

Evsyukova, Yulia; Rusche, Felix; Mill, Wladislaw

Working Paper

LinkedOut? A Field Experiment on Discrimination in Job Network Formation

CESifo Working Paper, No. 11433

Provided in Cooperation with:

Ifo Institute – Leibniz Institute for Economic Research at the University of Munich

Suggested Citation: Evsyukova, Yulia; Rusche, Felix; Mill, Wladislaw (2024) : LinkedOut? A Field Experiment on Discrimination in Job Network Formation, CESifo Working Paper, No. 11433, CESifo GmbH, Munich

This Version is available at:

<https://hdl.handle.net/10419/307363>

Standard-Nutzungsbedingungen:

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

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

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

Terms of use:

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

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

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

LinkedOut? A Field Experiment on Discrimination in Job Network Formation

Yulia Evsyukova (r) Felix Rusche (r) Wladislaw Mill

Impressum:

CESifo Working Papers

ISSN 2364-1428 (electronic version)

Publisher and distributor: Munich Society for the Promotion of Economic Research - CESifo GmbH

The international platform of Ludwigs-Maximilians University's Center for Economic Studies and the ifo Institute

Poschingerstr. 5, 81679 Munich, Germany

Telephone +49 (0)89 2180-2740, Telefax +49 (0)89 2180-17845, email office@cesifo.de

Editor: Clemens Fuest

<https://www.cesifo.org/en/wp>

An electronic version of the paper may be downloaded

- from the SSRN website: www.SSRN.com
- from the RePEc website: www.RePEc.org
- from the CESifo website: <https://www.cesifo.org/en/wp>

LinkedOut? A Field Experiment on Discrimination in Job Network Formation

Abstract

We assess the impact of discrimination on Black individuals' job networks across the U.S. using a two-stage field experiment with 400+ fictitious LinkedIn profiles. In the first stage, we vary race via AI-generated images only and find that Black profiles' connection requests are 13 percent less likely to be accepted. Based on users' CVs, we find widespread discrimination across social groups. In the second stage, we exogenously endow Black and White profiles with the same networks and ask connected users for career advice. We find no evidence of direct discrimination in information provision. However, when taking into account differences in the composition and size of networks, Black profiles receive substantially fewer replies. Our findings suggest that gatekeeping is a key driver of Black-White disparities.

JEL-Codes: J710, J150, C930, J460, D850.

Keywords: discrimination, job networks, labor markets, field experiment.

Yulia Evsyukova
Department of Economics
University of Mannheim / Germany
yulia.evsyukova@gess.uni-mannheim.de

Felix Rusche
Department of Economics
University of Mannheim / Germany
felix.rusche@uni-mannheim.de

Wladislaw Mill
Department of Economics
University of Mannheim / Germany
mill@uni-mannheim.de

October 2024

We thank Lawrence Katz, Andrei Shleifer, and five anonymous referees for their constructive comments, which greatly improved the paper. This paper benefited from discussions with Albrecht Glitz, Alex Imas, Amanda Pallais, Amelie Schiprowski, Amit Goldberg, Andrei Shleifer, Anke Becker, Antonio Ciccone, Arthur Schram, Ben Greiner, Cornelius Schneider, Christina Rott, Christine Exley, Christopher Roth, David Autor, David Rand, Drazen Prelec, Edward Chang, Edward McFowland III, Florian Heine, Frank Schilbach, Gianmarco Leon, Henning Hermes, Henrik Orzen, Iavor Bojinov, Ingo Isphording, Jean-Robert Tyran, Johannes Rincke, John Horton, Kai Barron, Katharina Brütt, Katherine B. Coffman, Klarita Gerxhani, Kevin Lang, Laura Gee, Lawrence Katz, Lumumba Seegras, Maria Balgova, Maria Petrova, Marina Kogłowski, Mashail Malik, Max Steinhardt, Menusch Khadjavi, Mohsen Mosleh, Oliver Falk, Philip Ager, Philipp Kircher, Raymond Fisman, Robert Livingston, Robin Ely, Ruben Enikolopov, Ryan Enos, Serena Does, Shaul Shavi, Simon Jäger, Simon Trenkle, Summer Jackson, Theo Offerman, Thomas Buser, Thomas Graeber, Vanya Georgieva, and Zoe Cullen. We also thank participants at several conferences and research seminars, including at the NBER Summer Institute: Labor Studies 2023, Advances with Field Experiments (University of Chicago), European Economic Association Annual Conference, European Association of Labor Economics Conference, Discrimination and Diversity Workshop (University of East Anglia), IZA, University of Jena, CRC Bonn/Mannheim, ETH Zürich, Pompeu Fabra University, University of Zürich, WU Vienna, CRC TR224 Young Researchers Workshop (Bonn/Mannheim), London Business School Transatlantic Doctoral Conference, German Economic Association Annual Conference, Leibniz Centre for European Economic Research (ZEW), Armenian Economic Association, CEREB Seminar Erfurt, ifo Institute Munich, HeiKaMaxY Workshop (Heidelberg), University of Mannheim, and University of Augsburg. Funding: Financial support from the German Research Foundation (Deutsche Forschungsgemeinschaft) through CRC TR224 (projects A01, A04, B02) and from the state of Baden-Wuerttemberg through bwHPC is gratefully acknowledged. F.R. also acknowledges financial support from the Joachim Herz Stiftung. Preregistration: The experiment's first and second stages were pre-registered on aspredicted.org ([#RDPZ67](#), [#8RRVLY](#)). Ethics approval: The study obtained ethics approval from the University of Mannheim's Ethics Committee (EK Mannheim 32/2021).

1. Introduction

[...] market-based explanations will tend to predict that racial discrimination will be eliminated. Since they are not, we must seek elsewhere for non-market factors [...] networks seem to be good places to start.

Kenneth J. Arrow, *Journal of Economic Perspectives*, 1998, p. 98

Around half of all jobs in the U.S. are found using information and referrals obtained through informal networks (Topa, 2011). Members of underrepresented groups rely on job networks as much as White individuals, but their networks include fewer high-status individuals with connections to high-paying jobs. They also provide less information, such as news about vacancies or insider perspectives (Fernandez and Fernandez-Mateo, 2006), which could help elucidate the worse labor market outcomes of underrepresented groups (Bayer and Charles, 2018; Coffman, Exley, and Niederle, 2021). Yet, existing research does not explain why networks of underrepresented groups are providing fewer benefits. Discrimination could play a pivotal role in the establishment and utilization of these networks. However, differences may also be confounded by other factors like self-selection and pre-existing inequalities, such as neighborhood segregation and socioeconomic background.

We causally investigate if and how discrimination affects the size, composition, and information provision of the job networks of Black Americans. To mimic real-world networks and their use, we conduct a pre-registered field experiment on LinkedIn – the world’s largest and most utilized online job networking platform with more than 900 million users (LinkedIn, 2023). Our field experiment consists of two stages. In the first stage, we build networks for 400+ fictitious profiles. We signal race (Black or White) solely via AI-generated profile pictures. In the second stage, we request job-related information from the networks formed in the first stage. Our novel research design allows us to resolve potential endogeneity in the networks formed in the first stage and separately identify discrimination in the second stage. Specifically, we can decompose discrimination in informational benefits into direct discrimination during Stage II and discrimination that occurs in Stage I.

A key feature of our field experiment is that we signal race exclusively through profile pictures. In particular, we create new AI-generated profile pictures and develop an algorithm that varies aspects of race inherently assigned by birth, like skin tone and facial features (i.e., the algorithm ‘morphs’ the race of a picture). To minimize behavioral responses due to stereotypes, the algorithm does not alter facial expression, hairstyle, clothing, or background. We validate our approach through an online experiment, which provides three main insights: 1) participants are not able to identify our pictures as fake, 2) the pictures clearly and

precisely signal race, and 3) the pictures of Black and White individuals are rated as highly comparable with regard to characteristics such as looks, authenticity, intelligence, etc.

Each profile in our experiment has a unique Black or White AI-generated profile picture, to ensure that our results are not driven by particular pictures. Further, each profile has a ‘twin’ of the other race with the same CV but a morphed profile picture. To make it realistic that our profiles joined LinkedIn only recently, their CVs represent them as young men who recently finished college and are otherwise similar to regular LinkedIn users. The profiles are furthermore assigned names that are both frequently used and racially ambiguous (e.g., Michael), such that race is exclusively signaled through the pictures.

To investigate how discrimination affects the formation of job networks, our profiles send a connection request to around 20,000 users during the first stage of the experiment. Each user receives requests from both a Black and a White profile with equivalent CVs and a time lag of four weeks between the requests. This experimental setup allows us to causally identify whether race affects the size of networks as twins differ only in their race and send connection requests to an identical number of users drawn from the same subject pool. Based on rich information gathered from users’ public CVs, we can further identify who discriminates.

We find a 13% lower connection acceptance rate for Black (23%) compared to White profiles (26%). In exploring who discriminates based on our rich set of individual-level characteristics, we observe discriminatory behavior to be widespread. In fact, there is little evidence of user groups that do not discriminate against Black profiles. That said, there is also substantial heterogeneity. Interestingly, men and older users show lower gaps in Black versus White acceptance rates relative to women and younger users, respectively. Additionally, we provide suggestive evidence that Black users discriminate as well, though to a lesser extent than non-Black individuals. Higher education and social status are only weakly associated with lower levels of discriminatory behavior. Gaps in Black versus White acceptance rates occur across almost all U.S. states. Within states, we find larger gaps for users who reside in more Republican counties.

In the second stage of the experiment, we assess the informational benefits of Black versus White job networks by asking the connections made in the first stage for advice. Importantly, our experimental design allows us to distinguish between disparities in informational benefits resulting from gatekeeping (Stage I) and discrimination in responses to information requests (Stage II). Before asking for advice, we swap half of the AI-generated Black profile pictures for White pictures and vice versa. As a result, half of the individuals who accepted the connection request of a White profile are asked for advice by a Black profile and, similarly, half of those who originally accepted a request of a Black profile are asked for advice by a White profile. The picture swap allows us to evaluate how much information Black and White

profiles would receive if they had access to the same networks. We also can examine whether swapping itself affects behavior, i.e., whether connections of swapped accounts notice the picture swap. Our results suggest that they do not, as we find no difference between swapped and non-swapped profiles in the number of views, profile blocking, connection dissolution, or types of responses. Moving to results on message response rates, we find no discrimination in responses if Black and White profiles are given access to the same networks. The zero result is precisely estimated and extends to messages’ content and usefulness.

Next, we assess the expected informational benefit provided by each profile’s network, accounting for the possibility of discrimination during both stages. Specifically, we estimate the expected number of responses for each profile, had they sent messages to their entire Stage I network. We find compelling evidence that the networks of White profiles provide substantially more informational benefits than those of Black profiles. Furthermore, differences in informational benefits emerging during Stage II can almost fully be accounted for by ‘gatekeeping’ in Stage I, as opposed to direct discrimination in Stage II. Our findings are consistent with models of rational inattention and potentially in-group preferences and are not driven by mechanical features of the experiment such as salience, exposure, etc.

This paper provides a number of new insights that expand and complement the existing literature. First, we provide **causal** evidence on discrimination in job network formation and information provision. Crucially, these informal networks have been shown to substantially benefit individuals’ careers (Pallais and Sands, 2016; Schmutte, 2015; Topa, 2011). To date, however, research on discrimination in network formation and information provision has largely relied on correlational analyses due to the challenges of causally studying a network formation process using observational data (Ioannides and Datcher Loury, 2004; McDonald, Lin, and Ao, 2009). Further, our two-stage design proved to be essential for identifying discrimination in outcomes, helping to disentangle sources of discrimination (Bohren, Hull, and Imas, 2022; Bohren, Imas, and Rosenberg, 2019).

In addition, our study advances correspondence studies’ methodology by introducing an AI algorithm to vary and signal race. The algorithm allows us to precisely and uniquely vary racial characteristics. Thus, it provides an alternative to relying on noisy proxies like names, which might convey unintended characteristics such as socioeconomic background (Bertrand and Mullainathan, 2004; Fryer Jr and Levitt, 2004; Gaddis, 2017), or skills and productivity (Abel and Burger, 2023; Kreisman and Smith, 2023). Our experiment further deviates from traditional correspondence studies (e.g., Acquisti and Fong, 2020; Agan and Starr, 2018; Kroft, Lange, and Notowidigdo, 2013) by examining discrimination in a novel setting characterized by a substantially more diverse target group, low decision-making costs, and targets that may desire network diversity for information benefits or virtue signaling (Angeli,

Lowe, Lowe, et al., 2023). We also contribute more generally to correspondence studies conducted on online platforms or social media (e.g., Ajzenman, Ferman, and Sant’Anna, 2023; Bohren, Imas, and Rosenberg, 2019; Doleac and Stein, 2013).

The first stage also adds important insight into who discriminates. We have the key advantage of observing individuals’ choices alongside a wide range of user characteristics. In contrast, classical audit and correspondence studies are typically conducted at the industry or firm-level (e.g., Kline, Rose, and Walters, 2022), while work focusing on individuals (mostly not in the context of labor markets) observe only a few individual-level characteristics, such as gender or race (e.g., Block et al., 2021; Edelman, Luca, and Svirsky, 2017).

Our second stage meanwhile shows that LinkedIn networks provide valuable information, thus contributing to the literature studying effects of professional networks on career outcomes (Cullen and Perez-Truglia, 2023; Gallen and Wasserman, 2021). The information provided in the second stage of the experiment also highlights the importance of weak ties (Gee et al., 2017; Gee, Jones, and Burke, 2017) and their role in providing valuable insights.

Finally, we contribute to recent theories on the nature of discrimination. In line with frontier contributions in the discrimination literature (Dobbie, Hull, and Arnold, 2022; Baron et al., 2024), we conduct a multistage experiment to understand where discrimination originates. We also combine the design with detailed information on the participants and a survey of LinkedIn users to contextualize the results in light of different theories of discrimination. The results are not compatible with simple theories of direct discrimination, such as taste-based and statistical discrimination, as usually considered in static correspondence studies. They are, however, well accounted for by rational inattention (Bartoš et al., 2016; Maćkowiak, Matějka, and Wiederholt, 2023) and potentially in-group bias (Akerlof and Kranton, 2000; Chen and Chen, 2011).

Overall, our study provides causal evidence on a previously understudied mechanism that helps to explain the worse labor market outcomes of Black individuals: the effect of discrimination on the size and information provision of job networks. Our two-stage experiment shows that differences in informational benefits emerge due to gatekeeping (i.e., during Stage I), rather than through differences in message response rates during Stage II. The findings furthermore offer crucial insights into potential ways to combat inequality in the labor market. Improving networking opportunities for Black individuals, e.g., through mentorship programs, could be an effective approach. Another could be diminishing the role of exclusive institutions such as ‘old boys clubs’ (Cullen and Perez-Truglia, 2023; Michelman, Price, and Zimmerman, 2022). Such steps would help improve information access for Black individuals and thus equitable access to job opportunities.

2. LinkedIn

With over 199 million U.S. users and 900 million worldwide, LinkedIn is the leading global online job networking platform ([LinkedIn, 2023](#)). Users create profiles that highlight professional experience, including work history, education, and additional customized information, such as skills and volunteer experience. The platform offers features for job hunting, networking, content sharing, and educational resources. Users build their professional networks by adding contacts or accepting connection requests. Firms also use LinkedIn extensively, creating profiles to post job openings, receiving applications, and using the platform for general promotion. Globally, 58.4 million companies have profiles ([LinkedIn, 2023](#)).

To shed light on why and how people use LinkedIn, we conducted an online survey with 500 U.S.-based LinkedIn users recruited through Prolific (see online Appendix I for details). Respondents are mostly motivated to use the platform for professional reasons, with 92% viewing LinkedIn as a job networking platform rather than social media. Moreover, in their rankings of the different reasons to use LinkedIn, they place all professional motives higher than any social ones. Job searching, networking, increasing one’s visibility to potential employers, and finding out which skills employers are looking for are ranked highest. In contrast, dating is ranked lowest. Regarding professional development, 69% consider LinkedIn useful or extremely useful for their career, and 53% agree or strongly agree that LinkedIn connections are useful for acquiring jobs. Regarding professional development, 53% agree or strongly agree that LinkedIn connections are useful for acquiring jobs, and 69% consider LinkedIn useful or extremely useful for their career. Specifically, between 60 and 82% indicate being likely or very likely to use the platform for purposes such as job searching, researching employers, contacting employees at a company of interest, or using LinkedIn’s ‘Easy Apply’ feature.

3. Experimental Design

To test the effects of discrimination on network formation and information provision, we begin by creating 400+ realistic profiles of early-career professionals on LinkedIn. To signal race, we create our own AI-generated pictures and develop an algorithm that transforms an AI-generated picture’s race while keeping other facial features stable. In the first stage of our experiment, these profiles develop networks by sending connection requests to 20,000 users. Each user receives requests from two statistically identical accounts, differing only in terms of race. In the second stage, connected users are asked for advice via a direct message, allowing us to explore potential gaps in the informational benefits of the resulting networks. Our design allows us to cleanly identify differences in the number of messages received as a

result of smaller or different networks (Stage I) or direct discrimination in message responses (Stage II). In what follows, we provide a detailed description of the experimental design.

3.1 Creating Realistic LinkedIn Profiles

This subsection explains how we brought our profiles "to life." That is, how we create profiles that resemble those of real users.

Basic Profile Features Each profile represents a male user born in the late nineties who recently graduated with a bachelor's degree in business administration and just started his first job. The beginning of one's career is a common time to start developing a professional network and opening a LinkedIn account. In fact, 54% of users in our survey report having opened their account during college or at an early career stage (see online Appendix I for more details on the survey).

Geography To increase external validity, we develop eight profiles in each of the 50 federal states and the District of Columbia. Profiles live in the largest city of the respective state (as shown in Figure 1), which ensures anonymity and reduces the chances of being identified as fake and blocked by LinkedIn.¹

Education Each of our profiles is a recent graduate of business administration, which is by far the most popular major among U.S. college graduates (Niche, 2019). Further, any type of firm – from hospitals to steel plants – employs individuals with business degrees, ensuring that we do not have to focus our study on a specific industry. Regarding the degree-granting institution, we assign each profile a college from their home state. To ensure that educational quality is comparable across states while avoiding additional signals through out-of-state education and experience, we refrain from assigning top universities, such as Harvard. Instead, we choose institutions offering business degrees from Niche.com's 2022 ranking of the 557 "Best Colleges for Business in America." We only assign universities ranked 70th and below. Details about universities chosen are provided in Section 3.3.

Jobs Profiles all have one year of work experience in one of five randomly assigned job titles. These include 'Buyer', 'Office Manager,' 'Administrative Assistant,' 'Marketing Assistant,' and 'Office Administrator.' The job titles are obtained from payscale.com by searching for entry-level jobs for bachelor graduates of business administration. Titles are chosen, given their suitability for early-career professionals and generality – i.e., almost any firm could employ someone with the titles above. All these positions are comparable regarding their skill level, with an average salary between 38,000 and 48,000 dollars. To fill the profiles with

¹Table A.2 in the online appendix provides a full list of the corresponding cities.

information, we also randomly assign them one job description.²

Employer Each profile is assigned an employer. We draw upon Statista’s Company Database to obtain the biggest employers in the U.S. (Statista, 2022), identify the 10 largest companies in each selected city, and randomly assign one of these to each profile. We choose large corporations as employers to make it less likely that ‘coworkers’ will encounter our profiles and realize they are fake. Moreover, larger employers are likely to have sufficient workforce turnover to remain anonymous, are more likely to have hired a recent graduate, and are very likely to hire business-related workers.

Names To avoid potential drawbacks of signaling race via names, first and last names of profiles are race-neutral. For first names, we focus on names that appear among the 50 most common names for both White and Black men (i.e., the intersection of popular Black and popular White names). We sort these by the relative popularity among Black Americans and take the 10 most popular names. The selected names are all among the top 30 names. For last names, we draw on race shares by last names from U.S. Census Bureau (2022) and choose those that are roughly equally likely to be of a Black and White individual and unlikely to be of any other race. We again pick relatively common names.³

Additional Details on Profiles To make profiles more realistic, we add appropriate details. In particular, LinkedIn allows users to signal skills such as ‘Teamwork’ or ‘Bookkeeping’. We create a collection of skills, drawing on LinkedIn’s 20 most commonly reported skills for each of the given job titles. From these 20 skills, we randomly assign five relevant skills to the specific job of each profile.

To further round out the profile with information, we add past volunteer experience. We chose organizations that are very popular in the U.S., non-partisan, and present across the country: ‘Big Brother and Sister,’ ‘Red Cross,’ and ‘Crisis Text Line.’ All of these organizations are within or close to the biggest city in a given state or can be contributed to remotely, ensuring that we do not have idiosyncratic results due to very specific volunteering experiences. In addition, all of these experiences do not require special skills, thus avoiding any signaling of differential information.⁴

²See Tables A.7 and A.8 in the online appendix for details.

³To obtain first names, we rely upon the most common first names of men born in 1997 in Georgia (Georgia Department of Public Health, 2022). Table A.5 in the online appendix provides an overview of all first names and their popularity. It also shows the rank of the first name for all baby names across the U.S. in 1997. Table A.6 shows the surnames, race shares, and rank across the U.S.

⁴Table A.9 in the online appendix lists the relevant skills by job title. Table A.10 provides an overview of the volunteering experience. It also includes descriptions of the tasks we created based on real profiles.

3.2 Varying Race via AI-Generated Pictures

The key variation in our study is race. To signal race, we create AI-generated pictures and develop an algorithm that transforms the picture’s race while holding other characteristics stable. The approach has two important features to account for ethical concerns. First, all pictures are AI-generated, thus avoiding privacy issues. Second, and more importantly, we do not define race characteristics ourselves. Rather, we take an agnostic approach. The transformation algorithm is defined as follows: we take all images of young Black men found among the 100 thousand images provided by StyleGAN2 (Karras et al., 2020). We translate these images into multidimensional vectors and do the same for a comparable number of White images. Next, we calculate the average Black and White image vectors and take their multidimensional vector difference. The transformation algorithm then simply adds this difference vector to a White image or subtracts it from a Black one.⁵

We conduct an experiment on Amazon’s MTurk ($n \approx 500$) to validate the pictures along a number of dimensions. In the first step, we test whether participants perceive the profile pictures to depict real humans rather than computer-generated ones. For this, participants are shown 20 pictures in a style that resembles a Google Captcha. They are told to select all computer-generated images, and given a monetary incentive to click on the correct ones. Among the pictures shown, ten are our A.I-generated images while another six depict real humans. The real human pictures are chosen from the set of training images of StyleGAN2 to fit the age, race, and gender category of our images. An additional four pictures show obviously computer-generated images, i.e., with strange hats, deformations, or unrealistic facial features. The results indicate that our White and Black images are *not* perceived as more likely to be computer generated (12% and 14%, respectively) than the images of real humans (15%), while obvious fakes are correctly identified in 84% of cases. These results align with a recent study by Nightingale and Farid (2022), which suggests that good AI-generated pictures are indistinguishable from real faces.

Following this exercise, each participant rates ten of our pictures along a number of dimensions. Importantly, both Black and White pictures are associated with the targeted race and gender. In addition, our results suggest that, despite potential biases of participants, pictures are rated similarly across several additional dimensions, including trust, appearance, authenticity, and intelligence, thus providing reassuring evidence that our algorithm keeps picture characteristics stable.⁶

⁵Online Appendix B describes the picture creation and the transformation algorithm in greater detail.

⁶See online Appendix F for more details, and Section F3.2 for the detailed results.

3.3 First Stage – Network Creation

The experiment’s first stage consists of building networks of our profiles to measure discrimination in network formation. Drawing on detailed information on individual user characteristics, we additionally identify which are most predictive of discriminating behavior (pre-registration: [#RDPZ67](#)). Below, we describe how we choose the users to whom we send connection requests.⁷

Creating (Twin) Profiles Using the characteristics described above, we create four ‘twin pairs’ in each U.S. state (and D.C.). A twin pair consists of a Black and a White profile with the same CV. Twins differ *only* in terms of their race, as signaled via the AI-generated profile picture. While they also have different names to ensure the profiles are not detected by LinkedIn, these are randomly assigned and do not signal race. Importantly, within each state, the characteristics above, such as names, job title, picture, etc. are drawn without replacement. As a result, profiles *within* a twin pair are identical except for the picture. Meanwhile, profiles *across* twin pairs differ in pictures and almost every aspect of the CV.

Collection of Targets After creating the profiles and before sending any connection requests, we identify relevant users with whom to connect, i.e., ‘targets.’ To this end, we collect roughly 150 contacts recommended by LinkedIn for each of our profiles. Drawing on these initial platform suggestions rather than, e.g., a random sample of all LinkedIn users in the U.S., has two advantages. First, these suggested contacts tend to be geographically relevant, i.e., they live close to our profiles. Figure 1 shows the locations of our profiles and selected targets as indicated on their LinkedIn profiles and geolocated using *Google Maps API*. We observe a median distance of 14km from our profiles. Second, they are professionally relevant; that is, they work in a similar industry or related job, attended the same university, or hold a related degree. Overall, following the LinkedIn suggestions closely mimics the construction of job networks of real users, as the platform aims to propose professionally relevant connections. After collecting the initial suggestions, we pool all of them by state (that is, over the eight profiles) and identify their race and gender based on their profile pictures and names.⁸ We then draw on these characteristics to create four exclusive pools per state with

⁷A timeline of our experiment is provided in Figure A.1 in the online appendix. In the pre-registration, we had initially planned to exclude LinkedIn users without profile pictures. During the construction of the target pool, it became apparent that a significant number lacked a profile picture. Excluding them would have severely reduced our target pool and induced potential selection bias. At the same time, the algorithm inferring demographic features from pictures turned out to be less reliable than anticipated, leading to our decision to instead rely mainly on first and last names to infer gender and race. We therefore included accounts without pictures but excluded individuals without a first name from the sample.

⁸This is done based on U.S. census and social security data on first and last names ([Kaplan, 2022](#)). For pictures, we use a face recognition software ([Taigman et al., 2014](#)).

96 targets each. Across all pools, we balance on gender to ensure that half of the targets are women. We further balance the pools in terms of race shares to have sufficient data on the behavior of underrepresented groups. For each state, we thus obtain four balanced pools (roughly 50-50 in terms of gender, and 70-30 in terms of White vs. non-White). Finally, we *randomly* assign targets to one of the pools, meaning that target characteristics are comparable in expectation across pools. Importantly, through the random assignment of the initial LinkedIn suggestions to our profiles, we resolve any potential endogeneity resulting from the initial suggestions of the algorithm.




Figure 1: Locations of profiles and targets.

Note: Profile locations and city names show the places where our profiles reside. Target locations represent unique geolocations using Google Maps API based on targets’ self-reported location. Each location includes one or multiple targets. The figure includes the number of targets by Core Based Statistical Area (CBSA). In cases where a given county does not belong to a CBSA, county borders are displayed instead.

Sending Connection Requests Next, we use our profiles to connect to targets. Specifically, each target receives two requests: one from a Black and one from a White profile.⁹ This leads to a sample size of 19,584 target profiles per group, allowing us to detect even

⁹We follow this strategy for internal validity and causal identification. As argued above, reaching out to users one does not know is very common on LinkedIn and the platform’s algorithm is tailored to suggest relevant connections. In general, however, decisions about whom to send a request to may be driven by other factors that we explicitly exclude, such as professional status, homophily, or behavioral reactions in anticipation of discrimination.

effect sizes of an order of magnitude smaller (Cohen’s D of 0.028) than what is usually considered small (i.e., a Cohen’s D below 0.2, see [Sawilowsky, 2009](#)), with high power (80%). In requests, profile pictures – and hence the race of profiles – are very salient. The pictures always appear next to the request, including, but not limited to an email the user may receive informing her of the request and on LinkedIn itself, where users can choose to ignore or accept a request (see example screenshots in Figure A.3 in the online appendix). As might be expected on a platform aimed at professional networking, receiving connection requests from users one does not know is very common. In our LinkedIn user survey, 63% report receiving such requests a few times a month or more often. That said, sending a message along with requests appears to be less common (only around 22% of surveyed LinkedIn users state that connection requests contain a message at least ‘most of the time’ while 55% say sometimes or never; see online Appendix I). For this reason, we send connection requests without an accompanying message.

We would ideally contact each target with two profiles that differ only in their race, i.e., two profiles from the same twin pair. Since this would likely raise suspicion, we instead send each user a request from two profiles from different twin pairs. Doing so ensures that the requests received by a single user are sufficiently different, given that the profiles differ in almost every aspect of their CV, including their underlying profile picture. In addition, by drawing all profile characteristics from the same distribution (without replacement), we ensure that the profiles sending the request are statistically identical.

Figure 2 shows the requesting procedure of two twin pairs A and B. Each consists of a Black and a White account. Potential contacts are two mutually exclusive pools of 96 targets each, Pool 1 and 2. All targets in the first pool receive requests from James and Joshua. Those in the second pool receive requests from Michael and Tyler. By using this approach, we can ensure two things. First, given that the pools are balanced and randomized, both twins contact people who are, in expectation, the same, allowing us to account for twin-fixed effects, keeping everything but race stable. Second, contacting targets using two profiles in combination with information on target characteristics means we can draw conclusions about who discriminates at the individual level. Specifically, we can observe which characteristics predict a higher acceptance rate gap between Black and White requests.

Receiving two requests from unfamiliar accounts at the same time may raise suspicion. To mitigate this concern, we implement a four-week lag between the first and the second request. The lag is introduced by contacting a subset of 12 targets per week and profile, hence running the experiment over a period of eight weeks. While both profiles contact the same pool and send the same number of requests each week, a given target receives the second request with a lag of a month. Sending a limited number of requests per week has

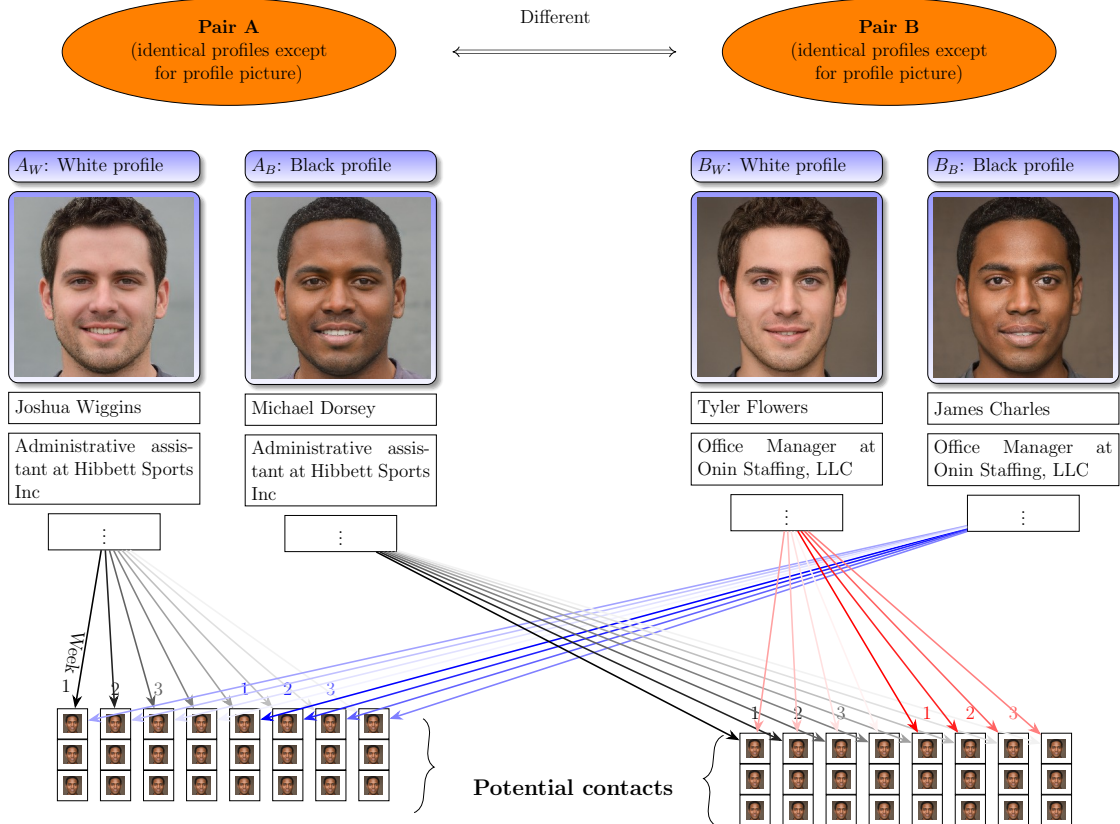


Figure 2: Requesting procedure

the additional advantage of reducing the chances of our accounts being blocked by LinkedIn.

Additional Variation While the key variation in our experiment is race, our design also allows for the analysis of variation across a number of additional dimensions. Specifically, we can explore heterogeneity in discriminatory behavior between users with different characteristics as well as measure differences in discriminatory behavior between geographic regions.

In line with existing studies (e.g., [Oreopoulos, 2011](#)), we also vary profile quality by assigning half of the profiles within a given city to a higher- and half to a lower-ranked university. High types attended an institution ranked 70-270 in [Niche.com](#)’s listing, while the universities of low types are not ranked but are present in the “Business and Management” category. In states that do not offer a suitable university, the institutions are chosen from a neighboring state for both types.¹⁰

Building our networks over eight weeks allows us to study dynamic effects. Initially, all

¹⁰We verify the quality signal in two ways. First, we ensure that Niche’s ranking is consistent with other popular rankings, such as Forbes and USA Today. Second, we verify the perceived ranking through our MTurk survey by asking individuals to identify the better-ranked universities (see online Appendix F). On average, subjects are able to correctly differentiate between the high- and low-ranked universities. Tables A.3 and A.4 in the online appendix provide details on the selected universities.

profiles start without connections. Through the staggered design, we can assess whether acceptance rates change as networks grow. We can also study differences in acceptance rates for profiles with and without connections, and hence mitigate concerns of starting our profiles without initial connections. Although starting without an initial network poses no threat to identification as the same holds for both Black and White profiles, users may find this odd. We nonetheless decided to start with zero connections as starting with an established network would only have been possible by creating a large number of fictitious profiles to befriend ours, increasing the risk of being blocked. Alternatively, we could have built an initial network with real users, which would, however, have made the setup less clean, as the network composition would have differed between our profiles, making dynamic effects likely. Zero contacts thus seemed like a valid starting point. Our survey of LinkedIn users supports this choice, showing that the majority of respondents opened their accounts at an early career stage or in college. Further, in an open-ended survey question, only very few users mention a small network as a ‘red flag’, suggesting that this is not something users are typically worried about. Their main concerns are suspicious CV or profile entries, strange or no profile pictures, and suspicious posting behavior (see Figure I.14 in the online appendix).

Profile Gender as a Limitation An important limitation is that we do not vary gender. While this choice has some obvious disadvantages, there are several reasons for this choice. First, given the technical difficulties of running a large-scale experiment on LinkedIn, we focus on varying just one dimension (i.e., race), keeping everything else constant. Adding female profiles would have doubled the experiment’s size. Responses to women’s requests may furthermore follow a different logic and thus necessitate an adjusted experimental setup to interpret results. Previous research shows, in fact, different reactions to online activities by men and women (e.g., [Bohren, Imas, and Rosenberg, 2019](#)). For instance, young women report being sexually harassed online much more frequently ([Vogels, 2021](#)). Varying treatment across more than one dimension would indubitably complicate the interpretation of results. Finally, a technical issue is that the morphing of pictures is more error-prone for female pictures, and the baseline sample of Black women is relatively small. While we believe studying the effects of job network formation for women is equally important, interpreting and studying the effects for women would require an adjusted experimental setup, warranting a dedicated independent study.

3.4 Second Stage – Information Provision

After finishing the first stage of the experiment, we explore whether valuable information can be obtained through online networks and whether the informational benefit differs between ‘Black’ and ‘White networks’ (pre-registration: [#8RRVLY](#)). LinkedIn allows users

to contact each other through private messages, which is a common approach when seeking employment-related information. In our LinkedIn user survey, 59% report having received messages at least a few times a month or more, while 51% report having received messages from users they did not know at least a few times a month (see online Appendix I).

Sample Selection We start by selecting eligible users. Since we want to investigate the value of the networks, the second stage of the experiment includes only those users who had accepted at least one contact request by the end of Stage I. In addition, the selected users meet certain criteria. Given that we ask questions about the job application process, we exclude all users for whom we do not have information on the company they work at or who are retired, self-employed, freelancers, or unemployed. To avoid raising suspicion, we also exclude users who work in companies with less than 50 employees. For such small companies, it is likely that no relevant positions are vacant or exist, which the user may be aware of given the company’s size. By design, we focus on users who are willing to engage in networking. While this constitutes a natural constraint of networks, it also means that users who accepted our requests may generally be more open-minded than those who rejected both requests. As a result, our findings are limited to such users. We can, however, distinguish between users who only accepted Black, White, or both requests during in Stage I.

Next, if a suitable user accepts only one of the requests of our profiles in the first stage, she is contacted by that profile. If she accepts both requests, she is randomly and with equal probability contacted by either the Black or White profile. After allocating the targets to our profiles, we exclude all users who work at the same company as the respective profile. We also exclude individuals who sent a message to our profile before the beginning of the second stage, as messaging such users without having answered their previous message might be perceived as rude or suspicious behavior. Each of our 400 (still active) profiles contacts up to 10 unique suitable targets from their network, with each target being contacted only once.¹¹ For profiles with more than 10 suitable contacts, we randomly select 10 targets to receive a message. If there are fewer than 10 suitable connections in the profile’s network, it contacts all of them. Overall, we sent 3,319 messages, which allows us to detect even very small effect sizes (Cohen’s D of <0.1) with high power (80%).

Message Variation To mimic the use of job networks along more than just a single dimension, and to reduce the risk that the results are driven by idiosyncratic message features, we design two messages. In one message, the profile signals to targets that they are interested in applying for a position at their company and asks for information about both the company

¹¹ During the first stage of the experiment, 8 of our original 408 profiles (i.e. $<2\%$) were blocked. This happened within the first few weeks of the experiment’s first stage and without any visible pattern as to why, e.g., regarding their race or location.

and the application process. We design the message such that it is sufficiently generic to fit any contacted user and, simultaneously, asks for information not contained in job postings (such as potential pitfalls and tips for the application process). In the other message, we ask targets for general career advice, given that our profiles just recently entered the labor market. Both messages are randomized at the level of our profiles, meaning around half of the contacts receive the general career advice message and the other half the job application message. The messages are shown in online Appendix A8. As in the first stage, race is very salient: the profile picture always appears next to the message (see example screenshots in Figure A.3 in the online appendix).

Resolving Endogeneity Concerns The composition of networks of Black and White profiles obtained in the first stage might be quite different in terms of their characteristics. For instance, users in Black networks might be less discriminatory and more responsive to messages. Thus, if we were to simply contact users within profiles’ networks, the results could be driven by (1) differences in networks originating in Stage I and (2) differences in response rates to messages from Black versus White profiles. Ideally, we would want to ensure that Black profiles have connections in their network who would typically rather accept White and not Black profiles; that is, eliminating the differences from (1) and only observing differences due to (2).

To disentangle these effects, we exploit a feature of our experiment, namely, the fact that twin pairs differ only in their race signaled through their picture. Instead of matching on observables, we put Black profiles into the network of people who would typically accept a White profile and vice versa. We achieve this by simply swapping the picture of Black profiles with their White twin’s picture and vice versa. Thus, people who have accepted a connection request from a White profile are now faced with a profile that is Black instead of White. Similarly, people who accepted a connection request from a Black profile now face a White profile. We make this switch two weeks after the end of Stage I for 200 of the 400 remaining profiles (i.e., 200 profiles retain their original profile picture), resulting in half of our Black profiles now having a White network. Similarly, half of our White profiles now have a Black network. Our approach has several advantages. First, it equalizes access to networks between Black and White profiles. As a result, on average, Black and White accounts have the same networks, allowing us to explicitly study direct discrimination during the second stage of the experiment. Second, combining insights from the first and second stages allows us to calculate total differences in expected informational benefits obtained through the profiles’ original networks. Specifically, we can estimate the expected total number of messages obtained from the first- and second-stage results. Finally, half of our profiles remain in their original networks, allowing us to explore whether the swapping itself

is detected by users, e.g., whether swapped accounts lose more contacts or are visited more frequently after the swap (we find no evidence of any behavioral changes due to swapping, see online Appendix G9 for a detailed discussion).

Expert Survey To compare our findings to the priors of experts working on labor economics and/or discrimination, we conducted an expert survey following our experiment. We reached out to 2,171 labor economists from the Institute of Labor Economics (IZA) network and participants of the NBER Summer Institute: Labor Studies (2021-22). In the survey, we briefly present experts with the key features of our experiment and ask them to predict the results (see online Appendix H for more details). Overall, 269 experts completed our survey.

3.5 Ethics

We briefly discuss the main ethical aspects of our study in light of the considerations put forward by Salganik (2019) in the context of correspondence studies (online Appendix E provides an extended reflection). First, field experiments should limit harm, i.e., costs to participants. The costs our experiment imposes are very low, especially compared to classical correspondence studies. They consist of answering a connection request and (at most) voluntarily answering a short message. Second, costs should be evaluated against “the great social benefit of having a reliable measure of discrimination” (p. 304). Given that, to our knowledge, this is the first paper to causally study discrimination in professional networking, our experiment provides important insights that can help explain differences in labor market outcomes between Black and White Americans. A third consideration is “the weakness of other methods of measuring discrimination” (p. 304). Existing observational data does not allow for the causal study of discrimination in job networking given major issues such as self-selection, omitted variable bias, and endogeneity (Fernandez and Fernandez-Mateo, 2006). Laboratory experiments meanwhile likely suffer from concerns of external validity. A field experiment thus presents the best solution. Running a field experiment is not without difficulty though, as LinkedIn is keen on preventing the creation of fake accounts. While this is beneficial to users, researchers must, amongst others, circumvent captchas, use proxy servers, and bypass phone and email verifications for each account. Overall, we argue that the social benefits of a reliable estimate of discrimination in the formation and information provision of job networks outweigh the low costs imposed on participants.

4. Results

In discussing our results, we first address whether discrimination is present in the formation of job networks and investigate heterogeneity in discriminatory behavior (Section 4.1). We then turn to informational benefits and disentangle differences due to discrimination originating

in Stage I and direct discrimination in responses in Stage II (Section 4.2). Finally, we discuss the mechanisms that may explain our results (Section 4.3).

4.1 Formation of Job Networks

4.1.1 Differences in Network Size

As Black and White profiles share the same observable characteristics and differ only in the racial signal conveyed through their picture, we can causally identify the impact of discrimination on the formation of job networks. Figure 3 reveals a clear difference between Black and White profiles in the number of contacts obtained by the end of the experiment’s first stage. White profiles have about 3 more connections than Black profiles – a considerable difference given a baseline of about 23 connections. In relative terms, White profiles have approximately 13% larger networks than Black profiles. This result is just slightly below the prediction of experts who expect a gap in connections of around 18% (see Chapter H in the online Appendix for more details).

Aside from visualizing differences in means, the graph further shows the number of connections obtained by each profile, as indicated by the blue and orange dots. Each dot is connected to the twin of the other race. The raw data provides two insights. First, there is substantial heterogeneity in the number of connections obtained by a given profile. Second, most lines (66%) are upwards-sloping and hence colored in red, suggesting that most White profiles made more connections than their Black twins. Only 29% are downward-sloping and colored blue. This is despite the fact that the only difference between the two profiles is their race, as signaled by the profile picture. These observations are confirmed using common regressions reported in Table 4.¹²

An advantage of our design is the possibility to study dynamic effects, geographical differences, and differences between high- and low-quality profiles. We observe that Black profiles are disadvantaged from the first week onward. While the absolute gap in connections widens over time, the relative difference remains stable at roughly 13%. This suggests that White profiles are not perpetually improving, but Black profiles are also unable to catch up over time. As an alternative way of looking at the dynamics, we study how the White-favoring gap in acceptance rates changes as a function of the size of the profile’s network. Here again, we observe a comparatively constant gap for any number of connections, suggesting it persists irrespective of the users’ network size (see online Appendix G1 for more details). Regarding geographical variation at the state level, we observe that Black profiles’ disadvantage appears in all but six states (see online Appendix G2 for more details). Finally, the Black-White

¹²Table J.1 in online Appendix J replicates this estimated difference while accounting for a variety of profile picture characteristics (e.g., rated looks, trust, etc.).




Figure 3: Number of contacts at end of the first stage by race of the profile

The figure depicts the number of contacts obtained individually by Black and White profiles by the end of Stage I. White profiles are depicted by orange dots and Black profiles by blue ones. Each dot represents a single profile, and twin pairs are connected by colored lines (where red vs. blue lines denote a gap in favor of White or Black profiles, respectively). Whiskers around the mean denote the corresponding 95% confidence intervals. T-tests are used to obtain the following significance levels: $p < 0.10$; $p < 0.05$; $p < 0.01$; $p < 0.001$. Table J.1 reports the corresponding regressions, and Figure G.4 further displays differences by profile quality (both in the online appendix).

gap is remarkably similar for profiles with a more/less prestigious university background (see online Appendix G3). Overall, discrimination faced by Black profiles is immediate, stable, and geographically ubiquitous.

4.1.2 Differences in Network Composition

Beyond the substantial gap in the number of connections, the networks of Black and White profiles might differ in their composition. Of the 6,213 users who accepted at least one connection request (about 33% of all contacted users), almost half accepted only one request, with 29% accepting only the White profile and 20% accepting only the Black profile. The network of Black profiles is thus not merely a subsample of the network of White profiles.

To compare the composition of Black and White networks, we enrich our data with detailed information from users' public CVs. The data includes estimates of race and gender from names (U.S. Census Bureau, 2022; U.S. Social Security Administration, 2022), links between reported education and college statistics (Forbes, 2021; IPEDS, 2022), the classification of degrees to infer age, and links to employer information through platform employer data. In addition, we assess salaries using job titles to find matches on salary estimation websites (glassdoor.com). We also geo-locate self-reported locations using Google Maps

API and match these with county/CBSA shapefiles (U.S. Census Bureau, 2020), as shown in Figure 1. Based on counties, we connect users to local vote shares (MIT Election Data and Science Lab, 2018), county-level demographics (Hopkins Population Center, 2020), and other county-level variables (Chetty et al., 2022a,b; Xu et al., 2022).¹³

A comparison of users who accepted either the connection request of the Black profile or the White profile reveals several differences between the resulting networks (Table 1).¹⁴ First, regarding the gender composition, Black networks have a substantially higher fraction of men. Second, users in Black networks are older and a bit better educated. Third, the connections of Black profiles are more engaged, in the sense that they have more contacts themselves, more followers, and post and share more than the connections of White users. Thus, although the networks are comparable, there are clear differences in their composition.¹⁵

4.1.3 Heterogeneity: Who is (Most) Discriminating?

The above results raise the question of who discriminates. In contrast to most correspondence studies, our design produces a rich set of characteristics from around 20,000 targets. We can use these to investigate which characteristics are most predictive of discriminatory behavior, i.e., result in greater gaps in acceptance rates.

Given the vast number of characteristics, we first focus on those that are most obvious: age, gender, job position, share of Republican votes in the home county, race, and education. The first five were explicitly pre-registered.¹⁶ Table 2 shows regression results on these characteristics. The first two rows show the total gap in acceptance rates for users with and without the respective characteristic. The last row displays the difference between these coefficients and its level of significance. All regressions include profile picture and target random effects. All results remain true conditional on the other observables (see Table J.10 in the online appendix).

We find all user groups to discriminate and favor White requests, as reflected by the exclusively negative coefficients in the first two rows. However, the degree of discrimination substantially varies by characteristics. Targets with higher job positions and better education show (slightly) less discrimination, though differences between these users and their counterparts are insignificant (Col. 4-5). Residents of more Republican counties are

¹³Online Appendix C provides information on the precise data preparation process. Online Appendix D provides summary statistics and a comparison with official statistics.

¹⁴In online Appendix G4, we further discuss predictors of connection request acceptance in detail and show the difference between Black and White networks if we include people who accepted both requests (Table J.5).

¹⁵In Section G11, we explore additional outcomes related to the value of profiles’ networks. We show that Black profiles have around 20% fewer profile visits and receive fewer unsolicited messages than White profiles.

¹⁶The pre-registration mentions the city/state-level Republican vote share. Given that we can observe people’s self-reported location, we draw on county-level data here.

Table 1: Differences in networks (Black vs. White) of users who accepted only one request

Network characteristics of			
	Black profiles (N=1223)	White profiles (N=1821)	p-value
USER DEMOGRAPHICS			
<i>Female (First Name)</i>	0.43 (0.50)	0.53 (0.50)	$\leq 0.001^{***}$
<i>Black (Last Name)</i>	0.07 (0.25)	0.05 (0.22)	0.131
<i>White (Last Name)</i>	0.73 (0.45)	0.72 (0.45)	0.898
<i>Asian (Last Name)</i>	0.07 (0.25)	0.08 (0.27)	0.226
<i>Hispanic (Last Name)</i>	0.13 (0.34)	0.13 (0.34)	0.743
<i>Age</i>	34.99 (10.73)	33.54 (10.58)	$\leq 0.001^{***}$
EMPLOYMENT AND PLATFORM USE			
<i>Salary</i>	93846.26 (57009.77)	84853.71 (53265.85)	$\leq 0.001^{***}$
<i>High Job Position</i>	0.17 (0.37)	0.15 (0.36)	0.356
<i>Works in HR</i>	0.09 (0.29)	0.09 (0.29)	0.81
<i>Number of Contacts</i>	341.86 (179.74)	313.08 (183.95)	$\leq 0.001^{***}$
<i>Number of Skills</i>	22.18 (13.95)	20.90 (13.40)	0.016*
<i>Number of Skill Verifications</i>	43.73 (58.32)	42.02 (231.83)	0.809
<i>Number of Posts</i>	0.57 (0.50)	0.51 (0.50)	0.002**
<i>Has Volunteering Experience</i>	0.23 (0.42)	0.20 (0.40)	0.025*
<i>Gender Pronouns Shown</i>	0.15 (0.35)	0.13 (0.34)	0.166
<i>Profile picture is happy</i>	0.80 (0.40)	0.82 (0.38)	0.106
<i>Follows a philanthropist</i>	0.03 (0.18)	0.03 (0.18)	0.9
EMPLOYER			
<i>Employees</i>	5073.61 (4533.03)	4820.76 (4521.65)	0.153
<i>Employees on Platform</i>	27714.60 (73301.18)	27353.04 (75176.30)	0.901
<i>Open Jobs on Platform</i>	1975.35 (6121.76)	2077.55 (8122.78)	0.723
HIGHER EDUCATION			
<i>None</i>	0.13 (0.34)	0.16 (0.37)	0.032*
<i>Some College</i>	0.11 (0.31)	0.12 (0.32)	0.531
<i>Associate</i>	0.03 (0.18)	0.04 (0.20)	0.313
<i>Bachelor</i>	0.46 (0.50)	0.43 (0.50)	0.133
<i>Master</i>	0.23 (0.42)	0.21 (0.41)	0.213
<i>PhD</i>	0.03 (0.17)	0.03 (0.18)	0.532
<i>Undergrads: White</i>	0.62 (0.19)	0.63 (0.18)	0.375
<i>Undergrads: Black</i>	0.09 (0.12)	0.09 (0.11)	0.925
COUNTY			
<i>Share Democrat (2020)</i>	0.61 (0.15)	0.60 (0.15)	0.053.
<i>Share White</i>	0.57 (0.19)	0.58 (0.19)	0.593
<i>Share Black</i>	0.17 (0.15)	0.17 (0.14)	0.684
<i>Dissimilarity Index (Black/White)</i>	54.88 (11.36)	54.59 (11.81)	0.508

The table reports the differences in the resulting networks between Black and White profiles. Each row represents a certain feature of the connected users. The table is restricted to users who only accepted either the Black or White connection request (i.e., we exclude all users who accepted both). T-tests are used to obtain the following significance levels: * $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$; **** $p < 0.001$.

considerably more discriminatory (Col. 6). Most surprisingly, women appear to discriminate substantially more than men, and younger individuals more than older ones (Col. 1-2). Specifically, we find a 2 pp. gap in acceptance rates for men and a gap more than double that in size for women. Interestingly, we also find that Black individuals discriminate, although less so than White users (Col. 3). Given that only around 6% of targets are Black, the coefficients are rather imprecise and we can neither exclude that they discriminate as much as White users (the last row, showing the interaction effect, is not significantly different from zero) nor that they do not discriminate (the negative coefficient in the first row is not significantly different from zero). The estimates do, however, provide tentative evidence that

Black users also discriminate against Black profiles, but to a lesser extent than non-Black users. Further, comparing the White favoring gap of White users (3.22 p.p. favoring) to the Black favoring gap of Black users (-1.41 p.p. favoring, i.e. disfavoring) clearly reveals that Black users do not prefer Black profiles to the same extent as White users prefer White profiles ($p = 0.0015$).¹⁷

Experts correctly anticipate the results for education, Republican vote shares, and the direction of the effect for race (see the online Appendix H). However, our findings regarding gender and age go against experts’ priors. Although experts expect women to show a smaller acceptance gap (women: 10% vs. men: 15%, in relative terms), we instead find a substantially larger one (women: 20% vs. men: 8%). Experts further predict that discrimination increases linearly in age (Boomers: 22%; GenZ: 6.5%), while we find the opposite to be the case (Boomers: 5%; GenZ: 16%).

We also explore heterogeneous treatment effects using causal forests (Athey, Tibshirani, and Wager, 2019; Wager and Athey, 2018). The results reveal additional heterogeneity and confirm our findings above. Specifically, gender and age are the most predictive variables that explain heterogeneity. Moreover, the causal forests provide conditional average treatment effects for each individual. Given the rich set of covariates fed into the forest, this provides an idea of the distribution of discriminatory behavior. Overall, 91% of individuals are predicted to discriminate (see Figure G.10 in the online appendix), indicating that while not everyone discriminates, discrimination is also not concentrated in singular groups. Intuitively, it suggests that even if we had focused our study on only a specific subgroup of targets, we would, in most cases, have found a White-favoring gap in acceptance rates.

In sum, the findings of the first stage of the experiment show a considerable difference in the propensity to accept connection requests from Black and White profiles. This difference emerges instantly and persists over time, ultimately resulting in a 13% gap in the number of connections. The networks formed by Black and White profiles also differ in their composition. For example, the networks of Black profiles comprise more men and users with a greater number of contacts. In addition, the detailed information on our targets means that we can study multiple correlates of discrimination. We find, for instance, evidence that women and younger users discriminate comparatively more than others. Finally, an explorative analysis based on a large set of covariates reveals that discrimination is almost ubiquitous.

4.2 Informational Benefits

In this section, we turn to Stage II of our experiment, in which we evaluate discrimination in information provision. Specifically, to elicit the informational benefit provided by the

¹⁷For a comprehensive analysis and a specific focus on individual characteristics see online Appendix G5.

Table 2: Drivers of discrimination

Connection request acceptance probability (in %) by user’s characteristic						
	Is Old (1)	Is female (2)	Is Black (3)	High Education (4)	High Job (5)	Republican (6)
Black profile \times User characteristic = True (Black-White acceptance rate gap for users with characteristic)	-2.15*** (0.45)	-4.25*** (0.41)	-1.33 (1.36)	-2.89*** (0.36)	-2.57*** (0.73)	-3.85*** (0.42)
Black profile \times User characteristic = False (Black-White acceptance rate gap for users without characteristic)	-4.68*** (0.46)	-1.90*** (0.43)	-3.22*** (0.32)	-3.66*** (0.49)	-3.27*** (0.32)	-2.44*** (0.42)
User characteristics = True (Difference in acceptance rate for users with characteristic)	-7.58*** (0.68)	-0.22 (0.63)	2.53 (1.51)	3.22*** (0.65)	-3.33*** (0.85)	1.19 (0.69)
Constant (Baseline acceptance rate)	32.09*** (0.54)	26.44*** (0.52)	26.00*** (0.42)	24.36*** (0.58)	26.97*** (0.42)	25.81*** (0.54)
Differences in Gaps						
Difference in Gaps (Difference in acceptance rate gap between users with and without characteristic)	2.53*** (0.64)	-2.35*** (0.59)	1.89 (1.40)	0.77 (0.61)	0.70 (0.79)	-1.41* (0.60)
Picture random effects	✓	✓	✓	✓	✓	✓
Target random effects	✓	✓	✓	✓	✓	✓
Observations	33,446	36,911	33,861	38,299	38,299	36,306
Notes: p<0.10;*p<0.05;**p<0.01;***p<0.001.						

The table estimates the decision to accept a connection as a function of the requesting profile’s race and user characteristics. The regressions are conducted at the target level and all follow Equation 1 in online Appendix J1. Each column denotes one regression that estimates the difference in the acceptance rate (gap) for users with and without the reported characteristic. For example, in Column 1 the user characteristic is being older than the median target. Columns 2-6 are defined as follows: target (2) is female, (3) is Black, (4) has at least a bachelor’s degree, (5) has a job title that includes “CEO”, “director”, or “senior”, (6) lives in a county with an above-median Republican vote share. The number of observations between columns differs as we do not observe all characteristics for every user. In Table J.11 in the appendix, we show that the results are robust to reducing the sample to targets for whom we observe all six characteristics. The first part of the table (rows 1-4) shows regressions that separately compute the acceptance rate gap between Black and White profiles for users with and without the respective characteristic (rows 1-2). We also report the baseline acceptance rate and the difference in the baseline rate for users with the respective characteristic (rows 3-4). The second part of the table (row 5) reports the difference between the coefficients in rows 1 and 2, i.e., the difference in acceptance rate gaps between users with and without the respective characteristic. The result stems from a separate regression with a simple interaction effect between the profile being Black and the user characteristic being true.

networks and akin to the use of real job networks, we send connected users a message asking for relevant career or job application advice. Before sending the message, we endow Black and White profiles with statistically equal networks by swapping half of the twin pairs’ profile pictures and hence their race. By giving Black and White profiles access to the same networks we can cleanly identify discrimination in information provision. We first study differences in responses and response rates, i.e., direct discrimination during Stage II. We

then estimate the informational benefits a Black versus White profile can expect to receive due to discrimination during both the first and second stage.

4.2.1 Replies

Overall, roughly 21% of all contacts who received a message responded to our inquiry. While their replies consist of 50 words on average, these range from just a few words to over half a page. Most respondents share their experience, information, or advice, and many give extensive and valuable replies. Some offer to meet or talk on the phone and some are even willing to act as a reference for future applications. Overall, almost 65% of the responses contain some useful content (offer a referral, to talk on the phone, non-generic detailed information, etc.).¹⁸

4.2.2 Discrimination in Responses

Prior to conducting our second stage of the experiment, we swapped half of the profile pictures. The swap levels the playing field by providing Black and White profiles with access to the same networks (i.e., half of the Black profiles have access to White networks and half of the White profiles have access to Black networks). By design, this approach resolves all endogeneity from the first stage of the experiment and allows for a clean investigation of racial preferences in the second stage. Our main analysis in this section compares response rates toward requests from Black versus White profiles. Independent of the first stage, any such difference would suggest that Black and White profiles are treated differently when asking for advice during the second stage.

Figure 4 reveals no significant difference in response probabilities (see also Table 4). The difference in response rates is remarkably close to zero, with just a 0.6 p.p. difference between responses toward requests of Black and White profiles, which is neither statistically significant nor economically meaningful. Even if we were to assume that the effect is indeed there, it would require a total sample of almost 140,000 observations at this stage to statistically detect it at the 5% level with a power of 80%, highlighting the small magnitude of the estimated effect. This small effect also contrasts with the predictions of experts. While our results suggest an insignificant 3% gap in response rates, experts predict a 13% gap, with 88% of experts predicting a gap in favor of White profiles. Overall, the findings indicate very little direct discrimination in information provision. Once Black and White profiles are (artificially) equipped with the same networks, response rates are very similar.

This result remains when computing differences as a function of profile quality, network type, or message content.¹⁹ We also analyze whether there are differences in how useful the

¹⁸See online Appendix G6 for examples of messages, Table J.15 for message summary statistics, and online Appendix G7 for a deeper discussion of usefulness. Online Appendix G8 provides further analysis of the predictors of (useful) replies.

¹⁹See Tables J.18 and J.19 in the online appendix.

responses are, as defined by three different proxies: the length of the message, whether the message is highly valuable (a referral or a meeting is offered), or how useful the message is considered by a large language model.²⁰ In line with the results above, we find no differences in the usefulness of the messages received by Black and White profiles. We similarly observe no evidence of discrimination for message content and only very little (marginally significant) discrimination for profiles attending worse universities. We can, in addition, move beyond the aggregation at the profile level and break down results based on target characteristics, as each profile sends multiple connection requests. Black and White profiles seem to be treated essentially identically across most target characteristics (age, gender, race, etc.) in this second stage of the experiment.²¹ An additional advantage of our design is that we can differentiate between targets who only accepted one connection request in Stage I and those who accepted both. While the latter show an almost identical response probability towards Black and White profiles, those who accepted only the White connection requests are just slightly more likely to respond to White profiles, while targets who accepted only the Black connection requests are conversely slightly more likely to respond to Black profiles. While this hints at targets’ preference to interact with profiles of the same race as the network-generating profiles, differences are small and insignificant. Overall, when employing a clean identification by granting Black and White profiles access to the same networks, we find no evidence of discrimination against Black profiles in Stage II.

Our experimental setup allows us to study differences in responses in even greater detail by leveraging the fact that the race of the network-creating profile in Stage I is independent of the race of the messaging profile in Stage II. Overall differences in information provision can be attributed to three separate effects. First, as shown above, *discrimination* during Stage II may affect the response rate, i.e., people may react differently to Black vs. White messages. Second, the *composition of networks* between those originally built by Black and White accounts in Stage I may differ, as suggested in Section 4.1.2. Third, responses may be affected by a profile’s *fit* into the users’ networks. Users may have a preference to interact with individuals with certain characteristics at a given point in time. In a natural setting, the ‘fit’ might depend on race as well as other components, such as the user’s network, skills, job interests, etc. Our second stage, however, solely switches the race of half of our profiles,

²⁰See online Appendix G7 for more details and validation of the measure of ‘usefulness’.

²¹See online Appendix G10 for a detailed discussion. Note that, by design, targets treated in this stage are selected by whether they accepted at least one request during Stage I, making them not comparable to Stage I targets. Additionally, the sample of targets receiving a message is slightly biased towards less discriminatory individuals, as we exclude self-employed, freelancers, or small company workers. Indeed we see that while 51% of users that accepted at least one connection request in Stage I accepted both requests, these users make up 57% of participants in Stage II. This increase is, however, relatively minor, and re-weighting our sample based on observables to match the original sample does not change any of the results.




Figure 4: Response rate by race of the profile

The figure depicts the response rate by the race of the requesting profiles separately. Orange dots denote White profiles, while blue ones denote Black profiles. Each dot represents one profile and twin pairs are connected through colored lines (where red versus blue lines indicate a gap in favor of White or Black profiles, respectively). Whiskers around the mean reflect the corresponding 95% confidence intervals. T-tests are used to obtain the following significance levels: $\cdot p < 0.10$; $* p < 0.05$; $** p < 0.01$; $*** p < 0.001$. Tables J.18 and J.19 report the corresponding regressions, and Figure G.5 further displays differences by profile quality (all in the online appendix).

leaving everything else intact. Thus, the ‘fit’ we measure captures first-stage preferences for interacting with a specific race. This is different from direct discrimination in the sense that ‘fit’ can be beneficial to both Black and White profiles. It simply describes whether people who accepted a profile of a specific race prefer to interact with a user of that race at a given point in time. Thus, ‘fit’ might be driven by homophily during Stage I but may also be explained by users’ current preferences for interacting with a certain race. For example, one might think of HR professionals who groom their network for potential future employees or diversity officers interested in the experiences of underrepresented groups. Thus, while ‘fit’ in our specific setting is only race-based, it is different from direct discrimination. The ‘composition’ and ‘fit’ are hence both functions of discrimination in Stage I. ‘Fit’ captures racial preferences in Stage I while ‘composition’ results from heterogeneity in discriminatory behavior in Stage I. ‘Discrimination’ meanwhile evolves exclusively in Stage II, i.e., after resolving endogeneity in network formation.²²

²²Table J.17 (in the online appendix) illustrates these effects in our setting, showing which are active for

In studies based on observational data, it is not possible to disentangle these typically endogenous and correlated effects. Our experimental design allows us to measure each of the three components. To operationalize the components, we run a regression with three independent variables. The first describes whether the requesting profile is Black in Stage II and captures discrimination, as shown in Figure G.5. The second indicates whether the profile’s network was constructed by a Black profile (as opposed to a White profile) during Stage I, thus capturing differences in the ‘composition’ of networks by comparing response rates between the two types of networks. Finally, the regression includes a dummy for whether a given profile’s picture was inserted into the network of its twin. This variable measures the ‘fit’ of the network and hence the preference for interacting with a profile that has the same race as the profile originally accepted into the network. Importantly, by design, all three variables are orthogonal to one another. For example, given that half of the profiles swapped their profile picture, knowledge of whether the messaging profile is White provides no information on whether it was swapped (‘fit’) or resides in a Black network (‘composition’). Similarly, knowing that a profile resides in a Black network (‘composition’) reveals no information on its Stage II race (‘discrimination’) or whether it was swapped (‘fit’).

Table 3 estimates each of these components separately (Columns 1-3) and jointly (Column 4). Given their orthogonality, the joint regression preserves the point estimates. As shown above, discrimination has no statistically or economically significant effect on response rates. In addition, users in Black networks are more likely to respond to a given message, as indicated by the ‘composition’ estimate, which is 1.8 p.p. though not significantly different from zero. Finally, the ‘fit’ component is the largest in magnitude and marginally significantly different from zero, indicating that networks are more likely to respond to a profile with the race that originally created the network. The increased response rate suggests some preference for interacting with the race originally accepted into the network.²³

In sum, we disentangle three drivers of responses: discrimination, the ‘composition’ of the network, and ‘fit’ into the network. The smallest of these drivers, once accounted for endogenous differences between Black and White networks, is race (only in terms of magni-

which type of network as well as average response rates by type of network.

²³Although interacting fit’ with race in Stage II might reveal separate discrimination coefficients for Black and White networks, it would conflate two effects: network-specific discrimination in Stage II (e.g., users in White networks may be more discriminatory) and first-stage preferences for certain races (fit’). Users who accepted a Black request in Stage I may now prefer Black profiles, making it hard to disentangle these factors. However, Table J.20 shows that users accepting both requests in Stage I show no significant differences in Stage II, while those accepting only one request show slight favoritism toward the race of that profile. This may suggest the ‘fit’ component being driven by individuals who only accepted a Black or White request in Stage I, but it is important to note again that these results do not solely reflect the ‘fit’ component, but are also driven by second-stage discrimination.

tude and statistically not significantly different from the other drivers). Overall, the results suggest that once Black and White profiles are endowed with the same networks, there are only minor and insignificant differences in users’ propensity to respond to their messages.

Table 3: Decomposing Stage II effects

	Response Rate (in %)			
	(1)	(2)	(3)	(4)
Fit (Race of network-generating and messaging profile is identical)	2.55 (1.52)			2.54 (1.52)
Composition (Network-generating profile was Black)		1.81 (1.33)		1.79 (1.33)
Discrimination (Profile is Black)			-0.66 (1.33)	-0.67 (1.33)
Constant	19.45*** (1.07)	19.82*** (1.01)	21.05*** (1.02)	18.88*** (1.43)
Picture specific random effects	✓	✓	✓	✓
Observations	400	400	400	400
Notes:	p<0.10;*p<0.05;**p<0.01;***p<0.001.			

The table estimates the response rate in Stage II (after swapping profile pictures). *Fit* denotes a dummy with a value of one if the profile is in the original network, and zero if the profile is in an alien network (i.e., the race of the network-generating profile and that of the messaging profile are identical). *Composition* denotes a dummy with a value of one if the network is built by a Black profile (i.e., has the composition of a Black network), and zero otherwise. *Discrimination* denotes a dummy with a value of one if the profile picture (in the current stage) depicts a Black person, and zero otherwise. Negative values, therefore, indicate discrimination against Black profiles. The regressions are conducted at the profile level and follow the mixed effects models of Equation 1 in online Appendix J1. To account for twin-profile-specific heterogeneity, we use a random effect at the twin-target level. Table J.21 in the online appendix reports estimates by profile quality.

4.2.3 Expected Informational Benefits

Finally, we are interested in the compound effect of discrimination, i.e., direct discrimination in Stage I and II. We obtain the compound effect by computing the *expected number* of responses for both Black and White profiles had they remained in their original network. The overall informational benefit of a network is a function of both the likelihood of a response and the size of a given network. As for network size, White profiles have an advantage. Black networks, however, are more responsive, as suggested by the composition result above.

To compute the overall informational benefit of a profile’s network, we first calculate the response probability for each target based on their characteristics and the profile’s race.²⁴

²⁴In more detail, we begin by using a stepwise regression builder to obtain the most important link-independent predictors of response. The main predictors are whether the user has an HR job, has obtained a bachelor’s degree, the number of contacts the user has, and whether the user resides in a Democratic county. Moreover, we also make use of the most salient demographic characteristics like gender, age, race, and whether the user has a senior job. Thereafter, we estimate the individual response probability of each connected user based on these features interacted with the race of the profile. Missing values (for example,

We then aggregate expected response probabilities of a profile’s acquired connections during Stage I. The result describes the expected number of responses, i.e., the expected informational benefit.

Figure 5 shows that White profiles can expect to receive roughly one more message than Black profiles. In other words, White profiles receive around 24% more messages than Black profiles.²⁵ Hence, the advantage of Black profiles in terms of having a more responsive network does not improve the response rate enough to overcome the disadvantage of having fewer connections.

Overall, we find that Black profiles can expect substantially lower informational benefits from their networks. Given that – in the aggregate – we find no evidence of second-stage discrimination, we conclude that Black profiles’ reduced informational benefits are driven by the experiment’s first stage, i.e., Black networks being substantially smaller than those of White profiles.²⁶

4.3 Mechanisms

In this section, we discuss different mechanisms that explain our results. We begin by eliminating several rather mechanical explanations for observing direct discrimination in the first but not in the second stage. First, users may have noticed the swapped profile picture and adjusted their behavior. Our results do not, however, change when focusing only on non-swapped profiles (see Table J.19 in the online appendix). Further, we find no difference in profile views and link suspensions between swapped and non-swapped accounts in a simple Diff-in-diff setup. Finally, we exploit the randomly induced variation in the time lag between the connection request and the message (from 4 to 13 weeks). The results show that the time users had to forget about the initial request did not affect their Stage II behavior. Overall, the swap itself does not seem to explain the results (see Appendix G9 for an in-depth discussion).

Second, users might meet more White people in their everyday lives and thus, when answering their connection requests, be more likely to believe they had met a given White rather than Black profile. In general, exposure asymmetry is a disadvantage inherent to being an underrepresented group and may also be relevant in other contexts of professional networking. However, we find several results speaking against this mechanism. Although

for users who do not have a job title) would lead to missing predictions, which in turn could bias our results as the composition of Black and White networks differ. We accordingly impute the missing values for all users (at this point only) with the mean of the respective variable. This simply ensures that we have a non-missing prediction for each user’s probability of responding to a message.

²⁵See Tables 4 and J.26 in online Appendix J for further analysis.

²⁶We can also zoom into the non-swapped and swapped profiles separately. When doing so, we find the same gap in the expected number of responses for swapped but also non-swapped profiles.




Figure 5: Number of ex-ante expected responses when creating a network

The figure depicts the ex-ante expected responses when creating a network for White and Black profiles separately. Orange dots denote White profiles, while blue ones indicate Black profiles. Each dot represents one profile and twin pairs are connected through colored lines (where red vs. blue lines reflect a gap in favor of White or Black profiles, respectively). Whiskers around the mean denote the corresponding 95% confidence intervals. T-tests are used to obtain the following significance levels: $\cdot p < 0.10$; $\ast p < 0.05$; $\ast\ast p < 0.01$; $\ast\ast\ast p < 0.001$. Table J.26 reports the corresponding regressions, and Figure G.6 further displays differences by profile quality (both in the online appendix).

we cannot directly observe users' full networks and whom they might have met, we can indirectly test for their exposure to Black Americans. Specifically, we observe the share of Black people in their home county and the share of Black students at the university they attended. We find no evidence of exposure to either a larger Black population or student body leading to less discrimination (see Table J.12 in the online appendix for a detailed analysis). If anything, the estimates suggest more discrimination with increased exposure (though coefficients are insignificant). This result is supported by our state-level analysis, which shows that states in the 'Black belt' (regions with a strong history of slavery and a large Black population) exhibit significantly higher rates of discrimination (see Appendix G2). If it were exposure to Black Americans rather than discrimination driving the results, we would expect the opposite result. Finally, we asked LinkedIn users about their willingness to accept a connection request from somebody they do not know and what they consider the up- and downsides of doing so (see Appendix I). Only half of all participants said they were unlikely or extremely unlikely to accept such a connection request. This suggests that acceptance rates are likely driven by users' willingness to accept strangers rather than their uncertainty about having met a given profile. These insights align with using LinkedIn as a professional website, where advancing one's career is the main motivation. Indeed, surveyed

Table 4: Main estimates

	Race effect					
	Number of Contacts (Stage I)		Response Rate (in %) (Stage II)		Informational Benefit (Stage I+II)	
	(1)	(2)	(3)	(4)	(5)	(6)
Profile is Black	-3.06*** (0.47) [0.45]	-3.07*** (0.54) [4.45]	-0.66 (1.33) [1.02]	-0.38 (1.37) [13.13]	-1.01*** (0.10) [0.10]	-0.99*** (0.11) [0.91]
Constant	26.13*** (0.44) [0.48]	35.66*** (5.14) [0.58]	21.05*** (1.02) [1.33]	27.23 (14.32) [1.47]	5.32*** (0.09) [0.10]	7.61*** (1.05) [0.12]
State Controls	×	✓	×	✓	×	✓
Job Controls	×	✓	×	✓	×	✓
Firstname Controls	×	✓	×	✓	×	✓
Lastname Controls	×	✓	×	✓	×	✓
Picture trait Controls	×	✓	×	✓	×	✓
Picture specific random effects	✓	✓	✓	✓	✓	✓
Observations	400	400	400	400	400	400

Notes:

p<0.10; *p<0.05; **p<0.01; ***p<0.001.

The table estimates the number of contacts a profile has by the end of Stage I as a function of their race in Columns 1 and 2. Columns 3 and 4 estimate the response rate in Stage II (after swapping profile pictures). Columns 5 and 6 estimate the expected informational benefit of the profiles. *Profile is Black* denotes a dummy with a value of one if the profile picture (in the current stage) depicts a Black person and zero otherwise. The regressions are conducted at the profile level, use various controls, and all follow the mixed effects models of Equation 1 in online Appendix J1. To account for twin-profile-specific heterogeneity, we use a random effect at the twin-target level. In square brackets, we further display robust standard errors clustered at the twin level.

users referred to personal benefits – like expanding networks and advancing professionally – as the main reasons for responding to connection requests or job messages (see Figures I.9, I.10, and I.13 in the online appendix). These arguments suggest that the exposure mechanism cannot explain the first-stage results.²⁷

Next, fast versus slow thinking (Kahneman, 2017; Shleifer, 2012) is another potential driver of our results. Specifically, people may apply fast thinking and hence simple heuristics when choosing whom to accept in Stage I, while making a more rational decision of whom to respond to in Stage II. Several arguments, however, speak against this mechanism being the main driver of the results. First, our survey participants report taking around 3 minutes to consider a connection request, with only 18% of users requiring only ‘seconds’ (see online Appendix I). Hence, a vast majority need substantially more time than is usually considered intuitive thinking. Second, the decision as to whether or not to respond can be equally fast in Stage II, where we do not document discrimination on either the intensive or

²⁷Another argument against the exposure mechanism is that people are better at recognizing faces of their own race (Hills and Pake, 2013). Hence, White users should more easily identify strangers among White profiles and thus accept them at a lower rate than Black profiles. We, however, observe the opposite.

extensive margin. Finally, slow thinking would also predict discrimination due to biases and statistical differences between Black and White Americans, such as Black Americans’ higher unemployment rate (U.S. Bureau of Labor Statistics, 2023), a difference that LinkedIn users are aware of.²⁸ To summarize, while slow thinking and simple heuristics may play a role in deciding which connection requests to accept, they cannot explain our null results in Stage II.

Another more mechanical concern could be a selection of non-discriminatory people into Stage II. While we naturally cannot observe response rates by people who do not want to engage in networking in the first place, our experimental design ensures that Stage II also captures individuals who accepted only one of our connection requests. As discussed above, we find no significant differences in response rates between Black and White message requests for users who only accepted the connection request of a Black or a White profile.²⁹

Finally, users might consider connection requests from profiles with zero initial connections to be odd. As argued in Section 3., starting the experiment with initial connections would come with serious drawbacks to the experimental design. In addition, our profiles represent professionals in the early stage of their careers, a time when the majority of users join the platform. Further, respondents to our LinkedIn user survey do not seem to view few connections as a ‘red flag’ when judging requesting profiles (see Figure I.14 in the online appendix). Importantly, our experimental design allows us to measure discriminatory behavior dynamics over the duration of the experiment’s first stage, i.e., conditional on profiles’ connections. We find a stable and constant gap in acceptance rates from day 1 until the end of the experiment, speaking against this concern (see online Appendix G1).

Next, we look at our results from the perspective of traditional theories of discrimination, which usually consider static settings. If either taste-based discrimination (Becker, 1957) or homophily (McPherson, Smith-Lovin, and Cook, 2001) were the main explanation for our findings, they would be present in Stages I and II since race is equally salient in both. However, we find virtually no differences in response rates during Stage II, even among those who only accepted the White request during Stage I (see Table J.20 in the online appendix for a detailed analysis). Furthermore, both theories are incongruent with suggestive evidence that Black users discriminate against Black profiles as well (or at least do not favor Black requests comparably to White users favoring White requests). Hence, both theories are

²⁸See Figure I.12 in the online appendix for details on perceived differences in the economic performance/well-being of Black and White Americans.

²⁹Yet another potential explanation for observing discrimination in Stage I but not in Stage II is that users’ decision frame may shift. In Stage I, users may primarily consider whether a potential connection is valuable to them. In Stage II, they might instead wonder whether or not they can provide helpful information. However, when asking respondents which factors most strongly affect their decision to accept/respond, users prioritize their own benefit for both decisions (see Figures I.5 and I.11 in the online appendix).

unlikely to be primary drivers of the results.³⁰

Statistical discrimination (Phelps, 1972) seems to be a more likely candidate. Our survey indicates that LinkedIn users prioritize the perceived value of a connection request when deciding to accept or decline, with 82% considering it highly relevant, outweighing all other factors (see online Appendix I). At the same time, participants expect to gain less value from connections with Black users. First, they expect Black users to be significantly less helpful than White ones. In addition, when asked about the career trajectories of workers resembling our profiles, participants predict a 13% income gap to emerge within the next 5 years, with virtually everyone expecting the White candidate to earn more. That said, statistical discrimination on its own cannot account for our null result in Stage II, as statistical differences between Black and White profiles would remain in Stage II.

We identify two potential effects that might consistently explain our results in both stages.³¹ First, users might have in-group preferences (Akerlof and Kranton, 2000; Chen and Chen, 2011). As famously documented by Tajfel and Turner (2004), even randomly defined groups can activate in-group favoritism. Thus, users that have something in common with our profiles may consider them part of their in-group during Stage I and consequently have higher acceptance rates. In support of this claim, we find substantially higher acceptance rates for users who attended the same university or worked for the same employer as our profiles. However, an argument against this mechanism is that Black users do not favor Black profiles; indeed, we find suggestive evidence of Black users also discriminating against Black profiles. Once in Stage II, all profiles in the user’s network may be considered part of her in-group, resulting in no discrimination. Our survey likewise indicates that the presence of a profile in their network is the primary factor influencing respondents’ decisions to answer a profile’s message (other than personally knowing the person).³² Hence, it seems plausible that the group identification changes from a very weak association based on race and other characteristics in Stage I to a stronger in-group association due to being part of the user’s network in Stage II, consistent with the vanishing discrimination in Stage II. However, as mentioned above, this mechanism alone likely cannot consistently explain all our results as we do not observe a Black-favoring acceptance gap by Black users in Stage I.

Another potential explanation is rational inattention (Bartoš et al., 2016; Maćkowiak,

³⁰In addition, our LinkedIn user survey reveals that only 8% of respondents rate ethnicity as very or extremely important when deciding on connection requests (see Figure I.8 in the online appendix).

³¹Stereotypes might be another possible explanation (Bordalo et al., 2019) as they can be context-dependent, and thus more or less pronounced in Stages I and II. Unfortunately, in our setup, it is not possible to test how stereotypes shape the behavior of targets and it is unclear why the effect would be different between the stages. This might, however, offer an interesting direction for future research.

³²Figure I.11 in the online appendix shows the weight people put on different characteristics when deciding whom to answer.

Matějka, and Wiederholt, 2023). Specifically, when receiving a connection request, users can decide to incur attention costs to acquire information before deciding whether to accept or ignore it. Users might accordingly make their decision based on the profile picture, or could also visit the profile page to examine all relevant CV characteristics to judge whether the potential connection is of sufficient value. Our survey shows that LinkedIn users are aware of statistical differences in economic performance/well-being between Black and White Americans. They may therefore allocate different levels of attention in searching CVs of White versus Black profiles, which could explain different acceptance rates. In line with this argument, we find that Black profiles have lower view rates than White profiles (see online Appendix G11). In contrast, when deciding whether to respond to a direct message in Stage II, users may not remember the information they collected in Stage I, if any. They can either decide to gather the information again or avoid attention costs and rely on the judgment of their Stage-I-self, who accepted the connection request and, hence, must have deemed the profile of sufficient value. Since the latter is a clear signal of the profile’s value, users may stick to past decisions and treat connections in their network equally. Both explanations (i.e., rational inattention and in-group preferences) could be interpreted as a ‘foot-in-the-door’ effect, suggesting that once a profile is part of a user’s network, race no longer plays a significant role in users’ decisions.

Our results can be viewed through the lens of ‘systemic discrimination’. Although most static correspondence studies focus on direct (e.g., taste-based) discrimination, much of the social sciences sees discrimination as a cumulative result of multistage processes (Onuchic, 2023). For example, high school teachers may discriminate against Black students in grading, leading to fewer college acceptance letters for those students, even if the admissions processes are race-neutral. In a recent contribution, Bohren, Hull, and Imas (2022) describe the idea that earlier discrimination affects outcomes at a given stage as ‘systemic discrimination’.³³ We observe a similar dynamic in our results. We find that targets are discriminating during Stage I of the experiment, leading Black profiles to have fewer connections. Despite no direct discrimination in Stage II, Black profiles receive substantially lower informational benefits due to systemic discrimination originating in Stage I.

Overall, our two-stage design shows that differences in total informational benefits emerge due to direct discrimination in network formation during Stage I. The findings are most consistent with ‘gatekeeping’ driven by rational inattention and potentially in-group bias, leading users to treat Black and White profiles differentially in Stage I and then similarly once they are part of their network in Stage II.

³³Recent contributions introduce this view into economics using both observational data (Dobbie, Hull, and Arnold, 2022; Baron et al., 2024) and experiments (Bohren, Hull, and Imas, 2022).

5. Summary and Discussion

Our experiment yields three main findings. First, we find substantial evidence of discrimination in the formation of job networks. Black profiles have a 13% lower acceptance rate for their connection requests than White profiles (23% vs 26%). This is very close to the 2-3 p.p. difference in employers’ callbacks found in previous studies (Bertrand and Duflo, 2017). Though heterogeneous, discrimination is widespread, both geographically and across individuals with different characteristics. Second, we find that Black users receive substantially fewer informational benefits during Stage II. Third, we find a precise null effect of race on response rates in Stage II itself. Differences in informational benefits are primarily driven by discrimination during the formation of networks in Stage I (i.e., gatekeeping). These findings are consistent with rational inattention (Maćkowiak, Matějka, and Wiederholt, 2023) and potentially in-group bias (Akerlof and Kranton, 2000). Both suggest a ‘foot-in-the-door’ interpretation of the results, indicating that once Black profiles are part of a user’s network, they are treated no differently.

Overall, our paper presents compelling evidence that discrimination plays a significant role in shaping the informational benefits provided by professional job networks. Our findings, thus, offer evidence regarding informal networks (Topa, 2011), which can help explain the disparities in labor market outcomes between underrepresented groups and White individuals in the U.S. labor market. Already, in 1987, a paper in the *American Economic Review* suggested that “*informal methods of search [...] account for 87-90 percent of the difference in youth employment probabilities between blacks and whites*” (p. 451 Holzer, 1987). Our paper provides the first causal evidence in this regard, showing that discrimination is effectively driven by gatekeeping. This underlines both the importance of creating inclusive institutions and breaking up ‘old boys’ clubs’ (Cullen and Perez-Truglia, 2023; Michelman, Price, and Zimmerman, 2022), as well as justifies the use of affirmative action strategies, including targeted networking events and workshops. We also shed light on the mechanisms through which professional job networking platforms, such as LinkedIn, help users advance their careers. We show that weak-tie networks, such as those of our users, provide substantial informational benefits in terms of mentorship advice and job application insights, complementing previous work on the strength of weak ties (Gee et al., 2017).

It is important to note that our experiment is designed to causally identify the effects of discrimination. For causal identification and internal validity, both Black and White profiles reach out to the same target pool. In reality, gaps may be driven by other factors as well. Aside from economic conditions and homophily, this may also entail adjusted networking strategies in anticipation of discrimination. Our causal forest results suggest that only 9%

of all users are not predicted to discriminate. Further, even experts cannot predict who (more strongly) discriminates. Hence, even if Black users tried to optimize their requests towards less discriminatory users, these are both rare and hard to identify. These insights are consistent with descriptive evidence from LinkedIn suggesting that Black LinkedIn users have about 15% smaller networks than White users (Baird et al., 2023a,b).

Our study opens up numerous avenues for subsequent research. To begin with, we are the first to causally study the effects of discrimination on network formation and information provision on LinkedIn. Although offline job networks may function somewhat differently than online networks, the platform nonetheless provides an ideal setting to cleanly study job networks in general and discrimination more specifically. Given that both offline and online networks have been shown to strongly affect labor market outcomes, evidence on other countries, underrepresented groups, and genders is sorely needed. Here, our approach of varying race via AI-generated pictures offers an alternative to using names as signals, which tend to be noisy and potentially biased (Kreisman and Smith, 2023). Our method is easily adaptable to different contexts, enabling researchers to modify a range of individual attributes, from race to gender or age. Importantly, to keep treatment stable across geographies, we do not substantially vary the complexion of Black and White profiles. However, by more gradually varying race, our approach opens pathways for studying race or gender in non-binary categories, something that is typically not possible using names. For example, it allows for research on colorism; that is, discrimination towards darker versus lighter skin even within the same ethnic or racial groups (see Bodenhorn, 2006; Dixon and Telles, 2017). Similarly, the gender could gradually be varied to present more stereotypically feminine or masculine profiles. Finally, our study varies just one dimension of profile pictures, namely race, such that the findings might not be directly generalizable to women. Given that women are, for example, more frequently the subject of sexual harassment on online platforms (Atske, 2021), future research might use an adjusted experimental design to focus on this population.

Broadly, our insights emphasize the importance of understanding who drives discrimination in job networking and elsewhere, including why people discriminate and where it originates. Greater knowledge in this regard is crucial for the design of effective and well-targeted policies to counter discrimination in the labor market and beyond.

References

- Abel, M. and R. Burger (2023). Unpacking name-based race discrimination. Working paper, IZA Discussion Papers.
- Acquisti, A. and C. Fong (2020). An experiment in hiring discrimination via online social networks. *Management Science* 66(3), 1005–1024.
- Agan, A. and S. Starr (2018). Ban the box, criminal records, and racial discrimination: A field experiment. *The Quarterly Journal of Economics* 133(1), 191–235.
- Ajzenman, N., B. Ferman, and P. C. Sant’Anna (2023). Discrimination in the formation of academic networks: A field experiment on# econtwitter. Technical report, Institute of Labor Economics (IZA).
- Akerlof, G. A. and R. E. Kranton (2000). Economics and identity. *The Quarterly Journal of Economics* 115(3), 715–753.
- Angeli, D., M. Lowe, M. Lowe, et al. (2023). Virtue signals. Working paper, CESifo.
- Arechar, A. A., S. Gächter, and L. Molleman (2018, Mar). Conducting interactive experiments online. *Experimental Economics* 21(1), 99–131.
- Athey, S. (2020, October). Susan Athey Guest Talk - Estimating Heterogeneous Treatment Effects.
- Athey, S., J. Tibshirani, and S. Wager (2019). Generalized random forests. *The Annals of Statistics* 47(2), 1148–1178.
- Athey, S. and S. Wager (2019). Estimating treatment effects with causal forests: An application. *Observational Studies* 5(2), 37–51.
- Atske, S. (2021, January). The State of Online Harassment — pewresearch.org. <https://www.pewresearch.org/internet/2021/01/13/the-state-of-online-harassment/>. [Accessed 02-Mar-2023].
- Baird, M., D. Kavanagh-Smith, O. Osoba, and Y. Wu (2023a, Sep 25). Disparities in U.S. economic network formation. Economic graph research note, LinkedIn Economic Graph.
- Baird, M., D. Kavanagh-Smith, O. Osoba, and Y. Wu (2023b, Sep 25). Economic network strength: Race, gender, and community gaps in the us. Economic graph research note, LinkedIn Economic Graph.
- Barabási, A.-L. and R. Albert (1999). Emergence of scaling in random networks. *Science* 286(5439), 509–512.
- Baron, E. J., J. Doyle, Joseph J, N. Emanuel, P. Hull, and J. Ryan (2024, 02). Discrimination in Multiphase Systems: Evidence from Child Protection*. *The Quarterly Journal of Economics* 139(3), 1611–1664.
- Bartoš, V., M. Bauer, J. Chytilová, and F. Matějka (2016). Attention discrimination: Theory and field experiments with monitoring information acquisition. *American Economic Review* 106(6), 1437–75.
- Bayer, P. and K. K. Charles (2018). Divergent paths: A new perspective on earnings differences

- between black and white men since 1940. *The Quarterly Journal of Economics* 133(3), 1459–1501.
- Becker, G. S. (1957). *The economics of discrimination*. University of Chicago Press.
- Berinsky, A. J., G. A. Huber, and G. S. Lenz (2012). Evaluating online labor markets for experimental research: Amazon.com’s mechanical turk. *Political Analysis* 20(3), 351–368.
- Bertrand, M. and E. Duflo (2017). Field experiments on discrimination. *Handbook of Economic Field Experiments* 1, 309–393.
- 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.
- Block, R., C. Crabtree, J. B. Holbein, and J. Q. Monson (2021). Are americans less likely to reply to emails from Black people relative to White people? *Proceedings of the National Academy of Sciences* 118(52), e2110347118.
- Bodenhorn, H. (2006). Colorism, complexion homogamy, and household wealth: Some historical evidence. *American Economic Review* 96(2), 256–260.
- Bohren, J. A., P. Hull, and A. Imas (2022). Systemic discrimination: Theory and measurement. Working paper, National Bureau of Economic Research.
- Bohren, J. A., A. Imas, and M. Rosenberg (2019). The dynamics of discrimination: Theory and evidence. *American Economic Review* 109(10), 3395–3436.
- Bordalo, P., K. Coffman, N. Gennaioli, and A. Shleifer (2019). Beliefs about gender. *American Economic Review* 109(3), 739–773.
- Brooke Auxier, M. A. (2021, April). Social Media Use in 2021. <https://www.pewresearch.org/2021/04/07/social-media-use-in-2021-3/>. [Accessed 02-Mar-2023].
- Chen, R. and Y. Chen (2011, October). The potential of social identity for equilibrium selection. *American Economic Review* 101(6), 2562–2589.
- Chernozhukov, V., M. Demirer, E. Duflo, and I. Fernandez-Val (2018). Generic machine learning inference on heterogeneous treatment effects in randomized experiments, with an application to immunization in india. Working paper 24678, National Bureau of Economic Research.
- Chetty, R., M. O. Jackson, T. Kuchler, J. Stroebel, N. Hendren, R. B. Fluegge, S. Gong, F. Gonzalez, A. Grondin, M. Jacob, et al. (2022a). Social capital i: measurement and associations with economic mobility. *Nature* 608(7921), 108–121.
- Chetty, R., M. O. Jackson, T. Kuchler, J. Stroebel, N. Hendren, R. B. Fluegge, S. Gong, F. Gonzalez, A. Grondin, M. Jacob, et al. (2022b). Social capital ii: determinants of economic connectedness. *Nature* 608(7921), 122–134.
- Coffman, K. B., C. L. Exley, and M. Niederle (2021, June). The role of beliefs in driving gender discrimination. *Management Science* 67(6), 3551–3569.
- Conzelmann, J. G., S. W. Hemelt, B. J. Hershbein, S. Martin, A. Simon, and K. M. Stange (2023).

- Grads on the go: Measuring college-specific labor markets for graduates. *Journal of Policy Analysis and Management* *n/a*(*n/a*), 1–23.
- Cullen, Z. B. and R. Perez-Truglia (2023). The old boys’ club: Schmoozing and the gender gap. *American Economic Review* *113*(7), 1703–40.
- Davis, A. Y. (1981). Rape, racism and the capitalist setting. *The Black Scholar* *12*(6), 39–45.
- Dixon, A. R. and E. E. Telles (2017). Skin color and colorism: Global research, concepts, and measurement. *Annual Review of Sociology* *43*, 405–424.
- Dobbie, W., P. Hull, and D. Arnold (2022). Measuring racial discrimination in bail decisions. *American Economic Review* *112*(9), 2992–3038.
- Doleac, J. L. and L. C. Stein (2013). The visible hand: Race and online market outcomes. *The Economic Journal* *123*(572), 469–492.
- Dustmann, C., A. Glitz, U. Schönberg, and H. Brücker (2016, April). Referral-based Job Search Networks. *The Review of Economic Studies* *83*(2), 514–546.
- Edelman, B., M. Luca, and D. Svirsky (2017). Racial discrimination in the sharing economy: Evidence from a field experiment. *American Economic Journal: Applied Economics* *9*(2), 1–22.
- Fernandez, R. M. and I. Fernandez-Mateo (2006). Networks, race, and hiring. *American Sociological Review* *71*, 42–71.
- Forbes (2021). Forbes America’s Top Colleges List 2021. <https://www.forbes.com/top-colleges/>.
- Fryer Jr, R. G. and S. D. Levitt (2004). The causes and consequences of distinctively Black names. *The Quarterly Journal of Economics* *119*(3), 767–805.
- Gaddis, S. M. (2017). How black are lakisha and jamal? racial perceptions from names used in correspondence audit studies. *Sociological Science* *4*, 469–489.
- Gallen, Y. and M. Wasserman (2021). Do male and female students use networks differently? *AEA Papers and Proceedings* *111*, 154–158.
- Gee, L. K., J. Jones, and M. Burke (2017). Social networks and labor markets: How strong ties relate to job finding on facebook’s social network. *Journal of Labor Economics* *35*(2), 485–518.
- Gee, L. K., J. J. Jones, C. J. Fariss, M. Burke, and J. H. Fowler (2017). The paradox of weak ties in 55 countries. *Journal of Economic Behavior & Organization* *133*, 362–372.
- Georgia Department of Public Health (2022). Most Popular Baby Names. <https://oasis.state.ga.us/babynames>.
- Halevy, N., O. Weisel, and G. Bornstein (2012). “in-group love” and “out-group hate” in repeated interaction between groups. *Journal of Behavioral Decision Making* *25*(2), 188–195.
- Hern, A. (2017). Faceapp forced to pull ‘racist’ filters that allow ‘digital blackface’. <https://www.theguardian.com/technology/2017/aug/10/faceapp-forced-to-pull-racist-filters-digital-blackface>.
- Hills, P. J. and J. M. Pake (2013). Eye-tracking the own-race bias in face recognition: Revealing the perceptual and socio-cognitive mechanisms. *Cognition* *129*(3), 586–597.

- Holzer, H. J. (1987). Informal Job Search and Black Youth Unemployment. *American Economic Review* 77(3), 446–452.
- Hopkins Population Center (2020). HPC Data Hub Phase 3 Release (11/16/2020). https://github.com/QFL2020/COVID_DataHub.
- Horton, J. J., D. G. Rand, and R. J. Zeckhauser (2011). The online laboratory: Conducting experiments in a real labor market. *Experimental Economics* 14(3), 399–425.
- Ioannides, Y. M. and L. Datcher Loury (2004). Job information networks, neighborhood effects, and inequality. *Journal of Economic Literature* 42(4), 1056–1093.
- IPEDS (2022). Integrated Postsecondary Education Data System (IPEDS). Department of Education, National Center, <https://nces.ed.gov/ipeds/>.
- Kahneman, D. (2017). *Thinking, fast and slow*. Farrar, Straus and Giroux.
- Kalmijn, M. (1998). Intermarriage and homogamy: Causes, patterns, trends. *Annual Review of Sociology* 24(1), 395–421.
- Kaplan, J. (2022). jacobkap/predictrace.
- Karras, T., S. Laine, M. Aittala, J. Hellsten, J. Lehtinen, and T. Aila (2020). Analyzing and improving the image quality of StyleGAN. In *Proc. CVPR*.
- Kline, P., E. K. Rose, and C. R. Walters (2022). Systemic Discrimination Among Large U.S. Employers*. *The Quarterly Journal of Economics* 137(4), 1963–2036.
- Kreisman, D. and J. Smith (2023). Distinctively Black names and educational outcomes. *Journal of Political Economy* 131(4), 877–897.
- Kroft, K., F. Lange, and M. J. Notowidigdo (2013). Duration dependence and labor market conditions: Evidence from a field experiment. *The Quarterly Journal of Economics* 128(3), 1123–1167.
- LinkedIn (2022). LinkedIn economic graph: Career explorer. <https://web.archive.org/web/20220126182908/https://linkedin.github.io/career-explorer/>, 2022-01-26.
- LinkedIn (2023, April). About Us. <https://web.archive.org/web/20230405044221/https://news.linkedin.com/about-us/#Statistics>, 2023-06-30.
- Maćkowiak, B., F. Matějka, and M. Wiederholt (2023). Rational inattention: A review. *Journal of Economic Literature* 61(1), 226–273.
- McClintock, E. A. (2010). When does race matter? race, sex, and dating at an elite university. *Journal of Marriage and Family* 72(1), 45–72.
- McDonald, S., N. Lin, and D. Ao (2009). Networks of opportunity: Gender, race, and job leads. *Social Problems* 56(3), 385–402.
- McPherson, M., L. Smith-Lovin, and J. M. Cook (2001). Birds of a feather: Homophily in social networks. *Annual Review of Sociology* 27(1), 415–444.
- Michelman, V., J. Price, and S. D. Zimmerman (2022). Old boys’ clubs and upward mobility among the educational elite. *The Quarterly Journal of Economics* 137(2), 845–909.

- MIT Election Data and Science Lab (2018). County Presidential Election Returns 2000-2020. Dataset, Harvard Dataverse.
- News, U. (2019). U.S. news education rankings. <https://www.usnews.com/best-colleges/rankings>. Accessed on 29-Jun-2023.
- Niche (2019). The most popular college majors. <https://www.niche.com/blog/the-most-popular-college-majors/>. Accessed on 29-Jun-2023.
- Nightingale, S. J. and H. Farid (2022). Ai-synthesized faces are indistinguishable from real faces and more trustworthy. *Proceedings of the National Academy of Sciences* 119(8), e2120481119.
- Nosek, B. A., F. L. Smyth, J. J. Hansen, T. Devos, N. M. Lindner, K. A. Ranganath, C. T. Smith, K. R. Olson, D. Chugh, A. G. Greenwald, et al. (2007). Pervasiveness and correlates of implicit attitudes and stereotypes. *European Review of Social Psychology* 18(1), 36–88.
- Onuchic, P. (2023). Recent contributions to theories of discrimination. Working Paper 2205.05994, arXiv. org Papers.
- Oreopoulos, P. (2011). Why do skilled immigrants struggle in the labor market? a field experiment with thirteen thousand resumes. *American Economic Journal: Economic Policy* 3(4), 148–171.
- Pallais, A. and E. G. Sands (2016, December). Why the Referential Treatment? Evidence from Field Experiments on Referrals. *Journal of Political Economy* 124(6), 1793–1828. Publisher: The University of Chicago Press.
- Pebesma, E. (2018). Simple Features for R: Standardized Support for Spatial Vector Data. *The R Journal* 10(1), 439–446.
- Peysakhovich, A., M. A. Nowak, and D. G. Rand (2014, 9). Humans display a ‘cooperative phenotype’ that is domain general and temporally stable. *Nature Communications* 5, 4939.
- Phelps, E. S. (1972). The statistical theory of racism and sexism. *American Economic Review* 62(4), 659–661.
- Quillian, L., A. Heath, D. Pager, A. H. Midtbøen, F. Fleischmann, and O. Hexel (2019). Do some countries discriminate more than others? evidence from 97 field experiments of racial discrimination in hiring. *Sociological Science* 6, 467–496.
- Rajkumar, K., G. Saint-Jacques, I. Bojinov, E. Brynjolfsson, and S. Aral (2022). A causal test of the strength of weak ties. *Science* 377(6612), 1304–1310.
- Rand, D. G., A. Peysakhovich, G. T. Kraft-Todd, G. E. Newman, O. Wurzbacher, M. A. Nowak, and J. D. Greene (2014, 4). Social heuristics shape intuitive cooperation. *Nature Communications* 5, 3677.
- Rosenfeld, M. J., R. J. Thomas, and S. Hausen (2019, August). Disintermediating your friends: How online dating in the united states displaces other ways of meeting. *Proceedings of the National Academy of Sciences* 116(36), 17753–17758.
- Salganik, M. J. (2019). *Bit by bit: Social research in the digital age*. Princeton University Press.
- Sawilowsky, S. S. (2009, November). New effect size rules of thumb. *Journal of Modern Applied Statistical Methods* 8(2), 597–599.

- Schmutte, I. M. (2015). Job referral networks and the determination of earnings in local labor markets. *Journal of Labor Economics* 33(1), 1–32.
- Shleifer, A. (2012). Psychologists at the gate: a review of Daniel Kahneman’s thinking, fast and slow. *Journal of Economic Literature* 50(4), 1080–1091.
- Silva, B. C. and S.-O. Proksch (2021). Fake it ‘til you make it: a natural experiment to identify European politicians’ benefit from twitter bots. *American Political Science Review* 115(1), 316–322.
- Sommerville, D. M. (1995). The rape myth in the old south reconsidered. *The Journal of Southern History* 61(3), 481–518.
- Statista (2022). Statista company data base. <https://www.statista.com/companies/about,2023-02-06>.
- Suri, S. and D. J. Watts (2011, 03). Cooperation and contagion in web-based, networked public goods experiments. *PLOS ONE* 6(3), 1–18.
- Taigman, Y., M. Yang, M. Ranzato, and L. Wolf (2014). Deepface: Closing the gap to human-level performance in face verification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 1701–1708.
- Tajfel, H. and J. C. Turner (2004). *The social identity theory of intergroup behavior*. Psychology Press.
- Tibshirani, J., S. Athey, R. Friedberg, V. Hadad, D. Hirshberg, L. Miner, E. Sverdrup, S. Wager, and M. Wright (2023, May). grf: Generalized Random Forests.
- Topa, G. (2011). Labor markets and referrals. In *Handbook of Social Economics*, Volume 1, pp. 1193–1221. Elsevier.
- U.S. Bureau of Labor Statistics (2022). Table F. Distribution of private sector employment by firm size class: 1993/Q1 through 2022/Q1, not seasonally adjusted. https://www.bls.gov/web/cewbd/table_f.txt.
- U.S. Bureau of Labor Statistics (2023). BLS report: Labor force characteristics by race and ethnicity, 2021.
- U.S. Census Bureau (2013). CBSA Shapefiles. <http://meta.geo.census.gov/data/existing/decennial/GEO/GPMB/TIGERline/TIGER2013/cbsa/>.
- U.S. Census Bureau (2020). County Shapefiles. <https://www2.census.gov/geo/tiger/TIGER2020/COUNTY/>.
- U.S. Census Bureau (2021). PINC-09. Source of Income-Number with Income and Mean Income of Specified Type of People 15 Years Old and Over, by Age, Race, Hispanic Origin, and Sex. <https://www.census.gov/data/tables/time-series/demo/income-poverty/cps-pinc/pinc-09.html>.
- U.S. Census Bureau (2022). Frequently Occurring Surnames from the Census. <https://www.census.gov/topics/population/genealogy/data.html>.

- U.S. Social Security Administration (2022). Babynames. <https://www.ssa.gov/oact/babynames/>.
- Utz, S. (2016). Is LinkedIn making you more successful? the informational benefits derived from public social media. *News media & society* 18(11), 2685–2702.
- Vogels, E. A. (2021). The state of online harassment. *Pew Research Center* 13, 625.
- Wager, S. and S. Athey (2018). Estimation and inference of heterogeneous treatment effects using random forests. *Journal of the American Statistical Association* 113(523), 1228–1242.
- Wheeler, L., R. Garlick, E. Johnson, P. Shaw, and M. Gargano (2022). LinkedIn (to) job opportunities: Experimental evidence from job readiness training. *American Economic Journal: Applied Economics* 14(2), 101–25.
- Xu, F. K., N. Lofaro, B. A. Nosek, A. G. Greenwald, J. Axt, L. Simon, and N. Frost (2022). Race IAT 2002-2022.
- Yadlowsky, S., S. Fleming, N. Shah, E. Brunskill, and S. Wager (2024). Evaluating treatment prioritization rules via rank-weighted average treatment effects. *Journal of the American Statistical Association* 0(0), 1–14.
- Zounlome, N. O., Y. J. Wong, E. M. Klann, and J. L. David (2021). “I’m already seen as a sexual predator from saying hello”: Black men’s perception of sexual violence. *Journal of Interpersonal Violence* 36(19-20), NP10809–NP10830.
- Zschirnt, E. (2019). Research ethics in correspondence testing: An update. *Research Ethics* 15(2), 1–21.

Online Appendix for

LinkedOut? A Field Experiment on Discrimination in Job Network Formation

by

Yulia Evsyukova

Ⓐ

Felix Rusche

Ⓐ

Wladislaw Mill

Table of Contents

A	Experimental Design	3
A1	Overview of Surveys, Experiments, and Timeline	3
A2	Geography	4
A3	Education	5
A4	Names	7
A5	Jobs, Skills, and Volunteering	8
A6	Firms	12
A7	Process of Profile Creation	13
A8	Message	15
A9	Screenshots of LinkedIn Treatments	16
B	Picture Creation	17
C	Preparation and Structuring of Data	20
C1	Demographics	20
C2	Employment and Platform	20
C3	Employer	22
C4	Education	22
C5	Location and County Information	23
D	Demographics & Salaries: Comparative Analysis	26
E	Ethical Considerations	29
E1	Costs to LinkedIn and Participants	29
E2	Further Ethical Considerations	31
E3	Benefits of our Approach	32
E4	Ethics: A Brief Summary and Conclusion	33
F	Validation Experiment	33
F1	Design of the Validation Experiment	33

F2	Procedure	36
F3	Results	37
G	Additional Data Analysis	41
G1	Dynamic Effects	41
G2	Geographical Variation	46
G3	Variation in Profile Quality	47
G4	Predictors of Acceptance	49
G5	Drivers of Discrimination	50
G6	Some Selected Replies	64
G7	Usefulness of Messages	65
G8	Predictors of Message Response	66
G9	Effects of Picture Swapping	67
G10	Heterogeneity in Responses	71
G11	Ancillary Outcomes	75
H	Expert Survey	75
I	Survey of LinkedIn Users	80
J	Tables	94
J1	Regressions	94
J2	Main experiment – First Stage	96
J3	Main Experiment – Second Stage	107
J4	Validation Experiment	116
J5	Expert Survey	120

A Experimental Design

A1 Overview of Surveys, Experiments, and Timeline

Survey	n	Timing	Main Objective	Details in Section
Validation Experiment (MTurk)	506	March 2022	Validation of pictures and universities	F
Stage I (LinkedIn)	19,481	May to July 2022	Main experiment: Stage I	3.3
Stage II (LinkedIn)	3,319	August 2022	Main experiment: Stage II	3.4
Expert Survey (E-Mail)	269 of 2,171	June 2023	Expert prediction of results	H
LinkedIn User Survey (Prolific)	500	January 2024	Survey on LinkedIn usage and potential mechanisms	I

Table A.1: Summary of surveys conducted

Note: The table above provides an overview of the surveys and experimental stages conducted as part of this study, their sample size, timing, and purpose.

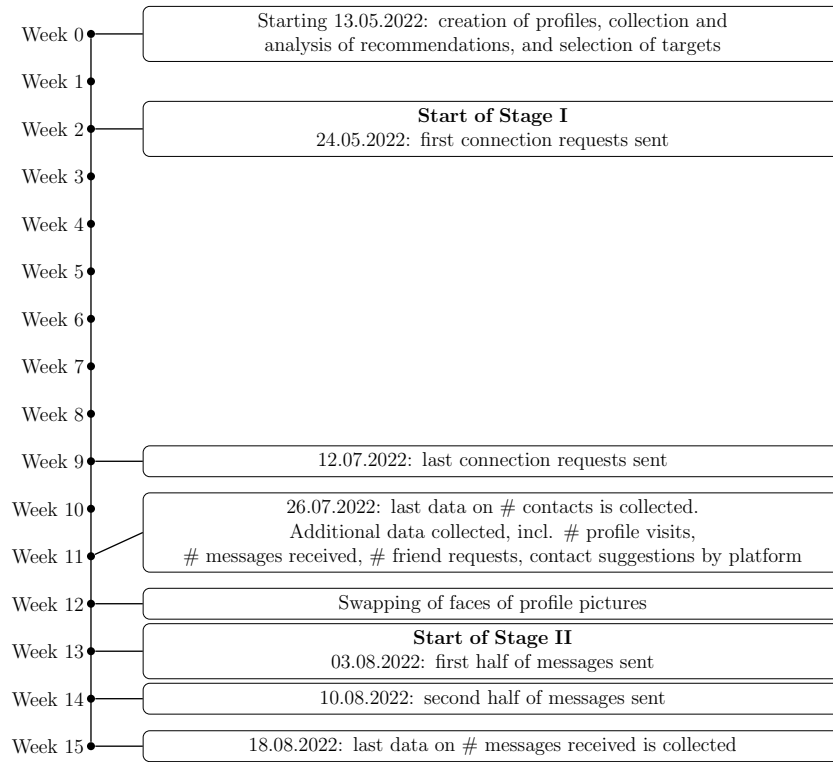


Figure A.1: Timeline of experiment

Note: Pre-registrations were done on 09.05.2022 and 28.08.2022 for the first and second stage respectively. During the first and second stage, we collect data on the main outcomes (# contacts & # messages) three times per week and end collection on 26.07.2022 and 18.08.2022 respectively. Data on targets is collected when these are sent a request to connect. Following the experiment, we answer the received messages with a short and personalized ‘thank you’ message.

A2 Geography

State	City	Population
Alaska	Anchorage	288.000
Alabama	Birmingham	209.403
Arkansas	Little Rock	197.312
Arizona	Phoenix	1.680.992
California	Los Angeles	3.979.576
Colorado	Denver	727.211
Connecticut	Bridgeport	144.399
District of Columbia	Washington	705.749
Delaware	Wilmington	70.166
Florida	Miami	467.963
Georgia	Atlanta	506.811
Hawaii	Honolulu	345.064
Iowa	Des Moines	214.237
Idaho	Boise	228.959
Illinois	Chicago	2.693.976
Indiana	Indianapolis	876.384
Kansas	Wichita	389.938
Kentucky	Louisville	617.638
Louisiana	New Orleans	390.144
Massachusetts	Boston	692.600
Maryland	Baltimore	593.490
Maine	Portland	654.741
Michigan	Detroit	670.031
Minnesota	Minneapolis	429.606
Missouri	Kansas City	495.327
Mississippi	Jackson	160.628
Montana	Billings	109.577
North Carolina	Charlotte	885.708
North Dakota	Fargo	124.662
Nebraska	Omaha	478.192
New Hampshire	Manchester	112.673
New Jersey	Newark	282.011
New Mexico	Albuquerque	560.513
Nevada	Las Vegas	651.319
New York	New York City	8.336.817
Ohio	Columbus	898.553
Oklahoma	Oklahoma City	655.057
Oregon	Portland	654.741
Pennsylvania	Philadelphia	1.584.064
Rhode Island	Providence	179.883
South Carolina	Charleston	137.566
South Dakota	Sioux Falls	183.793
Tennessee	Nashville	670.820
Texas	Houston	2.320.268
Utah	Salt Lake City	200.567
Virginia	Virginia Beach	449.974
Vermont	Burlington	42.819
Washington	Seattle	753.675
Wisconsin	Milwaukee	590.157
West Virginia	Charleston	137.566
Wyoming	Cheyenne	64.235

Table A.2: Cities where experiment is run

Note: We choose the biggest city in each U.S. State according to [U.S. Census 2019](#) estimates. In Florida, we replace Jacksonville with Miami.

A3 Education

University	Niche Ranking	Forbes	US News	Enrollment	City	State of Profile	If none in State, which other
University of North Alabama	No	No	No	5k	Mobile	Alabama	
Peninsula College	No	No	No	1k	Port Angeles	Alaska	Washington
University of Phoenix - Arizona	No	No	299-391	72k	Phoenix	Arizona	
University of Central Arkansas	No	542	299-391	8k	Conway	Arkansas	
Dominican University of California	No	572	No	1k	San Rafael	California	
University of Northern Colorado	No	444	No	12k	Greely	Colorado	
Sacred Heart University	No	526	No	5k	Fairfield	Connecticut	
Delaware State University	No	No	No	4k	Dover	Delaware	
Radford University	No	465	No	9k	Radford	Washington DC	Virginia
Barry University	No	570	No	3.5k	Miami	Florida	
University of Montevallo	No	No	No	2k	Montevallo	Georgia	Alabama
Whittier College	No	567	No	1.5k	Whittier	Hawaii	California
Eastern Oregon University	No	No	No	1.7k	Pocatello	Idaho	Oregon
Concordia University Chicago	No	No	No	1.5k	River Forest	Illinois	
University of Akron	No	591	299-391	12k	Akron	Indiana	Ohio
University of Northern Iowa	No	457	No	12k	Indianola	Iowa	
Rogers State University	No	No	No	2k	Claremore	Kansas	Oklahoma
Western Kentucky University	No	521	299-391	12k	Bowling Green	Kentucky	
McNeese State University	No	No	No	5k	Lake Charles	Louisiana	
Worcester State University	No	573	No	4k	Worcester	Maine	Massachusetts
Mount St. Mary's University	No	589	No	2k	Emmitsburg	Maryland	
Assumption University	No	559	No	2k	Worcester	Massachusetts	
Central Michigan University	No	454	No	13k	Mount Pleasant	Michigan	
Minnesota State University Moorhead	No	No	No	4k	Moorhead	Minnesota	
Delta State University	No	No	No	2k	Cleveland	Mississippi	
University of Central Missouri	No	530	No	8k	Warrensburg	Missouri	
Snow College	No	No	No	3k	Ephraim	Montana	Utah
Peru State College	No	No	No	1k	Peru	Nebraska	
Great Basin College	No	No	No	1k	Elko	Nevada	
Saint Anselm College	No	477	No	2k	Manchester	New Hampshire	
Saint Peter's University	No	531	No	2k	Jersey City	New Jersey	
Bryan University - Tempe	No	No	No	1k	Tempe	New Mexico	Arizona
SUNY Oneonta	No	527	No	6.5k	Oneonta	New York	
University of North Carolina Asheville	No	No	No	3.9k	Asheville	North Carolina	
Crown College	No	No	No	1k	Saint Bonifacius	North Dakota	Minnesota
Cleveland State University	No	474	299-391	17k	Cleveland	Ohio	
Mid-America Christian University	No	No	No	1k	Oklahoma City	Oklahoma	
Southern Oregon University	No	519	No	3k	Forest Grove	Oregon	
Washington & Jefferson College	No	480	No	1k	Washington	Pennsylvania	
Lasell University	No	No	No	1.6k	Newton	Rhode Island	Massachusetts
University of South Carolina - Beaufort	No	No	No	1.7k	Bluffton	South Carolina	
Waldorf University	No	No	No	1.6k	Brookings	South Dakota	Iowa
Carson-Newman University	No	No	299-391	1.5k	Jefferson City	Tennessee	
West Texas A&M University	No	579	No	10k	Canyon	Texas	
Dixie State University	No	No	No	6.5k	Saint George	Utah	
SUNY Oswego	No	529	No	6.6k	Oswego	Vermont	New York
Marymount University	No	553	No	2k	Arlington	Virginia	
Eastern Washington University	No	575	No	13k	Cheney	Washington	
Walsh University	No	No	No	2k	North Canton	West Virginia	Ohio
Illinois College	No	No	No	1k	Platteville	Wisconsin	Illinois
Fort Lewis College	No	No	299-391	4k	Durango	Wyoming	Colorado

Table A.3: Universities for low-ranked education profile

Note: All universities in the ranking are present in Niche’s Business and Management Category. Not all of them have a rank in the website’s “Best Colleges for Business in America 2022” ranking. The table further includes each institution’s rank in [Forbes \(2021\)](#) 600 ranking and [US News Ranking](#). In some cases, no high-ranked university is available from a given state. In this case, we choose both a high- and low-ranked university from a neighboring state. If a state has several suitable “neighbors”, we proceed by selecting a high-ranked university that is closest to the biggest city in the target state. We choose among universities that are second best ranked within the respective state (first best-ranked universities are assigned to the profiles within the respective state). For the low type, we then choose a suitable university from the same state. We also present information on the enrollment at each institution ([News, 2019](#); [Niche, 2019](#)).

University	Niche Ranking	Forbes	US News	Enrollment	City	State of Profile	If none in State, which other
The University of Alabama	111	233	148	29k	Tuscaloosa	Alabama	
Washington State University	128	175	179	23k	Pullman	Alaska	Washington
Arizona State University	74	121	117	41k	Tempe	Arizona	
University of Arkansas	138	190	162	21k	Fayetteville	Arkansas	
University of San Diego	80	132	93	8k	San Diego	California	
University of Denver	127	165	93	5k	Denver	Colorado	
University of Connecticut	233	70	63	18k	Storrs	Connecticut	
University of Delaware	96	108	93	18k	Newark	Delaware	
George Mason University	265	91	148	22k	Fairfax	Washington DC	Virginia
Florida International University	72	145	162	28k	Tampa	Florida	
Samford University	180	250	136	4k	Birmingham	Georgia	Alabama
Loyola Marymount University	88	124	75	9k	Los Angeles	Hawaii	California
University of Oregon	166	144	99	7k	Eugene	Idaho	Oregon
Loyola University Chicago	145	220	103	12k	Chicago	Illinois	
Miami University	165	120	55	17k	Oxford	Indiana	Ohio
University of Iowa	94	118	83	22k	Iowa City	Iowa	
Oklahoma State University	85	204	187	17k	Stillwater	Kansas	Oklahoma
University of Kentucky	142	209	127	21k	Lexington	Kentucky	
Tulane University	69	119	42	7k	New Orleans	Louisiana	
Brandeis University	189	128	42	3k	Waltham	Maine	Massachusetts
Loyola University Maryland	92	210	No	4k	Baltimore	Maryland	
University of Massachusetts - Amherst	77	141	68	22k	Amherst	Massachusetts	
Kalamazoo College	266	172	No	1.5k	Kalamazoo	Michigan	
Gustavus Adolphus College	140	264	No	2k	Saint Peter	Minnesota	
University of Mississippi	271	221	148	16k	University	Mississippi	
Saint Louis University	174	176	103	7k	Saint Louis	Missouri	
Utah State University	192	267	249	17k	Logan	Montana	Utah
University of Nebraska - Lincoln	197	193	136	19k	Lincoln	Nebraska	
University of Nevada - Reno	251	236	227	15k	Reno	Nevada	
University of New Hampshire	121	244	136	12k	Durham	New Hampshire	
Stevens Institute of Technology	179	158	83	4k	Hoboken	New Jersey	
University of Arizona	258	127	103	29k	Tucson	New Mexico	Arizona
Syracuse University	82	113	59	15k	Syracuse	New York	
University of North Carolina - Wilmington	185	269	187	12k	Wilmington	North Carolina	
University of St. Thomas - Minnesota	146	213	136	6k	Collegeville	North Dakota	Minnesota
John Carroll University	126	273	No	3k	University Heights	Ohio	
University of Oklahoma	75	125	127	21k	Norman	Oklahoma	
University of Portland	115	157	No	4k	Portland	Oregon	
Temple University	83	200	103	26k	Philadelphia	Pennsylvania	
Worcester Polytechnic Institute	272	135	63	5k	Worcester	Rhode Island	Massachusetts
Furman University	200	171	No	3k	Greenville	South Carolina	
Iowa State University	115	156	122	27k	Ames	South Dakota	Iowa
University of Tennessee	118	161	103	22k	Knoxville	Tennessee	
Baylor University	97	205	75	14k	Waco	Texas	
University of Utah	190	95	99	19k	Salt Lake City	Utah	
Skidmore College	89	170	No	3k	Saratoga Springs	Vermont	New York
James Madison University	249	96	No	20k	Harrisonburg	Virginia	
Gonzaga University	68	331	79	5k	Spokane	Washington	
Denison University	131	288	No	2k	Granville	West Virginia	Ohio
Wheaton College - Illinois	183	211	No	2k	Mequon	Wisconsin	Illinois
Colorado State University	147	199	148	25k	Fort Collins	Wyoming	Colorado

Table A.4: Universities for high-ranked education profile

Note: All universities in the ranking are ranked between the 68th and 272th place in Niche’s “Best Colleges for Business in America 2022” ranking. In cases where no high-ranked university from the respective state is available in Niche’s Ranking, we substitute with a university from a neighboring state, as indicated by the last column. If a state has several suitable “neighbors”, we proceed by selecting a high-ranked university that is closest to the biggest city in the target state. We choose among universities that are second best ranked within the respective state (first best-ranked universities are assigned to the profiles within the respective state). For the low type, we then choose a suitable university from the same state. The table further includes each institution’s rank in the Forbes 600 ranking and US News Ranking. We also present information on the enrollment at each institution

A4 Names

Name	Births White	% of White Births	Births Black	% of Black Births	Rank US
CHRISTOPHER	765	2	280	1.4	4
JOSHUA	662	1.7	278	1.4	5
BRANDON	551	1.4	285	1.4	8
MICHAEL	757	2	224	1.1	1
JORDAN	260	0.7	194	1	26
ANTHONY	216	0.6	180	0.9	18
JUSTIN	435	1.1	166	0.8	20
JAMES	682	1.8	135	0.7	17
TYLER	543	1.4	118	0.6	10
NICHOLAS	506	1.3	112	0.6	6

Table A.5: First names of profiles

Note: We obtain the most common first names of men born 1997 in Georgia from [Georgia Department of Public Health \(2022\)](#). We then focus on the names that are within the top 50 most common names for *both* White and Black men, i.e., the intersection of popular Black and popular White names. For these remaining names, we sort by popularity among Black Americans and take the 10 most popular ones. Aside from the number of share of births by race in Georgia in 1997, we also report the rank of the first name for all baby names in 1997 from [U.S. Social Security Administration \(2022\)](#). All chosen baby names are within the top 30 in the US in 1997.

No.	Name	Share White	Share Black	US Rank	Frequency (count)	name per 100k population
1	BANKS	39.3	54.5	292	105,833	35.9
2	JOSEPH	29.6	54.2	313	100,959	34.2
3	MOSLEY	40.5	53.2	730	47,963	16.3
4	JACKSON	39.9	53	19	708,099	240.1
5	CHARLES	33.7	53	548	61,211	20.8
6	DORSEY	41.8	52.2	793	43,631	14.8
7	RIVERS	40.5	50.9	897	38,662	13.1
8	GAINES	42.9	50.7	788	43,821	14.9
9	MAYS	54.8	39.7	854	40,408	13.7
10	WIGGINS	54.7	39.6	685	50,247	17
11	DIXON	54.3	39.3	167	159,480	54.1
12	FLOWERS	53.1	40.3	578	57,549	19.5
13	THOMAS	52.6	38.8	16	756,142	256.3
14	TERRELL	55.30	39	983	35,408	12
15	ROBERSON	51.3	42.8	605	56,180	19.1
16	BENJAMIN	49	41.6	850	40,590	13.8

Table A.6: Surnames of profiles

Note: We obtained the most common US last names from [U.S. Census Bureau \(2022\)](#). We choose names that are roughly equally likely to be of a Black and White individual and unlikely to be of any other race. We aimed to have a similar rank and proportion per 100,000 population across names. We further choose names that are relatively common.

A5 Jobs, Skills, and Volunteering

Job Title	Average	10%	90%
Office Manager	48,971	34,000	70,000
Buyer	56,005	42,000	76,000
Administrative Assistant	39,968	29,000	57,000
Office Administrator	47,077	32,000	77,000
Marketing Assistant	38,949	30,000	51,000

Table A.7: Job titles and average pay according to [Payscale.com](https://www.payscale.com) (2022)

Job	Description Items
Office Manager	
Description 1	<ol style="list-style-type: none"> 1. Perform methodological and extensive preparation of financial reports, management reports, and ad hoc reporting 2. Identify business challenges and shaped effectual benchmarked solutions in meeting companies objectives 3. Function as primary liaison to customers and ensured a consistently positive customer experience 4. Regularly assess office productivity and making team adjustments as needed
Description 2	<ol style="list-style-type: none"> 1. Oversee diverse roles in accounting, HR, finance, logistics and sales operation while implementing strategies 2. Facilitate information management while effectively collaborating with the CEO for operational improvements 3. Implement and maintained company protocols to ensure smooth daily activities 4. Direct all office staff in the processing and submitting of payroll
Office Administrator	
Description 1	<ol style="list-style-type: none"> 1. Develop relationships with customers, vendors, and guests to present the company in a professional manner. 2. Support office staff by organizing company events, meetings, and scheduling. 3. Release reports and other data requested by accounting, sales and warehouse departments 4. Create PowerPoint presentations used for business development
Description 2	<ol style="list-style-type: none"> 1. Provide strategic administrative and development support 2. Design electronic file systems and maintained electronic and paper files 3. Draft meeting agendas, supply advance materials, and execute follow-up for meetings and team conferences 4. Properly route agreements, contracts and invoices through the signature process
Buyer	
Description 1	<ol style="list-style-type: none"> 1. Worked with internal customers to gain a deep understanding of supply needs. 2. Analyzed price proposals, conducted detailed performance reports, and developed and co-managed annual purchasing budget. 3. Assisted in the strategic sourcing management, identified and evaluated potential suppliers and business partners, and negotiated contracts. 4. Responsible for the placement, management, and data entry of purchase orders.
Description 2	<ol style="list-style-type: none"> 1. Monitor and analyze everyday business operations, purchased quality goods for the company, and managed and monitored inventories. 2. Serve as point of contact for vendors and other buyers with questions about purchase order discrepancies 3. Conduct research to formulate new sales strategies. 4. Maintain and updated daily retail purchase records for submission to senior buyer.
Administrative Assistant	
Description 1	<ol style="list-style-type: none"> 1. Developed positive relations with external vendors and clients 2. Streamlined processes to effectively track, order, and maintain inventory 3. Oversaw calendar maintenance, appointment scheduling and expense report preparation 4. Compose and proofread memos, letters, reports, and presentations, providing accurate, concise, and error-free communication
Description 2	<ol style="list-style-type: none"> 1. Manage executive calendars, strategically coordinating meetings, appointments, events, and travel arrangements. 2. Strategically manage complex calendars, organizing meetings, appointments, and travel arrangements, and proactively identifying and adjusting conflicting events 3. Extract information from registrations, applications and executed contracts, contract information and action memorandum 4. Greet and proactively assist visitors in a timely manner
Marketing Assistant	
Description 1	<ol style="list-style-type: none"> 1. Helped to coordinate client reports at the end of each study and also helped audit final information. 2. Utilized time tracking software for accurate project and time management. 3. Assisted with development and implementation of marketing strategies. 4. Keep the marketing database up-to-date by inputting new data, updating old records and performing cross checks
Description 2	<ol style="list-style-type: none"> 1. Use lead generation software to create organised lists of prospective customers. 2. Coordinate a wide range of marketing communications. 3. Prepare company documents, proposals, reports and presentations. 4. Carry out the daily administrative tasks that keep the marketing department functioning.

Table A.8: Job descriptions

Job descriptions are taken from CV examples on websites like [ideed.com](https://www.Indeed.com), [monster.com](https://www.monster.com), etc. We exclude descriptions that are company- or industry-specific. Each description contains information from multiple example-resumés.

No.	Buyer	Office Manager	Administrative Assistant	Marketing Assistant	Office Administrator
1	Purchasing	Office Administration	Administrative Assistance	Social Media Marketing	Office Administration
2	Procurement	QuickBooks	Office Administration	Marketing	Administrative Assistance
3	Inventory Management	Accounts Payable	Data Entry	Social Media	QuickBooks
4	Supply Chain Management	Accounts Receivable (AR)	Event Planning	Digital Marketing	Data Entry
5	Retail Buying	Payroll	Administration	Adobe Photoshop	Accounts Payable
6	Merchandising	Administrative Assistance	Time Management	Facebook	Accounts Receivable (AR)
7	Negotiation	Invoicing	Customer Service	Adobe InDesign	Invoicing
8	Strategic Sourcing	Data Entry	Social Media	Email Marketing	Administration
9	Retail	Bookkeeping	Research	Event Planning	Payroll
10	Forecasting	Human Resources (HR)	Teamwork	Advertising	Event Planning
11	Manufacturing	Accounting	Phone Etiquette	Adobe Illustrator	Time Management
12	Material Requirements Planning (MRP)	Customer Service	Executive Administrative Assistance	Marketing Strategy	Customer Service
13	Continuous Improvement	Event Planning	Organization Skills	Teamwork	Human Resources (HR)
14	Visual Merchandising	Budgeting	QuickBooks	Adobe Creative Suite	Bookkeeping
15	Product Development	Sales	Microsoft Access	Google Analytics	Social Media
16	Trend Analysis	Office Operations	Public Speaking	Graphic Design	Phone Etiquette
17	Lean Manufacturing	Team Building	Travel Arrangements	Time Management	Sales
18	Inventory Control	Administration	Clerical Skills	WordPress	Accounting
19	Fashion	Accounts Payable & Receivable	Community Outreach	Public Relations	Microsoft Access
20	Apparel	Time Management	Nonprofit Organizations	Search Engine Optimization (SEO)	Marketing

Table A.9: Skills assigned to profiles

Note: To each profile, we randomly assign five of the 20 most commonly mentioned skills by platform users with the respective job title. We obtain this information directly from LinkedIn's Economic Graph Career Explorer ([LinkedIn, 2022](#)).

Organization	American Red Cross
Role	Blood Donor Ambassador
Cause	Health
Description	Engaged in promoting and enhancing blood donation process via communication with donors.
Organization	American Red Cross
Role	Blood Donor Ambassador
Cause	Health
Description	Provided organisational support in blood donation process, ensured comfort and safety of donors.
Organization	American Red Cross
Role	Blood Donor Ambassador
Cause	Health
Description	Maintained blood donation process, promoted blood donation commitment of donors.
Organization	Big Brothers and Big Sisters of America
Role	Volunteer Big Brother
Cause	Children
Description	Acted as a mentor of a child by providing guidance and support to the Little.
Organization	Big Brothers and Big Sisters of America
Role	Volunteer Big Brother
Cause	Children
Description	Served as a positive role model for at-risk youth, guiding through activities.
Organization	Big Brothers and Big Sisters of America
Role	Volunteer Big Brother
Cause	Children
Description	Mentored a child by building relationships based on trust and providing support and encouragement to my little brother.
Organization	Crisis Text Line
Role	Volunteer Crisis Counselor
Cause	Disaster and Humanitarian Relief
Description	Provided psychological support to people who were facing mental health issues like depression, anxiety, bullying, among others, via text messaging.
Organization	Crisis Text Line
Role	Volunteer Crisis Counselor
Cause	Disaster and Humanitarian Relief
Description	Involved in text communication with individuals in crisis, providing them mental and emotional support, assisting in developing an action plan to cope with a current crisis.

Table A.10: Volunteer work indicated in profile

Note: descriptions are taken from CV examples on websites like [ideed.com](https://www.ideed.com), [monster.com](https://www.monster.com), etc.

A6 Firms

To obtain employers, we first used [Statista's Company Data Base](#) to identify the largest employers in each city. If the city is unique in the USA, we use the largest employers as our companies. For cities with too few employers or cities with multiple mentionings, we search for local information on the largest employers. We use the following sources (click on the source to get to the website):

- [Jackson \(MS\)](#)
- [Portland \(OR\)](#) [Source1](#) and [Source 2](#)
- [Providence \(RI\)](#)
- [Sioux Falls \(SD\)](#)
- [Nashville \(TN\)](#)
- [Burlington \(VT\)](#)
- [Cheyenne \(WY\)](#)
- [Charleston \(SC\)](#)
- [Charlotte \(NC\)](#)
- [Wilmington \(DE\)](#)

We further tried to avoid the following employers in general: Universities, school districts, hospitals (only if sufficiently many employers were found), and religious institutions. We further tried to avoid similar-sounding companies (Liberty Mutual Insurance Company; Liberty Mutual Holding Company Inc.; Liberty Mutual Group Inc.).

The resulting firms are shown in Table [A.11](#).

No	State	Employer	No	State	Employer	No	State	Employer
1	AL	Encompass Health Corp	69	KY	Kentucky Hospital	137	ND	Wells Fargo & Co.
2	AL	Hibbett Sports Inc	70	KY	Yum Brands Inc.	138	ND	Sanford
3	AL	Onin Staffing, LLC	71	KY	Pharmacia Corporation	139	ND	Rdo Holdings Co.
4	AL	Questor Partners Fund II, L.P.	72	KY	Humana Inc.	140	ND	Titan Machinery Inc
5	AK	Asrc Energy Services, LLC	73	LA	Southern Theatres, L.L.C.	141	OH	St Francis Health, LLC
6	AK	Afognak Native Corporation	74	LA	Jazz Casino Company, L.L.C.	142	OH	Couche-Tard U.S. Inc
7	AK	Saexploration, Inc.	75	LA	Weiser Security Services, Inc.	143	OH	Express Topco LLC
8	AK	Veco Corporation	76	LA	Vss-Southern Theatres LLC	144	OH	American Electric Power Company Inc.
9	AZ	Phoenix Parent Holdings Inc.	77	ME	WEX LLC	145	OK	Braum's, Inc.
10	AZ	Avnet Inc.	78	ME	Unum	146	OK	Integrus Health, Inc.
11	AZ	Knight Transportation, Inc	79	ME	Td Bank US Holding Company	147	OK	Chesapeake Operating, L.L.C.
12	AZ	ON Semiconductor Corp.	80	ME	Amatos	148	OK	Devon Oei Operating, Inc.
13	AR	Dillard's Inc.	81	MD	Edge Acquisition, LLC	149	OR	Precision Castparts Corp.
14	AR	Baptist Health	82	MD	Abacus Corporation	150	OR	Columbia Sportswear Co.
15	AR	Montaire Corporation	83	MD	T. Rowe Price Group Inc.	151	OR	Esco Group LLC
16	AR	Windstream Services	84	MD	Dla Piper LLP	152	OR	Legacy Health
17	CA	Lowe Enterprises, Inc.	85	MA	Fmr LLC	153	PA	Independence Health Group, Inc.
18	CA	AECOM	86	MA	Mass General Brigham Incorporated	154	PA	Aramark
19	CA	Guess Inc.	87	MA	National Financial Services LLC	155	PA	Comcast Corp
20	CA	Forever 21, Inc.	88	MA	General Electric Co.	156	PA	Axalta Coating Systems Ltd
21	CO	Gates Industrial Corporation plc	89	MI	Henry Ford Health System	157	RI	conrail inc
22	CO	Digital First Media, LLC	90	MI	Vhs of Michigan, Inc.	158	RI	Lifespan Finance
23	CO	Aimco Properties, L.P.	91	MI	Michigan Bell Telephone Company	159	RI	San Francisco Toyota
24	CO	The Anschutz Corporation	92	MI	DTE Energy Co.	160	RI	Dsi, Inc.
25	CT	St. Vincent'S Health Services Corporation	93	MN	Buffalo Wild Wings, Inc.	161	SC	Ingevity
26	CT	Xylem Dewatering Solutions, Inc.	94	MN	General Mills, Inc.	162	SC	Volvo Car USA LLC
27	CT	Goodwill of Western & Northern Connecticut, Inc.	95	MN	Medtronic Usa, Inc.	163	SC	Igor
28	CT	Schrader-Bridgeport International Inc.	96	MN	Target Corp	164	SC	Nucor Steel
29	DE	AstraZeneca	97	MS	Nissan	165	SD	Citi
30	DE	ING Direc	98	MS	Delphi Auto Systems	166	SD	Sanford Health
31	DE	Bank of America	99	MS	Cal-Maine Foods	167	SD	Billion Automotive Companies
32	DE	Delmarva Power/PEPCO	100	MS	Kroger	168	SD	Meta Financial Group
33	DC	Danaher Corporation	101	MO	Dst Systems, Inc.	169	TN	Randstad
34	DC	Fannie Mae	102	MO	Reorganized Fli, Inc.	170	TN	HCA Healthcare Inc.
35	DC	Kipp DC	103	MO	Cerner Corp.	171	TN	The Kroger Co.
36	DC	FTI Consulting	104	MO	Burns & McDonnell, Inc.	172	TN	Bridgestone Americas
37	FL	Freepart-Memoran Miami Inc.	105	MT	First Interstate BancSystem Inc.	173	TX	National Oilwell Varco Inc.
38	FL	Lennar Corp.	106	MT	Talen Montana, LLC	174	TX	Sysco
39	FL	Norwegian Cruise Line Holdings Ltd	107	MT	The Tire Guys Inc	175	TX	Baker Hughes Co
40	FL	Lenzing AG	108	MT	Kamgrounds of America, Inc.	176	TX	Schlumberger Limited
41	GA	UHS of Peachford LP	109	NE	HDR Engineering, Inc.	177	UT	Overstock
42	GA	Home Depot, Inc.	110	NE	Hdr, Inc.	178	UT	Avalon Health Care, Inc.
43	GA	Coca	111	NE	Peter Kiewit Sons', Inc.	179	UT	AlSCO Inc.
44	GA	Delta Air Lines, Inc.	112	NE	Intrado Corporation	180	UT	SendOutCards
45	HI	Hawaiian Airlines, Inc.	113	NV	Cannae Holdings Inc	181	VT	G.S. Blodgett Company
46	HI	Hawaiian Electric Industries, Inc.	114	NV	MGM Resorts International	182	VT	Gardener's Supply
47	HI	The Queen's Health Systems	115	NV	Mandalay Resort Group	183	VT	Bruegger's Enterprises
48	HI	Td Food Group, Inc.	116	NV	Las Vegas Sands, LLC	184	VT	IDX systems
49	ID	American Stores Company, LLC	117	NH	Elliot Health System	185	VA	Naval Air Station Oceana-Dam Neck
50	ID	Wincos Foods, LLC	118	NH	Easter Seal New Hampshire, Inc.	186	VA	Amerigroup (Anthem)
51	ID	Wincos Holdings, Inc.	119	NH	Legacy Echm, Inc.	187	VA	DOMA Technologies
52	ID	AB Acquisition LLC	120	NH	Bob's Discount Furniture, LLC	188	VA	Lockheed Martin Corporation
53	IL	Mondelez International Inc.	121	NJ	Black & Decker Inc.	189	WA	Amazon.com Inc.
54	IL	Boeing Co.	122	NJ	Ecco, Inc.	190	WA	Starbucks Corp.
55	IL	Commonspirit Health	123	NJ	Prudential Financial Inc.	191	WA	Carrix, Inc.
56	IL	AON Corporation	124	NJ	Pruco Securities, LLC	192	WA	Safeco
57	IN	Lilly(Eli) & Co	125	NM	Laguna Development Corporation	193	WV	AMFM
58	IN	Anthem Insurance Companies, Inc.	126	NM	Optumcare New Mexico, LLC	194	WV	Eastern Associated Coal
59	IN	Steak N Shake Inc.	127	NM	National Technology & Engineering Solutions of Sandia, LLC	195	WV	Dow Chemical Co
60	IN	American United Mutual Insurance Holding Company	128	NM	PNM Resources Inc	196	WV	Thomas Health
61	IA	Catholic Health Initiatives - Iowa, Corp.	129	NY	JPMorgan Chase	197	WI	Aurora Health Care, Inc.
62	IA	Berkshire Hathaway Energy Company	130	NY	Pfizer	198	WI	Marcus Corp.
63	IA	Allied Group, Inc	131	NY	Philip Morris International	199	WI	Johnson Controls
64	IA	Meredith Corp.	132	NY	Christian Dior	200	WI	Ascension Wisconsin
65	KS	Restaurant Management Company of Wichita, Inc.	133	NC	Goodrich Corporation	201	WY	Union Pacific Railroad
66	KS	Learjet Inc.	134	NC	Compass Group USA	202	WY	Echo Star Communications
67	KS	Ascension Via Christi Health, Inc	135	NC	JELD	203	WY	Sinclair Marketing, Inc.
68	KS	Koch Industries, Inc.	136	NC	Nucor Corp.	204	WY	Wallick & Volk, Inc.

Table A.11: List of employers indicated in profiles

A7 Process of Profile Creation




Figure A.2: Profile creation: example for Birmingham (Alabama)

Note: The graph describes the profile creation process. As described in the text, job titles and companies are assigned without replacement within a given city/state. Further, for each state, we collect one more prestigious and one less prestigious university. Finally, pictures are assigned without replacement across the entire collection of pictures.

A8 Message

Below we provide the messages that the targets receive as the experimental treatment.

Treatment 1: Job-Application Message

“Hi {YOUR NAME}, Thanks for accepting my connection. I’m thinking of applying as an {POSITION} at {COMPANY NAME} and would really appreciate your advice. For instance, are there any qualities your company is particularly looking for in applicants? And are there any pitfalls to avoid during the interview process? I want to make sure that my application stands out and gets noticed. Thank you for your time. I hope to hear from you soon.”

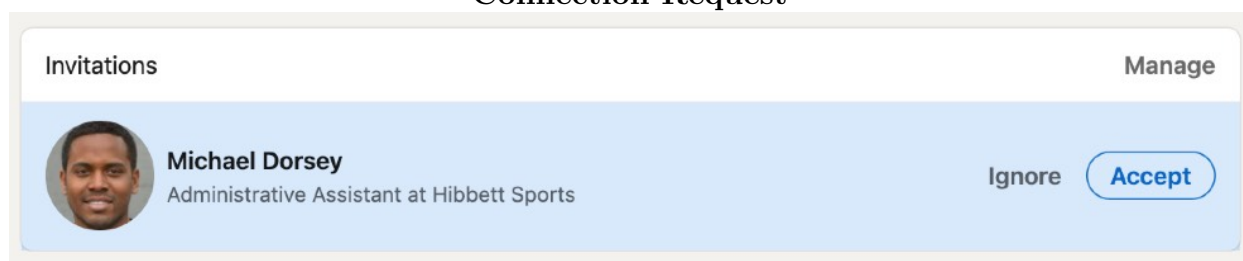
Treatment 2: Mentorship Message

“Hi {YOUR NAME}, Thanks for accepting my connection. As a young professional, I am currently trying to build a professional network and I’m looking for career advice. Do you have any insights on how to succeed in this business? For instance, do you have any recommendations on what kind of skills and qualities to acquire or develop? And are there any particular pitfalls to avoid? Thank you for your time. I hope to hear from you soon.”

In the messages, {YOUR NAME} and {COMPANY NAME} are replaced by the first name of a target and the name of the company she works at, respectively. {POSITION} is replaced by the job position of the contacting profile.

A9 Screenshots of LinkedIn Treatments

Connection Request



Message Received

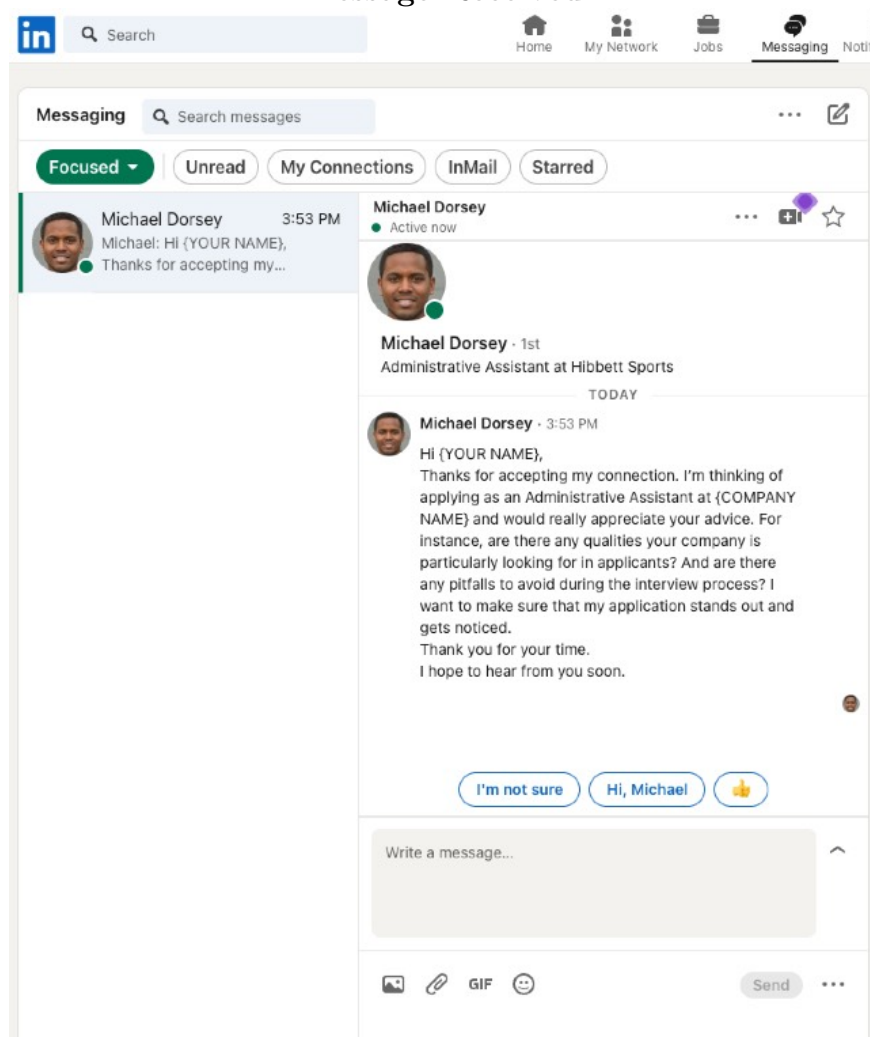


Figure A.3: Example screenshots of the connection request and messaging windows on LinkedIn

Note: In the messages to real users, {YOUR NAME} and {COMPANY NAME} are replaced by the first name of a target and the name of the company she works at, respectively. Further, the name, firm, and job position of the profile sending the message are adjusted to the respective profile's CV.

B Picture Creation

To signal race, this study creates pictures and an algorithm that can transform pictures’ race, while holding other characteristics stable. This section aims to explain the procedure to create pictures. The creation has two aims: first, we provide each twin pair with a unique input image, which is then transformed into the other race. A unique image is obtained to ensure that the results are not driven by specific pictures’ characteristics. Second, half of the input images in each state should be Black and half White. This guarantees that the results are not due to any bias introduced by the transformation algorithm. Thus, overall, 102 Black and 102 White input images are required, which are then transformed to create 408 unique pictures. All operations to obtain these are based on NVIDIA’s StyleGAN2, an image modeling algorithm (Karras et al., 2020).

The picture creation and validation process, as visualized in Figure B.1, is done in seven steps:

1. First, we obtain 100,000 AI-generated images provided by the creators of StyleGAN2 (Karras et al., 2020)
2. These are sorted using DeepFace (Taigman et al., 2014), a facial recognition algorithm, to obtain information on the age, ethnicity, and gender of each image. We use these characteristics to select pictures that fit the target group of young Black and White men. We find a total of 157 Black and 1652 White suitable images. This strong bias is likely driven by StyleGAN’s training data, which is primarily made up of White and only very few Black individuals. We sort through the 70k training images using DeepFace (Taigman et al., 2014) and find that around 4.9% of images are Black, while 57.4% are classified as White.
3. Next, we go through the Black images by hand and sort out misclassifications, such as images representing women, older individuals, children, or pictures with weird deformations. This leaves us with a total of 42 Black images. We select a similar number of White images.
4. Given that 102 pictures of each race are required to create a unique picture for each profile pair, we use the images obtained through the procedure described above to create additional ones. More specifically, we, first, utilize StyleGAN2 to represent each image as a latent vector. Using these, we create ‘grandchildren’ of the input images, meaning that we calculate the average vector for each of four unique picture combinations of the same race. To ensure that pictures do not look too similar, we only create grandchildren that share at most two ‘grandparents’ with any other picture created. We do so until we obtain a total of 2,310 pictures of each race.
5. These images are then transformed into the other race. We do so using a simple algorithm that does not require us to define race features. More specifically, we simply take the 42 Black and 51 White images’ vector representations from Step (2) and calculate the average vectors for Black and White images. We then take the difference between the average White and Black image to obtain a transformation vector. Simply adding this difference to a Black image results in a White one. Similarly, subtracting it from a White image results in a transformation to a Black one.
6. We use the vector to translate all 4,620 images obtained in Step (4) to the other race.
7. Given that we only need 204 pairs of Black and White images, we next analyze the pictures using DeepFace (Taigman et al., 2014) regarding their gender, age, race, etc. and choose pairs that are most similar to one another in characteristics other than race (Taigman et al., 2014). This results in around 700 images that we use for further analysis.

8. Finally, these images are evaluated by humans using Amazon MTurk (the survey experiment is described in Chapter F). Only images that have the smallest difference between the potential White and Black profile in terms of picture characteristics are used in the final sample.³⁴

³⁴Note that reducing the difference between the potential White and Black profile does not mean we reduce the distance between Black and White profiles in terms of race. In fact, the difference in how Black the picture is considered between the White and Black profile pictures is 65.11 before and 66.18 after we exclude all those pairs with the biggest difference in terms of non-race picture characteristics ($t(359.9) = 0.7$, $p = 0.4795$).

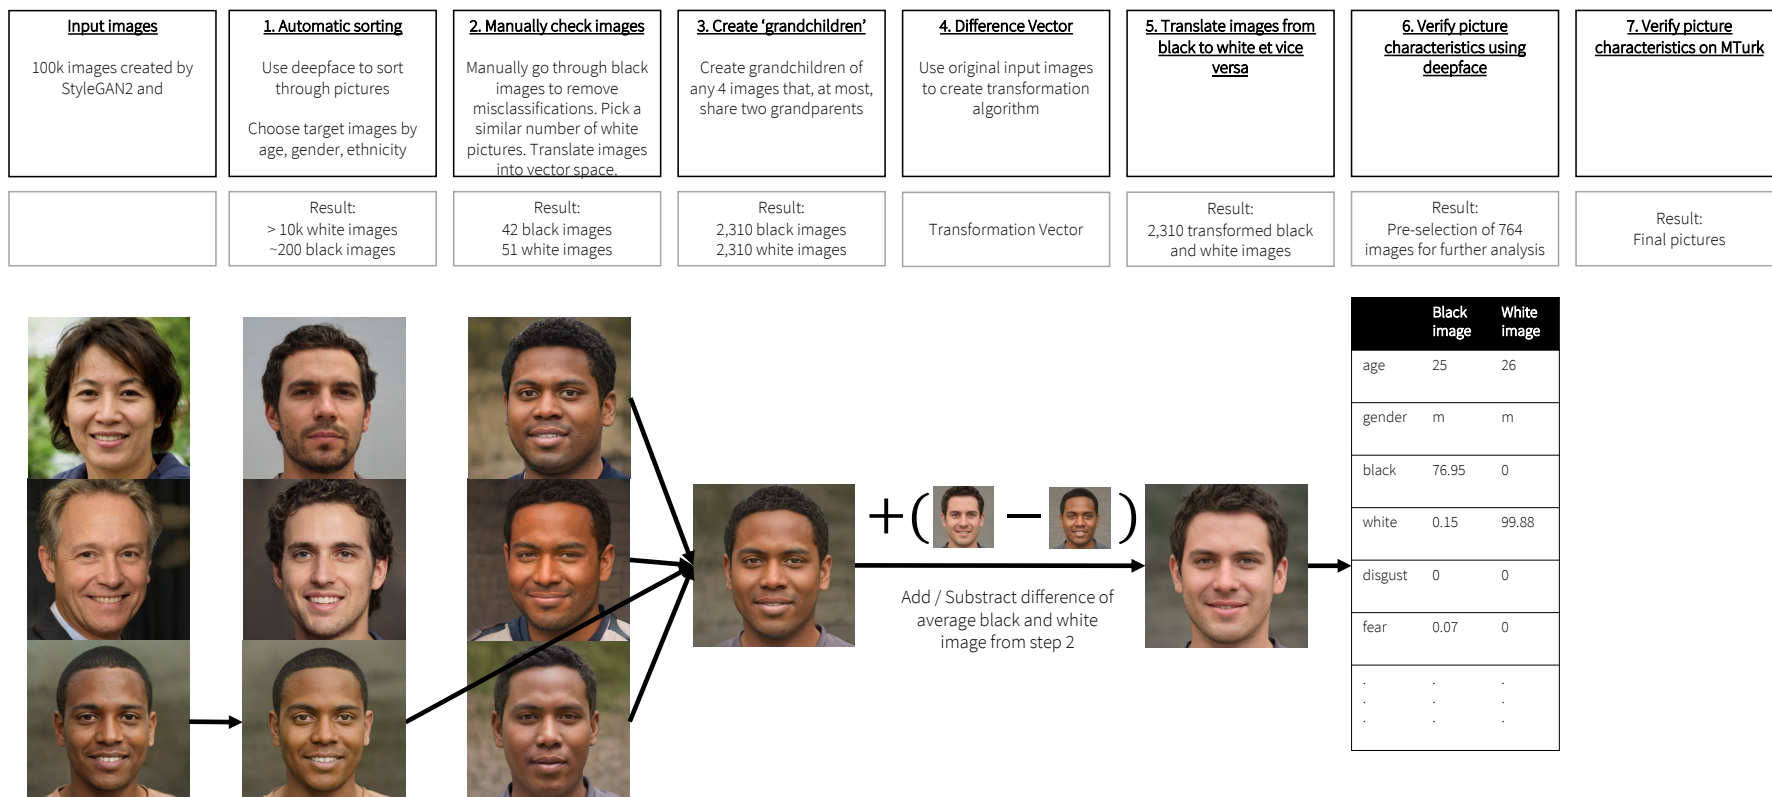


Figure B.1: Picture creation: visualization of data processing, selection, and validation of pictures

C Preparation and Structuring of Data

This section describes the preparation of data on targets, including their employers, their places of residence, their education, and demographics. To obtain information, a number of data sets are connected to targets through their publicly available CVs. These are obtained before sending targets a connection request, thus ensuring that we only draw on information targets made publicly available, i.e., to users not connected to them. Table C.2 describes the sources of data connected to targets’ CVs. Further, Table C.1 provides summary statistics on the main variables.

C1 Demographics

Age First, we estimate targets’ age using information on their level of education (as explained below) and their graduation year. We calculate age as follows: $Age = 2022 - Graduation_Year + 18 + Degree_Duration$, where degree duration is defined as 0, 2, 2, 4, 6, and 10 years respectively for the following degrees: none, some college, associate, bachelor, master, and PhD. The average target is 34 years old as shown in Table C.1.

Gender To obtain information on targets’ gender, data from the United States Social Security Administration is drawn upon. The data provide information on the gender share of each first name. It only includes men and women as potential genders. Given our balancing, around half of the targets are women.

Race A similar approach is used to estimate individuals’ race: here, U.S. census data provides information on the race share of each last name. This provides an unconditional probability of an individual with a given last name being of a certain race. Using a simple majority rule to classify individuals by race shows that 69% of targets are white, 10% Asian, 13% Hispanic, and 6% Black.

All operations regarding gender and race are done using the [predictrace](#) package in R ([Kaplan, 2022](#)). As an alternative, we analyze profile pictures using DeepFace ([Taigman et al., 2014](#)) to obtain information on race, age, and gender.

C2 Employment and Platform

Salary We estimate individuals’ salaries through their job titles. Something that is both an advantage and challenge in our context is that job titles are often very unique, e.g., instead of “Human Resources Manager”, individuals state titles, such as “Regional Human Resources Manager / Sr. HR Manager” or “National Recruitment Manager”. In total, we observe 10,509 unique titles, meaning that each title is held by an average of fewer than two targets. The distribution of mentions of job titles follows a power law distribution with the first 100 and 200 titles accounting for 32 and 36% of targets respectively. Observing many job titles has the advantage that it allows us to more precisely estimate earnings based on job titles. To obtain these, we draw on job title-specific salary estimates by [glassdoor.com](#) and [payscale.com](#). The websites draw on millions of reported salaries, providing median salaries for specific job titles. Drawing on these data has the substantial advantage that titles implicitly include information such as tenure, career advancement, and ability.

However, given the specificity of job titles, the websites do not have a specific estimate for each title. To find the closest match, we employ [google.com](#)’s search. More specifically, we restrict Google to search on [glassdoor.com](#) and [payscale.com](#). To obtain links to the annual pay within the US, we include “us annual salary” in the search term, followed by the job title. The full search term is:

“site:payscale.com OR site:glassdoor.com salary annual us [JOB TITLE]”

While doing so, we use VPNs located in the US to keep Google from reporting results for a location

outside the US. We then collect the first ten links presented on Google’s first page and their text. The first link usually includes the most precise match, e.g., for the first title listed above, it links to Glassdoor’s estimated earnings for Regional Human Resources Managers. Regarding the second title, it links to estimates for National Recruitment Managers. Overall, the estimates are highly precise.

Rather than scraping the links Google presents, we can directly draw on Google’s search results to obtain the estimate. Given that the search command includes the terms “annual”, “salary”, and “us”, Google’s snippet of the website automatically returns the median base salary estimates. Thus, we draw upon the snippet to obtain the necessary information. For example, the snippet for “Senior Vice President (SVP) & Chief Marketing Officer” reads:

“Senior Vice President (SVP) & Chief Marketing Officer . . . : 07.03.2023 — The average salary for a Senior Vice President (SVP) & Chief Marketing Officer (CMO) is \$225047. Visit PayScale to research senior vice . . .”

\$225,047 is thus the estimate we use. Most estimates we obtain stem from [glassdoor.com](https://www.glassdoor.com) (18,469 of 19,572 targets’ estimated salaries are from the site). In total, searching for 10,509 job titles yields results linking to 8,236 websites with 7,756 unique job titles, suggesting that for a number of job titles, the website provides the same links to multiple different job titles, such as linking both ‘Senior VP and CNO’ and ‘Vice President and Chief Nursing Officer’ to the same salary estimate.

Works in Human Resources To identify targets working in human resources, we create a dictionary on HR-related jobs and apply it to targets’ latest job titles. The dictionary contains the following terms: “recruit”, “recruiter”, “recruitment”, “human”, “payroll”, “talent”, “hr”, “hris”, “employment”, “employ”, “headhunter”, and “personnel”. In total, 8% of our targets work in HR-related jobs.

Senior Job Position To identify targets in senior job positions, we search targets’ latest job titles for the following terms: “ceo”, “senior”, “president”, and “director”. In total, 17.7% of targets work in senior job positions.

Employment Status We draw on the description and title of individuals’ latest jobs to identify those currently working, retired, and self-employed. *Employed* are those that do not list an end date of their current employment, that mention “today” as the end date, and that are not retired. *Self-Employed* are those whose firm- or job-title or employment type includes any of the following terms: “self-employed”, “owner”, “freelance”, or “founder”.³⁵ *Retired* are those that mention “retired” or “former” in their latest job title or “retired” as their firm. 97% of targets are currently working, 2% are self-employed, and 0.3% are retired.

LinkedIn Specific Variables A number of LinkedIn-specific variables are obtained from targets’ profiles: on average, these have 286 contacts, though this number is an underestimate as the number of reported contacts is capped at 500. Further, users can list skills and allow other users to verify these. Targets list an average of 20 skills. We observe the number of verifications of their top three skills. On average, these are verified 37 times by other platform users. Finally, 69% of profiles have a profile picture.

³⁵We also include the German translations of the respective terms (“selbstständig”, “freiberuflich”, “besitzer”, “betreiber”), as data was scraped with German browser setting. This causes the employment type to automatically be translated, though it does not affect job titles or firms .

C3 Employer

Firms can create their own profiles on LinkedIn, which they can use to advertise open positions, receive applications, advertise, increase their visibility, and for other purposes. Firm profiles include a rich set of variables. Amongst others, this includes information on their industry, the number of employees on the platform, and the number of jobs advertised on the platform. The information further includes the total number of employees in bins. These are defined as follows: 0-2, 0-10, 11-50, 51-200, 201-500, 501-1,000, 1,001-10,000, and $\geq 10,001$. We report the lower bound of each bin.

One important feature of firm sites is that users can directly link these with their current or former employment. We focus on firms targets are currently employed by and scrape information on these. Overall, 86% of targets’ current employers have a profile on the platform, a total of 7,259 unique companies. Targets work at rather large firms, with a median of 3,367 employees on the platform and 5,001 employees in total. Here, it’s important to note that the number of employees is an underestimate, given that firm size is reported in bins and the number reported corresponds to the respective bin’s lower bound. Targets thus work at rather large firms, given that our profiles were designed to work in each city’s biggest corporations, making targets also more likely to work here.

C4 Education

Degree The most recently listed degree in CVs is analyzed using a dictionary approach.³⁶ We remove punctuation from titles and move upper to lower case letters. Then the following dictionary is used to classify degrees. **Associate**: “associate”, “associates”; **Bachelor**: “bach”, “bsc”, “bachelors”, “bachelor”, “undergraduate”; **Master**: “masters”, “master”, “msc”; **PhD**: “phd”, “doctor”. In addition, individuals that we match with a college, as described below, but that do not list a degree, are assumed to have attended “some college”.

University Statistics To obtain information on the college individuals attended, we match individuals’ last attended educational institution with data on U.S. colleges. Precisely, we match university names with 2,832 degree-granting institutions in the Integrated Postsecondary Education Data System (IPEDS). We follow [Conzelmann et al. \(2023\)](#) and include the 2,832 institutions that (1) offer at least an associate’s degree and (2) were required to submit the survey every year from 2010 to 18. This suggests that they participated in any federal financial assistance program according to Title IV. Among all institutions that submitted data to IPEDS, these were responsible for 99% of undergraduate degrees according to [Conzelmann et al. \(2023\)](#). Matching is done in two steps: first, we try to directly match the university names reported by targets with those in the IPEDS list. Here we only take perfect matches. Second, we use a method we term ‘google matching’: we obtain the first 10 Google search results of each university name from both lists. Next, we reduce these to their domain and match the two lists using ‘.edu’ addresses. For the few remaining ones that we could not match with the two methods above, we use fuzzy matching if there is a close match. The majority of universities are matched using the second approach. In total, 72% of targets reporting a degree are matched to a college in this way. Table C.1 shows a few variables of this rich data, namely the share of women, Black, and White students at targets’ colleges.

College Rankings We also merge [Forbes \(2021\)](#)’ 600 ranking of top U.S. colleges to the list of targets’ universities. We do so using Fuzzy matching and then correct results and non-matches by hand. In total, around half of the targets attended a top 600 college, with a median rank of 188.

³⁶We draw on the first listed degree, which is typically the most recent and highest one.

C5 Location and County Information

Geocoding and Distance to Profile Targets’ profiles include reported locations. These are drawn upon to locate targets using Google Maps API (see Figure 1). In total, 93% of individuals are geolocated. Similarly, we ascribe our profiles coordinates using the API. Finally, we calculate the distance between our profile and any target it sends a connection request to. Targets are located close to profiles. In fact, the median target lives 14.1km (8.7 miles) from its associated profile.

CBSA- and County-Level Information Next, we use the coordinates and match these with county and CBSA (commuting area) shapefiles from the U.S. Census Bureau. We then draw on county codes to connect targets to further county-level information. First, this includes county-level vote shares in the 2020 presidential elections from [MIT Election Data and Science Lab \(2018\)](#). Second, we obtain demographics from the [COVID-19 Data Hub of the Hopkins Population Center](#). Third, we connect measures of social capital from [Chetty et al. \(2022a\)](#) and [Chetty et al. \(2022b\)](#) using both university identifiers from IPEDS and county codes. Finally, we connect average county-level race IAT scores ([Xu et al., 2022](#)).

All geographic operations are done using the SF package in R ([Pebesma, 2018](#)).

Edge-Level Information Finally, some of the information we utilize is collected on the edge level. An edge is a connection between a target and one of our profiles. Most importantly, 8% of targets attended the same university and 9% work at the same firm as the profile they are contacted by.

Variable	n	mean	sd	median	min	max
DEMOGRAPHICS						
Female (First Name)	18,756	0.52	0.50	1	0	1
Black (Last Name)	17,191	0.05	0.22	0	0	1
White (Last Name)	17,191	0.72	0.45	1	0	1
Asian (Last Name)	17,191	0.10	0.30	0	0	1
Hispanic (Last Name)	17,191	0.13	0.34	0	0	1
Other Race (Last Name)	17,191	0.93	0.26	1	0	1
Age	16,964	34.47	10.99	32	10	82
EMPLOYMENT AND PLATFORM USE						
Salary	19,450	87,278.37	56,763.05	67,196	10,625	951,257
High Job Position	19,450	0.16	0.37	0	0	1
Works in HR	19,450	0.08	0.28	0	0	1
Employed	19,450	0.97	0.18	1	0	1
Retired	19,450	0.003	0.06	0	0	1
Self-Employed	19,450	0.02	0.14	0	0	1
Number of Contacts	18,898	286.52	194.03	295	1	500
Number of Followers	19,427	861.92	28,358.97	287	0	3,700,031
Number of Skills	15,997	20.68	13.28	18	1	50
Number of Skill Verifications	15,997	37.16	94.27	15	0	9,090
Number of Posts	19,450	0.84	1.10	0	0	6
Has Volunteering Experience	19,450	0.19	0.39	0	0	1
Has Profile Picture	19,450	0.69	0.46	1	0	1
Gender Pronouns Shown	19,450	0.12	0.33	0	0	1
EMPLOYER						
Employees	16,817	4,993.01	4,551.74	5,001	0	10,001
Employees on Platform	16,761	29,550.63	75,648.07	3,367	0	962,414
Open Jobs on Platform	16,986	2,075.85	7,009.10	104	0	107,974
HIGHER EDUCATION						
None	19,450	0.20	0.40	0	0	1
Some College	19,450	0.11	0.32	0	0	1
Associate	19,450	0.04	0.19	0	0	1
Bachelor	19,450	0.40	0.49	0	0	1
Master	19,450	0.21	0.41	0	0	1
PhD	19,450	0.03	0.17	0	0	1
Undergrads: White	13,933	0.62	0.19	0.66	0.0004	1.00
Undergrads: Black	13,933	0.09	0.12	0.06	0	0.98
Undergrads: Female	13,933	0.54	0.08	0.54	0	1
Forbes Rank	9,703	231.07	164.59	188	1	599
COUNTY						
Distance to Profile (km)	18,657	385.03	886.26	14.29	0	8,068.31
Share Democrat (2020)	18,424	0.60	0.15	0.60	0.09	0.89
Share White	18,655	0.57	0.19	0.57	0.06	0.98
Share Black	18,655	0.16	0.15	0.12	0.002	0.82
Pop. Density	18,655	1,925.47	5,411.57	531.11	0.78	27,755.40
Dissimilarity Index (Black/White)	18,484	54.58	11.82	53	4	85
Dissimilarity Index (Non-White/White)	18,645	40.85	11.85	41	1	81
County: Avg. Race IAT	18,651	0.32	0.05	0.31	-0.33	0.73
EDGES						
Same University	38,299	0.08	0.28	0	0	1
Same Firm	38,299	0.09	0.29	0	0	1

Table C.1: Summary statistics

Outcomes	Description	Source
Demographics		
Age	Age of individual	Platform CV - Estimated based on degree and years of work experience & Deepface (Taigman et al., 2014)
Sex	Estimated sex of individual	First name and profile picture of individual & sex shares in first names based on data from U.S. Social Security Administration (2022) & Deepface (Taigman et al., 2014)
Race	Estimated race of individual	Last name and profile picture of individual & U.S. Census Bureau (2022) on race shares of last names & Deepface (Taigman et al., 2014)
Employment & Platform		
Salary	Salary estimate based on individual's job title	Platform CV & Glassdoor.com / Payscale.com
Works in Human Resources	Individual's job title indicates a job in HR	Most recent job title in platform CV and own dictionary
Employment Status	Employed, retired, self-employed	Most recent job title and its tenure in platform CV and own dictionary
Platform Specific Variables	e.g., # skills, contacts, skill verifications	Platform CV
Employer		
Firm's Employees	Number of firm's employees	Firm's site on the platform (lower bound of employee count which is reported in bins)
Employees on Platform and Open Positions	Number of open positions and employees of the firm on platform	Firm's site on the platform
Education		
Degree	Indicator for degree (none, some college, associate, bachelor, master, PhD)	Latest education in platform CV and own dictionary & matched degree institution from CV with IPEDS (2022) data
University Statistics	Statistics on degree-granting institution, e.g., size of university, race shares of student body, etc.	Latest education in platform CV matched with IPEDS (2022) data
University Ranking	Rank of attended university in Forbes ranking of the US' top 600 colleges	Latest education in platform CV & Forbes (2021)
County		
Distance to Profile County & CBSA	Distance between profile and reported location of individual County & CBSA in which individual lives	Reported Location in platform CV & Google Maps API Reported Location in platform CV & Google Maps API & Shapefiles on CBSA and County from U.S. Census Bureau (2013) and U.S. Census Bureau (2020)
Vote Shares County-Level Demographics	Vote shares by county in 2020 presidential election general demographics: population, population density, race shares, and dissimilarity on county-level	MIT Election Data and Science Lab (2018) from Hopkins Population Center (2020)
Social Capital Implicit Racial Attitudes	Measures of social capital on county and college level Average Race IAT Score by County	Social Capital Atlas based on Chetty et al. (2022a) and Chetty et al. (2022b) Project Implicit (Xu et al., 2022); County-level estimates by Liz Redford

Table C.2: Data sources

D Demographics & Salaries: Comparative Analysis

In this subsection, we briefly compare our targets and their characteristics to two data sets: a survey of US LinkedIn users by [Brooke Auxier \(2021\)](#) and data from the US Census. Thereafter, we present our salary estimates for different demographics and groups of LinkedIn users and compare these to data on personal incomes based on the [U.S. Census Bureau \(2021\)](#). Starting with demographics, Table D.1 shows these across the three sources.

The estimated age of the average user in our data is 32 and, thus, lower than that of the general population, but in line with LinkedIn users. This is likely driven by the fact that adoption rates among those above the age of 65 are comparatively low ([Brooke Auxier, 2021](#)).

Regarding gender, our data consists of about as many women as men, which is explained by our balancing. LinkedIn users, on the other hand, are more likely to be male.

Moving to race, compared to LinkedIn, our data slightly overrepresents the White population and underrepresents the Black one, while the data regarding Hispanics and other groups is consistent with LinkedIn’s demographic. These differences are likely driven by the fact that we create an equal number of profiles in each state, many of which are less racially diverse than, e.g., the average LinkedIn user’s hometown. As the comparison to U.S. Census data shows, targets consist of relatively many White Americans, as is expected given the comparatively high LinkedIn adoption rate in this demographic ([Brooke Auxier, 2021](#)).

Regarding education, targets have, on average, obtained a higher education than the average population. This is in line with the education of an average LinkedIn user.

Finally, we compare the average employer of targets to the average employer across the American workforce. Targets work at rather large firms: in 2022, only around 42% of the population worked at firms with a size of 1,000 or more ([U.S. Bureau of Labor Statistics, 2022](#)). In comparison, targets work at firms with a median of 3,367 employees on the platform and 5,001 employees in total.³⁷ This is likely driven by the fact that our profiles work in the biggest corporations in each city, meaning that suggestions are also more likely to work at these.

Moving to salary estimates, Table D.2 provides summary statistics of wages across different groups of targets. We obtain salary estimates for almost all targets (19,572 out of 19,619). The median salary of targets is \$67k with a higher average of \$87k. Starting with job titles, those whose titles include the terms “CEO”, “President”, “Director”, or “Manager” have above-average salaries, while assistants have below-average ones. Further, salaries increase by education, showing that those with a Ph.D. earn the most, followed by those with a Master’s and Bachelor’s degree. Further, those who went to higher-ranked colleges have higher wages.

Interestingly, LinkedIn variables are very good predictors of higher wages as well. Targets with more skill verifications by other users, more listed skills, and a higher number of contacts earn substantially more.

Moving to demographics, wages increase with age. Further, men make substantially more than women. Similarly, White users earn more than Black individuals, with Asian individuals having the highest wages.

Finally, Figure D.1 compares the income distribution in our data with the personal income distribution according to the Current Population Census 2021 ([U.S. Census Bureau, 2021](#)). As visible, users in our sample earn substantially more than the average individual in the US population. This is strongly driven by the fact that we find very few targets with estimated earnings below \$35k. Overall, the average wage of an individual in our sample lies at \$87k, while the average earnings of

³⁷It’s important to note that the number of employees is an underestimate, given that firm size is reported in bins and the number reported corresponds to the respective bin’s lower bound.

Category	Measure	This Study	LinkedIn USA	US Census
Age	Median	32	30-49	38.8
	Share 18-29	40.5%	23.9%	15.7%
	Share 30-49	47.6%	34.8%	19.0%
	Share 50-64	10.4%	31.9%	19.0%
Gender	65+	1.4%	9.3%	16.7%
	Female	52.4%	46.1%	50.5%
	Male	47.6%	53.9%	49.5%
Race	White	71.7%	63.8%	57.8%
	Black	5%	12.0%	12.1%
	Hispanic	13.2%	13.7%	18.9%
	Other	8.9%	10.5%	11.2%
Education	Highschool or Less	20.1%	11.5%	36.8%
	Some College	11.5%	13.0%	14.9%
	College +	68.4%	75.5%	48.4%

Table D.1: Comparison of demographics between LinkedIn users in our study, in the USA, and the general US population

Note: The data on LinkedIn users stems from [Brooke Auxier \(2021\)](#). The survey was only conducted on adults above the age of 18. Further, the survey only includes information on, e.g., the share of 18-29-year-olds who use the platform. To obtain a rough estimate of the share of LinkedIn users in this age range, this share is multiplied by the number of individuals in the age range according to the US Census. This is done for the other three age categories as well to obtain the total number of LinkedIn users. Finally, the number in each age range is divided by the estimated total number of LinkedIn users. We proceed in the same way for race groups and education. The following assumptions are made when estimating the demographics of LinkedIn users: (1) there are no LinkedIn users below the age of 18. (2) as the survey does not collect data on races other than Hispanic, Black, and White, we assume that the share of users of those of ‘Other Races’ using the platform is equal to the average of the above three groups.

individuals in the CPS lie at \$64k when only considering those earning at least \$20k ([US Current Population Survey \(2021\)](#)).

Group	Mean	Median	SD	N
All	87,307	67,243	56,821	19,318
SALARIES BY CAREER LEVEL				
CEO	153,399	181,804	55,976	149
President	176,554	172,410	70,150	961
Director	130,248	127,992	62,847	1,835
Senior	112,583	94,723	58,783	1,258
Assistant	56,763	41,926	40,579	2,136
SALARIES BY EDUCATION				
Degree: None	70,683	56,400	45,372	3,893
Degree: Some College	78,989	60,045	53,194	2,210
Degree: Associate	64,943	52,920	39,278	734
Degree: Bachelor	86,680	67,681	56,387	7,738
Degree: Master	108,291	91,581	62,515	4,139
Degree: PhD	116,314	117,646	61,402	604
Forbes: Top 100	113,532	96,735	67,687	2,316
Forbes: Top 200	101,794	80,964	62,921	5,272
Forbes: Ranked	96,277	76,554	61,042	9,641
Forbes: Not Ranked	83,545	64,956	53,420	5,784
SALARIES BY LINKEDIN VARIABLES				
Num. Skills Verified: >Median	105,434	87,033	61,134	8,004
Num. Skills Verified: <Median	72,918	59,392	47,693	7,884
Num. Skills: >Median	97,565	79,232	58,678	8,147
Num. Skills: <Median	80,598	61,919	54,334	7,741
Num. Contacts: >Median	104,918	85,165	63,150	9,425
Num. Contacts: <Median	70,908	58,408	43,915	9,345
SALARIES BY DEMOGRAPHICS				
Age: <30	72,177	59,014	45,672	6,787
Age: 30-39	90,942	72,386	56,736	5,681
Age: 40-49	108,457	90,861	65,963	2,377
Age: >50	113,573	96,735	68,348	2,007
Female	75,596	59,392	50,071	9,754
Male	100,191	79,813	61,219	8,873
Black	79,298	63,769	49,191	851
White	88,624	67,683	58,576	12,277
Asian	96,065	78,479	61,495	1,682
Hispanic	76,273	59,422	48,224	2,252

Table D.2: Salary statistics based on job titles and glassdoor.com / payscale.com

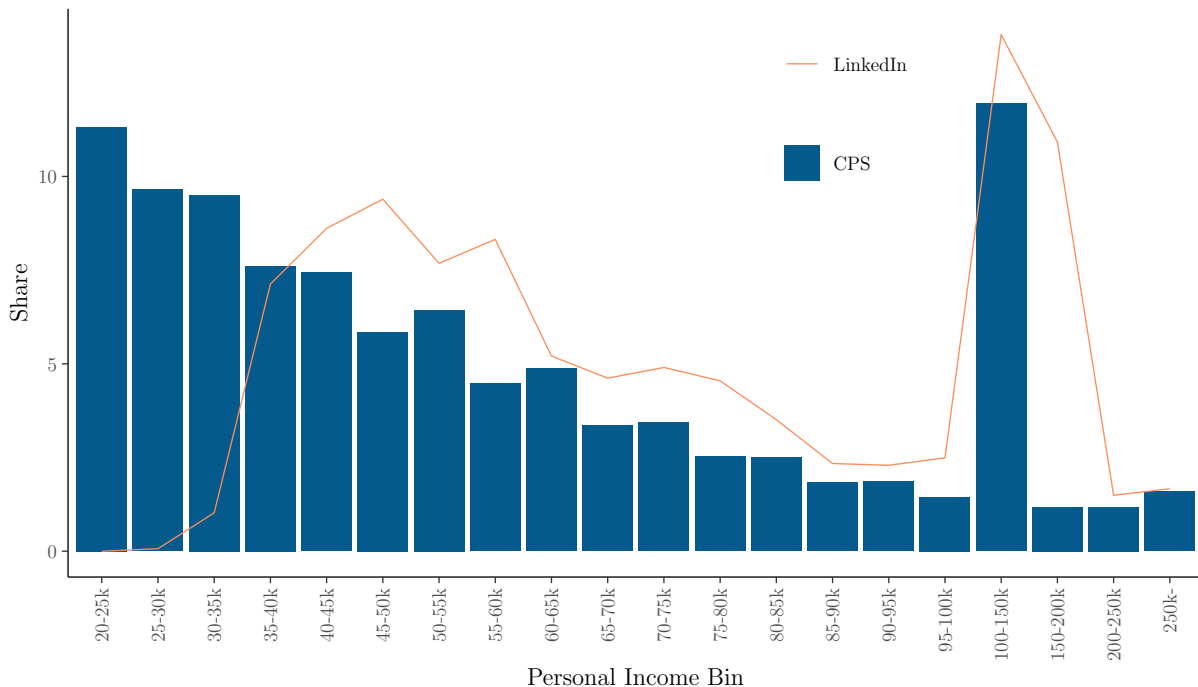


Figure D.1: Income distribution: LinkedIn sample vs. census

Note: Comparison of Personal Income Distribution in CPS and estimated salaries of targets. Source of Personal Income: [US Current Population Survey \(2021\)](#). Participants (aged 15 and above) were asked to report their personal income. To exclude part-time workers, the distribution displayed here only displays individuals with an income of at least \$20k.

E Ethical Considerations

Multiple ethical considerations have to be made in our experiment. In the main part of the paper (Section 3.5), we have briefly mentioned and argued why we believe that the benefits of our experiment outweigh the costs. Here we will address each of the ethical questions in more detail.

Our experiment has multiple avenues through which participants and non-participants might incur costs. We will first discuss the costs to the platform and participants before addressing the issues not directly affecting participants.

E1 Costs to LinkedIn and Participants

We need first to differentiate between the potential costs to LinkedIn and then the potential costs to targets.

Costs to LinkedIn In the process of creating profiles, we might impose some costs on the platform provider as we add bots to the sample of users. However, we believe these costs to be negligible given the vast number of (active and non-active) profiles on this platform: in total, we create 408 profiles on a platform with almost 200 million users in the US alone.³⁸ Moreover, fake profiles are a feature of most social media (e.g., [Silva and Proksch, 2021](#)). While LinkedIn is likely to have a much lower share of fake accounts than, say, Twitter, there exist professional sites selling fake contacts on the platform. For example, [linked500.com](#) sells 500 contacts for \$27.99 as of April 2023. Thus, the creation of our fake profiles does not considerably change the number of

³⁸see [LinkedIn’s Statistics Page \(2023\)](#)

users, and it does not burden the server capacity in a relevant way. Further, it seems unlikely that our experiment will substantially shift the users' prior to believing that the platform has too many bots.

A credible concern LinkedIn might have is that we would reveal how to create fake profiles on that website successfully. To alleviate that concern, we describe the exact creation of profiles abstractly without revealing in detail how to circumvent all the barriers and without explaining what strategies the company seems to employ to detect fake profiles.

Finally, social media and job networking platforms have become vital elements of the public sphere, including spaces for public debate and job networking (e.g., [Utz, 2016](#); [Wheeler et al., 2022](#)). Nevertheless, most platforms provide civil society and researchers with little access to data. Regarding job networking platforms, we are, in fact, only aware of one published study, which was initially run internally and later published ([Rajkumar et al., 2022](#)). We thus follow the arguments of other researchers³⁹ and, increasingly, lawmakers⁴⁰, that platforms should enable researchers to conduct independent studies on the respective platforms to justify our experiment further.

Costs to Participants As is inherent to a field experiment, the participants in our experiment are not volunteers who are aware that they are taking part in the study but are subjects who did not consent to take part in the study. Thus, they deserve special consideration and protection. These participants might involuntarily bear some costs.

The first potential cost is time spent on deciding whether to accept our profile's connection request or not. However, the cost of this decision is very minor as it does not take users long to decide whether to accept a connection request or not.⁴¹ Further, being contacted and making decisions upon connection requests is inherent to the platform, and therefore, participants at least consent to receive connection requests. Moreover, connecting with our profiles might, in fact, be beneficial for targets, as they at least increase their network. In the results, we will see that the number of connection requests is correlated with multiple advantageous outcomes (for example, the probability of receiving a message response). Thus, the mere connection decision has a tiny cost but might even have benefits associated with it, which is why we believe this intervention to be innocuous.

The more severe intervention is asking the new contacts for advice. This request indeed might have some costs as targets have to read our request and potentially draft an answer. To reduce these costs, we design our message as relatively short. However, this stage of the experiment might indeed pose non-negligible costs to participants. Yet, it might be helpful to compare these costs to costs associated with typical correspondence studies. In a typical correspondence study, the participants are HR professionals at firms, and researchers apply for jobs posted at the firm. The costs of participants in these typical studies are substantially higher than in our study. These professionals have to read the CV carefully and potentially respond to the application. They also do not have the option of simply ignoring the request. Thus, while the costs to our participants are likely non-negligible, they are substantially lower than the costs incurred in typical correspondence studies ([Bertrand and Duflo, 2017](#); [Quillian et al., 2019](#)).

Another concern participants might have is privacy. Specifically, participants might not consent to link their personal data to their connection decision and to make this data publicly available. We minimize the risks to participants' privacy. First, we access only data that is accessible to all platform users. More specifically, we scrape data before sending a request, thus not seeing information individuals make accessible only to their contacts. Thus, all the data we obtain is

³⁹see, e.g., [Jeff Hemsley's comment in the Columbia Journalism Review, 2019](#).

⁴⁰see [Center for Democracy and Technology, 2023](#).

⁴¹see Appendix I for details on how long a decision takes

data participants voluntarily made public. Second, we will make the data public as soon as the manuscript is accepted. However, we will do so after careful consideration of included variables to ensure that subjects cannot be identified. Thus, we will omit all variables that could identify a specific person, and we will reduce the set of target-specific characteristics to ensure sufficient scope for uncertainty.

E2 Further Ethical Considerations

In this section, we want to discuss multiple further ethical issues arising from our experiment.

Costs to Non-Participants Many correspondence studies pose, beyond the costs to the firms, also costs to non-participants. Specifically, in classical correspondence studies, other applications might be sorted out due to the (better) fake CVs. Specifically, if recruitment professionals aim at a specific target of how many people to invite for interviews, real applicants might be crowded out by fake applicants, thus potentially imposing non-negligible costs on non-participating subjects. In our setting, costs to not-contacted users seem highly unlikely. This concern could be valid if the number of contacts was restricted. However, no such restriction is present, and in fact, many users try to increase their number of contacts, making it unlikely that accepting our profiles will reduce the chance of accepting real profiles.

Deception Deception is inherent to most correspondence studies and many field experiments (Bertrand and Duflo, 2017). Nevertheless, the issue of deception needs to be addressed. A typical concern, in particular among experimental economists, is that the subject pool might start to be suspicious and not respond honestly to the questions asked, consequently posing a threat to the internal validity of future studies. However, this concern mostly applies to subject pools repeatedly used for experiments. In our setting, however, targets are typically not used for standard economics experiments, and thus they are unlikely to pose a threat to the internal validity of future studies. Another argument against the concern of deception is that fake profiles are a feature of most social media (e.g., Silva and Proksch, 2021), and therefore, participants could potentially anticipate being deceived on the platform. Hence, on the one hand, deception is expected, and on the other hand, deception is unlikely to spill over into future studies. Therefore, we consider the issue of deception to be minor in our setting. Finally, it is worth noting that, in the context of correspondence studies both previous research and lawmakers have acknowledged the need for deception, as informing participants would invalidate the results (Zschirnt, 2019).

Debriefing An important point to discuss is the debriefing of participants. Debriefing is rather common in psychology, in particular, if deception of the participant is involved. However, debriefing after field experiments is rather uncommon. Even though we did send a kind thank-you message to those who answered our message in the second stage, we decided not to debrief participants. There are two main reasons for that decision. The first is a mere technical one, as most website users only accept messages from contacts. Given that not all users accept our requests, we would not have been able to contact all. The more important reason is that we believe that debriefing would induce considerable costs to both the participants and the platform and would clearly outweigh the potential benefits of debriefing. First, debriefing participants would make it very salient that bots are created and used on this platform. While this is implicitly assumed on a social media platform, it is different if participants are actively made aware of this issue. Thus, debriefing might have a negative impact on the platform. The other reason is the costs to participants. The one avenue of costs is the mere reading of such a debriefing, which costs time. The other is more implicit. Information about having participated in an experiment on discrimination may impose psychological costs on users, e.g., if they believe in having behaved discriminatively. Another

problem arising from debriefing could be that participants lose their trust in users and might be less likely to respond to messages in the future, thus posing further costs for users and the platform. Thus, both targets and the platform would face considerable costs of debriefing, while the benefits of debriefing in a field experiment setting are less clear.

Change of Ethnicity A final and important ethical aspect of our study is the use of pictures and, in particular, our race transformation algorithm. We have carefully considered its use, especially given recent controversies around apps like *FaceApp*, which offered filters that allowed users to change their ethnicity.⁴² Our algorithm differs in a number of important aspects: first, none of the pictures we use are of real human beings. Thus, we do not ‘dress anyone up’ in another race. Rather, all pictures are computer-generated and are essentially vectors translated into images. Second, we swap pictures in both directions. Third, our algorithm is agnostic in the sense that we do not make any choices as to what constitutes the features of Black or White individuals (see Section 3.2). Lastly, we do not use the algorithm for entertainment purposes but merely for scientific reasons and, more specifically, to study discrimination in a setting that, arguably, requires the use of profile pictures. Thus, we believe that using the race transformation algorithm is necessary and justifiable in our setting.

E3 Benefits of our Approach

After having discussed, in detail, the costs of our experiment and how we tried to elevate ethical concerns, we need to argue that our setting is necessary and adds value to a better understanding of discrimination and that the research question warrants the costs imposed upon the platform and the participants.

Social Value of the Research Labor networks play a very important role in labor markets and those with good networks have been shown to strongly benefit from these connections (e.g., [Dustmann et al., 2016](#)). Moreover, underrepresented groups are often in worse networks ([Fernandez and Fernandez-Mateo, 2006](#)). However, in comparison to hundreds of correspondence studies on discrimination in the formal labor market ([Bertrand and Duflo, 2017](#); [Quillian et al., 2019](#)), there are no causal studies on the role of discrimination in the formation and information provision of job networks. Our study helps to fill this research gap by providing direct evidence on whether access to job networks and the benefits obtained through these are driven by discrimination. In addition, we provide direct evidence on the characteristics and geography of discrimination, i.e., answering the questions of who is more likely to discriminate and where discrimination is more likely. The results thus provide evidence that may directly support policymakers in targeting anti-discrimination policies and inform the public debate regarding the issue.

Necessity of the Employed Setting While other methods, such as the use of observational data, would impose lower costs on participants, we argue that such methods are not viable for studying discrimination in our context. More specifically, previous research noted that the use of existing data, such as representative samples, does not allow for a causal study of the effect of discrimination on job networks (see discussion in Chapter 1. and [Fernandez and Fernandez-Mateo \(2006\)](#)). Further, designing a laboratory study with externally valid results and without experimenter demand bias or other biases is hard to imagine.

⁴²See for example [Hern \(2017\)](#)

E4 Ethics: A Brief Summary and Conclusion

We conclude that – as with almost any field experiment – our experiment does create some costs to participants. However, these costs are very low, ranging from answering a connection request to voluntarily writing a couple of sentences in response to our message. Compared to more classical correspondence studies, which require the thorough study and evaluation of applications and may impose costs on third parties, the costs associated with participating in our study are very low. At the same time, this is, to our knowledge, the first study to provide causal evidence on discrimination in the formation of job networks. Given that around half of all jobs are found through informal networks, studying discrimination in their formation is important to better understand differences in unemployment rates and wages between Black and White Americans and, more generally, underrepresented groups and the mostly White population. We thus conclude that the benefits of obtaining causal evidence on discrimination through a field experiment strongly outweigh the very low costs imposed on participants.

F Validation Experiment

To validate our pictures and the universities we conducted an experiment in April 2022. The goal of this validation experiment was first, to validate that our pictures are not easily recognizable as fake, second, to validate that pictures of Black and White profiles are recognized as such (i.e., opposed to other races), third, to ensure that there are no major differences between pictures of Black and White profiles, and lastly to validate that people recognize better-ranked universities as such. To achieve our goal we conducted a three-stage experiment.

F1 Design of the Validation Experiment

The validation experiment consists of three stages.

First Stage The first stage is designed to validate that our pictures are not easily recognizable as fake. Specifically, participants are presented with a Captcha-like screen where they are asked to select all images created by a computer program. The screen contains 20 pictures.

As we anticipated that some people might guess and randomly pick pictures we require a baseline to compare the indicated number of computer-generated pictures. We choose two baselines. First, we present participants with obviously computer-generated pictures. Specifically, we choose four pictures that had either weird artifacts or contained unusual features to make it obvious that these pictures are computer generated. The second baseline contains real pictures. Here we choose six real pictures of men of the same demographic as our pictures. Following [Nightingale and Farid \(2022\)](#), we choose these from the pictures used to train the StyleGAN2 algorithm ([Karras et al., 2020](#)).

The remaining 10 pictures are our own AI-generated pictures. To ensure that all of our pictures are indeed validated we randomize, on the participant level, which of our pictures are presented. A sample screenshot of the task is shown in Figure F.1.

To incentivize this task we pay 20 cents to participants if they are able to select all computer-generated pictures. We, on purpose, choose a relatively low pay for this task to make participants less suspicious of the task and to roughly reflect the decision-making process on job-networking websites.

In this simple task you are asked to select all the pictures which are computer generated (i.e. created by an artificial intelligence (AI)).

Please select any pictures you believe are created by a computer program.

If you correctly choose all pictures created by a computer program you will receive a bonus payment of 20 cents.




Figure F.1: Screenshot of the captcha task

The figure shows a screenshot of the Captcha task. Four pictures are obviously fake, six pictures are real, and ten pictures are our AI-generated pictures.

Second Stage The second stage is designed to validate that 1) pictures of Black and White profiles are recognized as such (i.e., opposed to another race), and 2) there are no major differences between pictures of Black and White profiles.

To achieve this goal, we have to resolve two challenges: first, we need to validate more than 700 pictures, and second we aim to ensure that participants pay attention and that the data is useful.

To resolve the first issue, every participant is shown ten random pictures out of our pictures (the same pictures as in the Captcha task). To resolve the second issue, we asked participants to validate one obviously fake picture, which also clearly wears a hat. As none of our profiles wears a hat, we are able to capture participants not paying sufficient attention, or making random decisions through that question.

Participants are asked to rate all ten of our pictures plus the one obviously fake picture with respect to ten characteristics. Specifically, we ask them to estimate the age, and to rate how likely the person is to be a woman, Asian, African American, White, trustworthy, intelligent, authentic, good-looking, and to wear a hat (all on a scale from 0-100). A screenshot of this individual rating task for the obviously fake picture is shown in Figure F.2.

As we ask primarily for the perception of participants, and there is no objective answer to most of the questions and we do not incentivize the question. However, if participants do indicate that our picture has a hat, we take it as a sign of lacking attention or random decision-making. In the main analysis of the validation experiment, we, thus, exclude all participants who either indicate that one of our profiles has a hat (i.e., rated the probability of the picture having a hat as more than 50%) or indicate that the obviously fake pictures do not have a hat (less than 50%). However, all the key insights of the validation experiment remain (even though with substantially more noise) if we do not exclude these participants.

Picture #6 out of 11
 In this task you are asked to judge the pictures below with regard to the following questions:



14 (years old)	How old is the person in this picture?	99 (years old)
0 (not at all likely)	How likely is the person a female?	100 (very likely)
0 (not at all likely)	How likely is the person in this picture Asian?	100 (very likely)
0 (not at all likely)	How likely is the person in this picture African American?	100 (very likely)
0 (not at all likely)	How likely is the person in this picture white?	100 (very likely)
0 (not trustworthy at all)	How trustworthy do you think is the person in this picture?	100 (very trustworthy)
0 (not intelligent at all)	How intelligent do you think is the person in this picture?	100 (very intelligent)
0 (not authentic at all)	How authentic do you think is the person in this picture?	100 (very authentic)
0 (not good looking at all)	How good looking do you think is the person in this picture?	100 (very good looking)
0 (I see no hat at all)	How clearly can you see a hat in the picture?	100 (I clearly see a hat)

Figure F.2: Screenshot of the individual rating task

The figure shows a screenshot of the individual rating task for the obviously fake picture.

Third Stage The third stage aimed at validating that people can differentiate between better- and worse-ranked universities. For that purpose, participants are asked to indicate which university within a given state is better ranked. For every state participants are confronted with two selected options. For each correct guess, participants receive one cent. A screenshot of the task is shown in Figure F.3.

In this simple task you are asked to select the better ranked university

Below you see two universities for some of the US states. Please select the university you think is better ranked for each state out of the two options shown.
For each correct choice, you will receive 1 cent.

State	Options	
Alabama	<input checked="" type="radio"/> University of North Alabama	<input type="radio"/> The University of Alabama
Washington	<input type="radio"/> Peninsula College	<input checked="" type="radio"/> Washington State University
Arizona	<input checked="" type="radio"/> Arizona State University	<input type="radio"/> University of Phoenix - Arizona
Arkansas	<input type="radio"/> University of Central Arkansas	<input type="radio"/> University of Arkansas
California	<input type="radio"/> Dominican University of California	<input type="radio"/> University of San Diego
Colorado	<input type="radio"/> University of Denver	<input type="radio"/> University of Northern Colorado
Connecticut	<input type="radio"/> University of Connecticut	<input type="radio"/> Sacred Heart University

Figure F.3: Screenshot of the university ranking task

F2 Procedure

The validation experiment was implemented using Qualtrics. We recruited subjects online via Amazon’s Mechanical Turk (MTurk). On Mturk, registered individuals can choose to work on so-called “human intelligence tasks” (HITs) and are paid by the requester after performing the task. Most assignments are relatively simple and quick tasks like answering surveys, transcribing data, classifying images, etc. (Berinsky et al., 2012; Horton et al., 2011).

One reason for recruiting participants via MTurk is that the samples tend to be more representative of the US population than convenient student samples and consequently, social scientists established this platform as a frequent subject pool for conducting experiments (Peysakhovich et al., 2014; Rand et al., 2014; Suri and Watts, 2011). Several studies show that the data obtained on MTurk is very reliable and very similar to data typically obtained in laboratory experiments Arechar et al. (2018). The main reasons for us to conduct the experiment online was to recruit US-based workers and to receive ratings from a more representative sample.

We implemented a couple of measures and checks to ensure a high-qualitative sample. We were only interested in ratings of US workers, as we conducted the experiment in the US context. Specifically, non-US workers would likely not be able to rate universities and might also have different perceptions of race. Thus, we recruited only US-based workers, verified through IP addresses in MTurk. We further implemented basic measures such as limiting the visibility of our survey to participants who signed up at MTurk with a US address and asking to confirm participants’ US residency in the consent form. As a “gate-keeper” and to double-check the self-indicated location,

we used a third-party web service that identified participants using a tool to mask their location outside the US (i.e., VPS, VPN, or proxy).

Further, we set up eligibility criteria to ensure that participants understand the task and pay attention. As is common with MTurk-experiments, we restricted recruitment to individuals with an MTurk approval rate of 97% or higher and a history of more than 500 approved HITs.⁴³ Individuals were not allowed to take part via mobile phone or VPN clients, also as a safeguard against multiple participation. Furthermore, participants had to pass a Google-CAPTCHA to take part. Subsequently, we designed an attention check which visually resembled a typical straightforward lottery-choice task. Readers of the text were instructed to select one specific option. Selecting any other option resulted in direct exclusion from the experiment to safeguard against inattentive participants. Finally, we prevented workers from participating in our study more than once.

The experiment was publicized as an MTurk HIT with a fixed payment of \$2 and a potential bonus payment of up to 70 cents. After accepting the HIT, participants were directed to Qualtrics, where they were first asked to answer some basic demographic questions. Subsequently, participants had to pass the attention check before going through stages one, two, and three of the experiment. After finishing all the rating tasks, participants were asked whether they were able to understand the instructions and were presented with their bonus payment for the experiment.

The experiment was conducted in April 2022. 506 participants finished our experiment. However, only 307 participants were considered reliable as they indicated that none of our profiles wears a hat and indicated that the obviously fake picture does wear a hat.

F3 Results

In the analysis below we will restrict the sample to only those participants who paid sufficient attention, i.e., participants who indicated that none of our profiles wears a hat and indicated that the obviously fake picture does wear a hat. However, most of the insights reported below remain – with more noise – if we were to use the responses of all the participants who finished the survey instead.

F3.1 Captcha

The first goal of the validation experiment is to show that our pictures are not easily recognized as fake (computer generated). Figure F.4 depicts the frequencies of a figure being selected as fake. Obviously, computer-generated pictures have been selected as such rather frequently. Real and our AI-generated pictures have been selected significantly less often as fake compared to obviously computer-generated pictures. More importantly, our AI-generated pictures are not considered to be more fake than real pictures. If anything, they are considered *less* often to be fake than real pictures – however, this difference is not significantly different from zero. This finding is in line with [Nightingale and Farid \(2022\)](#), who show that well-designed AI-generated pictures are sometimes considered less fake than real pictures.

⁴³Requesters can review the work done by MTurkers and decide to approve or reject the work. Approved work is paid as indicated in the contract, and rejected work is not paid. Hence, higher approval rates of workers indicate a higher quality of work.




Figure F.4: Captcha-task: detecting fake pictures

This figure displays how often a given picture was classified as computer generated. Whiskers denote the corresponding 95% confidence intervals. T-tests are used to obtain the following significance levels: $p < 0.10$; $*p < 0.05$; $**p < 0.01$; $***p < 0.001$.

Table J.27 reports on the probability of a picture being selected as fake for the three types of pictures and interacted with multiple rater characteristics. Essentially we find that, under all specifications, our AI-generated pictures are as often selected as fake as real pictures. Obviously, computer-generated pictures, on the other hand, are substantially more often selected to be fake. In terms of heterogeneity, we see that non-White raters select our AI-generated pictures significantly less often as fake. Also, older raters are better at selecting obviously computer-generated pictures as fake, and Democrats are less likely to select obviously computer-generated pictures as fake.

One straightforward question is whether pictures of Black and White people (both real and AI) are selected as fake at different rates. Table J.28 reports upon regressions tackling this question. In essence, pictures of real White and real Black people are equally likely to be selected as fake. Our AI-generated pictures of a Black person are also not more likely to be selected as fake compared to pictures of a real Black person. Only pictures of our AI-generated pictures of a White person are slightly less likely to be selected as fake compared to pictures of a real White person. This difference, however, is not very large and not robust to controls. Hence, we take these results as evidence that our pictures are considered real, and this insight does not differ between Black and White profiles.

One possible concern a reader might have is that our Mturk sample might differ from the actual sample in the field experiment. Therefore, our insights might not hold with the sample of LinkedIn users. To deal with that issue, we can re-weight our sample based on observable demographic characteristics to resemble the sample of LinkedIn users. The corresponding regression using a weighted sample is reported in the last column of Tables J.27 and J.28. Essentially, we find no relevant difference between the two “samples”. Thus, it is likely that users of LinkedIn will consider our profiles as fake at the same rate as they would consider real profiles to be fake.

Summarizing the insight from the first part of the validation experiment, we provide evidence that our AI-generated pictures are not easily recognized as fake. Most raters consider our profiles

fake at the same rate as they would consider real profiles to be fake. Further, no subgroups seem to be systemically better equipped to correctly differentiate between our AI-generated pictures and real pictures.

F3.2 Individual Rating

In the validation experiment, we validated 764 pictures. As explained in section B, we restricted the final sample of pictures to those 408 pictures with the smallest difference between twins. In this section, we report only on the sample of those pictures we actually use in the field experiment.

The two key questions of the second part of the validation experiment are: 1) are Black and White profiles recognized as Black and White, and 2) are there major differences between Black and White profiles in terms of rated characteristics? Figure F.5 displays the rated characteristics of our AI-generated pictures, and the difference between a Black and White person being shown in the picture. The first important insight is that a Black person is considered to be very likely Black and a White person is considered to be very likely White – thus, the manipulation of our algorithm clearly works, as raters are able to correctly identify the race of the person presented on the picture. Further, our profiles are clearly considered male. In terms of demographic differences between our Black and White profiles, we find some slight divergence. Black profiles are considered to be more likely Asian than our White profiles, and to be slightly older. In terms of attributed differences between Black and White pictures, we also find some slight variation. Black pictures are, on average, considered to be slightly more trustworthy, intelligent, and authentic, while White profiles are considered slightly better looking. However, it is noteworthy that these differences are rather small and we cannot clearly disentangle whether the differences in ratings are due to tastes or driven by actual changes due to our algorithm.

Table J.29 and Table J.30 report upon the differences between White and Black profiles with regard to their demographic and trait characteristics, respectively. Both tables also investigate rate heterogeneity, account for controls, and also reweigh the sample to better resemble the population at the job-networking website. For most characteristics, we find rather little heterogeneity differences. Non-White raters and Democrats are slightly more likely to consider a Black person to be more trustworthy, and Non-White raters consider a Black person to be more intelligent – but the general patterns remain throughout the regressions. The biggest differences are found in classifying a person in a picture as White and Black. Here, non-White raters had more distinct perceptions. Specifically, non-White raters considered a Black person to be significantly more Black than a White rater would, and similarly, non-White raters considered the White person to be significantly more White than a White rater would. Raters who indicated to be Democrats were going in the other direction and had less distinct perceptions of race. Specifically, Democrats considered a Black person to be significantly less Black than a non-Democrat rater would, and similarly, Democrats considered a White person to be significantly less White than a non-Democrat rater would.

Summarizing the insight from the second part of the validation experiment: First, we provide evidence that Black and White profiles are very reliably recognized as Black and White. Second, we find some, mostly minor, differences between Black and White profiles in terms of demographic and assigned traits. Most of the differences, however, are even in favor (trustworthiness and authenticity) of Black profiles. Thus, we conclude that the manipulation of our algorithm mostly worked: it primarily changes the race of the profile without majorly changing other characteristics of the person in the picture.

F3.3 Universities

The third stage aimed at validating that people recognize better-ranked universities as such. Figure F.6 displays the propensity to correctly identify the better university as such for all 51 states. On average, participants correctly identified the better universities in 37 states. In 8 states,




Figure F.5: Average classification of profile pictures

Here, we look at how a given picture was classified. The bars indicate the responses to the questions: How old is the person in this picture (Age)? How likely is the person a woman (Gender)? How likely is the person in this picture Asian/African American/White (Asian, Black, White)? How trustworthy/intelligent/authentic/good-looking do you think is the person in the picture (Trust, Authentic, Intelligence, Looks)? Whiskers denote the corresponding 95% confidence intervals.

participants were not able to disentangle lower- and higher-ranked universities (i.e., the confidence interval of the average rating contained the 50% mark). However, there are some states where participants actually rated the lower-ranked university as better. The most striking example is Michigan, where participants rated Central Michigan University as better than Kalamazoo College, even though most rankings place Kalamazoo College higher. In total, participants systematically rated the lower-ranked university as better in 6 states (Michigan, Minnesota, New Jersey, Ohio, South Carolina, and Wisconsin). Averaging over all the states, participants considered the better-ranked university as better 67% of the time.

Table J.31 depicts multiple potential predictors of correctly identifying the better universities, as well as the results after reweighing the sample. Essentially we find that better universities are systematically recognized as better, and there is little variation based on demographic characteristics. Older raters are better at correctly identifying better universities, while Democrats are doing significantly worse. To interpret the results, it is worth noting that the high-ranked universities have a national rank of 91-331 in Forbes' ranking. Thus, we would expect that individuals from the respective states, i.e., those that are primarily treated by the respective profiles, are better able to distinguish between local universities.

Summarizing the results of the third stage: Participants were able to correctly identify the better-ranked universities in most states. On average, participants, most of the time, correctly recognized the better universities. Thus, our signal of quality, while noisy, is likely to be informative.




Figure F.6: Propensity of correctly identifying the better ranked university as such
The bars indicate the propensity to correctly select the better university in each state. The red line denotes the 50% line. Whiskers denote the corresponding 95% confidence intervals.

G Additional Data Analysis

G1 Dynamic Effects

One advantage of our design is the possibility of studying dynamic effects. In particular, we can observe whether Black profiles are able to catch up at some point, or whether White profiles are perpetually improving over time.

Figure G.1 shows the difference in the number of contacts between Black and White profiles over time. The figure also shows the bootstrapped difference in the number of contacts between Black and White profiles relative to the number of contacts of Black profiles, over time. The figure reveals that discrimination kicks in almost immediately. Already in the first week of the experiment, White profiles receive more connections than Black profiles. The absolute gap between Black and White profiles is also increasing over time. However, the relative gap stays rather constant. Hence, White profiles are not perpetually improving, but Black profiles are also not able to catch up over time. Essentially, it does not seem like having an established network is additionally beneficial for Black people. Discrimination is ubiquitous when starting off and also when being already established.⁴⁴ The figure also shows that having attended a better university does not shield one

⁴⁴Obviously, it might be the case that the non-linearity of the effect will show up at substantially higher levels of contacts. Our data cannot exclude this possibility as we cannot speak to well-established large networks. But what our results do show is that White profiles *starting off* have a clear advantage over Black profiles also starting off. The existing literature on (online) networks further suggests that, if anything, the gap should be expected to widen: the number of connections in a network follows a scale-free power-law

from this experience. The general pattern and the gap in connections are present for both types of profiles.

distribution. This is typically driven by connections being preferentially made with others that already possess many connections ([Barabási and Albert, 1999](#)). However, we cannot rule out that our setting constitutes a special case in which these insights do not apply.

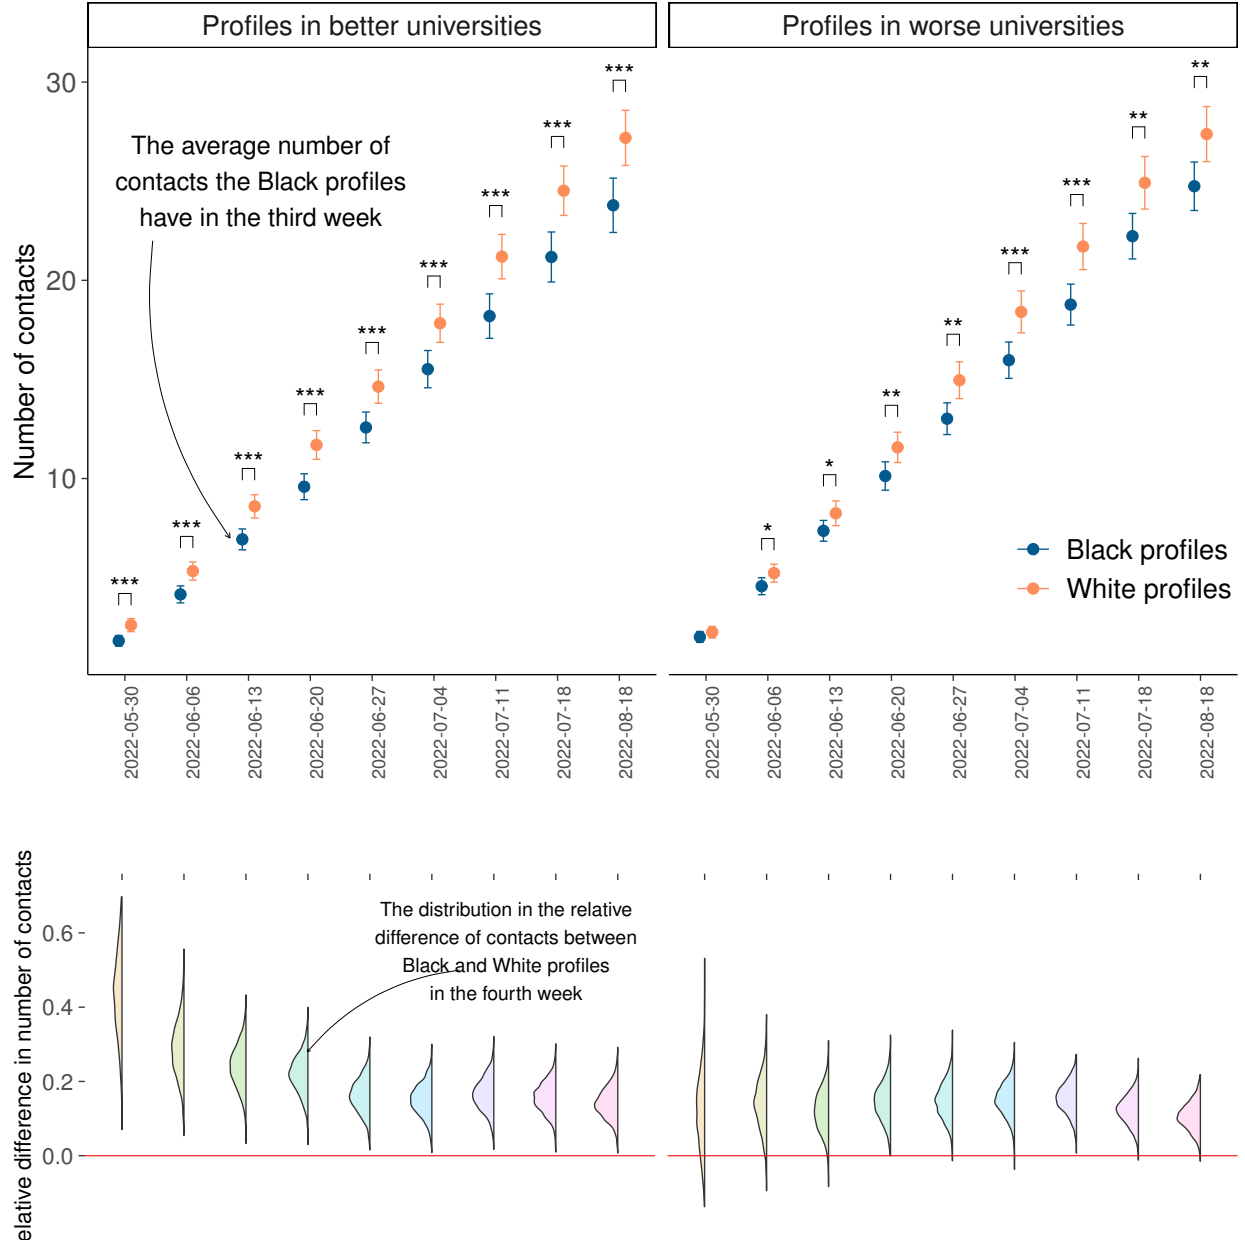


Figure G.1: The evolution of the number of contacts by race and profile quality

The figure depicts the number of contacts by the week of the experiment for Black and White profiles separately. The left panel denotes results for profiles attending worse universities, while the right panel denotes profiles indicating attendance at a better university. Orange and blue dots denote the aggregate number of contacts of White and Black profiles, respectively. Whiskers around the mean denote the corresponding 95% confidence intervals. T-tests are used to obtain the following significance levels: $p < 0.10$; $p < 0.05$; $p < 0.01$; $p < 0.001$.

The bottom panel depicts the distribution of the relative difference (i.e., the gap in the number of contacts between Black and White profiles relative to the average number of contacts Black profiles have) for every week of the experiment.

In Tables J.2 and J.3 we report the relevant regressions and further investigate the dynamic effects. The regressions support the findings reported in the figure above. The connection gap emerges directly at the beginning of the experiment and increases over time (see Table J.2). However, we do not find any evidence that the gap-increase changes over time – thus, the connection gap

grows constantly (see Table J.3). Both tables also show that the connection gap is not influenced by the quality of the profile indicated by the university attended. These results remain consistent across various control variables and model specifications.

An alternative way to view the dynamic effect is to examine how the gap in acceptance probabilities changes as the network grows (i.e., the number of connections increases). Figure G.2 shows the acceptance probability of connection requests from White and Black profiles based on their number of connections at the time of the request (top panel). Counterintuitively, the acceptance probability slightly decreases with the number of connections until around 20 connections, then increases. The acceptance gap between Black and White profiles remains relatively constant at about 4 percentage points (see middle and bottom panels), with only a slight reduction for profiles in the top 25% of connections (more than 15 connections), though this sample is relatively small. This suggests that having more connections does not disproportionately benefit Black profiles, nor does it disproportionately benefit White profiles. Black profiles are consistently less likely to be accepted, regardless of their network size (i.e., for any number of connections).

However, an important caveat of this approach is that the variable of interest, i.e., the number of connections, is endogenous. Specifically, it is driven by three factors (and their interactions): first, it is driven by the timing. Acceptance rates of observations with 20 connections are far more likely to have been collected by the end of the experiment than, say, observations with 5 connections. At the same time, acceptance rates may vary over time, as targets go on vacation or similar. Second, race directly affects the number of connections. Hence, comparing acceptance rates between a Black and a White profile with the same number of connections tends to compare a Black profile later in the experiment to a White profile earlier in the experiment. This is especially problematic if acceptance rates vary over time. Finally, race-independent profile characteristics (e.g., profile pictures, CVs, location, etc.) drive acceptance rates. For this reason, acceptance rate gaps between Black and White profiles with 30 connections are only driven by the most successful profiles, while unsuccessful ones more strongly drive acceptance rates at the lower bounds of the distribution. This may explain the apparent fall in overall acceptance rates for profiles with a medium number of connections. It also means that comparing Black and White profiles with the same number of connections may effectively compare a Black profile from a more successful twin pair to a White profile from a less successful one. To partially deal with this issue, we control for profile picture random effects when estimating the effect (see Table J.13). Here, we again find no effect of the number of connections on the acceptance rate gap. However, this result also remains endogenous to the other factors mentioned above.

Given that comparing acceptance rate gaps over time does not suffer from the same issues, this remains our preferred specification to study dynamics. Nevertheless, the results on acceptance rates conditional on the number of connections provide interesting insights as well. The results are similar to our findings on acceptance rate gaps over time, and suggest that gaps can be described as insensitive to the number of connections a profile has.

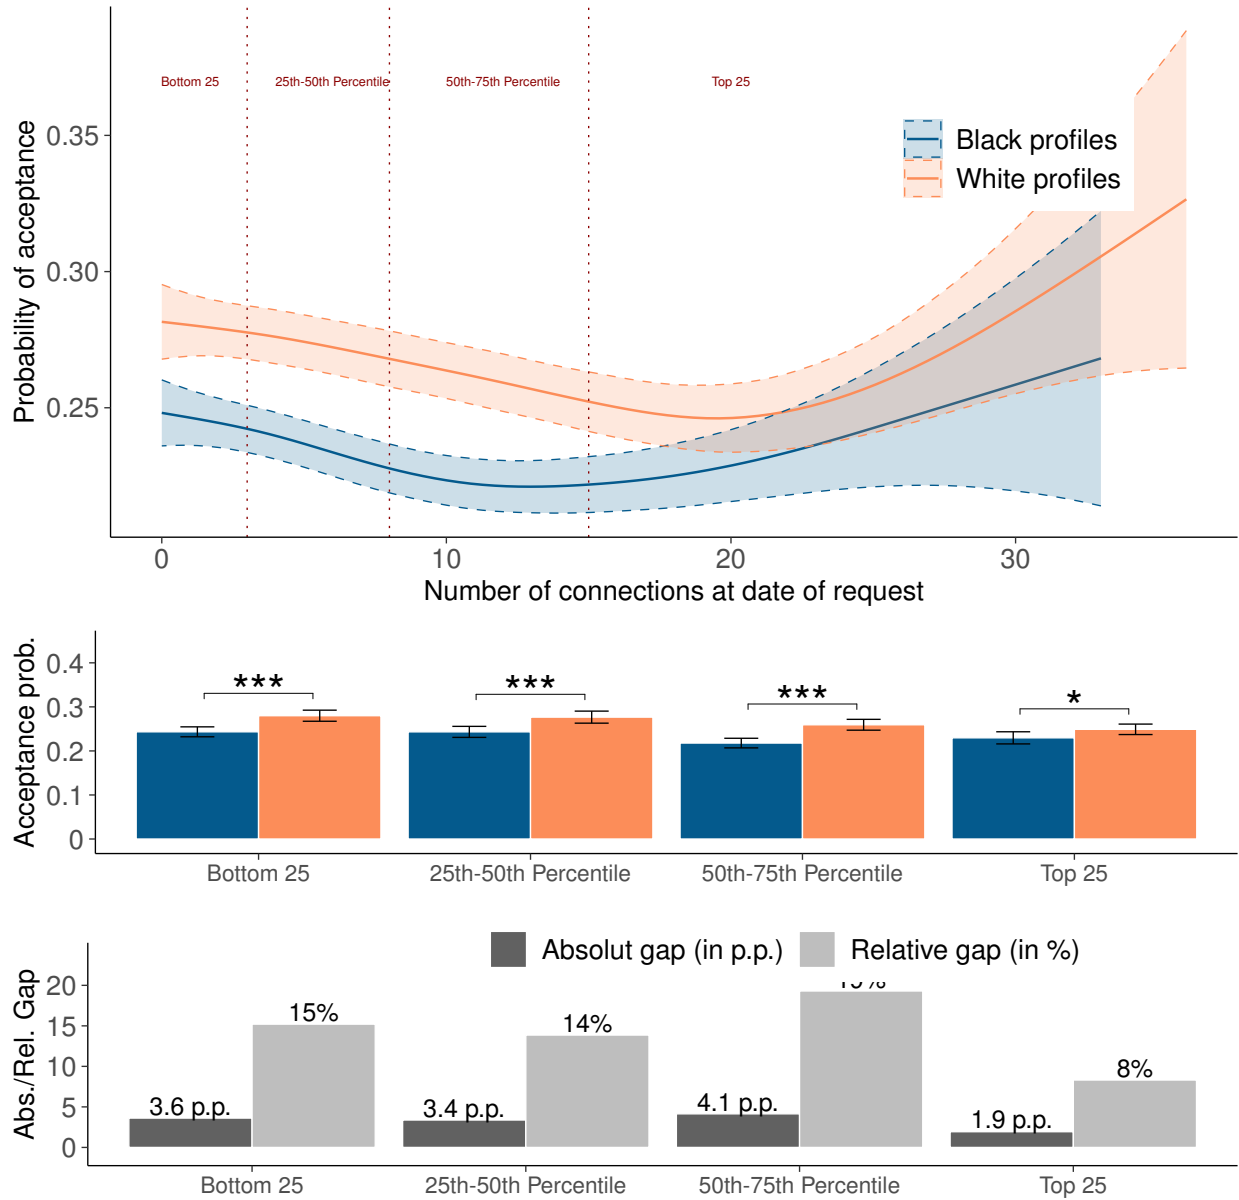


Figure G.2: Acceptance probability of a connection request by a Black and White profile as a function of the profile's number of connections at request

The figure illustrates the acceptance probability of a connection request as a function of the number of connections the profile has at the time of request. The top panel illustrates the fitted acceptance probability as a function of the profile's number of connections, including the four quartiles separated through vertical dotted lines. The middle panel depicts the acceptance probability by the profile's number of connections in quartiles. The bottom panel illustrates the relative gap (i.e., accounting for the acceptance probability of a White profile's request) in acceptance probability by the profile's number of connections in quartiles. Orange and blue objects denote the White and Black profiles, respectively. Whiskers around the mean, and bands around the spline, denote the corresponding 95% confidence intervals. T-tests are used to obtain the following significance levels: $p < 0.10$; $*p < 0.05$; $**p < 0.01$; $***p < 0.001$.

G2 Geographical Variation

As we do not only vary the timing but also the place of the experiment we can study geographical variation. However, focusing on state-level differences reduces the sample size to 8 profiles per state (4 profile pairs), making inference rather noisy and less reliable. Therefore, we discuss the geographical variation at this level in the Appendix. In the results 4.1.3, we further provide evidence of geographic variation by drawing on targets' home counties.

Figure G.3 displays the difference in the number of contacts between White and Black profiles for each state. We see that in most states (43 out of 51) White profiles have more connections than Black profiles. However, for the majority of states, the difference in the number of connections between the Black and White profiles is not significantly different from zero, as the inference builds on 4 observations per state (4 profile pairs). In Table J.4, we also study whether the state-level differences in the number of contacts between White and Black profiles correlate with some relevant state-level characteristics. However, given the very noisy measure, we find very little relevant variation. The only significant predictor (at the 5% level) of more discrimination is whether a state is part of the so-called Black Belt consisting of states in the south of the US where a large number of Black slaves have been exploited before the Civil War. Specifically, we find that the difference in the number of contacts between White and Black profiles doubles in the so-called Black Belt.

Summarizing this section, we essentially observe that, on the state-level, the disadvantage of Black profiles is rather stable across space, however, with some variation. For the most part, we do not find a clear pattern explaining this variation, which, again, is most likely driven by the small number of independent observations per state.




Figure G.3: Difference in the number of contacts by race in each state

G3 Variation in Profile Quality

In line with existing studies (e.g., [Oreopoulos, 2011](#)), we vary profile quality by assigning half of the profiles within a given city to a higher- and half to a lower-ranked university (see Section 3.3 for more details). In this section, we briefly discuss the results as a function of profile quality. Figure G.4 displays the acceptance probability as a function of whether the profile attended a better or worse university. We find the same pattern as observed in the pooled results reported in the main part of the paper: Black profiles have substantially fewer connections than White profiles under both quality conditions. We also see that the gap under both conditions is comparable. Table J.1 reveals that the gap is small and statistically not different between the two conditions.




Figure G.4: Number of contacts by the end of the experiment by race and quality of the profile

The figure depicts the number of contacts obtained by Black and White profiles individually at the end of the experiment. The left panel displays the results for profiles from lower-ranked universities, while the right panel represents profiles indicating attendance at more prestigious universities. White profiles are depicted by orange objects, and Black profiles are denoted by blue ones. Each dot on the graph represents a single profile, and twin pairs are connected by gray lines. Whiskers around the mean denote the corresponding 95% confidence intervals. T-tests are used to obtain the following significance levels: $\cdot p < 0.10$; $\ast p < 0.05$; $\ast\ast p < 0.01$; $\ast\ast\ast p < 0.001$.

Figure G.5 displays the response rate in Stage II as a function of whether the profile attended a better or worse university. We find that Black profiles are slightly more likely to receive a response than White profiles if they indicate to have attended a better university. This difference is, however, not significantly different from zero. On the other hand, White profiles are slightly more likely to receive a response than Black profiles if they indicate to have attended a worse university. This difference is barely significant at the 5% level. Tables J.18 and J.19 reveal that the difference in discrimination between profiles attending a better or worse university is only significant at the 5%

level and can only be found for the response rate and not for other message characteristics like the length and usefulness of the message. Overall, these results suggest that there is little evidence for discrimination in Stage II.




Figure G.5: Response rate by race and quality of the profile

The figure depicts the response rate by the race of the requesting profiles separately. The left panel displays the results for profiles from lower-ranked universities, while the right panel represents profiles indicating attendance at more prestigious universities. Orange objects denote White profiles, while blue objects denote Black profiles. Each dot represents one profile and twin pairs are connected through colored lines (where red vs. blue lines denote a gap in favor of White or Black profiles, respectively). Whiskers around the mean denote the corresponding 95% confidence intervals. T-tests are used to obtain the following significance levels: $p < 0.10$; * $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$.

Figure G.6 finally displays the ex-ante informational benefit of the networks of Black and White profiles as a function of whether the profile attended a better or worse university. We find the same pattern as observed in the pooled results reported in the main part of the paper: Black profiles are expected to receive substantially fewer responses than White profiles under both quality conditions. We also see that the gap under both conditions is comparable. Table J.26 reveals that the gap is very small and statistically not different between the two conditions.

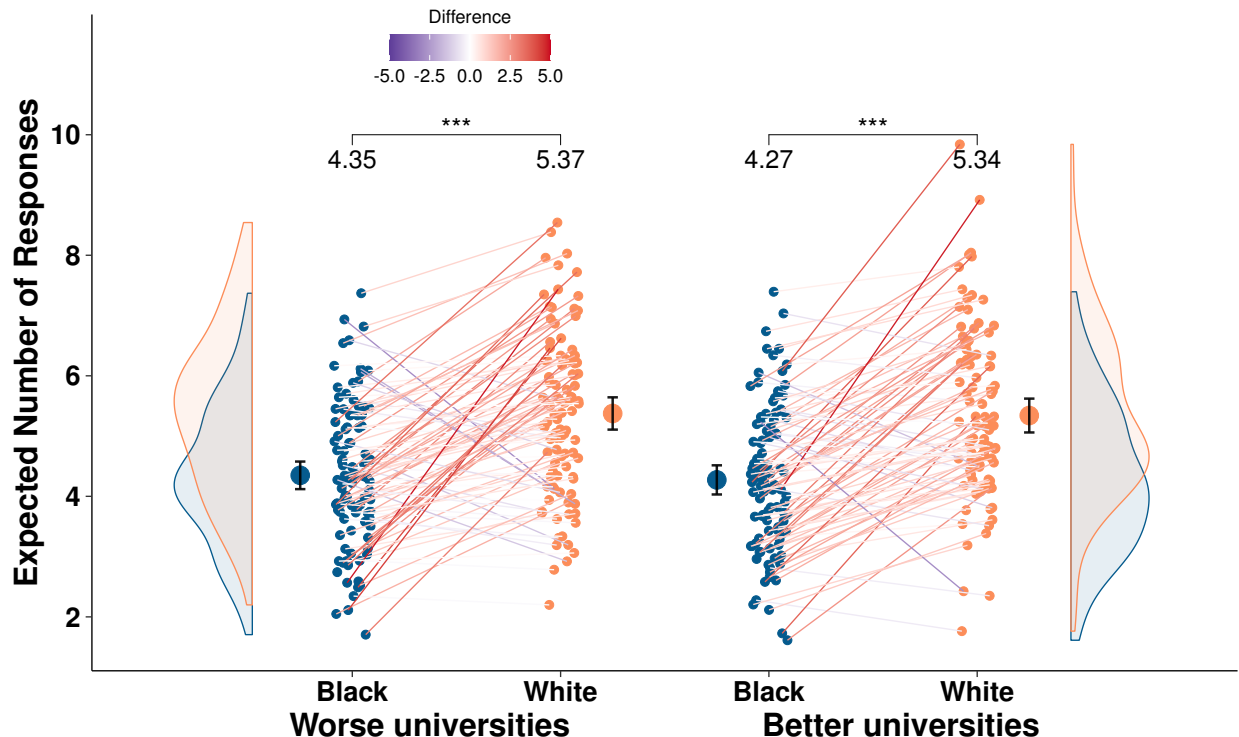


Figure G.6: Number of ex-ante expected responses when creating a network

The figure depicts the ex-ante expected responses when creating a network for White and Black profiles separately. The left panel denotes results for profiles attending worse universities, while the right panel denotes profiles indicating attendance at a better university. Orange objects denote White profiles, while blue objects denote Black profiles. Each dot represents one profile and twin pairs are connected through colored lines (where red vs. blue lines denote a gap in favor of White or Black profiles, respectively). Whiskers around the mean denote the corresponding 95% confidence intervals. T-tests are used to obtain the following significance levels: $\cdot p < 0.10$; $* p < 0.05$; $** p < 0.01$; $*** p < 0.001$.

G4 Predictors of Acceptance

In this section, we discuss predictors of a contact request. Specifically, we are interested in understanding who is more likely to accept a contact request, as it might be useful when targeting potential contacts. We also want to understand what characteristics of the targeted person as well as the characteristics of our profile predict acceptance. Figure G.7 reports upon multiple relevant characteristics and how they are associated with accepting a request of our profile.

Focusing first on the demographic characteristics of the target, as well as their educational attainment, reveals that gender, education, and age highly predict whether a target will accept the connection request. Specifically, better-educated targets (i.e., users who have either an associate degree or a bachelor's degree as their lowest degree) have an almost 5 p.p. higher probability of accepting a contact request, and comparably people without a degree have a 6 p.p. lower probability of accepting a contact request. Women also seem to have a slightly lower probability of accepting a request (about 1.5 p.p.). Further, a one standard deviation increase in age reduces the probability of accepting a request by roughly 4 p.p.

In terms of the targets' usage of LinkedIn, we find that variables indicating actual engagement on the platform are highly predictive of accepting a profile. Specifically, one standard deviation increase in the log of the number of followers, and similarly, one standard deviation increase in the number of contacts increases the probability of accepting a request by roughly 7 p.p. and 4 p.p.,

respectively. Also, users who decided to display volunteering experience are slightly more likely to accept a contact request.

Interestingly, most job characteristics have little predictive power in the probability of acceptance. Reassuringly, people who have an HR job are 5 p.p. more likely to accept contact requests. In contrast, retired people and users who have a management position (director, president) are almost 10 p.p. and 5 p.p. less likely to accept a connection request, respectively.

The area in which the person lives also has little predictive power over acceptance rates. However, multiple characteristics of our profile (and the link between our profile and the target) predict whether the target will accept. The most striking predictors are whether the target and our profile have something in common. Specifically, if both attended the same university or currently have the same employer, they have a 13 p.p. higher probability of accepting our connection request. Two other important characteristics of our profile that predict acceptance are whether our profile is White and how likely the person on the profile picture is considered Black (which is directly a function of whether our profile is White). In case our profile is White, the probability of accepting is 3 p.p. higher, and one standard deviation increase in the likelihood the person on the profile picture is considered Black reduces the acceptance rate by 2 p.p. Notably, the quality of the university our profile attended does not impact the acceptance probability.

G5 Drivers of Discrimination

In the main part of the paper, we have mentioned multiple relevant correlates of discrimination. In this section, we first discuss and analyze the main and pre-registered correlates in detail (Section G5.1). In a second step, we explore heterogeneity in discriminatory behavior across a wider range of covariates using causal machine learning methods proposed by [Wager and Athey \(2018\)](#). Finally, Section G5.3 zooms in on the strongest predictors of discrimination: age, gender, and race.

G5.1 Heterogeneity by Key Demographics

First, we restrict our attention to major and obvious characteristics. These include age, gender, job position, share of Republican votes, race, and education. The first five were explicitly pre-registered.⁴⁵ In a second step, we then explore additional heterogeneity in discriminatory behavior, applying methods proposed by [Wager and Athey \(2018\)](#).

To examine variations in how people react to connection requests, we assess whether the difference in how a target responds to Black and White profile requests is connected to the individual’s age, race, and gender. We determine age from their CV, race from their last name, and gender from their first name. Further, we draw on targets’ home counties and include a dummy for an above median Republican vote share in the 2020 presidential elections. For education, we use a dummy variable to indicate whether a target has obtained at least a bachelor’s or master’s degree. We further include two variables related to an individual’s job position: (1) whether a target’s job title suggests she is a president, director, CEO, or senior employee,⁴⁶ and (2) whether the residualized income of the target is above the sample’s median income. We residualize income by running a regression of log income on an individual’s age, gender, race, and education.⁴⁷

⁴⁵The pre-registration mentions the city/state level Republican vote share. Given that we can observe more precise data, namely people’s self-reported location, we draw on county-level data here.

⁴⁶As shown in Table D.2, these job titles are related to substantially higher incomes than the average target.

⁴⁷More specifically, the variables included are a second-order polynomial of age, indicator variables for each type of the highest degree achieved (none, associate, some college, bachelor, master, doctoral), an indicator for a likely female first name, and a race dummy variable based on the user’s last name (consisting of eight different race categories). Given that income estimates are based on average wages for specific job titles across the entire US, we do not control for regional wage levels. We then use the difference between




Figure G.7: Predictors of acceptance

The figure illustrates the β -coefficients of the following regression: $accepts_{i,j} = \alpha + \beta \cdot \text{Variable} + \epsilon_i + \epsilon_j + \epsilon_{i,j}$. ϵ_i and ϵ_j are user and profile picture random effects with $(\epsilon_i \sim \mathcal{N}(0, \sigma_1^2), \epsilon_j \sim \mathcal{N}(0, \sigma_2^2), \epsilon_{i,j} \sim \mathcal{N}(0, \sigma_3^2))$. *Variable* denotes the z-scored variable, if the original variable is not binary. The regression thus computes the acceptance rates as a function of one feature of the target/or profile, while accounting for the fact that each target made two decisions, and the fact that each profile has contacted multiple targets.

Generally, we find that all groups of users discriminate, i.e., react more favorably to a White than a Black request. However, there is substantial heterogeneity between different groups as shown in Figure G.8, which plots the coefficients of the interaction term only.⁴⁸ Several of the correlates are in line with what one might expect. In particular, we find that Black individuals are less likely

the actual log income and an individual's predicted log income and create a dummy for having an above and below median residual income. One could interpret the income residual in terms of outperforming others with similar demographics and education. However, it may still include a number of unobserved characteristics, such as personal preferences driving career choices, and the coefficient should thus be interpreted with some caution.

⁴⁸See Table 2 for results on the full gap.

to discriminate, although they still do in absolute terms, as shown in Table 2. Given the low share of Black individuals among targets, the confidence interval is rather large, though. Appendix G5.3 investigates this further, showing that it is Black women who discriminate less, while Black and White men discriminate to a similar extent against Black profiles. This result is similar to Edelman et al. (2017) who show that, on Airbnb, Black men discriminate more than Black women, though the difference is insignificant. These findings are in contrast to Block et al. (2021), who show that other than the rest of the population, Black Americans do not discriminate when asked to participate in a survey.

Our results further suggest that targets reporting to reside in more Republican counties discriminate more strongly. This is in line with Block et al. (2021), who document stronger discrimination for registered Republicans, as well as studies on IAT, observing stronger racial bias for more conservative individuals (Nosek et al., 2007).⁴⁹ Finally, we document that targets with a higher job position and income residual discriminate to a slightly lower extent, though the results are insignificant.

Perhaps surprisingly, the two strongest predictors show that women discriminate significantly *more* than men and that older targets discriminate significantly *less* than younger ones. We investigate both in more detail in Appendix G5.3. Starting with gender, we find that the effect remains, even after controlling for a host of other target characteristics. We further document that the results are driven by White women, while we find little evidence of discrimination among Black women. This suggests that dating preference might be one possible explanation for the observed pattern, given a majority heterosexual population and homophily in dating. However, our survey among LinkedIn users clearly shows that dating is the least relevant reason for using LinkedIn. More than 90% of users rarely or never received romantic advances from users on LinkedIn, and 99% of users indicate to rarely or never use LinkedIn to search for romantic partners (see Appendix I), all of which is speaking against dating preference driving the result. An alternative explanation might be stereotypes against Black men specifically held by or salient for White women (e.g., Davis, 1981; Sommerville, 1995; Zounlome et al., 2021). Both explanations are in line with our data, but future research is required to investigate the underlying reasons. Interestingly, Edelman et al. (2017) also document a higher gap in response rates of White women in comparison to White men towards Black men on Airbnb.

Moving to age, we document a substantially higher level of discrimination for young individuals. In Appendix G5.3, we show that this is particularly driven by Gen Z and Gen Y and explore some potential explanations. Regarding other studies, Edelman et al. (2017) show that young hosts do not discriminate less on Airbnb.

We find some suggestive evidence for discrimination slightly decreasing with education, as indicated by the point estimates for holding at least a master’s or bachelor’s degree. However, the results are not very strong, which suggests that education is only weakly associated with lower levels of discrimination.⁵⁰

⁴⁹Acquisti and Fong (2020) also show that discrimination in hiring against Muslims is increased in more Republican states.

⁵⁰The regressions above all test for absolute differences in the acceptance rate between Black and White profiles’ requests. Here, we essentially follow the literature (e.g., Block et al., 2021; Edelman et al., 2017). The coefficients’ interpretation becomes challenging if baseline acceptance rates differ substantially. For instance, consider that group A accepts 20% of Black and 25% of White requests, while group B accepts 50% and 60%, respectively. Based on the regressions above, we would conclude that group B has a higher acceptance rate gap (5 vs. 10%). However, one could also argue that, compared to the baseline, being White increases the acceptance rate by 25% for group A and 20% for group B. To account for this, we proceed as follows: first, we calculate the propensity to accept the *Black* request based on the user’s characteristic alone




Figure G.8: Correlates of discrimination

The figure illustrates the degree of association between specific user characteristics and discrimination, with smaller values indicating stronger associations and larger values indicating weaker associations. To estimate heterogeneity, we run the following regression with the above figure showing β_1 : $accepts_{i,j} = \beta_0 + \beta_1 Black_i \times characteristic_j + \beta_2 Black_i + \beta_3 characteristic_j + \gamma_{P(i)} + \omega_j + u_{i,j}$ where the dependent variable indicates whether target j accepted the request to connect from profile i . β_1 is the coefficient of interest, i.e., the interaction effect between a target's characteristic and whether the profile sending the request is Black. ω_i is a target-specific intercept and $\gamma_{P(i)}$ is a separate intercept for the (transformed) profile picture, i.e., a twin-specific control. The red dashed line denotes a null effect. Blue dots denote the interaction effect between race and the variable on the y-axis (e.g., "Age>Median" indicates that users above the median age of users are less likely to discriminate against a Black profile). Whiskers denote the corresponding 95% confidence intervals.

G5.2 Causal Machine Learning to Obtain Heterogeneous Treatment Effects

We use causal forests to estimate heterogeneity in treatment effects based on 18 variables. To implement causal forests, we employ the *grf* package in *R* (Tibshirani et al., 2023). While we do observe substantially more variables, we restrict our selection for two reasons: first, causal forests perform worse if too many covariates are included (Chernozhukov et al., 2018; Wager and Athey, 2018). This is particularly the case if variables are strongly correlated. For instance, including the share of the Black and White population in a county separately may make it harder to distinguish the effects. Second and most importantly, many variables are not available for the entire set of observations. In total, 77.45% of observations have full data and are thus included in this section.⁵¹ Regarding the outcome of interest, the full data does not substantially differ from the reduced data set, with a raw difference in the acceptance rate between Black and White profiles of 3.2 p.p. in the full and 3.4 p.p. in the reduced data set.⁵²

To choose variables to focus on, evaluate model fit, and the presence of heterogeneity, we start by splitting the data into a training and test data set, with 75% of observations in the training

and use the result to predict the probability of accepting the request for the entire data set. Next, we divide the decision to accept by the predicted probability. We then re-run all regressions using the resulting value as the dependent variable. The reported coefficient now measures the relative increase in the acceptance rate, i.e., it would suggest a gap of 25% and 20%, respectively. Overall, these are very similar to those before, with all coefficients going in the same direction.

⁵¹While causal forests can also include missings, forests treat their status as missing as informative. As this makes interpretation more difficult, we, instead, only include observations with full data.

⁵²The following individual-level variables are included in the forest: probability of being female according to first name, race is Black (by name and, where missing, picture), Bachelor+, high job position, age, number of connections, number of skills verified, number of skills, works in HR, visited same university as our profile, works at the same firm as our profile, number of posts, signals pronouns in profile. Further, the following county-level variables are included: Republican vote share (2020), share Black, Black/White dissimilarity index (segregation), average IAT Score, and economic connectedness index.

data. We then separately estimate causal forest with 50k trees, including all covariates for each data set. Given that we observe each observation twice, once treated and once untreated, this allows us to provide the forest with both a propensity score, which is 0.5 across all observations, and an estimate of the dependent variable. The latter is obtained by simply taking the average contact request acceptance rate for each user.

While our results on pre-registered variables show substantial heterogeneity, we run further tests for heterogeneity and model fit to validate the forest. The first is informative but rather qualitative, as pointed out by [Athey and Wager \(2019\)](#). Here, we first predict CATE in the test data using the training forest. We then compare average treatment effects according to the test data forest between individuals predicted to have above vs. below median CATE. We show that the ATE in both groups differ significantly (difference: 0.027; $p < 0.01$). In a similar fashion, we estimate a Rank-Weighted Average Treatment Effect ([Yadlowsky et al., 2024](#)). The method was originally designed to evaluate treatment prioritization rules of policies. In our setting, the RATE is used to evaluate heterogeneity in treatment effects, i.e., asking the question: what would the gap in acceptance rates look like if we were only to send connection requests to observations whose CATE is below that of the q^{th} quintile of CATE. Note that in our context, lower CATE signify larger gaps in acceptance rates. To implement this, we use the test data’s CATE estimates based on the training forest to order observations by CATE. We then create a Target Operating Characteristic (TOC) curve, which compares the total effect on acceptance rate gaps of only treating those in the q^{th} quintile of predicted CATE to treating everyone, i.e., the Average Treatment Effect (ATE). The curve is shown in Figure G.9 and suggests substantial heterogeneity in treatment effects, with those in the highest quintile exhibiting substantially higher gaps in acceptance rates than suggested by the ATE across the entire data. The curve significantly differs from zero, again, suggesting heterogeneity (estimate: -0.027; $p < 0.001$).




Figure G.9: Target operating characteristic (TOC) curve

Note: The above graph compares the average treatment effect on treating everyone with predicted CATE above the q^{th} quintile to the average treatment effect (ATE). 95% confidence intervals are shown in light grey.

Finally, we also test the calibration of our training forest on the test data ([Chernozhukov et al., 2018](#)). More specifically, we run a linear regression on the difference between the actual decision to accept and a target’s average acceptance rate on two independent variables: the first is the difference between the treatment status and propensity score (which is 0.5) multiplied by the

average treatment effect. The second is the same difference multiplied by the difference between the edge-specific CATE minus the average treatment effect. We predict both the ATE and CATE using the training forest on the test data. The first coefficient shows a value of 0.94 ($p < 0.001$). This suggests that, on average, the ATE is precisely estimated on the test data. The second coefficient is equal to 0.97 ($p < 0.001$). The coefficient’s significance suggests the presence of heterogeneity, given that an increase in predicted CATE is associated with an increase in the difference in acceptance rates. The coefficient’s size close to 1 suggests that the heterogeneity is well-estimated. Based on the tests above, we conclude that the causal forest does suggest the presence of heterogeneity and is well-calibrated.

We continue by analyzing CATE. For this, we follow [Athey \(2020\)](#) and use the training forest to obtain out-of-bag CATE estimates. Thus, for each observation in the training data, we obtain a predicted gap in acceptance rates between requests coming from White or Black accounts based on all included covariates. The result is shown in Figure G.10, suggesting substantial heterogeneity with CATE, i.e., predicted gaps in acceptance rates between White and Black profiles, ranging from around -0.10 to 0.05. It is important to note that the estimates strongly depend on included variables, i.e., true CATE could be both more or less widespread.




Figure G.10: Distribution of conditional average treatment effects based on causal forest

Note: This figure shows the distribution of predicted CATE based on the Causal Forest with 50k trees. The distribution is chosen to include 100 bins. The red dotted line shows the average treatment effect, i.e., the average gap in response rates between Black and White users. Note that while the figure does provide an idea of predicted treatment heterogeneity, it does not represent the true heterogeneity in the data. The reported CATE strongly depends on the included variables. True CATE could be both more and less widely distributed, as noted by [Athey and Wager \(2019\)](#).

In the next step, we follow [Athey and Wager \(2019\)](#) by treating the forest above as a pilot forest to identify variables to focus on. More specifically, we estimate each variable’s importance, i.e., the share of splits these are responsible for when growing the forest.⁵³ We then restrict our attention to nine variables with above-median variable importance. Age and gender have the, by far, highest variable importance, being responsible for around 25% of splits each.⁵⁴ Based on these variables, we estimate a second causal forest with 50,000 trees on all observations.

⁵³Note, that this does not necessarily imply that variables with low importance are not responsible for heterogeneity.

⁵⁴The number verified skills is responsible for around 12% of splits. All other variables have a variable importance between 3 and 6%.

Figure G.11 shows the variables included in the final forest. The results for age, gender, and Republican vote share are in line with our findings above. Looking at age results, the figure suggests that the average age of individuals predicted to have the lowest CATE is 28. Given that CATE reflect the predicted gap in acceptance rates between Black and White users, users in lower quantiles discriminate the most. Moving to the right, the average age increases to around 40. Thus, amongst users predicted to discriminate the least (based on all 9 included variables), the average age is 40. Further, a number of county-level variables are among the variables with the highest variable importance. First, a county’s Black-White race segregation shows a U-shaped relationship, suggesting higher segregation in both the most and least discriminating counties. Further, we find that individuals from counties with lower economic connectedness discriminate more strongly. This relates to the results by Chetty et al. (2022a,b), who show that, on the county level, more diverse counties show lower levels of economic connectedness. Our results are indicative that this is, at least partially, driven by discrimination. As a direct measure of local-level implicit discrimination, we further document that local-level race IAT scores increase in CATE, i.e., a measure of implicit negative stereotypes towards Black individuals. Finally, three strong predictors of lower CATE are a higher number of connections, skills, and skill verifications obtained via LinkedIn.

G5.3 Zooming in on the Main Predictors of Discrimination

In this section, we zoom in on the three most important predictors of discrimination: age, gender, and race.

Age as a Predictor of Discrimination In this paragraph, we zoom in on the effect of age on discrimination. Figure G.12 illustrates the probability of accepting a contact request as a function of the target’s age as a continuous variable and categorized into generations. We clearly see that the probability of accepting a contact request is highly decreasing in age. Specifically, the acceptance rate drops from almost 40% for 20-year-old users to less than 20% for users older than 60. Categorizing age by generation we find that the connection gap is most pronounced for Gen Z users (users born between 1996 and 2010) and Millennials (users born between 1981 and 1996) (see the middle panel of the figure). For Gen X (users born between 1965 and 1981) and boomers (users born before 1965) the gap is not significantly different from zero. At the same time, we have fewer observations for older users and, more importantly, the acceptance baseline is different. Therefore, in the bottom panel of Figure G.12, we focus on the absolute gap and the relative connection gap accounting for the acceptance probability of a White profile request. We see that the biggest absolute and relative gap is found for Gen Zs. While the second biggest gap is found for Millennials the relative gap of Millennials and Boomers is indistinguishable, and we find there is an 11% gap in the acceptance rate of a White and Black user’s connection request. Gen X has the lowest absolute and relative connection gap (which is still different from zero).

Thus, we clearly find that age is predictive of discrimination, with the lowest connection gap found for Gen Xs, and the biggest gap for Gen Zs. However, the question is what is driving this result. There are multiple possible reasons why age is predictive of discrimination.

For one, it might be that younger users employ LinkedIn not only to develop a network and focus on work-related aspects, but might also be using it as a social media website by posting content, commenting, and responding to content. If that were to be true, then younger users’ racial preferences would be weighted higher in their utility function as more interaction is anticipated. To speak to this explanation, we can look at how age is related to the probability of posting/commenting, and also to the number of followers a user has. Counter to our expectation, we find that older users are *more* likely to be engaged on the platform ($\beta=0.04, t(16964)=4.80, p \leq 0.001$), and also have more followers than young users ($\beta=0.23, t(16939)=18.75, p \leq 0.001$). Thus, it seems unlikely that younger users discriminate based on their anticipated interaction with a new contact.

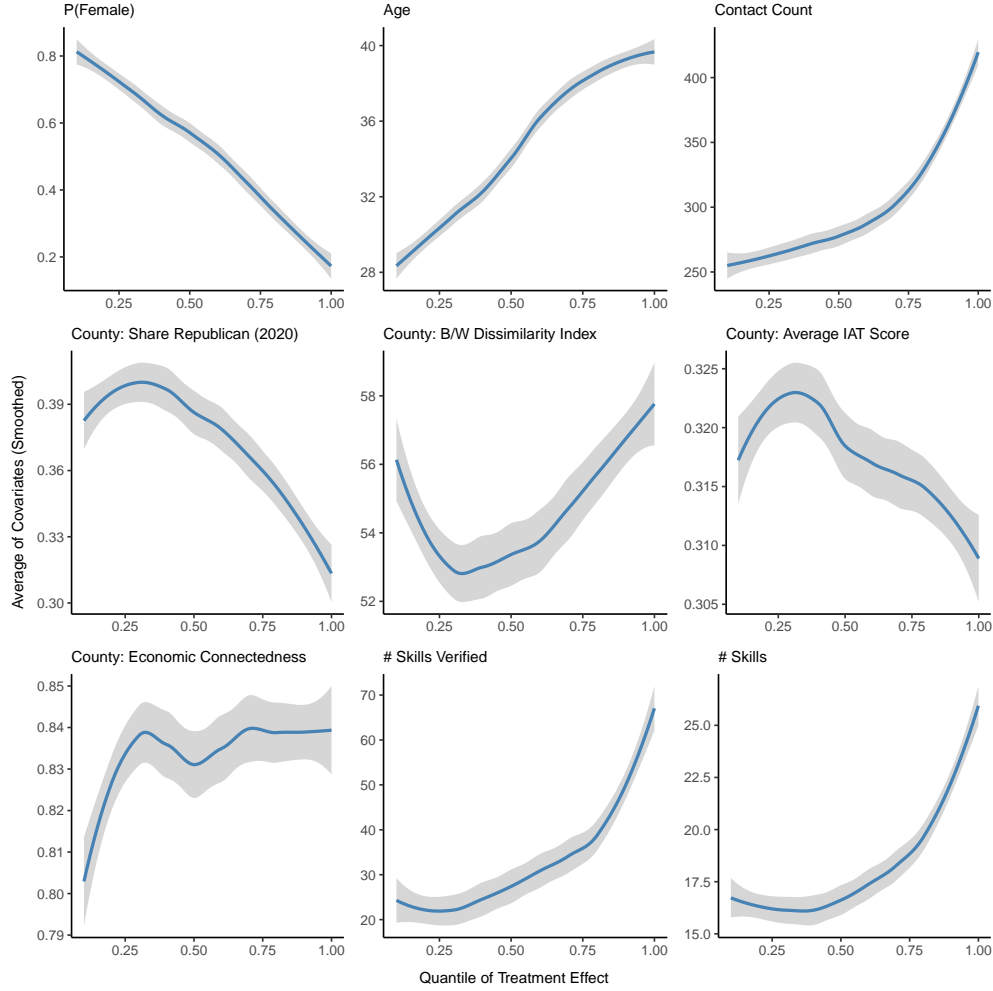


Figure G.11: Quantile of estimated CATE and conditional mean of covariate

Note: This visualization follows [Athey \(2020\)](#). On the x-axis, it shows the quantile of out-of-bag CATE estimates across all targets based on the causal forest trained on the nine included variables with above-median variable importance. The y-axis shows the (smoothed) conditional mean of covariates and 95% confidence intervals using a local polynomial regression (LOESS). Estimates are based on a causal forest with 50k trees.

One possible alternative explanation would be that older users differ from younger users in terms of observable characteristics. To account for them, we estimate a model where we account for multiple target characteristics, reported in Table J.6. We find that the age results remain rather unchanged by controlling for those features. Thus, the age result is not merely a byproduct of some other characteristics. A related explanation could be that younger users more strongly view our profiles as competition in the labor market, given that these have a similar age and might thus compete for the same jobs. Conflict theory suggests that stronger competition between groups increases prejudice against out-groups and/or in-group favoritism ([Halevy et al., 2012](#)).

A final plausible explanation for the age effect in discrimination would be differential selection into LinkedIn usage. Specifically, it seems plausible that younger users are less selected as it is more common for younger users to have and use LinkedIn. Older people, on the other hand, are less likely to use LinkedIn in the first place, and therefore older LinkedIn users might be more selected. However, while this may be the case in our sample, a representative survey of the US

population suggests little differences in the use of LinkedIn across different age groups 18-29 (30%), 30-49 (36%), 50-64 (33%) ([Brooke Auxier, 2021](#)).

Overall, we do not feel comfortable in drawing any conclusions regarding the reason for the heterogeneity with respect to the age we observe in our paper. Above, we have put forward a number of possible explanations though we acknowledge that there might be more.

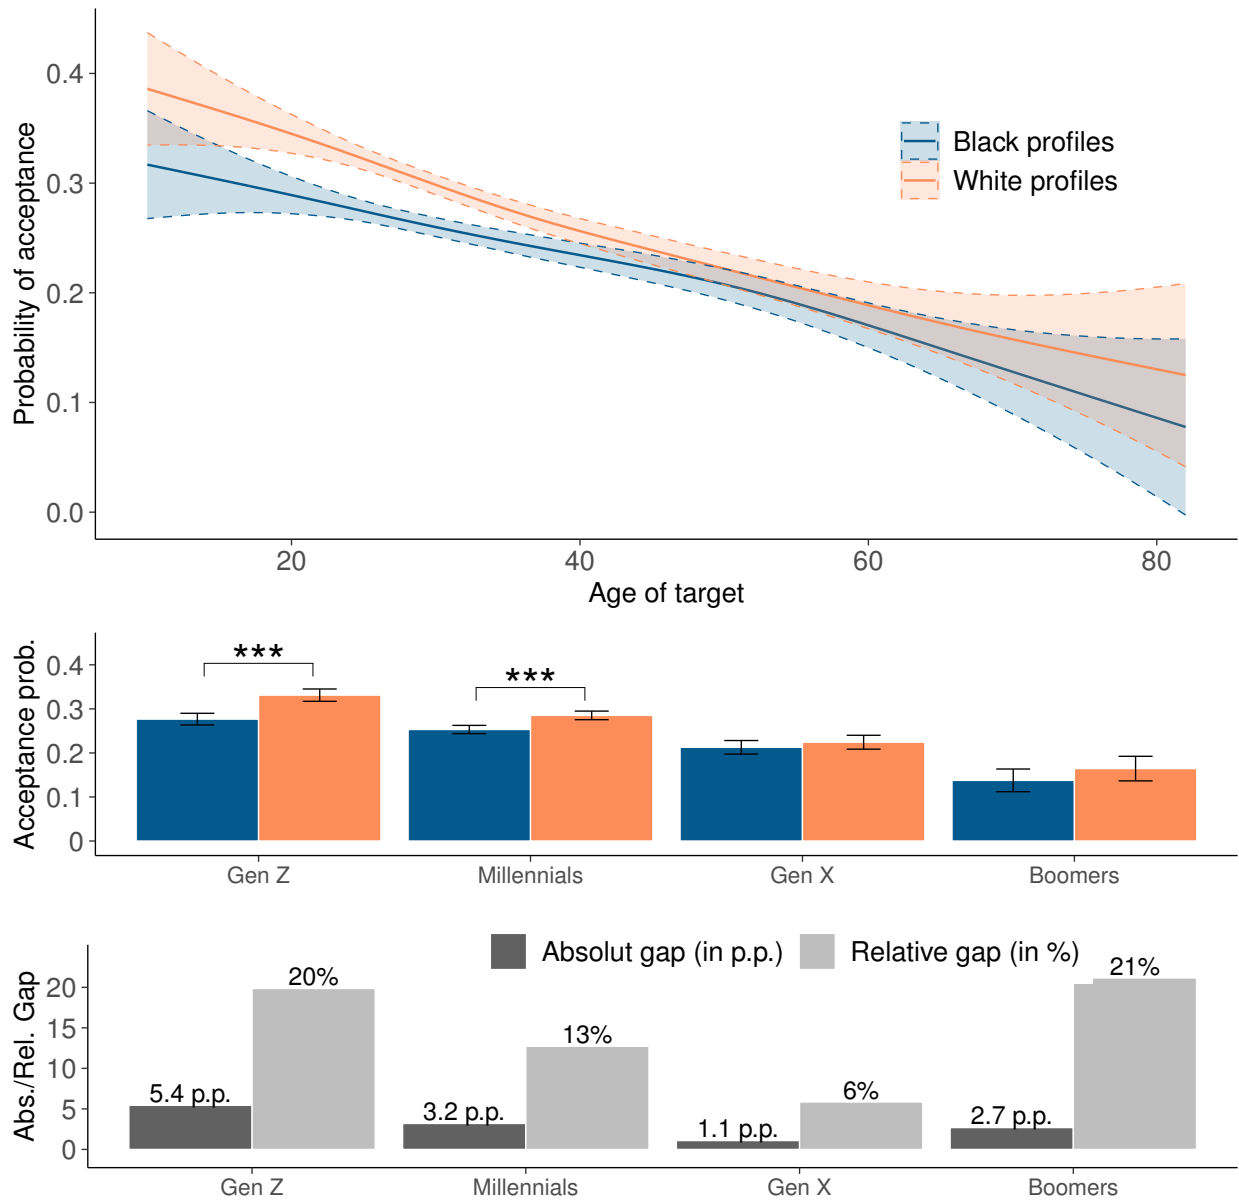


Figure G.12: Acceptance probability of a connection request by a Black and White profile as a function of the target's age

The figure illustrates the acceptance probability of a connection request as a function of the target's age. The top panel illustrates the fitted acceptance probability as a function of the target's age. The middle panel depicts the acceptance probability by generation of the target. The bottom panel illustrates the relative gap (i.e., accounting for the acceptance probability of a White profiles request) in acceptance probability by generation of the target. Orange and blue objects denote the White and Black profiles, respectively. Whiskers around the mean, and bands around the spline, denote the corresponding 95% confidence intervals. T-tests are used to obtain the following significance levels: $\cdot p < 0.10$; $\ast p < 0.05$; $\ast\ast p < 0.01$; $\ast\ast\ast p < 0.001$.

Gender as a Predictor of Discrimination The best predictor of discrimination is gender. As shown in Figure G.8, women are discriminating more, not less, than men. Here we take a closer look at how exactly women differ from men in terms of discrimination and study some possible explanations of this phenomenon.

Figure G.13 illustrates the acceptance probability of a connection request of a Black and White profile as a function of the target’s probability of being a woman based on their first name. White profiles have a constant probability of roughly 26% of being accepted independent of the target gender. Thus, men and women seem to accept White profiles to the same extent. For Black profiles, that pattern changes. Black profiles have an acceptance rate of 24% if the target is likely a man. This acceptance probability, however, reduces monotonically in the target’s probability of being a woman and reaches almost 22% if the target is very likely a woman. Thus, we observe a clear connection gap already for men, but this gap is magnified by women.

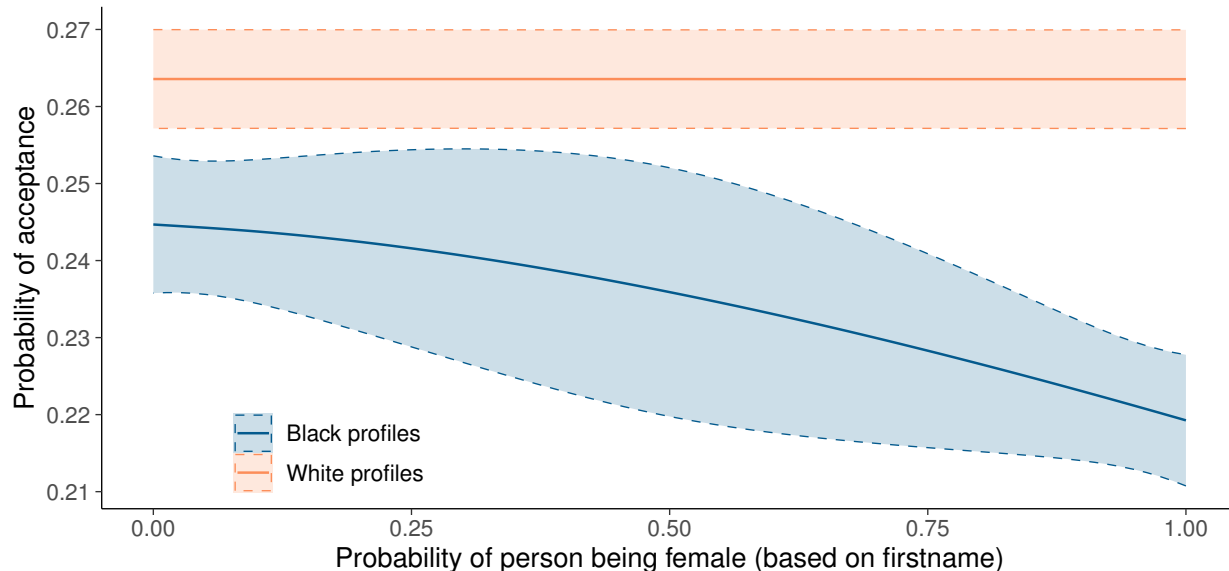


Figure G.13: Acceptance probability of a connection request by a Black and White profile as a function of the target’s gender (based on the first name)

The figure illustrates the fitted acceptance probability of a connection request as a function of the target’s gender (determined by their first name). Orange and blue objects denote the White and Black profiles, respectively. Bands around the spline denote the corresponding 95% confidence intervals.

One possible explanation could be that women anticipate more interaction (e.g., due to sexual harassment) and this anticipated increase in interaction (and the corresponding costs) could be driving the differences. This, however, would also imply that women are less likely to accept a connection request in the first place if ‘romantic’ advances are anticipated. This possible explanation that women are just generally less likely to accept a connection request can be alleviated by Figure G.13. It clearly shows that women have the same acceptance probability of a White profile as men. The gender difference is fully driven by the behavior toward Black profiles.

A similar explanation could be that other correlated features might be driving this effect. Therefore, in Table J.7, we account for the number of contacts and for a multitude of other target characteristics. The effect remains essentially unchanged. Thus, the gender effect is not a mere byproduct of another target characteristic.

One plausible alternative explanation could be the target’s romantic interest. Even though LinkedIn clearly is not primarily a platform for finding romantic partners – in fact, romantic advances are a violation of LinkedIn’s Professional Community Policies – workplace relationships are not uncommon. Some articles even suggest that LinkedIn might be a good website to find

partners, as all information is public and vetted.⁵⁵ Further, [Rosenfeld et al. \(2019\)](#) show that 11% of people meet their partner through or as a coworker, and almost 40% meet their partner online. If some of our targets perceive a connection request not only as a professional connection but also as a romantic advance, dating preferences might affect their decision to accept the profile. If that were to be the case, we would expect the race of the target to also affect that decision as dating preferences in the US are rather clearly split by race ([Kalmijn, 1998](#); [McClintock, 2010](#); [McPherson et al., 2001](#)). Further, having only male profiles we expect race to matter more for women than for men given a majority heterosexual population. Figure [G.14](#) illustrates the acceptance probability of a connection request of a Black and White profile as a function of the target’s probability of being a woman and the target’s probability of being White/Black. First, we see that men do not differ in their behavior towards a Black and White profile as a function of their own race. Specifically, the gap remains rather constant as a function of the target being White or the target being Black. We also clearly see that the estimates are very imprecise if the race of the target is increasingly likely Black, which is driven by a rather small sample of targets conclusively estimated to be Black. For women, we see that the target’s probability of being White or Black does not affect the acceptance probability of a White profile. The acceptance probability is relatively stable at 26%. However, the acceptance probability of a Black profile does clearly change with the probability of the target’s race. Female targets who have a higher probability of being White (bottom panel) have a decreasing probability of accepting a Black profile. The opposite observation is true with an increasing probability of the target being Black. Here we see that Black profiles are increasingly likely to be accepted and are even more likely to be accepted than White profiles as a function of the target’s probability of being Black (top panel). However, the acceptance rate is very imprecisely estimated for increasing the target’s probability of being Black. When looking at the interaction effect (see [Table J.8](#)) we find that the estimates are very much in line with the figure. However, the standard errors are too large to reject the null hypothesis. Thus, the picture can be taken as suggestive evidence that dating preferences might explain the connection-gap difference between male and female targets. However, as shown in [Figure I.3](#), our survey clearly shows that LinkedIn is very rarely used for dating and dating preferences are not important for LinkedIn usage. For instance, 99% of users indicate to rarely or never use LinkedIn to search for romantic partners while more than 90% of users rarely or never received romantic advances or experienced harassment on LinkedIn. An alternative explanation, in line with our data, might be stereotypes against Black men specifically held by or salient to White women (e.g., [Davis, 1981](#); [Sommerville, 1995](#); [Zounlome et al., 2021](#)).

⁵⁵See e.g., [Insider post](#) or [LinkedIn Blog](#).




Figure G.14: Acceptance probability of a connection request by a Black and White profile as a function of the target's gender (based on first name) and race (based on last name)

The figure illustrates the fitted acceptance probability of a connection request as a function of the target's race (determined by their last name) and gender (determined by their first name). The top/bottom panels illustrate the fitted acceptance probability as a function of the probability of the last name being Black and White, respectively. The left panels illustrate the behavior of women, while the right panels illustrate the behavior of men. Orange and blue objects denote the White and Black profiles, respectively. Bands around the spline denote the corresponding 95% confidence intervals.

Race as a Predictor of Discrimination As highlighted in the previous section and also shown in Figure G.8, race predicts discrimination. Here we focus in more detail on this effect by splitting our measures of race. Figure G.15 illustrates the acceptance probability of a connection request of a Black and White profile as a function of the target's probability of being White/Black. First, we see that the gap does not change as a function of the probability of the target being White

(bottom panel). Thus, on average, White and non-White targets discriminate to the same extent.

Things are different if we focus on Black vs. non-Black participants. The top panel displays how the connection-gap changes as a function of the target’s probability of being Black. We find that there is a considerable connection gap if the probability of the target being Black is relatively small ($<.2$). However, with an increasing probability of the target being Black, we find that the gap reduces and basically disappears if the target is very likely Black. This observation is primarily driven by the behavior toward Black profiles. Specifically, targets slightly increase their acceptance probability of White profiles in their probability of being Black, but they increase their likelihood of acceptance of a Black profile even more. Thus, the absolute gap is small, and the relative gap is even smaller, as targets with a higher probability of being Black are even more likely to accept a profile.

Table J.9 shows estimations of the effect. In line with the figure, we find no change in the connection gap as a function of the probability of the target being White. We, however, do find that the connection gap is reducing in the probability of the target being Black. As shown in Figure G.14 above, this is, interestingly, driven by Black women. Black men, on the other hand, retain a similar absolute acceptance rate gap.

Overall, we conclude that Black targets are less discriminating. This effect is primarily driven by the behavior toward Black profiles, and we do not find a similar result when focusing on non-White users. Further, we see from the section above that this race result is mostly driven by women.




Figure G.15: Acceptance probability of a connection request by a Black and White profile as a function of the target's race (based on last name)

The figure depicts the acceptance probability of a connection request as a function of the target's race, determined by their last name. The left panels illustrate the fitted acceptance probability as a function of the probability of the last name being Black (top panels) and White (bottom panels). The right panels use the binary variable instead of the continuous variable (x-axis). Orange and blue objects denote the White and Black profiles, respectively. Whiskers around the mean, and bands around the spline, denote the corresponding 95% confidence intervals. T-tests are used to obtain the following significance levels: $p < 0.10$; $*p < 0.05$; $**p < 0.01$; $***p < 0.001$.

G6 Some Selected Replies

As mentioned in the main part of the paper, most respondents shared some experience, information, or generic advice, while others provided substantially more elaborate and valuable responses.

Those new connections offered to meet or talk on the phone, refer our profiles to another more knowledgeable co-worker, and were even willing to function as a reference for future applications. Overall, almost 65% of the responses contained some useful content (offered a referral, shared detailed information, etc.).

However, the true value of these messages is obscured by statistics. Providing some specific examples can give a clearer insight into their content. Below, we show four (for privacy reasons slightly adjusted) messages, which aim to show how valuable the response might be for an application.

“Thanks for reaching out. **I would connect with [Name of a Person] and feel free to mention my name.** We have a lot of people that are really motivated and driven to succeed. My advice for any interview is to highlight your ambitions and be confident. Best of luck.”

“Hi [Name], Glad to connect. [Company] looks for people who have an entrepreneurial mindset and are looking to pave their own way in their careers. The interview process will vary between the person/group. My interview experience was much more of a conversation about what I was looking for, how I felt my experience could benefit [Company] and my questions for the interviewer, rather than a typical set of interview questions. I’d make sure your resume includes all the softwares/programs you’ve used, as recruiters will look for certain keywords when reviewing resumes. **I’m happy to submit you in as a referral if you like. This will help get you to the front of the line for applicants.**”

“Hi [Name] - That’s great! A couple tips ... depending on which part of the business you’re looking to support, admin roles can vary a bit, however, some common skills and experiences that we look for are: organized, proactive, taking initiative, experience with systems like outlook, workday, and zoom, comfortable with reporting and learning new technology, resourceful, and building strong relationships across organizational lines. Our company values are rooted in connection, inclusivity and drive. So, speaking to your experiences and how you get work down through that lens will also be helpful. If you’re interested in a role supporting our field and store teams, we have some movement on our admin team in my region, and **I’d be happy to pass your resume along to our recruiter.** Let me know!”

“[...] I left [Company] after nearly 13 yrs, I needed a change. Great company but just like all mortgage cos right now **they are downsizing.** Good luck wherever u wind up.”

All messages highlight the value of engaging with new contacts. The first messages offer crucial details about the application process and required skills, with offers to support the application or submit a referral. Even the last message, though short, is important as it signals company downsizing, which might be highly informative when thinking of applying.

G7 Usefulness of Messages

To determine the usefulness of a message we use three proxies: 1) the length of the message, 2) whether the person offers a referral or a meeting (coded by two RAs), and 3) how useful the message is considered by a large language model.

To obtain a usefulness rating of a large language model, we used OpenAI’s ChatGPT-4 and prompted the following command: *“The following is an answer to a young professional who asked*

another LinkedIn user for advice regarding the application process at his firm. Please rate the following answer on a scale from 0 to 10, where 0 means ‘not useful’ and 10 means ‘highly useful’. The rating should take into account both the message’s content and whether the user asks follow-up questions or offers to get in touch. As an answer only return a single number. Do not elaborate on your choice. Message: [...]”.⁵⁶ For the mentorship treatment, the first sentence was replaced by “The following is an answer to a young professional who asked another more experienced LinkedIn user for general career advice.”. The average message consists of 60 words, 6% of responses offer a referral or a meeting, and the median response has a usefulness rating of 7 (on a scale between 0 and 10).

To validate all three approaches, two RAs hand-coded all 681 messages, identified which characteristics the message contained (like whether a referral or a meeting was offered), and rated the usefulness of the messages (see Table J.15 for further characteristics). Table J.16 reveals that a one-standard deviation increase in the length of the message increases the usefulness by 1.2 points (from a baseline of 6.22), similarly a referral or a meeting increases the usefulness by more than one point, and a one-standard deviation increase in usefulness rating of RAs is associated with a 1.3 points increase in the usefulness rating of the language model. These effects indicate that first the length and a referral or message increase the usefulness of the message, and second that the usefulness rating of the language model is highly associated with the usefulness rating of RAs, validating our approach.

G8 Predictors of Message Response

In this section, we discuss predictors of a message response. Specifically, we are interested in understanding who is more likely to respond to our message inquiry. Figure G.16 reports upon multiple relevant characteristics and how they are associated with responding to the message of our profile.

Given the substantial reduction in sample size compared to the first stage, most estimates are rather noisy and do not differ significantly from zero. Essentially the only relevant predictors of a response are the education of the target, how active the target is on LinkedIn, and whether the target and our profile attended the same university. Specifically, better-educated targets (i.e., users who have either an associated degree or a bachelor’s degree as their lowest degree) have an almost 5 p.p. higher probability of responding to our message, and comparably people without a degree have a 6 p.p. lower probability of responding. If our profile and the target attended the same university then they have a 7 p.p. higher probability of responding to our message. Finally, one standard deviation increases in the log of the number of followers, and similarly, one standard deviation increase in the number of contacts increases the probability of responding by roughly 3 p.p. One of the strongest predictors of whether a target responded is whether they decided to display volunteering experience on their CV. Those with volunteering experience are almost 6 p.p. more likely to respond to our message than those targets without volunteering experience. Note that other characteristics, in particular our profile characteristics, do not predict response behavior. Specifically, whether our profile is White in the first stage or whether our profile is White in the second stage does not have a significant impact on the probability of getting a response to the message.

⁵⁶For better replicability, the model’s ‘temperature’ was set to 0.




Figure G.16: Characteristics predicting message response

The figure illustrates the β -coefficients of the following regression: $response_{i,j} = \alpha + \beta \cdot Variable + \epsilon_i + \epsilon_{i,j}$. ϵ_i is the user random effect with $(\epsilon_i \sim \mathcal{N}(0, \sigma_1^2), \epsilon_{i,j} \sim \mathcal{N}(0, \sigma_3^2))$. $Variable$ denotes the z-scored variable, if the original variable is not binary. The regression thus computes the response rates as a function of one feature of the target/or profile, while accounting for the fact that each profile has contacted multiple targets.

G9 Effects of Picture Swapping

A possible concern a reader might have is that the swapping of profile pictures after the first stage of the experiment might taint our results. Specifically, the concern could be that targets realize that a former White profile is now Black (or the other way around), and that could lead to biased results. To alleviate that concern, we provide three pieces of evidence that speak against it. First, we will show that the number of profile views does not change differently over time for profiles whose picture has been swapped and for profiles whose picture has not been changed. Second, we will show that the number of suspended ties (i.e., removed profiles from the own contacts) is not affected by the picture change. Third, we will show that responses (in terms of probability, length, and usefulness) do not change as a result of the picture swapping.

Profile View Frequency One might anticipate that changes in profile pictures could pique a target’s curiosity, compelling them to visit the profile’s website for closer examination of other potential changes. However, Figure G.17 paints a different picture. It visualizes the frequency of profile views before (25.07) and after (08.08) the swapping of profile pictures, which took place between July 28th and August 1st, 2022. The left panel illustrates the difference in views between profiles whose pictures have and have not been changed. We find no difference in views between these two groups. Prior to the swap, the to-be-swapped group received, on average, a marginal 0.23 ($p = 0.752$) more views compared to the non-swapped group (relative to a baseline of 36 views). This minimal difference endured post-swap (One and three weeks after: 0.33, 0.52), and remained insignificant even after all messages were dispatched ($p = 0.66$, $p = 0.485$). Regression estimates reported in Table J.22 reinforce these findings, indicating no discernible difference in visit frequencies or dynamic changes over time between swapped and non-swapped profiles.

Thus, the face-swapping seems not to have been suspicious enough for targets to view our profiles’ sites.

Suspended Ties However, one could be concerned that targets realized that something was off with the profile and just directly suspended the connection. While this is only possible when visiting the profile’s site, we still want to investigate this. However, suspension of connection was extremely rare as of all the 9523 established links, merely 101 were suspended. Delving deeper, we found most suspensions were enacted by individuals severing ties with both Black and White profiles they connected with, possibly signaling their exit from LinkedIn. Moreover, the suspension rates between swapped and non-swapped profiles are virtually identical (49 vs. 52 suspensions, or 1.02% vs. 1.1% suspensions, $p = 0.723$).

Figure G.18 also illustrates the suspension probabilities of all connected targets (left panels) for swapped and non-swapped profiles by their race on the picture after the swapping. However, one still could be concerned that targets don’t realize that the race of the new connection changes as long as they are not messaged. However, after receiving a message, these targets are made aware of the change. Therefore, we split the sample into those targets who have received a message (see the middle panels) and those who have not (see the right panels) in Figure G.18). The negligible suspension probability persists across profiles that swapped pictures and those that didn’t. This trend remains unchanged when we split by profiles now exhibiting a Black or White picture post-swap. These observations are confirmed by regressions reported in Table J.23.

To wrap up, connection suspensions are a rare phenomenon, potentially indicative of targets exiting LinkedIn rather than reacting to profile picture swaps.

Response Characteristics Despite the minor variations observed, lingering concerns may persist that targets, having noticed the changes, opt against severing connections outright, choosing instead not to engage in message responses or questioning profile changes within their responses. Yet, considering responses are an outcome in and of itself, we cannot simply draw comparisons between swapped and unswapped profiles. A concern is that the original network and discrimination might interact, making it difficult to interpret whether the mere swapping is responsible for changes or whether the characteristics of the network interact with the race of the asking subject. More specifically, a network mismatch may simply result in less responsiveness. Figure G.19 presents the response probability, normalized response length (in characters), the likelihood of the response being highly useful (i.e., replies that offer a referral or a meeting), and the usefulness of the message. In line with the match-network hypothesis, the probability of responding is marginally lower, however, only at the 5% level and only for better universities.

However, to cleanly isolate the effect of face-swapping on responses we can leverage the time passed between accepting a connection request and receiving a message. The idea here is to use

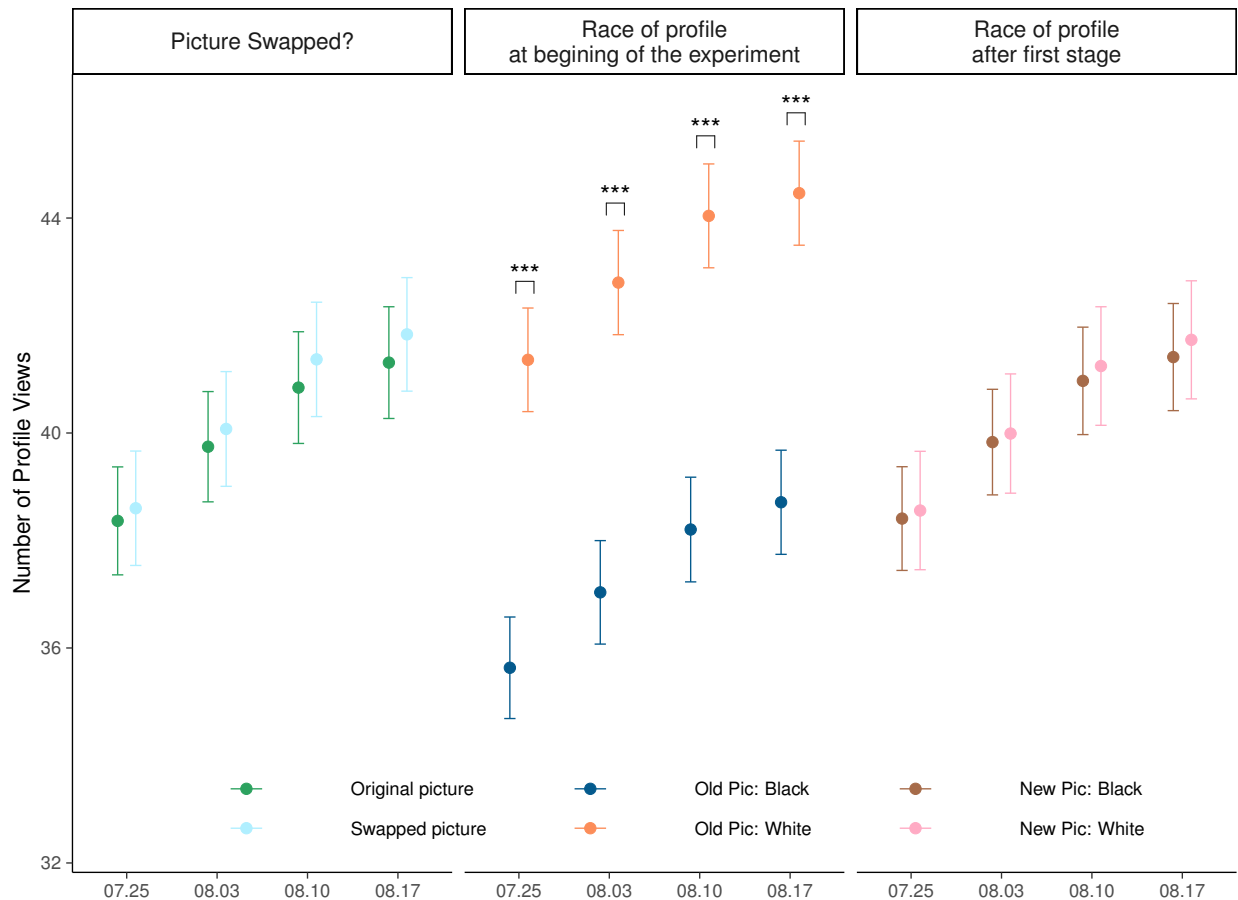


Figure G.17: Number of profile views before and after face-swapping

The figure depicts the number of profile views as a function of profile characteristics. The left panel compares profiles whose picture has not been swapped (i.e., their original picture) in green and profiles whose picture has been swapped to their twin's pictures (i.e., a formerly White profile uploaded a picture of their Black twin, and vice versa) in blue. The middle panel compares originally Black and White profiles. The right panel compares profiles based on their second-stage race, i.e., profiles that have (or will have) a Black or White profile picture in the second part of the experiment. Whiskers around the mean denote the corresponding 95% confidence intervals. T-tests are used to obtain the following significance levels: · $p < 0.10$; * $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$.

the fact that targets were contacted in waves in the first stage of the experiment. Consequently, for some of the targets more than 8 weeks have passed between seeing the profiles the last time (when accepting) and receiving a message, and for other targets, only two weeks have passed. If targets were to observe the swapping and react to it, we would expect the time passed between accepting a connection request and receiving a message to affect the response. Table J.24 illustrates the corresponding regressions. We find no evidence that the time passed between accepting a connection request and receiving a message impacts the response probability, the length of the message, or the value of the message. Thus, mere face-swapping does not affect the responses.

Taking everything into account, the lack of any differences in the number of profile views, connection suspensions, and response traits suggests there is no evidence to support concerns that face-swapping has significantly altered target behavior.⁵⁷

⁵⁷While our evidence strongly suggests that users do not realize that the profile's picture changed, there nonetheless remains the possibility that some users did. However, such users would only affect our overall




Figure G.18: Suspension probabilities by swapped and original profile pictures

The figure depicts the probability of a connection suspension for profiles whose picture has not been swapped (i.e., their original picture) in green and profiles whose picture has been swapped to their twin's pictures (i.e., a formerly White profile uploaded a picture of their Black twin, and vice versa) in blue. The top panel reports the outcome for originally Black profiles, while the bottom panel illustrates originally White profiles. The left panel aggregates over all targets, while the middle and right panel illustrates the responses of targets who have **not** and have been contacted by a message, respectively. Whiskers around the mean denote the corresponding 95% confidence intervals. T-tests are used to obtain the following significance levels: $p < 0.10$; $p < 0.05$; $p < 0.01$; $p < 0.001$.

results if realizations of the swap were to happen with a different likelihood for Black and White profiles. Given that users in our experiment and on LinkedIn are predominantly White, such a realization would likely be more salient for profiles being Black after the swap. As we did not receive any messages asking about the swap itself, it is likely these users simply chose not to respond. A realization would thus show up in a lower response rate. As a result, a differential realization of the swap would bias results for discrimination upward. Given that we, nevertheless, find no evidence of discrimination in the experiment's second stage, this is further evidence that the swap is not being realized.

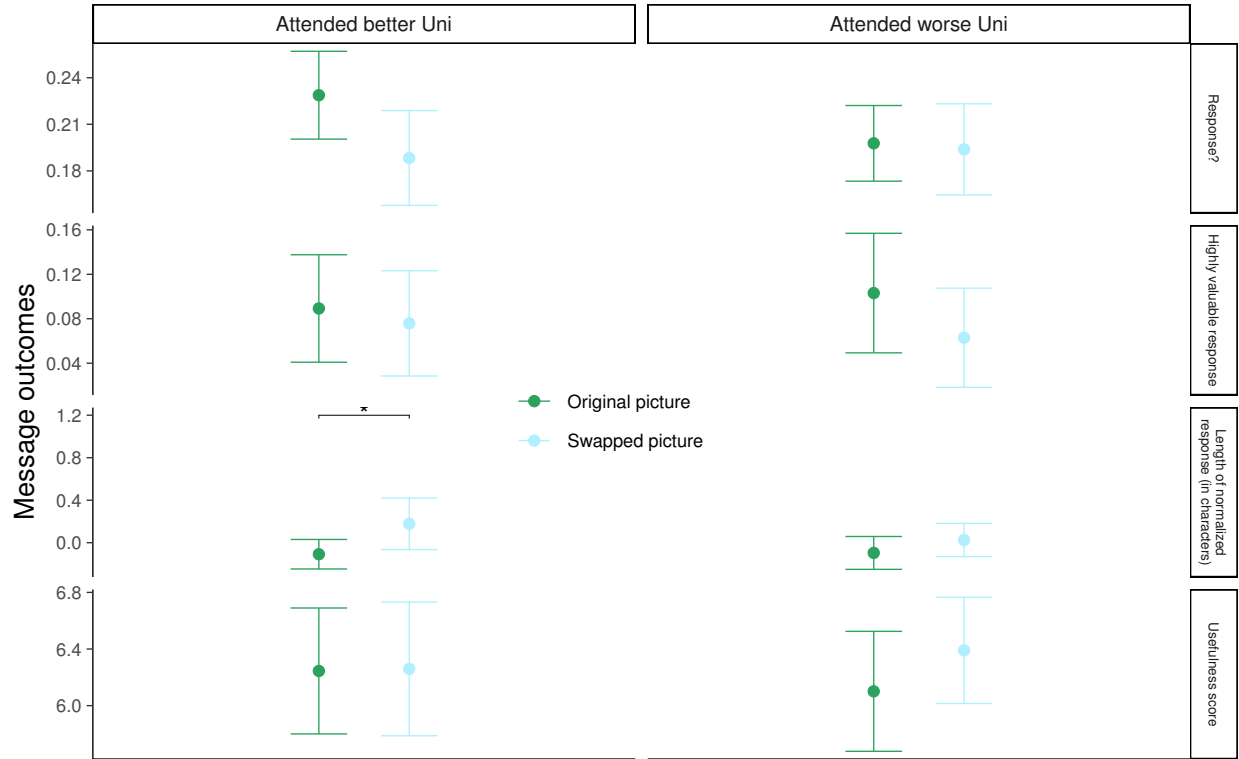


Figure G.19: Response characteristics by swapped and original profile pictures

The figure depicts several response characteristics for profiles whose picture has not been swapped (i.e., their original picture) in green and profiles whose picture has been swapped to their twin's pictures (i.e., a formerly White profile uploaded a picture of their Black twin, and vice versa) in blue. The top panel reports the probability of a response, the subsequent panel illustrates the probability of a response being highly valuable, the panel second to last illustrates the normalized length of the response, and the panel at the bottom illustrates the usefulness of the message. Whiskers around the mean denote the corresponding 95% confidence intervals. T-tests are used to obtain the following significance levels: $\cdot p < 0.10$; $* p < 0.05$; $** p < 0.01$; $*** p < 0.001$.

G10 Heterogeneity in Responses

As previously shown we essentially find no difference in responses towards Black and White requests. Nonetheless, discrimination may still be present within certain target groups. Given that we essentially randomly allocated the race of our profiles *after* our profiles have been accepted, we can focus on each subgroup without being concerned that self-selection is driving behavior. However, when comparing subgroups (e.g., men versus women reacting to Black or White profile requests) we should keep in mind that there was some self-selection in forming a tie. For example, hypothetically it could be that women who do accept a link are generally more helpful, while men are always accepting, and therefore are less selected.

As a first step, we examine subgroups to pinpoint discriminatory responses, focusing on five key discrimination predictors: age, gender, education, network size, and political leaning.⁵⁸ Figure G.20 shows no substantial response disparity to Black or White profile messages across these five subgroups, except for targets with less than a Bachelor's degree, who respond slightly less often to Black profiles.

⁵⁸We exclude race here due to sample size limitations.

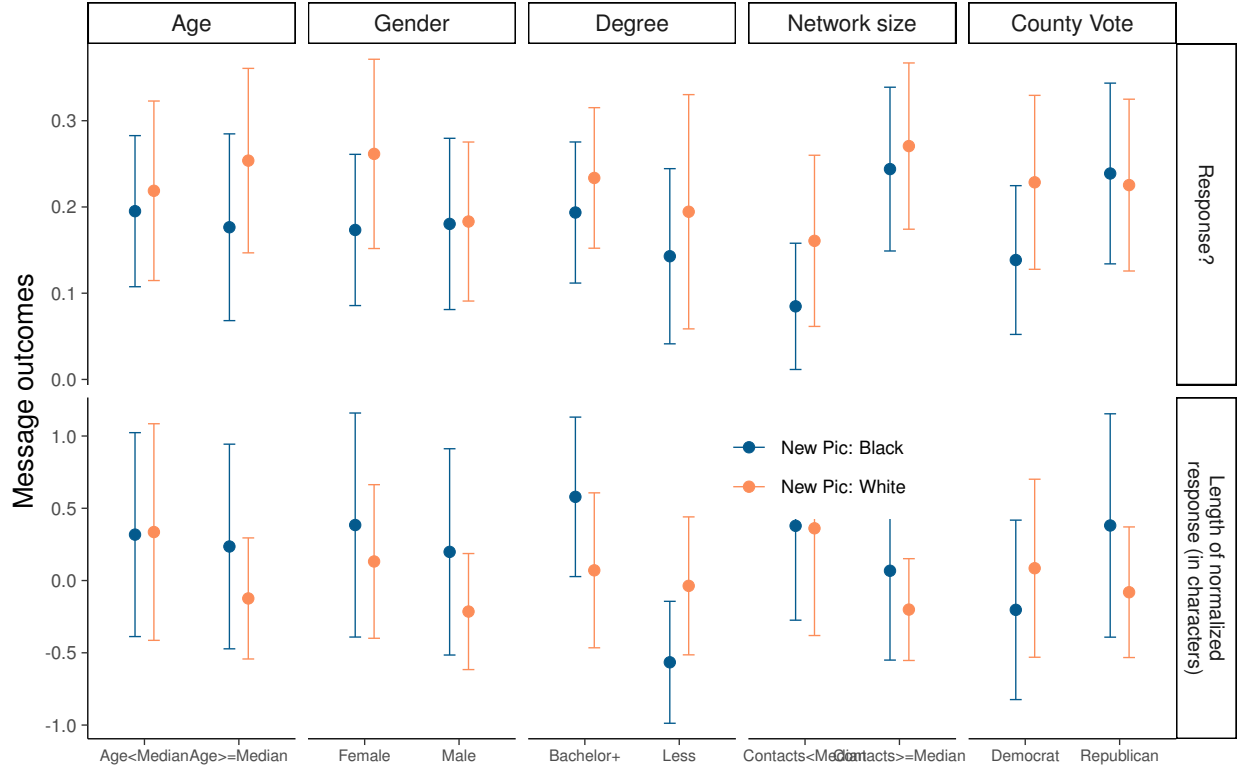


Figure G.20: Response characteristics by race of the profile pictures and basic target characteristics in the second stage of the experiment

The figure depicts several response characteristics based on their second-stage race, i.e., profiles that have (or will have) a Black or White profile picture in the second part of the experiment. The top panel reports the probability, and the bottom panel illustrates the normalized length of the response. Whiskers around the mean denote the corresponding 95% confidence intervals. T-tests are used to obtain the following significance levels: $\cdot p < 0.10$; $* p < 0.05$; $** p < 0.01$; $*** p < 0.001$.

As a second step, we compare how behavior differs between targets who accepted only a Black profile, only a White profile, or both profiles' requests in Stage I, as a broader marker of discrimination. Figure G.21 reveals that there is again not much heterogeneity. Targets who accepted both in Stage I do not differentiate in their responses between Black and White profiles at all. Things are slightly different for those who accept the request of the Black or the White profile only. Those targets who accepted the request of a Black profile are slightly more likely to respond, and to respond more helpfully to a Black profile (surprisingly, they write longer messages to White compared to Black profiles). The opposite is true for those who accepted originally a White profile (all these effects are not significant). If we focus on the interaction (i.e., is the direction of discrimination different between targets who accepted originally the Black or the White profile only), we find indeed some evidence of a difference (see Table J.20). The response probability towards a White profile is identical between these two groups of targets, but they slightly differ in how likely they are to respond to a Black profile. However, these outcomes are not very robust and are primarily suggestive.




Figure G.21: Response characteristics by race of the profile pictures in the second stage of the experiment

The figure depicts several response characteristics based on their second-stage race, i.e., profiles that have (or will have) a Black or White profile picture in the second part of the experiment. *Both* denotes targets who accepted the connection requests of both (the Black and the White profiles), while *Only Black/Only White* denote targets who accepted the connection requests of the Black/White profile only. The top panel reports the probability of a response, the subsequent panel illustrates the probability of a response being highly valuable, the panel second to last illustrates the normalized length of the response, and the panel at the bottom illustrates the usefulness of the message. Whiskers around the mean denote the corresponding 95% confidence intervals. T-tests are used to obtain the following significance levels: $\cdot p < 0.10$; $* p < 0.05$; $** p < 0.01$; $*** p < 0.001$.

To further explore heterogeneity in response discrimination, we ran regressions integrating race with factors of interest. For instance, we used a regression to analyze how likely a response was, interacting the profile's race with user's gender. Figure G.22 shows the interaction estimates. We find little heterogeneity. Some characteristics interact with the profile's race, but we do not find a clear pattern.

Which characteristics predict discrimination in the message responses (stage 2)?




Figure G.22: Correlates of discrimination in message responses

The figure illustrates the β -coefficients of the following regression: $ResponseChar_{i,j} = \alpha_0 + \alpha_1 \cdot BlackProfile + \alpha_2 \cdot Variable + \beta \cdot Variable \cdot BlackProfile + \epsilon_i + \epsilon_{i,j}$. ϵ_i is the user random effect with $(\epsilon_i \sim \mathcal{N}(0, \sigma_1^2), \epsilon_{i,j} \sim \mathcal{N}(0, \sigma_2^2))$. $BlackProfile$ denotes a dummy with value one if the profile messaging the target is Black, and zero otherwise. $Variable$ denotes the z-scored variable, if the original variable is not binary. $ResponseChar$ denotes one of the following response characteristics: the probability of responding, the length of normalized response (in characters), the probability of the response being highly valuable, and the usefulness of the message. The regression thus computes how certain response characteristics are a function of a specific feature of the target interacted with the profile's race while accounting for the fact that each profile has contacted multiple targets. For example, the negative value of "Jobtitle-President" in the left panel indicates that targets whose job title indicates "president" are less likely to discriminate against a Black profile than targets whose job title does not indicate "president".

In conclusion, there is minimal heterogeneity in the discrimination observed in the second stage of the experiment. Black profiles generally receive comparable treatment to White profiles.

G11 Ancillary Outcomes

In this section, we take a look at some ancillary outcomes. Specifically, whether our profiles received contact requests, received unsolicited messages, and, more importantly, how often our profiles have been viewed. All data are obtained after the end of the experiment’s first stage, i.e., before swapping profile pictures.

Table J.14 compares Black and White profiles with regard to these measures. We find that, on average, our profiles receive one contact request, which, however, does not differ between Black and White profiles. Further, every fourth Black profile received an unsolicited message, whereas White profiles received slightly more messages. However, accounting for the number of contacts resolves these differences, suggesting that messages mostly stem from contacts. The most important ancillary outcome is the times a profile is viewed, i.e., visited by LinkedIn users. First, we see that profiles were viewed relatively often in the past 90 days (which is the number LinkedIn reports). On average, every profile receives almost 36 views. Importantly, White profiles are substantially more likely to be viewed. Part of this difference can be explained by the difference in network size, as a one-standard-deviation increase in the number of contacts, increases the number of views by 3.

H Expert Survey

To contrast our findings to the priors of researchers working in the field, we conducted an expert survey in early June 2023. The aim of this survey was for experts working on labor economics and/or discrimination to predict the results of our experiment. The goal is to see, where our results align with experts’ priors and where they diverge. To not bias our participants, we did not have a working paper version online before the survey. However, we had presented the paper multiple times by June 2023 and have spoken to several people. Still, 90% of survey participants indicated to not have heard about the project and only 1% indicated to have heard the results of the paper.

To grasp the perspective of the most relevant audience for this project, we sent the survey to 2,143 labor economists. These were chosen from two sources. First, we contacted all economists in the Institute for Labor Economics’ (IZA) network. This includes a total of 2,091 labor economists by June 2023. Second, we obtain the email addresses of all 109 participants in the ‘NBER’s Summer Institute: Labor Studies’ from 2021 and 2022.⁵⁹ Given some overlap, this results in 2,143 Emails sent.⁶⁰ We purposefully designed the survey to be very short to have a relatively high response rate and, indeed, the median time participants required to finish the survey was about 6 minutes. Aside from demographic questions, we ask participants to predict the result of the first stage and second stage of our study and how some selected groups of users discriminate in the first stage of the experiment. The screenshots of the questions are shown in Figures H.2a, H.2b, H.2c. After having sent the invitation email once, we waited two weeks to collect the data. Responses thereafter were not collected for analysis.

Overall, 269 (12.6%) experts have taken part and finished the survey. Roughly 27 % of the participants indicated to be a woman and 71% to be a man. 25% indicate living in the US. The vast majority of experts are White (86%), 7% Asian, 3% Hispanic, 2% Middle Eastern, and 1% are Black. 82% of respondents specify to have a professorial position (assistant, associate, or full professor), and 97 % disclose to have published in a peer-reviewed journal. 93% of participants consider themselves to be labor economists and 57 % indicated to do research on discrimination.

By the end of the survey, we also asked participants to indicate how confident they were in their

⁵⁹We only contact NBER participants with a linked NBER account from where we obtain email addresses.

⁶⁰Of these, 23 Emails could not be delivered.

assessment. Only 5 participants, indicate to feel very or extremely confident in their assessment. The median expert indicated feeling slightly confident, with 29% being not at all confident in their estimate. Thus, one way to interpret this low confidence is that professional experts know that it is difficult to predict the results of academic studies. An alternative interpretation, however, could be that experts were genuinely unsure about the results.

Before discussing the results, it should be pointed out that the predictions of experts are extremely homogeneous. Specifically, we do not find any consistent heterogeneous differences between different groups of experts.⁶¹ This finding is striking as experts consistently predict the same behavior and essentially agree on all questions. It is also true that no group of experts does better in predicting than other subgroups. For example, male and female labor economists are extremely similar in their predictions, do not differ significantly in their prediction of any task, and also do not differ in terms of correctly predicting the results of our study. We also measured how often the prediction of experts falls within the bootstrapped 95% confidence interval of each actual result. The effect estimated by Experts was, on average, 6 times within the 95% confidence interval of the actual effect. No group of experts is significantly better at predicting our results. In particular, gender, race, experience in publishing, experience in discrimination research, etc., all do not mediate how well experts predict the results.

Experts' responses to the questions are depicted in Figure H.1.

Stage 1 – Overall Prediction Starting with Stage 1, we observe that experts clearly predict that White profiles will fare substantially better than Black profiles. In the first stage of the experiment, participants expect White profiles to have, on average, 18.4% more contacts relative to Black profiles. This number is relatively close to the actual gap of 13%. To have a better understanding of how experts predict some common demographics to explain discrimination, we asked them to predict the relative gap between White and Black profiles for multiple subgroups of users. In the question, we inform them of the actual gap across the entire sample, which is 13%. Again, experts were similar in their predictions.

Stage 1 – Age Prediction In terms of age, experts clearly predict boomers to behave most preferential towards White profiles, followed by Gen X, Gen Y, and finally Gen Zs. This decrease seems to be predicted almost linearly from 22.5% for boomers to 6.2% for Gen Zs. However, as we know from our results, this relationship is almost reversed in our data with Gen Zs and Gen Ys preferring a White profile relative to a Black profile at 16% compared to Gen Xs, who “only” have a 5% relative gap.

Stage 1 – Education Prediction Experts also predict that educational attainment is positively associated with less discrimination. Specifically, the average expert predicts users who have not attended college to prefer White profiles at a 16.9% relative rate, while they expect this relative acceptance gap to be only 9.1% for users who have attended college. These predictions are very close to our actual observation where users who have not attended college prefer White profiles at a 18% relative rate, and users who have attended college prefer White profiles at a 12% rate.

Stage 1 – Race Prediction Interestingly, experts predict the result of the first stage to be almost fully driven by non-Black users. Specifically, they expect Black users to treat White and Black profiles roughly the same. In fact, on average they even expect Black users to have a slight preference for Black profiles (-1% relative gap), while they expect non-Black users to highly prefer

⁶¹The only differences we find are professors who expect a slightly smaller gap with regard to education than non-professors, and experts who work on discrimination expect a smaller gap in discrimination as a function of the user's race.

White profiles with a 14.7% gap relative to Black profiles. These numbers are very similar to our actual results, at least for non-Black users, where we observe a relative gap of 14.7%. However, different from what experts predict, we observe that Black users also discriminate against Black profiles. They do seem to discriminate less, but not zero (the relative gap is 7.7%).

Stage 1 – Gender Prediction Another case where experts clearly predict a different result is the association of gender and discrimination. Experts predict men to discriminate substantially more than women. Specifically, they expect that male users have a relative gap of 15.6 % in favor of White profiles and they expected this gap to reduce to 10.3% ($p \leq 0.001$) for female users. In our data, however, the reverse is true. Men display a relative acceptance gap towards White profiles of 8% relative to Black profiles, which is significantly smaller than the predicted value of experts ($p \leq 0.001$). On the other hand, women display a relative acceptance gap towards White profiles of 20% relative to Black profiles, which is significantly higher than the predicted value of experts ($p \leq 0.001$). As before, we find this pattern for all groups of experts (i.e., women, men, professors, non-professors, etc.) and we do not see any group of experts predicting this gap correctly. In fact, only 17% of experts were correct with respect to the direction of the effect. In short: Experts predict men to discriminate substantially more than women – while we find the exact opposite in our data.

Stage 2 – Overall Prediction Finally, we wanted to understand how well experts predict the results of our second stage. Specifically, we wanted to understand whether experts correctly anticipate that, once a profile has access to a job network and all endogeneity is accounted for, there will be no discrimination against Black profiles. However, this seems not to be the case. 87% of all experts predict a higher response rate towards White profiles relative to Black profiles. On average, experts expect White profiles to receive 12.9% more message responses relative to Black profiles. This prediction is substantially different from what we actually observe, as the actual relative gap is 3.7% ($p \leq 0.001$). Once again this finding is very robust to a variety of heterogeneity analyses. Professors, experts who work on discrimination, men, and women, all predict that White profiles will receive substantially more responses than Black profiles after accounting for differences in networks originating from the first stage.

In summary, we see that experts do well in predicting some of our results. Their prediction of the relative gap between White and Black profiles is very close to the actual gap for the first stage. Experts are also correct about the effect of education on discrimination, and they are somewhat correct in the prediction of how race affects discrimination (even though they predict no discrimination of Black profiles by Black users, which is different than what we find). Strikingly, however, our experiment revealed multiple unexpected findings. First, experts predict discrimination to almost linearly increase in age – we, however, find that it is mostly the younger generations who discriminate more than the older generations (in particular Gen Y and Z vs. Gen X). Further, experts predict men to discriminate substantially more than women. As shown in the main part of the paper, we find the exact opposite: it is women who discriminate substantially more than men. Finally, they expect the effect of discrimination to continue to prevail during the second stage and after removing endogeneity in networks from the first stage. In particular, they expect White profiles to receive more responses than Black profiles. We, however, find that White profiles do not receive significantly more responses, and our actual relative gap is substantially and significantly smaller than predicted by experts.

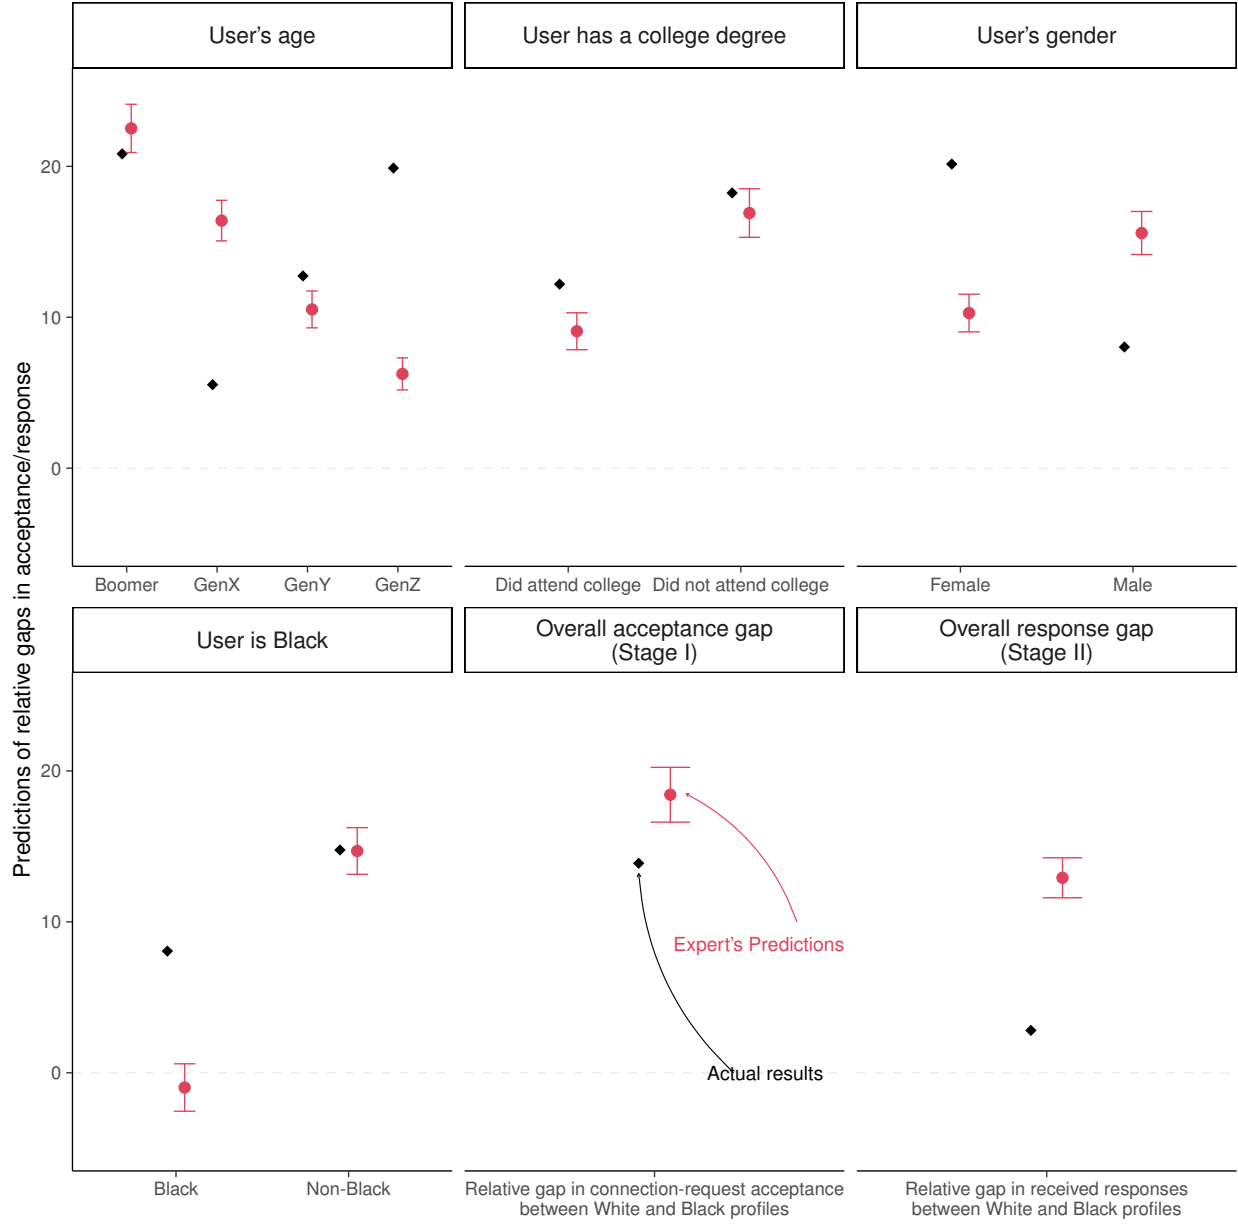


Figure H.1: Experts' predictions of discrimination on LinkedIn

The figure depicts experts' predictions and the actual discrimination in our setting. Red dots denote the average predictions of experts. Black diamonds denote the actual results of our paper. The Y-axis denotes the relative difference between White and Black profiles. The x-axis denotes the group of users whose relative acceptance rate is predicted by experts. Whiskers show the 95% CI. The CI for the relative gaps in our data is obtained from bootstrapping our sample 1000 times.

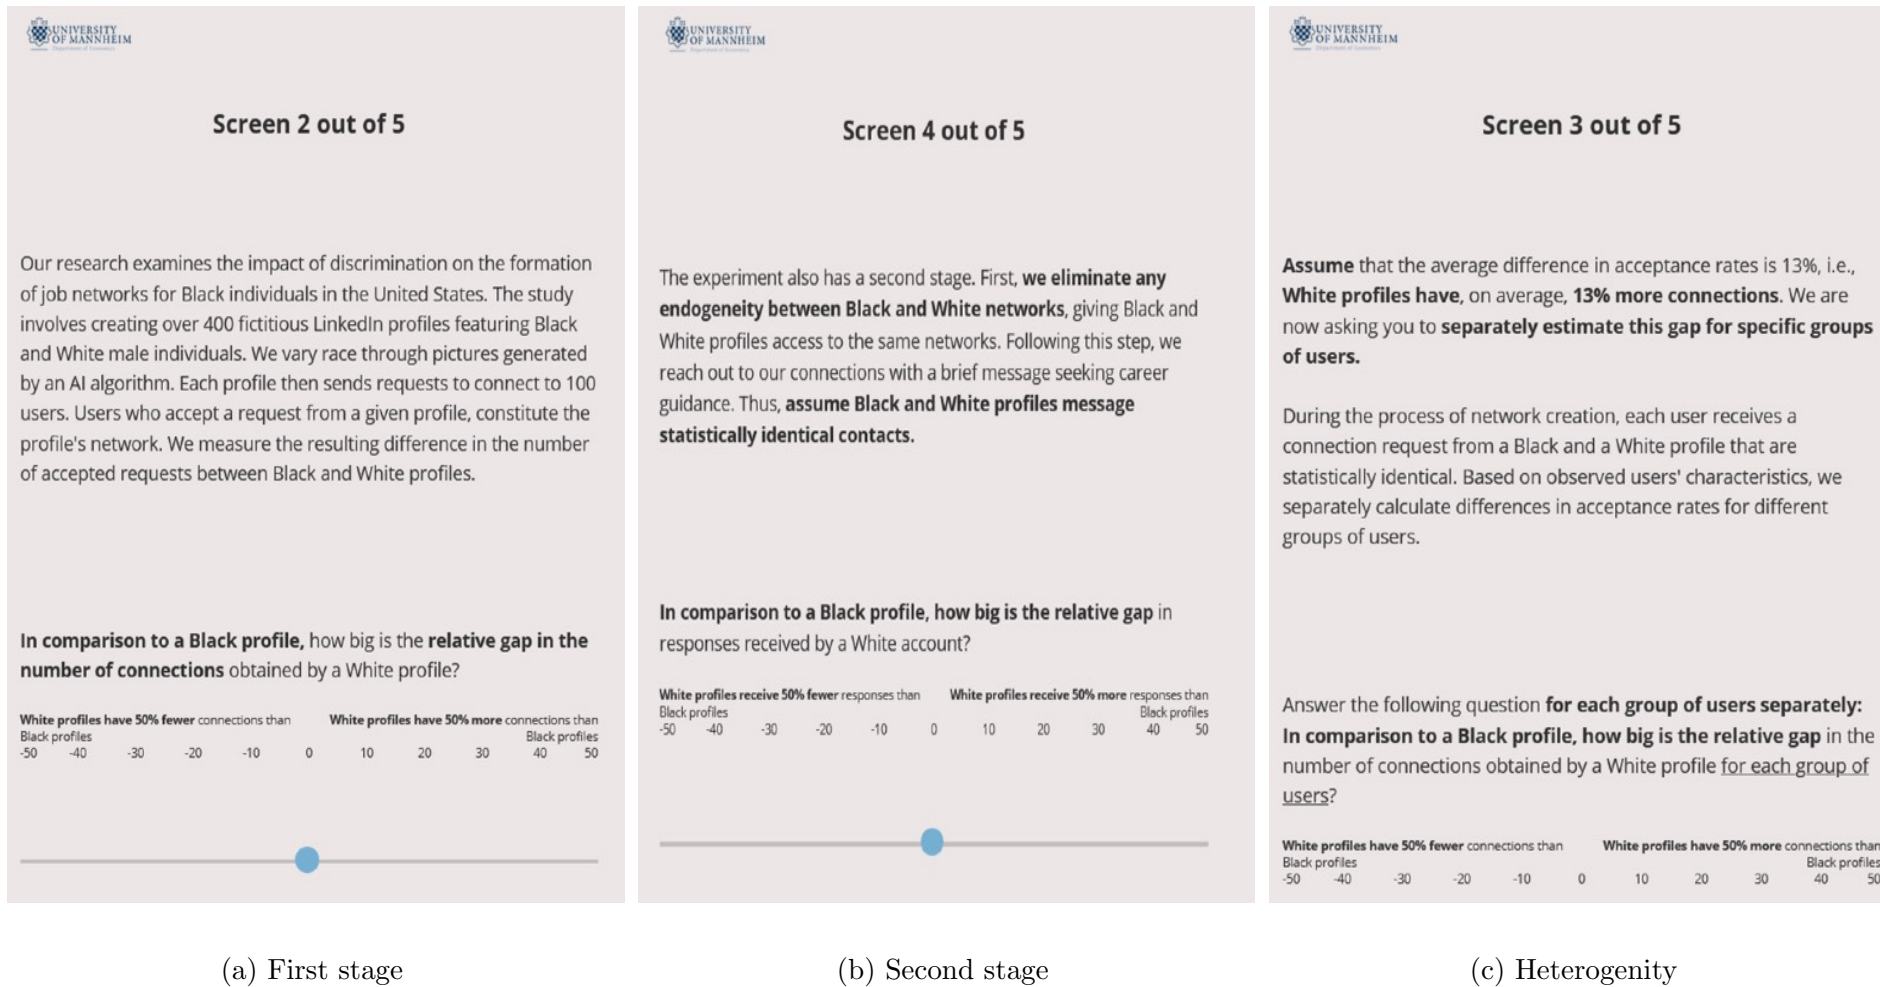


Figure H.2: Screenshots of the expert survey prediction tasks

I Survey of LinkedIn Users

To better understand the behavior on LinkedIn, the attitudes and motives of LinkedIn users, and to contrast some of our design choices with actual experience of users, we conducted a survey among LinkedIn users in early 2024. Aside from demographic questions, we ask participants a variety of questions concerning their experience, behavior, and observations on LinkedIn, some questions concerning knowledge with regard to the labor market gap between Black and White Americans, and their personal attitudes towards certain groups of people. As the survey is relatively comprehensive, individual questions are available from the authors upon request.

We created a survey on Qualtrics and distributed it via prolific. We restrict the convenience sample to employed US-based and fluent English-speaking users with at least 500 survey submissions and an approval rate of more than 95%. The median time participants required to finish the survey was about 27 minutes.

Overall, 500 random LinkedIn users have participated in and finished the survey. Approximately 60 % of the participants indicated to be a woman and 40% to be a man. The majority of respondents are White (66%), 9% are Asian, 5% Hispanic, and 20% are Black. The median respondent is 41 years old. 96 % of respondents reside in the US, and 22 % of respondents live in an urban area. 95 % of respondents specify to be employed (the remainder is either self-employed or studying), and they have an average of 192 connections on LinkedIn.

In what follows, we discuss the survey’s results in detail. Further, Table [I.1](#) shows the survey’s main results as cited in the paper’s main text.

	Main Question (Partially Shortened)	Sub-Options (Partially Shortened)	Scale	Outcome	Detail
Experience on LinkedIn	How useful, do you think, is your LinkedIn network for your career?		extr. useless - extr. useful	69% (extremely) useful	Page 82
	How LinkedIn changed how you search for jobs	more rely on LinkedIn connections for job-search assistance than non-LinkedIn ones LinkedIn social connections are/were useful to get jobs more likely to apply for/find jobs in other industries more likely to apply for/find jobs different from own job	strongly disagree - strongly agree	45% (strongly) agree 53% (strongly) agree 53% (strongly) agree 61% (strongly) agree	Fig. I.1
	If you were looking for a new position/job, how likely would you be to use LinkedIn	to reach out to employees at a firm you are interested in? Easy Apply feature to apply to an open position? to find a job? to find information about other employers?	very likely - very unlikely	60% (very) likely 70% (very) likely 80% (very) likely 82% (very) likely	Fig. I.2
	What is the importance of the following purposes in your use of LinkedIn?	job searching and career opportunities networking with professionals increasing your visibility to potential employers finding skills employers are looking for recruiting or hiring purposes knowledge sharing and professional discussions building a personal brand staying connected with friends look at what (old) friends and schoolmates are doing entertainment dating	not important (0) - very important (100)	68.61 (1.3) 67.46 (1.31) 63.39 (1.41) 55.33 (1.41) 49.93 (1.58) 48.01 (1.46) 40.34 (1.5) 31.07 (1.36) 27.79 (1.31) 11.94 (0.96) 5.6 (0.75)	Fig. I.3
	Questions on romantic advances: How often have you...	used LinkedIn to search for romantic partners? received romantic advances from own connections? received romantic advances from unconnected users? perceived actions or communications of others as uncomfortable romantic or sexual advances	never - always	99% never or rarely 93% never or rarely 91% never or rarely 91% never or rarely	Fig. I.4
	Do you consider LinkedIn mainly a professional networking website or a social media network?		professional networking or social media	92% professional networking	p. 84
	When considering whether to accept a connection request, how relevant are the following factors to your decision to accept?	The person might be helpful to you the person might post career-relevant information the person increases your network size you might be helpful to the person the person might post entertaining content the person might be a potential romantic partner	not relevant - extremely relevant	82 % v. or extr. relevant 71 % v. or extr. relevant 52 % v. or extr. relevant 52 % v. or extr. relevant 22 % v. or extr. relevant 13 % v. or extr. relevant	Fig. I.5
	At which stage of your career did you create your LinkedIn profile?		College/education; First job; Early Career; Mid Career; Established; Career	54% in early stage of career or college	p. 85
	When somebody tries to connect with you: How common is it that they send you a personalized message alongside the request?		always - never	55% sometimes or never	p. 85
Behavior on LinkedIn	How often do you send messages? ... send connection requests? ... receive messages from strangers? ... receive messages? ... use LinkedIn? ... receive connection requests from strangers? ... receive connection requests?	Never; Rarely; Once a month or less; A few times a month; 1-2 times a week; 3-4 times a week; once a day; multiple times a day	48 % > monthly 54 % > monthly 74 % > monthly 76 % > monthly 83 % > monthly 83 % > monthly 87 % > monthly	Fig. I.6
	How likely are you to accept a connection request from somebody who ...	might be able to provide you with career advice is a recruiter or in human resources attended the same educational institution as you have a large LinkedIn network you definitely don't know looks like a potential romantic partner	From extremely unlikely to extremely likely	71 % likely or extr. likely 68 % likely or extr. likely 61 % likely or extr. likely 55 % likely or extr. likely 23 % likely or extr. likely 17 % likely or extr. likely	Fig. I.7
	When you receive a connection request, how important are the following profile characteristics for your decision to accept or ignore the request?	industry skills job title education recent posts network size profile picture place of residence age gender ethnicity	Don't look at info; not important; slightly important moderately important; very important; extremely important	43 % v. or extr. important 39 % v. or extr. important 29 % v. or extr. important 23 % v. or extr. important 21 % v. or extr. important 18 % v. or extr. important 13 % v. or extr. important 13 % v. or extr. important 11 % v. or extr. important 9 % v. or extr. important 8 % v. or extr. important	Fig. I.8
	Name the biggest upsides (downsides) of accepting a request by somebody you do not know.	upsides downsides	Open Question	Top 3: Own network gets bigger, Potential for collaborations or networking with requesting user, Potential for future jobs Top 3: Potential for scams and fraud, Unwanted communication, Spam and unsolicited messages	Fig. I.9
	Name the most important features determining whether you will accept the connection request of somebody you do not know.		Open Question	Top 3: industry relevance, senders intentions, job title and experience	Fig. I.10
	How much time do you spend deciding whether to accept a connection request from a person you do not know?		Secs; <1 Min; 1-2 Min; 3-5 Min; 6-10 Min; >10 Min;	17 % Secs; 52% more than one minute	p. 90
	When considering whether to respond to job-related messages on LinkedIn, please indicate to what extent each of the following factors positively influences your decision to respond	you know personally is in your network might be able to help you in the future is a recruiter or in human resources attended the same educational institution as you lives/works closely might benefit from your advice has a large LinkedIn network has a nice profile picture has the same ethnicity as you looks like a potential romantic partner has the same gender as you	From not relevant (1) to extremely relevant (5)	88 % v. or extr. relevant 80 % v. or extr. relevant 74 % v. or extr. relevant 63 % v. or extr. relevant 47 % v. or extr. relevant 42 % v. or extr. relevant 41 % v. or extr. relevant 39 % v. or extr. relevant 23 % v. or extr. relevant 13 % v. or extr. relevant 12 % v. or extr. relevant 11 % v. or extr. relevant	Fig. I.11
	Name the most important features determining whether you will respond to a message by somebody you do not know.		Open Question	Top 3: message content and purpose, personal benefit or gain or value proposition, relevance to professional interests	Fig. I.13
	Which features of a profile are considered red flags?		Open Question	Top 3: Suspicious CV or Profile, Suspicious or Missing Photo, Suspicious Posts or Content	Fig. I.14
Perceptions	How likely is the following person to be helpful in your job search through personal interactions on LinkedIn?	White person Black person	extremely unlikely (-2) - extremely likely (2)	0.81 (0.04) 0.62 (0.04)	Fig. I.12
	What do you think the following person will earn in 5 years (in 1,000\$)	White person with BA and current income of \$55k Black person with BA and current income of \$55k	\$0 - \$ 200k	90.038 (1.411) 78.660 (1.260)	
	What is the current US unemployment rate for	White population Black population	0% - 100 %	13.89 (0.93) 19.80 (1.02)	
	What is the future US unemployment rate for	White person Black person	0% - 100 %	14.80 (0.99) 19.98 (1.08)	

Table I.1: Overview over key questions in LinkedIn User survey, as cited in main text.

Note: The table above includes answers to the main questions from the LinkedIn user survey, as cited in the paper's main text. Some answers and questions were slightly adjusted from their original version to fit into the table. Their meaning is preserved. For original questions and answers, see Chapter I or follow the links in Column 'Detail'.

Usefulness of LinkedIn To understand how useful LinkedIn users consider their network for their career, we asked “In general, how useful, do you think, is your LinkedIn network for your career?”. 69% of users respond that they consider LinkedIn useful (48%) or extremely useful (21%) for their career, while only 17% of users consider LinkedIn useless (12%) or extremely useless (5%). We also asked respondents whether LinkedIn connections are useful for acquiring jobs (see top right panel of Figure I.1), whether respondents are more likely to rely on LinkedIn connections for job search assistance rather than non-LinkedIn sources (see top left panel of Figure I.1), and whether LinkedIn has been helpful to find out about different jobs (see bottom right panel of Figure I.1) and jobs from different industries (see bottom left panel of Figure I.1). 53% of respondents agree or strongly agree that LinkedIn connections are useful to acquire jobs, 45% of respondents agree or strongly agree that they would rather rely on LinkedIn connections than other connections, while only 28% disagree or strongly disagree, and 61%, and 53% agree or strongly agree that LinkedIn has been helpful to find out about different jobs and jobs from different industries, respectively.




Figure I.1: Whether LinkedIn has changed how you search for jobs

In addition, we asked respondents how likely they are to use LinkedIn to 1) find a job (see top left panel of Figure I.2), 2) find information about other employers (see top right panel of Figure I.2), 3) reach out to employees at a firm they are interested in (see bottom left panel of Figure I.2), and 4) to apply to an open position via LinkedIn’s Easy Apply feature (see bottom right panel of Figure I.2). The (vast) majority of respondents indicate to be likely or very likely to use LinkedIn for all four purposes.

Reason for Using LinkedIn To better understand the incentives to accepting connections, we asked why people use LinkedIn. Specifically, we asked them how important several potential purposes of LinkedIn are to their usage of LinkedIn. These purposes range from mainly social (like dating, and looking at what (old) friends and schoolmates are doing) to professional (like finding out which skills and experiences employers are looking for, increasing visibility to potential employers,




Figure I.2: If you were looking for a new position/job, how likely would you be to use

and searching for jobs). Figure I.3 reports the importance of these purposes. Respondents rank all professional reasons higher than any social reason. Job searching, networking, increasing one's visibility to potential employers, and finding out which skills employers are looking for are ranked highest. Dating, on the other hand, is ranked lowest.




Figure I.3: What is the importance of the following purposes in your use of LinkedIn?
Whiskers around the mean denote the corresponding 95% confidence intervals.

Next, we zoom into how likely LinkedIn users are to use LinkedIn to search for romantic partners, have received romantic advances on LinkedIn, and had situations where they perceived actions or communications of others as uncomfortable romantic or sexual advances (see Figure I.4). 99% of users indicate to rarely or never use LinkedIn to search for romantic partners, more than 90% of users rarely or never received romantic advances from users on LinkedIn, and 91% of users rarely or never experienced harassment on LinkedIn. All of these results show that dating preferences are unlikely to be a driving force for behavior on LinkedIn.

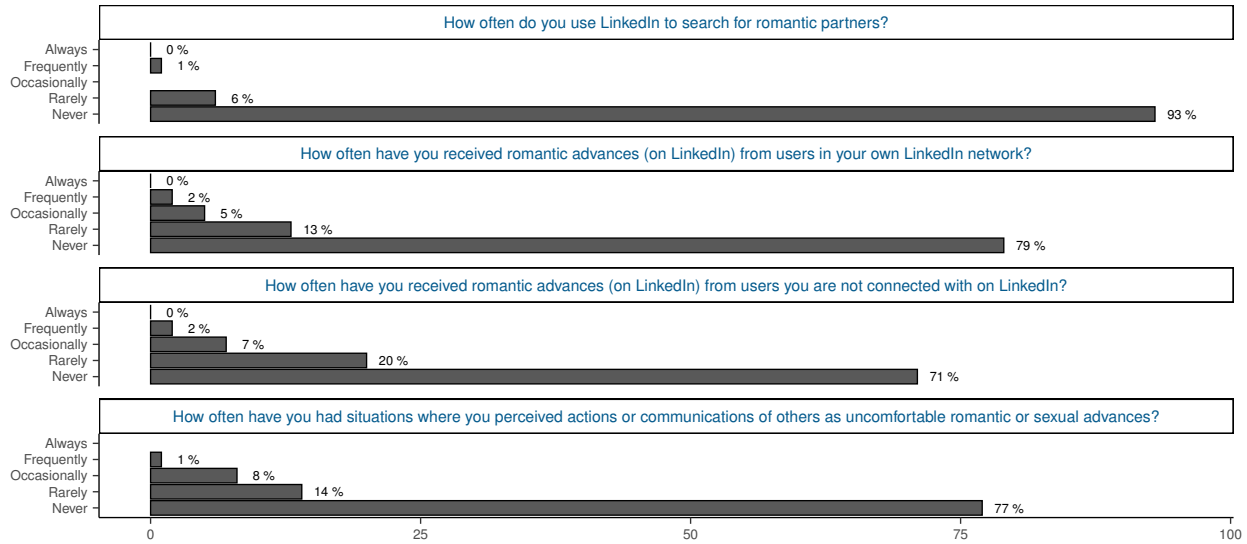


Figure I.4: Questions on romantic advances

LinkedIn: A Professional and not Social Network To put our results in perspective and understand how representative our results are for professional networks, we ask LinkedIn users “Do you consider LinkedIn mainly a professional networking website (i.e., the main focus of interactions is job-related) or a social media network (i.e., the main focus is to consume and post content, and interact with friends, etc.)?”. The vast majority (92 %) of users indicated to consider LinkedIn a professional website, which is in line with the reasons for using LinkedIn (see Figure I.3) and the factors LinkedIn users consider when thinking to accept a connection request (see Figure I.5, showing that professional factors are very important and social factors are not important). Overall, LinkedIn is clearly considered a professional and not a social network.




Figure I.5: When considering whether to accept a connection request, how relevant are the following factors to your decision to accept?

Behavior on LinkedIn One of the motivations for conducting the survey among LinkedIn users is to observe common behavior and to validate some of our design choices. First, we ask when users typically open their LinkedIn account to see whether it is congruent with our profiles joining LinkedIn shortly after finishing college and starting their first job. We ask: “At which stage of your career did you create your LinkedIn profile?”. 55% of respondents indicate to have created their account in the early stages of their career (23 % in college, 13 % in their first job, and 18 % in early stages of their career), while only 21 % indicate to have created it once they had an established career (which is mostly driven by users older than 50 years).

Similarly, to validate whether sending a message along with a connection request is uncommon. We ask participants: “In case somebody tries to connect with you: How common is it that they send you a personalized message along with the connection request?”. The majority indicates that this only happens sometimes (46 %) or never (9 %), with a minority stating it happens most of the time (19 %) or always (3 %). This suggests that not sending a message along with a connection request is common and would unlikely be considered odd.

In addition, we wanted to understand 1) how often users engage on LinkedIn, send connection requests and messages, and 2) how common it is to receive connection requests or messages from people they do not know. Figure I.6 displays the responses. While users seem to send connection requests and messages relatively rarely, it was not unusual to receive connection requests and messages from people they did not know. Specifically, 63 % of respondents indicate to receive connection requests from people they do not know at least a few times a month or more, and more than half of respondents indicate to receive messages from people they do not know at least a few times a month or more. This suggests that reaching out is not considered unusual or fake (which

is also supported by the relatively high acceptance and response rate in the experiment).




Figure I.6: How often do you ...

To get at features predicting acceptance rates, we also asked users "How likely are you to accept a connection request from somebody who ...". We again see that professional features like being able to provide career advice, being in HR, and having a large network strongly influence the probability of acceptance (see Figure I.7). Having attended the same university also increases the acceptance rate, while being a romantic partner is largely irrelevant. Interestingly, only 51% of participants indicate to be unlikely or extremely unlikely to accept someone who they definitely don't know, suggesting that many users are willing to accept someone they don't know.




Figure I.7: How likely are you to accept a connection request from somebody who ...

Further, we wanted to understand what factors are relevant for users to accept a connection request. In terms of factors relevant for accepting a connection request (see Figure I.8), we can see that skills and industry are the most relevant, while gender and ethnicity are the least relevant.




Figure I.8: When you receive a connection request, how important are the following profile characteristics for your decision to accept or ignore the request?

Before asking users the questions in Figure I.8, we posted open-ended survey questions on the key up- and downsides when accepting a request of someone one does not know. Starting with upsides, users post a variety of reasons, which are categorized in Figure I.9. Aside from simply increasing one's network, users primarily care about direct networking with the requesting user,

the potential of the connection to lead to a future job, and information they may gain through the user, such as information on job openings, industry news, or learning opportunities. With regard to downsides, users are primarily concerned with scams, spam, security, and privacy, as well as unwanted or low-quality communication.⁶²




Figure I.9: Open-ended survey question: Name the biggest upsides (downsides) of accepting a request by somebody you do not know.

In another open-ended survey question, we also asked users what they consider the most important feature when considering whether to accept a connection request from an unknown user. Users primarily care about professional reasons, including the requester’s professional relevance and status, as shown in Figure I.10. They also explicitly mention that they consider whether they expect mutual benefits. In addition, mutual connections or intentions of the sender play a role.

⁶²We omitted “other responses” from the graphs on open-ended survey questions.




Figure I.10: Open-ended survey question: Name the most important features determining whether you will accept the connection request of somebody you do not know.

To study how long users need to make a decision on whether to accept a connection request from a person they do not know, we ask participants: “How much time do you spend deciding whether to accept a connection request from a person you do not know?”. The majority indicate that they need more than one minute to decide (23 % need 1-2 minutes, 15 % need 3-5 minutes, 6 % need 6-10 minutes, 9 % need more than 10 minutes), 30 % indicate to need less than one minute, and only 18 % indicate to need only seconds to make the decision. This highlights that a majority of users likely don’t make very fast decisions when receiving a connection request.

When focusing on factors relevant to answering job-related messages on LinkedIn (see Figure I.11), we find that gender, and ethnicity are the least relevant factors, while being in the network and knowing the person personally are the most important ones.




Figure I.11: When considering whether to respond to job-related messages on LinkedIn, please indicate to what extent each of the following factors positively influences your decision to respond

To understand whether LinkedIn users anticipate relevant differences between Black and White profiles, we ask them to predict several outcomes of Black and White men (see Figure I.12). LinkedIn users anticipate that Black men will be less helpful in terms of job search, will have a lower income in 5 years, and are more likely to be unemployed currently and in the future.




Figure I.12: Perceived differences between Black and White Americans

Note: The question posed to participants in the top right reads: “Think of a young person who just finished a Bachelor’s in Business Administration and started working at a large firm. In their new position, they earn \$55,000 a year. What do you think they will earn in 5 years (in 1,000\$)?”

Regarding messages, we, again, post an open-ended survey question before the above questions as shown in Figure I.13. Similar to connection requests, it asks users what they consider the most important features when considering whether to respond to a message from an unknown user. The most frequently mentioned features refer to the message’s content, purpose, or relevance and to how the user herself can benefit from responding. In addition, users take the sender’s profile, tone, and authenticity into account.




Figure I.13: Open-ended survey question: Name the most important features determining whether you will respond to a message by somebody you do not know.

Red Flags of LinkedIn profiles To understand what makes a profile realistic or suspicious on LinkedIn, we asked users in an open-ended survey question what they consider to be red flags in LinkedIn profiles. Users were allowed to provide multiple responses. We summarize the 847 red flags users provided into seven categories, as shown in Figure I.14. The predominant red flags mentioned refer to suspicious CV or profile entries, profile pictures, or content/posts.⁶³ Much less frequently, users mention poor grammar or spelling followed by few network connections and having too many connections.⁶⁴

⁶³Suspicious CV entries include: Too little information or incomplete profile/CV; too much information in profile/CV; exaggerated, inconsistent, or suspicious work history; exaggerated, inconsistent, or suspicious education history; and lack of recommendations.

Suspicious posts or content includes: Lack of engagement or activity; unprofessional or inappropriate content or excessive posting; and overly personal or unprofessional language.

⁶⁴47 users did not provide a valid answer.




Figure I.14: Open-ended survey question: Which features of a profile are considered red flags?

J Tables

J1 Regressions

To test how discrimination affects behavior, we use a set of standard regressions. We can run the regression on two levels: the profile level and the target level.

The standard regression we are using is:

$$\begin{aligned}
 Outcome_{i,j} = & \alpha_0 + \beta \cdot BlackProfile + \gamma \cdot Variable + \delta \cdot Variable \cdot BlackProfile + \\
 & \sum_{k=1}^K \omega_k X_k^{State} + \sum_{l=1}^L \eta_l X_l^{Job} + \sum_{m=1}^M \lambda_m X_m^{Firstname} + \\
 & \sum_{q=1}^Q \phi_q X_q^{Lastname} + \sum_{t=1}^T \rho_t X_t^{Picture} + \\
 & \epsilon_i + \epsilon_j + \epsilon_{i,j}
 \end{aligned} \tag{1}$$

BlackProfile denotes a dummy with value one if the profile picture (in the current stage) depicts a Black person, and zero otherwise. *Variable* denotes a variable of interest, most often “attended worse Uni”. In many regressions, we do not estimate an interaction effect, i.e., $\gamma = \delta = 0$. X_k denote possible control variables. X_k^{State} , X_l^{Job} , $X_m^{Firstname}$, and $X_q^{Lastname}$ denote fixed effects for the state, the job title, the first name, and the last name of the profile, respectively. $X_t^{Picture}$ denote the fixed effects for picture-specific characteristics like how fake, trustworthy, intelligent, authentic, and good-looking the profile is considered, as well as how old and how likely the person on the picture is a woman or Asian.⁶⁵

$Outcome_{i,j}$ denotes the behavior of target i towards profile j with regard to an outcome. The most common outcomes are: a dummy indicating whether a connection request has been accepted, a dummy indicating whether a message was answered, the length of the normalized response (in

⁶⁵We do not control for how Black and how White the picture is considered as this is highly correlated with the race of the profile.

characters), the probability of the response being highly valuable (i.e., replies that offer a referral or a meeting), and the usefulness of the message. In case we run the regression on the profile level, we first aggregate $Outcome_{i,j}$ to $Outcome_j$ on the profile level. If, for example, we focus on the probability of responding, we would aggregate the number of positive responses on the profile level to run the corresponding regressions on the profile level.

ϵ_i and ϵ_j are target and profile picture random effects with $(\epsilon_i \sim \mathcal{N}(0, \sigma_1^2), \epsilon_j \sim \mathcal{N}(0, \sigma_2^2))$, which account for the fact that each profile reaches out to multiple people and for the fact that each target is contacted twice, allowing us to control for target-specific acceptance rates (in the first stage of the experiment). Note that in the second stage of the experiment, we do not account for the latter, as each target receives a message only once (i.e., $\epsilon_i = 0$). Also note, that regressions on the profile level do not account for target-specific effects as they are already aggregated on the profile level (and given the random assignment of targets to profiles there is also no need to account for selection, etc.).

$\epsilon_{i,j} \sim \mathcal{N}(0, \sigma_3^2)$ denotes the residual.

J2 Main experiment – First Stage

J2.1 Aggregate Results

Panel A: Aggregate difference in number of contacts							
	Number of Contacts						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Profile is Black	-3.06*** (0.47)	-3.05*** (0.47)	-3.06*** (0.47)	-3.13*** (0.46)	-3.06*** (0.49)	-2.97*** (0.55)	-3.07*** (0.54)
Constant	26.13*** (0.44)	39.28*** (1.91)	25.96*** (1.06)	26.02*** (0.87)	24.91*** (1.15)	22.70*** (5.06)	35.66*** (5.14)
State Controls	×	✓	×	×	×	×	✓
Job Controls	×	×	✓	×	×	×	✓
Firstname Controls	×	×	×	✓	×	×	✓
Lastname Controls	×	×	×	×	✓	×	✓
Picture trait Controls	×	×	×	×	×	✓	✓
Picture specific random effects	✓	✓	✓	✓	✓	✓	✓
Observations	400	400	400	400	400	400	400

Panel B: Differences in number of contacts accounting for profile quality							
	Number of Contacts						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Profile is Black	-3.26*** (0.68)	-3.27*** (0.67)	-3.26*** (0.68)	-3.54*** (0.67)	-3.24*** (0.70)	-3.12*** (0.74)	-3.49*** (0.73)
Profile attended worse Uni	0.16 (0.88)	0.16 (0.72)	0.17 (0.89)	-0.08 (0.89)	0.02 (0.90)	0.27 (0.89)	-0.09 (0.76)
Profile is Black and attended worse Uni	0.40 (0.95)	0.42 (0.95)	0.40 (0.95)	0.79 (0.94)	0.36 (0.98)	0.28 (0.96)	0.80 (0.96)
Constant	26.06*** (0.63)	39.20*** (1.95)	25.84*** (1.17)	26.04*** (0.97)	24.90*** (1.23)	22.28*** (5.14)	35.25*** (5.22)
State Controls	×	✓	×	×	×	×	✓
Job Controls	×	×	✓	×	×	×	✓
Firstname Controls	×	×	×	✓	×	×	✓
Lastname Controls	×	×	×	×	✓	×	✓
Picture trait Controls	×	×	×	×	×	✓	✓
Picture specific random effects	✓	✓	✓	✓	✓	✓	✓
Observations	400	400	400	400	400	400	400

Notes: p<0.10;*p<0.05;**p<0.01;***p<0.001.

Table J.1: Number of contacts by race and education of profiles

The table estimates the number of contacts a profile has by the end of Stage I as a function of their race. Panel A focuses only on race, while Panel B additionally reports the interaction between profile quality and race. The regressions are conducted on the profile level, use various controls, and all follow Equation 1.

J2.2 Dynamic Effects

Panel A: Aggregate difference in number of contacts over time							
	Number of Contacts over time						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Profile is Black	-0.51** (0.19)	-0.51** (0.19)	-0.51** (0.19)	-0.57** (0.19)	-0.47* (0.19)	-0.49* (0.21)	-0.56** (0.21)
Week	3.24*** (0.03)	3.24*** (0.03)	3.24*** (0.03)	3.24*** (0.03)	3.24*** (0.03)	3.24*** (0.03)	3.24*** (0.03)
Profile is Black x Week	-0.37*** (0.05)	-0.37*** (0.05)	-0.37*** (0.05)	-0.37*** (0.05)	-0.37*** (0.05)	-0.37*** (0.05)	-0.37*** (0.04)
Constant	2.06*** (0.24)	9.14*** (1.12)	1.89*** (0.56)	1.94*** (0.30)	0.66 (0.36)	0.20 (1.48)	4.62* (1.96)
State Controls	×	✓	×	×	×	×	✓
Job Controls	×	×	✓	×	×	×	✓
Firstname Controls	×	×	×	✓	×	×	✓
Lastname Controls	×	×	×	×	✓	×	✓
Picture trait Controls	×	×	×	×	×	×	✓
Picture specific random effects	✓	✓	✓	✓	✓	✓	✓
Observations	3,200	3,200	3,200	3,200	3,200	3,200	3,200
Panel B: Differences in number of contacts accounting for profile quality over time							
	Number of Contacts over time						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Profile is Black	-0.84** (0.28)	-0.84** (0.28)	-0.84** (0.28)	-1.06*** (0.27)	-0.76** (0.28)	-0.78** (0.29)	-0.94*** (0.28)
Profile attended worse Uni	-0.36 (0.47)	-0.36 (0.40)	-0.34 (0.48)	-0.51 (0.47)	-0.44 (0.47)	-0.37 (0.47)	-0.57 (0.41)
Week	3.18*** (0.05)	3.18*** (0.05)	3.18*** (0.05)	3.18*** (0.05)	3.18*** (0.05)	3.18*** (0.05)	3.18*** (0.05)
Profile is Black and attended worse Uni	0.66 (0.39)	0.67 (0.39)	0.66 (0.39)	0.97* (0.38)	0.58 (0.39)	0.57 (0.39)	0.74 (0.38)
Profile is Black x Week	-0.31*** (0.07)	-0.31*** (0.07)	-0.31*** (0.07)	-0.31*** (0.06)	-0.31*** (0.07)	-0.31*** (0.07)	-0.31*** (0.06)
Profile attended worse Uni x Week	0.13 (0.07)	0.13 (0.07)	0.13 (0.07)	0.13* (0.06)	0.13* (0.06)	0.13* (0.07)	0.13* (0.06)
Profile is Black x attended worse Uni x Week	-0.12 (0.09)	-0.12 (0.09)	-0.12 (0.09)	-0.12 (0.09)	-0.12 (0.09)	-0.12 (0.09)	-0.12 (0.09)
Constant	2.24*** (0.33)	9.32*** (1.14)	2.04** (0.62)	2.18*** (0.38)	0.88* (0.42)	0.35 (1.51)	4.84* (1.98)
State Controls	×	✓	×	×	×	×	✓
Job Controls	×	×	✓	×	×	×	✓
Firstname Controls	×	×	×	✓	×	×	✓
Lastname Controls	×	×	×	×	✓	×	✓
Picture trait Controls	×	×	×	×	×	✓	✓
Picture specific random effects	✓	✓	✓	✓	✓	✓	✓
Observations	3,200	3,200	3,200	3,200	3,200	3,200	3,200
Notes: p<0.10;*p<0.05;**p<0.01;***p<0.001.							

Table J.2: Number of contacts over time by race and education of profiles

The table estimates the number of contacts a profile has over time as a function of their race. Panel A focuses on how the total number of contacts changes over time as a function of the profile's race. Panel B additionally reports how the profile quality interacts with the dynamic effect. The regressions are conducted on the profile level, use various controls, and all follow Equation 1. The number of observations results from 8 weeks of requests for each of the 400 profiles.

Panel A: Weekly Change in the number of contacts over time

	Weekly Change in the number of contacts						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Profile is Black	-0.47*** (0.11)	-0.47*** (0.11)	-0.47*** (0.11)	-0.48*** (0.11)	-0.47*** (0.11)	-0.48*** (0.11)	-0.48*** (0.11)
Week	0.11*** (0.02)	0.11*** (0.02)	0.11*** (0.02)	0.11*** (0.02)	0.11*** (0.02)	0.11*** (0.02)	0.11*** (0.02)
Profile is Black x Week	0.03 (0.03)	0.03 (0.03)	0.03 (0.03)	0.03 (0.03)	0.03 (0.03)	0.03 (0.03)	0.03 (0.03)
Constant	2.72*** (0.08)	4.21*** (0.24)	2.71*** (0.14)	2.73*** (0.12)	2.57*** (0.15)	2.43*** (0.62)	3.87*** (0.64)
State Controls	×	✓	×	×	×	×	✓
Job Controls	×	×	✓	×	×	×	✓
Firstname Controls	×	×	×	✓	×	×	✓
Lastname Controls	×	×	×	×	✓	×	✓
Picture trait Controls	×	×	×	×	×	✓	✓
Picture specific random effects	✓	✓	✓	✓	✓	✓	✓
Observations	3,200	3,200	3,200	3,200	3,200	3,200	3,200

Panel B: Weekly Change in the number of contacts over time accounting for profile quality

	Weekly Change in the number of contacts						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Profile is Black	-0.50** (0.15)	-0.50*** (0.15)	-0.50** (0.15)	-0.53*** (0.15)	-0.50** (0.15)	-0.50** (0.16)	-0.54*** (0.16)
Profile attended worse Uni	0.004 (0.17)	0.002 (0.15)	0.01 (0.17)	-0.03 (0.17)	-0.02 (0.17)	0.02 (0.17)	-0.03 (0.16)
Week	0.11*** (0.03)	0.11*** (0.03)	0.11*** (0.03)	0.11*** (0.03)	0.11*** (0.03)	0.11*** (0.03)	0.11*** (0.03)
Profile is Black and attended worse Uni	0.06 (0.21)	0.06 (0.21)	0.06 (0.21)	0.11 (0.21)	0.06 (0.22)	0.04 (0.21)	0.11 (0.22)
Profile is Black x Week	0.03 (0.04)	0.03 (0.04)	0.03 (0.04)	0.03 (0.04)	0.03 (0.04)	0.03 (0.04)	0.03 (0.04)
Profile attended worse Uni x Week	0.01 (0.04)	0.01 (0.04)	0.01 (0.04)	0.01 (0.04)	0.01 (0.04)	0.01 (0.04)	0.01 (0.04)
Profile is Black x attended worse Uni x Week	-0.01 (0.05)	-0.01 (0.05)	-0.01 (0.05)	-0.01 (0.05)	-0.01 (0.05)	-0.01 (0.05)	-0.01 (0.05)
Constant	2.72*** (0.12)	4.21*** (0.26)	2.70*** (0.17)	2.74*** (0.15)	2.58*** (0.17)	2.38*** (0.63)	3.82*** (0.65)
State Controls	×	✓	×	×	×	×	✓
Job Controls	×	×	✓	×	×	×	✓
Firstname Controls	×	×	×	✓	×	×	✓
Lastname Controls	×	×	×	×	✓	×	✓
Picture trait Controls	×	×	×	×	×	✓	✓
Picture specific random effects	✓	✓	✓	✓	✓	✓	✓
Observations	3,200	3,200	3,200	3,200	3,200	3,200	3,200

Notes:

p<0.10;*p<0.05;**p<0.01;***p<0.001.

Table J.3: Weekly change in the number of contacts by race and education of profiles

The table estimates the weekly change of the number of contacts a profile as a function of their race (i.e., the relative change in the number of contacts). Panel A focuses on how the number of contacts changes per week over time as a function of the profile's race. Panel B additionally reports how the profile quality interacts with the dynamic effect. The regressions are conducted on the profile level, use various controls, and all follow Equation 1. The number of observations results from 8 weeks of requests for each of the 400 profiles.

J2.3 Geographical Variation

	Difference in the number of contacts										
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
Absolute Male	0.0000 (0.0000)										
Edu: Share Bachelor		0.0000 (0.0000)									
Absolute White			0.0000 (0.0000)								
Share White				-1.79 (2.45)							
Share African-American					6.26 (3.45)						
Share Democratic						-3.64 (2.95)					
GDP per Capita (current USD)							-0.0000 (0.0000)				
In Bible Belt								1.46* (0.70)			
In Rust Belt									-1.12 (0.96)		
In Mormon Belt										0.22 (1.10)	
In Black Belt											2.35** (0.77)
Constant	2.89*** (0.47)	2.93*** (0.47)	2.82*** (0.53)	4.23* (1.70)	2.33*** (0.51)	4.79** (1.47)	4.62*** (1.00)	2.44*** (0.44)	3.19*** (0.38)	2.99*** (0.38)	2.46*** (0.37)
Observations	51	51	51	51	51	51	51	51	51	51	51
Notes:	p<0.10;*p<0.05;**p<0.01;***p<0.001.										

Table J.4: Regression estimates on the difference in the number of contacts between White and Black profiles by state characteristics

The table estimates the difference in the number of contacts between White and Black profiles for every US state + DC (i.e., N=51). The difference is calculated on the twin level and then aggregated to the state level. Each regression simply estimates the linear relationship between state-level differences in the number of connections between White and Black profiles and certain state characteristics. *Absolute Male* denotes the absolute number of men in the state. *Edu: Share Bachelor* denotes the share of people with a bachelor degree in a state. *Absolute White* denotes the absolute number of White people in the state. *Share White/Share African-American* denotes the relative number of White/Black people in the state. *Share Democratic* denotes the share of Democrats in the state. *GDP per Capita (current USD)* denotes state-level GDP. *In Bible Belt*, *In Rust Belt*, *In Mormon Belt*, and *In Black Belt* denote dummy variables with value one if the state is in the corresponding category. The Bible Belt is known for its religious states. The Mormon Corridor is characterized by a high proportion of Mormons. The Rust Belt consists of old-industrial states, and the Black Belt is historically associated with black slavery in the southern US before the Civil War.⁶⁶

J2.4 Differences in Networks

	Network characteristics of		
	Black profiles (N=4477)	White profiles (N=5046)	p-value
USER DEMOGRAPHICS			
<i>Female (First Name)</i>	0.50 (0.50)	0.52 (0.50)	0.019*
<i>Black (Last Name)</i>	0.06 (0.24)	0.06 (0.23)	0.344
<i>White (Last Name)</i>	0.71 (0.45)	0.71 (0.45)	0.92
<i>Asian (Last Name)</i>	0.07 (0.26)	0.08 (0.26)	0.632
<i>Hispanic (Last Name)</i>	0.15 (0.36)	0.15 (0.35)	0.869
<i>Age</i>	32.91 (9.97)	32.67 (9.97)	0.254
EMPLOYMENT AND PLATFORM USE			
<i>Salary</i>	84324.20 (54751.97)	82242.41 (53257.01)	0.061.
<i>High Job Position</i>	0.14 (0.35)	0.14 (0.35)	0.921
<i>Works in HR</i>	0.10 (0.30)	0.10 (0.29)	0.749
<i>Number of Contacts</i>	319.88 (185.70)	311.97 (185.89)	0.039*
<i>Number of Skills</i>	20.92 (13.77)	20.58 (13.54)	0.24
<i>Number of Skill Verifications</i>	37.36 (57.41)	37.57 (146.00)	0.93
<i>Number of Posts</i>	0.55 (0.50)	0.53 (0.50)	0.078.
<i>Has Volunteering Experience</i>	0.21 (0.40)	0.20 (0.40)	0.215
<i>Gender Pronouns Shown</i>	0.17 (0.37)	0.16 (0.36)	0.227
<i>Profile picture is happy</i>	0.81 (0.39)	0.82 (0.38)	0.377
<i>Follows a philanthropist</i>	0.03 (0.18)	0.03 (0.18)	0.775
EMPLOYER			
<i>Employees</i>	4739.38 (4529.35)	4699.56 (4521.37)	0.686
<i>Employees on Platform</i>	25714.65 (69553.80)	25781.80 (70590.17)	0.965
<i>Open Jobs on Platform</i>	1903.24 (6516.28)	1939.82 (7209.71)	0.806
HIGHER EDUCATION			
<i>None</i>	0.15 (0.36)	0.16 (0.37)	0.274
<i>Some College</i>	0.12 (0.32)	0.12 (0.33)	0.806
<i>Associate</i>	0.04 (0.20)	0.05 (0.21)	0.666
<i>Bachelor</i>	0.45 (0.50)	0.44 (0.50)	0.375
<i>Master</i>	0.21 (0.41)	0.21 (0.40)	0.645
<i>PhD</i>	0.03 (0.17)	0.03 (0.17)	0.713
<i>Undergrads: White</i>	0.63 (0.19)	0.63 (0.19)	0.695
<i>Undergrads: Black</i>	0.10 (0.12)	0.10 (0.12)	0.968
COUNTY			
<i>Share Democrat (2020)</i>	0.60 (0.15)	0.60 (0.15)	0.312
<i>Share White</i>	0.58 (0.19)	0.58 (0.19)	0.926
<i>Share Black</i>	0.17 (0.15)	0.17 (0.15)	0.993
<i>Dissimilarity Index (Black/White)</i>	54.73 (11.62)	54.61 (11.78)	0.636

Table J.5: Differences in resulting networks (Black vs. White)

The table reports upon the differences in the resulting networks between White and Black profiles. Each row represents a certain feature of the connected users. T-tests are used to obtain the following significance levels: ·p<0.10;*p<0.05;**p<0.01;***p<0.001.

⁶⁶States in the Bible belt are: Alabama, Arkansas, Georgia, Kentucky, Louisiana, Mississippi, Missouri, North Carolina, Oklahoma, South Carolina, Tennessee, Texas, Virginia, West Virginia, Florida, Illinois, Indiana, Kansas, New Mexico, Ohio. States in the Rust belt are: Michigan, Wisconsin, Indiana, Illinois, Ohio, Pennsylvania, West Virginia, and Kentucky. States in the Mormon Corridor are: Arizona, California, Idaho, Nevada, Utah, and Wyoming. States in the Black belt are: Alabama, Arkansas, Florida, Georgia, Louisiana, Maryland, Mississippi, North Carolina, South Carolina, Tennessee, Texas, and Virginia.

J2.5 Individual Predictors of Discrimination

	Connection request acceptance probability (in %)		
	(1)	(2)	(3)
Profile is Black	-3.39*** (0.32)	-3.44*** (0.33)	-3.50*** (0.36)
Z-Scored Age	-4.46*** (0.34)	-5.02*** (0.34)	-4.92*** (0.38)
Profile is Black x Z-Scored Age	1.18*** (0.32)	1.18*** (0.32)	1.22*** (0.35)
Constant	28.23*** (0.42)	23.52*** (0.65)	24.58*** (1.05)
Picture random effects	✓	✓	✓
Target random effects	✓	✓	✓
Nbr of contacts	×	✓	✓
Other fixed effects	×	×	✓
Observations	33,446	32,928	26,246
Notes:	p<0.10;*p<0.05;**p<0.01;***p<0.001.		

Table J.6: Age as a driver of discrimination

The table estimates the decision to accept a profile as a function of the profile's race and the user's (z-scored) age. The regressions are conducted on the target level, use various controls, and all follow Equation 1. Other fixed effects include: the number of contacts the target has; a dummy on whether the target has a bachelor's degree (or more); a dummy variable on whether the target has a senior position, is a CEO or a director; a dummy on whether the target is Black; a dummy on whether the target is a woman; the share of Republicans in the target's county. The number of observations denotes the number of acceptance decisions made by users. Columns controlling for specific user characteristics, reduce the sample to users without missing entries in those characteristics.

	Connection request acceptance probability (in %)								
	All users			White users			Black users		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Profile is Black	-1.90*** (0.43)	-1.88*** (0.44)	-1.91*** (0.49)	-2.21*** (0.52)	-2.23*** (0.53)	-2.35*** (0.60)	-1.91 (2.31)	-2.26 (2.37)	-1.68 (2.67)
Female	-0.22 (0.63)	1.20 (0.65)	0.24 (0.73)	-1.35 (0.78)	0.36 (0.80)	-0.27 (0.91)	-4.59 (3.19)	-2.52 (3.30)	-2.62 (3.67)
Profile is Black x Female	-2.35*** (0.59)	-2.52*** (0.60)	-2.83*** (0.68)	-1.96** (0.73)	-2.10** (0.75)	-2.35** (0.85)	-0.11 (3.04)	-0.08 (3.14)	-1.17 (3.57)
Constant	26.44*** (0.52)	20.44*** (0.71)	27.66*** (1.04)	26.55*** (0.61)	19.33*** (0.86)	26.67*** (1.31)	31.53*** (2.43)	25.64*** (3.31)	34.92*** (5.03)
Picture random effects	✓	✓	✓	✓	✓	✓	✓	✓	✓
Target random effects	✓	✓	✓	✓	✓	✓	✓	✓	✓
Nbr of contacts	×	✓	✓	×	✓	✓	×	✓	✓
Other fixed effects	×	×	✓	×	×	✓	×	×	✓
All subjects	✓	✓	✓	×	×	×	×	×	×
Only White subjects	×	×	×	✓	✓	✓	×	×	×
Only Non-white subjects	×	×	×	×	×	×	✓	✓	✓
Observations	36,911	35,794	29,973	23,792	23,130	19,183	1,617	1,550	1,285

Notes:

p<0.10;*p<0.05;**p<0.01;***p<0.001.

Table J.7: Gender as a driver of discrimination

The table estimates the decision to accept a profile as a function of the profile's race and the user's gender. The regressions are conducted on the target level, use various controls, and all follow Equation 1. Other fixed effects include: the number of contacts the target has; a dummy on whether the target has a bachelor's degree (or more); a dummy variable on whether the target has a senior position, is a CEO or a director; a dummy on whether the target is Black; a dummy variable indicating whether the target has above-median age; the share of Republicans in the target's county. The first three columns focus on the whole sample, while Columns (4)-(6), and Columns (7)-(9) restrict the sample to White and Black targets. The number of observations denotes the number of acceptance decisions made by users. Columns controlling for specific user characteristics, reduce the sample to users without missing entries in those characteristics.

	Connection request acceptance probability (in %)											
	Race: increasing prob of being non-white						Race: increasing prob of being black					
	Male users		Female users		All users		Male users		Female users		All users	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Profile is Black	-2.38*** (0.67)	-2.64*** (0.77)	-4.37*** (0.61)	-4.98*** (0.72)	-4.37*** (0.63)	-4.96*** (0.74)	-2.37*** (0.57)	-2.40*** (0.65)	-4.60*** (0.50)	-5.30*** (0.59)	-4.60*** (0.52)	-5.29*** (0.61)
Prop(Target \neq White)	-1.61 (1.46)	-1.87 (1.67)	2.09 (1.33)	1.36 (1.56)	2.11 (1.34)	1.22 (1.56)						
Profile is Black x Prop(Target \neq White)	0.74 (1.38)	1.55 (1.60)	0.65 (1.19)	0.62 (1.41)	0.62 (1.24)	0.55 (1.46)						
Prop(Target = Black)							5.05 (3.09)	5.19 (3.47)	1.14 (2.76)	2.88 (3.24)	0.68 (2.79)	2.71 (3.25)
Profile is Black x Prop(Target = Black)							2.31 (2.94)	2.60 (3.37)	4.18 (2.48)	4.84 (2.95)	4.10 (2.58)	4.71 (3.06)
Male					1.83 (0.97)	0.85 (1.12)					0.08 (0.81)	-0.42 (0.94)
Profile is Black x Male					2.01* (0.90)	2.31* (1.05)					2.26** (0.75)	2.89** (0.88)
Male x Prop(Target \neq White)					-3.51 (1.95)	-2.85 (2.25)						
Profile is Black x Male x Prop(Target \neq White)					0.15 (1.81)	1.04 (2.12)						
Male x Prop(Target = Black)											4.31 (4.12)	2.31 (4.72)
Profile is Black x Male x Prop(Target = Black)											-1.87 (3.82)	-2.04 (4.46)
Constant	26.91*** (0.76)	27.71*** (1.62)	24.97*** (0.71)	27.58*** (1.45)	24.98*** (0.73)	27.31*** (1.19)	25.78*** (0.66)	26.29*** (1.51)	25.62*** (0.59)	27.86*** (1.34)	25.70*** (0.62)	27.43*** (1.08)
Picture random effects	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Target random effects	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Other fixed effects	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓
All subjects	×	×	×	×	✓	✓	×	×	×	×	✓	✓
Only Women	×	×	✓	✓	×	×	×	×	✓	✓	×	×
Only Men	✓	✓	×	×	×	×	✓	✓	×	×	×	×
Observations	15,771	12,897	17,068	13,349	32,839	26,246	15,771	12,897	17,068	13,349	32,839	26,246

Notes:

p<0.10;*p<0.05;**p<0.01;***p<0.001.

Table J.8: Gender and race (White vs. non-White targets) as a driver of discrimination

The table estimates the decision to accept a profile as a function of the profile's race, the user's gender, and the user's race. The regressions are conducted on the target level, use various controls, and all follow Equation 1. Other fixed effects include: the number of contacts the target has; a dummy on whether the target has a bachelor's degree (or more); a dummy variable on whether the target has a senior position, is a CEO or a director; a dummy variable indicating whether the target has above-median age; the share of Republicans in the target's county. The first six columns estimate the target's race with a continuous variable indicating how likely the person is not White. Columns (7)-(12) estimate the target's race with a continuous variable indicating how likely the person is Black. Columns (1), (2), (7), and (8) restrict the sample to male targets only. Columns (3), (4), (9), and (10) restrict the sample to female targets only. Columns (5), (6), (11), and (12) use the whole sample and interact the target's race with their gender. The number of observations denotes the number of acceptance decisions made by users. Columns controlling for specific user characteristics, reduce the sample to users without missing entries in those characteristics.

Connection request acceptance probability (in %) by user's race				
	(1)	(2)	(3)	(4)
Profile is Black	-3.25*** (0.45)	-3.62*** (0.52)	-3.52*** (0.37)	-3.89*** (0.43)
Prop(Target ≠ White)	0.83 (0.96)	0.23 (1.11)		
Profile is Black x Prop(Target ≠ White)	0.35 (0.88)	0.50 (1.03)		
Prop(Target = Black)			2.86 (2.02)	4.19 (2.32)
Profile is Black x Prop(Target = Black)			3.56 (1.87)	4.06 (2.18)
Constant	25.82*** (0.55)	27.62*** (1.08)	25.80*** (0.47)	27.22*** (1.00)
Picture random effects	✓	✓	✓	✓
Target random effects	✓	✓	✓	✓
Other fixed effects	×	✓	×	✓
Observations	33,861	27,159	33,861	27,159

Notes: p<0.10;*p<0.05;**p<0.01;***p<0.001.

Table J.9: Differences in discrimination based on user's race

The table estimates the decision to accept a profile as a function of the profile's race and the user's race. The regressions are conducted on the target level, use various controls, and all follow Equation 1. Other fixed effects include: the number of contacts the target has; a dummy on whether the target has a bachelor's degree (or more); a dummy variable on whether the target has a senior position, is a CEO or a director; a dummy variable indicating whether the target has above-median age; a dummy on whether the target is a woman; the share of Republicans in the target's county. The first two columns estimate the target's race with a continuous variable indicating how likely the person is not White, while the last two columns estimate the target's race with a continuous variable indicating how likely the person is Black. The number of observations denotes the number of acceptance decisions made by users. Columns controlling for specific user characteristics, reduce the sample to users without missing entries in those characteristics.

Connection request acceptance probability (in %) by user's characteristic																		
	Is Old			Is female			Is Black			High Education			High Job			Republican		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)	(16)	(17)	(18)
Profile is Black	-4.68*** (0.46)	-4.97*** (0.53)	-4.97*** (0.53)	-1.90*** (0.43)	-2.10*** (0.52)	-2.10*** (0.52)	-3.22*** (0.32)	-3.51*** (0.37)	-3.51*** (0.37)	-3.66*** (0.49)	-4.37*** (0.68)	-4.38*** (0.68)	-3.27*** (0.32)	-3.53*** (0.40)	-3.54*** (0.40)	-2.44*** (0.42)	-2.60*** (0.51)	-2.61*** (0.51)
User Characteristic	-7.58*** (0.68)	-7.98*** (0.77)	-8.88*** (0.78)	-0.22 (0.63)	0.33 (0.77)	0.13 (0.78)	2.53 (1.51)	2.29 (1.79)	2.61 (1.78)	3.22*** (0.65)	-1.88* (0.86)	-3.74*** (0.88)	-3.33*** (0.85)	-3.69*** (1.03)	-4.28*** (1.05)	1.19 (0.69)	2.08* (0.81)	3.18*** (0.82)
Profile is Black x User Characteristic	2.53*** (0.64)	2.89*** (0.72)	2.89*** (0.72)	-2.35*** (0.59)	-2.65*** (0.72)	-2.65*** (0.72)	1.89 (1.40)	1.25 (1.68)	1.24 (1.68)	0.77 (0.61)	1.30 (0.80)	1.30 (0.80)	0.70 (0.79)	0.51 (0.97)	0.52 (0.97)	-1.41* (0.60)	-1.67* (0.72)	-1.67* (0.72)
Constant	32.09*** (0.54)	32.62*** (0.61)	28.89*** (1.11)	26.44*** (0.52)	28.25*** (0.60)	27.44*** (1.11)	26.00*** (0.42)	28.31*** (0.47)	28.16*** (1.10)	24.36*** (0.58)	29.77*** (0.77)	28.59*** (1.13)	26.97*** (0.42)	29.04*** (0.49)	26.38*** (1.00)	25.81*** (0.54)	27.38*** (0.61)	27.70*** (1.11)
Picture random effects	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Target random effects	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Other fixed effects	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×
Observations	33,446	26,246	26,246	36,911	26,246	26,246	33,861	26,246	26,246	38,299	26,246	26,246	38,299	26,246	26,246	26,246	26,246	26,246
Notes:	p<0.10; *p<0.05; **p<0.01; ***p<0.001																	

Notes: p<0.10;*p<0.05;**p<0.01;***p<0.001.

Table J.10: Drivers of discrimination

The table estimates the decision to accept a profile as a function of the profile's race and user's characteristics. The regressions are conducted on the target level, and all follow Equation 1. Each column denotes one regression estimating the difference in connection request acceptance between users who do and do not have the specific characteristic. For example, in Column 1 the user characteristic is being older than the median target. Columns 4-18 are defined as follows: target... (4-6) is a woman, (7-9) is Black, (10-12) has at least a bachelor's degree, (13-15) has a job title that includes "CEO", "director", or "senior", (16-18) lives in a county with an above-median Republican vote share. The second column of each characteristic restricts the sample to all users where we can determine all six characteristics (26,246 observations). The last column of each characteristic controls for the other five characteristics as a fixed effect (i.e., column 3 controls for gender, race, education, high job, and Republican).

Connection request acceptance probability (in %) by user's characteristic						
	Is Old (1)	Is female (2)	Is Black (3)	High Education (4)	High Job (5)	Republican (6)
Black profile \times User characteristic = True (Black-White acceptance rate gap for users with characteristic)	-2.08*** (0.50)	-4.75*** (0.51)	-2.26 (1.64)	-3.08*** (0.43)	-3.02*** (0.88)	-4.28*** (0.51)
Black profile \times User characteristic = False (Black-White acceptance rate gap for users without characteristic)	-4.97*** (0.53)	-2.10*** (0.52)	-3.51*** (0.37)	-4.37*** (0.68)	-3.53*** (0.40)	-2.60*** (0.51)
User characteristics = True (Difference in acceptance rate for users with characteristic)	-7.98*** (0.77)	0.33 (0.77)	2.29 (1.79)	-1.88* (0.86)	-3.69*** (1.03)	2.08* (0.81)
Constant (Baseline acceptance rate)	32.62*** (0.61)	28.25*** (0.60)	28.31*** (0.47)	29.77*** (0.77)	29.04*** (0.49)	27.38*** (0.61)
Differences in Gaps						
Difference in Gaps (Difference in acceptance rate gap between users with and without characteristic)	2.89*** (0.72)	-2.65*** (0.72)	1.25 (1.68)	1.30 (0.80)	0.51 (0.97)	-1.67* (0.72)
Picture random effects	✓	✓	✓	✓	✓	✓
Target random effects	✓	✓	✓	✓	✓	✓
Restricted sample	✓	✓	✓	✓	✓	✓
Observations	26,246	26,246	26,246	26,246	26,246	26,246
Notes:				p<0.10;*p<0.05;**p<0.01;***p<0.001.		

Table J.11: Drivers of discrimination

The table estimates the decision to accept a connection as a function of the requesting profile's race and user characteristics. The regressions are conducted at the target level and all follow Equation 1. Each column denotes one regression that estimates the difference in the acceptance rate (gap) for users with and without the reported characteristic. For example, in Column 1 the user characteristic is being older than the median target. Columns 2-6 are defined as follows: target (2) is female, (3) is Black, (4) has at least a bachelor's degree, (5) has a job title that includes "CEO", "director", or "senior", (6) lives in a county with an above-median Republican vote share. The table restricts the sample to all users for whom we have data on all six characteristics (26,246 observations). The first part of the table (rows 1-4) shows regressions that separately compute the acceptance rate gap between Black and White profiles for users with and without the respective characteristic (rows 1-2). We also report the baseline acceptance rate and the difference in the baseline rate for users with the respective characteristic (rows 3-4). The second part of the table (row 5) reports the difference between the coefficients in rows 1 and 2, i.e., the difference in acceptance rate gaps between users with and without the respective characteristic. The result stems from a separate regression with a simple interaction effect between the profile being Black and the user characteristic being true.

Connection request acceptance probability (in %) by user's characteristic			
	Share of Black students (1)	Share of Black citizens (2)	Segregation (3)
Profile is Black	-2.62*** (0.50)	-2.83*** (0.42)	-3.52*** (0.44)
User Characteristic	2.38** (0.77)	2.05** (0.70)	-0.30 (0.69)
Profile is Black x User Characteristic	-0.65 (0.71)	-0.56 (0.59)	0.78 (0.60)
Constant	26.52*** (0.59)	25.33*** (0.54)	26.53*** (0.55)
Picture random effects	✓	✓	✓
Target random effects	✓	✓	✓
Observations	27,496	36,761	36,424
Notes: p<0.10;*p<0.05;**p<0.01;***p<0.001.			

Table J.12: Exposure as a driver of discrimination

The table estimates the decision to accept a profile as a function of the profile's race and user's characteristic. The regressions are conducted on the target level, and all follow Equation 1. Each column denotes one regression estimating the difference in connection request acceptance between users who do and do not have the specific characteristic. For example, Column 1 denotes a target who attended a university with above-median many Black students. Column 2 denotes a target who lives in a county with above median many Black Americans, and Column 3 denotes a target who lives in a county with above median levels of segregation. The number of observations denotes the number of acceptance decisions made by users. As each column controls for different user characteristics (and thus restricts the sample to users without missing entries in those characteristics), the sample changes between columns.

Connection request acceptance probability (in %)			
	(1)	(2)	(3)
Profile is Black	-3.55*** (0.29)	-3.64*** (0.30)	-3.82*** (0.37)
Z-Scored Number of Profile's Connections	-1.90*** (0.25)	-1.95*** (0.26)	-1.76*** (0.31)
Profile is Black x Z-Scored Number of Profile's Connections	-0.33 (0.43)	-0.30 (0.44)	-0.19 (0.52)
Constant	26.67*** (0.44)	21.40*** (0.61)	28.34*** (1.11)
Picture random effects	✓	✓	✓
Target random effects	✓	✓	✓
Nbr of contacts	×	✓	✓
Other fixed effects	×	×	✓
Observations	38,142	37,010	26,150
Notes: p<0.10;*p<0.05;**p<0.01;***p<0.001.			

Table J.13: Profile's number of connections as a driver of discrimination

The table estimates the decision to accept a profile as a function of the profile's race and the profile's number of connections at the time of request. The regressions are conducted on the target level, use various controls, and all follow Equation 1. Other fixed effects include: the number of contacts the target has; a dummy on whether the target has a bachelor's degree (or more); a dummy variable on whether the target has a senior position, is a CEO or a director; a dummy variable indicating whether the target has above-median age; a dummy on whether the target is Black; a dummy on whether the target is a woman; the share of Republicans in the target's county. The number of observations denotes the number of acceptance decisions made by users. Columns controlling for specific user characteristics, reduce the sample to users without missing entries in those characteristics.

J2.6 Ancillary Outcomes

	#Views		# unsolicited messages		# unsolicited contact requests	
	(1)	(2)	(3)	(4)	(5)	(6)
Profile is Black	-5.72*** (0.51)	-4.30*** (0.48)	-0.09* (0.04)	-0.06 (0.04)	0.01 (0.10)	0.04 (0.10)
Nbr. of Contacts		3.00*** (0.32)		0.08** (0.03)		0.06 (0.07)
Constant	41.38*** (0.48)	40.68*** (0.45)	0.31*** (0.04)	0.29*** (0.04)	0.96*** (0.09)	0.94*** (0.09)
Observations	400	400	400	400	400	400
<i>Notes:</i> p<0.10;*p<0.05;**p<0.01;***p<0.001.						

Table J.14: Additional outcomes by race of the profiles

The table estimates several additional profile-level outcomes as a function of the profile's race. *# Views* denote how often the profile has been viewed in the last six weeks. *# unsolicited messages* denotes the number of unsolicited messages, and *# unsolicited contact requests* denotes the number of unsolicited contact requests. The regressions are conducted on the profile level and all follow Equation 1.

J3 Main Experiment – Second Stage

J3.1 Message Summary Statistics

Statistic	N	Median	Mean	St. Dev.	Min	Max
Nbr. messages answered	400	2	1.70	1.18	0	6
Response rate	400	.20	.21	.14	0	.67
Nbr. words	338	49.35	58.37	39.85	3	300
Nbr. characters	338	282	330.35	227.52	13	1,805
Friendliness	338	3.50	3.50	.48	2.25	5
Mentioned referral	338	0	.03	.12	0	1
Mentioned reference to other	338	0	.05	.16	0	1
Offered meeting	338	0	.05	.16	0	1
Shared experience	338	.17	.23	.27	0	1
Shared materials	338	0	.06	.18	0	1
Shared information	338	.25	.30	.33	0	1
Generic advise	338	.50	.44	.35	0	1
Mere response	338	0	.27	.35	0	1
Offers to keep in touch	338	0	.15	.28	0	1
Usefulness (GPT-4)	338	7	6.24	1.99	0	10

Table J.15: Summary statistics of the responses received on profile level

The table reports basic summary statistics of the responses received on profile level. As some profiles receive zero responses, we have only 338 profiles for the summary statistics following row three.

Usefulness Rating by ChatGPT			
Length of message (in words)	1.19*** (0.08)		
Referral or Meeting		1.04** (0.38)	
RAs usefulness rating			1.57*** (0.07)
Constant	6.22*** (0.08)	6.15*** (0.10)	6.22*** (0.07)
Observations	664	664	664
Notes:	p<0.10;*p<0.05;**p<0.01;***p<0.001.		

Table J.16: Validation of ChatGPT’s usefulness rating

Length of message (in words) denotes the z-scored length of the message. *Referral or Meeting* denotes a dummy with value one if the message offers a referral or a meeting. *RAs usefulness rating* denotes the z-scored average usefulness rating of the two RAs. The number of observations denotes all responses to messages (664).

J3.2 Message Responses

Response Rate (in %)				
	Network generated by Black profile		Network generated by White profile	
	Black profile	White profile	Black profile	White profile
Estimate	22.56*** (1.41)	20.68*** (1.60)	18.22*** (1.40)	21.45*** (1.27)
Observations	101	100	100	99
Network characteristics				
Fit	✓	×	×	✓
Composition	✓	✓	×	×
Discrimination	✓	×	✓	×
Notes:	p<0.10;*p<0.05;**p<0.01;***p<0.001.			

Table J.17: Effects of Stage II

The table estimates the response rate in Stage II by the race of the profile and the race of the profile generating the network. Networks generated by Black profiles have the feature “*composition*”, indicating potential differences in member characteristics accepting Black profiles. *Discrimination* denotes a dummy with value one if the profile picture (in the current stage) depicts a Black person, and zero otherwise. Finally, the feature “*fit*” denotes that the race of the profile writing the message is the same as the race of the profile generating the network. The number of observations denotes the number of profiles in each category. For example, we have 100 profiles who had a Black profile picture in Stage I, and have a White profile picture in Stage II.

Panel A: Aggregate difference in messages (response rate, length, and usefulness)

	Response Rate (in %)					Message Length (in char)					Highly Useful Message?					Usefulness score				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)	(16)	(17)	(18)	(19)	(20)
Profile is Black	-0.66 (1.33)	-0.67 (1.34)	-0.64 (1.34)	-0.61 (1.35)	-0.38 (1.37)	-10.23 (7.00)	-9.90 (7.26)	-8.75 (7.19)	-10.47 (7.04)	-8.72 (7.66)	-0.01 (0.02)	-0.02 (0.02)	-0.02 (0.02)	-0.01 (0.02)	-0.02 (0.02)	-0.13 (0.22)	-0.11 (0.22)	-0.10 (0.22)	-0.11 (0.22)	-0.08 (0.23)
Constant	21.05*** (1.02)	36.93*** (5.82)	16.07*** (3.45)	20.12 (12.15)	27.23 (14.32)	83.21*** (4.96)	91.59*** (25.29)	76.25*** (18.38)	26.13 (58.19)	14.76 (73.87)	0.08*** (0.02)	0.13 (0.08)	0.01 (0.06)	0.06 (0.19)	0.05 (0.23)	6.31*** (0.15)	5.69*** (0.78)	6.25*** (0.57)	3.00 (1.79)	2.09 (2.19)
State Controls	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×
Job Controls	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×
Firstname Controls	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×
Lastname Controls	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×
Picture trait Controls	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×
Picture specific random effects	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Observations	400	400	400	400	400	338	338	338	338	338	338	338	338	338	338	338	338	338	338	338

Panel B: Differences in messages accounting for profile quality

	Response Rate (in %)					Message Length (in char)					Highly Useful Message?					Usefulness score				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)	(16)	(17)	(18)	(19)	(20)
Profile is Black	2.69 (1.87)	2.68 (1.88)	2.67 (1.88)	2.79 (1.91)	3.03 (1.93)	-11.75 (9.95)	-12.26 (10.35)	-10.37 (10.25)	-13.05 (10.11)	-13.19 (10.97)	-0.02 (0.03)	-0.03 (0.03)	-0.03 (0.03)	-0.02 (0.03)	-0.04 (0.03)	-0.36 (0.31)	-0.33 (0.31)	-0.36 (0.32)	-0.41 (0.31)	-0.38 (0.33)
Profile attended worse Uni	1.79 (2.02)	1.80 (2.04)	2.36 (2.06)	2.35 (2.09)	2.83 (2.15)	-9.80 (9.95)	-12.76 (10.31)	-6.73 (10.39)	-9.13 (10.35)	-10.69 (11.43)	-0.02 (0.03)	-0.04 (0.03)	-0.01 (0.03)	-0.03 (0.03)	-0.04 (0.03)	-0.26 (0.31)	-0.27 (0.32)	-0.28 (0.32)	-0.31 (0.32)	-0.35 (0.34)
Profile is Black and attended worse Uni	-6.63* (2.63)	-6.64* (2.64)	-6.63* (2.68)	-6.71* (2.69)	-6.79* (2.75)	2.11 (14.02)	3.40 (14.56)	2.70 (14.71)	4.29 (15.76)	8.21 (15.62)	0.03 (0.04)	0.03 (0.04)	0.02 (0.05)	0.03 (0.04)	0.04 (0.05)	0.46 (0.43)	0.43 (0.44)	0.51 (0.44)	0.59 (0.44)	0.61 (0.47)
Constant	20.14*** (1.44)	36.09*** (5.93)	15.07*** (3.55)	19.48 (12.30)	26.43 (14.50)	88.40*** (7.24)	98.76*** (25.90)	80.15*** (19.30)	38.55 (59.68)	30.68 (75.88)	0.09*** (0.02)	0.15 (0.08)	0.02 (0.06)	0.09 (0.19)	0.10 (0.23)	6.44*** (0.22)	5.83*** (0.80)	6.39*** (0.59)	3.29 (1.84)	2.47 (2.25)
State Controls	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×
Job Controls	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×
Firstname Controls	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×
Lastname Controls	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×
Picture trait Controls	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×
Picture specific random effects	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Observations	400	400	400	400	400	338	338	338	338	338	338	338	338	338	338	338	338	338	338	338

Panel C: Differences in messages accounting for message type

	Response Rate (in %)					Message Length (in char)					Highly Useful Message?					Usefulness score				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)	(16)	(17)	(18)	(19)	(20)
Profile is Black	-1.50 (1.99)	-1.49 (1.99)	-1.51 (2.01)	-1.47 (2.00)	-1.20 (2.02)	-51.99 (27.75)	-42.17 (28.23)	-47.90 (28.38)	-54.31 (27.97)	-43.79 (29.77)	-0.01 (0.03)	-0.004 (0.03)	-0.02 (0.03)	-0.001 (0.03)	-0.01 (0.03)	-0.34 (0.25)	-0.31 (0.26)	-0.33 (0.26)	-0.35 (0.25)	-0.31 (0.27)
Mentor message	-15.34*** (2.00)	-15.34*** (2.00)	-15.34*** (1.99)	-15.34*** (2.00)	-15.34*** (2.00)	-30.72 (30.48)	-27.44 (31.42)	-21.31 (30.84)	-32.31 (30.44)	-18.06 (31.62)	-0.01 (0.03)	-0.0000 (0.03)	-0.01 (0.03)	-0.01 (0.03)	-0.0003 (0.03)	0.04 (0.28)	0.004 (0.29)	0.13 (0.29)	-0.01 (0.28)	0.06 (0.30)
Profile is Black and Mentor message	1.54 (2.82)	1.54 (2.82)	1.54 (2.81)	1.54 (2.82)	1.54 (2.82)	81.66 (43.07)	82.12 (44.77)	71.07 (43.59)	85.01* (42.97)	75.11 (45.05)	0.01 (0.05)	-0.01 (0.05)	0.01 (0.05)	0.01 (0.05)	-0.01 (0.05)	0.64 (0.40)	0.63 (0.42)	0.51 (0.41)	0.69 (0.40)	0.56 (0.42)
Constant	28.17*** (1.44)	44.74*** (5.81)	22.75*** (3.70)	29.33* (12.34)	36.99* (14.49)	359.82*** (19.39)	272.38*** (81.35)	358.67*** (62.20)	116.36 (198.13)	7.23 (244.22)	0.08*** (0.02)	0.14 (0.07)	0.02 (0.06)	0.01 (0.20)	0.02 (0.23)	6.34*** (0.17)	6.26*** (0.68)	6.07*** (0.55)	3.69* (1.74)	2.46 (2.12)
State Controls	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×
Job Controls	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×
Firstname Controls	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×
Lastname Controls	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×
Picture trait Controls	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×
Picture specific random effects	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Profile specific random effects	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Observations	800	800	800	800	800	463	463	463	463	463	463	463	463	463	463	463	463	463	463	463

Notes:

p<0.10;*p<0.05;**p<0.01;***p<0.001.

Table J.18: Response rate and message characteristics

The table estimates several response characteristics as a function of profiles race in the second stage of the experiment. Panel A focuses only on race, Panel B additionally reports the interaction between profile quality and race, and Panel C reports the interaction between the type of message and race. As every profile sends two types of requests (mentor and application), we have double the sample size in Panel C. The regressions are conducted on the profile level, use various controls, and all follow Equation 1. The number of observations is 800 in Panel C (columns 1-5) as each profile has sent two types of requests. The number of observations reduces in some columns as some profiles received no responses.

Panel A: Aggregate difference in messages (response rate, length, and usefulness)

	Response Rate (in %)				Message Length (in char)				Highly Useful Message?				Usefulness score			
	Native	Alien			Native	Alien			Native	Alien			Native	Alien		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)	(16)
Profile is Black	1.12 (1.82)	0.50 (2.39)	-2.46 (1.95)	-1.88 (2.45)	1.15 (8.88)	1.89 (12.66)	-22.73* (10.85)	-16.07 (14.19)	-0.002 (0.03)	0.001 (0.04)	-0.02 (0.03)	-0.06 (0.04)	0.13 (0.31)	0.06 (0.43)	-0.41 (0.28)	-0.55 (0.34)
Constant	21.43*** (1.35)	31.54 (22.75)	20.68*** (1.51)	36.34 (25.40)	73.47*** (6.28)	91.59 (123.13)	93.94*** (7.75)	-155.01 (140.68)	0.08*** (0.02)	-0.10 (0.40)	0.07*** (0.02)	0.39 (0.38)	6.11*** (0.22)	0.38 (4.17)	6.53*** (0.21)	5.24 (3.74)
State Controls	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓
Job Controls	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓
Firstname Controls	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓
Lastname Controls	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓
Picture trait Controls	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓
Picture specific random effects	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Observations	200	200	200	200	176	176	162	162	176	176	162	162	176	176	162	162

Panel B: Differences in messages accounting for profile quality

	Response Rate (in %)				Message Length (in char)				Highly Useful Message?				Usefulness score			
	Native	Alien			Native	Alien			Native	Alien			Native	Alien		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)	(16)
Profile is Black	5.67* (2.51)	4.96 (3.20)	-0.31 (2.76)	0.30 (3.25)	5.28 (12.48)	2.64 (17.13)	-32.96* (15.69)	-32.64 (19.64)	-0.01 (0.04)	-0.04 (0.05)	-0.04 (0.05)	-0.08 (0.05)	-0.23 (0.43)	-0.14 (0.58)	-0.52 (0.41)	-0.96* (0.48)
Profile attended worse Uni	1.17 (2.65)	2.93 (3.16)	2.39 (3.02)	3.05 (3.48)	1.81 (12.62)	8.34 (16.92)	-25.60 (15.57)	-30.97 (20.04)	-0.01 (0.05)	-0.12* (0.05)	-0.04 (0.04)	-0.09 (0.05)	-0.50 (0.44)	-0.60 (0.57)	-0.04 (0.43)	-0.36 (0.53)
Profile is Black and attended worse Uni	-8.99* (3.53)	-8.20* (4.07)	-4.26 (3.89)	-4.38 (4.34)	-8.57 (17.85)	-2.82 (22.60)	16.87 (21.76)	29.34 (26.73)	0.01 (0.06)	0.09 (0.07)	0.04 (0.06)	0.03 (0.07)	0.72 (0.62)	0.46 (0.76)	0.21 (0.57)	0.88 (0.66)
Constant	20.84*** (1.88)	32.90 (23.42)	19.46*** (2.16)	32.59 (25.91)	72.57*** (8.92)	75.02 (128.00)	108.33*** (11.68)	-101.76 (144.48)	0.09** (0.03)	0.12 (0.40)	0.09** (0.03)	0.54 (0.39)	6.36*** (0.31)	1.45 (4.32)	6.55*** (0.33)	5.65 (3.82)
State Controls	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓
Job Controls	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓
Firstname Controls	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓
Lastname Controls	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓
Picture trait Controls	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓
Picture specific random effects	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Observations	200	200	200	200	176	176	162	162	176	176	162	162	176	176	162	162

Panel C: Differences in messages accounting for message type

	Response Rate (in %)				Message Length (in char)				Highly Useful Message?				Usefulness score			
	Native	Alien			Native	Alien			Native	Alien			Native	Alien		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)	(16)
Profile is Black	0.92 (2.73)	0.25 (3.17)	-3.95 (2.91)	-3.33 (3.28)	-36.98 (32.71)	-29.11 (42.02)	-72.09 (46.10)	-56.57 (59.84)	0.004 (0.04)	0.02 (0.05)	-0.02 (0.03)	-0.07 (0.04)	-0.06 (0.34)	-0.11 (0.45)	-0.69 (0.36)	-0.84 (0.45)
Mentor message	-15.39*** (2.74)	-15.39*** (2.78)	-15.28*** (2.91)	-15.28*** (2.93)	-23.20 (37.66)	5.62 (41.08)	-47.16 (49.15)	-42.98 (53.98)	0.02 (0.05)	0.02 (0.05)	-0.05 (0.04)	-0.07 (0.04)	0.26 (0.39)	0.35 (0.41)	-0.26 (0.41)	-0.30 (0.44)
Profile is Black and mentor message	1.05 (3.86)	1.05 (3.92)	2.03 (4.11)	2.03 (4.15)	66.12 (52.78)	22.78 (59.01)	106.35 (70.01)	140.95 (79.40)	0.02 (0.07)	-0.02 (0.08)	0.003 (0.06)	0.04 (0.06)	0.70 (0.55)	0.44 (0.59)	0.60 (0.58)	0.70 (0.65)
Constant	28.32*** (1.96)	44.56 (23.25)	28.02*** (2.09)	43.22 (25.37)	325.47*** (22.78)	35.16 (376.64)	401.98*** (32.50)	-260.85 (503.85)	0.08** (0.03)	-0.07 (0.44)	0.08*** (0.02)	0.34 (0.40)	6.08*** (0.24)	0.62 (4.11)	6.66*** (0.26)	6.15 (3.80)
State Controls	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓
Job Controls	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓
Firstname Controls	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓
Lastname Controls	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓
Picture trait Controls	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓
Picture specific random effects	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Profile specific random effects	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Observations	400	400	400	400	247	247	216	216	247	247	216	216	247	247	216	216

Notes:

p<0.10;*p<0.05;**p<0.01;***p<0.001.

Table J.19: Response rate and message characteristics by network

The table estimates several response characteristics as a function of profiles race in the second stage of the experiment and differentiates between native and alien networks. Panel A focuses only on race, Panel B additionally reports the interaction between profile quality and race, and Panel C reports the interaction between the type of message and race. As every profile sends two types of requests (mentor and application), we have double the sample size in Panel C. Columns (1), (2), (5), (6), (9), (10), (13), and (14) restrict the sample to all profiles who remain in their native network (i.e., their picture is not swapped), while Columns (3), (4), (7), (8), (11), (12), (15), and (16) restrict the sample to all profiles who move to an alien network (i.e., their picture is swapped). The regressions are conducted on the profile level, use various controls, and all follow Equation 1. The number of observations reduces in some columns as some profiles received no responses.

	Response Rate (in %)										Message Length (in char)										Highly Useful Message?										Usefulness score																			
	Black only					White only					Both					Not both					All targets					Black only					White only					Both					Not both					All targets				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)	(16)	(17)	(18)	(19)	(20)	(21)	(22)	(23)	(24)	(25)	(26)	(27)	(28)	(29)	(30)	(31)	(32)	(33)	(34)	(35)	(36)	(37)	(38)	(39)	(40)										
Profile is Black	2.59 (3.25)	3.73 (3.99)	-4.63 (2.56)	-3.65 (1.93)	0.72 (2.01)	0.84 (3.09)	2.58 (3.28)	3.91 (3.28)	0.72 (1.85)	1.04 (1.90)	-0.52** (0.17)	-1.64*** (0.47)	0.03 (0.23)	-0.23 (0.52)	0.01 (0.10)	0.01 (0.11)	-0.51* (0.21)	-0.26 (0.11)	0.005 (0.04)	0.006 (0.04)	0.02 (0.11)	-0.14 (0.05)	-0.04 (0.04)	0.08 (0.04)	0.01 (0.03)	-0.02 (0.03)	0.02 (0.05)	-0.05 (0.07)	0.01 (0.03)	-0.01 (0.03)	-0.56 (0.47)	-4.99*** (1.48)	-0.004 (0.49)	0.03 (1.05)	-0.01 (0.23)	0.10 (0.27)	-0.57 (0.49)	-0.87 (0.61)	-0.01 (0.24)	-0.05 (0.26)										
Stage I: Accepted Only Black											-3.85 (2.69)	-3.53 (2.77)									0.23 (0.15)	0.18 (0.16)																		0.21 (0.37)	0.25 (0.40)									
Stage I: Accepted Only White											-0.51 (2.92)	0.40 (3.33)		-4.33 (2.39)	-4.16 (2.45)						-0.27 (0.21)	0.13 (0.28)		-0.03 (0.14)	-0.06 (0.15)																-1.01* (0.48)	-0.53 (0.65)	-0.77* (0.33)	-0.92* (0.37)						
Profile is Black											1.87 (3.76)	1.87 (3.85)									-0.49* (0.21)	-0.49* (0.22)																			-0.58 (0.51)	-0.77 (0.55)								
Stage I: Accepted Only Black																																																		
Profile is Black																																																		
Stage I: Accepted Only White											-7.24 (4.05)	-8.87* (4.49)		-5.39 (3.34)	-5.06 (3.42)						0.55 (0.29)	0.23 (0.39)	0.02 (0.20)	0.10 (0.22)																		0.64 (0.68)	0.85 (0.92)	0.01 (0.49)	0.23 (0.53)					
Constant	18.12*** (2.34)	75.79 (41.36)	17.62*** (1.84)	-7.87 (30.36)	-7.87 (1.35)	21.96*** (19.34)	34.91 (2.22)	21.96*** (20.94)	31.89* (1.30)	0.26* (14.07)	0.13 (0.13)		-11.44 (9.41)	-0.02 (0.16)	2.00 (4.63)	0.003 (0.07)	-0.53 (1.07)	0.24 (0.16)	-1.72 (1.71)	0.002 (0.07)	-1.25 (0.87)	0.06 (0.04)	0.99 (0.40)	0.07*** (0.03)	-1.29 (0.03)	0.08*** (0.04)	0.23 (0.02)	0.06 (0.31)	0.19 (0.03)	0.08*** (0.04)	0.19 (0.03)	6.66*** (0.42)	-61.90* (0.22)	5.63*** (0.33)	-2.00 (0.35)	6.39*** (0.28)	6.27*** (0.30)	-3.06 (0.36)	6.39*** (0.39)	2.92 (0.17)	2.52 (0.24)									
State Controls	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x									
Job Controls	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x								
Firstname Controls	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x								
Lastname Controls	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x								
Picture trait Controls	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x								
Picture specific random effects	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x								
Profile specific random effects	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x								
Observations	596	596	829	829	1,855	1,855	1,425	1,425	3,280	3,280	116	116	126	126	414	414	242	242	656	656	116	116	126	126	414	414	242	242	656	656	116	116	126	126	414	414	242	242	656	656	656	656								
Notes	p<0.10* p<0.05** p<0.01*** p<0.001****																																																	

Notes:

p<0.10;*p<0.05;**p<0.01;***p<0.001.

Table J.20: Response rate and message characteristics by first-stage behavior

The table estimates several response characteristics as a function of profile race in the second stage of the experiment and differentiates between users who have accepted only the White, only the Black, or both profiles. The regressions are conducted on the target level, use various controls, and all follow Equation 1. The number of observations varies depending on the user's behavior in the first stage, the race of the profile, and the user's second-stage behavior (e.g., as not all requests are answered we have a reduced sample size in analyzing the replies).

	Response Rate (in %)										By quality			
	Low quality					High quality					(10)		(11)	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)		(10)	(11)	(12)	
Fit	0.74				0.79	4.41*		4.38*	4.40*				4.36*	
	(2.13)				(2.13)	(2.15)		(2.16)	(2.15)				(2.15)	
Composition		0.62		0.62		3.02		2.98		3.03		3.00		
		(1.74)		(1.70)		(2.01)		(2.01)		(1.89)		(1.87)		
Discrimination			-3.95*	-3.95*			2.67	2.67			2.69	2.69		
			(1.69)	(1.70)			(2.02)	(2.01)			(1.87)	(1.87)		
Profile attended worse Uni									0.29	-0.34	1.79	4.77		
									(2.14)	(2.03)	(2.02)	(2.84)		
Fit x														
Profile attended worse Uni									-3.66			-3.58		
									(3.03)			(3.03)		
Composition x														
Profile attended worse Uni										-2.41		-2.38		
										(2.66)		(2.63)		
Discrimination x														
Profile attended worse Uni											-6.63*	-6.64*		
											(2.63)	(2.63)		
Constant	19.59***	19.65***	21.94***	21.23***	19.30***	20.00***	20.15***	16.48***	19.30***	19.99***	20.14***	16.46***		
	(1.51)	(1.38)	(1.36)	(1.93)	(1.52)	(1.49)	(1.50)	(2.08)	(1.52)	(1.44)	(1.44)	(2.02)		
Picture specific random effects	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓		
Observations	202	202	202	202	198	198	198	198	400	400	400	400		

Notes:

*p<0.10;*p<0.05;**p<0.01;***p<0.001.

Table J.21: Decomposing Stage II effects by profile quality

The table estimates the response probability in Stage II (after swapping profile pictures). Columns 1 and 4 focus on all profiles, while columns 2 and 3 estimate the effects for low and high-quality profiles separately. *Fit* denotes a dummy with value one if the profile is in the original network, and zero if the profile is in an alien network. *Composition* denotes a dummy with value one if the picture is in a network built by a Black profile (i.e., has the composition of a Black network), and zero otherwise. *Discrimination* denotes a dummy with value one if the profile picture (in the current stage) depicts a White person, and zero otherwise. Positive values, therefore, indicate discrimination against Black profiles. *Profile attended worse Uni* denotes a dummy with value one if the profile indicates attendance at a worse university. The regressions are conducted on the profile level and follow the mixed effects models of Equation 1. To account for twin-profile-specific heterogeneity, we use a random effect on the twin-target level. 202 of the 400 remaining profiles were assigned a lower-ranked university, while 198 of the 400 remaining profiles were assigned a higher-ranked university.

J3.3 Consequences of Swapping the Profile Picture

	Views			
	(1)	(2)	(3)	(4)
Picture swapped	0.33 (0.83)	0.55 (0.61)	0.26 (0.83)	0.43 (0.62)
Weeks after swapping			0.99*** (0.04)	0.99*** (0.04)
Picture swapped x Weeks after swapping			0.11 (0.06)	0.11 (0.06)
Constant	39.22*** (0.58)	38.73*** (5.92)	38.57*** (0.58)	37.63*** (5.92)
State Controls	×	✓	×	✓
Job Controls	×	✓	×	✓
Firstname Controls	×	✓	×	✓
Lastname Controls	×	✓	×	✓
Picture trait Controls	×	✓	×	✓
Picture specific random effects	✓	✓	✓	✓
Profile specific random effects	✓	✓	✓	✓
Observations	1,599	1,599	1,599	1,599
Notes:	p<0.10;*p<0.05;**p<0.01;***p<0.001.			

Table J.22: Profile views and picture swapping

The table estimates the number of profile views as a function of whether the profile picture has been swapped or not. *Weeks after swapping* is a continuous variable from zero (28.07) to three (17.08). *Picture swapped* is a dummy variable with value one if the profile picture has been swapped, and zero otherwise. The regressions are conducted on the profile level, use various controls, and all follow Equation 1. The number of observations reflects the 400 profiles at four points in time minus one profile where a technical error prevented to record the number of views on the 3rd of August.

	Probability of connection being suspended (in %)																							
	All connected targets												Connected targets who did not receive message											
	All profiles						Black profiles						White profiles						All profiles					
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)	(16)	(17)	(18)	(19)	(20)	(21)	(22)	(23)	(24)
Picture Swapped	-0.04 (0.10)	-0.17 (0.11)	-0.08 (0.14)	-0.39 (0.21)	-0.01 (0.15)	-0.41* (0.20)	0.10 (0.14)	-0.04 (0.15)	-0.11 (0.20)	-0.41 (0.26)	0.29 (0.21)	-0.13 (0.25)	0.12 (0.14)	-0.03 (0.17)	-0.14 (0.15)	-0.37 (0.22)	0.35 (0.23)	-0.22 (0.33)	-0.26 (0.32)	-0.23 (0.34)	-0.26 (0.52)	-0.50 (0.69)	-0.32 (0.36)	-0.78 (0.47)
Target received message							0.39 (0.21)	0.39 (0.21)	0.32 (0.30)	0.35 (0.30)	0.39 (0.30)	0.37 (0.30)												
Target received message x Picture Swapped							-0.47 (0.34)	-0.45 (0.34)	0.10 (0.48)	0.09 (0.49)	-0.96* (0.48)	-0.91 (0.48)												
Constant	1.14*** (0.14)	2.09 (1.44)	1.31*** (0.21)	3.32 (2.23)	0.98*** (0.18)	2.59 (2.35)	1.01*** (0.16)	1.97 (1.44)	1.18*** (0.23)	3.25 (2.23)	0.86*** (0.21)	2.42 (2.35)	0.98*** (0.16)	1.49 (1.89)	1.10*** (0.22)	1.73 (2.35)	0.89*** (0.23)	3.18 (3.46)	0.98*** (0.23)	3.40 (3.26)	1.28*** (0.38)	8.24 (5.97)	0.70** (0.25)	2.52 (4.39)
State Controls	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓
Job Controls	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓
Firstname Controls	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓
Lastname Controls	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓
Picture trait Controls	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓
Picture specific random effects	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Target specific random effects	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Observations	9,523	9,523	4,724	4,724	4,799	4,799	9,523	9,523	4,724	4,724	4,799	4,799	6,204	6,204	3,055	3,055	3,149	3,149	3,319	3,319	1,669	1,669	1,650	1,650
Notes:	p<0.10;*p<0.05;**p<0.01;***p<0.001.																							

Table J.23: Connection suspension and picture swapping

The table estimates the probability of a connection being suspended as a function of whether the profile picture has been swapped or not. *Target received message* is a dummy with value one if the target received a message, and zero otherwise. *Picture swapped* is a dummy variable with value one if the profile picture has been swapped, and zero otherwise. The regressions are conducted on the target level, use various controls, and all follow Equation 1. Columns (1), (2), (7), (8), (13), (14), (19), and (20) focus on all profiles, while Columns (3), (4), (9), (10), (15), (16), (21), and (22), and Columns (5), (6), (11), (12), (17), (18), (23), and (24) focus on Black and White profiles only, respectively. Columns (1)-(12) focus on all targets. Columns (13)-(18) restrict the sample to those targets that have *not* been messaged, while Columns (19)-(24) restrict the sample to those targets that have been messaged.

	Response Rate (in %)				Message Length (in char)				Highly Useful Message?				Usefulness score			
	Swapped Only		All		Swapped Only		All		Swapped Only		All		Swapped Only		All	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)	(16)
Time passed between 28.07 and acceptance	1.85 (1.12)	2.17 (1.14)	1.83 (1.15)	2.21 (1.16)	0.01 (0.07)	0.01 (0.08)	0.01 (0.07)	0.01 (0.07)	-0.01 (0.02)	-0.01 (0.02)	-0.01 (0.02)	-0.01 (0.02)	0.04 (0.16)	0.07 (0.19)	0.03 (0.16)	0.04 (0.17)
Picture not swapped			2.74 (1.41)	2.57 (1.52)			-0.19* (0.09)	-0.26** (0.10)			0.01 (0.02)	0.02 (0.02)			-0.10 (0.20)	-0.27 (0.22)
Picture not swapped x Time passed between 28.07 and acceptance			-0.83 (1.63)	-1.59 (1.65)			-0.03 (0.09)	-0.05 (0.10)			0.02 (0.02)	0.01 (0.03)			0.03 (0.23)	0.004 (0.24)
Constant	18.39*** (1.01)	36.98 (23.34)	18.37*** (0.99)	30.41* (13.95)	0.10 (0.07)	-2.57 (1.84)	0.10 (0.06)	-1.60 (0.88)	0.07*** (0.02)	0.48 (0.42)	0.07*** (0.02)	0.20 (0.22)	6.27*** (0.15)	6.34 (3.84)	6.26*** (0.14)	2.81 (2.02)
State Controls	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓
Job Controls	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓
Firstname Controls	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓
Lastname Controls	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓
Picture trait Controls	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓
Picture specific random effects	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Observations	1,685	1,685	3,326	3,326	315	315	664	664	315	315	664	664	315	315	664	664
Notes: p<0.10;*p<0.05;**p<0.01;***p<0.001.																

Table J.24: Response rate and message characteristics by the time difference between swapping and accepting

The table estimates how several response characteristics are affected by whether the profile picture has been swapped or not and the time difference between swapping and accepting. *Time passed between 28.07 and acceptance* is a continuous measure of the time passed (in days) between accepting the profile's connection request in Stage I and the 28th of July when the profile pictures of half of the sample have been swapped. *Picture not swapped* is a dummy variable with value one if the profile picture has not been swapped, and zero otherwise. The regressions are conducted on the target level, use various controls, and all follow Equation 1. Columns (1), (2), (5), (6), (9), (10), (13), and (14) focus only on those targets that have been messaged by profiles whose picture has been swapped. The remaining columns focus on all targets. We have a reduced sample size in analyzing the replies, as not all requests are answered.

J3.4 Response Probabilities and Heterogeneity

	Response Rate (in %)		Message Length (in char)		Highly Useful Message?		Usefulness score	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
White	-0.44 (1.80)	-0.05 (2.23)	0.15 (0.10)	0.23 (0.12)	-0.02 (0.03)	-0.03 (0.04)	0.12 (0.25)	0.04 (0.29)
Male	0.23 (1.61)	0.04 (1.89)	-0.03 (0.09)	-0.12 (0.10)	0.07** (0.03)	0.08** (0.03)	-0.17 (0.22)	-0.30 (0.25)
Age	0.05 (0.08)	0.02 (0.10)	-0.002 (0.005)	-0.001 (0.01)	-0.001 (0.001)	-0.001 (0.001)	-0.01 (0.01)	-0.02 (0.01)
Bachelor+	2.05 (1.83)	-0.53 (2.40)	0.12 (0.11)	-0.02 (0.13)	0.02 (0.03)	0.01 (0.04)	0.39 (0.26)	0.02 (0.32)
Contact Count	0.02*** (0.005)	0.02** (0.01)	-0.001 (0.0003)	-0.001* (0.0003)	0.0002* (0.0001)	0.0002* (0.0001)	0.0004 (0.001)	-0.0001 (0.001)
HR.Job		-3.29 (3.01)		0.04 (0.17)		0.05 (0.05)		0.84* (0.41)
Same Uni		3.96 (2.57)		0.09 (0.13)		-0.01 (0.04)		0.35 (0.32)
UniWhite		5.74 (9.00)		-0.32 (0.45)		-0.10 (0.13)		-1.58 (1.08)
UniBlack		-1.78 (6.21)		0.09 (0.34)		0.03 (0.10)		-0.08 (0.80)
Share Democrat		-8.27 (6.27)		0.58 (0.35)		0.19 (0.10)		-0.55 (0.83)
Nbr. of Profile's friends		0.38* (0.15)		-0.001 (0.01)		-0.003 (0.002)		-0.01 (0.02)
Constant	12.03*** (3.28)	10.92 (7.94)	0.05 (0.19)	-0.10 (0.44)	0.02 (0.05)	-0.02 (0.12)	6.27*** (0.46)	7.85*** (1.03)
Observations	2,550	1,956	514	406	514	406	514	406

Notes: p<0.10;*p<0.05;**p<0.01;***p<0.001.

Table J.25: Response probability and usefulness by target characteristics

The table reports upon regressions estimating several response characteristics as a function of multiple target characteristics. The number of observations varies depending on the user's characteristics (we do not have all characteristics for all users), and the user's second-stage behavior (e.g., as not all requests are answered we have a reduced sample size in analyzing the replies).

J3.5 Ex-Ante Informational Benefit

Panel A: Aggregate difference in the ex-ante informational benefit of the network

	Ex-ante informational benefit						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Profile is White	-1.01*** (0.10)	-1.01*** (0.10)	-1.01*** (0.10)	-1.02*** (0.09)	-1.00*** (0.10)	-0.98*** (0.11)	-0.99*** (0.11)
Constant	5.32*** (0.09)	8.03*** (0.40)	5.24*** (0.22)	5.34*** (0.18)	5.07*** (0.23)	4.98*** (1.04)	7.61*** (1.05)
State Controls	×	✓	×	×	×	×	✓
Job Controls	×	×	✓	×	×	×	✓
Firstname Controls	×	×	×	✓	×	×	✓
Lastname Controls	×	×	×	×	✓	×	✓
Picture trait Controls	×	×	×	×	×	✓	✓
Picture specific random effects	✓	✓	✓	✓	✓	✓	✓
Observations	400	400	400	400	400	400	400

Panel B: Differences in the ex-ante informational benefit of the network accounting for profile quality

	Ex-ante informational benefit						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Profile is Black	-1.03*** (0.14)	-1.03*** (0.14)	-1.03*** (0.14)	-1.09*** (0.13)	-1.03*** (0.14)	-0.99*** (0.15)	-1.06*** (0.15)
Profile attended worse Uni	0.03 (0.18)	0.03 (0.15)	0.04 (0.18)	-0.02 (0.18)	-0.001 (0.18)	0.06 (0.18)	-0.03 (0.15)
Profile is Black and attended worse Uni	0.05 (0.19)	0.05 (0.19)	0.05 (0.19)	0.13 (0.19)	0.06 (0.20)	0.03 (0.19)	0.15 (0.20)
Constant	5.30*** (0.13)	8.02*** (0.41)	5.21*** (0.24)	5.35*** (0.20)	5.07*** (0.25)	4.90*** (1.05)	7.55*** (1.07)
State Controls	×	✓	×	×	×	×	✓
Job Controls	×	×	✓	×	×	×	✓
Firstname Controls	×	×	×	✓	×	×	✓
Lastname Controls	×	×	×	×	✓	×	✓
Picture trait Controls	×	×	×	×	×	✓	✓
Picture specific random effects	✓	✓	✓	✓	✓	✓	✓
Observations	400	400	400	400	400	400	400

Notes:

·p<0.10;*p<0.05;**p<0.01;***p<0.001.

Table J.26: Ex-ante informational benefit of the network by race and education of profile

The table estimates the ex-ante informational benefit a profile would have as a function of their race. Panel A focuses only on race, while Panel B additionally reports the interaction between profile quality and race. The regressions are conducted on the profile level, use various controls, and all follow Equation 1.

J4 Validation Experiment

J4.1 First Stage: Captcha Task

	Picture selected as computer generated						
Our AI-Pictures	-0.03 (0.02)	-0.01 (0.02)	-0.04 (0.02)	-0.03 (0.02)	-0.02 (0.03)	-0.03 (0.02)	-0.02 (0.05)
Obvious Fake	0.69*** (0.03)	0.69*** (0.04)	0.71*** (0.04)	0.69*** (0.03)	0.75*** (0.04)	0.69*** (0.04)	0.72*** (0.07)
Our AI-Pictures x Rater is non-White		-0.08*** (0.02)					
Obvious Fake x Rater is non-White		0.005 (0.03)					
Our AI-Pictures x Rater is female			0.03 (0.02)				
Obvious Fake x Rater is female			-0.03 (0.03)				
Our AI-Pictures x Age of Rater				0.01 (0.01)			
Obvious Fake x Age of Rater				0.03** (0.01)			
Our AI-Pictures x Rater is a democrat					-0.01 (0.02)		
Obvious Fake x Rater is a democrat					-0.09*** (0.03)		
Constant (Real Picture)	0.15*** (0.02)	0.14*** (0.02)	0.14*** (0.02)	0.15*** (0.02)	0.14*** (0.03)	0.08* (0.03)	0.14*** (0.05)
Controls	×	×	×	×	×	✓	×
Weighted Sample	×	×	×	×	×	×	✓
Sbj.Spec.Effects	✓	✓	✓	✓	✓	✓	✓
Pic.Spec.Effects	✓	✓	✓	✓	✓	✓	✓
Main effects omitted	✓	✓	✓	✓	✓	✓	✓
Observations	6,141	6,141	6,141	6,141	6,141	6,141	6,141

Notes: p<0.10; *p<0.05; **p<0.01; ***p<0.001.

Table J.27: Regressions estimating the likelihood of a picture being selected as fake

The table estimates whether a given picture is selected as computer-generated as a function of whether the profile is AI-generated. *Sbj.Spec.Effects* and *Pic.Spec.Effects* denote subjects and picture-specific random effects accounting for the fact that each rater has seen multiple pictures and the fact that each picture has been rated multiple times. Whenever interactions are estimated, we omit the main effect for brevity. Controls include the rater's age, gender, education, income, ethnicity, and political preference. In the last column, we reweight our sample to match the sample characteristics of LinkedIn users. The sample size results from 307 reliable raters who decided on 20 pictures each.

	Picture selected as computer generated						
Our AI-Picture (AI)	0.02 (0.03)	0.04 (0.03)	0.005 (0.03)	0.02 (0.03)	0.04 (0.04)	0.02 (0.03)	0.04 (0.07)
Picture of White Person (PWP)	0.04 (0.04)	0.04 (0.05)	0.04 (0.05)	0.04 (0.04)	0.04 (0.05)	0.04 (0.04)	0.06 (0.09)
PWP x AI	-0.10* (0.04)	-0.09* (0.05)	-0.09* (0.05)	-0.10* (0.04)	-0.12* (0.05)	-0.08* (0.05)	-0.12 (0.10)
AI x Rater is non-white		-0.09** (0.03)					
PWP x Rater is non-white		-0.01 (0.04)					
PWP x AI x Rater is non-white		0.02 (0.05)					
AI x Rater is female			0.03 (0.03)				
PWP x Rater is female			-0.002 (0.03)				
PWP x AI x Rater is female			0.01 (0.04)				
AI x Age of Rater				0.02* (0.01)			
PWP x Age of Rater				0.02* (0.01)			
PWP x AI x Age of Rater				-0.03* (0.02)			
AI x Rater is a democrat					-0.04 (0.03)		
PWP x Rater is a democrat					-0.001 (0.03)		
PWP x AI x Rater is a democrat					0.06 (0.04)		
Constant (Real Picture of Black Person)	0.13*** (0.03)	0.12*** (0.03)	0.12*** (0.03)	0.13*** (0.03)	0.12*** (0.03)	0.08* (0.04)	0.12* (0.07)
Controls	×	×	×	×	×	✓	×
Weighted Sample	×	×	×	×	×	×	✓
Sbj.Spec.Effects	✓	✓	✓	✓	✓	✓	✓
Pic.Spec.Effects	✓	✓	✓	✓	✓	✓	✓
Main effects omitted	✓	✓	✓	✓	✓	✓	✓
Observations	4,913	4,913	4,913	4,913	4,913	4,913	4,913

Notes: p<0.10; *p<0.05; **p<0.01; ***p<0.001.

Table J.28: Regression estimates on the differences in the likelihood of a picture being selected as fake by race of the person in picture

The table estimates whether a given picture is selected as computer-generated as a function of whether the profile is AI-generated and the race of the person on the picture. *Sbj.Spec.Effects* and *Pic.Spec.Effects* denote subjects and picture-specific random effects accounting for the fact that each rater has seen multiple pictures and the fact that each picture has been rated multiple times. Whenever interactions are estimated, we omit the main effect for brevity. Controls include the rater's age, gender, education, income, ethnicity, and political preference. In the last column, we reweight our sample to match the sample characteristics of LinkedIn users. The sample size results from 307 reliable raters who decided on 16 pictures (20 minus 4 obvious fake ones) each.

J4.2 Second Stage: Individual Rating Task

Panel A: Age of person in the picture								
	How old is the person in this picture?							
Picture of Black Person (PBP)	0.91*** (0.21)	0.82*** (0.24)	0.50* (0.27)	0.93*** (0.21)	1.04** (0.35)	0.96*** (0.23)	0.90*** (0.21)	1.20*** (0.22)
PBP x Rater is non-white		0.37 (0.50)						
PBP x Rater is female			1.04* (0.44)					
PBP x Age of Rater				-0.23 (0.21)				
PBP x Rater is a democrat					-0.22 (0.44)			
PBP x Rater rated picture as fake						-0.66 (0.68)		
Constant	31.53*** (0.27)	31.58*** (0.30)	31.35*** (0.33)	31.51*** (0.27)	31.34*** (0.42)	31.46*** (0.27)	32.19*** (1.07)	31.46*** (0.33)
Controls	x	x	x	x	x	x	x	x
Weighted Sample	x	x	x	x	x	x	x	x
Sbj.Spec.Effects	✓	✓	✓	✓	✓	✓	✓	✓
Pic.Spec.Effects	✓	✓	✓	✓	✓	✓	✓	✓
Main effects omitted	✓	✓	✓	✓	✓	✓	✓	✓
Observations	1,740	1,740	1,740	1,740	1,740	1,740	1,740	1,740
Panel B: Gender of person in the picture								
	How likely is the person in this picture female?							
Picture of Black Person (PBP)	0.05 (0.23)	-0.11 (0.27)	0.05 (0.29)	0.06 (0.23)	0.03 (0.38)	0.04 (0.25)	0.05 (0.23)	0.08 (0.15)
PBP x Rater is non-white		0.63 (0.54)						
PBP x Rater is female			-0.01 (0.48)					
PBP x Age of Rater				-0.14 (0.23)				
PBP x Rater is a democrat					0.03 (0.48)			
PBP x Rater rated picture as fake						0.52 (0.74)		
Constant	1.82*** (0.49)	2.36*** (0.57)	2.00** (0.63)	1.81*** (0.50)	1.50* (0.82)	1.94*** (0.50)	1.15 (2.30)	1.82*** (0.49)
Controls	x	x	x	x	x	x	x	x
Weighted Sample	x	x	x	x	x	x	x	x
Sbj.Spec.Effects	✓	✓	✓	✓	✓	✓	✓	✓
Pic.Spec.Effects	✓	✓	✓	✓	✓	✓	✓	✓
Main effects omitted	✓	✓	✓	✓	✓	✓	✓	✓
Observations	1,740	1,740	1,740	1,740	1,740	1,740	1,740	1,740
Panel C: Person in the picture is black								
	How likely is the person in this picture African American?							
Picture of Black Person (PBP)	89.24*** (0.71)	87.93*** (0.81)	90.48*** (0.90)	88.95*** (0.70)	93.62*** (1.16)	89.51*** (0.76)	89.24*** (0.71)	92.18*** (0.57)
PBP x Rater is non-white		5.44*** (1.64)						
PBP x Rater is female			-3.23* (1.45)					
PBP x Age of Rater				3.34*** (0.69)				
PBP x Rater is a democrat					-6.90*** (1.46)			
PBP x Rater rated picture as fake						-1.51 (2.19)		
Constant	2.98*** (0.57)	3.39*** (0.66)	2.61*** (0.72)	3.07*** (0.57)	1.39 (0.93)	3.06*** (0.60)	0.63 (2.17)	1.47* (0.68)
Controls	x	x	x	x	x	x	x	x
Weighted Sample	x	x	x	x	x	x	x	x
Sbj.Spec.Effects	✓	✓	✓	✓	✓	✓	✓	✓
Pic.Spec.Effects	✓	✓	✓	✓	✓	✓	✓	✓
Main effects omitted	✓	✓	✓	✓	✓	✓	✓	✓
Observations	1,740	1,740	1,740	1,740	1,740	1,740	1,740	1,740
Panel D: Person in the picture is white								
	How likely is the person in this picture White?							
Picture of Black Person (PBP)	-92.34*** (0.54)	-91.43*** (0.61)	-92.83*** (0.68)	-92.24*** (0.54)	-95.31*** (0.88)	-92.02*** (0.57)	-92.34*** (0.54)	-93.94*** (0.45)
PBP x Rater is non-white		-3.78** (1.25)						
PBP x Rater is female			1.32 (1.10)					
PBP x Age of Rater				-1.24* (0.53)				
PBP x Rater is a democrat					4.68*** (1.11)			
PBP x Rater rated picture as fake						-3.24* (1.66)		
Constant	95.91*** (0.43)	95.62*** (0.50)	96.57*** (0.55)	95.87*** (0.44)	97.08*** (0.71)	95.58*** (0.46)	95.20*** (1.63)	96.48*** (0.52)
Controls	x	x	x	x	x	x	x	x
Weighted Sample	x	x	x	x	x	x	x	x
Sbj.Spec.Effects	✓	✓	✓	✓	✓	✓	✓	✓
Pic.Spec.Effects	✓	✓	✓	✓	✓	✓	✓	✓
Main effects omitted	✓	✓	✓	✓	✓	✓	✓	✓
Observations	1,740	1,740	1,740	1,740	1,740	1,740	1,740	1,740
Panel E: Person in the picture is asian								
	How likely is the person in this picture Asian?							
Picture of Black Person (PBP)	4.11*** (0.55)	4.68*** (0.63)	4.48*** (0.70)	4.08*** (0.56)	3.32*** (0.91)	3.66*** (0.59)	4.12*** (0.55)	3.87*** (0.57)
PBP x Rater is non-white		-2.35* (1.29)						
PBP x Rater is female			-0.95 (1.14)					
PBP x Age of Rater				0.32 (0.54)				
PBP x Rater is a democrat					1.26 (1.15)			
PBP x Rater rated picture as fake						2.41 (1.76)		
Constant	4.59*** (0.83)	5.22*** (0.96)	3.89*** (1.06)	4.64*** (0.84)	2.91* (1.37)	4.48*** (0.85)	8.74* (3.72)	4.91*** (0.94)
Controls	x	x	x	x	x	x	x	x
Weighted Sample	x	x	x	x	x	x	x	x
Sbj.Spec.Effects	✓	✓	✓	✓	✓	✓	✓	✓
Pic.Spec.Effects	✓	✓	✓	✓	✓	✓	✓	✓
Main effects omitted	✓	✓	✓	✓	✓	✓	✓	✓
Observations	1,740	1,740	1,740	1,740	1,740	1,740	1,740	1,740
Notes:								
p<0.10;*p<0.05;**p<0.01;***p<0.001.								

Table J.29: Regression estimates on the differences in rated demographic characteristics of the pictures

The table estimates several characteristics as a function of the race of the person in the picture. *Sbj.Spec.Effects* and *Pic.Spec.Effects* denote subjects and picture-specific random effects accounting for the fact that each rater has seen multiple pictures and the fact that each picture has been rated multiple times. Whenever interactions are estimated, we omit the main effect for brevity. Controls include the rater's age, gender, education, income, ethnicity, and political preference. In the last column, we reweight our sample to match the sample characteristics of LinkedIn users. The number of observations reflect that each of the 307 reliable participants rated multiple pictures, but not all of these pictures were finally used in the experiment, which is why we have 1740 observations per characteristic.

Panel A: Trustworthiness of person in the picture

	How trustworthy do you think is the person in this picture?							
Picture of Black Person (PBP)	2.47*** (0.54)	1.76** (0.62)	3.32*** (0.69)	2.49*** (0.55)	0.91 (0.89)	2.55*** (0.58)	2.45*** (0.54)	5.96*** (0.61)
PBP x Rater is non-white		2.91* (1.27)						
PBP x Rater is female			-2.23* (1.12)					
PBP x Age of Rater				-0.22 (0.53)				
PBP x Rater is a democrat					2.46* (1.12)			
PBP x Rater rated picture as fake						0.93 (1.73)		
Constant	69.30*** (0.91)	70.45*** (1.05)	68.96*** (1.16)	69.16*** (0.90)	69.54*** (1.51)	69.65*** (0.92)	62.79*** (4.11)	67.59*** (1.07)
Controls	×	×	×	×	×	×	✓	×
Weighted Sample	×	×	×	×	×	×	×	✓
Sbj.Spec.Effects	✓	✓	✓	✓	✓	✓	✓	✓
Pic.Spec.Effects	✓	✓	✓	✓	✓	✓	✓	✓
Main effects omitted	✓	✓	✓	✓	✓	✓	✓	✓
Observations	1,740	1,740	1,740	1,740	1,740	1,740	1,740	1,740

Panel B: Intelligence of person in the picture

	How intelligent do you think is the person in this picture?							
Picture of Black Person (PBP)	0.19 (0.47)	-0.48 (0.54)	0.73 (0.60)	0.15 (0.47)	-0.27 (0.77)	0.18 (0.50)	0.18 (0.47)	1.60*** (0.47)
PBP x Rater is non-white		2.78* (1.09)						
PBP x Rater is female			-1.40 (0.96)					
PBP x Age of Rater				0.39 (0.46)				
PBP x Rater is a democrat					0.73 (0.97)			
PBP x Rater rated picture as fake						0.97 (1.50)		
Constant	73.31*** (0.82)	74.21*** (0.94)	72.29*** (1.04)	73.21*** (0.81)	73.07*** (1.35)	73.53*** (0.83)	68.14*** (3.69)	72.74*** (0.90)
Controls	×	×	×	×	×	×	✓	×
Weighted Sample	×	×	×	×	×	×	×	✓
Sbj.Spec.Effects	✓	✓	✓	✓	✓	✓	✓	✓
Pic.Spec.Effects	✓	✓	✓	✓	✓	✓	✓	✓
Main effects omitted	✓	✓	✓	✓	✓	✓	✓	✓
Observations	1,740	1,740	1,740	1,740	1,740	1,740	1,740	1,740

Panel C: Authenticity of person in the picture

	How authentic do you think is the person in this picture?							
Picture of Black Person (PBP)	3.40*** (0.54)	3.35*** (0.62)	3.84*** (0.69)	3.37*** (0.54)	2.22* (0.89)	3.44*** (0.58)	3.39*** (0.54)	5.17*** (0.58)
PBP x Rater is non-white		0.23 (1.26)						
PBP x Rater is female			-1.14 (1.11)					
PBP x Age of Rater				0.35 (0.53)				
PBP x Rater is a democrat					1.87* (1.12)			
PBP x Rater rated picture as fake						0.89 (1.73)		
Constant	70.01*** (0.93)	70.64*** (1.07)	69.73*** (1.19)	69.86*** (0.92)	71.71*** (1.54)	70.28*** (0.94)	59.13*** (4.19)	68.96*** (1.06)
Controls	×	×	×	×	×	×	✓	×
Weighted Sample	×	×	×	×	×	×	×	✓
Sbj.Spec.Effects	✓	✓	✓	✓	✓	✓	✓	✓
Pic.Spec.Effects	✓	✓	✓	✓	✓	✓	✓	✓
Main effects omitted	✓	✓	✓	✓	✓	✓	✓	✓
Observations	1,740	1,740	1,740	1,740	1,740	1,740	1,740	1,740

Panel D: Looks of person in the picture

	How good looking do you think is the person in this picture?							
Picture of Black Person (PBP)	-4.73*** (0.55)	-4.17*** (0.64)	-4.76*** (0.71)	-4.84*** (0.56)	-6.78*** (0.91)	-5.12*** (0.59)	-4.74*** (0.55)	-3.54*** (0.57)
PBP x Rater is non-white		-2.31* (1.29)						
PBP x Rater is female			0.06 (1.14)					
PBP x Age of Rater				0.94* (0.54)				
PBP x Rater is a democrat					3.24** (1.15)			
PBP x Rater rated picture as fake						3.83* (1.77)		
Constant	75.22*** (1.02)	76.17*** (1.17)	74.21*** (1.29)	75.10*** (1.01)	74.79*** (1.68)	75.55*** (1.03)	68.06*** (4.58)	74.49*** (1.14)
Controls	×	×	×	×	×	×	✓	×
Weighted Sample	×	×	×	×	×	×	×	✓
Sbj.Spec.Effects	✓	✓	✓	✓	✓	✓	✓	✓
Pic.Spec.Effects	✓	✓	✓	✓	✓	✓	✓	✓
Main effects omitted	✓	✓	✓	✓	✓	✓	✓	✓
Observations	1,740	1,740	1,740	1,740	1,740	1,740	1,740	1,740

Notes:

*p<0.10; **p<0.05; ***p<0.01; ****p<0.001.

Table J.30: Regression estimates on the differences in rated traits of the person in the picture. The table estimates several characteristics as a function of the race of the person in the picture. *Sbj.Spec.Effects* and *Pic.Spec.Effects* denote subjects and picture-specific random effects accounting for the fact that each rater has seen multiple pictures and the fact that each picture has been rated multiple times. Whenever interactions are estimated, we omit the main effect for brevity. Controls include the rater's age, gender, education, income, ethnicity, and political preference. In the last column, we reweight our sample to match the sample characteristics of LinkedIn users. The number of observations reflect that each of the 307 reliable participants rated multiple pictures, but not all of these pictures were finally used in the experiment, which is why we have 1740 observations per characteristic.

J4.3 Third Stage: University Rating Task

The better universities are correctly identified as such										
Rater is non-white	0.01 (0.01)								0.03*** (0.01)	
Rater is female		-0.02* (0.01)							-0.02* (0.01)	
Age of Rater			0.03*** (0.01)						0.03*** (0.004)	
Rater has at least a bachelor				0.01 (0.01)					0.01 (0.01)	
Rater's homestate					0.01 (0.03)				0.01 (0.03)	
Rater's household income < 75k						-0.01 (0.01)			-0.01 (0.01)	
Rater is a democrat							-0.02*** (0.01)		-0.02* (0.01)	
Constant	0.67*** (0.02)	0.66*** (0.02)	0.67*** (0.02)	0.66*** (0.02)	0.66*** (0.03)	0.67*** (0.02)	0.67*** (0.02)	0.68*** (0.02)	0.67*** (0.02)	0.67*** (0.03)
Weighted Sample	×	×	×	×	×	×	×	×	×	✓
Sbj.Spec.Effects	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
State.Spec.Effects	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Observations	15,657	15,657	15,657	15,657	15,657	15,657	15,657	15,657	15,657	15,657

Notes:

p<0.10;*p<0.05;**p<0.01;***p<0.001.

Table J.31: Regression estimates on the propensity of correctly identifying the better ranked university

The table estimates the propensity of correctly identifying the better-ranked university as a function of a host of rater characteristics (like age, gender, education, etc.). To account for the fact that two universities were rated per state and each person rated multiple universities, we include state and subject-specific random effects. In the last column, we reweight our sample to match the sample characteristics of LinkedIn users. The number of observations reflect that each of the 307 reliable participants decided 51 times between a better and worse-ranked university.

J5 Expert Survey

	How many more connections do White profiles have relative to Black profiles?								How many more responses do White profiles receive relative to Black profiles?							
Knows this research	-4.88 (4.28)								5.61 (3.11)							
Knows results	-10.80 (6.16)								5.80 (4.51)							
Is Female	1.61 (1.62)								1.29 (1.18)							
Works on Discrimination	-0.22 (1.86)								0.35 (1.36)							
Has published	10.56 (5.76)								-0.37 (4.24)							
Is Prof	-1.43 (2.39)								0.28 (1.75)							
Is White	-4.26 (2.65)								-2.72 (1.98)							
Constant	18.42*** (0.92)	18.76*** (0.97)	18.77*** (0.94)	16.28*** (2.34)	18.55*** (1.40)	8.14 (5.69)	19.59*** (2.16)	22.44*** (2.45)	12.92*** (0.67)	12.54*** (0.70)	12.74*** (0.69)	11.21*** (1.71)	12.73*** (1.02)	13.29** (4.18)	12.69*** (1.58)	15.17*** (1.84)
Observations	269	269	269	269	269	269	269	254	269	269	269	269	269	269	269	254

Notes: .p<0.10; *p<0.05; **p<0.01; ***p<0.001.

Table J.32: Regression estimates experts' predictions of the first and second stage of the experiment

The table reports the average prediction of experts with regard to the average Stage I and Stage II results. The first eight columns denote the prediction of the relative gap between White profiles relative to Black profiles in terms of connections. Columns (9)-(16) denote the prediction of the relative gap between White profiles relative to Black profiles in terms of received responses. *Is Female* indicates whether the expert is a woman, *Knows results* indicates whether the expert has heard of the results, *Knows this research* indicates whether the expert has heard of this research, *Has published* indicates whether the expert has ever published in a peer-reviewed journal, *Works on Discrimination* indicates whether the expert works themselves on discrimination research. *Is Prof* indicates whether the expert has a professorial position (assistant, associate, or full professor), *Is White* indicates whether the expert indicated to be White (15 experts provided no answer to this demographic).

	How many more connections do White profiles have relative to Black profiles?															
	By user's age								By user's gender							
GenX	-6.11*** (0.54)	-6.09*** (0.56)	-6.17*** (0.55)	-6.69*** (1.36)	-6.38*** (0.81)	-4.29 (3.32)	-6.78*** (1.26)	-5.53*** (1.47)								
GenY	-12.00*** (0.54)	-11.63*** (0.56)	-11.79*** (0.55)	-11.67*** (1.36)	-12.19*** (0.81)	-7.29* (3.32)	-13.67*** (1.26)	-11.97*** (1.47)								
GenZ	-16.27*** (0.54)	-16.12*** (0.56)	-16.36*** (0.55)	-16.95*** (1.36)	-16.91*** (0.81)	-10.14** (3.32)	-18.37*** (1.26)	-16.14*** (1.47)								
Knows this research		-2.25 (3.11)													-6.47* (3.14)	
GenX:Knows this research		-0.36 (2.49)														
GenY:Knows this research		-5.25* (2.49)														
GenZ:Knows this research		-2.21 (2.49)														
Knows results			-8.93* (4.46)												-7.51* (4.52)	
GenX:Knows results			1.75 (3.59)													
GenY:Knows results			-6.53* (3.59)													
GenZ:Knows results			3.06 (3.59)													
Is Female				1.07 (1.18)											0.18 (1.19)	
GenX:Is Female				0.43 (0.94)												
GenY:Is Female				-0.24 (0.94)												
GenZ:Is Female				0.51 (0.94)												
Works on Discrimination					-1.26 (1.35)										0.12 (1.37)	
GenX:Works on Discrimination					0.47 (1.08)											
GenY:Works on Discrimination					0.34 (1.08)											
GenZ:Works on Discrimination					1.13 (1.08)											
Has published						10.21* (4.20)									7.92* (4.26)	
GenX:Has published						-1.87 (3.37)										
GenY:Has published						-4.84 (3.37)										
GenZ:Has published						-6.29* (3.37)										
Is Prof							-3.81* (1.73)									-0.43 (1.76)
GenX:Is Prof							0.81 (1.39)									
GenY:Is Prof							2.05 (1.39)									
GenZ:Is Prof							2.57* (1.39)									
Is White								-1.79 (1.94)								-4.51* (1.97)
GenX:Is White								-0.68 (1.58)								
GenY:Is White								-0.01 (1.58)								
GenZ:Is White								-0.15 (1.58)								
Male:Knows this research									0.84 (3.60)							
Male:Knows results									-6.78 (5.19)							
Male:Is Female										0.94 (1.36)						
Male:Works on Discrimination											-1.49 (1.56)					
Male:Has published												-1.16 (4.87)				
Male:Is Prof													-3.45* (2.00)			
Male:Is White														3.39 (2.33)		
Male														2.39 (2.16)		
Constant	22.52*** (0.67)	22.67*** (0.70)	22.80*** (0.68)	21.09*** (1.70)	23.23*** (1.02)	12.57** (4.14)	25.63*** (1.57)	24.08*** (1.80)	5.30*** (0.77)	5.24*** (0.81)	5.52*** (0.79)	4.05* (1.97)	6.15*** (1.17)	6.43 (4.81)	8.12*** (1.81)	14.19*** (1.83)
Sbj.Spec.Effects	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Observations	1,076	1,076	1,076	1,076	1,076	1,076	1,076	1,016	538	538	538	538	538	538	538	508

Notes: p<0.10;*p<0.05;**p<0.01;***p<0.001.

Table J.33: Regression estimates experts' predictions of how age and gender affect discrimination

The table reports the average prediction of experts with regard to how the relative connection gap between White and Black profiles differs as a function of the users' age and the users' gender. The first eight columns denote the prediction of how age affects discrimination. Columns (9)-(16) denote the prediction of how gender affects discrimination. *GenX*, *GenY*, and *GenZ* denotes dummy variables indicating whether the user is part of the generation Gen X, Gen Y, or Gen Z. *Male* denotes a dummy indicating whether the user is a man or a woman. *Is Female* indicates whether the expert is a woman, *Knows results* indicates whether the expert has heard of the results, *Knows this research* indicates whether the expert has heard of this research, *Has published* indicates whether the expert has ever published in a peer-reviewed journal, *Works on Discrimination* indicates whether the expert works themselves on discrimination research. *Is Prof* indicates whether the expert has a professorial position (assistant, associate, or full professor), *Is White* indicates whether the expert indicated to be White (15 experts provided no answer to this demographic). As all experts have been asked multiple questions, we account for subject-specific heterogeneity by using a subject-specific random effect. The number of observations reflect that each of the 269 experts had to judge user's age on four categories, and gender on two.

	How many more connections do White profiles have relative to Black profiles?															
	By user's race								By user's level of education							
NonBlack	15.67*** (1.12)	14.90*** (1.17)	15.16*** (1.14)	14.53*** (2.85)	18.58*** (1.69)	3.14 (6.93)	19.35*** (2.62)	15.25*** (3.04)								
Knows this research		-8.09* (3.66)							-4.43 (3.37)							
Uni									-7.83*** (0.97)	-8.09*** (1.02)	-8.12*** (0.99)	-8.03** (2.46)	-8.82*** (1.47)	-2.29 (6.00)	-13.92*** (2.23)	-6.56* (2.61)
NonBlack:Knows this research		11.22* (5.18)														
Knows results			-14.16** (5.28)													
NonBlack:Knows results			16.26* (7.47)								-9.98* (4.85)					
Is Female				-0.09 (1.39)								1.40 (1.27)				
NonBlack:Is Female				0.86 (1.97)												
Works on Discrimination					1.59 (1.59)									-1.78 (1.46)		
NonBlack:Works on Discrimination					-5.15* (2.25)											
Has published						-5.69 (4.97)									9.15* (4.54)	
NonBlack:Has published						12.86* (7.02)										
Is Prof							1.50 (2.05)									-6.78*** (1.86)
NonBlack:Is Prof							-4.49 (2.90)									
Is White								-0.90 (2.32)								-1.00 (2.08)
NonBlack:Is White								0.47 (3.28)								
Uni:Knows this research									3.72 (4.49)							
Uni:Knows results										9.12 (6.47)						
Uni:Is Female											0.15 (1.70)					
Uni:Works on Discrimination												1.75 (1.95)				
Uni:Has published													-5.70 (6.08)			
Uni:Is Prof														7.44** (2.47)		
Uni:Is White															-1.87 (2.81)	
Constant	-0.97 (0.79)	-0.42 (0.83)	-0.53 (0.81)	-0.85 (2.01)	-1.87 (1.20)	4.57 (4.90)	-2.20 (1.85)	-0.17 (2.15)	16.91*** (0.73)	17.21*** (0.76)	17.22*** (0.74)	15.05*** (1.84)	17.91*** (1.10)	8.00* (4.48)	22.45*** (1.68)	18.03*** (1.93)
Sbj.Spec.Effects	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Observations	538	538	538	538	538	538	538	508	538	538	538	538	538	538	538	508

Notes:

p<0.10;*p<0.05;**p<0.01;***p<0.001.

Table J.34: Regression estimates experts' predictions of how race and education affect discrimination

The table reports the average prediction of experts with regard to how the relative connection gap between White and Black profiles differs as a function of the users' race and the users' education. The first eight columns denote the prediction of how race affects discrimination. Columns (9)-(16) denote the prediction of how education affects discrimination. *NonBlack* denotes a dummy indicating whether the user is non-Black. *Uni* denotes a dummy indicating whether the user has attended college. *Is Female* indicates whether the expert is a woman, *Knows results* indicates whether the expert has heard of the results, *Knows this research* indicates whether the expert has heard of this research, *Has published* indicates whether the expert has ever published in a peer-reviewed journal, *Works on Discrimination* indicates whether the expert works themselves on discrimination research. *Is Prof* indicates whether the expert has a professorial position (assistant, associate, or full professor), *Is White* indicates whether the expert indicated to be White (15 experts provided no answer to this demographic). As all experts have been asked multiple questions, we account for subject-specific heterogeneity by using a subject-specific random effect. The number of observations reflect that each of the 269 experts had to judge user's race and education on two categories each.