

Diversity and Faculty Recruitment

A report to the Economics department by the EDI committee

January 2023

1. Introduction

There is increasing awareness around the lack of diversity in academia in terms of gender, race, ethnicity, and various other social dimensions, which is, of course, a reflection of this phenomenon in society more broadly. This occurs at the level of hiring as well as promotion, and the gaps are particularly large at higher and more senior positions. The flagship associations of the academic economics profession such as the AEA and the RES are increasingly recognizing the problem and publishing statistics on it in regular basis, as well as taking various measures to promote diversity and combat explicit and implicit bias.

To the extent biases exist against under-represented groups (URG) that limit diversity and make the body of academic economists not representative of society at large, there are many downsides to that. First, this is undesirable from the point of view of fairness or equal treatment of equals. Second, it also leads to inefficient outcomes as there is insufficient tapping of the potential talent pool. Third, a diverse and societally representative student and faculty body creates a better learning environment, fosters creativity and innovation in research as well as creating a positive campus environment by drawing in different perspectives and life experiences in the curriculum and research agenda, with resulting impact on policy. Finally, through role model effects and breaking down of negative stereotypes it can have a continuing positive effect by attracting a better and more diverse pool of talent in the student and faculty body.

The plan of this paper is as follows. In section 2 we discuss some broad summary statistics about the representation of women and minorities among academic economists, as well as in the student body. In section 3.1 we summarise the theoretical literature on discrimination that leads to inefficiencies and henceforth the benefits from diversity. In section 3.2 we review the empirical literature on discrimination in the Economics profession and other related professions, highlighting the possibility of the URG candidate group representing a higher quality, which either arises passively or actively by positive selection into the pool. In section 4 we provide a summary of a first look at the data on our hiring practices in the last 7 years. In section 5 we discuss policies to increase diversity, looking at quotas and alternatives. In section 5.1 we provide a discussion of the theoretical literature of the benefits of quotas or policies that seem to positively discriminate in favour of URG. In section 5.2 we provide a discussion of the

empirical literature on the usefulness of quotas in the short and long term, as well as a discussion of alternatives to quotas. Finally, in Section 6 we provide some recommendations.

2. A Snapshot of Diversity in the Economics Profession

The Committee on the Status of Women in the Economics Profession (CSWEP) of the AEA in its recent report (2022) shows that in 2021, only 13.1% of Full Professors were women in Top-20 schools, with the corresponding fractions for Associate Professor and Assistant Professor being 21.2% and 22.7%, and for overall tenure track faculty, the fraction was 16.7%. In non-tenure track faculty, the proportion of women was higher (39.5%), giving a total fraction of 20.9% for all faculty (tenure and non-tenure track). However, there is an upward trend, with data from 1994-97 indicating that only 4.3% of Full Professors were women, and the corresponding figures for Associate, Assistant, All Tenure Track Faculty, All Non-Tenure Track Faculty, and All Faculty being 11.9%, 18%, 9%, 37.3%, and 10.2%. Interestingly, the growth is the sharpest at higher ranked positions (Full and Associate Professor). The study also provides data on placement of women among new PhDs from top ten and twenty departments, which shows an upward trend over this period (1994-97 to 2021). In particular, the fraction of women among new PhDs from top 20 schools who landed jobs in US PhD granting departments rose from 24% in 1994-97 to 30.8% in 2021. Interestingly, 2019 was a year with particularly high rates (41%) but it is not clear whether that was an outlier.

A recent report by the Committee on the Status of Minority Groups in the Economics Profession (CSMGEP) of the AEA (2022) present data from the annual survey, the Universal Academic Questionnaire (UAQ), of approximately 850 degree granting institutions for the percentage of economics faculty by race/ethnicity in the academic year 2020-21. With some cautionary note about the data since it is survey-based and the response rates are not high (about 41%), the report provides some interesting findings. It reports that among institutions included in the survey, representation of minority faculty in economics (across all academic positions) totals about 8.2% which is far less than the 31.9% that is the Black, Latino and Native Americans proportion in the population. Only 6.9% of full-professor positions are held by such minorities, with Hispanic and Black economists making up 4.2% and 2.3%, respectively.

A recent study (Li and Koedel, 2017) looked at faculty at forty selective public universities ranked highly by the 2016 U.S. News & World Report and looked at six academic departments (biology, chemistry, economics, educational leadership and policy, English, and sociology) in the 2015–2016 academic year. They find

that Black, Hispanic, and female faculty are underrepresented relative to their population shares in the US, whereas White, Asian, and male faculty are overrepresented, and when broken down by field the gap is largely coming from underrepresentation in STEM fields (which in this study includes biology, chemistry, and economics). Looking at faculty rank, the study finds that Assistant Professors are less likely to be White and more likely to be Asian and Hispanic, and less likely to be male than Associate and Full professors, which suggests a potentially positive trend for the future, with the exception of Black faculty – they are as underrepresented in junior faculty as they are for senior faculty. Since the study focuses on public universities in the US, it can use publicly available salary data and finds that annual earnings by race-ethnicity and gender reflect gaps for Black and Hispanic faculty relative to Whites and to a lesser extent Asian faculty. They find that academic field, work experience and research productivity almost entirely explain the race-ethnic gaps but while they explain a large part of the gender gap too, they do not fully explain it.

Turning to the UK, according to a RES2021 report (Bateman et al, 2021) women are substantially more likely to be employed at lower academic ranks. In 2018, women made up 33 percent of lecturers, 27 percent of senior lecturers/readers and 15 percent of professors. The study reports that the overall growth in women's representation is upward, but it has slowed down since 2012, particularly among lecturers and professors, whereas the growth has been stronger among senior lecturers/readers. Interestingly, the report points out that women have been gaining ground in teaching-only and research-only positions but have made slower progress obtaining positions with both responsibilities. The report also notes that female academic economists disproportionately originate from outside the UK -- of the 455 female economists employed on standard academic contracts in 2018, only 28% were UK nationals.

A recent report (Advani et al, 2020) on ethnic diversity in UK using data from the Higher Education Statistics Agency (HESA) presents some statistics on the representation of various ethnic groups among academic economic researchers in the UK. They find that the share of non-White ethnic minorities in academic economists doing research and teaching was 19% in 2012-13 and this has increased to 24% in 2018-19. This is higher than their proportion in the UK population (13% among individuals aged 25-64 as of 2011) as well as higher than in UK academia in general (17% in 2018-19). However, the report notes that this is largely driven by individuals coming from abroad. Among academic economists doing research in the UK, only 40% are UK nationals. Among non-UK nationals, 29% are non-White, while among UK nationals 16% are non-White.

Clearly, the status of women and minorities in the faculty of Economics departments is a reflection of the pipeline of students who enter graduate school in Economics, and indeed, even at an earlier stage, namely, at the undergraduate level. The AEA CSWEP Report (2022) notes that of the total PhDs in economics awarded in 2020, 32.8% were women, while the corresponding figure for 1994-97 was 24.7%. The corresponding figures for undergraduates with majors in Economics who graduated were 34.7% in 2020, and 32% in 1994-97.

Data on the race and ethnicity of US Economics PhDs is available only for those who are U.S. citizens or permanent residents, and according to the AEA CSMGEP (2022) report, within this group of Economics PhDs 6.9% were Latinos, and 4.3% were Black, with the overall share of minorities being 11.64%. For all Economics degrees (BA, MA, and PhD), the share of minorities was slightly higher at 17.65%.

In the UK, according to the RES2021 report, in 2018 women constituted 32% of undergraduates, 52% of Master's students, and 39% of Ph.D. students in economics. The report notes that the representation of women is worse among UK nationals – for example, the proportion of UK-domiciled economics undergraduates who are women was 27% in 2018, while for Master's students, it was 31%.

The report by Advani et al (2021) on ethnic diversity in Economics in the UK notes that among British nationals, in 2012-13 students from ethnic minorities made up 33% of economists at the undergraduate level, 33% at the Master's level, but only 18% of economics PhDs. In 2018-19, the proportion of ethnic minorities who were undergraduate students among British nationals went marginally up to 37%, the share at the Master's level did not change, but for PhDs it went up to 28%. Once again, the report focuses on students of British nationality, who accounted for 58% of those studying economics and 67% of undergraduate economics students in 2018–19. Ethnicity information is not available for a majority of international students.

While these reports do not in general present data that look at the distribution of these numbers by ethnicity *and* gender, the RES2021 report (Bateman et al, 2021) offers some interesting numbers. The reports states that in 2018, 8% of standard academic posts in economics were held by Black and minority ethnic (BME) women, whereas BME men constituted 17% of all standard academic contracts, which suggests that the combination of race and gender makes the gap starker.

To sum up, both in the US and the UK, women and ethnic minorities are generally well under-represented both in the faculty and in the student body.

Looking at our own Department, the percentage of women by rank is as follows, 11% for Full Professors, 23% for Associate Professors, 41% for Assistant Professors, and 43% for Teaching Track faculty. Overall, for All Research Faculty 14 out of 64 or 22% are women. Comparing with the UK average for 2018, LSE Economics is ahead only at the rank of Assistant Professor – for both Full and Associate Professor levels, it is below the UK average.

If we look at ethnic/racial minorities (both genders), then at LSE Economics 8 out of a total research faculty of 64 or 12% would be fall under the category of ethnic minorities. This is lower than the average for UK academia (17% in 2018–19).

Clearly, with promotions and recruitment, the gender imbalance at the top (Full and Associate Professor levels) will be corrected to some extent but we cannot expect a large change in the near future unless we take steps now.

3. Why it is beneficial to have diversity.

3.1 Theoretical models on discrimination and the implied benefits of diversity

In this first part we summarise (mostly recent) theoretical arguments pointing at different types of discrimination and their implications. The inefficiencies highlighted in this literature, mostly due to loss of talent and inefficient allocation of talent, indicate that reversing wrong discrimination (which we consider as analogous to increasing diversity) has welfare benefits.

3.1.1 Taste-based discrimination

We will not discuss in details theories based on “taste” for discrimination. These can be interpreted as implying that discrimination is efficient, and this is arbitrary as in the same vein individuals may have a “taste” for diversity. However, this can also be interpreted as a “tax” or negative externality on those who are discriminated against face in the labour market, which would justify corrective policies. An optimistic implication of the early models of taste-based discrimination is that in competitive markets, those employers who indulge in discrimination will be at a disadvantage relative to neutral employers. However, with market-power on the part of employers discrimination can persist – after all, monopsonists can discriminate in both senses of the term. Moreover, once one adds labour market frictions, such as search frictions, discriminatory equilibria can persist (Lang, Manove and Dickens, 2005).

Let us henceforth refer to a worker of type A as a privileged type in society (we will define privilege depending on the environment we consider) and a type B a disadvantaged type.

3.1.2 Classic Statistical discrimination models

The literature on statistical discrimination originates with Phelps (1972) and Arrow (1973), where firms view type B's distribution over quality as either having a lower mean, or a higher variance and hence these types are viewed as a more risky prospect. These and the literature that follows show how self-fulfilling expectations arise in equilibria: If employers expect type B to have a lower probability of being qualified, then they will have a lower expected return of 'hiring' type B. As in this case qualified type B workers will be paid less than type A, indeed type B best responds by making lower investment in human capital in the first place. This implies that "discrimination" is justified when one conditions on type, and equilibrium beliefs are self-fulfilling. Coate and Loury (1993) builds on Arrow's discriminatory model where employers receive only a noisy signal of the human capital investment. In the event of an ambiguous signal, employers make inferences based on the worker's social group. Thus workers from social groups with worse human capital are at a disadvantage, which weakens their incentive to invest in the first place. This behaviour justifies the employers' stereotyping of social groups. A negative stereotyping equilibrium is again justified if a particular social group initially had low levels of human capital for (exogenous) historical reasons.

In the above, self-fulfilling equilibria are inefficient as equilibria with higher investment in human capital would lead to higher production and higher welfare in society.

We proceed to highlight more recent literature on statistical discrimination, which focuses on the possibility that type A and type B have the same distribution over quality, but are still perceived to have different distribution by employers. Again, discrimination is always inefficient here.

3.1.3 Statistical discrimination due to learning frictions (technological/myopic social learning):

In the first model we describe, discrimination arises due to potentially wrong prior beliefs on type B, combined with a need of an employer to invest in a costly signal to ascertain a worker's quality. Specifically, Cavounidis and Lang (2015) is a recent example of a model of statistical discrimination that aims to reproduce a number of stylised facts regarding labour market discrimination, including a wage

differential, as well as longer unemployment spells for black workers. The basic idea is that if worker quality is unobserved, and employers can assess quality with noise at a cost, then they will engage in costly assessment only if their prior beliefs about the worker quality is sufficiently low. Workers who are detected to be of low quality are fired and returned to the pool of unemployed workers. If employers start off with very negative beliefs about the quality of workers from a certain social group, they will undertake costly investments to observe their quality which, in turn, will worsen the pool of unemployed workers from that social group. This can result in an equilibrium where black and white workers with the same quality distribution are treated differently, e.g. white workers do not face costly assessment, and the pool of unemployed workers is of high quality, and black workers face costly assessment and the pool of unemployed workers is of low quality.

Small differences in initial priors can also imply long term differences as employers allocate workers to different tasks in which they learn differentially about agents' types: Bardhi, Gue and Strulovici (2019) characterise environments in which a small difference in the perceived ability of types A and B (which may be wrongly perceived) can be translated into large differences in the future stream of wages/promotion under particular allocation of workers to different tasks with different signalling technologies. For example, type A worker on whom there is a slightly higher prior, may be allocated to a more risky task where the employer can learn about her quality and potentially increase their belief, whereas a type B may be allocated to a safe task, where no further learning is possible. In an extension, they show how this will create a differential incentive to invest in ability. See also Fosgerau, Sethi and Weibull (2021). A similar phenomenon can arise not due to signalling technology per se, but in a social learning environment in which each firm myopically hires the perceived better worker of type A— thus there is too little experimentation with type B worker. See also Li, Raymond and Bergman (2020).

Discrimination can also arise due to different signalling technologies available to types A and B themselves (rather than when the employers choose how to learn about workers): Chambers and Echenique (2021) and Escudé, Onuchic, Sinander and Valenzuela-Stookey (2022) characterize environments where type A and type B have the same skill distribution, but different skill-signalling technologies to show this skill. Chambers and Echenique's (2021) result implies that discrimination is inevitable. More specifically, whenever there are distinct populations with the same skill distribution, but different signalling possibilities, there will be discrimination between them.

3.1.4 Statistical discrimination with wrong beliefs, their persistence, and effects:

We now describe different models in which the authors consider how wrong beliefs on type B relative to A can persist, evolve and possibly change, in environments of social learning and a host of potentially cognitive biases. Bohren, Haggag and Imas (2019) propose a model and show experimentally how evaluation of women and men on math advice websites suffers from wrong beliefs about agent's types. Their key contribution is to consider a dynamic environment, in which statistical discrimination in the first phase, can be reversed by those that are not discriminating and are aware of such first-phase behaviour by other evaluators. In particular, such second-phase evaluators are aware that a woman who receives the same first phase evaluation as a man, is of higher innate quality due to the first-phase discrimination against women. Their experiment's results show how such reversal arises, a reversal which is not consistent with taste-based discrimination.

Levy and Razin (2017) illustrate how wrong beliefs about type A versus B workers -that imply discrimination in the labour market- can perpetuate in the long term; They consider an environment in which individuals choose their networks endogenously, anticipating the potential influence of others' beliefs on themselves. This implies that they prefer to be in a network with individuals with similar beliefs. In such endogenous "chambers", beliefs can "echo" and become extreme, as individuals ignore their selection bias and learn from the beliefs of others in their chamber. Discrimination in such an "old boys network" can become *structural* as it is worthwhile for individuals with the same wrong beliefs to coordinate on their network, and entrepreneurs may also step in to form such networks.

In a different model of wrong beliefs, Heidhues, Köszegi and Strack (2019) analyze a model in which individuals have wrong beliefs about themselves – specifically they are wrongly overconfident (with pointwise beliefs that do not change, and so they have a misspecified model). This implies that they interpret the success of themselves and people similar to them as a result of their high ability -and failure of the other group as the result of inability. When their group fails though, they attribute it to discrimination, and become even more sure of this when they are themselves of low ability.

Finally, Bordalo, Coffman, Gennaioli and Shleifer (2016) propose a specific model of wrong beliefs arising from a cognitive bias of interpreting distributions over qualities where small differences become amplified.

3.1.5 Specific literature on the effects of wrong statistical discrimination in Academia: While many of the above can be applied to the academic profession,

two papers highlight issues relating to the academic profession, including co-authorship and the refereeing process:

Onuchic and Ray (2022) study signalling through team formation and find that discriminatory outcomes may arise where team members that belong to different identities may systematically receive different credit for team outcomes (e.g., in the case of co-authors). One implication of their work is that type B workers will keep the best ideas for their own work or co-authored with other type B workers.

Siniscalchi and Veronesi (2021) analyze a model in which male referees have a self-image bias and appreciate the research more when it is conducted by a male researcher, while research characteristics across the two populations can be different (creativity, technical ability) but of the same total quality. They show that if initially the population of referees is very unbalanced and mostly male, the long-term equilibrium in society may become extreme with almost only male remaining the profession following successful evaluations, which implies a large loss of talent. As they show, researchers' career concerns and institutions' practices can exacerbate such talent loss.

We have provided only a partial review of recent literature on statistical discrimination; for more, see an excellent review in Onuchic (2022).

3.1.6 Benefits of diversity given initial different characteristics:

Above we focused on a literature that indicates that the main benefit of “diversity” in the sense of the reversal of wrong discrimination is **a more efficient talent pool**. Mainly, inefficient discrimination implies that: (i) in the case of equally talented types A and B, firms/organisations do not tap into the full talent pool and may draw upon weaker A types compared with better B types, due to wrong beliefs or different signalling technologies. (ii) unequal quality distribution of types A and B may arise partly due to statistical discrimination providing weaker incentives for type B to invest and so fixing this will potentially create more efficient incentives. This is naturally harder for one firm to fix given equilibrium considerations.

There are reasons however for benefits from diversity even when there are initially different characteristics, as described below:

Production considerations: Prior discrimination may potentially lead to different characteristics developed by individuals that are not directly summarised by “quality” as above, or potentially to situations in which type B has lower quality, as in the original literature on statistical discrimination. In some environments, technology may imply that diversity or alternatively homophily

may yield more efficient production. Some examples can be team production where outputs can be all additive or alternatively strongly complementary. While people who have similar background may be able to work better together (assortative matching), negative sorting can also be efficient when output/ability is not a scalar (quick problem solving vs creativity, driven but self-centred vs those with social skills who can make a team work, people with very different backgrounds may teach each other softer skills and knowledge). Such latter environments may be more pronounced in the classroom or in the academic world.

A particularly interesting recent contribution is Sethi and Somanathan (2022) who consider an environment in which production is a function of ability and training. Due to less privilege, type B individuals have less training compared with type A. If production is higher for <high ability, low training> compared with <low ability, high training>, then diversity according to training level (which is what firms observe) can be more efficient. This indicates that **reducing entry requirements for type B is efficient**. See more in Section 5.1.

Two papers consider the idea of the role of **role models** (that is, the benefit of diversity at the “top” where diversity exists at the “bottom”) and their ability to increase the quality of type B workers. Athey, Avery and Zamsky (2000) consider a production environment in which to utilize talent efficiently for employees at a lower level, the firm needs a diverse set of top-level employees. This is derived with an assumption that mentoring is done by type and there are decreasing returns of having many mentors of the same type. Multiple steady states are possible in the long run -if mentoring is very important in total production, then firms may become homogeneous (with type A if this is the starting point). Chung (2000) also considers a role model environment in which individuals of type B provide a signal to type B workers that they can be successful. See more in Section 5.1.

3.2 Discrimination in the Economics Profession and similar ones: Empirical

The representation of women and minorities in faculty groups in economics has increased since the 90s but, contrary to popular belief, has not been incrementally improving over the past 10 years (CSMGEP 2017 Report). There is empirical evidence from our own field suggesting that unequal treatment of female scholars is a likely contributing factor. Among other studies, Heather Sarsons’ work on coauthorship in economics suggests that female economists are implicitly given less credit for co-authored papers relative to male economists and, accordingly, receive tenure at a lower rate relative to their male counterparts for the same research output (Sarsons et al, JPE 2019). In that same paper, the authors are

unable to perform the same analysis for racial minorities in economics because there are simply too few racial minorities in top departments to analyse with classical statistical methods. Relatedly, Koffi (2021) uses machine learning algorithms to identify similarity across papers and establish which papers should be cited; She shows that papers omitted from references are 15% more likely to be female-authored than male-authored.

Even in cases where women are tenured, unequal treatment permeates economics culture. In a recent NBER working paper, Dupas et al (2021) conduct a massive observational survey of economics seminars (primarily using graduate student ratings) and find that female presenters in economics seminars are treated differently than their male counterparts. They elaborate: “Women are asked more questions during a seminar and the questions asked of women presenters are more likely to be patronizing or hostile. These effects are not due to women presenting in different fields, different seminar series, or different topics, as our analysis controls for the institution, seminar series, and JEL codes associated with each presentation.”

A very recent NBER Working Paper by Card et al (2022) finds that female scholars in all three fields of Economics, Psychology, and Mathematics are significantly more likely to be inducted in the National Academy of Sciences and the American Academy of Arts and Sciences, than male scholars. Within economics, some of this can be explained by the larger pool of men in the profession, suggesting greater competition, while much fewer women have entered. It can also imply an attempt to correct the problem of diversity by some visible measures. Finally, this can also reflect selection by unobservables. While the paper attempts to control for research productivity through number of papers and citation counts, the quality of publications is not weighed. The authors also caution that this metric does not account for, for example, women facing greater obstacles in achieving similar records in publishing or citation, some of which are highlighted in the research in the paragraphs above. This would suggest that women could be more positively selected on unobservables than equivalent-quality men are into the profession, thereby justifying positive discrimination based on observables.

A recent working paper by Ashraf, Bandiera, Minni and Quintas-Martínez (2022) identifies such positive selection in a professional environment. The authors use a novel dataset of 100K employees doing the same job, with the same qualifications, in over 100 countries to quantify selection on unobserved ability into the labour force. In almost every country studied, they find positive selection by women into the overall labour force; this positive selection on ability is higher

the lower female labour force participation is in a given country, and the higher the barriers to entry for women are as measured by gender norms. Using structural estimation, they find that equalizing these barriers to labour force participation by women would increase productivity on average by 32%. The main insight of this paper is that the observed minorities in an applicant pool are likely positively selected, given the barriers they needed to overcome in order to be in the applicant pool. This has a strong managerial policy implication for additional consideration given to minority candidates, and an understanding by evaluators that selection into the applicant pool may well be at play despite observables looking equal (in our Department, this logic may help in perceptions of women and minorities who come through for their interviews and job talks). The paper runs counterfactual policy simulations to increase labour force participation by women and finds that providing greater returns to ability through steeper salary curves would draw in significantly more women. This is a different logic than needing to compete for the few women who are in the labour force; it is, rather, a compensation for the underlying differential ability (see further recommendations along these lines in Sections 5.2 and 6).

4. Hiring practices at the department of Economics at LSE

This section reports on data from junior recruitment in the LSE Economics Department from 2015 onwards, to shed light on where the greatest challenges to gender parity might be. Our Department undertook two empirical exercises using our own data. We first constructed a funnel of the applicant pool until job offers. The percent of women applying for junior positions at the LSE has been remarkably stable over the past 7 years, between 25-30% of the full applicant pool. At every stage since- in long-listing, short-listing, interviewing, flyouts, and offers- women make up increasing proportions of the pool, suggesting either positive underlying (self)selection, or positive discrimination. The Department instated a 50/50 gender quota in gender recruiting for interviews in 2016-2017. Prior to this, the shortlist for interviews consisted of 44% women, and flyouts remained at 50%. Since the quota, the interviews and flyouts have been at least at 50%, increasing by year (with the notable exception of 2019-2020 where only 27% of interviews, 25% of flyouts, and 17% of offers were to women).

Indeed, in 2021, 75% of the flyouts were women. Women do not appear to be significantly less likely to accept our offers, although we are looking at very small numbers in this domain. We are now doing the same exercise for minority, non-white candidates.

Keeping track of these overall pipeline numbers can help us as a Department understand which levers might be most helpful to pull. In particular, increasing

applicant pool numbers, proactively seeking women- particularly in lateral hires- and, as discussed below, offering benefits that could encourage even more positive selection, could be actions the Department could take.

One other exercise we undertook was to go back to those candidates who gave flyouts but who did not receive offers, and tracking their career progress via publications, citation count, and present position (data attached). There are two insights that emerged from this exercise: the first is that there have certainly been notable candidates who have had extensive success since we passed on them- they are most likely to be lateral hires one or two years out but which didn't have publications at the time they interviewed with us. This fact in itself encourages a degree of humility towards our ability to perfectly predict quality and subsequent performance. The second is that, at least in the very small numbers we have and using this post-metric of realized success, it appears that we were just as likely to pass on good female candidates as male ones.

5. How to increase diversity: Quotas and alternatives

5.1 Quotas as a useful way to address the need for diversity: Theory

We now build on the theoretical models described in Section 3.1 to assess the effects of quotas favouring type B, and whether these can increase efficiency.

If one considers the classical statistical discrimination models (as in 3.1.2), and takes as given an equilibrium with self-fulfilling expectations in which type B has lower quality on average than type A, then a measure such as quota favouring type B will lead on average to lower quality. For example, Fryer and Loury (2013) assume that type B has higher investment cost and hence invests less and is less represented; so necessarily, affirmative action reduces welfare. Coate and Loury (1993) show that for certain parameter values, affirmative action policies can adversely affect investment incentives within the disadvantaged groups, such that differential levels of human capital across social groups is maintained in equilibrium even with such measures.

Turning to models described in 3.1.3-3.1.5, environments with equally qualified candidates that allow for discrimination due to wrong beliefs or differential signaling of talent, lend themselves easily to the conclusion that quotas increase efficiency due to organisations tapping into a better talent pool. Papers that specifically analyse this measure and show how efficiency increases are Siniscalchi and Veronesi (2021) and Fosgerau, Sethi and Weibull (2021). Siniscalchi and Veronesi (2021) analyses how self-image bias allows male referees to favour male researchers at the expense of the different but equally

productive female researchers. A quota easily corrects this and allows society to evolve to the most efficient and fair equilibrium. Fosgerau, Sethi and Weibull (2021) discuss how differential learning on types A and B that stems from wrong beliefs about the superiority of type A leads to inefficient incentives of type B to invest in quality. Importantly in their model, a selector chooses how much to invest in screening, and so low beliefs on type B imply little learning (which yields the inefficient investment incentives in quality). Quotas improve the type B quality distribution through more active screening. The authors show that candidates accepted under a quota from type B are more likely to be qualified than those rejected (reflecting active screening). At some point the quota becomes non-binding, and the group is treated in the same way as one for which no quota has been imposed. More subtly, at no point beliefs about type B (when the quote is binding) are more optimistic (conditional on acceptance or rejection) than the corresponding beliefs about type A.

Turning to models described in 3.1.6, Sethi and Somanathan (2022) assume a production function in which individuals with high ability but lower training (which is what the firm observes, but can be lower to type B due to less privileged position) can produce more than individuals with low ability and high training. They show then that hiring according to a perceived quality measure such as training is sometimes inefficient; to capture these individuals with high ability and low training, the organisation needs to reduce its requirements in terms of training or observable quality measure, which can be attained by a quota for those with low training. Following from Chung (2000)'s role model environment, quotas or other measure of affirmative action can trigger a switch to type B learning that they can be successful, which can increase human capital investment and efficiency.

Interestingly, if there is a direct preference for diversity, then affirmative action also leads to higher quality: Chan and Eyster (2003) show how an admissions office that cares both about quality and diversity admits the best-qualified candidates from each group when allowed to use quotas. Under a ban of affirmative action, it may promote diversity by partially ignoring candidates' qualifications and therefore not admitting the best-qualified candidates from either group.

5.2 Quotas as a useful way to address the need for diversity: Evidence

In this section we review the empirical literature on quotas. Throughout, the literature shows that there is no substantial cost in terms of quality and in fact a benefit in terms of quality can arise. The benefits arise through the mechanisms discussed above in the theoretical literature: Through better learning on type B in

environments in which this was not facilitated due to wrong priors or different signaling technologies, through potentially overcoming inefficient discrimination, and through positive role model effects. The literature also indicates that quotas can be a temporary measure that becomes non-binding in the long term.

5.2.1 Evidence on the Direct Impact of Quotas on Quality

The empirical economics literature speaks directly to the concern that affirmative action via quotas means compromising on quality, and finds overall evidence against this concern. The most striking evidence comes from our own colleague: Besley et al (2017) study Swedish politicians and find two important results. First, “Contrary to the expectations of quota skeptics, women’s competence did not go down but stayed roughly constant.” This is congruent with there being a large pool of untapped/unused talent among female politicians at the margin. Moreover, there were beneficial spillover effects onto the quality of the selected male politicians: “Far from being at odds with meritocracy,” they conclude, “this quota raised the competence of male politicians where it raised female representation the most.” This finding is consistent with the notion that male applicants are overrepresented due to positive discrimination in their favour.

Balafoutas and Sutter (2012) find a similar result in a controlled laboratory experiment with undergraduate students across different disciplines. They set up several controlled experiments in order to study four interventions: quotas, where one of two winners of a competition must be female; two variants of preferential treatment, where a fixed increment is added to women’s performance; and repetition of the competition, where a second competition takes place if no woman is among the winners. They evaluated the results on team performance relative to no gender-related influence on selection. Their findings suggested there was no drop in quality at all arising from the gender quota regardless of which way it was implemented: “Compared with no intervention, all interventions encourage women to enter competitions more often, and performance is at least equally good, both during and after the competition.”

In Bagde, Epple and Taylor (2016) the authors study a quota system for caste minorities in the Indian engineering colleges and find no evidence that the applicants selected due to the quota were in any way underprepared or underperforming in their college: “in short”, they conclude at the end of their introduction “we find no evidence of mismatch.” (The paper focuses on the harm such mismatch would do to the students selected, but the evidence that there is no mismatch is of course relevant to concerns that lower quality students would eventually lower the average competence of their professions.) Bertrand, Hanna and Mullainathan (2009) find similar lack of mismatch in a similar setting,

although they note that affirmative action along one dimension can worsen representation issues along another (in their setting, it is caste representation versus female representation) which suggests prioritisation is necessary when making these decisions; one does not get a “free lunch” from quotas along one axis of disadvantage, and the department may need to decide which of affirmative action for female applicants or ethnic minority applicants is its priority. Earlier literature in the USA, such as Loury and Garman (1993), had found more evidence of mismatch among elite college students selected via racial affirmative action, but the effect was not large enough to outweigh the gains in performance and human capital, and the selected students still performed well enough to out-earn their counterparts.

5.2.2 Persistence in Effects of Temporary Affirmative Action

Another important and more recent empirical finding is that even temporary affirmative action tends to have persistent positive gains. Miller (2017) studies US federal mandates for racial affirmative action and finds that the share of black employees continued to rise even after the program ended, which he attributes to employers’ improved ability to effectively screen candidates from racially underrepresented backgrounds, due to their greater experience with these candidates during the mandated period. This result is consistent with a bandit-style problem in which historical underrepresentation of a certain group in a certain sector leads to so much uncertainty over the distribution of talent in that group that employers will find it safer to hire from the group they have experience with. Card and Krueger (2005) find in the USA that racial minorities continue to apply to elite colleges that have had affirmative action policies even after those policies have ended.

Overall this suggests that the department will not need to engage in quotas forever – it should be viewed as a temporary corrective mechanism in order to break our path dependence from an undesirable initial equilibrium in labour allocation.

5.2.3 Spillovers from Affirmative Action for Faculty onto Students

Bettinger and Long (AER, 2005) find evidence in the US college setting that female faculty can increase interest in their subject matter among female students in quantitative scientific fields. A more recent study of MIT undergraduate instruction found that female students benefit more from being mentored by women and receiving class instruction from women relative to male instructors (Russell, 2017). This is congruent with earlier research showing that students in engineering tend to prefer mentors who share their race and gender (Chesler and Chesler 2002). There is also evidence that these students who do persist in these cases also tend to perform well in the field, although the mechanism is complex

and multifaceted: According to Thomas, Willis and Davis' 2007 qualitative study, "These [same-gender faculty mentoring] relationships are especially important for racial minorities who often lack access to informal networks and information that is required to be successful in academic and professional environments in which they are under-represented." Thus, they conclude: "The lack of mentors for minority graduate students is important to consider given the potential impact of this experience for minority graduate students' retention and subsequent success, but also for the future diversity of the discipline (especially its instruction and research)."

5.2.4 Evidence that Alternative to Affirmative Action are Ineffective

The literature also suggests that alternatives to direct affirmative action are unlikely to be successful. For example, Long (2003) finds that affirmative action in US college admissions that focuses on the highest-achieving students from historically underrepresented high schools cannot achieve the same racial equity outcomes as affirmative action on historically underrepresented races directly. This result is congruent with the idea that discrimination is operating within every level of the human capital accumulation process.

Alternatively, Ashraf, Bandiera, Minni and Quintas-Martinez (2022) discuss as above the usefulness of compensation policies for women as triggering positive selection. In line with this paper, more creative solutions could also involve specially designed research funding support for URG or offering packages which could be particularly appealing for women, and then help overcome structural disparities, for example, money to take children to conferences. In addition, understanding the particular intra-household constraints that women often have to face may lead us to look into smaller departments where women may be more likely to attend.

6. Recommendations

We have provided in the introductory section different motivations to increase the diversity of the faculty in the department (some of these reasons can provide motivation to increase the diversity of the student body and professional staff). We focused in the rest of the paper on the efficiency motivation to increase diversity in the department. We highlighted in this paper the following: (i) There is indication that URG are discriminated in the Economics profession; (ii) Discrimination can potentially take the forms highlighted in the theoretical literature, such as stemming from wrong beliefs and different signalling technologies; (iii) Such discrimination implies that organisations do not tap into the largest talent pool and hence quotas can increase efficiency. Quotas as we

highlight are also beneficial given the other motivations to increase efficiency mentioned in the introduction, such as fairness, role models for students, direct benefits from diversity, and others.

We therefore recommend that the department continues in its previous effort to increase diversity by employing a target for URG at every stage of the junior hiring process. We see no reason to use different targets than we had used before and therefore recommend a target of 2/3 at each stage of the process.

Transcending the exact quota numbers, or even of a quota per se, is the logic of making precise the benefits that diversity can bring for the productivity of the Department, and for each sub-field: due to positive selection, due to role model effects, and due to valued inputs in the academic research and teaching production function. Internalizing this logic could help guide discussions within fields as well.

Importantly, we recommend looking into supplementing any current policies with an effort to also make the department attractive to URG, who may potentially have different needs as they faced, and continue to face, barriers to succeed in the profession. Such measures can constitute differential pay or other forms of compensation, differential research packages, assistance in travelling to conferences and others.

References

Advani, A. Sonkurt Sen, and Ross Warwick (2020): “Ethnic diversity in UK economics”, IFS Briefing Note BN307, The Institute for Fiscal Studies.

Arrow, K. J. (1973). “The Theory of Discrimination.” In Orley Ashenfelter and Albert Rees, eds. *Discrimination in Labour Markets*: 3-33.

Ashraf, Nava, Oriana Bandiera, Virginia Minni, Victor Quintas-Martínez (2022): “Gender Roles and the Misallocation of Labour Across Countries: Evidence from Personnel Data”, Working Paper.

Athey, Avery and Zamsky (2000), “Mentoring and Diversity”, *AER*, 765-786.

Bagde, Surendrakumar, Dennis Epple, and Lowell Taylor (2016): “Does Affirmative Action Work? Caste, Gender, College Quality, and Academic Success in India”, *American Economic Review*, 106(6), 1495-1521.

Balafoutas Loukas and Matthias Sutter (2012): “Affirmative Action Policies Promote Women and Do Not Harm Efficiency in the Laboratory”, *Science*, Vol 335(6068), 579-582.

Bardhi, A., Y. Guo, and B. Strulovici (2019). "Early-Career Discrimination: Spiraling or Self-Correcting?" Working paper.

Bateman, Victoria, Danula Kankanam Gamage, Erin Hengel, and Xianyue Liu (2021): "The Gender Imbalance in UK Economics", Royal Economic Society, Silver Anniversary Women's Committee Report.

Becker, G. (1957). "The Economics of Discrimination." University of Chicago Press.

Bernheim D. (1994). "A Theory of Conformity." *Journal of Political Economy*, 102: 841-77.

Bertrand, Marianne & Hanna, Rema & Mullainathan, Sendhil. (2010): "Affirmative action in education: Evidence from engineering college admissions in India," *Journal of Public Economics*, 94(1-2),16-29.

Besley, Timothy, Olle Folke, Torsten Persson, and Johanna Rickne (2017): "Gender Quotas and the Crisis of the Mediocre Man: Theory and Evidence from Sweden", *American Economic Review*, 107 (8), 2204-2242.

Bettinger, Eric P. and Bridget Terry Long (2005): "Do Faculty Serve as Role Models? The Impact of Instructor Gender on Female Students", *American Economic Review*, 95(2),152-157.

Black, D. A. (1995). "Discrimination in an Equilibrium Search Model." *Journal of Labour Economics*, 13: 309-334.

Bordalo, Pedro, Katherine Coffman, Nicola Gennaioli, Andrei Shleifer (2016). "Stereotypes." *Quarterly Journal of Economics*, 131: 1753-1794.

Bohren, J. Aislinn, Alex Imas, and Michael Rosenberg. (2019). "The Dynamics of Discrimination: Theory and Evidence." *American Economic Review*, 109: 3395-3436.

Bramoullé, Yann, and Sanjeev Goyal. (2016) "Favoritism." *Journal of Development Economics*, 122: 16-27.

Card, D., & Krueger, A. B. (2005). Would the Elimination of Affirmative Action Affect Highly Qualified Minority Applicants? Evidence from California and Texas. *ILR Review*, 58(3), 416-434.

David Card, Stefano DellaVigna, Patricia Funk & Nagore Iriberry (2022): "Gender Gaps at the Academies", NBER Working Paper 30510.

Cavounidis and Lang (2015), "Discrimination and Worker Evaluation", NBER Working Paper 21612.

Chan, J. and E. Eyster (2003), "Does Banning Affirmative Action Lower College Student Quality?", *AER* 858-872.

Chambers, Christopher, and Federico Echenique (2021) "A Characterisation of 'Phelpsian' Statistical Discrimination." *Economic Journal*, 131: 2018-2032.

Chesler, Naomi C., Mark A. Chesler (2002): "Gender-Informed Mentoring Strategies for Women Engineering Scholars: On Establishing a Caring Community", *Journal of Engineering Education*, 49.

Chevalier, Judy (2022): "Report: Committee on the Status of Women in the Economics Profession (CSWEP)." *AEA Papers and Proceedings*, 112: 746-67.

Chung, K. (2000) "Role Models and Arguments for Affirmative Action", *AER* 640-648.

Coate, S., and G. Loury (1993). "Will Affirmative Action Eliminate Negative Stereotypes?" *American Economic Review*, 83: 1220-1240.

Dupas, Pascaline , Alicia Sasser Modestino, Muriel Niederle, Justin Wolfers & The Seminar Dynamics Collective. (2021). "Gender and the Dynamics of Economics Seminars," NBER Working Paper 28494

Escudé, Onuchic, Sinander and Valenzuela-Stookey (2022), "Statistical Discrimination and Statistical Informativeness", working paper.

Fosgerau, Sethi and Weibull (2021), "Equilibrium Screening and Categorical Inequality", working paper.

Fryer, R. and G. Loury (2013), "Valuing Diversity", *Journal of Political Economy*.

Heidhues, Paul, Botond Köszegi, and Philipp Strack. (2019). "Overconfidence and Prejudice," working paper.

Hoover, Gary, A., and Ebonya Washington (2022): "Report: Committee on the Status of Minority Groups in the Economics Profession (CSMGEP)." *AEA Papers and Proceedings*, 112: 768-85.

Koffi, M. (2021), "Innovative Ideas and Gender Inequality", mimeo.

Lang, K. M. Manove and W. Dickens (2005), "Racial Discrimination in Labour Markets with Posted Wage Offers", *AER*, PP.1327-1340.

Levy and Razin (2017), "The Coevolution of Segregation, Polarized Beliefs, and Discrimination: The Case of Private versus State Education", *American Economic Journal: Microeconomics*, pp. 141-170.

Li, D., L. R. Raymond, and P. Bergman (2020). “Hiring as Exploration,” working paper.

Li, D. and C. Koedel (2017): “Representation and Salary Gaps by Race-Ethnicity and Gender at Selective Public Universities”, *Educational Researcher*, Vol. 46 No. 7, pp. 343–354.

Loury, Linda Datcher and David Garman (1993): “Affirmative Action in Higher Education,” *The American Economic Review Papers and Proceedings*, 83(2), 99-103.

Long, Mark C. (2004): “Race and College Admissions: An Alternative to Affirmative Action,” *The Review of Economics and Statistics*, 86(4), 1020–1033.

Miller, C. (2017). “The Persistent Effect of Temporary Affirmative Action”. *American Economic Journal: Applied Economics*, 9(3), 152-190.

Onuchic, P. r D. Ray (2022). “Signaling and Discrimination in Collaborative Projects,” *AER*.

Onuchic, P. (2022), *Recent Contributions to Theories of Discrimination*, working paper.

Phelps, E. S. (1972). “The Statistical Theory of Racism and Sexism.” *American Economic Review*, 62: 659-61.

Report of the the Committee on the Status of Minority Groups in the Economics Profession (CSMGEP) (2017), American Economic Association.

Siniscalchi, Marciano, and Pietro Veronesi. (2021). “Self-Image Bias and Lost Talent,” working paper.

Sethi R. and R. Somanathan (2022), “Meritocracy and Representation”, forthcoming, *Journal of Economic Literature*.

Thomas, Kecia M., Leigh A. Willis, Jimmy Davis (2007): “Mentoring minority graduate students: issues and strategies for institutions, faculty, and students”, *Equal Opportunities International*, 26(3), 178-192.