# Monitoring hiring discrimination through online recruitment platforms

Dominik Hangartner, [1,2,3,⋆] Daniel Kopp,[3,4,⋆] Michael Siegenthaler[3,4,⋆]

[1] *Public Policy Group, ETH Zurich, 8092 Switzerland*

[2] *Department of Government, London School of Economics and Political Science, London WC2A 2AE, UK*

[3] *Immigration Policy Lab, Stanford University, Stanford, CA 94305 and ETH Zurich, 8092 Zurich, Switzerland*

[4] *KOF Swiss Economic Institute, ETH Zurich, 8092 Switzerland*

⋆ *These authors contributed equally to this work.*

**Women (compared to men) and ethnic minorities (compared to natives) face inferior labor market outcomes in many economies[1,2], but the extent to—and the channels through—which discrimination is responsible for these effects remains debated[3,4]. While correspondence tests[5], where researchers send fictitious resumes that are identical except for the randomised minority trait to be tested (e.g. Black vs. White-sounding names), are an increasingly popular method to quantify discrimination in hiring practices[6,7], they can usually study only a few applicant characteristics in select occupations at a particular point in time. To overcome these limitations, we leverage a new approach to investigating hiring discrimination that combines tracking of recruiters' search behavior on employment websites and supervised machine learning to control for all relevant jobseeker characteristics that are visible to recruiters. We apply this methodology to the online recruitment platform of the Swiss public employment service and find that, depending on their country of origin, ethnic minorities face 4–19% lower contact rates than otherwise identical natives. Women face a penalty of**

**7% in male-dominated professions, and the opposite pattern emerges for men in female-dominated professions. We find no evidence that recruiters spend less time evaluating ethnic minorities' profiles. Our methodology provides a widely applicable, non-intrusive, and cost-efficient tool that researchers and policy-makers can employ to continuously monitor hiring discrimination, to illuminate some of the drivers of discrimination, and to inform approaches to counter it.**

Labor market outcomes such as wages and unemployment differ markedly across socio-demographic groups defined by immutable characteristics such as gender, ethnicity, or race. A vast literature has explored how differences in workers' skills, abilities, or preferences contribute to such employment disparities[1,2]. Because of its fundamental repercussions for equality of opportunity, prior research has paid particular attention to the extent to which these disparities are driven by discrimination[3,5–7]—understood as the decision to hire a person or pay a wage based not on the individual's merit but on his or her membership of a particular group, defined, e.g., by gender or ethnicity[8].

Earlier studies on discrimination relied on observational data such as labor market surveys to estimate a "minority coefficient" that captures the differences in wage and employment outcomes between, for example, Black and White people, controlling for other observable worker characteristics. The limitation of this regression-based approach is that the minority coefficient typically does not identify the extent of discrimination, since it is plausibly confounded by productivity signals that are unobserved by the researcher[1]. To overcome this limitation, correspondence studies have been proposed[5], where researchers send fictitious resumes that are identical except for the minority trait to be tested (e.g. Black- vs. White-sounding names). Over the last decade, correspondence studies conducted in dozens of countries have consistently provided evidence of ethnic discrimination in hiring practices (for overviews, see[6,7]). Due to ethical concerns[9] and logistical constraints, correspondence studies usually randomize only a few applicant characteristics and focus on a small number of entry-level jobs at a particular point in time. In addition, most correspondence studies have a limited ability to differentiate between competing theories of differential treatment (for an exception, see[4]). Thus, although there is a large body of research on this topic, we still lack the ability to comprehensively monitor the level of labor market discrimination associated with a range of immutable characteristics, across occupations, and over time. The results of such monitoring are required in order to determine when, against whom, and through which mechanisms discrimination operates.

Here, we take important steps towards such monitoring by advancing a broadly applicable methodology, which consists of tracking recruiter search behaviour on employment websites and analyzing this wealth of click data using supervised machine-learning algorithms. Employment websites from private companies and public providers are rapidly growing in number and size around the world[10]. Typically, employment websites provide a platform for recruiters to post job openings, which jobseekers can search for and apply to, and for jobseekers to post their resumes, which recruiters can search in order to fill vacancies. Our approach leverages the data generated by the second functionality to infer hiring biases among recruiters. For this, we build on [11], who study hiring decisions on an online crowdsourcing marketplace. We expand on existing research[12] by focusing on online recruitment for the offline labor market, which strengthens the generalizability and external validity of our approach. Furthermore, we not only leverage information on contact attempts, but also on the duration and time when recruiters visited a jobseeker profile, which helps us to shed light on some of the channels of discrimination.

We illustrate this approach by analyzing the search behaviour of recruiters on the online recruitment platform of the Swiss public employment service. On this platform, recruiters can enter the search criteria (e.g. occupation and place of work), request a list of jobseekers that match these criteria, inspect jobseekers' resumes, and click on a button to obtain their contact information (see Extended Data Figures 1a–1c). We use the click on the contact button as our binary outcome measure of request for a job interview. We validate this outcome by linking jobseekers' profiles to the unemployment register and showing that every click on the contact button increases the probability to leave unemployment in the next three months by 2.1% (95% CI: 1.7, 2.4; ordinary least squares regression, see Methods section for details). In addition, we also record the time that recruiters spent looking at each profile.

In the following, we compare profiles that appeared in the same search and use ordinary least squares regressions to control for all jobseeker characteristics (and their first-order interactions) that are predictive of contact or ethnicity (or gender, respectively). The Methods section and Supplementary Information (SI) provide more detail on the Lasso-based[13] post-double selection method[14] used to select the features that entered the regression (3,729 covariates for the ethnicity regressions, 1,927 for the gender regressions). Note that case workers, rather than jobseekers, enter and validate the information about unemployed jobseekers on this platform (see SI for further details). Because we can observe the same information that is available to the recruiter in a given search, our methodology allows us to identify the causal effects of immutable characteristics such as gender or ethnicity on being contacted. Extended Data Figure 3a validates this assumption and shows that there is no effect of ethnicity on contact rates for the subset of recruiters that are not able to view this information, suggesting that the Lasso-selected control variables account for all factors correlated with ethnicity and contact rates.

**Testing for ethnic discrimination**

We showcase the inferential power of this approach by focusing on ethnicity- and gender-based discrimination. Figure 1A shows the effects of ethnic and immigrant origin on the contact rate. The results demonstrate that recruiters treat otherwise identical jobseekers who appear in the same search list differently depending on their ethnicity. Except for jobseekers from Southern Europe, ethnic minorities face a substantial penalty compared to Swiss natives. The penalty is 4.2% (95% CI: 3.4, 5.0) for candidates from Western and Northern Europe, 6.2% (95% CI: 4.8, 7.7) for those from Central and Eastern Europe, 6.4% (95% CI: 4.6, 8.2) for those from North and South America, 12.6% (95% CI: 11.4, 13.7) for those from the Balkans, 13.5% (95% CI: 12.2, 14.9) for those from the Middle East and North Africa, 17.1% (95% CI: 15.3, 18.9) for those from sub-Saharan Africa, and 18.5% (95% CI: 16.5, 20.4) for those from Asia.

These penalties are not only statistically significant but also substantial in economic terms.

Figure 1A shows that applicants with four or more years of work experience (the highest category) have a 11.5% (95% CI: 10.5, 12.5) higher contact rate compared to otherwise similar applicants with no experience. Thus, our results suggest that even acquiring the highest level of work experience is insufficient to offset the ethnic penalty faced by many immigrant jobseekers. Furthermore, Extended Data Figures 7a and 7b document that discriminatory hiring is similarly prevalent among different types of recruiters. Together, these results confirm and generalize the findings from correspondence studies that focus on select ethnic or racial groups in a few occupations[15]. They provide strong evidence that discrimination is a major factor in explaining ethnic minorities' inferior labor market outcomes and that ethnic penalties are larger for jobseekers from more marginalised immigrant groups[17].

**Leveraging variation in time**

Next, we illustrate how our approach can leverage information on when and for how long recruiters evaluate jobseeker profiles to shed light on some of the drivers of discrimination. One mechanism suggests that recruiters use lexicographic search[5], i.e. they only proceed to fully read the jobseeker's CV if the ethnic origin matches their heuristic, and immediately screen out jobseekers if it does not. In order to test for lexicographic search, we regress the log time spent on profiles on the ethnicity indicators, again controlling for our rich set of covariates as well as search fixed effects. Figure 1, Panel B presents the results. We find only very small differences in the time spent on ethnic minority profiles relative to natives: the largest difference is between applicants from Southern Europe and Swiss natives, but this is only 2.7% (95% CI: 2.2, 3.2), or less than 0.3 seconds. These

Figure 1: **Effects of jobseekers' ethnicity on contact rate and time on profile.** Panel A shows the effects of jobseekers' characteristics on the probability of being contacted (*n*=3,251,303). Panel B shows the effects of jobseeker characteristics on the log time that recruiters are looking at the profile (*n*=3,191,818). Dots with horizontal lines indicate point estimates with cluster-robust 95% confidence intervals from ordinary least squares regression. The hollow dots on the zero line denote the reference category for each jobseeker attribute. The effect of work experience provides a benchmark for the size of ethnic penalties on contact likelihood, and profile length for time on profile. Extended Data Table 4 contains the numerical estimates.

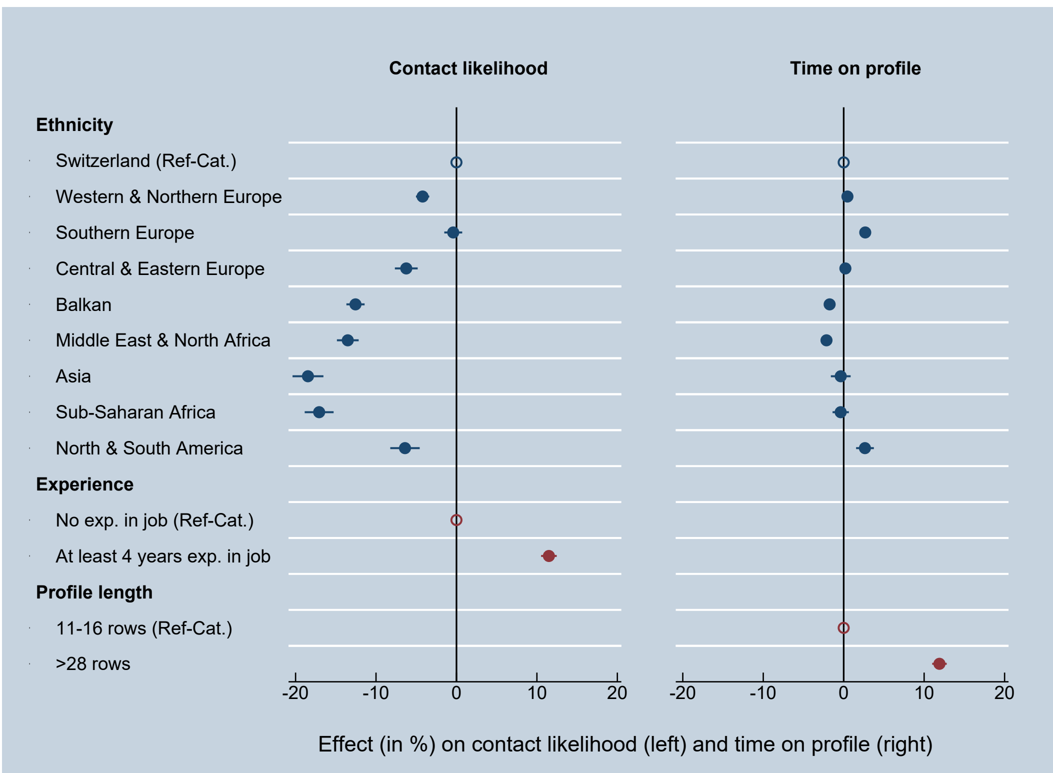

results therefore do not support theories that postulate that discriminatory recruiters use ethnicity as a shortcut to screen out applicants in this context.

Next, we explore how both time spent on profile and ethnic discrimination vary during the workday. To estimate the trajectory of ethnic penalties, we focus on searches conducted by the same recruiter and for the same occupation at different times of the day, while also controlling for all observable candidate characteristics. Figure 2A shows that the average time that recruiters spend evaluating a profile decreases from 10.5 seconds for searches conducted between 9–10am, when most recruiters start using the platform, to 9.7 seconds just before noon (11–11.59am; p-value=0.00, here and below always from a two-sided test against the corresponding coefficient from 9–10am). After the lunch break, average time is back to 10.8 seconds (p-value=0.01), but decreases again to 10.1 seconds (4-5pm; p-value=0.01) and 9.5 seconds towards the end of the workday (5-6pm; p-value=0.00). Figure 2B shows the corresponding evolution of ethnic penalties that non-European minorities face compared to natives. These ethnic penalties increase from 11.5% between 9–10am to 13.4% between 11–11.59am (p-value=0.08); are reset to 11.5% (p-value=0.97) after the lunch break, and rise again to 13.6% (4-5pm; p-value=0.07) and 14.7% (5-6pm; p-value=0.03). Figure 2C shows that the relative discrimination pattern for European minorities (compared to natives) is very similar albeit the differences in penalties are smaller in absolute size and less often statistically significant (SI Table 8 provides detailed estimates). The finding that ethnic minorities experience about 20% larger penalties later in the day when recruiters spend less time evaluating each profile resonates with theories of implicit bias[18–21], as well as research in psychology showing that sequential choices between alternatives depletes mental resources and

Figure 2: **Variation in ethnic penalties during the workday.** Panel A reports the average time recruiters look at jobseekers' profiles during the workday ($n$=3,281,297). Panel B shows the effect of ethnicity for non-European immigrants compared to natives during the workday. Panel C shows the effect of ethnicity for European immigrants compared to natives ($n$=3,341,209). Dots with vertical lines indicate point estimates with cluster-robust 95% confidence intervals from ordinary least-squares regressions.

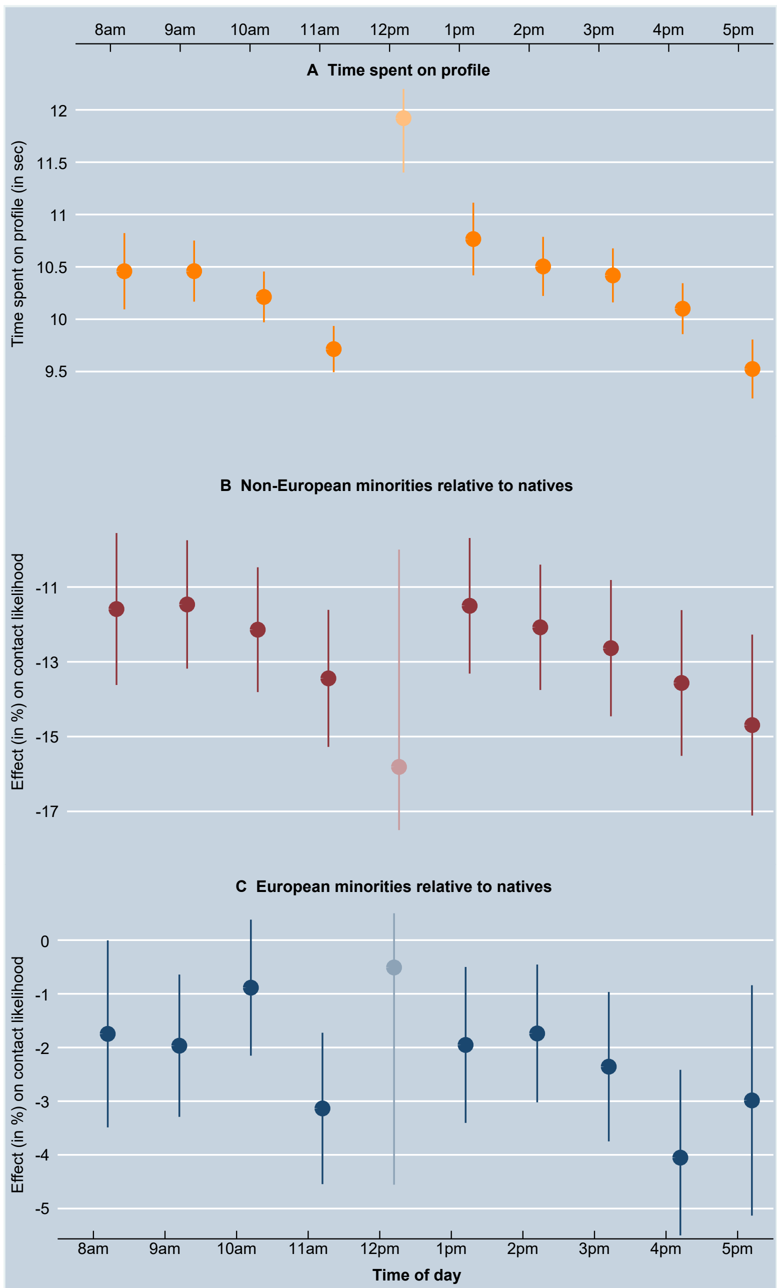

leads to an increase in intuitive decision-making in a range of contexts[22–24].

### Heterogeneity in gender discrimination

We also apply our approach to shed new light on gender discrimination. While prior research has found sizeable disparities between the wages paid to men vs. women[2,25], research on the gender gap in hiring decisions is more mixed[26]. In contrast to previous studies that focus on a limited number of occupations, our fine-grained data allows us to estimate occupation-specific gender penalties for all 323 jobs covered by the recruitment platform (see Extended Data Table 6), and thereby provide a comprehensive assessment of the heterogeneous nature of gender-based hiring discrimination in the labor market.

Figure 3 shows the gender penalty across occupations[27] plotted against the share of female jobseekers in the same search (Extended Data Figure 4a shows that the pattern is virtually identical if we use the share of women employed in these occupations). Consistent with previous research[8,11,28], we find that on average, women face no hiring discrimination compared to men. This average null effect, however, masks sizeable heterogeneity. While women face a penalty of 6.7% in male-dominated occupations (weighted average of gender penalties in the five occupations with the lowest female share in the search list), men face a penalty of 12.6% in typically-female occupations (weighted average of gender penalties in the five occupations with the highest female share in the search list). The relationship between the genderdness of the occupation and the (dis)advantage that women face can be approximated by a linear regression, which shows that a

Figure 3: **Gender-based hiring disparities across occupations.** The figure shows occupation-specific gender penalties in hiring (*n*=17,369,372). The circles show the (dis)advantage that women face compared to men in a given occupation ($y$-axis), plotted against the average share of female job seekers in the same search in that occupation ($x$-axis). The circumference of the circle is proportional to the the number of searches in each occupation. The color of the circles indicate the ISCO-1 level occupation classification. The solid black line shows the weighted least squares regression of the estimated, occupation-specific gender effect against the share of female workers in the searched occupation. The dashed black lines show 95% confidence intervals. SI Table 10 displays the underlying numerical estimates. While the linear trend is estimated on all data, the figure does not show the 50% hiring penalty for women in the occupation "skilled forestry, fishery and hunting workers".

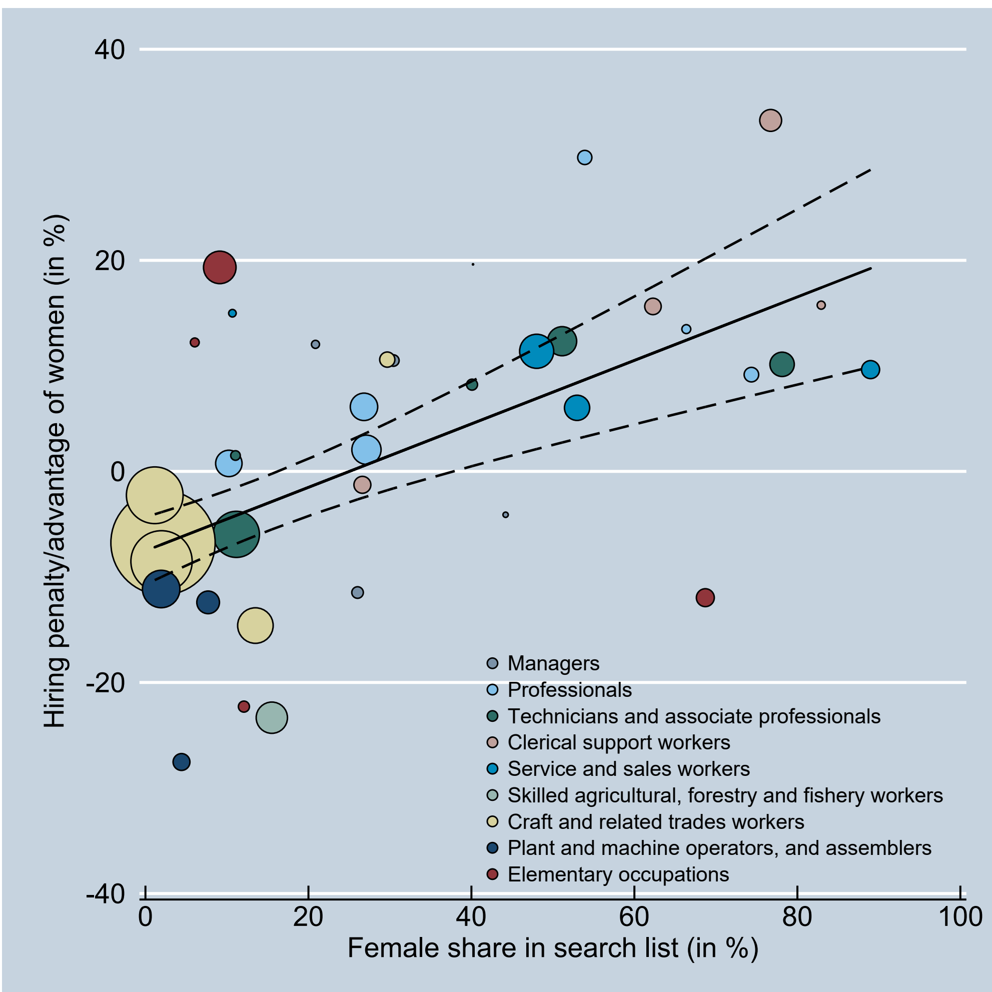

1% increase in the occupational share of women is associated with a 0.30% (95% CI: 0.17, 0.43) increase in the difference in the contact rate between women and men[29].

These findings can reconcile some of the conflicting results on the extent of gender discrimination from previous research. In line with a few prior studies (e.g.[28]), we find no average gender penalty for hiring decisions, but important variation across occupations. These findings are consistent with theories of gender stereotypes that infringe on the economic opportunities of both men and women[30] and suggest that firm recruitment reinforces gender segregation in the labor market.

## Discussion

While identifying the causes of discrimination is notoriously difficult [3,31], the data generated by our approach allow us to take some steps towards shedding light on potential drivers of labor market discrimination. First, the finding that more-marginalised immigrant groups experience greater ethnic penalties is consistent with psychological theories of ethnic hierarchies[32].

Second, several aspects of our results are harder to explain with economic theories of belief-based (statistical) discrimination[33,34] that postulate that recruiters are not prejudiced against ethnic minorities but instead act on imperfect information about the productivity of individual jobseekers[35]. To fully account for the documented pattern of ethnic hierarchies, unobserved skills would have to be lower for groups that are more severely discriminated against. Two tests suggest that this is unlikely to be the case. Extended Data Table 5 shows that ethnic minorities that experience more discrimination have often higher levels of productivity-related characteristics that are unobservable to recruiters (but observable to us researchers), including age and previous wage. When using these two variables together with the unobserved marital status to construct an occupation-specific "unobserved employability index", we find that ethnic penalties are very similar for jobseekers with low, medium or high unobserved employability (see Extended Data Figure 5a). We can also leverage the unstructured text field in the CV, which allows the case worker to provide further information on the skills and experience of the jobseeker. Extended Data Figure 5b compares otherwise similar jobseekers with and without additional information on skills and experience to test for statistical discrimination[3]. The results show that more information in this text field reduces ethnic penalties by up to 20% for jobseekers from Asia and Sub-Saharan Africa, while having no discernible effects for other ethnic minorities. Furthermore, the finding that men face a significant disadvantage in typically-female occupations is also difficult to square with standard belief-based explanations, which assume that women are more likely to leave the firm or take on family responsibilities[3].

Together, these results leave large scope for taste-based discrimination[36]. Yet, within-recruiter variation in hiring discrimination across the work-day suggests that conscious prejudice—which would be constant over a day—cannot entirely explain our findings either. This is further corroborated by the lack of evidence of lexicographic search, which would lead recruiters to immediately screen out minority applicants. Rather, the findings suggest that implicit biases and stereotypes[18,37] may also play a role in driving recruiter behaviour.

## Concluding remarks

The findings generated by our approach to studying hiring discrimination also have implications for policy. First, they corroborate previous studies that emphasise the role of discrimination in explaining labor market disparities between immigrants and native-born citizens, and men and women. Second, as employment websites are becoming an increasingly important channel for matching employers and workers in many sectors, our approach provides researchers and policy-makers with a non-intrusive, practical, and cost-efficient tool with which to continuously monitor labor market discrimination. Third, in addition to providing a laboratory to study hiring decisions, recruitment platforms may also be used to increase equal access to employment opportunities. For example, by redesigning the structure of the jobseeker's resume, productivity-relevant signals such as experience and skills can be strengthened, and other characteristics such as gender and ethnicity can be downplayed. Another design option are side-by-side comparisons of candidates, which have been shown to reduce discrimination compared to individual evaluations[38]. Lastly, employment websites may also be leveraged for targeted interventions, for example to identify occupations that experience high levels of discrimination and offer online implicit bias training for recruiters hiring in these occupations.

## Online content

Any methods, additional references, Nature Research reporting summaries, supplementary information, acknowledgements, peer review information; author contributions and competing interests; and statements of data and code availability are available at https://doi.org/10.1038/XXXXX.

1. Altonji, J. G. & Blank, R. M. Race and gender in the labor market. *Handbook of Labor Economics* **3**, 3143–3259 (1999).

2. Blau, F. D. & Kahn, L. M. The gender wage gap: extent, trends, and explanations. *Journal of Economic Literature* **55**, 789–865 (2017).

3. Neumark, D. Experimental research on labor market discrimination. *Journal of Economic Literature* **56**, 799–866 (2018).

4. Bartoš, V., Bauer, M., Chytilová, J. & Matějka, F. Attention discrimination: theory and field experiments with monitoring information acquisition. *American Economic Review* **106**, 1437–75 (2016).

5. Bertrand, M. & Mullainathan, S. Are Emily and Greg more employable than Lakisha and Jamal? a field experiment on labor market discrimination. *American Economic Review* **94**, 991–1013 (2004).

6. Bertrand, M. & Duflo, E. Field experiments on discrimination. In *Handbook of Economic Field Experiments*, vol. 1, 309–393 (Elsevier, 2017).

7. Baert, S. Hiring discrimination: an overview of (almost) all correspondence experiments since 2005. In *Audit studies: behind the scenes with theory, method, and nuance*, 63–77 (Springer, 2018).

8. Kuhn, P. & Shen, K. Gender discrimination in job ads: evidence from China. *The Quarterly Journal of Economics* **128**, 287–336 (2012).

9. For instance, the submission of fictitious applications involves deception and could plausibly interfere with the outcomes of real applications; it also increases the workload of human resources personnel. An alternative experimental design that in contrast to correspondence studies does not involve deception is proposed by [39].

10. Horton, J. J. & Tambe, P. Labor economists get their microscope: big data and labor market analysis. *Big data* **3**, 130–137 (2015).

11. Chan, J. & Wang, J. Hiring preferences in online labor markets: evidence of a female hiring bias. *Management Science* **64**, 2973–2994 (2017).

12. Besides [11], another closely related study is [40], who examine hiring decisions on the Swedish 'My CV' recruitment platform. In contrast to [40], we track recruiters entire search behavior on the employment website and observe the same information about the job seekers as recruiters do, which allows us to rule out confounding by unobserved factors.

13. Tibshirani, R. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Methodological)* **58**, 267–288 (1996).

14. Belloni, A., Chernozhukov, V. & Hansen, C. Inference on treatment effects after selection among high-dimensional controls. *The Review of Economic Studies* **81**, 608–650 (2014).

15. Extended Data Table 3 shows that the estimated penalties are somewhat smaller in size compared to a correspondence study fielded in the same labor market[16].

16. Zschirnt, E. & Fibbi, R. Do Swiss citizens of immigrant origin face hiring discrimination in the labour market? Tech. Rep., Université de Neuchâtel (2019).

17. Hainmueller, J. & Hangartner, D. Who gets a swiss passport? A natural experiment in immigrant discrimination. *American Political Science Review* **107**, 159–187 (2013).

18. Greenwald, A. G. & Banaji, M. R. Implicit social cognition: attitudes, self-esteem, and stereotypes. *Psychological Review* **102**, 4–27 (1995).

19. Chugh, D. Societal and managerial implications of implicit social cognition: why milliseconds matter. *Social Justice Research* **17**, 203–222 (2004).

20. As pointed out by an anonymous reviewer, this pattern might also be consistent with conscious discrimination if, for example, recruiters rationally decide to focus more on signals such as nationality when they are more in a hurry.

21. The finding that recruiters discriminate more just before noon and towards the end of the work day does not imply that they also become more discriminatory as the duration of the search session for a particularly hard-to-fill vacancy drags on (see Extended Data Figure 6a) nor that fast-deciding recruiters discriminate more than slow deciders (see Extended Data Figure 6b).

22. Muraven, M. & Baumeister, R. F. Self-regulation and depletion of limited resources: does self-control resemble a muscle? *Psychological Bulletin* **126**, 247–259 (2000).

23. Pocheptsova, A., Amir, O., Dhar, R. & Baumeister, R. F. Deciding without resources: resource depletion and choice in context. *Journal of Marketing Research* **46**, 344–355 (2009).

24. Danziger, S., Levav, J. & Avnaim-Pesso, L. Extraneous factors in judicial decisions. *Proceedings of the National Academy of Sciences* **108**, 6889–6892 (2011).

25. Olivetti, C. & Petrongolo, B. The evolution of gender gaps in industrialized countries. *Annual Review of Economics* **8**, 405–434 (2016).

26. Azmat, G. & Petrongolo, B. Gender and the labor market: what have we learned from field and lab experiments? *Labour Economics* **30**, 32–40 (2014).

27. Following the International Standard Classification of Occupations (ISCO).

28. Riach, P. A. & Rich, J. An experimental investigation of sexual discrimination in hiring in the english labor market. *Advances in Economic Analysis & Policy* **5** (2006).

29. Extended Data Figure 4b explores the relationship of the occupation-specific gender penalty with the average (log) wage of the occupation and shows that the correlation between average wages and hiring advantage for women is slightly positive.

30. Eagly, A. H. & Steffen, V. J. Gender stereotypes stem from the distribution of women and men into social roles. *Journal of Personality and Social Psychology* **46**, 735 (1984).

31. Pager, D. & Shepherd, H. The sociology of discrimination: Racial discrimination in employment, housing, credit, and consumer markets. *Annu. Rev. Sociol* **34**, 181–209 (2008).

32. Hagendoorn, L. Intergroup biases in multiple group systems: the perception of ethnic hierarchies. *European Review of Social Psychology* **6**, 199–228 (1995).

33. Phelps, E. S. The statistical theory of racism and sexism. *American Economic Review* **62**, 659–661 (1972).

34. Arrow, K. *et al.* The theory of discrimination. *Discrimination in labor markets* **3**, 3–33 (1973).

35. Note that similar to typical correspondence studies, our data does not allow us to distinguish between statistical discrimination based on accurate or inaccurate beliefs[41,42].

36. Becker, G. S. *The economics of discrimination* (University of Chicago press, 1957).

37. Greenwald, A. G. & Krieger, L. H. Implicit bias: scientific foundations. *California Law Review* **94**, 945–967 (2006).

38. Bohnet, I., Van Geen, A. & Bazerman, M. When performance trumps gender bias: Joint vs. separate evaluation. *Management Science* **62**, 1225–1234 (2016).

39. Kessler, J. B., Low, C. & Sullivan, C. D. Incentivized resume rating: Eliciting employer preferences without deception. *American Economic Review* **109**, 3713–44 (2019).

40. Eriksson, S. & Lagerström, J. Detecting discrimination in the hiring process: evidence from an internet-based search channel. *Empirical Economics* **43**, 537–563 (2012).

41. Fershtman, C. & Gneezy, U. Discrimination in a segmented society: An experimental approach. *The Quarterly Journal of Economics* **116**, 351–377 (2001).

42. Bohren, J. A., Haggag, K., Imas, A. & Pope, D. G. Inaccurate statistical discrimination. Tech. Rep., National Bureau of Economic Research (2019).

## Methods

**Online recruitment platforms as a laboratory for studying hiring decisions.** Extended Data Figures 1a–1c illustrate the search process on the Swiss government-affiliated online recruitment platform "Job-Room". First, the recruiter completes the search mask, typically by entering the occupation for which a vacancy exists and place of work, from which a ranked list of candidates who meet these criteria is generated. From this list, the recruiter selects the profiles, which provide detailed productivity-related information akin to a standard resume about the jobseekers' education, work experience, and skills, as well as gender and markers of ethnicity (registered users observe jobseekers' language skills, nationality, and name; unregistered users only observe language skills, see SI for details on the coding of ethnicity). The bottom of the profile view contains a button that reveals the contact information, typically a jobseeker's phone number, email address, and contact details of the employment office. Extended Data Figures 1a–1c also provide screenshots of the actual search mask, along with examples for result lists and profile views. Both the search process and the information contained in resumes on Job-Room is similar to those of many recruitment platforms in other countries (examples are provided in Extended Data Table 2).

**Data collection and outcome measures.** From March to December 2017, we collected on Job-Room data on 43,352 recruiters, 452,729 searches, 17.4 million profiles that appeared in the search list, 3.4 million profile views, the time that recruiters spent looking at each profile, and their decisions about whether to contact a jobseeker (see Extended Data Table 1 for details). About 80% of all jobseekers that are registered as unemployed in Switzerland have a profile on Job-Room, and Extended Data Table 6 shows this population covers about 65% (94% when weighted by employment) of all ISCO 4-digit occupations. The SI provides more details on the sample and the information collected. We use the click on the contact button as our binary outcome measure of request for a job interview. We validate this measure by linking the click data with the unemployment register. Extended Data Figure 2 shows that, controlling for jobseeker characteristics that are visible to employers, the jobseeker's ranking on the candidate list, and information on unemployment duration (the latter two variables are not observable to recruiters), every click on the contact button increases the jobseeker's probability of leaving unemployment within the next 90 days by 2.1% (95% CI: 1.7, 2.4). SI Table 7 shows that these effects are similar across origin groups. Furthermore, Extended Data Figure 3b shows that the estimated ethnic penalties are similar when using two more restrictive outcome measures. We also analyse the time that recruiters spend on the candidate profile (i.e. the logarithm of the time between opening a profile and either contacting the candidate or leaving the profile page). We top-code this outcome at 120 seconds.

**Statistical analysis.** We estimate the effect of ethnicity and gender on recruiter contact (or the time on profile) using ordinary least squares (OLS) regressions (see SI for detailed model specifications). SI Table 12 shows that the main results are very similar when using binary logistic regression instead of OLS to model contact. The regressions control for all 3,729 jobseeker characteristics (and their first-order interactions) that are predictive of contact or ethnicity (or gender; for the gender regressions the number of covariates selected is 1,927). We also control for search fixed effects and the list rank of the jobseeker in a given search (or, as a robustness test, an employability ranking relative to the other applicants in the same search, see SI Table 6). We therefore only compare profiles that appeared in the same search and control for all recruiter and search characteristics constant within a search. The SI provides more detail on the Lasso-based[13] post-double selection method[14] used to select the features that entered the regression. Because we can observe the same information that is available to the decision-maker in a given search and can control for all characteristics that are correlated with ethnicity (gender) and contact rate, our methodology allows us to identify the causal effects of immutable characteristics such as gender or ethnicity—and their interaction with a wide range of other jobseeker characteristics—on being invited to a job interview. In SI Table 3 and Extended Data Figure 3a, we use a series of placebo tests to validate this assumption and to confirm that productivity and ethnicity indicators that are not observed by recruiters do not predict contact.

**Acknowledgements** We acknowledge funding from the Swiss National Science Foundation (grant No. 162620). We are grateful to the Swiss State Secretariat for Economic Affairs for collaborating on this project and sharing their data. We are indebted to Severin Thöni and Matthias Bannert for programming support.

**Author contribution** D.H., D.K. and M.S. designed the research, D.H., D.K. and M.S. performed the research and analysed the data and D.H., D.K. and M.S. wrote the paper.

**Competing interests** The authors declare that they have no competing financial interests.

**Data availability** To obtain access to the anonymized data, researchers have to sign a data sharing agreement with the ETH Zurich, KOF Micro Data Centre, Leonhardsstrasse 21, 8092, Zurich, Switzerland and the Swiss State Secretariat for Economic Affairs, Labour Directorate, Holzikofenweg 36, 3003 Bern, Switzerland.

**Code availability** Code to replicate all analyses presented here is available at the publicly accessible Harvard Dataverse (doi:10.7910/DVN/) [link will be updated before publication].