
Discrimination in Developing Countries

Discussion Paper no. [2024-03](#)**Pushkar Maitra and Ananta Neelim****Abstract:**

In this chapter, we provide an account of the discrimination literature in economics and discuss two widely pursued policies to tackle discrimination in developing countries. We present a discussion of the main classical theories of discrimination: taste-based discrimination and statistical discrimination and extend our discussion to the new and developing literature on norm-based discrimination. We briefly discuss the measurement approaches to discrimination and alternative policies to reduce it. We focus on two main policy prescriptions: increasing contact between adversarial groups (the contact hypothesis); and legislating diversity (through affirmative action policies including quotas).

Keywords: Discrimination, Inter-group Contact, Gender, Minority-Majority, Experiments**JEL Classification:** C9, D9, J1, J7

Pushkar Maitra: Department of Economics, Monash University, Clayton Campus, VIC 3800 Australia (email: Pushkar.Maitra@monash.edu); Ananta Neelim: Tasmanian School of Business and Economics, University of Tasmania, Hobart, Tasmania 7001, Australia (email: ananta.neelim@utas.edu.au).

Discrimination in Developing Countries*

Pushkar Maitra

Department of Economics

Monash University

Clayton Campus

VIC 3800

Australia

Pushkar.Maitra@monash.edu

Ananta Neelim

Tasmanian School of Business and Economics

University of Tasmania

Hobart

Tasmania 7001

Australia

ananta.neelim@utas.edu.au

December 2023

Chapter for *Handbook of Experimental Development Economics*. U. Dasgupta and P. Maitra (eds).

Key Words: Discrimination, Inter-group Contact, Gender, Minority-Majority, Experiments, Developing Countries

JEL Codes: C9, D9, J1, J7

*We thank Abhijeet Singh and Russell Smyth for their comments on an earlier version of this chapter.

1 Introduction

Traditionally, discrimination refers to situations where an individual, or a group of individuals, receive differential treatment because of their background (gender, race, religion etc.) or personal characteristics (sexual orientation, disability status, political views etc.). Based on this definition, discrimination is widespread and persistent in many spheres of life. According to UN Secretary-General's annual report on the Sustainable Development Goals in 2020, one in five people in the 31 surveyed countries reported facing some form of discrimination. For example, in almost all labour markets in the world, women earn a fraction of and are less likely to be represented in leadership positions relative to men. Similarly, in criminal justice systems or in housing markets individuals from minority backgrounds are more likely to face discrimination and harassment relative to their majority counterparts.

From the middle of the 20th century, the reduction of discrimination has been an area of significant policy focus. Many interventions, drawn from economics and other social science research, were rolled out with different degrees of success. For example, affirmative action policies in the form of quotas have been implemented in workplace, educational and political settings. Similarly, interventions based on the psychology literature around improving inter-group interactions have also been implemented. Finally, in the last decade, many interventions to reduce the impact of behavioural biases have been pursued. Evaluations of these interventions have provided us with an evidence base of what works and more importantly, have allowed for a richer understanding of why and how (*i*) discrimination occurs and persists and (*ii*) which interventions succeed and under what conditions.

Economists have investigated differential labour market outcomes across groups for decades with an early emphasis on understanding its various causes. Starting with Gary Becker's seminal work *The Economics of Discrimination* (Becker, 1957), early papers modelled whether differential outcomes observed across groups were due to prejudice or differences in expected productivity. Differential outcomes based on prejudicial preferences was defined as *taste-based* discrimination. On the other hand, differential outcomes across groups that arose from firms trying to efficiently solve informational issues around expected productivity was termed as *statistical* discrimination. Insights from the labour market discrimination models were, over time, adapted for other markets. In recent years, there has been a resurgence in how economists have modelled discrimination going beyond the differential outcomes observed. This new research often draws on findings from the psychology and sociology literature to understand why discrimination occurs and incorporates the role of cognitive biases and social norms that impact the behaviour of different decision-makers and enact systemic barriers for certain groups in society.

In this chapter, our goal is to provide an account of the discrimination literature in economics and discuss two widely pursued policies to tackle discrimination in developing countries. Section 2 presents a discussion of the theories of discrimination: taste-based discrim-

ination (Section 2.1), statistical discrimination (Section 2.2). We extend our discussion to the new and developing literature on norm-based discrimination (Section 2.3). Section 3 briefly discusses the different approaches to measuring discrimination. Section 4 discusses alternative policies to reduce discrimination. We focus on two main policy prescriptions: (a) increasing contact between adversarial groups (the contact hypothesis); and (b) legislating diversity (through affirmative action policies including quotas). Finally Section 5 concludes with a discussion of directions for future research.

2 Theories of Discrimination

Traditionally, economists have modelled discrimination to explain differential outcomes attained by majority (White or male) over minority (non-White or female) workers in markets. Early models assumed that employers (i) hold prejudicial preferences for employing and promoting majority workers and/or (ii) hold beliefs (correctly or incorrectly) that majority workers are more productive than minority workers. These models explore the effect of existing discrimination on equilibrium behaviour in the short and long run. For example, the long-run impact of discrimination on the human capital accumulation of majority and minority workers can explain the persistence of discrimination over time. Recently, newer models of discrimination have investigated other avenues through which discrimination arises and persists. The most prominent of these avenues is in the form of social norms, which prescribe individuals from majority and minority groups differently in markets. In the next three subsections, (Sections 2.1–2.3), we briefly discuss each of these motivations. For more detailed reviews, we refer readers to the excellent recent papers by [Lang and Kahn-Lang Spitzer \(2020\)](#) and [Onuchic \(2022\)](#).

2.1 Preference-based discrimination

In his seminal work, [Becker \(1957\)](#), posits that discrimination arises because employers hold prejudicial preferences (or a taste for discrimination) against minority workers. For these types of employers, minority workers of equal expected productivity will only be hired (or promoted) if they are willing to accept lower wages (or produce more output at the same wages) to compensate for the disutility they cause on the prejudiced employer. The taste-based models of discrimination predict worker segregation, at least in the short run: discriminatory employers will exclusively hire majority workers and the non-discriminatory employers will exclusively hire minority workers. Further, these models predict wage differentials between majority and minority workers arise if the fraction of the prejudicial employers is relatively large. Otherwise, wages will equalise over time as in the long run: the marginal productivity of workers from majority and minority groups will have to equalise as more efficient non-discriminatory employers will enter the market.

The predictions of early taste-based models are, however, typically not supported by empirical evidence. Very rarely do we observe complete segregation in the labour market and wage gaps between minority and majority workers usually persist even after accounting for productivity differences.

Models based on Becker's original work explore channels to reconcile the difference between the theoretical predictions and the empirical findings, particularly around the persistence of wage gaps in the labour market. For example, [Black \(1995\)](#) introduces labour market frictions *via* a costly job search process to explain the persistence of discrimination in equilibrium. He argues that in such labour markets, reservation wages are not only determined by the productivity characteristics of workers, but also the costs associated with finding work. At the margin, for any wage offer, workers consider the difference between that wage and the expected payoffs from continuing with the costly search process. They only accept the offer if the difference is positive. Going back to our example, the presence of discriminatory employers lowers the expected payoff from continuing in the search process, but only for minority workers. This would imply that for the same productivity characteristics, minority workers will be willing to accept lower wage offers because of their relatively lower reservation wages. Employers, discriminatory or otherwise, can exploit the resulting differences in reservation wages across worker types. Accordingly, wage gaps across minority and majority workers will persist in the long run.

More recently, researchers have tried to identify the behavioural underpinnings of prejudicial preferences assumed in taste-based models. They appeal to established theories in psychology that explain inter-group interactions. The most prominent theories are that of social categorization and social identity ([Tajfel and Turner, 1978](#)). In short, these theories posit that human beings innately categorise people into groups based on their characteristics (gender, ethnicity, profession etc.) and treat individuals who share the same identity as them better than those who do not. In a seminal paper, [Chen and Li \(2009\)](#) systematically validates the main predictions of the social identity theory in economic settings. Using a minimal group paradigm approach, where group identities are artificially manufactured in a laboratory setting, they find that participants are significantly more charitable and less envious of in-group partners compared to out-group partners across a set of social preference games. More importantly, the original psychology theories and their adaptations in the economics literature thus created new avenues through which discrimination (particularly in-group bias) can be reduced. We revisit some of these avenues in [Section 4](#).

2.2 Statistical discrimination

The starting point of many labour market models of discrimination is the assumption that hiring decisions are based solely on the productivity of workers. However, often these decisions are made by employers facing asymmetric information: they can only observe a noisy signal of the productivity of workers. In the simplest statistical models, (see,

for example [Phelps, 1972](#), [Arrow et al., 1973](#)) employers are assumed to use observable characteristics of workers (gender or race) in conjunction with the noisy signal (GPA or test-scores) to make inferences about their expected productivity. This implies that identical signals across worker types may receive different expected productivity inferences, which in turn can lead to discrimination in the form of differential wage offers. In other words, unlike models based on prejudicial preferences, statistical models posit that discrimination arises because of the efforts of the employer to efficiently solve a signal extraction problem.

There are two ways in which the identity of an applicant can influence the firm's inference of her expected productivity given a realised signal. First, if the underlying productivity distribution across worker types at the population level is different: (i) the average productivity of the majority workers higher than the minority workers and (ii) the variance of productivity of the minority workers is higher than that of the majority workers, and therefore the expected productivity levels will be lower for the minority workers given a realised signal. Second, even if there are no productivity differences across majority and minority workers at the population level, expectations about differential productivity may arise if the signal is noisier for one of the worker types. For example, suppose GPA scores are less informative about female productivity than male productivity, then the inferences about productivity made are more responsive to the GPA scores for males, who produce less noisy signals. It is important to note that if discrimination arises through this channel, it occurs at an individual level. If there are no underlying productivity differences at the population level, on average no discrimination will exist at the mean or the group level. In other words, some minority workers will fare better, while others will fare worse. To explain group-level differential wages, [Aigner and Cain \(1977\)](#) propose a model where employers are assumed to be risk-averse and hence care about the noisiness of the signal received. Accordingly, they offer wages that account for this noisiness and in doing so advantage groups that have higher signal accuracy. Therefore group-level statistical discrimination is possible even without any differences in the productivity distributions across groups.

Since productivity is central to statistical models, a significant amount of research has also focused on the dynamic effects of discrimination and the persistence of wage gaps in the labour market. Specifically, this research has modelled workers' human capital investments prior to entering the job market as a way to explain the persistence of wage gaps. Intuitively, these models argue that if workers believe that investment is not going to be worthwhile due to the presence of discrimination, they will refrain from doing so. More formally, [Coate and Loury \(1993\)](#) model the complementarity between the outcomes observed in the labour market and worker investment. Their model is characterised by the presence of multiple equilibria with different investment levels on the part of the workers. Assuming employers hold biased priors about the expected productivity of minority workers and discriminate on the basis of that, in equilibrium minority workers would invest less in skill acquisition relative to majority workers. This will perpetuate the negative stereotypes and the resulting discrimination (see, for example [Arrow et al., 1973](#)).

The newer models that build on original statistical discrimination incorporate the impact

of behavioural biases in the signal extraction problem outlined above. Inferring expected productivity from realised signals is often subject to behavioural biases, which lead to inefficient learning on the part of employers as well as workers. Amongst others, these biases include (i) the use of the representative heuristics (Kahneman and Tversky, 1972, Bordalo et al., 2016), (ii) engaging in causal misperceptions (Spiegler, 2020), (iii) committing fundamental attribution errors (Chauvin, 2018) and (iv) misspecified/inaccurate beliefs (Bohren et al., 2023). The consequences of these biases are inefficient and inaccurate discrimination, which in a dynamic setting contribute towards its persistence. For a more thorough review of these new developments in this line of research, we refer readers to the review article by Onuchic (2022).

2.3 Norm-based theories of discrimination

Social norms prescribe and proscribe behaviour in social settings: actions undertaken by individuals may then be evaluated against these rules/norms (Krupka and Weber, 2013). Violations of these rules generate psychological costs (in the form of guilt or shame) or social costs (in the form of sanctions or backlash), which, in turn, sustain these norms over time (Young, 2015). There are several ways in which social norms can influence discriminatory employer behaviour in labour markets. Let us assume the presence of discriminatory social norms in the hiring environment. Here, hiring a minority worker leads to additional costs to the employer due to breaking the aforementioned norms. Employers respond to these costs by discriminating against minority workers. However, it is important to note that this discrimination is not motivated by prejudice against or lower expected productivity about minority workers, rather it is simply the result of adherence to norms.

Another way in which social norms can influence labour market outcomes is by generating identity-specific expectations (Akerlof and Kranton (2000)). These expectations can pose minority individuals with different costs and constraints in social and economic interactions compared to majority individuals. These differential constraints influence skill acquisition from an early age and as a result, can influence which sectors of the economy minority and majority workers will work in. For example, Bian et al. (2017) shows that by as early as six years of age, girls think of boys as smarter and steer themselves away from games that are intended for smarter kids. Buser et al. (2014) shows that these patterns persist in the choice of career tracks. Even at the job application stage, how one presents themselves matter. In an interesting paper, Bursztyn et al. (2017) shows that female MBA students lower their expected salary in response to gender prescriptions. In another study with children from low and high-caste school students in India, Hoff and Pandey (2006) demonstrate that performance difference between low-caste and high-caste students are generated when caste-identity is made public, i.e. norm-based behaviour is made salient. Again, note that the channel through which these behavioural differences arise is different from those highlighted in Section 2.2. Finally, there is voluminous evidence that inform us how behavioural preferences across worker type can impact the wage determination process

(choice of contract and negotiation patterns) which ultimately leads to the persistence of wage gaps in the labour market.

3 Measuring Discrimination

Economists have measured discrimination using different approaches. The purpose of this section is to introduce these approaches and to briefly highlight their advantages and disadvantages. For a more comprehensive discussion, we refer readers to the review article by [Neumark \(2018\)](#).

The *first* approach uses observable data and utilises decomposition techniques to explain wage gaps across groups. In this approach, wage regressions are estimated by controlling for productivity-related differences across minority and majority groups. In doing so the observed wage gap is decomposed to the *explained* and *unexplained* components ([Blinder, 1973](#), [Oaxaca, 1973](#)). The latter is interpreted as a measure of discrimination against the minority group. The biggest critique of this approach is that it is almost impossible to effectively control for all productivity characteristics in wage regressions, especially when they are unmeasured or unobserved. As a result, the unexplained component can under/overestimate the true level of discrimination in the market. For example, in a recent study, [Maitra et al. \(2021\)](#) used lab-in-the-field experiments to measure preferences for negotiation and risk for workers in Vietnam and show that including these preferences (which are unobserved) as additional controls, can reduce the unexplained component of the wage gap by up to 15.5%. Despite the drawbacks of the regression-based decomposition approach, its biggest advantage is that it provides direct evidence of discrimination from the real world and therefore provides a benchmark for tracking discrimination in the labour market over time.

The *second* approach utilises laboratory experiments. This approach prioritises achieving control over the decision-making environment and is therefore aimed at establishing causal relationships between factors (including minority status) and discrimination. Typically, participants are randomly assigned to different conditions, which allows researchers to make such causal inferences. In this space, earlier studies utilised vignettes (a short hypothetical description of situations involving a central person) in labour market contexts to elicit participant judgments regarding these scenarios (for example hiring/promotion decisions). Across vignettes, only the ethnicity or gender of the central person is randomly varied and any difference in participant evaluation due to this manipulation is attributed to discrimination. The biggest critique of vignette studies is that the decisions are hypothetical and are often prone to measurement issues and hypothetical bias ([Harrison, 2014](#)). Note that the goal of vignette studies is to identify the existence of discrimination against specific groups and not to tease apart various causes of discrimination.

More recent laboratory-based studies try to address the aforementioned issues by: (i) making choices of participants incentive compatible; and (ii) administering experimental manipulations that tease out the impact of the various causes of discrimination. The first set of these studies examines the presence of in-group bias in social-preference games. The simplest of the games used is the Dictator game (Forsythe et al., 1994), where the first player decides how much of his endowment he wants to send/share with a second player. In this game, discrimination against second players not sharing the first player's identity is solely driven by preferences and hence can be attributed to taste-based discrimination. The second, more elaborate, game utilised is the Trust or Investment game (Berg et al., 1995), where the first player has to decide how much of his endowment to invest with the second player knowing that (i) the investment will be multiplied by a factor of three and (ii) the return on investment is dependent on the decision of the second player. In this interaction, the first players' discrimination is driven not only on the basis of taste, but also based on expectations of the behaviour of the second player. Therefore, any discrimination observed in this game is a composite effect of both taste-based and statistical discrimination.

In a seminal paper, Fershtman and Gneezy (2001) use both these games to investigate ethnic discrimination in the segmented Israeli Jewish population. Specifically, they examine whether there was systematic discrimination against Eastern Jews (Asian and African immigrants and their Israeli-born off-springs), who often achieve lower levels of earnings and education relative to Ashkenazic Jews (European and American immigrants and their Israeli-born off-springs). Using the Trust game, they established that there was systematic discrimination against men of Eastern origin, who received lower contributions as a second player. Further, their Dictator game results showed that this discrimination was not driven by taste, i.e. the amount shared by the first player was not systematically influenced by the ethnicity of the second player. There are several papers that have followed similar designs to Fershtman and Gneezy (2001) to investigate discrimination in other segmented societies around the world. The results vary considerably. For example, while Chuah et al. (2013) report taste-based discrimination across Hindus and Muslims in India, Burns (2012) find discrimination against Blacks in South Africa on the basis of mistaken expectations (statistical discrimination). However, these variations are not surprising as the underlying social preferences being measured in these games have been shown to vary across cultures (Falk et al., 2018). Nevertheless, studies have tried to investigate why such variations arise. For example, Gupta et al. (2018) shows the important role status (majority vs minority) plays in determining inter-religious interactions in Bangladesh and India.

Another set of laboratory experiments take a different approach. In these studies, participants are allocated to different roles (Worker or Manager) and asked to undertake different decisions mimicking real-world scenarios. Workers perform specific tasks, the Managers observe their performance and status and make an evaluation about hiring (or giving bonus) to them. For example, Reuben et al. (2014) conduct a labour market experiment, where Workers undertake a mathematical task. Across different conditions, Employers make hiring decisions with or without knowing how Workers perform in these tasks. When performance is not known, male workers are more likely to be hired. The provision of performance

information (which on average is the same across gender) reduces the gender hiring gap. These results identify the different channels (beliefs about productivity at the population level vs the informativeness of the signal) through which statistical discrimination impacts hiring decisions in a controlled environment.

The *third* approach to study discrimination is through field experiments. The most prevalent form is audit and correspondence studies. Relative to laboratory-based experiments, field-based studies provide higher external validity by trading off some of the control over the decision-making environment (Harrison and List, 2004, Roe and Just, 2009). In a typical audit study, fake job candidates (of different statuses) participate in job interviews. These fake candidates often have similar CVs and go through extensive training to control for candidate-level heterogeneity beyond their different statuses. The outcome variable in typical audit studies is job offers made. In contrast, in correspondence studies, the fake candidate only applies for jobs and does not actually interact with potential employers. In other words, the fake candidate only exists on paper. The outcome variable is call-back rates. Early audit and correspondence studies were interested in establishing the presence of discrimination in the labor market. In a meta-analysis conducted with sixty-seven field experiments on discrimination since 2000, Rich (2014) report significant discrimination against (ethnic) minority individuals across different markets in countries around the world. Later studies have investigated the nature of the discrimination observed. Typically these studies (a) vary the information content of CVs and see whether differences in call-back rates reduce across groups when superior information is provided or (b) observe discrimination rates across occupations or cohorts of minorities. For the first group of studies, the contention is that if the discrimination is eliminated then it is entirely driven by statistical discrimination. Rich (2014) finds that the evidence is mixed: some studies find that the provision of superior information eliminated discrimination (Petit, 2007, Kaas and Manger, 2012), while other studies did not provide compelling evidence for statistical discrimination (Lahey, 2008, Albert et al., 2011). Similarly, studies that look at the nature of discrimination across cohorts and occupations suggest the importance of statistical differences to be limited. For example, Bertrand and Mullainathan (2004) find that the extent of discrimination against ethnic minorities was the same across occupations, while Carlsson and Rooth (2012) show that discrimination against minorities was higher in areas of Sweden that held more prejudicial views about immigrants.

The use of field experiments is increasingly receiving traction in the context of developing countries. For example, Hanna and Linden (2012) conducts a audit study where teachers are recruited to evaluate the performance of high and low-caste students. Another notable example is Rao (2019), who uses a team-selection experiment to measure discrimination. The experiment creates a trade-off for wealthy students in India between choosing a high-ability teammate (and thus increasing their own expected payoff) or choosing a lower-ability teammate with whom they would prefer to socialize. The task used is a relay race, where ability is observed and measured from individual sprints. Teammates ran a relay race together and also were required to spend time socializing with each other. Rao (2019) finds significant discrimination against the poor on average (the rich and the poor are identified

with their school uniforms). Specifically, poor students competing against rich students were systematically discriminated against. However, the extent of discrimination was lower as the costs of discriminating increased: in the experiment discrimination decreased as the stakes (monetary prize for winning the race) increased.

The final approach is to utilise natural or *quasi*-experiments to measure discrimination. The first set of studies utilises the natural (random) variation in interactions between individuals to infer the existence of discrimination. For example, (Price and Wolfers, 2010) shows the presence of discrimination when players are evaluated by referees of the opposite ethnicity in professional basketball. Another set of studies utilises changes in policies over time to infer the presence of discrimination. For example, Goldin and Rouse (2000) measures the impact of blind auditions in orchestras on discrimination against female musicians. In the same spirit, companies are increasingly moving towards using Artificial Intelligence (AI) in their recruitment process to reduce discrimination. In a recent paper, Avery et al. (2023) investigate whether using such tools helps reduce gender discrimination in a male-dominated technology sector. They find that indeed using such tools increases the gender diversity of the applicant pool as applicants believe that the process is less biased.

In summary, there are many ways in which researchers have measured discrimination. While the debate is still ongoing as to how to best measure discrimination, the focus of empirical literature has been on understanding channels through which discrimination can be reduced. We discuss two prominent streams of such research in Section 4. While the number of studies is very large, our focus will be on studies/experiments/interventions/natural experiments from developing countries.

4 Policies to reduce Discrimination

Mechanically, discrimination is likely to breed more discrimination because decision-making is typically concentrated among the majority. For example, male managers (typically most managers are male) are likely to recruit and promote other men – individuals are more likely to be comfortable engaging with and hence favour those that are similar to them. This can be used to explain the dearth of women on corporate boards or in higher echelons of academia.

Policies to tackle discrimination need to be multidimensional as the motivations behind them are quite varied. In this section, we will discuss two main policy prescriptions to address prejudice and discrimination. The first is increasing inter-group contact and the second is legislating diversity in leadership positions. It needs to be made clear that these are not mutually exclusive or exhaustive, but they do represent some of the ways in which policymakers have tried to address the problems of prejudice and discrimination.

4.1 Inter-group Contact

The idea of inter-group contact and how it can reduce discrimination is not new. The idea goes back to Allport (1954), who argued that under *appropriate* conditions (for example, equal footing, everyday/regular setting), interpersonal contact is one of the most effective ways to reduce discrimination. This is known as the *Contact Hypothesis*. The idea is that if members of different groups have the opportunity to communicate, each group will be better able to appreciate and understand the other. In consequence, prejudice should decrease. This literature has been extended to examine whether *pleasant and cooperative contact* (Hewstone and Browne, 1986) between antagonistic groups in a conflict setting can reduce prejudice and induce positive outcomes.

There is a vast and growing literature that tests whether the contact hypothesis (or its extensions) actually works. One of the earliest studies that tested the contact hypothesis were the railroad studies by Cook (1971, 1978). They simulated inter-racial workplace contact by hiring racially prejudiced young White adults to work on a railroad company management task with two co-workers: the first being a Black and the second being another White confederate. Participants were hired in a real part-time job and worked together for a month under optimal conditions of the contact hypothesis (equal status, shared goals, support of authorities and absence of competition). If the contact hypothesis is valid, this increased interaction between the two groups should lead to reduced prejudice. Indeed, at the end of the study, the two White participants rated their Black co-worker highly in terms of attractiveness, likeability and competence and these views persisted over time. Several months later these participants expressed less racial prejudice than corresponding controls when responding to an unrelated questionnaire on race relations and race-relevant social policies.

However, studies on the contact hypothesis have traditionally been restricted to the psychology literature and have often relied on using observational data. This makes the findings hard to interpret: it is often difficult to causally identify the effect of inter-group contact on discrimination and prejudice as the results potentially reflect reverse causation or selection because less-prejudiced individuals are more likely to engage in contact with the other group in the first place (Paluck et al., 2019). Accordingly, social scientists and economists have started using experimental methods to test the contact hypothesis. In this setup, groups are randomly formed and this implies that the results are not contaminated by selection in particular. Not surprisingly, this literature has found a fertile ground in developing countries, where conflicts of different kinds (ethnicity, caste, religion) are highly prevalent. We can subdivide this literature into those that use field experiments or interventions; and those that utilize natural experiments or policy changes at a macro level. In the rest of this section, we will briefly summarize some of this literature. Our focus will be on evidence from developing countries.

Evidence from Field Experiments/Interventions

One of the best-known examples of the use of the contact hypothesis is to examine outcomes relating to discrimination and prejudice to the random matching of roommates in college dormitories. While [Sacerdote \(2001\)](#) was possibly the first to utilize this random assignment of roommates (here the outcome variable was the study of peer effects on test scores), this technique has recently been extensively used to examine the evolution of inter-group relationships. [Boisjoly et al. \(2006\)](#) utilized the random assignment of roommates at Harvard University to examine the impact of contact on attitudes towards affirmative action policies. They find that White students who were randomly assigned African-American roommates were significantly more likely to support and endorse affirmative action. The effects also persist over time, with White students who were assigned roommates from any minority group being more likely to continue engaging socially with members of other ethnic groups after one year. [Carrell et al. \(2019\)](#) use similar data from the United States Air Force Academy (USFA). They exploit the fact that freshman students to the academy are randomly assigned to peer groups (or squadrons), with whom they live, eat, and train. They then examine whether exposure to African American peers within squadrons affects how White students behave and interact with African-American students subsequently. [Carrell et al. \(2019\)](#) find that increased cross-race interaction can lead to meaningful changes in behavior toward minorities: exposure to higher aptitude African-American peers, as well as a larger number of African-American peers, during freshman year, leads to subsequent changes in behaviour toward African-Americans.

[Corno et al. \(2022\)](#) present one of the earliest field tests of the contact hypothesis in a developing country. Their context is South Africa, where, the history of apartheid ensures that conflict between Whites and Blacks persists even today (even though apartheid officially ended in 1993 with the election of Nelson Mandela as the first Black President of South Africa). Black South Africans are also relatively prone to marginalization and stereotyping. The authors take advantage of a policy introduced and implemented by the University of Cape Town (UCT), one of the largest, oldest and most prestigious universities in the country, with the aim of promoting racial integration. The policy is the same as in [Boisjoly et al. \(2006\)](#): there is a random allocation of students across university residences and in some of the residences to roommates. This policy then provides a unique opportunity to estimate the causal effect of the contact hypothesis on attitudes and behaviour. Specifically, [Corno et al. \(2022\)](#) examine whether interaction with someone of a different race can change individual stereotypes towards that race. They administered a set of Implicit Association Tests (IATs) to freshmen living in double rooms. This provides an objective measure of prejudice and stereotype. They find that exposure to members of a different race led to statistically significant changes in stereotypes: Whites became relatively less prejudiced against Blacks.

Recall that the *contact hypothesis* argues that the effects of inter-group contact on prejudice should depend on the type of contact. The integration will reduce discrimination and prejudice only if specific conditions are met. An emerging literature suggests that naturally occurring integration often has negative effects (for example the negative political effects of immigrants and refugees in many parts of Europe or the virulence against the *bohiragotos*

or outsiders in national, state and local elections in India). So the question that arises is: how important is the type of contact? To answer this question, [Lowe \(2021\)](#) conducts an innovative field experiment. Specifically, he uses cricket (the most popular sport in India) to integrate men of different castes to form teams. The experiment is designed in a way that enables a distinction between collaborative and adversarial contacts. Out of a sample of 1261 men, 800 were randomized to play in 8 one-month long cricket leagues; the rest were controls. Of those assigned to play, 35 percent were assigned to homogeneous caste teams and the remaining to mixed caste teams. The ultimate aim was to win: so those in the same team had to play collaboratively: integrated groups with common goals – in this case winning – have incentives to cooperate with one another. The relationship with the opposing teams was adversarial: groups with competing goals have incentives to undermine each other. Behavioural outcomes were measured 1–3 weeks after each league ended. The results show that the two types of contacts have opposite effects, particularly on self-reported cross-caste friendships: having all other-caste teammates (as opposed to being in a homogeneous caste team) increases the number of other-caste friends by 1; on the other hand having all other-caste opponents instead of none reduces the number of other-caste friends by 3.5. The results are consistent with the argument that contacts are most effective when they are *pleasant and cooperative* in nature. Collaborative contact also resulted in a large and statistically significant reduction in own caste favouritism measured through a voting exercise where each team member voted to determine which of their teammates could receive professional coaching. While contact only improves inter-group relations when the groups have common goals, the intervention resulted in a reduction in overall inter-group differences. This suggests that in such settings, the positive aspects of inter-group contact can generally offset any negative aspects.

[Mousa \(2019\)](#) conduct a similar experiment to investigate whether inter-group contact, through sports, can build social cohesion in a conflict environment. She randomly assigned Iraqi Christian football (soccer) players who had been displaced by the ISIS to either homogeneous Christian teams or mixed teams, consisting of Christian and Muslim players. She finds evidence of significant changes in the behaviour of Christian players who were assigned to mixed teams: they were significantly more likely to vote for a Muslim player to receive a sportsman award, sign up to play for mixed teams and train with Muslim players six months into the future. There is also a partial change in attitudes, with Christian players in mixed teams being more likely to believe that co-existence is possible. However, the contact did not change the overall attitude of Christians towards Muslims, as captured by their attitude towards Muslim strangers. There was no change in prejudice.

The work by [Lowe \(2021\)](#) and [Mousa \(2019\)](#) are part of the growing literature that argues that sports can be vital in promoting interaction between different groups in society (see, for example [Putnam, 2000](#)). Sport is a close-contact social activity and can therefore be transformative in nature. Contact in the sports field can be an effective way of overcoming differences and decreasing discrimination between multiple groups ([Amir, 1969](#)). Researchers are increasingly pointing to the potential for sports to build relationships and social cohesion across religious, ethnic and social dimensions and this transformative nature

of sports is increasingly recognized at a policy level: for example, sports can provide an important pathway to integration for young people with refugee backgrounds in resettlement countries (particularly in Europe).

The notion of how *pleasant and cooperative* contact can causally reduce prejudice between different groups has recently been implemented in different settings. [Jha and Shayo \(2019\)](#) conduct a field experiment to examine whether even in situations of conflict contact (in this case participation in financial markets) can result in individuals re-evaluating the costs of conflict and changing their political attitudes (even their votes). Prior to the 2015 Israeli elections, the authors randomly assigned Palestinian and Israeli financial assets to likely voters, and incentivized them to actively trade for up to seven weeks. No political messages or non-financial information were included. The results show that the intervention caused a shift in vote choices towards parties that were more supportive of the peace process. The authors argue that this happens because the intervention reduces opposition to concessions for peace, and increases awareness of the broader economic risks of conflict.

More recently a set of papers have examined whether *pleasant and cooperative* contact through participation in vocational education programs can improve the relationship between groups in environments that are characterized by a long history of conflicts. If well designed, many vocational education programs can provide opportunities for *pleasant and cooperative* contact: they can be designed to be hands-on in nature with a major focus on group activities and discussion in focus groups with participants from other groups. [Scacco and Warren \(2018\)](#) examine whether inter-group contact between Christian and Muslim men participating in an urban youth vocational training program reduces prejudice toward the other religious group in Nigeria. The intervention brought together a random sample of Christian and Muslim young men from disadvantaged neighborhoods in Kaduna, Nigeria, a city that has experienced repeated episodes of severe communal violence, for sixteen weeks of computer training. The experimental design randomized recruitment into the computer training program; assignment to a religiously homogeneous or heterogeneous classroom and assignment to a co-religious or non co-religious learning partner within the classroom. Prejudice is measured through survey-based assessments of agreement with negative and positive stereotypes, and discrimination was measured through experimental tasks that the participants engaged in after the intervention was complete. They find that prejudice is resistant to change, i.e., attitudes are entrenched, but contact can change behaviour even in such an environment. After the end of the training course, subjects assigned to heterogeneous classes discriminated significantly less against out-group members than subjects assigned to homogeneous classes. The result that prejudice is resistant to change is consistent with the result obtained by [Mousa \(2019\)](#).

[Maiti et al. \(2022\)](#) conduct a similar experiment where they randomly assign Hindu and Muslim participants into groups, in which they interact over the course of a beautician training program. The experiment is conducted in North India, where there is a long history of conflict between the majority Hindu and the minority Muslims. Inter-religious attitudes are particularly important in this experimental setting because they have a history of being

manipulated to fuel violence. Communal violence in contemporary India is a genuine threat and against the backdrop of the threat of violence, a common method to increase physical security is to live and interact among one's own religious group. They find that inter-group contact reduces the prejudice of both Hindu and Muslim participants toward members of the other religion one week after the training program concluded and the positive effects on Hindu attitudes toward Muslims are found to persist even after 6-months of program completion.

In another study, [Gu et al. \(2019\)](#) conduct three framed field experiments to investigate whether pleasant and cooperative contact is effective in reducing prejudice between native Malawian and Chinese migrant entrepreneurs. For context, large inflows of Chinese migrants in recent decades in Malawi led to increased tensions with local and migrant businesses. They find that the effects of contact are asymmetric: while there is no change in the attitude of Malawian shopkeepers toward their Chinese counterparts, contact facilitates a positive attitudinal change in Chinese shopkeepers toward their Malawian counterparts and these effects persist for at least 10 days.

While all of these papers support the basic tenet of the contact hypothesis (increased contact between the different groups will under specific circumstances result in a reduction of prejudice and discrimination against the other-group), what about the intensity of contact? The question of intensity of contact is examined by [Ghosh \(2022\)](#), who implements a field experiment in India to understand the effects of religious diversity on productivity and attitudes and whether this depends on a firm's production technology. Hindu and Muslim workers in a manufacturing plant in West Bengal, India are randomly assigned to religiously mixed or homogeneous teams. The production tasks they are engaged in vary in intensity: they are categorized as high- or low-dependency based on the degree of continuous coordination required for production. [Ghosh \(2022\)](#) finds that mixed teams are less productive than homogeneous teams in high-dependency tasks, but this difference vanishes completely after four months. On the other hand, diversity has no effect on productivity in low-dependency tasks. Mixing improves the attitude of the majority (Hindu) workers towards minority (Muslims) when engaged in high-dependency tasks; with little or no corresponding effects in low-dependency tasks. However, the burden of integration is on the minority (Muslim) workers who need to initiate and pay for the cost of integration in high-intensity tasks, which ultimately leads to improvements in production and attitudes. The intensity of interactions, therefore, has an important role to play in individuals from different groups learning to work together and overcoming existing inter-group differences.

Evidence from Natural Experiments/Policy Changes at a Macro level

While thus far, we have discussed how interventions (field experiments) that randomize contact find positive effects of increased contacts on out-group discrimination (increased contact leads to reduced out-group discrimination), researchers have also used policy changes at a macro level (natural experiments) that lead to increased contact between groups to investigate whether the contact hypothesis works. We provide three examples from devel-

oping countries.

[Barnhardt \(2009\)](#) exploits a natural experiment in the Indian city of Hyderabad where Hindus and Muslims were randomly assigned units (apartments) in a public housing complex with physically distinct clusters (floors) consisting of four units. This was done to re-settle slum dwellers when a large slum on the south side of the city of Hyderabad in the state of Andhra Pradesh was the scene of a serious fire in February 2005, destroying nearly half of the houses. Against a backdrop of the threat of communal violence, a common method to increase physical security (particularly for the minority, in this case, Muslims) is to reside in religious enclaves, resulting in limited contact between the two groups. This policy of re-settlement and randomly assigning units forced contact between Hindus and Muslims. The policy thus created exogenous variation in the extent of religious diversity across clusters in the same complex. Using a survey that collected information on implicit and explicit attitudes about members of the other religion and willingness to co-exist, [Barnhardt \(2009\)](#) finds that location (proximity) indeed influences interactions, which in turn affects attitudes. Greater exposure to Muslims improves the attitudes of Hindus towards and their willingness to co-exist with Muslims. Overall there is a convergence of attitudes across religious groups.

[Rao \(2019\)](#) exploits a natural policy experiment in India to examine the effect of increased contact on behavioural outcomes of students from different socio-economic groups. His focus is not on caste or religion but on socioeconomic status. He utilizes the staggered timing of the adoption of a policy change in India, which required elite private schools to offer free positions to students from low socio-economic households, who would otherwise not have been able to attend these schools. In most schools, cohorts starting in 2007 have many poor students while older cohorts consist of purely rich students. Some schools implemented the policy late and some were exempt due to historical/legacy reasons. This allows [Rao \(2019\)](#) to use a Difference-in-Difference estimation strategy by comparing both within school (treated vs untreated cohorts) and within cohort (treated vs untreated schools) to estimate the average effect on wealthy students of having poor classmates. He also utilizes a second exogenous variation in the composition of peer groups within the classroom since some schools assign students to study or work groups based on alphabetical order (of first name). This leads to plausible exogenous variation in personal interactions between wealthy and poor students. [Rao \(2019\)](#) finds that having poor classmates makes wealthy students more pro-social, more generous, more egalitarian and they are significantly less likely to discriminate against poor students. The effects are driven largely by increased personal interaction between rich and poor students.

[Bazzi et al. \(2019\)](#) utilize Indonesia's Transmigration Program to understand how inter-group contact affects nation-building. During the period 1978–1988, the Indonesian government relocated two million voluntary migrants (hereafter, transmigrants) from the Inner Islands of Java and Bali to newly created agricultural villages in the Outer Islands of the country. The program offers *plausibly* exogenous variation in ethnic diversity and segregation. The authorities had limited (or little) ability to systematically assign transmigrants;

nor could transmigrants choose their destinations: the ethnic mix of Inner Islanders in the new villages was randomly determined by the position in the queue (in transit camps) and the timing of opening of settlements in the outer islanders. Additionally, upon arrival, all settlers received houses and farms by lottery. Full ownership rights were transferred after 5–10 years, but because land markets were imperfect, this effectively tied migrants to their initial settlement areas (and plots). [Bazzi et al. \(2019\)](#) find that even after three decades the Transmigration villages exhibit significantly greater ethnic diversity and less within-village segregation than other organically settled villages in the Outer Islands.

Taste-based or Statistical Discrimination?

While much of the economics literature on discrimination has focused on taste-based versus statistical discrimination (see Section 2), many of the papers that we survey in this section, do not delve into much detail on which aspect of discrimination is affected by the intervention. Exceptions are [Hanna and Linden \(2012\)](#), [Lowe \(2021\)](#), [Corno et al. \(2022\)](#).

[Hanna and Linden \(2012\)](#), in their experiment on caste-based discrimination by teachers against students (discussed earlier in this chapter, see Section 3), find that the teachers' behaviour are consistent with statistical discrimination. Teachers appear to discriminate more against children who are graded early in the evaluation process. This suggests that teachers are using demographic characteristics when the testing instrument or grade distribution are more uncertain. If the teachers were engaged in taste-based discrimination, there would be little reason to expect that discrimination would vary by the order in which they graded the exam.

[Lowe \(2021\)](#) finds that both cooperative (being on the same team) and adversarial (being opponents) reduces ability-based statistical discrimination: other-caste men are more likely to be chosen as teammates in future matches when prize money is at stake. This is because both types of contacts reveal information about the ability of other caste group members. On the other hand, it is not clear whether contact actually reduces prejudice. This is because when players are asked to choose teammates, for matches without involving prize money, the effects are smaller for both types of contact, though smaller for adversarial contact.

[Corno et al. \(2022\)](#) uses two different Implicit Association Tests (or IATs) – population IATs and academic IATs – to distinguish between taste-based and statistical discrimination. The population IATs involved pairing positive and negative attributes with racial categories (White and Black South Africans) and the authors argue that they embed more elements of taste. On the other hand academic IATs involved eliciting an association between race and academic ability. Specifically participants were asked to match pictures of individuals of different races with different percentiles of the grade distribution. The authors argue that academic IATs are more likely to provide statistical information on relative performance. They find that White students are more prejudiced against Blacks when the population IATs are concerned but there is no difference between the races in the academic IATs. In

their analysis, the driving mechanism is taste-based discrimination.

4.2 Diversity in Leadership Positions

One major consequence of discrimination against historically disadvantaged groups (defined as groups that have been systematically excluded from institutions and cultural practices that provide skills and resources) is that few leaders emerge from these marginalized groups. As discussed in Section 2, this has several important and long-term consequences. First, since decisions continue to be made by those in the majority, such discrimination will persist. This, in turn, might: (i) reinforce the beliefs of both the majority and the disadvantaged groups that the minority is incapable of making decisions since they have rarely observed decision-making by the minority; and (ii) the minority might “give-up”, believing that they can never succeed and never make a difference. There are strong moral and economic arguments suggesting that it is in the interest of society to improve the economic standing of such groups and that historical discrimination against a group should not be allowed to perpetuate itself and inhibit the groups’ right to well-being.

In addition to improving equity, policies that improve the economic standing and welfare of these groups can, by improving the talent allocation across different occupations, enhance efficiency (Holzer and Neumark, 2000). In democracies, the use of legislative policy to bring about such improvements remains contingent on legislator behaviour. Arguably, a significant barrier to the introduction of such policies is the political under-representation of individuals belonging to minority groups who might vote in their own interest. Cross-country evidence documents the fact that members of minority groups are less likely to get selected as candidates by parties, and are therefore underrepresented in the legislature (Rule and Zimmerman, 1994). The discrimination against the minority (and the underrepresented) is perpetuated.

The purpose of quotas for underrepresented groups (minorities, disadvantaged groups or women) is to increase the representation of these groups in the decision-making process. The idea is that having more diversity in the decision-making body and process could lead to the development of more targeted policies, those that will benefit the condition of these groups. It is expected that increased contact between different groups will, over time, result in reduced discrimination. In the rest of this section, we focus on the effect of affirmative action policies on *outcomes*. Ultimately, if leaders are able to deliver they effectively signal their ability and this, in turn, could lead to a reduction in statistical discrimination against specific groups as this information allows the majority to update their beliefs about their ability; even if prejudice against the minorities does not change, i.e., tastes remain unaffected.

4.2.1 Reservations for Women

Much of the literature on inclusive leadership has focused on addressing the gender gap in leadership positions. While the second half of the 20th century has witnessed large gains for women as leaders in government, the corporate sector and academia, women are still far from achieving parity with men in leadership positions. Having more women in leadership positions is considered not just a moral imperative, but yields several tangible benefits since women often make different policy decisions compared to men (Eagly et al., 1995, Eagly and Carli, 2003). For example, female leaders are more likely to prioritize spending on social protection programs, as well as on public services such as health and education (Lott and Kenny, 1999, Edlund and Pande, 2001, Chattopadhyay and Duflo, 2004, Clots-Figueras, 2011, 2012, Pande and Ford, 2012, Bhalotra and Clots-Figueras, 2014, De Paola et al., 2014, Braga and Scervini, 2017). Firms with a larger fraction of women on their boards put more emphasis on long-term rather than short-term considerations (Adams and Ferreira, 2009, Kenneth and Dittmar, 2012, Matsa and Miller, 2013). Several experimental studies suggest that the choices women make once in decision-making positions may be more socially oriented than those of men (Eckel and Grossman, 1998, Andreoni and Vesterlund, 2001, Gneezy et al., 2003) and consistent with this experimental evidence, it has been observed that appointing more women to leadership positions improves the quality of governance Dollar et al. (2001), Swamy et al. (2001), Gokcekus and Mukherjee (2002), World Bank (2002). Women are increasingly viewed as *political cleaners*, less corrupt than men and more likely to act as whistle-blowers when faced with unethical behaviour (Goetz, 2007, Brollo and Troiano, 2016). The rationale is that women are, on average, observed to be better stewards of resources and contribute more towards public goods than men; so women who hold public positions and leadership roles might also exhibit these traits, leading to improved social outcomes. There is, therefore, increasing urgency to ensure that more women are in leadership positions.

Policy response in this context has taken the form of affirmative action policies (or quotas), which require that a specific proportion of leadership positions be exclusively reserved for women (or for that matter any historically disadvantaged group).

Affirmative Action in Political Leadership Positions

Evidence from developed countries on the effect of gender quotas on female political engagement is mixed. On the one hand O'Brien and Rickne (2016), De Paola et al. (2010), Baltrunaite et al. (2014) document positive effects of gender quotas on women's selection into and survival in top political posts (in Sweden), increased women's representation in politics (in Italy) and improved the quality of politicians elected to office (again in Italy). On the other hand Bagues and Campa (2020), using data from Spain are unable to find any evidence of gender quotas leading to an increase in the probability that women reach powerful positions such as party leader or mayor and Lassébie (2020) find that the gender quota law in France resulted in an increase in the share of female candidates and the

share of female elected politicians, but find limited evidence of women being promoted to leadership positions. Women are more likely to resign and leave politics than men.

In developing countries, on the other hand, the majority of the evidence suggests that quotas for women in leadership positions have generally had positive effects. This literature has generally focused on the effects of the 73rd Constitutional Amendment in India (ratified in 1993). Through this amendment, the Indian government also introduced reservations for women: in each village council, at least one-third of the councillor positions were to be reserved for women (in these reserved seats or wards only women could stand as candidates, both men and women are eligible to vote). Across all GPs in the state, at least one-third of the head of the village council would be women. Similar patterns of reservation were introduced for sub-district and district-level councils (the other two levels of rural local government). Importantly seats to be reserved and village councils that were to have female heads were randomly determined before each election. Higher levels of the legislature (national parliament and state assemblies) were exempt from such a gender-based quota.

There is now a large literature that examines the impacts of the introduction of this affirmative action policy. Evidence suggests that female politicians (elected because of this affirmative action policy) have had beneficial effects on society as a whole. These include increased government spending on female-friendly investments ([Chattopadhyay and Duflo, 2004](#)), more reporting of crimes ([Iyer et al., 2012](#)), a reduction in the adolescent gender gap in aspirations ([Beaman et al., 2012](#)), improved voters' perceptions about female candidates ([Beaman et al., 2009](#)), increases in the number of new women-owned establishments ([Ghani et al., 2014](#)), increases in the number of births in government and private health facilities and lower child mortality rates ([Kumar and Prakash, 2017](#)) and an increase in the survival of higher order girls ([Kalsi, 2017](#)). The effects are, however, temporally non-monotonic. Both [Deininger et al. \(2015\)](#) and [Gangadharan et al. \(2016\)](#) find that in the short-term, residents of reserved constituencies have a more negative view of the quality of the head of the village council (or the Pradhan) and the value of public goods and that show that there may be an initial backlash by males in response to female Pradhans, however this attenuates over time. [O'Connell \(2018\)](#) and [Maitra and Rosenblum \(2022\)](#) show that there are also considerable upstream effects of this downstream quota on female political participation at the state and national level, where there are no quotas for women. See [Clots-Figueras and Iyer \(2023\)](#) in this Handbook for an extensive survey of Gender and Natural Experiments in Developing Countries.

Affirmative Action in the Corporate World

The presence of woman directors on corporate boards has been increasingly recognized as a good governance practice. The existing literature suggests, women have attributes, strengths and experiences, distinct from that of men, which add value to board deliberations and monitoring of management. In recent years, many countries have initiated affirmative action through the introduction of mandates to hire women (gender quotas) to ensure a minimum presence of woman directors on their company boards. One of the earliest

examples of affirmative action policies in the corporate sector was the quota law passed in Norway in December of 2005, which required that women make up a minimum of 40% of corporate boards. Corporations had 2 years to comply with the law. As a result of the law, the fraction of women directors went from 5% in 2001 to 40% in 2008. In 2011 France introduced a law that required corporate boards to be at least 20% female by January 2014. As of 2019, Germany, Netherlands and the US (California only) had introduced some form of mandate to ensure that boards are gender diverse. The only developing country that has a similar mandate in India. See [Catalyst](#). The evidence on the effects of women on boards on firm performance, however, remains mixed: effects have been estimated to be positive, negative or none.

Legislative efforts to address the persistently low presence of female directors on company boards in India first took root when a gender quota was mooted in the Draft Companies Bill, December 2011, followed by the enactment of the Companies Act, 2013, which came into effect from August 30, 2013, which mandated that the Board of Directors of every company or classes of companies, as may be prescribed, must have *at least one woman director*. Note that the Indian law specifies a minimum number of woman directors on board as opposed to a minimum percentage of women directors. With listed companies in India having an average board size of around nine, the quota of having a minimum of one woman director would lead women accounting for 11% of directors on a company board on average. While there were fears that this number quota would simply encourage firms to increase the size of the board (fears of tokenism), such fears have not eventuated. Indeed estimates available for Indian companies two years into the implementation of gender quotas show that the percentage of women on boards for the top 500 listed companies increased from the pre-quota level of 5% to around 13% in 2016.

4.2.2 Reservations for Other Disadvantaged Groups

Many countries have experimented with affirmative action which seeks to increase the representation of historically disadvantaged groups in both the economic and political process. It is important to recognize, however, that the historically disadvantaged groups need not necessarily belong to the minority. For example in South Africa affirmative action policies have been introduced to address the inequity and discrimination that had been the result of the codified apartheid policies (introduced in 1948) that were officially dismantled only in the 1990s. The policy of apartheid institutionalized racial discrimination, classifying the people of South Africa racially into either White, Colored, Asian or Indian, and Black (African). The rules specified where and how the different “races” could live, travel, work, be educated, get married, and mingle. The legacy of apartheid was the deep-rooted differential treatment of the “non-White” population of South Africa, resulting in imbalances and inequality in terms of the type of housing, employment opportunity, education, medical care, and other public services. It was expected that affirmative action would transform society following the apartheid governments and bring about equality and social justice for

all in South Africa. Interestingly, however, affirmative action policies in South Africa did not take the form of quotas (or reservations). The Broad-Based Black Economic Empowerment (BBBEE) program, which is now the key integration program to reconcile South Africans and address the inequalities, encourages businesses to integrate Black people in the workspace, support Black businesses, and give back to poor Black communities affected by land repossession. Businesses are awarded points that they can claim on a BBBEE certificate which entitles them to a greater chance of obtaining government contracts. Whether this policy has been successful or not, however remains an open question.

Malaysia's New Economic Policy, implemented in 1971 is another example of an affirmative action policy that provided intensified preferential treatment in education, employment, business, and asset ownership for the economically disadvantaged but majority population group in the country: the Malays. Ong (2012) shows that the policy (with a combination of the expansion of the number of schools and quotas for Malays in public institutions) increased the educational attainment of the Malays, which in turn led to an increase in the Malay-Chinese wage gap among tertiary graduates and increased the likelihood of Malay tertiary graduates being employed in the public sector, relative to their Chinese peers.

The Indian experiment on using public policy to end caste-based discrimination, and to improve the economic status of disadvantaged groups is one of the most radical. A centerpiece of this endeavor has been the implementation of the constitutional mandate which ensures the presence of legislators belonging to minority groups in state and national legislatures. Indeed reservations for disadvantaged caste groups both in the state and national legislatures were first introduced in 1951. Seats are reserved for candidates of disadvantaged castes and tribes (Scheduled Castes and Scheduled Tribes). Only members of the group that benefit from reservation can stand for election, but the entire electorate votes over the set of candidates. The effect of the mandate is to alter legislator identity without affecting voter identity. The reservation of seats for specific groups, increases the political representation afforded to these groups in the legislature. Political reservation has had a profound effect on the Indian political landscape. About 25% of all legislators in India, at both the national and state level, come from reserved seats. In 1993, as a part of the 73rd amendment to the Indian constitution, this reservation of seats on the basis of caste was extended to rural local elections. The caste-based reservations were orthogonal to the gender reservations (discussed above). Reservations were over time extended to *Other Backward Castes* (or OBCs), though as with gender reservations, reservations for the OBCs have not yet been extended to the national or state legislature.

Pande (2003) identifies strong positive effects on redistribution: she shows that political reservation in Indian states has increased redistribution of resources in favour of the groups that benefit from the political reservation. She argues that reservation can enhance a group's influence on policy-making and that legislators belonging to minority groups have used this influence to increase the incidence of targeted redistribution. However, whether such changes, away from general redistribution programs and toward targeted programs improve the well-being of either the minority groups or the polity at large is, however, an

open question. [Jensenius \(2015\)](#), on the other hand, shows that more than three decades of mandated reservations have had little or no detectable constituency-level effect on overall development or redistribution to the marginalized groups, neither on the literacy rates or employment patterns, nor on village amenities in reserved constituencies. [Chin and Prakash \(2011\)](#) find that the effects are quite heterogeneous: they show that an increase in the share of seats reserved for Scheduled Tribes significantly reduces poverty while an increase the share of seats reserved for Scheduled Castes has no impact on poverty.

Moving beyond the reservations at the state and national level legislatures, [Besley et al. \(2004\)](#) show that politician identity impacts on public good provision at the village level, but the patterns vary by type of public good. They find that for high spillover public goods residential proximity to elected representatives matters; however for low spillover public goods sharing the politician's group identity is what is relevant. This is consistent with the evidence provided by [Chattopadhyay and Duflo \(2004\)](#), who find that SC hamlets (in India within villages there is residential sorting by caste) received more public investment when the head of the village council (or the Pradhan) also belongs to a SC.

While much has been written about reservation in the political sphere in India, what is often not talked about is that the Constitution of India allows for similar reservations for the marginalized groups (SCs, STs and OBCs) in educational institutions, in government jobs. There is little that arouse more passion than the question of job reservations. There is a demand to extend reservation to persons who are not from the SC or the ST but who, nevertheless, belong to economic and socially backward groups – the OBC and there is a demand from the SC and the ST to extend reservation to private sector jobs.

The obvious question that arises, how effective has this policy been in improving the outcomes of the targeted groups. [Borooah et al. \(2007\)](#) using data from the 55th round of the National Sample Survey data find that the boost provided by the job reservation policies was around 5 percentage points. [Prakash \(2020\)](#) finds that a one-percentage point increase in the employment quota for Scheduled Castes increases the likelihood of obtaining a salaried job by 0.6-percentage points for male SC members residing in the rural sector, but has no employment impact for ST members and importantly there is no evidence of elite capture among the SCs.

What about the effect of quotas for members of the disadvantaged caste group in educational institutions? [Deshpande and Ramachandran \(2015\)](#) estimate the impact of the employment quota policy for the OBC and find that the policy increased the percentage of the OBCs obtaining public sector jobs and finishing secondary education by 2.6 and 4 percentage points, respectively. [Bertrand et al. \(2010\)](#) estimate the impact of the OBC quota in engineering colleges and find that lower caste group applicants benefit from attending these colleges (which they would not have been able to attend without the quota). Their results come with a caveat though: the benefit is greater for the marginal admitted applicant from a high caste group than it is for the marginal admitted applicant from a low caste group, implying that the policy of reserving seats for disadvantaged groups leads to inefficiencies

in the allocation of educational slots. [Bagde et al. \(2016\)](#) find that the education quota policy (that fixes percentage quotas, common across a large number of engineering colleges in India) increases college attendance for the targeted students, particularly at relatively higher quality institutions.

Of course, affirmative action policies targeting specific (historically disadvantaged) groups are not specific to developing countries. There is a large literature that estimates the effects of employment-related affirmative action policies on the outcomes of targeted groups in the US (see, for example [Holzer and Neumark, 2000](#)). In this context, one set of studies has focused on the U.S. federal contractor program. Under this program, targeted groups (including African Americans and women) are given preferential treatment when bidding for business from the federal government. The consensus in the literature is that the program has had somewhat limited impacts on the economic outcomes for the key targeted groups (the African-Americans).

Affirmative Action Policies and Lasting Changes in Attitudes

Have affirmative action programs created lasting change in prejudicial attitudes? Affirmative action policies and quotas are often thought of as being transitory, needed to improve the condition of particular disadvantaged groups until they can take care of themselves: the idea being that targeted groups could secure representation for themselves after affirmative action policies are dismantled. This is because the constraints that previously prevented them from being in leadership positions are (expected to be) mitigated, either by the policy or changes in the broad social environment resulting from the policy. That does not seem to be happening despite these policies being in place for considerable periods of time. The usefulness of affirmative action policies are thus increasingly being questioned.

What happens when affirmative action policies are removed? Evidence from developed countries generally suggests that quotas had a long-term effect and these effects persisted even after the quotas were withdrawn. [De Paola et al. \(2010\)](#) examine the case of gender quotas (reserving a percentage of the places in the party lists for Municipal Council elections for female candidates) that were introduced by law in Italy in 1993 and were in force until 1995. Because of the short period covered by the reform, some municipalities never voted under the gender quota regime. They find that women's representation in politics after the reform increased significantly more in municipalities that were affected by the reform than in municipalities that were not affected. In the context of the US, [Miller \(2017\)](#) examines the dynamic effects of federal affirmative action regulation. He exploits the variation in the timing of regulation and deregulation across work establishments and finds that affirmative action increases the African American share of employees over time: in 5 years after an establishment is first regulated, the African American share of employees increases by an average of 0.8 percentage points and this share continues to grow at a similar pace even after an establishment is deregulated. Finally [Antonovics and Sander \(2013\)](#) consider the effect if California's Proposition 209, which led to the 1998 ban on the use of racial preferences in admissions at the University of California (UC). They do not find any evidence that

the repeal of affirmative action disincentivized minorities from attending UC schools and if anything the repeal of affirmative action increased the signaling value of attending UC schools for minorities

Evidence from India in this domain is, however, mixed. In the context of gender-based quotas, [Bhavnani \(2009\)](#), using data from three municipal wards in the city of Mumbai and found that the probability of a woman winning office conditional on the constituency being reserved for women previously are approximately five times the probability of a woman winning office if the constituency had not been previously reserved for women. He argued that reservations work because of a supply-side effect i.e., they introduce into politics women who are able to win elections even after reservations are withdrawn. However, [Sekhon and Titiunik \(2012\)](#) obtains effects that are considerably smaller and [Maitra \(2021\)](#), using data from the universe of village councils in West Bengal, political parties exhibit reluctance to re-nominate women who had won in reserved seats, when the reservation is withdrawn. In the context of caste-based reservations, [Bhavnani \(2017\)](#) argue that electoral quotas for India's SCs fail to boost SCs' chances of winning office after they are discontinued. These results can be explained by beliefs about the ability of candidates and are consistent with statistical discrimination against specific groups be it women or disadvantaged caste groups.

The original idea behind the introduction of affirmative action policies and quotas was to change the norms relating to negative stereotypes and low investment in skill acquisition. While quotas have had many positive effects, whether or not they have been able to reduce discrimination and prejudice continues to remain an open question. There are several issues at play here. First, legislating diversity often creates tensions around fairness: beneficiaries of such policies can be viewed as people who are not deserving on the basis of merit. Second, it also encourages individuals to take on roles that are not in line with their preferences or against existing social norms. In both cases, the fear of or the actual backlash faced can pose severe limits on the dynamic effects of affirmative action policies in changing norms and investments in skill acquisition.

5 Conclusion

Since the middle of the 20th century, the reduction of discrimination has been a major policy focus. Often this required input from academics in the social sciences. Accordingly, researchers investigated the issue of discrimination by examining its roots and its impact over time. Interventions based on this research have been developed and implemented and have resulted in modest decreases in discrimination over time. Nevertheless, discrimination has not been eradicated. Disparities in outcomes achieved based on gender, ethnicity, disability status etc. persist. At the same time, it is also an exciting time to be a researcher in the area of discrimination: there is (i) appetite from stakeholders to address the issue (ii)

an evidence base of what the drivers of discrimination and what works and what does not (*iii*) access to large rich data sets (*iv*) technological advances that allow for new innovative solutions to be forged.

In the first part of the chapter, we presented how the economics literature has conceptualised and measured discrimination over the years. It is apparent from this account that the mechanics of why discrimination occurs and its equilibrium implications have been quite well established. In stark contrast to other disciplines, the economics literature shows that discrimination can occur even without agents holding prejudicial preferences. Nevertheless, the discrimination literature in economics is undergoing significant changes led by insights gained from behavioural economics: how do behavioural biases in decision-making, systematic differences in behavioural preferences and social norms impact the behaviour of various agents in the labour market. The second part of the chapter focused on providing an account of the performance of two sets of policies in reducing discrimination. The first set of policies, that promote inter-group contact, is rooted in the psychology literature. The contention is that making people interact with each other will allow them to appreciate and understand each other and subsequently reduce prejudice. Empirical evidence validates the basic claim of the hypothesis. It also shows that the impact of contact is higher when the nature of the contact is cooperative and pleasant in nature and the intensity of the contact is high. The second set of policies is based on affirmative action. The contention is that making minority individuals more visible in leadership positions has two broad effects: (*i*) members of minority groups can provide a different type of leadership leading to better outcomes and (*ii*) the inter-generational link between discrimination and human capital choices and negative stereotypes will be broken. Our review, which looks at affirmative action policies across many different domains, shows mixed results on both fronts.

So what should be our agenda going forward? *First*, we believe that economists should try to systematically incorporate evidence and accounts from other disciplines into economic models. To this end, research in behavioural economics has attempted to bridge this gap between economics and psychology and the sociology literature. This research has highlighted the limitations of the standard economic models of discrimination. It is also important to note that some of this research has filtered through the existing theoretical frameworks, e.g. impacts of heuristics and biases are now incorporated in statistical models of discrimination. Nevertheless, newer models are needed: see, for example, [Small and Pager \(2020\)](#). A deeper understanding of the roots of discrimination would also help us understand why certain interventions work and others do not.

Second, improvements in technology and the emergence of Artificial Intelligence (AI) have provided us with new tools to tackle discrimination. Companies are increasingly using these tools in their hiring process. The assumption is that the decision-making using AI tools would not be as susceptible to behavioural biases that often lead to discrimination. Indeed, recent research has shown that the usage of AI can indeed improve diversity in application pools. Nevertheless, the implementation of such tools can be problematic. Typically, the algorithms that are the basis of AI tools are often written by humans. To

avoid biases or perceptions of biases, often the tools incorporate objective criteria against which performance is evaluated. As discussed in the first part of the chapter, individuals from different groups may face different constraints in both human capital formation and signalling and the objective criteria set can create bias in the hiring process. Further, the evaluations made by AI can often act as strong anchor points, which decision-makers would be less willing to override even when they believe discrimination is being taken place.

Third, even if we successfully develop AI tools to be unbiased, discrimination may persist. If discrimination observed is based on social norms, addressing the issue of prejudice and bias will not eradicate it. For example, in an interesting paper, [Cook et al. \(2021\)](#) shows that even when pay decisions are made objectively (i.e., formula-based), choices made across gender can generate a gender pay gap. In fact, there is voluminous evidence in psychology and economics that shows that behaviour that is perceived as breaking social norms is often avoided due to fear of backlash.

Finally, we believe that the issue of discrimination can only be solved after taking into account the various constraints people face in markets. This will often require policies to be conducted in tandem. As discussed at the end of Section 4, the effectiveness of legislating diversity will depend on whether or not beneficiaries are perceived to be deserving. Similarly, policies on the other end of the spectrum, which put the onus on individuals to change their behaviour (to often break social norms) have also been shown to have limited impact. An interesting avenue for research is to develop interventions (or policies within organisations) that recognise and provide solutions for the different constraints by individuals belonging to different groups. For example, [Erkal et al. \(2022\)](#) show that in a leadership selection process, which is often norm-incongruent for females, changing the default option from opt-in to opt-out can greatly reduce the gender gap in leadership. [Mancuso et al. \(2023\)](#), show how the impact of social modesty norms can be mitigated when individuals communicate their achievements and good deeds to their peers.

References

- Adams, R. and D. Ferreira (2009). Women in the boardroom and their impact on governance and performance. *Journal of Financial Economics* 94(2), 291 – 309.
- Aigner, D. J. and G. G. Cain (1977). Statistical theories of discrimination in labor markets. *Ilr Review* 30(2), 175–187.
- Akerlof, G. A. and R. E. Kranton (2000). Economics and identity. *The quarterly journal of economics* 115(3), 715–753.
- Albert, R., L. Escot, and J. A. Fernández-Cornejo (2011). A field experiment to study sex and age discrimi-

- nation in the madrid labour market. *The International Journal of Human Resource Management* 22(02), 351–375.
- Allport, G. (1954). *The Nature of Prejudice*. Oxford: Addison-Wesley.
- Amir, Y. (1969). Contact Hypothesis in Ethnic Relations. *Psychological Bulletin* 71(5), 319–342.
- Andreoni, J. and L. Vesterlund (2001). Which is the Fair Sex? Gender Differences in Altruism. *Quarterly Journal of Economics* 116(1), 293 – 312.
- Antonovics, J. L. and R. H. Sander (2013). Affirmative Action Bans and the Chilling Effect. *American Law and Economics Review* 15(2), 252 – 299.
- Arrow, K. J., O. Ashenfelter, and A. Rees (1973). Discrimination in labor markets. *The Theory of Discrimination*, 3–33.
- Avery, M., A. Leibbrandt, and J. Vecchi (2023). Does Artificial Intelligence Help or Hurt Gender Diversity? Evidence from Two Field Experiments on Recruitment in Tech. Technical report, Monash University and Gothenburg University.
- Bagde, S., D. Epple, and L. Taylor (2016). Does Affirmative Action Work? Caste, Gender, College Quality, and Academic Success in India. *American Economic Review* 106(6), 1495 – 1521.
- Bagues, M. F. and P. Campa (2020). Can gender quotas in candidate lists empower women? evidence from a regression discontinuity design. Technical report, CEPR Discussion Paper No. DP12149.
- Baltrunaite, A., P. Bello, A. Casarico, and P. Profeta (2014). Gender quotas and the quality of politicians. *Journal of Public Economics* 118, 62 – 74.
- Barnhardt, S. (2009). Near and Dear? Evaluating the Impact of Neighbor Diversity on Inter-Religious Attitudes. Technical report, J-PAL.
- Bazzi, S., A. Gaduh, A. D. Rothenberg, and M. Wong (2019). Unity in Diversity? How Intergroup Contact can Foster Nation Building. *American Economic Review* 109(11), 3978–4025.
- Beaman, L., R. Chattopadhyay, E. Duflo, R. Pande, and P. Topalova (2009). Powerful women: Does exposure reduce bias? *Quarterly Journal of Economics* 124(4), 1497–1540.
- Beaman, L., E. Duflo, R. Pande, and P. Topalova (2012). Female leadership raises aspirations and educational attainment for girls: A policy experiment in India. *Science* 335(6068), 582 – 586.
- Becker, G. S. (1957). *The economics of discrimination*. University of Chicago press.
- Berg, J., J. Dickhaut, and K. McCabe (1995). Trust, reciprocity, and social history. *Games and economic behavior* 10(1), 122–142.
- Bertrand, M., R. Hanna, and S. Mullainathan (2010). Affirmative Action in Education: Evidence from Engineering College Admissions in India. *Journal of Public Economics* 94(1-2), 16 – 29.
- Bertrand, M. and S. Mullainathan (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.
- Besley, T., R. Pande, L. Rahman, and V. Rao (2004). The Politics of Public Good Provision: Evidence from Indian Local Governments. *Journal of European Economic Association* 2(2/3), 416 – 426.
- Bhalotra, S. and I. Clots-Figueras (2014). Health and the political agency of women. *American Economic Journal: Economic Policy* 6(2), 164–197.

- Bhavnani, R. (2009). Do electoral quotas work after they are withdrawn? Evidence from a natural experiment in India. *American Political Science Review* 103(1), 23 – 35.
- Bhavnani, R. (2017). Do the effects of temporary ethnic group quotas persist? evidence from india. *American Economic Journal: Applied Economics* 9(3), 105 – 123.
- Bian, L., S.-J. Leslie, and A. Cimpian (2017). Gender stereotypes about intellectual ability emerge early and influence children’s interests. *Science* 355(6323), 389–391.
- Black, D. A. (1995). Discrimination in an equilibrium search model. *Journal of Labor Economics* 13(2), 309–334.
- Blinder, A. S. (1973). Wage discrimination: reduced form and structural estimates. *Journal of Human Resources*, 436–455.
- Bohren, J. A., K. Haggag, A. Imas, and D. G. Pope (2023). Inaccurate statistical discrimination: An identification problem. *Review of Economics and Statistics*, 1–45.
- Boisjoly, J., G. J. Duncan, M. Kremer, D. M. Levy, and E. J. (2006). Empathy or antipathy? the impact of diversity. *American Economic Review* 96(5), 1890 – 1905.
- Bordalo, P., K. Coffman, N. Gennaioli, and A. Shleifer (2016). Stereotypes. *The Quarterly Journal of Economics* 131(4), 1753–1794.
- Borooh, V. K., A. Dubey, and S. Iyer (2007). The Effectiveness of Jobs Reservation: Caste, Religion and Economic Status in India. *Development and Change* 38(3), 423 – 445.
- Braga, M. and F. Scervini (2017). The performance of politicians: The effect of gender quotas. *European Journal of Political Economy* 46, 1 – 14.
- Brollo, F. and U. Troiano (2016). What happens when a woman wins an election? Evidence from close races in Brazil. *Journal of Development Economics* 122, 28–45.
- Burns, J. (2012). Race, diversity and pro-social behavior in a segmented society. *Journal of Economic Behavior & Organization* 81(2), 366–378.
- Burszty, L., T. Fujiwara, and A. Pallais (2017). ‘acting wife’: Marriage market incentives and labor market investments. *American Economic Review* 107(11), 3288–3319.
- Buser, T., M. Niederle, and H. Oosterbeek (2014). Gender, competitiveness, and career choices. *The quarterly journal of economics* 129(3), 1409–1447.
- Carlsson, M. and D.-O. Rooth (2012). Revealing taste-based discrimination in hiring: a correspondence testing experiment with geographic variation. *Applied Economics Letters* 19(18), 1861–1864.
- Carrell, S. E., M. Hoekstra, and J. E. West (2019). The Impact of College Diversity on Behavior toward Minorities. *American Economic Journal: Economic Policy* 11(4), 159–182.
- Chattopadhyay, R. and E. Duflo (2004). The impact of reservation in the Panchayati Raj: Evidence from a nationwide randomized experiment. *Economic and Political Weekly* 39(9), 979 – 986.
- Chauvin, K. (2018). A misattribution theory of discrimination. Technical report, working paper.
- Chen, Y. and S. X. Li (2009). Group identity and social preferences. *American Economic Review* 99(1), 431–57.

- Chin, A. and N. Prakash (2011). The redistributive effects of political reservation for minorities: Evidence from India. *Journal of Development Economics* 96(2), 265 – 277.
- Chuah, S.-H., R. Fahoum, and R. Hoffmann (2013). Fractionalization and trust in india: A field-experiment. *Economics Letters* 119(2), 191–194.
- Clots-Figueras, I. (2011). Women in politics: Evidence from the Indian States. *Journal of Public Economics* 95(7–8), 664 – 690.
- Clots-Figueras, I. (2012). Are female leaders good for education? Evidence from India. *American Economic Journal: Applied Economics* 4(1), 212 – 244.
- Clots-Figueras, I. and L. Iyer (2023). Gender and natural experiments in developing countries. In U. Dasgupta and P. Maitra (Eds.), *Handbook of Experimental Development Economics*. Edward Elgar.
- Coate, S. and G. C. Loury (1993). Will affirmative-action policies eliminate negative stereotypes? *The American Economic Review*, 1220–1240.
- Cook, C., R. Diamond, J. V. Hall, J. A. List, and P. Oyer (2021). The gender earnings gap in the gig economy: Evidence from over a million rideshare drivers. *The Review of Economic Studies* 88(5), 2210–2238.
- Cook, S. W. (1971). The effect of unintended interracial contact upon racial interaction and attitude change. Technical report, Proj. No. 5–1320, Final Rep. Washington, DC: U.S. Dep. Health Educ. Welfare.
- Cook, S. W. (1978). Interpersonal and attitudinal outcomes in cooperating interracial groups. *Journal of Research & Development in Education* 12(1), 97 – 113.
- Corno, L., E. La Ferrara, and J. Burns (2022). Interaction, Stereotypes, and Performance: Evidence from South Africa. *American Economic Review* 112(12), 3848–75.
- De Paola, M., V. Scoppa, and M. A. De Benedetto (2014). The impact of gender quotas on electoral participation: Evidence from italian municipalities. *European Journal of Political Economy* 35, 141 – 157.
- De Paola, M., V. Scoppa, and R. Lombardo (2010). Can gender quotas break down negative stereotypes? evidence from changes inelectoral rules. *Journal of Public Economics* 94, 344 – 353.
- Deininger, K., S. Jin, H. K. Nagarajan, and F. Xia (2015). Does female reservation affect long-term political outcomes? evidence from rural india. *The Journal of Development Studies* 51(1), 32–49.
- Deshpande, A. and R. Ramachandran (2015). Affirmative Action, Political Representation and Caste Disadvantage: Mapping changes in post-Mandal India. Technical report, Delhi School of Economics.
- Dollar, D., R. Fisman, and R. Gatti (2001). Are women really the “fairer” sex? Corruption and women in government. *Journal of Economic Behavior & Organization* 46(4), 423 – 429.
- Eagly, A. and L. Carli (2003). The female leadership advantage: An evaluation of the evidence. *The Leadership Quarterly* 14, 807 – 834.
- Eagly, A., S. Karau, and M. Makhijani (1995). Gender and the effectiveness of leaders: A meta-analysis. *Psychological Bulletin* 117(1), 125 – 145.
- Eckel, C. and P. Grossman (1998). Are women less selfish than men? evidence from dictator experiments. *Economic Journal* 108(448), 726 – 735.

- Edlund, L. and R. Pande (2001). Why have women become left-wing? The political gender gap and the decline of marriage. *Quarterly Journal of Economics* 117(3), 917 – 961.
- Erkal, N., L. Gangadharan, and E. Xiao (2022). Leadership selection: Can changing the default break the glass ceiling? *The Leadership Quarterly* 33(2), 101563.
- Falk, A., A. Becker, T. Dohmen, B. Enke, D. Huffman, and U. Sunde (2018). Global evidence on economic preferences. *The Quarterly Journal of Economics* 133(4), 1645–1692.
- Fershtman, C. and U. Gneezy (2001). Discrimination in a segmented society: An experimental approach. *The Quarterly Journal of Economics* 116(1), 351–377.
- Forsythe, R., J. L. Horowitz, N. E. Savin, and M. Sefton (1994). Fairness in simple bargaining experiments. *Games and Economic behavior* 6(3), 347–369.
- Gangadharan, L., T. Jain, P. Maitra, and J. Vecci (2016). Social identity and governance: The behavioral response to female leaders. *European Economic Review* 90, 302 – 325.
- Ghani, E., W. R. Kerr, and S. D. O’Connell (2014). Political reservations and women’s entrepreneurship in india. *Journal of Development Economics* 108, 138 – 153.
- Ghosh, A. (2022). Religious divisions and production technology: Experimental evidence from india. Technical report, BRIQ - Institute on Behavior and Inequality.
- Gneezy, U., M. Niederle, and A. Rustichini (2003). Performance in Competitive Environments: Gender Differences. *Quarterly Journal of Economics* 118(3), 1049 – 1074.
- Goetz, A. (2007). Political cleaners: Women as the new anti-corruption force? *Development and Change* 38(1), 87 – 105.
- Gokcekus, O. and R. Mukherjee (2002). Public sector corruption and gender: Perception of public officials from six developing and transition countries. Technical report, World Bank.
- Goldin, C. and C. Rouse (2000). Orchestrating impartiality: The impact of” blind” auditions on female musicians. *American economic review* 90(4), 715–741.
- Gu, J., A. Mueller, I. Nielsen, J. Shachat, and R. Smyth (2019). Improving Intergroup Relations Through Actual and Imagined Contact: Field Experiments with Malawian Shopkeepers and Chinese Migrants. *Economic Development and Cultural Change* 68(1), 273–303.
- Gupta, G., M. Mahmud, P. Maitra, S. Mitra, and A. Neelim (2018). Religion, minority status, and trust: Evidence from a field experiment. *Journal of Economic Behavior & Organization* 146, 180–205.
- Hanna, R. N. and L. L. Linden (2012). Discrimination in Grading. *American Economic Journal: Economic Policy* 4(4), 146–168.
- Harrison, G. W. (2014). Real choices and hypothetical choices. In *Handbook of Choice Modelling*, pp. 236–254. Edward Elgar Publishing.
- Harrison, G. W. and J. A. List (2004). Field experiments. *Journal of Economic literature* 42(4), 1009–1055.
- Hewstone, M. and R. Browne (1986). Contact is not enough. An intergroup perspective on the ‘contact hypothesis’. In M. Hewstone and R. Browne (Eds.), *Contact and Conflict in Intergroup Encounters*, pp. 1 – 44. New York: Blackwell.
- Hoff, K. and P. Pandey (2006). Discrimination, social identity, and durable inequalities. *American economic review* 96(2), 206–211.

- Holzer, H. and D. Neumark (2000). Assessing Affirmative Action. *Journal of Economic Literature* 38(3), 483 – 568.
- Iyer, L., A. Mani, P. Mishra, and P. Topalova (2012). The power of political voice: women’s political representation and crime in india. *American Economic Journal: Applied Economics* 4(4), 165–193.
- Jensenius, F. R. (2015). Development from representation? a study of quotas for scheduled castes in india. *American Economic Journal: Applied Economics* 7, 196 – 220.
- Jha, S. and M. Shayo (2019). Valuing peace: the effects of financial market exposure on votes and political attitudes. *Econometrica* 87(5), 1561 – 1588.
- Kaas, L. and C. Manger (2012). Ethnic discrimination in germany’s labour market: A field experiment. *German economic review* 13(1), 1–20.
- Kahneman, D. and A. Tversky (1972). Subjective probability: A judgment of representativeness. *Cognitive psychology* 3(3), 430–454.
- Kalsi, P. (2017). Seeing is believing- can increasing the number of female leaders reduce sex selection in rural india? *Journal of Development Economics* 126, 1–18.
- Kenneth, R. A. and A. K. Dittmar (2012). The changing of the boards: The impact on firm valuation of mandated female board representation. *Quarterly Journal of Economics* 127(1), 137 – 197.
- Krupka, E. L. and R. A. Weber (2013). Identifying social norms using coordination games: Why does dictator game sharing vary? *Journal of the European Economic Association* 11(3), 495–524.
- Kumar, S. and N. Prakash (2017). Effect of political decentralization and female leadership on institutional births and child mortality in rural bihar, india. *Social Science & Medicine* 185, 171–178.
- Lahey, J. N. (2008). Age, women, and hiring an experimental study. *Journal of Human resources* 43(1), 30–56.
- Lang, K. and A. Kahn-Lang Spitzer (2020). Race discrimination: An economic perspective. *Journal of Economic Perspectives* 34(2), 68–89.
- Lassébie, J. (2020). Gender quotas and the selection of local politicians: Evidence from French municipal elections. *European Journal of Political Economy* 62.
- Lott, J. and L. Kenny (1999). Did women’s suffrage change the size and scope of government? *Journal of Political Economy* 107(6), 1163 – 1198.
- Lowe, M. (2021). Types of contact: A field experiment on collaborative and adversarial caste integration. *American Economic Review* 111(6), 1807 – 1844.
- Maiti, S. N., D. Pakrashi, S. Saha, and R. Smyth (2022). Don’t Judge a Book by its Cover: The Role of Intergroup Contact in Reducing Prejudice in Conflict Settings. *Journal of Economic Behavior & Organization* 202, 533–548.
- Maitra, P. (2021). Affirmative Action and Female Representation in Politics: Evidence from Local Politics in India. Technical report, Monash University.
- Maitra, P., A. Neelim, and C. Tran (2021). The role of risk and negotiation in explaining the gender wage gap. *Journal of Economic Behavior & Organization* 191, 1–27.
- Maitra, P. and D. Rosenblum (2022). Upstream effects of female political reservations. *European Journal of Political Economy* 102061.

- Mancuso, J., A. Neelim, and J. Vecci (2023). Gender Differences in Self-promotion: Understanding the Female Modesty Constraint. Technical report, LaTrobe University, University of Tasmania and Gothenburg University.
- Matsa, D. and A. Miller (2013). A female style in corporate leadership? Evidence from quotas. *American Economic Journal: Applied Economics* 5(3), 136 – 169.
- Miller, C. (2017). The Persistent Effect of Temporary Affirmative Action. *American Economic Journal: Applied Economics* 9(3), 152 – 190.
- Mousa, S. (2019). Creating Coexistence: Intergroup Contact and Soccer in Post-ISIS Iraq. Technical report, Stanford University.
- Neumark, D. (2018). Experimental research on labor market discrimination. *Journal of Economic Literature* 56(3), 799–866.
- Oaxaca, R. (1973). Male-female wage differentials in urban labor markets. *International economic review*, 693–709.
- O’Brien, D. and J. Rickne (2016). Gender quotas and women’s political leadership. *American Political Science Review* 110(1), 112 – 126.
- O’Connell, S. D. (2018). Can quotas increase the supply of candidates for higher-level positions? evidence from local government in India. *Review of Economics and Statistics Forthcoming*.
- Ong, E. (2012). *Education and Labor Market Outcomes of Malaysia’s Affirmative Action Policies*. Ph. D. thesis, Department of Economics, Brown University.
- Onuchic, P. (2022). Recent contributions to theories of discrimination. *arXiv preprint arXiv:2205.05994*.
- Paluck, E., S. Green, and D. Green (2019). The contact hypothesis re-evaluated. *Behavioural Public Policy* 3(2), 129 – 158.
- Pande, R. (2003). Can Mandated Political Representation Increase Policy Influence for Disadvantaged Minorities? Theory and Evidence from India. *American Economic Review* 93(4), 1132 – 1151.
- Pande, R. and D. Ford (2012). Gender quotas and female leadership. World Development Report: Gender Equality and Development.
- Petit, P. (2007). The effects of age and family constraints on gender hiring discrimination: A field experiment in the french financial sector. *Labour Economics* 14(3), 371–391.
- Phelps, E. S. (1972). The statistical theory of racism and sexism. *The american economic review* 62(4), 659–661.
- Prakash, N. (2020). The Impact of Employment Quotas on the Economic Lives of Disadvantaged Minorities in India. Technical report, IZA Discussion Papers, No. 13847, Institute of Labor Economics (IZA), Bonn.
- Price, J. and J. Wolfers (2010). Racial discrimination among nba referees. *The Quarterly journal of economics* 125(4), 1859–1887.
- Putnam, R. (2000). *Bowling Alone: The Collapse and Revival of American Community*. New York: Simon and Schuster.
- Rao, G. (2019). Familiarity does not Breed Contempt: Generosity, Discrimination, and Diversity in Delhi Schools. *American Economic Review* 109(3), 774 – 809.

- Reuben, E., P. Sapienza, and L. Zingales (2014). How stereotypes impair women's careers in science. *Proceedings of the National Academy of Sciences* 111(12), 4403–4408.
- Rich, J. (2014). What do field experiments of discrimination in markets tell us? a meta analysis of studies conducted since 2000.
- Roe, B. E. and D. R. Just (2009). Internal and external validity in economics research: Tradeoffs between experiments, field experiments, natural experiments, and field data. *American Journal of Agricultural Economics* 91(5), 1266–1271.
- Rule, W. and J. Zimmerman (1994). *Electoral Systems in Comparative Perspective: Their Impact on Women and Minorities*. Westport, CT: Greenwood Press.
- Sacerdote, B. (2001). Peer Effects with Random Assignment: Results for Dartmouth Roommates. *The Quarterly Journal of Economics* 116(2), 681–704.
- Scacco, A. and S. S. Warren (2018). Can Social Contact Reduce Prejudice and Discrimination? Evidence from a Field Experiment in Nigeria. *American Political Science Review* 112(3), 654–677.
- Sekhon, J. S. and R. Titiunik (2012). When Natural Experiments are Neither Natural nor Experiments. *American Political Science Review* 106(1), 35–57.
- Small, M. L. and D. Pager (2020). Sociological perspectives on racial discrimination. *Journal of Economic Perspectives* 34(2), 49–67.
- Spiegler, R. (2020). Behavioral implications of causal misperceptions. *Annual Review of Economics* 12, 81–106.
- Swamy, A., Y. Knack, Y. Lee, and O. Azfar (2001). Gender and corruption. *Journal of Development Economics* 64(1), 25 – 55.
- Tajfel, H. and J. C. Turner (1978). Intergroup behavior. *Introducing social psychology*, 401–466.
- World Bank (2002). Integrating gender into the World Banks work: A strategy for action. Technical report, World Bank.
- Young, H. P. (2015). The evolution of social norms. *Annual Review of Economics* 7(1), 359–387.