that a lot of harassment is being under-reported, and point to policies explicitly prohibiting sexual harassment and outlining the consequences of such behaviour as a first step to tackle this issue. As an alternative to the explicit regulation and punishment

of overt prejudice, statements about a company's culture and accepted behaviour can go a long way in setting the 'rules of the game'. These lighter-touch interventions do not necessarily offer procedures to deal with problematic behaviour, or punishment, but they send a strong signal to discriminating actors about what will be tolerated and accepted (Azmat and Boring, 2020). The economics profession itself is taking its first steps in this direction. The major economic associations have only recently approved codes of conduct making harassment inappropriate behaviour in the profession, with the Royal Economic Society doing so in December 2019.⁶

Partly because of the challenges faced by policy-makers, traditional interventions have focused on compensating women *ex post*. As discussed in Bryson *et al.* (2020) and Azmat and Boring (2020), positive discrimination practices and quotas are an example of an attempt to correct things *ex post*, although the jury is still out on their success regarding the outcomes for women more generally (Bertrand *et al.*, 2019). Maternity leave policies and flexible work initiatives, including working from home and reduced hours of work, also fall within the type of policies aimed at compensating women *ex post* rather than attempting to change norms. However, as Corekcioglu *et al.* (2020) recognize, although family-friendly policies make it easier for women to have children and continue working, they may have negative career effects by perpetuating the traditional social norms. For example, working from home may allow women to continue working while simultaneously contributing a lot of time to home production and care of children, but it may perpetuate gender inequalities in earnings as working from home carries an earnings penalty because of increased time constraints (Stratton (2001) and Bryan and Sevilla (2011)) and reduced productivity (Adams-Prassl, 2020). Corekcioglu *et al.* (2020) find that longer maternity leave possibly decreases the probability of mothers moving to the top of the ladder as it incentivizes women to stay away from the labour market for longer, and favour shorter maternity leave as a proven effective way of getting women back to work. Childcare provision and subsidies can also be an effective policy, particularly in contexts of low labour supply and high childcare costs (Costa Dias *et al.*, 2020).

⁶ <https://www.res.org.uk/resources-page/code-of-conduct-pdf.html>

(ii) Tackling statistical discrimination

Statistical discrimination theory leaves policy-makers and organizations with two options for closing gender gaps. Taking the example from the labour market in the previous section as a proof of concept, the two options entail either solving the asymmetric information problems between employers and employees, or changing the existing negative relationship between caring and labour market productivity.

The first option, trying to enforce contracts between employers and potential female employees on the basis of full information disclosure has proven to be a real challenge. One reason highlighted in the literature is the fact that women themselves find it hard to predict their labour market attachment, making it difficult to set the conditions for such a contractual agreement. For example, Kuziemko *et al.* (2018) show that women on average, and more educated women in particular, tend to underestimate the effect of children and to report that parenthood is harder than they expected. Furthermore, growing research about stereotyping suggests that even when the principal is given full information about an agent, beliefs may not be correctly updated. For example, Reuben *et al.* (2015) show in a laboratory experiment that employers with a higher unconscious gender bias have a higher probability of sub-optimally updating their beliefs after the true ability of the candidate is revealed.

An alternative to deal with asymmetric information problems that is often discussed is to increase the level of uncertainty for all parties involved, so that employers cannot use gender as a source of information when hiring or promoting individuals. Based on the quasi-experimental literature covered in the previous sections, blind auditions and blind CVs are examples of such an intervention. In practice, it is challenging to apply this intervention throughout the entire hiring or promotion process. Also, employers may use other clues to infer information about candidates, which may lead to worse results. One such example is discussed in Bryson *et al.* (2020). A field experiment sent out fictitious CVs with typically black and white names, before and after a law was passed that restricted employers from asking about applicants' criminal histories on job applications.⁷ The authors found that before the restriction was in place, white

⁷ 'Ban the Box' (BTB) was the name of a campaign in the US by advocates for ex-offenders, aimed at removing the tick-box from job application forms where applicants are asked to declare a criminal record.

applicants received 7 per cent more call-backs than similar black applicants, but BTB increased this gap to 43 per cent. Effectively, incrementing the level of uncertainty along one dimension (felony convictions) may have led to a higher proportion of ex-offenders being hired, but increased discrimination along the racial dimension as employers relied on exaggerated impressions regarding felony conviction rates of blacks and whites (Agan and Starr, 2018).

Other policies and interventions are aimed at changing the existing negative relationship between caring responsibilities and labour market productivity. Considerations about the production function are crucial here. An interesting case study comes from the pharmaceutical sector, which has seen a rise in female participation and an elimination of the gender pay gap in the last decades as the computerization of drug prescription systems, which enhanced the ability of pharmacists to pass on clients, made pharmacists effectively perfect substitutes for each other (Goldin and Katz, 2012). In the UK job-sharing, a form of flexible working which enables two employees to voluntarily share the responsibilities and duties of one full-time job, is a policy along these lines.⁸ Job-sharing has proven difficult to implement in practice, particularly in the private sector, precisely because the substitutability of labour inputs that exists in the pharmaceutical sector is not common in other industries. Some of the findings in Cortes *et al.* (2020), about the role of routinization technology and the gender wage gap, suggest that supporting efforts to change the technology of production through computerization and other technical innovations rather than simply subsidizing a job-sharing model across the board may be a fruitful way forward.

(iii) Tackling explicit stereotypes

Successful interventions to tackle discriminatory behaviour arising from explicit stereotypes call for a focus on directly changing biased beliefs about the role of women in society. As pointed out by Giuliano (2020), gender roles are amenable to change, and a relatively recent line of enquiry looks at how long it takes and under what conditions cultural beliefs change. Answers to attitudinal questions have been widely used in gender economics to assess the cultural beliefs on the role of women in society (Fortin, 2005). A quick analysis from the World Value Survey (WVS), comparing the

⁸ Refer: <https://www.gov.uk/government/collections/job-sharing-in-the-civil-service>

answers given by respondents regarding their level of agreement to the statement ‘when a mother works for pay, the children suffer’, reveals that while in 1990 about 50 per cent of the population in the US agreed with that statement, in 2017 just 17 per cent of the population sampled did.

Shocks to the sex ratio as a result of wars have proven effective at changing gender roles. For example, following the Second World War many women entered the labour market, changing cultural beliefs about the role of women in society (Acemoglu *et al.*, 2004; Fernandez *et al.*, 2004; Goldin and Olivetti, 2013). However, these kind of shocks do not generally offer a promising or practical model for policy-makers.

Other active attempts to change social norms about the gender division of labour are paternity leave policies, which allow fathers to take leave when children are born. Paternity leave policies directly challenge the norm that women should be the main care-givers and men the breadwinners. The quasi-experimental evidence on the effects of paternity leave on household specialization is not clear-cut. Farré and González (2019) and Tamm (2019) show that paternity leave leads to a persistent increase in fathers’ involvement in childcare in the cases of Spain and Germany, respectively. However, Ekberg *et al.* (2013) do not find an effect of ‘daddy months’ in Sweden on fathers’ likelihood of taking medical leave to care for children. COVID-19 has proven to be a natural experiment for many households, as lockdown measures in many countries meant that men and women had to stay at home with their children for longer periods than the usual paternity leave mandate would allow. The short-run effects, however, do not show responsibilities being shared more equally (Sevilla and Smith, 2020; Adams-Prassl, 2020).

Giuliano (2020) emphasizes social learning as a way in which social norms may change. For example, when inactive women see that more women are at work, they tend to supply their labour (Fogli and Veldkamp, 2011). Similarly, exposure of women from a non-socialist regime to women from a socialist regime altered their behaviour and gender-role attitudes (Schmitz and Weinhardt, 2019). Role model and mentoring initiatives are two policies stemming from cultural learning models, and more research is needed to understand the pathways that make these interventions successful (Azmat and Boring, 2020). However, cultural values and beliefs are deeply rooted in societies,

and beliefs about the role of women in society in particular evolve very slowly. Interventions based on information, on the other hand, have proven to be quite effective at changing social norms. Based on the idea from behavioural economics that individuals are biased in their beliefs about others' perceptions and values, an experiment in Saudi Arabia that provided information about the true beliefs of other men about their willingness to let their wives work changed the attitudes of the treated group of men favourably towards their wives working, and increased the wives' of the treated group propensity to work (Bursztyn *et al.*, 2018).

(iv) Tackling implicit 'stereotypes'

When unconscious bias is at play, neither coercive policies targeting animus nor attempts to change explicit cultural beliefs may be effective at changing discriminatory behaviour. Instead, interventions involving small nudges can be effective, cheap, and easy to implement, and can limit the amount of discrimination without explicitly prohibiting agents from taking decisions against their will (Bohnet, 2016). Such interventions can include changing the context in which decisions are made so that more conscious processes kick in. Leaving more time for reviewers or implementing structured processes based on task cues rather than social cues are other examples (Bertrand *et al.*, 2005). Promoting increasing contact with the discriminated group can also reduce implicit bias (Allport, 1958). Spreading the decision-making process across several individuals can also help avoid stereotyping, although having more women in a hiring committee does not seem to lead to a higher probability of success for a female candidate (Bagues and Esteve-Volart, 2010). Outsourcing the decision-making process to a priori unbiased algorithms was initially thought to be a convenient solution, yet there are now concerns that black-box algorithms that train on historical data may reproduce and reinforce existing discrimination and biases (Rambacham *et al.*, 2020).

A popular tool for addressing unconscious-bias processes in the decision-maker consists in revealing the extent of implicit bias to increase awareness about implicit race and gender associations. Unconscious bias is often measured by an Implicit Association Test (IAT), which is a computer-based tool developed by social psychologists, designed to minimize the risk of social desirability bias (Greenwald *et al.*, 2009). An increasing number of firms and institutions, including Harvard

University, administer the IAT to their employees.⁹ In the educational context, Alesina *et al.* (2018) recently showed that revealing stereotypes to teachers, stereotypes measured by the IAT, could be a powerful intervention to decrease discrimination in grading. However, when provided feedback about their own implicit associations, individuals may react defensively and question the validity of the IAT (O'Brien *et al.*, 2010; Howell *et al.*, 2015; Sukhera *et al.*, 2018). Della Giusta and Bosworth (2020) enumerate the principles that a well-designed unconscious bias training must have, and their experiment reveals that making individuals aware of their bias may lead to moral licensing behaviour that counteracts the intervention, particularly for women.

As with social norms, stereotypes may also lead to self-fulfilling prophecies by influencing the behaviour of discriminated groups in the direction of the stereotypes, and have been argued to be at the root of the finding that girls are generally less competitive, for example. Individuals exposed to bias towards their own group may have less self-confidence (Bordalo *et al.*, 2018). Carlana (2019) shows that teachers' stereotypes affect the gender gap in maths, track choice, and self-confidence in mathematical abilities for girls in middle school. Glover *et al.* (2017) provide evidence that exposure to managers with stronger implicit bias negatively affects the performance of minorities in the workplace.

Interventions to prevent the unconscious internalizing of stereotypes should start early in childhood, and interventions to teach girls how to code may increase resilience and grit (Carlana and Fort, 2019). Small nudges like changing the words in a job advertisement may compensate for the fact that women tend to be more cautious before putting an application forward. Similarly, recognizing the role of language as a transmission mechanism of gender norms and gender inequality described in the previous section, policy-makers have recently incentivized the use of gender neutral language. In English, the use of police officer rather than policeman is an example of such policy intervention. Moving towards neutrality is more challenging in sex-based grammatically gendered languages such as Spanish, where a generic masculine is the default when it is not specified whether the subject is male or female.

⁹ Harvard University strongly encourages 'every search committee member to take at least one Implicit Association Test (IAT)' (<https://faculty.harvard.edu/recruitment-best-practices>)—Alesina *et al.* (2018).

V. Conclusion

This assessment has documented recent findings on gender inequality. In doing so it has identified the new lines of research emerging from the discrimination literature, and described the practical implications for the implementation of policies targeting gender inequality. The emergence of ‘big data’ and artificial intelligence has documented the ubiquitous nature of taste-based discrimination and has brought the interest in taste-based discrimination models back to the surface. Coercive policies directly tackling overt discrimination, such as professional codes of conduct, are back in vogue. Influence from social psychology has led to a growing recognition among economists that discriminatory behaviour may arise not from animus, but from implicit stereotyping. Nudges that attempt to tackle unconscious bias have gained in popularity, particularly in the corporate sector. Development in the economics of culture has contributed to our understanding of where stereotypes come from and how they are formed, paving the way to interventions that aim to change traditional gender roles more directly.

This assessment has reviewed interventions aimed at the individual discriminating, and also at the individual being discriminated against. These latter interventions may work in either changing girls’ and women’s perceptions away from stereotypes (for example, by using gender-neutral language). However, breaking with the stereotype may not always pay off for women. In the same way that returns to schooling are different for men and women (Charles and Luoh, 2003; Dougherty, 2005; Hubbard, 2011), Exley *et al.* (2020) have recently found in a laboratory experiment that women avoid negotiations more often than men because negotiating more does not pay off for them.

Because of the challenges in implementing some of these interventions, policy initiatives that attempt to compensate women *ex post* rather than trying to change norms continue to be necessary. Yet, although these policies may compensate women in the short run, they may have negative effects in the long run as they reinforce the social norm by either promoting it, or by instigating discriminatory behaviour through backlash as the other group may feel it is unfair.

References

- Acemoglu, D., Autor, D., and Lyle, D. (2004), 'Women, War and Wages: The Effect of Female Labor Supply of the Wage Structure at Midcentury', *Journal of Political Economy*, **112**(3), 497–550.
- Adams-Prassl, A. (2020), 'The Gender Wage Gap on an Online Labour Market: The Cost of Interruptions', CEPR Working Paper DP14294.
- Agan, A., and Starr, S. (2018), 'Ban the Box, Criminal Records, and Racial Discrimination: A Field Experiment', *The Quarterly Journal of Economics*, **133**(1), 191–235.
- Aizer, A. (2011), 'Poverty, Violence, and Health: The Impact of Domestic Violence During Pregnancy on Newborn Health', *Journal of Human Resources*, **46**(3), 518–38.
- Alesina, A., Giuliano, P., and Nunn, N. (2013), 'On the Origins of Gender Roles: Women and the Plough', *The Quarterly Journal of Economics*, **128**(2), 469–530.
- Carlana, M., Ferrara, E. L., and Pinotti, P. (2018), 'Revealing Stereotypes: Evidence from Immigrants in Schools', NBER Working Papers No. w25333, National Bureau of Economic Research.
- Allport, G. W. (1958), *The Nature of Prejudice: Abridged*, Doubleday.
- Almond, D., Currie, J., and Duque, V. (2018), 'Childhood Circumstances and Adult Outcomes: Act II', *Journal of Economic Literature*, **56**(4), 1360–446.
- Anderberg, D., and Moroni, G. (2020), 'Exposure to Intimate Partner Violence and Children's Dynamic Skill Accumulation: Evidence from a UK Longitudinal Study', *Oxford Review of Economic Policy*, **36**(4).
- Arrow, K. J. (1973), 'The Theory of Discrimination', in O. Ashenfelter and A. Rees (eds), *Discrimination in Labor Markets*, Princeton, NJ, Princeton University Press, 3–33.
- Azmat, G., and Boring, A. (2020), 'Gender Diversity in Firms', *Oxford Review of Economic Policy*, **36**(4).
- Petrongolo, B. (2014), 'Gender and the Labor Market: What Have We Learned from Field and Lab Experiments?', *Labour Economics*, **30**, 32–40.

- Bagues, M. F., and Esteve-Volart, B. (2010), ‘Can Gender Parity Break the Glass Ceiling? Evidence from a Repeated Randomized Experiment’, *The Review of Economic Studies*, **77**(4), 1301–28.
- Becker, G. S. (1965), ‘A Theory of the Allocation of Time’, *The Economic Journal*, **75**(299), 493–517.
- (1957), *The Economics of Discrimination*, Chicago, IL, University of Chicago Press.
- Bertrand, M., and Mullainathan, S. (2004), ‘Are Emily and Greg more Employable than Lakisha and Jamal? A Field Experiment on Labor Market Discrimination’, *American Economic Review*, **94**(4), 991–1013.
- Chugh, D., and Mullainathan, S. (2005), ‘Implicit Discrimination’, *American Economic Review*, **95**(2), 94–8.
- Black, S. E., Jensen, S., and Lleras-Muney, A. (2019), ‘Breaking the Glass Ceiling? The Effect of Board Quotas on Female Labour Market Outcomes in Norway’, *The Review of Economic Studies*, **86**(1), 191–239.
- Black, D. A. (1995), ‘Discrimination in an Equilibrium Search Model’, *Journal of Labor Economics*, **13**(2), 309–34.
- Haviland, A. M., Sanders, S. G., and Taylor, L. J. (2008), ‘Gender Wage Disparities among the Highly Educated’, *Journal of Human Resources*, **43**(3), 630–50.
- Blau, F. D., and Kahn, L. M. (2017), ‘The Gender Wage Gap: Extent, Trends, and Explanations’, *Journal of Economic Literature*, **55**(3), 789–865.
- Bohnet, I. (2016), *What Works*, Cambridge, MA, Harvard University Press.
- Bohren, J. A., Haggag, K., Imas, A., and Pope, D. G. (2019), ‘Inaccurate Statistical Discrimination’, NBER Working Papers No. w25935, National Bureau of Economic Research.
- Bordalo, P., Coffman, K., Gennaioli, N., and Shleifer, A. (2016), ‘Stereotypes’, *The Quarterly Journal of Economics*, **131**(4), 1753–94.
- — — (2018), ‘Beliefs about Gender’, *American Economic Review*, **109**(3), 739–73.
- Boserup, E. (1970), *Woman’s Role in Economic Development*, London, George Allen & Unwin, reprinted in 2007 by Earthscan, Abingdon and New York.
- Bowlus, A. J., and Eckstein, Z. (2002), ‘Discrimination and Skill Differences in an Equilibrium Search Model’, *International Economic Review*, **43**(4), 1309–45.

- Brown, C., and Corcoran, M. (1997), ‘Sex-based Differences in School Content and the Male–Female Wage Gap’, *Journal of Labor Economics*, **15**(3), 431–65.
- Bryan, M. L., and Sevilla-Sanz, A. (2011), ‘Does Housework Lower Wages? Evidence for Britain’, *Oxford Economic Papers*, **63**(1), 187–210.
- Bryson, A., Joshi, H., Wielgoszewska, B., and Wilkinson, D. (2020), ‘A Short History of the Gender Wage Gap in Britain’, *Oxford Review of Economic Policy*, **36**(4).
- Bursztyjn, L., Gonzalez, A., and Yanagizawa-Drott, D. (2018), ‘Misperceived Social Norms: Female Labor Force Participation in Saudi Arabia’, NBER WP 24736.
- Carlana, M. (2019), ‘Implicit Stereotypes: Evidence from Teachers’ Gender Bias’, *The Quarterly Journal of Economics*, **134**(3), 1163–224.
- Fort, M. (2019), ‘Girls Code It Better’, unpublished manuscript.
- Cavaglia, C., Machin, S., McNally, S., and Ruiz-Valenzuela, J. (2020), ‘Gender, Achievement, and Subject Choice in English Education’, *Oxford Review of Economic Policy*, **36**(4).
- Charles, K. K., and Luoh, M.-C. (2003), ‘Gender Differences in Completed Schooling’, *Review of Economics and Statistics*, **85**(3), 559–77.
- Guryan, J., and Pan, J. (2018), ‘The Effects of Sexism on American Women: The Role of Norms vs Discrimination’, NBER Working Papers No. w24904, National Bureau of Economic Research.
- Chiappori, P. A. (1997), ‘Introducing Household Production in Collective Models of Labor Supply’, *Journal of Political Economy*, **105**(1), 191–209.
- Corekcioglu, G., Francesconi, M., and Kunze, A. (2020), ‘Do Generous Parental Leave Policies Help Top Female Earners?’ , *Oxford Review of Economic Policy*, **36**(4).
- Cortes, G. M., Oliveira, A., and Salomons, A. (2020), ‘Do Technological Advances Reduce the Gender Wage Gap?’ , *Oxford Review of Economic Policy*, **36**(4).
- Costa Dias, M., Joyce, R., and Parodi, F. (2020), ‘The Gender Pay Gap in the UK: Children and Experience in Work’, *Oxford Review of Economic Policy*, **36**(4).
- Cunha, F., Heckman, J. J., and Navarro, S. (2005), ‘Separating Uncertainty from Heterogeneity in Life Cycle Earnings’, The 2004 Hicks Lecture, *Oxford Economic Papers*, **57**(2), 191–261.
- Della Giusta, M., and Bosworth, S. (2020), ‘Bias and Discrimination: What Do We Know?’ , *Oxford Review of Economic Policy*, **36**(4).

- Dougherty, C. (2005), ‘Why are the Returns to Schooling Higher for Women than for Men?’, *Journal of Human Resources*, **40**(4), 969–88.
- Dupas, P., Modestino, A., Niederle, M., and Wolfers, J. (2020), ‘Gender and the Dynamics of Economics Seminars’, presentation at lunchtime seminar, Economic Society of Australia.
- Ekberg, J., Eriksson, R., and Friebel, G. (2013), ‘Parental Leave—A Policy Evaluation of the Swedish “Daddy-Month” Reform’, *Journal of Public Economics*, **97**, 131–43.
- Ewens, M., Tomlin, B., and Wang, L.C. (2014), ‘Statistical Discrimination or Prejudice? A Large Sample Field Experiment’, *Review of Economics and Statistics*, **96**(1), 119–34.
- Exley, C. L., Niederle, M., and Vesterlund, L. (2020), ‘Knowing When to Ask: The Cost of Leaning In’, *Journal of Political Economy*, **128**(3), 816–54.
- Farré, L., and González, L. (2019), ‘Does Paternity Leave Reduce Fertility?’, *Journal of Public Economics*, **172**, 52–66.
- Fernandez, R., Fogli, A., and Olivetti, C. (2004), ‘Mothers and Sons: Preference Formation and Female Labor Force Dynamics’, *Quarterly Journal of Economics*, **119**(4), 1249–99.
- Fogli, A., and Veldkamp, L. (2011), ‘Nature or Nurture? Learning and the Geography of Female Labor Force Participation’, *Econometrica*, **79**(4), 1103–38.
- Fortin, N. M. (2005), ‘Gender Role Attitudes and the Labour-market Outcomes of Women across OECD Countries’, *Oxford Review of Economic Policy*, **21**(3), 416–38.
- Fryer, Jr, R. G., and Levitt, S. D. (2010), ‘An Empirical Analysis of the Gender Gap in Mathematics’, *American Economic Journal: Applied Economics*, **2**(2), 210–40.
- Galor, O., Özak, Ö., and Sarid, A. (2020), ‘Linguistic Traits and Human Capital Formation’, in *AEA Papers and Proceedings*, **110**, 309–13.
- Gamage, D. K., Sevilla, A., and Smith, S. (2020), ‘Women in Economics: A UK Perspective’, *Oxford Review of Economic Policy*, **36**(4).
- Giuliano, P. (2020), ‘Gender and Culture’, *Oxford Review of Economic Policy*, **36**(4).
- Glover, D., Pallais, A., and Pariente, W. (2017), ‘Discrimination as a Self-fulfilling Prophecy: Evidence from French Grocery Stores’, *The Quarterly Journal of Economics*, **132**(3), 1219–60.

- Goldin, C. (2015), ‘Gender and the Undergraduate Economics Major: Notes on the Undergraduate Economics Major at a Highly Selective Liberal Arts College’, manuscript, 12 April.
- Katz, L. F. (2012), ‘The Most Egalitarian of All Professions: Pharmacy and the Evolution of a Family-friendly Occupation’, NBER Working Papers No. w18410, National Bureau of Economic Research.
- Olivetti, C. (2013), ‘Shocking Labor Supply: A Reassessment of the Role of World War II on Women’s Labor Supply’, *American Economic Review*, **103**(3), 257–62.
- Rouse, C. (2000), ‘Orchestrating Impartiality: The Impact of “Blind” Auditions on Female Musicians’, *American Economic Review*, **90**(4), 715–41.
- Greenwald, A. G., Poehlman, T. A., Uhlmann, E. L., and Banaji, M. R. (2009), ‘Understanding and Using the Implicit Association Test: III. Meta-analysis of Predictive Validity’, *Journal of Personality and Social Psychology*, **97**(1), 17.
- Guiso, L., Sapienza, P., and Zingales, L. (2006), ‘Does Culture Affect Economic Outcomes?’, *Journal of Economic Perspectives*, **20**(2), 23–48.
- Monte, F., Sapienza, P., and Zingales, L. (2008), ‘Culture, Gender, and Math’, *Science*, **320**(5880), 1164–5.
- Guryan, J., and Charles, K. K. (2013), ‘Taste-based or Statistical Discrimination: The Economics of Discrimination Returns to its Roots’, *The Economic Journal*, **123**(572), F417–F432.
- Hubbard, W. H. J. (2011). ‘The Phantom Gender Difference in the College Wage Premium’, *Journal of Human Resources*, **46**(3), 568–86.
- Howell, J. L., Gaither, S. E., and Ratliff, K. A. (2015), ‘Caught in the Middle: Defensive Responses to IAT Feedback among Whites, Blacks, and Biracial Black/Whites’, *Social Psychological and Personality Science*, **6**(4), 373–81.
- Hsieh, C., Hurst, E., Jones, C. I., and Klenow, P. J. (2019), ‘The Allocation of Talent and US Economic Growth’, *Econometrica*, **87**(5), 1439–74.
- Kahneman, D. (2011), *Thinking, Fast and Slow*, New York, Farrar, Straus and Giroux.
- Keane, M. P., and Wolpin, K. I. (1997), ‘The Career Decisions of Young Men’, *Journal of Political Economy*, **105**(3), 473–522.
- Knowles, J., Persico, N., and Todd, P. (2001), ‘Racial Bias in Motor Vehicle Searches: Theory and Evidence’, *Journal of Political Economy*, **109**(1), 203–29.

- Kolovich, L., Malta, V., Newiak, M., and Robinson, D. (2020), ‘Gender Equality and Macroeconomic Outcomes: Evidence and Policy Implications’, *Oxford Review of Economic Policy*, **36**(4).
- Kuziemko, I., Pan, J., Shen, J., and Washington, E. (2018), ‘The Mommy Effect: Do Women Anticipate the Employment Effects of Motherhood?’, NBER Working Papers No. w24740, National Bureau of Economic Research.
- Lang, K., Manove, M., and Dickens, W. T. (2005), ‘Racial Discrimination in Markets with Announced Wages’, *American Economic Review*, **95**(4), 1327–40.
- List, J. A. (2004), ‘The Nature and Extent of Discrimination in the Marketplace: Evidence from the Field’, *The Quarterly Journal of Economics*, **119**(1), 49–89.
- Lundberg, S., and Pollak, R. A. (1993), ‘Separate Spheres Bargaining and the Marriage Market’, *Journal of Political Economy*, **101**(6), 988–1010.
- Manning, A., and Swaffield, J. (2008), ‘The Gender Gap in Early-career Wage Growth’, *The Economic Journal*, **118**(530), 983–1024.
- Marcén, M., Morales, M., and Sevilla, A. (2020), ‘Gender Stereotyping in Sports’, IZA Working Paper.
- Mayhew, K., and Wills, S. (2019), ‘Inequality: An Assessment’, *Oxford Review of Economic Policy*, **35**(3), 351–67.
- Nelson, T. D. (ed.) (2009), *Handbook of Prejudice, Stereotyping, and Discrimination*, Washington, DC, Psychology Press.
- Nollenberger, N., Rodríguez-Planas, N., and Sevilla, A. (2016), ‘The Math Gender Gap: The Role of Culture’, *American Economic Review*, **106**(5), 257–61.
- O’Brien, L. T., Crandall, C. S., Horstman-Reser, A., Warner, R., Alsbrooks, A., and Blodorn, A. (2010), ‘But I’m No Bigot: How Prejudiced White Americans Maintain Unprejudiced Self-images’, *Journal of Applied Social Psychology*, **40**(4), 917–46.
- Phelps, E. S. (1972), ‘The Statistical Theory of Racism and Sexism’, *The American Economic Review*, **62**(4), 659–61.
- Rambachan, A., Kleinberg, J., Ludwig, J., and Mullainathan, S. (2020), ‘An Economic Perspective on Algorithmic Fairness’, in *AEA Papers and Proceedings*, Vol. 110, 91–5.
- Reuben, E., Sapienza, P., and Zingales, L. (2015), ‘Taste for Competition and the Gender Gap among Young Business Professionals’, NBER Working Papers No. w21695, National Bureau of Economic Research.

- Rodgers, W. M. (2009), *Handbook on the Economics of Discrimination*, Cheltenham, Edward Elgar.
- Schmitz, S., and Weinhardt, F. (2019), ‘Immigration and the Evolution of Local Cultural Norms’, IZA DP 12509.
- Sevilla, A., and Smith, S. (2020), ‘Baby Steps: The Gender Division of Childcare during the COVID-19 Pandemic’, *Oxford Review of Economic Policy*, **36**(S1).
- Stevenson, B. (2007), ‘Title IX and the Evolution of High School Sports’, *Contemporary Economic Policy*, **25**(4), 486–505.
- Stiglitz, J. E. (1973), ‘Approaches to the Economics of Discrimination’, *The American Economic Review*, **63**(2), 287–95.
- Stratton, L. S. (2001), ‘Why Does More Housework Lower Women’s Wages? Testing Hypotheses Involving Job Effort and Hours Flexibility’, *Social Science Quarterly*, **82**(1), 67–76.
- Sukhera, J., Milne, A., Teunissen, P. W., Lingard, L., and Watling, C. (2018), ‘The Actual versus Idealized Self: Exploring Responses to Feedback about Implicit Bias in Health Professionals’, *Academic Medicine*, **93**(4), 623–9.
- Tamm, M. (2019), ‘Fathers’ Parental Leave-taking, Childcare Involvement and Labor Market Participation’, *Labour Economics*, **59**, 184–97.
- Wu, A. H. (2018), ‘Gendered Language on the Economics Job Market Rumors Forum’, *AEA Papers and Proceedings*, **108**, 175–9.