Discussion Paper Series – CRC TR 224

Discussion Paper No. 482
Project A 01, A 04, B 02

LinkedOut? A Field Experiment on Discrimination in Job Network Formation

Yulia Evsyukova[1]
Felix Rusche[2]
Wladislaw Mill[3]

December 2023

[1] University of Mannheim, Email: yulia.evsyukova@gess.uni-mannheim.de
[2] University of Mannheim, Email: felix.rusche@uni-mannheim.de
[3] University of Mannheim, Email: mill@uni-mannheim.de

Support by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) through CRC TR 224 is gratefully acknowledged.

Collaborative Research Center Transregio 224 - www.crctr224.de
Rheinische Friedrich-Wilhelms-Universität Bonn - Universität Mannheim

LinkedOut? A Field Experiment on Discrimination in Job Network Formation*

Yulia Evsyukova† ⓡ Felix Rusche‡ ⓡ Wladislaw Mill§

November 9, 2023

Abstract

We assess the impact of discrimination on Black individuals' job networks in the U.S. using a two-stage field experiment with 400+ fictitious LinkedIn profiles. Varying race via A.I.-generated images only, we find that Black profiles' connection requests are accepted at significantly lower rates (Stage I) and their networks provide less information (Stage II). Leveraging our experimental design to eliminate first-stage endogeneity, we identify gatekeeping as the key driver of Black-White disparities. Examining users' CVs reveals widespread discrimination across different social groups and – contrary to expert predictions – less discrimination among men and older users.

Keywords: Discrimination, Job Networks, Labor Markets, Field Experiment
JEL Classifications: J71, J15, C93, J46, D85

* **Acknowledgements:** This paper benefited from discussions with Albrecht Glitz, Amanda Palais, Amelie Schiprowski, Amit Goldberg, Andrej Shleifer, Anke Becker, Antonio Ciccone, Arthur Schram, Ben Greiner, Cornelius Schneider, Christina Roth, Christine Exley, Christopher Roth, David Autor, David Rand, Drazen Prelec, Edward Chang, Edward McFowland III, Florian Heine, Frank Schillbach, Gianmarco Leon, Henrik Orzen, Iavor Bojinov, Jean-Robert Tyran, Johannes Rincke, John Horton, Kai Barron, Katharina Brütt, Katherine B. Coffman, Klarita Gerxhani, Kevin Lang, Laura Gee, Lawrence Katz, Lumumba Seegars, Maria Petrova, Marina Koglowski, Mashail Malik, Max Steinhardt, Menusch Khadjavi, Mohsen Mosleh, Philip Ager, Raymond Fisman, Robert Livingston, Robin Ely, Ruben Enikolopov, Ryan Enos, Serena Does, Shaul Shavi, Summer Jackson, Theo Offerman, Thomas Buser, Thomas Graeber, Vanya Georgieva, and Zoe Cullen. The paper further benefited from discussions with attendants at several conferences and research seminars at 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), 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, HeiKaMaxY Workshop (Heidelberg), University of Mannheim, and University of Augsburg. **Funding:** Financial support by the German Research Foundation (DFG) through CRC TR224 (projects A01, A04, B02) and by the state of Baden-Wuerttemberg through bwHPC is gratefully acknowledged. **Preregistration:** The experiment's first and second stages were pre-registered on aspredicted.org (#RDPZ67, #8RRVLY). **Ethics approval:** The study obtained ethics approval from University of Mannheim's Ethics Committee (EK Mannheim 32/2021)

†University of Mannheim, Department of Economics, yulia.evsyukova@gess.uni-mannheim.de
‡University of Mannheim, Department of Economics, felix.rusche@uni-mannheim.de
§University of Mannheim, Department of Economics, mill@uni-mannheim.de, Telephone +49 621 181-1897

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*

Most jobs in the U.S. are found using information and referrals provided through informal networks (Dustmann et al., 2016; Topa, 2011). Minorities rely on job networks as much as majorities, but their networks are of lower quality, providing less information and fewer referrals (Fernandez and Fernandez-Mateo, 2006; Ioannides and Datcher Loury, 2004). This could help elucidate the worse labor market outcomes of minorities (Bayer and Charles, 2018; Coffman et al., 2021; Holzer, 1987). Yet, existing research does not explain why minorities' networks are of lower quality. Discrimination might play a pivotal role in the establishment and utilization of these networks. However, differences in existing networks may also be confounded by other factors like self-selection and pre-existing inequalities, such as neighborhood segregation and socio-economic background.

In this paper, we causally investigate if and how discrimination affects the size, composition, and information provision of the job networks of minorities. 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). LinkedIn members use the platform to advance their careers by building networks, obtaining and sharing information on job opportunities, and increasing their visibility to potential employers, which positively affects their labor market outcomes (Rajkumar et al., 2022; Wheeler et al., 2022). Our field experiment has two stages. In the first stage, we build networks of 400+ fictitious profiles. We signal race (Black or White) solely via A.I.-morphed 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 arising in the first stage and separately identify discrimination in the second stage. This enables us to study the causal effects of discrimination on the job-relevant information networks provide.

A key feature of our field experiment is that we signal race exclusively through pictures. In particular, we use A.I. to generate profile pictures that vary aspects of race inherently assigned by birth, like skin tone and facial features. We do not alter facial expression, hairstyle, clothing, and background, to minimize behavioral responses due to stereotypes (Bordalo et al., 2016). We validate our approach using an online experiment, which provides three main insights: 1) participants are not able to identify our A.I.-generated and morphed 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 like looks, authenticity, intelligence, etc.

Every profile in our experiment has a unique Black or White A.I.-generated profile picture, to ensure that results are not driven by the singular 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 only recently joined LinkedIn, their CVs represent them as young males who recently finished college and are otherwise similar to usual LinkedIn users. To ensure that race is signaled exclusively through pictures, profiles are assigned names that are both frequently used and ambiguous in terms of race (e.g., Michael). Following the literature, we additionally vary the quality of the CV across twin-pairs (Bertrand and Mullainathan, 2004).

To investigate how discrimination affects the formation of job networks, our profiles send a connection request to more than 20,000 users during the first stage of the experiment. Each user receives requests from both a Black and a White profile with CVs of equal quality and a time lag of four weeks between the requests. This experimental setup allows us to evaluate discrimination in job networks along multiple dimensions. First, we can causally identify whether race affects the size of networks as twins only differ in their race and send connection requests to an identical number of users drawn from the same subject pool. Second, we can identify who discriminates based on rich information we gather from users' publicly available CVs.

Our main result demonstrates significantly lower acceptance rates for Black than White profiles—23% versus 26%, which implies a 13% higher acceptance rate for White profiles. When we examine who discriminates, we find evidence of discriminatory behavior across diverse social groups. In fact, despite our rich set of individual-level characteristics, we find little evidence of user groups that do not discriminate against Black profiles. Nevertheless, there is substantial heterogeneity in discriminating behavior. Interestingly, men and older users show lower gaps in Black-White acceptance rates in comparison to women and younger users, respectively. Black users also discriminate, but to a lesser extent than non-Black individuals. Higher education and social status are only weakly associated with lower levels of discriminatory behavior. We also find gaps in Black-White acceptance rates across almost all U.S. states. Within states, we find larger gaps for users residing in more Republican counties. Drawing on a rich set of user characteristics, we employ a causal forest to provide additional insights on heterogeneity in discriminatory behavior (Athey et al., 2019; Wager and Athey, 2018). This yields higher Black-White acceptance rate gaps in counties with lower economic connectedness (Chetty et al., 2022a,b).

In the second stage of our experiment, we assess the informational benefits of Black versus White job networks by asking first-stage connections for advice. We mimic the use of actual job networks and request either mentorship/career advice or information regarding the application process at the company where the respective user works. Overall, 21% of messages are answered. The vast majority of responses offer valuable information, like details about the company's application process and even referrals.

Importantly, our experimental design allows us to distinguish between disparities in informational benefits resulting from gatekeeping (stage I) and discrimination in response to information requests (stage II). Before asking for advice, we swap half of the A.I.-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 profile that is Black and vice versa. This feature of our experiment allows us to evaluate how much information Black and White profiles would receive if they had access to the same networks. Importantly, we 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. We find no difference between swapped and non-swapped profiles in terms of the number of views, profile blocking, connection dissolution, or the types of responses. Overall, we find close to no discrimination in responses if Black and White profiles are given access to the same networks.

Next, we assess the expected informational benefit provided by each profile's network, taking into account the possibility of discrimination during both stages. To do so, we estimate the expected number of responses for each profile had they sent a message to their entire network. We find compelling evidence that the networks of White profiles provide substantially more informational benefits than the networks of Black profiles. Our results indicate that discrimination is mostly driven by the experiment's first stage. Based on back-of-the-envelope estimates, the observed disparity in network size and subsequent difference in information benefit corresponds to an extrapolated wage difference between Black and White users of around $200 per month.

To examine the extent to which our findings challenge or validate the prior beliefs of experts, we conducted a survey following the end of the experiment. In the survey, 269 experts in labor economics and discrimination predicted the outcomes of our experiment. The results suggest that experts do well in predicting discrimination during the first stage, but expect it to continue to a similar extent during the second stage, which we do not find. Regarding heterogeneity, we reveal that experts correctly predict Black users and those from more Democratic counties to discriminate less. They do, however, expect both men and older users to exhibit higher gaps in acceptance rates, the opposite of what we document.

Overall, our study provides causal evidence on a previously understudied mechanism through which labor market outcomes of minorities may be explained, namely the effect of discrimination on the size and information provision of minorities' job networks. Given that discrimination mostly takes place in the formation of job networks (rather than their provision of information), it also offers crucial insights into potential ways to combat inequality in labor market outcomes. Specifically, our results suggest that improving networking opportunities for Black individuals, e.g., through mentorship programs could be an effective approach. By creating inclusive environments that facilitate networking and fostering diverse connections, organizations could help mitigate the disparities in the network sizes of Black and White individuals. At the same time, the results underline the importance of reducing the role of exclusive institutions, which potentially strengthen inequalities in network formation. Such steps would help to increase informational benefits available for Black individuals and, thus, enhance equitable access to job opportunities.

This paper provides a number of new insights that expand and complement the existing literature. The key contributions of our study are as follows: first, we provide **causal** evidence on discrimination in job network formation and information provision. Even though around half of all jobs are found through informal networks (Dustmann et al., 2016; Topa, 2011), previous studies have primarily relied upon correlational analyses (Fernandez and Fernandez-Mateo, 2006; McDonald et al., 2009). Further, the first stage advances discrimination studies' methodology by applying an A.I. algorithm to signal race, allowing us to directly, precisely, and uniquely depict racial characteristics, instead of relying on noisy proxies such as names (Bertrand and Duflo, 2017; Edwards et al., 2015; Gaddis, 2017; Hum et al., 2011; Kreisman and Smith, 2023; Quillian et al., 2019). Our study also deviates from traditional correspondence studies by examining discrimination in a novel setting, with a substantially more diverse target group, where decision-making carries low costs and users may desire network diversity for informational benefits or virtue signaling (Acquisti and Fong, 2020; Angeli et al., 2023). Further, the first stage adds important insights into who discriminates. We have the key advantage of observing individuals' choices and a large range of individual-level characteristics. While classical audit and correspondence studies are typically conducted on the industry or firm-level (Kline et al., 2022), previous studies treating individuals only observe a few individual-level characteristics, such as gender or race (e.g. Block et al., 2021; Edelman et al., 2017). Finally, our results link to recent research on economic connectedness (Chetty et al., 2022a,b) by offering evidence that discrimination drives network formation and that it is related to lower county-level measures of economic connectedness.

The second stage shows that LinkedIn networks provide valuable information. It, thus, offers direct evidence for how individuals can benefit from job networks (Cullen and Perez-Truglia, 2023; Dustmann et al., 2016; Gee et al., 2017a,b) and, more specifically, LinkedIn, which have only been studied indirectly before (Rajkumar et al., 2022; Wheeler et al., 2022). Our findings reveal a significant racial disparity in network benefits, shedding light on a potential mechanism behind worse labor market outcomes of minorities. Furthermore, the paper highlights the importance of weak ties (Gee et al., 2017a,b; Granovetter, 1983; Rajkumar et al., 2022), demonstrating their ability to provide valuable insights. Our paper also highlights the need to better understand where discrimination originates, including through multi-level experiments (Bohren et al., 2022).

Our two-stage experiment shows that differences in informational benefits mostly emerge during the first stage of the experiment, i.e., gatekeeping, rather than differences in response rates. We, thus, illustrate how discrimination in network formation can proliferate at later stages, even with little direct discrimination during the second stage. The results may also be interpreted as a 'foot-in-the-door' effect, suggesting that once Black profiles are artificially given access to White networks, they face little discrimination. This underlines the importance of creating inclusive institutions and dissolving exclusive ones, such as 'old-boys-clubs' (Cullen and Perez-Truglia, 2023; Michelman et al., 2021) to lower discrimination in outcomes.

2 Contribution & Literature

This paper makes multiple contributions to several strands of literature.

Correspondence studies – Methodically Methodologically, we contribute to experimental research on discrimination. To cleanly identify the causal effect of discrimination, many studies – including ours – rely upon a correspondence study approach, sending applications to firms or landlords while varying characteristics such as race (for reviews see Bertrand and Duflo, 2017; Neumark, 2018; Quillian et al., 2019). Most studies on racial discrimination use names to indicate race, like distinctly Black or White names. However, this approach has potential drawbacks, as stereotypical Black names may convey unintended characteristics, such as lower socioeconomic background. (Bertrand and Mullainathan, 2004; Gaddis, 2017; Fryer Jr and Levitt, 2004; Kreisman and Smith, 2023) or lower skills and productivity (Abel and Burger, 2023; Kreisman and Smith, 2023), leading treated individuals to respond "[...] in a way that they might not have to a more typical Black candidate" (Doleac and Stein, 2013, p. 3). Moreover, signaling race through names is noisy as the perception of a name as Black depends both on the name and the individual evaluating it (Gaddis, 2017). Our study suggests a way to resolve these issues by using pictures as a signal of race. More specifically, we develop and validate an A.I. algorithm that changes a picture's race keeping stable other attributes like background, facial expression, emotions, and intelligence. Our algorithm solely modifies skin tone and race-specific facial features

Our approach has several advantages: first, the pictures allow us to signal race directly, rather than via a proxy. Being a noisy signal, names allow for motivated reasoning and give space for a wiggle room. Pictures resolve this uncertainty, ensuring that treatment strength is independent of recipient characteristics. Second, compared to existing studies utilizing real pictures (e.g. Kaas and Manger, 2012; Mejia and Parker, 2021), using A.I.-generated images allows us to keep picture characteristics other than race stable. Third, it enables us to create a unique image for each of the 408 profiles, thus making the results less dependent on specific image characteristics. Fourth, our approach allows us to keep names neutral regarding their race signal. Finally, using pictures is advantageous on online social media websites, where profiles without pictures are typically perceived as less credible (Edwards et al., 2015; Hum et al., 2011).

Correspondence studies – Contentually Aside from studying job networks, our setting is quite different from that of usual correspondence studies (e.g. Acquisti and Fong, 2020; Agan and Starr, 2018; Bertrand and Mullainathan, 2004; Kline et al., 2022; Kroft et al., 2013; Neumark et al., 1996). In our context, the costs and benefits of decision-makers differ significantly from those in typical job-application-correspondence studies. Unlike human resource professionals and recruiters who might seek candidates similar to the existing workforce, LinkedIn users aim to maximize the informational benefits of their networks. This means they might value diversity within their network to access a wider range of information sources, and may even engage in virtue signaling by showcasing a diverse network (Angeli et al., 2023). Consequently, our context may promote positive rather than negative discrimination in contrast to traditional settings. Next, the cost structure between the settings is very different. Recruiters and HR professionals make high-stakes hiring decisions with financial consequences. Thus, they have to make an informed decision relying on all kinds of information (including statistical) they can acquire (Acquisti and Fong, 2020; Manant et al., 2013). In contrast, our context involves a low-stakes decision of accepting or rejecting connection requests. The associated costs are minimal, mainly limited to the potential inconvenience of receiving unwanted content or messages, which could easily be reverted by dissolving and/or reporting the connection. This ensures a low-cost environment that reduces the likelihood of discrimination.

Aside from correspondence studies in the labor context, some explore discrimination on social media and online platforms. These include studies on Twitter (Ajzenman et al., 2023; Angeli et al., 2023), classified advertisements websites (Doleac and Stein, 2013), Airbnb (Edelman et al., 2017), and online Q&A websites (Bohren et al., 2019a). To the best of our knowledge, our paper is the first to perform an independent correspondence study on LinkedIn. In comparison to the existing works, we investigate the effect of discrimination in network formation and their information provision in the broad context of general professional networks. Lastly, we are able to directly test the usefulness of the online professional networks in the second stage of our experiment, estimating the informational benefits that could be obtained through them.

Finally, the targeted population is substantially more diverse than in usual correspondence studies. By design, most research involves sending applications to HR departments. Therefore, only a specific group of people makes a decision. We study around 20,000 real users with a broad range of backgrounds, professions (from self-employed individuals to CEOs), and levels of seniority (from fresh college graduates to retired individuals). Finally, by design, most correspondence studies cannot identify individual-level characteristics of decision-makers. In contrast, we have access to a rich set of publicly displayed individual-level characteristics of decision-makers, enabling us to examine correlates of discrimination.

Predictors of Discriminatory Behavior As hinted above, our study enriches the literature by shedding light on predictors of discrimination using data from 20,000 individuals. Existing correspondence studies tend to observe only firm-level data of the organizations where treated individuals are employed (Bertrand and Duflo, 2017; Kline et al., 2022). In our work, we obtain more fine-grained information, which contains individual characteristics of users, including age, gender, race, educational history, place of residence, employment history, and platform-specific variables. In line with existing research, our results indicate that targets residing in Republican areas are more likely to discriminate than those from Democratic counties (Acquisti and Fong, 2020; Block et al., 2021). We also find that Black individuals exhibit less discriminatory behavior (Block et al., 2021; Goncalves and Mello, 2021). Somewhat surprisingly, we find that females are substantially more likely to discriminate against Black profiles, with the effect being mostly driven by White females. Other studies find no gender disparity in attitudes (Hughes and Tuch, 2003), while Edelman et al. (2017) provide suggestive evidence of White women more strongly discriminating against Black men. However, it is important to note that our context is different from previous works. While LinkedIn is primarily a job-networking website, it can also be considered a form of social media where users may engage in dating and viewing connections as potential partners. The empirical literature suggests that individuals often exhibit strong racial homophily in their relationship preferences (Kalmijn, 1998; McClintock, 2010; McPherson et al., 2001). Thus, White female users may favor White profiles, possibly seeing them as potential dates, while male users might not (as all of our profiles are, on purpose, male). The findings could, however, also be linked to gender- and race-specific stereotypes about Black males (Davis, 1981; Sommerville, 1995; Zounlome et al., 2021). Another somewhat puzzling result is that older LinkedIn users tend to accept Black profiles more often than young users, which differs from the pattern observed in previous studies (Coenders and Scheepers, 1998; Davidov and Meuleman, 2012). Finally, we leverage our rich set of covariates by exploring discrimination through causal forests (Athey et al., 2019; Wager and Athey, 2018). Beyond revealing varying acceptance rate gaps across multiple dimensions, the analysis also demonstrates that almost no group is expected to have no gaps in acceptance rates, even when allowing for conditional average treatment effects across a high number of covariates. This underlines the conclusion that discrimination is extremely widespread.

Networks and Discrimination Our work is closely related to the research focusing on discrimination in the formation of job networks. Previous studies rely mostly on descriptive analyses (Fernandez and Fernandez-Mateo, 2006; McDonald, 2011) or theoretical insights (e.g. Galenianos, 2020; Mortensen and Vishwanath, 1994). As noted by Fernandez and Fernandez-Mateo (2006) and McDonald (2011), existing studies using observational data face circular causality issues in studying job networks, unable to separate race from socioeconomic factors like residential locations or socioeconomic background. Further, prior research reveals that, despite using networks to a similar extent as majorities, minorities often obtain lower-paying jobs and access less valuable information (Ioannides and Datcher Loury, 2004). This raises the question of whether it is direct discrimination or unobserved characteristics that drive Black individuals to sort into 'wrong networks' (Fernandez and Fernandez-Mateo, 2006).[1]

Our study adds to this literature in several ways. First, it offers the first causal evidence on how discrimination affects the formation and information provision of professional job networks, testing whether discrimination drives the selection of minorities into smaller or 'wrong networks'. Second, we provide direct evidence for differences in the informational benefits obtained through networks, a critical driver of labor market outcomes. The second stage closely relates to Gallen and Wasserman (2021a,b), who conduct an experiment where real university students request career advice from professionals. The authors find no gender effect on response rates, yet professionals emphasize work/life balance more strongly when interacting with female students.[2] Overall, our

[1] In addition to empirical studies a number of theoretical studies have formalized the mechanisms underlying the segregation of job networks and their effects on economic outcomes (Arrow and Borzekowski, 2004; Calvo-Armengol and Jackson, 2004; Galenianos, 2020; Mortensen and Vishwanath, 1994).

[2] More generally, the second stage also relates to the literature on discrimination in helping behavior. Previous

two-stage setup unveils that access to the 'right' professional network is crucial for receiving valuable job-related information and is affected by discrimination. Importantly, once minorities have access to the networks of majorities, they are as likely to receive helpful information and referrals as majorities. This also adds to the job and referral network literature (Dustmann et al., 2016; Pallais and Sands, 2016; Schmutte, 2015), highlighting the role of online networks (Wheeler et al., 2022), and the literature on the conditional benefits of referral networks (Beaman and Magruder, 2012; Beaman et al., 2018) by cleanly identifying the effect of race.

Social tie formation Our study fits into broader research on friendship formation, particularly in U.S. college contexts. These studies use random dorm assignments to explore factors impacting friendships, assessed via email exchanges (Mayer and Puller, 2008) or Facebook messages (Marmaros and Sacerdote, 2006). They highlight race as a significant determinant of social interaction and network segmentation.[3] More closely related, Michelman et al. (2021) and Cullen and Perez-Truglia (2023) study the career effects of social ties. Michelman et al. (2021) study Harvard's old boys clubs in the 1920s and 30s. Leveraging random allocation to dorms, they show that exposure to high-status peers increases membership and accelerates careers, though only present for those from private feeder schools, with no effect on minorities. Cullen and Perez-Truglia (2023) explore manager rotations and gender/smoking habits in a financial institution, revealing that smoking men switching to a smoking manager benefit from increased social interactions. They also note that men receive more promotions when transitioning from a female to a male manager, while this effect doesn't apply to women in the reverse situation.

We contribute to the literature in three ways. First, whereas prior research examines the role of homophily, i.e., the tendency of individuals to become friends with people with a set of similar attributes, our approach focuses specifically on discrimination by comparing network formation between individuals differing *only* in their race, cleanly isolating this single feature. Second, we shift the focus from social ties formed based on shared preferences to job-related ties that prioritize information and opportunity access. This distinction underscores our unique design setting. Third, an important distinction lies in the strength of ties. In the seminal paper, Granovetter (1983) distinguishes between 'strong ties' (friends and relatives) and 'weak ties' (colleagues and acquaintances), arguing that the latter are more helpful in the job search. The literature offers mixed findings on the usefulness of weak ties in networking (Gee et al., 2017a,b; Utz, 2016). Rajkumar et al. (2022) recently demonstrated an inverted U-shaped relationship between tie strength and job transmission, with weak ties contributing to higher job mobility. Our study causally shows that networks consisting of weak ties provide access to valuable job-related information and job referrals.

Economic connectedness and inequality of opportunity In previous research, Chetty et al. (2022a) find that high-socioeconomic status (SES) friends predict upward income mobility. Chetty

studies indicate that minorities are less likely to receive assistance (Block et al., 2021; Giulietti et al., 2019).

[3] Similarly, based on an observed high school network, Goeree et al. (2010) find that being of the same race is the strongest predictor for tie formation.

et al. (2022b) identify key factors for economic connectedness: exposure to high-SES individuals and friendship bias, both contributing to the disconnection between low- and high-SES groups. At the same time, economic connectedness decreases with a higher Black population share, limiting cross-class interaction. Our work adds causal evidence, highlighting race-based differences in the formation of job networks, revealing disadvantages for Black individuals. Our second-stage results align with Chetty et al. (2022a), indicating information disparities arising during network formation. Finally, our investigation of heterogeneity in discrimination using causal forests (Athey et al., 2019; Wager and Athey, 2018) shows that higher gaps in response rates are also associated with lower county-level economic connectedness.

3 Experimental Design

3.1 Challenges, Choices, and Ethics

This subsection discusses four main points: First, we argue why we believe a field experiment to be the right approach to causally study this research question. Second, we describe how LinkedIn is being used. Third, we discuss the ethical questions of running this field experiment. Finally, and most importantly, we describe our main choices and the target function of creating realistic profiles.

Why a Field Experiment As discussed in the introduction, previous studies using observational data show that minorities use job networks to the same extent as majorities. However, they tend to use the 'wrong' networks, i.e. more dense networks with lower quality contacts (Fernandez and Fernandez-Mateo, 2006; Ioannides and Datcher Loury, 2004). Previous studies have faced obstacles in establishing a causal link between discrimination and the formation of job networks. In particular, the finding could be driven by discriminating behavior on the labor market or during the job network formation process (Galenianos, 2020). Further, it is potentially driven by minorities selecting themselves into less advantageous networks and accelerated by pre-existing inequalities (McDonald et al., 2009). Observational data thus has a major issue of self-selection, omitted variable bias, and endogeneity. An experiment circumvents all these threats to a causal interpretation by using a random assignment to treatment. However, a perfectly controlled environment like a laboratory experiment, which ensures high internal validity, has its own drawbacks. Specifically, laboratory experiments would be very artificial in this context and, thus, have little external validity. Further, previous studies suggest that by making subjects aware of being part of an experiment and their choices being monitored, their discriminatory behavior decreases (Baker and Grimm, 2021). A field experiment offers the perfect synthesis of the two objectives: it has high external validity and ensures causal insights through sufficient internal validity.

Running such a field experiment is not without difficulty though, as LinkedIn is keen on preventing the creation of fake accounts, requiring researchers to, amongst others, circumvent captchas, use proxy servers, and bypass phone and email verifications for each account. While beneficial for users, this creates additional hurdles for researchers, which might explain the lack of other independent large-scale field experiments on LinkedIn.[4]

LinkedIn With over 199 million US users and 900 million worldwide, LinkedIn is a leading global online job networking platform (LinkedIn, 2023). Users create profiles highlighting professional experience, including work history, education, and additional customized information like 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. When users receive a request, LinkedIn contacts them via email and displays the request upon logging in. The request contains a link to the profile, the profile picture, the name, and shared similarities such as workplace, education institution, or common connections. Users' current jobs and employers are typically shown alongside the request. Users can either accept or ignore the request. If a user chooses the former, the connecting profile is added to their network. Deciding to 'ignore', she can also report the connecting profile. The user sending the request isn't explicitly notified when ignored. The users communicate with others through direct messages. However, only direct contacts can be contacted via messages.[5] While LinkedIn is not targeted toward any specific occupation, the user base primarily consists of white-collar professionals and tends to attract educated and higher-income individuals (Brooke Auxier, 2021).

Firms also use LinkedIn extensively. They create profiles to post job openings, receive applications, and use the platform for general promotion. Globally, LinkedIn (2023) reports that 58.4 million companies have profiles. Having a profile also allows (former) employees to link their profiles to the firm. Like users, firms can generate, comment on, and share content. Firms also use the platform for headhunting, directly contacting potential candidates for job openings.

Ethics In this section, we briefly discuss the main ethical aspects of our study. More specifically, we discuss the main ethical considerations in the context of correspondence studies and apply these to the context of our study (Salganik, 2019). Appendix E further includes a more thorough reflection on the ethics of running our experiment, including additional ethical aspects to consider when running a study like ours.

Salganik (2019) propose a number of conditions to consider when studying ethnic discrimination using field experiments (see p. 304-7). The first condition states that field experiments should limit harm, i.e., costs, to participants. In our context, answering two contact requests, which is likely to take seconds, does not represent a significant cost for participants. This point extends to the smaller sample of profiles (16% of those that initially receive a request to connect) whom we contact

[4]See Rajkumar et al. (2022) and Gee (2019) for two internal experiments that resulted in publicly available research papers. Further, Wheeler et al. (2022) indirectly study the effects of training people to use LinkedIn on unemployment.

[5]Users can also send a short message alongside a request to connect. To contact others, a rather costly premium account (amounting to several hundred USD per year) is required, which allows for a handful of messages to others per month.

via messages in the experiment's second stage, though these entail slightly higher costs. At the same time, all users we ask to connect can obtain a small potential benefit by having two additional contacts for a while. Overall, we consider the costs our experiment imposes to be very low, especially when compared to those of classical correspondence studies, which usually involve subjects reading, evaluating, and possibly answering an application for a job position.

The second condition put forward by Salganik (2019) states that costs should be evaluated against "the great social benefit of having a reliable measure of discrimination" (p. 304). As argued in Section 1, job networks play a decisive role in labor markets and their outcomes. Further, descriptive evidence suggests that minorities 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 (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.

Salganik (2019) suggest a third requirement, namely "the weakness of other methods of measuring discrimination". As discussed above, the use of existing data, such as representative samples, does not allow for a causal study of the effect of discrimination on job networks. Further, designing a laboratory study with externally valid results is hard to imagine. Therefore, our setting requires a field experiment.

Overall, we thus argue that the social benefits of a reliable estimate of discrimination in the formation and information provision of job networks clearly outweigh the low costs imposed on participants.

Aims and Choices Regarding Profile Creation The profiles we create have three main aims: first, given that the key outcome of interest is discrimination, they should keep everything but race constant. Second, they are supposed to look like profiles of real human users. Third, they should allow for conclusions regarding discrimination across industries and states to increase external validity. At the same time, we face two constraints. First, given the complexity of the setup and technical details, we can only create a limited number of accounts. Second, to avoid both the dynamic effects of pre-existing networks and being blocked by the platform, all profiles have to start with zero contacts. To moderate the aims and constraints, we make the following choices.

First, we signal exclusively race via pictures instead of names. This has three main advantages: (1) it makes profiles more realistic, (2) pictures are a direct signal without noise, and (3) using particularly Black names to signal race has been shown to have clear downsides, such as their association with low socioeconomic background (Gaddis, 2017; Kreisman and Smith, 2023).

Second, we only create male profiles. While this choice has some disadvantages, we do so for the following reasons: (1) given the technical constraints, we focus on varying one dimension, namely race, keeping everything else constant. Here, it is important to note that adding females would have doubled the experiment's size. (2) Responses to requests of women may follow a different logic than male requests and may require an adjusted experimental setup to interpret results.

Previous research shows different reactions to online activities by men and women (e.g. Bohren et al., 2019b). Further, especially young women report much more frequently being sexually harassed online (Vogels, 2021). Thus varying the treatment across more than one dimension would complicate the interpretation of results. For instance, Ajzenman et al. (2023) find higher follow-back rates of women in comparison to men on academic Twitter. They, however, cannot answer whether this is beneficial for women, as it may both be driven by professional and 'social' reasons. Finally, a technical issue is that the morphing of pictures is more error-prone for female pictures, and the baseline sample of Black females is relatively scarce. Thus, while we believe that studying the effects of job network formation for women is just as important as for men, we argue that interpreting and studying the effects between genders would require an adjusted experimental setup and would warrant a paper of its own.

Third, all profiles start without initial contacts. This has the potential disadvantage that users may find it odd. Even though we do not believe that this would pose any threat to our identification – as having no initial contacts is true for both our Black and White profiles – we decided to have no initial contacts for several reasons. Starting with contacts would only have been possible in two ways: first, we could have built networks before the start of the experiment. This would, however, make the setup less clean, as the network composition would differ between profiles, making dynamic effects likely. Especially, discrimination might endogenously affect results. Second, we could have created fake accounts to serve as initial contacts or connected our accounts with one another. This, however, would have substantially increased the risk of being blocked by the platform. In addition, any number of prior contacts would have been ad-hoc. Starting off with zero contacts is inherent to creating a network on any online website – thus, zero contacts seemed like a valid starting point. Finally, our setup allows us to directly compare the acceptance rate of profiles with zero or more contacts (due to the dynamic nature of our design), to test whether discrimination changes with the number of contacts.

Fourth, profiles live in the respective state's biggest city, work in large firms, have a business degree and job, and recently graduated. This is done for the following reasons. (1) We create profiles in each state's biggest city to ensure that results have external validity and we can add a sufficient number of contacts. (2) Profiles work at large firms to ensure that these are not easily identified as fake by co-workers. In large firms, it is unlikely that people know all their co-workers. (3) Profiles are assigned a business degree. This is done, given that any type of firm, from hospitals to steel plants, employs individuals with business degrees. It ensures that we do not have to focus our analysis on firms in specific industries. Further, a business major is, by far, the most popular major among US college graduates (Niche, 2019). (4) Our profiles have recently graduated, which makes it seem more realistic that these have only recently joined the platform, and, therefore, have no contacts, yet.

3.2 Profile Design

Our aim is to create realistic profiles, keeping the quality of profiles constant both within quality conditions and across treatment areas. Regarding the latter, we run the experiment across all 50 U.S. states plus Washington D.C. We do so for two reasons: First, it increases the external validity when compared to running studies within a single city. Second, this allows us to conduct heterogeneity analyses at the county/state level. More specifically, we choose to run our experiment within the most populous city in each state.[6] This is done to ensure that profiles remain anonymous and that there is a sufficient number of users available for us to add. The treatment cities are listed in Table A.1 and displayed in Figure 2.

Basic profile features Each of the profiles represents a male user born in the late nineties, who has recently graduated with a bachelor's degree in Business Administration and just started his first job. The beginning of one's career seems to be a reasonable time to start developing their professional network, which helps to explain why our profiles do not have any contacts on the website yet. Creating profiles that are more advanced in their career, might have caused suspicion: one would expect such users not just to start their website usage but rather to have well-developed online networks.

Jobs Each profile pair is assigned one of five job titles. These are selected without replacement and include 'Buyer', 'Office Manager', 'Administrative Assistant', 'Marketing Assistant', and 'Office Administrator'. The job titles are obtained from `payscale.com` by searching for jobs for bachelor graduates of business administration. Titles are chosen, given their generality, i.e., almost every firm could employ someone with the titles above. All titles are comparable regarding their skill level with an average salary between 38 000 and 48 000 dollars, according to Payscale (see Table A.6). What is also important, is that the job positions are accessible for those starting their careers and can be occupied by graduates with varying educational backgrounds. To fill profiles with information, we further randomly assign each profile pair one job description, as shown in Table A.7.

Education In line with existing studies (Oreopoulus, 2011), we vary the university the profile attended to signal profile quality. Within each city, four profiles are assigned to a low and four to a high-quality condition. To ensure that educational quality is comparable across states while avoiding adding additional signals through an out-of-state education and experience, we refrain from assigning top universities, such as Harvard. Rather, to choose universities offering business degrees, we draw on `Niche.com`'s 2022 ranking of the 557 "Best Colleges for Business in America". Within each state, we identify one low and one high-ranked university: for high types, we choose a university ranked 70-270; for low-quality profiles, we assign a university not included in Niche's

[6] The only exception is Florida where we use Miami instead of Jacksonville.

ranking but present in the category "Business and Management".[7] Some states do not offer a suitable university. If this is the case, universities for both types are chosen from a neighboring state.[8] Tables A.2 and A.3 list all universities, their states, rankings, and student population.

We verify the quality signal in two ways: first, we ensure that the Niche ranking is consistent with other popular rankings. In the Forbes and USA Today rankings, high types have average ranks of 177 and 116, respectively. Only few universities of low type are included in the rankings (25 and 7 of those, respectively), all being ranked below 444 and 228 in the corresponding ranking. Second, we verify the perceived ranking by conducting a survey among a convenience sample of Americans (n=500). Here, we ask individuals from the U.S. to identify the better-ranked universities (see Appendix F). On average, subjects are able to correctly differentiate between the high- and low-ranked universities, suggesting that the chosen universities convey the desired signal.

Employer Next, an employer is assigned to each profile. For this, we draw upon Statista's Company Database to obtain the biggest employers in the U.S. We identify the 10 largest companies in each selected city and randomly assign one of these to each profile.[9] We use large corporations as employers, as this makes it less likely that 'coworkers' will encounter our profiles and realize that these are fake. Moreover, large employers are likely to have sufficient turn-around in workers, to remain anonymous. Large employers are also more likely to have hired a recent graduate and are very likely to hire business-related workers.

Names Each profile is assigned a name. To avoid potential drawbacks of signaling race via names, both first and last names of the profiles are chosen to be race-neutral. To obtain such first names, we rely upon the most common first names of males born in 1997 in Georgia. We focus on names that appear among the 50 most common names for both White and Black males (i.e., the intersection of popular Black and popular White names). We sort these remaining names by the relative popularity among Black Americans and take the 10 most popular names. Table A.4 provides an overview of all names and their popularity. It also shows the rank of the first name for all baby names across the U.S. in 1997. All names are among the top 30 names. For last names, we draw on race shares by last names from U.S. Census Bureau (2022) and choose names that are roughly equally likely to be of a Black and White individual and unlikely to be of any other race. We further choose names that are relatively common. Table A.5 shows the names, race shares, and rank of each name across the US.

[7]We control for some additional characteristics. We exclude universities with an undergraduate enrollment of less than 1000. Since both Black and White profiles within each condition are assigned to the same university, we also exclude universities with outlying shares of Black or White students as well as historically Black colleges and universities.

[8]If a state has several suitable 'neighbors', we proceed by selecting the university closest to the biggest city in the target state, choosing among universities that are second-best ranked.

[9]For cities with too few employers or cities with multiple mentions (e.g., Charleston), we searched for local information on the largest employers.

Additional details on profiles To make profiles more realistic, we specify additional details. In particular, the website 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. We proceed by randomly assigning five skills of the 20 skills relevant to the specific job to each profile. Table A.8 lists the relevant skills by job.

To further fill the profile with information, we also assign each profile past volunteering experience. We choose organizations that are very popular in the U.S., are non-partisan, and are present throughout the country: 'Big Brother and Sister', 'Red Cross', and 'Crisis Text Line'. All these organizations are available within or close to the biggest city in a given state or can be done remotely ensuring that we don't have idiosyncratic results due to very specific volunteering experience. Further, all these experiences do not require special skills to ensure that no differential information is signaled between profiles. Table A.9 provides an overview of the volunteering experience. It also includes descriptions of the tasks, which we created based on real profiles.

Using the process described above, we create 8 profiles in each U.S. state (and Washington D.C.). Half of the profiles are high (i.e., attended a better university) and half are low types (i.e., attended a worse university). In addition, within each quality condition, there are two 'twin pairs', i.e., pairs of profiles that have the same CV (except for differing, but race-neutral names) with profile pictures that keep all characteristics other than race constant.

3.3 Creating and Validating Pictures to Vary Race

To signal race, we create A.I.-generated pictures and develop an algorithm that transforms the pictures' race while holding other characteristics stable. Appendix B describes the picture creation process and the transformation algorithm in much detail. The transformation algorithm has two important features to account for ethical concerns: first, all pictures are A.I.-generated, thus avoiding any privacy issues. Second and most important, we do not define race characteristics ourselves, which would be highly problematic. Rather, we take an agnostic approach. Shortly, the transformation algorithm is defined as follows: we take all images of young Black males we could find among the 100k images provided by StyleGAN2 (Karras et al., 2020).[10] We translate the images into multidimensional vectors and do so for a comparable number of White images. Next, we calculate the average Black and White image vector and take their multidimensional vector difference. The transformation algorithm then simply works by adding this difference vector to a White image or subtracting it from a Black one.

Using this approach also allows us to account for two concerns: first, we provide each twin pair with a unique input image, which is then transformed into the other race. This ensures that the results are not driven by specific pictures' characteristics. Second, half of the input images are Black and half are White. This guarantees that the results are not due to any bias introduced by the transformation algorithm.

[10] StyleGAN2 is a machine-learning model that generates highly realistic images by combining characteristics from different source images.

We conduct an experiment on Amazon MTurk (n=500) to validate the pictures along a number of dimensions.[11] In the first step, we test whether participants perceive them to depict real humans rather than computer-generated ones. To do so, the participants are shown 20 pictures in a style that resembles a Google Captcha (see Figure F.1). They are told to select all computer-generated images and are provided a monetary incentive to click on the correct ones. Among the pictures shown, ten are our A.I-generated images, while another six pictures depict real humans. The real human pictures are chosen from the set of training images of StyleGAN2 and they are chosen to fit the age, race, and gender category of our own images. An additional four pictures show obviously computer-generated pictures, i.e., images with 'weird' 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%) (see Figure F.4 in the Appendix). This goes in line with a recent study by Nightingale and Farid (2022), which suggests that good A.I.-generated pictures are indistinguishable from real faces.

Following this exercise, each participant rates 10 of our pictures along a number of dimensions (see Appendix F). Most importantly, both Black and White pictures are associated with the targeted race and gender (see Figure F.5 in the Appendix). In addition, our results suggest that pictures are rated similarly across a number of additional dimensions, including trust, appearance, authenticity, and intelligence. While we did not expect the same scores in these categories, given potential biases of participants, the results provide reassuring evidence that our algorithm keeps pictures' characteristics rather stable.

Nevertheless, we restricted our initial sample of 700 potential profile pictures, based on the survey, to those profile pictures with the lowest difference between the Black and White twin pictures. Thus, in the final creation of profiles, we used 408 pictures with the smallest difference between the Black and White profile pictures.

3.4 First Stage – Network creation

To create networks, we send contact requests to website users from each of the fictitious profiles (pre-registration: #RDPZ67).[12] The first stage has two aims: first, we aim to measure differences in the network size between Black and White profiles. Second, we aim to draw on detailed information on individual users to identify which user characteristics are most predictive of discriminating behavior. The timeline of our experiment is shown in Figure A.1.

[11] Before doing so, we go through the pre-selected images by hand to manually remove small issues, such as earrings or minimal but frequent deformations of the earlobes.

[12] We initially planned to exclude LinkedIn users without profile pictures from the list of potential targets as we anticipated only few instances of users without pictures. However, during the construction of the target pool, it became apparent that a significant number lacked profile pictures. This posed a methodological challenge, as excluding them would 1) severely shrink our target pool, and 2) induce potential selection bias in the sample. At the same time, the algorithm inferring demographic features from the profile pictures turned out to be less reliable than anticipated, which is why we decided to instead mainly rely on first and last names to infer gender and race respectively. Therefore, we include accounts without pictures, but excluded individuals without a first name from the sample.

Creation of Profiles First, we create our profiles. More specifically, we create eight profiles within each U.S. state and Washington DC. Of those profiles, half are Black and half are White. Further, half are of a high and half of a low-quality condition, i.e., having visited a higher or lower-ranked university. Most importantly, profiles are created in pairs, i.e., we always create two profiles with the same CV. These profiles only differ in their race (as signaled via the profile picture) and name (which does not signal race).

Collection of Targets Next, we need to identify relevant connections with whom to connect, i.e., 'targets'. To do so, we collect roughly 150 contacts recommended by LinkedIn with each of our profiles. Drawing on these initial platform suggestions rather than, e.g., a random sample of all LinkedIn users in the US, has two advantages: first, they tend to be geographically relevant, i.e., live close to our profiles. Second, they are professionally relevant, i.e., work in a similar industry or have a related job, have visited the same university, or obtained a similar degree.[13] After collecting the initial suggestions, we pool all of them by state (i.e., over all eight profiles) and identify their race and gender based on their profile pictures and names.[14] We then draw on these characteristics to create four exclusive pools per state with 96 targets each. Across all pools, we balance on gender to ensure that half of the targets are female. We further balance on race shares across pools to get sufficient data on the behavior of minorities. As a result, for each state, we obtain four balanced pools (roughly 50-50 balance on gender, and 70-30 on White vs. non-White targets) with 96 targets each. Further, given that we randomly assign targets to one of the pools, the target characteristics are comparable in expectation. Thus, we randomly allocate the initial suggestions of LinkedIn to our profiles, resolving any endogeneity resulting from the initial suggestions of the algorithm.

Sending Connection Requests Next, we use our profiles to connect to targets. More specifically, each target receives two requests: one from a Black and one from a White user. Combining this with target characteristics then allows us to measure which target characteristics are associated with gaps in acceptance rates between Black and White requests. While we would ideally want to contact each target with two profiles that only differ in their race, this would likely raise suspicion, given the profiles' similarities. This is especially true with respect to their profile pictures, which keep stable everything but race.

To reduce suspicion while still ensuring that targets receive two requests each, one from a Black and one from a White profile, we proceed as follows: we generate two distinct groups of profiles for each "quality" condition, for simplicity, we call them "A" and "B" (for example, see the left and right group in Figure 1). Within each group, we create 'twin pairs'. As discussed above, a twin pair consists of a Black and a White profile with the same CV. Both twins *only* differ in terms of their race, as signaled via the A.I. generated profile picture, which keeps other facial characteristics

[13] Moreover, we are not aware of any way of obtaining a truly random sample of LinkedIn users as we would essentially always rely on some sorting of profiles in suggestions.

[14] This is done based on U.S. census and social security data on first and last names (SSA, 2022; U.S. Census Bureau, 2022) using the *predictrace* package in R (Kaplan, 2022). For pictures we use DeepFace, a face recognition software (Taigman et al., 2014).

stable. While they also have a different name, names are randomly assigned and do not signal race. The second twin pair in the other group differs in all other aspects of their CV, including their job, firm, picture, and other information. The only characteristic that all profiles of groups A and B for a given quality condition have in common is the university they attended. Thus, the twins in group A are identical except for their pictures, and they differ from the twins in group B in most aspects. Statistically, however, the twins in group A are not distinguishable from the twins in group B, as their characteristics are drawn from the same distributions of characteristics. Figure 1 shows two twin pairs of the same quality condition. We replicate the same procedure for the other quality condition so that we have 2 (quality conditions) x 2 (race of profiles) x 2 (set of twins) = 8 profiles per city.

To obtain repeated observations per target, we then proceed by sending each target a request to connect from one Black and one White profile. Both requests are from the same quality condition but stem from profiles from a different race and twin pair (i.e., one request from a group-A-twin and one request from a group-B-twin). For example, in Figure 1, a target would receive a request from Joshua and James. Hence, each target is contacted by a White and a Black profile, who are sufficiently different.

We generally create four balanced pools of 96 targets within each state, two of which are shown in Figure 1. All targets in the first pool will receive requests from James and Joshua, while those in the second one receive requests from Michael and Tyler. This ensures two things: first, given that the pools are balanced and randomized, both twins contact people who are, in expectation, the same. This allows us to account for twin-fixed effects, keeping everything but race stable. Second, contacting targets by two profiles of the same quality condition in combination with information on target characteristics allows us to draw conclusions on who discriminates. More specifically, we can observe which characteristics predict a higher acceptance rate gap between Black and White requests.

A final issue is that receiving two requests from unfamiliar accounts may raise suspicion. To mitigate this, each profile only contacts a subset of 12 targets each week, running the experiment over a period of eight weeks. While both profiles contacting the same pool send the same number of requests each week, a given target thus receives the second request with a lag of 4 weeks. This has several additional advantages: first, it allows us to study the dynamics of discrimination, second, sending a limited number of requests per week reduces the chances of our accounts being blocked, third, it reduces the chance of targets realizing they have been contacted by a similar account before.

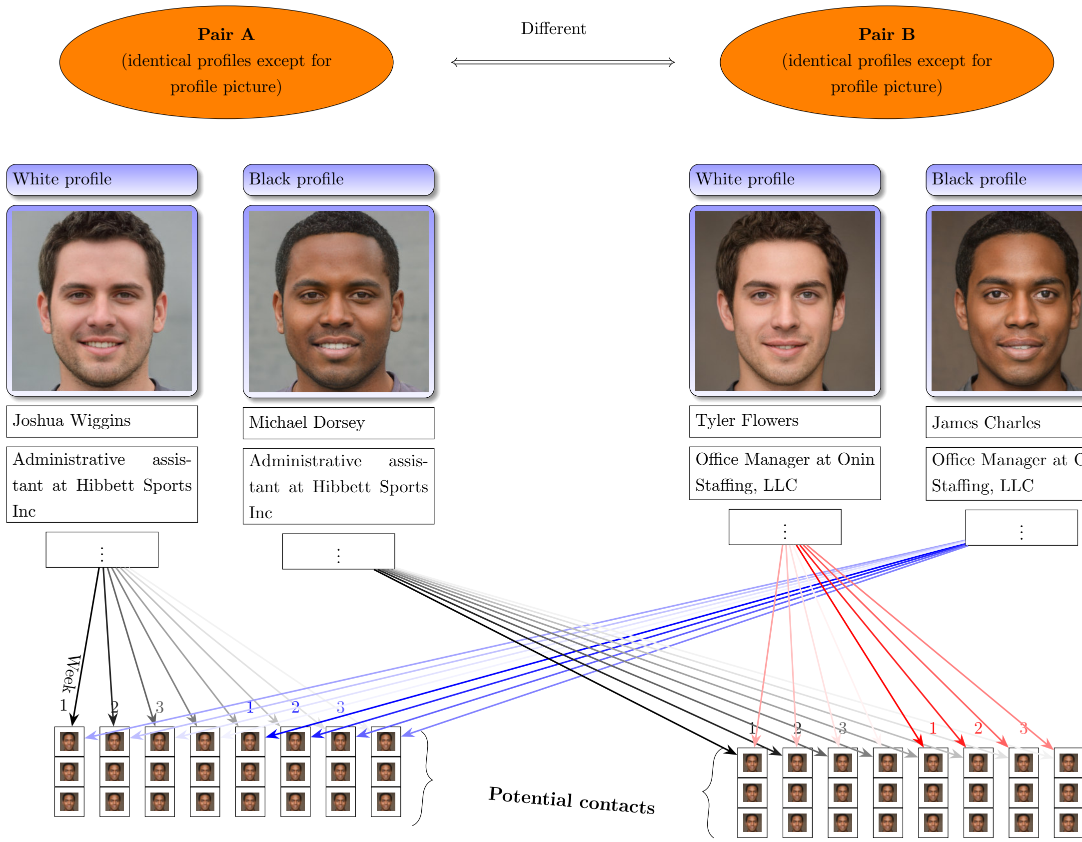

Figure 1: Requesting procedure

3.5 Second Stage – Information provision

After finishing the first stage of the experiment, we test the usefulness of the resulting networks in the second stage of our experiment (pre-registration: #8RRVLY). In particular, we examine whether valuable information can be obtained through online networks and whether the informational benefit provided by 'Black' and 'White' networks differs. LinkedIn allows users to contact each other through private messages. It is not uncommon to use messages to seek out job-related information. We use such messages to elicit the informational benefit provided by the networks. More specifically, each profile sends messages to the members of its network. Depending on the treatment condition, target users are asked for job-application advice for the company they work at or for career advice in general (see Appendix A.8 for the messages).

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 people who accepted at least one contact request of our profiles by the end of the first stage (i.e., by July 26, 2022). Further, we needed to ensure that the selected targets satisfied certain criteria. First, since we ask questions about the job application process in one treatment, 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. Second, in order to not raise suspicion of the targets, we exclude users who work in companies with less than 50 employees. For such small companies, it is likely that no relevant positions are vacant, which the target could be aware of given the company size. This, in turn, might make her perceive our message as generic or fictitious, biasing the results.

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 allocated to be contacted by either the Black or the White profile. After allocating the targets, we exclude all users who work in the same company as our profile. We also exclude individuals who sent a message to our profile before the beginning of the second stage, as messaging such users without answering their previous message might be perceived as rude or suspicious behavior, potentially, introducing bias to the results. Each of our (still active) 400 profiles[15] is contacting up to 10 unique suitable targets from her network, with each target being contacted only once. 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 (only one profile had less than 10 connections).

Conditions Each of the targets is assigned to one of the two conditions. In the first treatment, the subject is asked to provide information about the company she works at as well as for advice regarding the interview process (job application message). In the second treatment, the target is asked for career advice (mentorship message). Both treatments are randomized on the level of our profiles, meaning around half of the contacts that receive a message fall into either treatment group. The messages are displayed in Appendix A.8.

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 from the experiment's first stage and (2) differences in response rates towards Black and White profiles' messages. Ideally, however, we would want to ensure that Black profiles have connections in their network who typically would rather accept White profiles and not Black profiles, i.e., eliminating the differences from (1) and only observing differences due to (2).

In order to disentangle these effects, we draw upon a feature of our experiment, namely the fact that twin pairs only differ in their race signaled through their picture. Instead of matching on observables, we rather "input" Black profiles into the network of people who would typically accept a White profile and vice versa. We can 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

[15] During the first stage of the experiment, 8 of our profiles were blocked.

White. We do so two weeks after the end of stage one for 200 out of 400 remaining profiles (i.e., 200 profiles retain their original profile picture). Using this approach results in half of our Black profiles now having a White network and vice versa. Similarly, some users who accepted a Black connection request now have a White contact. This 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 for us to explicitly study discrimination during the experiment's second stage. 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. More specifically, it allows us to estimate the expected total number of messages obtained as a result of first and second stage results. Third, half of our profiles remain in their original networks. This allows us to study 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 Appendix G.6 for a detailed discussion).

3.6 Data Preparation

In addition to variables directly obtained via the experiment, such as whether a target accepted a request or answered a message, we obtain some additional data on targets. In particular, we download targets' public CVs just before sending the request to connect. This allows us to derive information on platform-specific variables, such as their contact count. We further structure the information and connect it to a rich set of covariates from other data sets. This chapter discusses the main sources of information on targets shown in Table C.2 and briefly discusses summary statistics in Table C.1. Appendix C provides information on the precise process of data preparation and a more detailed discussion, as well as a comparison to official statistics and users of professional job networking sites.

Regarding demographics, we draw on targets' first names and official data on gender shares of first names to obtain information on their gender (SSA, 2022). Similarly, we obtain information on their race through U.S. Census data on race shares by last name (U.S. Census Bureau, 2022).

Moving to education, we connect targets' latest education to official college statistics from IPEDS (2022) and the Forbes (2021) ranking of the 600 top colleges in the US. We also classify the college degree and use this information to estimate age.

For information regarding job and employment status, we mainly draw on information from the platform. First, we classify employment status and whether the target works in human resources based on her most recent job title. Next, we draw on rich information from linked employer sites on the platform for information regarding firms' employees, open job positions, etc.

We also obtain information on targets' salaries based on their job titles. Overall, targets list more than 10,000 unique titles. We use Google to search for the closest match for each title on `glassdoor.com`. Glassdoor provides salary estimates for specific job titles based on millions of reported salaries.[16] Table D.2 provides summary statistics of salaries by education, age, gender,

[16]Around 5% of the estimates stem from `payscale.com`, a very similar service.

and job title.

Finally, to obtain information on a target's location and surroundings, we geolocate self-reported locations using *Google Maps API* and match these with county shapefiles (U.S. Census Bureau, 2020), as shown in Figure 2. Using counties, we connect targets to local vote shares (MIT Election Data and Science Lab, 2018) and county-level demographics from the Hopkins Population Center (2020). Finally, we also include local measures of social capital based on Chetty et al. (2022a) and Chetty et al. (2022b) as well as average county-level race IAT scores (Xu et al., 2022).

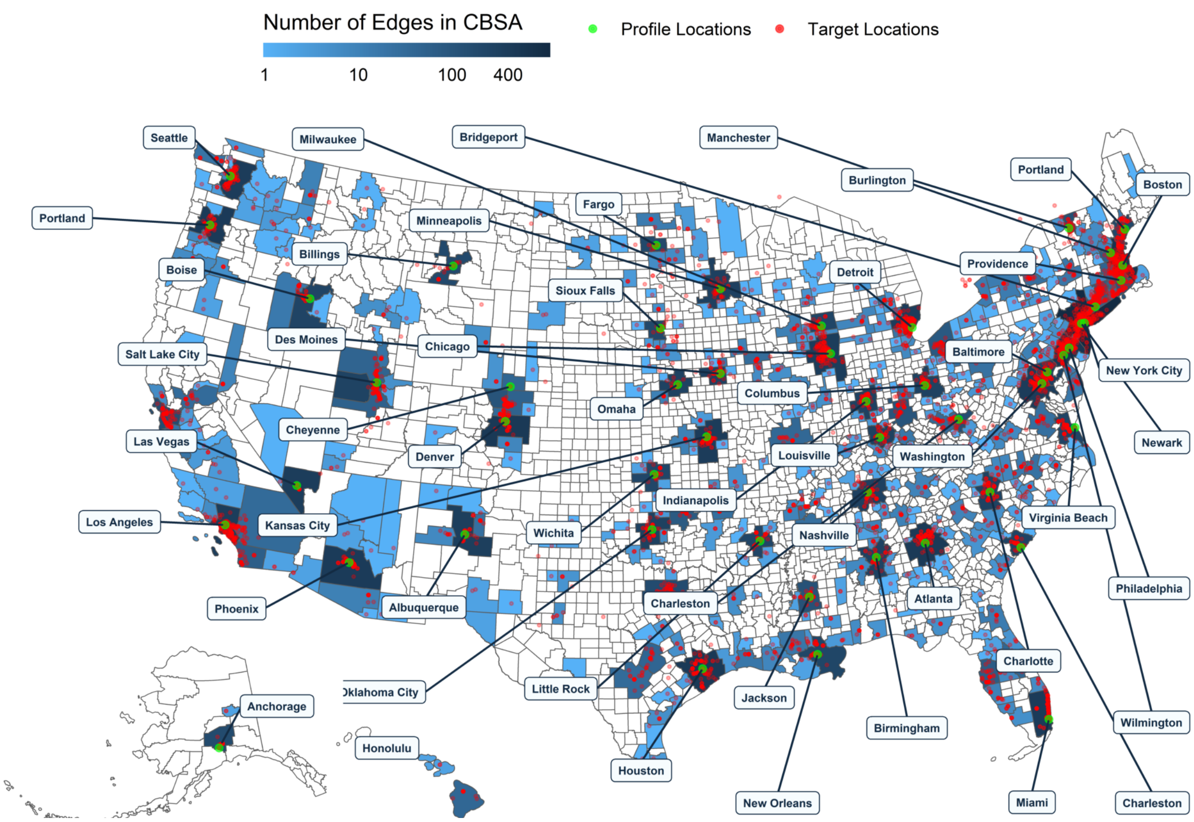

Figure 2: Locations of profiles and targets.

Note: Profile locations and city names show the cities in which our profiles indicate to reside. 'Target locations' represent unique geolocated locations using Google Maps API based on self-reported locations of targets. Each location include one or multiple targets. The figure further includes number of targets by Core Based Statistical Area (CBSA). In case a given county does not belong to a CBSA county borders are displayed instead.

4 Results

The main goal of this paper is to study whether discrimination is present in the formation of job networks and, if so, which consequences it has on the informational benefits provided by the network. To answer these questions, we split the results into two parts. First, we discuss whether discrimination is present in the formation of job networks and investigate heterogeneity in discriminatory behavior. Next, we focus on informational benefits and differentiate between differences in informational benefits due to gatekeeping (stage I) and differences in responses during stage II.

4.1 Formation of job networks

Difference in the number of contacts The main question to be answered in this section is whether the formation of job networks differs between Black and White profiles. As all of them share the same observable characteristics but differ only in the racial signals conveyed through their picture, we can causally identify the impact of discrimination on the formation of job networks.

Figure 3 displays the difference between Black and White profiles in terms of the number of contacts at the end of the experiment's first stage. Several insights can be taken from the figure. First, we see a clear difference in the number of contacts. White profiles have about 3 more connections than Black profiles, which is a considerable difference given a baseline of about 23 connections. Overall, White profiles have approximately 13% more contacts than Black profiles. This difference is roughly the same for both types of profile qualities.

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 its twin of the other race. Raw data provides two insights: first, there is substantial heterogeneity regarding the number of connections obtained by a given profile. Second, most lines are upwards-sloping, suggesting that most White twins have more contacts than their Black counterparts.[17] This is despite the fact that the only difference between the two profiles is their race, as assigned by the profile picture. All these observations are confirmed using common regressions reported in Table 2.

[17]Table J.1 in Appendix J reports upon this estimated difference while accounting for all kinds of profile characteristics (including the rated looks, trust, etc.). We again find that White profiles clearly have the edge over Black profiles under all specifications.

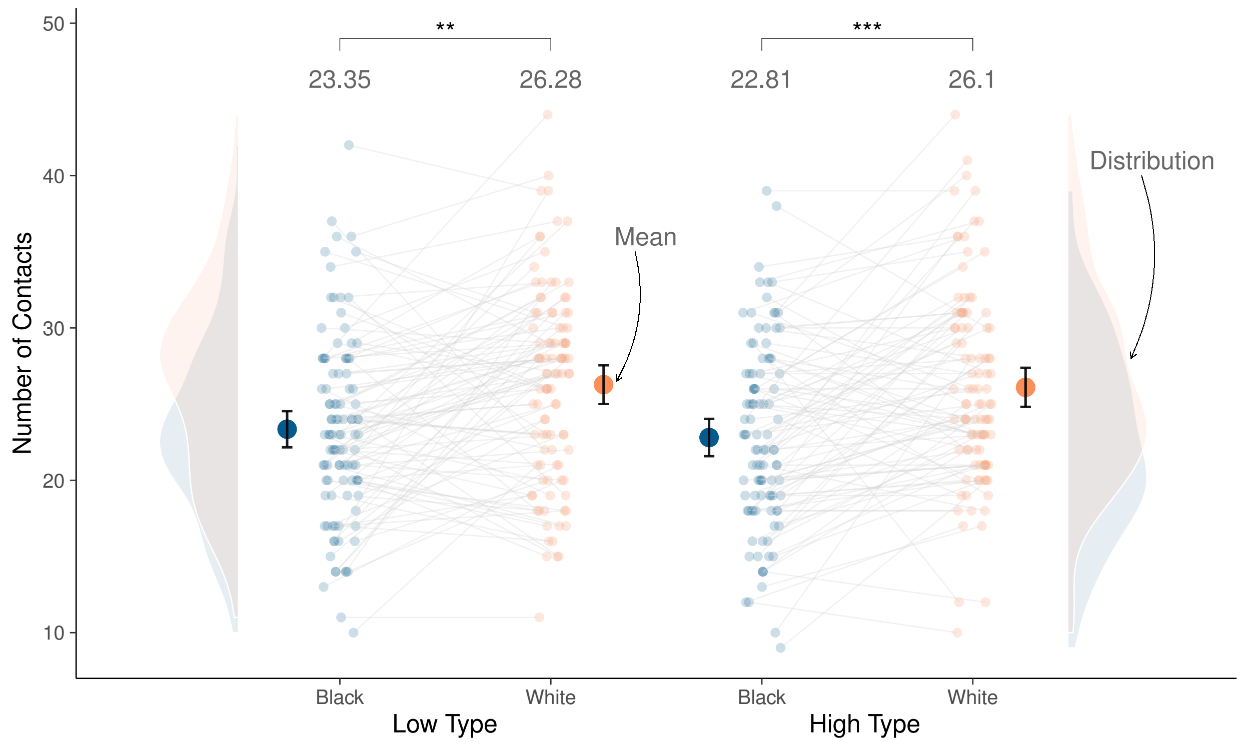

Figure 3: Number of contacts by the end of the experiment by race and quality of the profile.
The figure depicts the number of contacts obtained by Black and White profiles individually at the end of the experiment. The left panel displays the results for profiles from lower-ranked universities, while 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;*p<0.05;**p<0.01;***p<0.001.

One advantage of our design is the possibility of studying dynamic effects and geographical differences. 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. This suggests that White profiles are not perpetually improving, but Black profiles are also not able to catch up over time (see Appendix G.1 for more details). In terms of geographical variation on the state level, we observe that Black profiles' disadvantage is rather stable across space (see Appendix G.2 for more detail). Overall, discrimination faced by Black profiles is instant, stable, and geographically ubiquitous.

> **Result 1a:** We find a substantial gap in the number of connections. White profiles have 13% more contacts than Black profiles.

Difference in networks: Aside from the substantial gap in the number of connections, do Black and White profiles' networks also differ in their composition? Table J.5 reports upon multiple characteristics of the resulting networks.[18] Networks do not differ substantially in their structure. However, it's important to highlight two notable differences. Firstly, the gender composition of new connections in the network varies between Black and White networks. Specifically, Black networks have a slightly higher fraction of males. Second, the connections of Black profiles are more engaged as they have slightly more contacts themselves, have more followers, and are more active on the platform in terms of posting and sharing than the contacts of White users. Thus, although networks are comparable, there is a clear difference in their composition.[19]

Result 1b: The composition of the networks of Black and White profiles differs.

4.2 Who is (most) discriminating?

The above results raise the question of who discriminates. Our design offers a major advantage over most correspondence studies, as it allows us to obtain a rich set of characteristics from our targets while simultaneously treating each target twice – once with a connection request from a Black profile and once from a White profile. As such, we are able to investigate which characteristics are most predictive of discriminatory behavior, i.e., higher gaps in acceptance rates.

Given the vast number of characteristics, we proceed as follows: 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.[20] In a second step, we then explore additional heterogeneity in discriminatory behavior, applying methods proposed by Wager and Athey (2018).

Heterogeneity based on key demographics 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,[21] and (2) whether the residualized income of the target is above the sample's median

[18] In Appendix G.3, we further discuss predictors of connection request acceptance in detail and depict them in Figure G.3.

[19] In Section G.7, we explore additional outcomes related to the value of profiles' networks. More specifically, we test whether the number of unsolicited messages, connection requests, and views received (within 90 days prior to the end of the first stage) differ between Black and White profiles. Our results show that Black profiles have around 20% fewer profile visits and receive fewer unsolicited messages than White profiles. This is suggestive of Black profiles being less visible to other network members and potential employers. However, given that we only observe the number of profile visits once after the first stage, these results should be interpreted with some caution.

[20] The 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.

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

income. We residualize income by running a regression of log income on an individual's age, gender, race, and education.[22]

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 4, which plots the coefficients of the interaction term only.[23] Several of the correlates are in line with what one might expect. In particular, we find that Black individuals are less likely to discriminate, although they still do in absolute terms, as shown in Figure I.2. Given the low share of Black individuals among targets, the confidence interval is rather large, though. Appendix G.4.2 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).[24] 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 females discriminate significantly *more* than males and that older targets discriminate significantly *less* than younger ones. We investigate both in more detail in Appendix G.4.2. 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 explanation for the observed pattern. An alternative explanation might be stereotypes against Black males specifically held by or salient for White females (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.

[22] More 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 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.

[23] See Figure I.2 for a visualization of the full gap.

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

Moving to age, we document a substantially higher level of discrimination for young individuals. In Appendix G.4.2, 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.[25] However, the results are not very strong, which suggests that education is only weakly associated with lower levels of discrimination.[26]

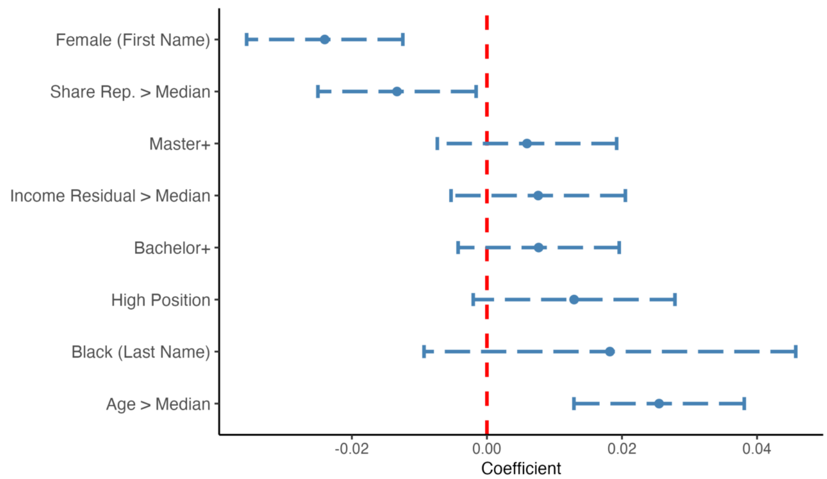

Figure 4: Correlates of discrimination.

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.

[25] Figure I.3 shows separate effects for each type of degree.

[26] The 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 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. The results are shown in Figure I.1. Overall, these are very similar to those before, with all coefficients going in the same direction.

Result 1c: There are multiple correlates of discrimination. Two of the most unexpected ones show that males and older users discriminate less.

Exploring Discrimination In order to make use of the rich data on targets, we further explore heterogeneous treatment effects using causal forests (Athey et al., 2019; Wager and Athey, 2018).[27] Following Athey and Wager (2019), we start by estimating a training forest on 18 variables with high coverage.[28] Appendix G.4.1 provides further details on the application.

To get a first idea of treatment effect heterogeneity, Figure 5 provides the distribution of conditional average treatment effects (CATE) based on the causal forest using all input variables. Two things are immediately visible: first, despite the high number of input variables, only around 8% of observations are predicted to have a CATE below zero. This can be interpreted as follows: based on all included covariates, only 8% of users are predicted to treat Black users more favorably than White users. Given that covariates cover a range of professional, personal, and geographic characteristics, we interpret this as additional evidence for discrimination being very widespread. This does not suggest that 92% of individuals do discriminate. Rather, it indicates that while not everyone discriminates, discrimination is also not concentrated in singular groups. Intuitively, this suggests that even if we had only focused our study on a specific subgroup of targets, we would have found a gap in acceptance rates in most cases. Second, the graph provides suggestive evidence of heterogeneity in the CATE, the presence of which we validate in Appendix G.4.1.[29]

[27] Given that the method remains fairly new, the analysis should be interpreted with some caution.

[28] Here, we include variables related to individuals' demographics (age, gender, race, and whether the profile signals its gender, e.g., "she/her"), education (Bachelor or above), and job position (holds a senior position, works in human resources). Further, we include variables on individuals' LinkedIn use (number of contacts, number of skills listed, number of skill verifications by other users, and number of posts). Additionally, county-level covariates are included: Republican vote share, share of Black population, local level segregation between Black and White population, Economic Connectedness index based on Chetty et al. (2022a), and average race IAT estimates by county (Xu et al., 2022)). Finally, we include dummies for similarities with our profiles, including having visited the same university or working at the same firm.

[29] We analyze this heterogeneity across covariates in Table G.1, where we plot averages of covariates by CATE tile, e.g., the share of females among those with the highest predicted CATE based on covariates (Athey et al., 2020).

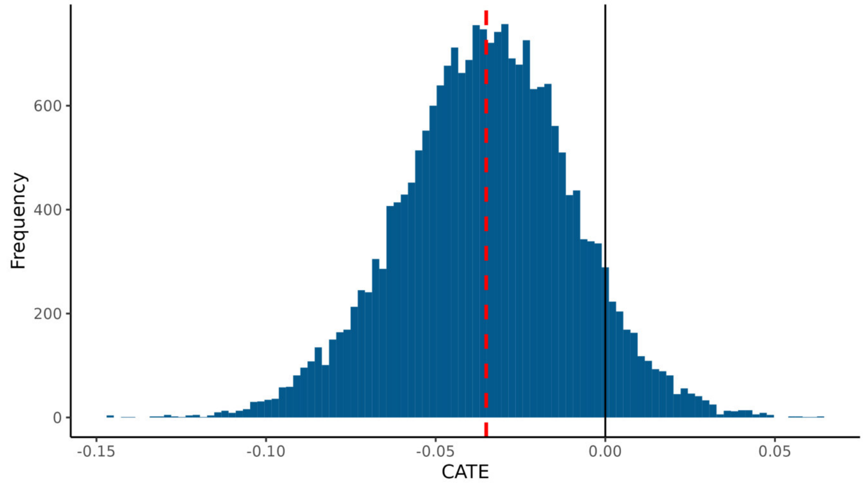

Figure 5: Distribution of Conditional Average Treatment Effects based on Causal Forest

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).

Following Athey and Wager (2019), we restrict our attention to variables with above-average variable importance, i.e., those that are responsible for the highest share of splits when building trees. Focusing on nine variables with above-median variable importance, we develop a second causal forest.[30]

Figure 6 plots the average of each variable against the quantile of the predicted treatment effect. The prediction suggests that the lowest quantiles exhibit the strongest discrimination. The figure is ordered by the variable importance, suggesting that the probability of the first name being female is responsible for the highest share of splits. The results for age, gender, Republican vote share, and Black are in line with our findings above. Further, a number of county-level variables are among the variables with the highest variable importance. First, a county's share of the Black population shows a U-shaped relationship, suggesting higher shares 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 of this being, 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, two strong predictors of lower CATE are a higher number of skills and skill verifications obtained via LinkedIn. Both suggest substantially lower levels of discrimination.

[30]Note that Athey et al. (2019) show good performance of Causal Forests for eight variables or less. As noted by Chernozhukov et al. (2018), the method only produces a consistent estimator for $num\_covariates < log(n)$, where $log(n) \approx 10$ in our study.

Result 1d: Discrimination is very widespread, both geographically and across individuals with different characteristics.

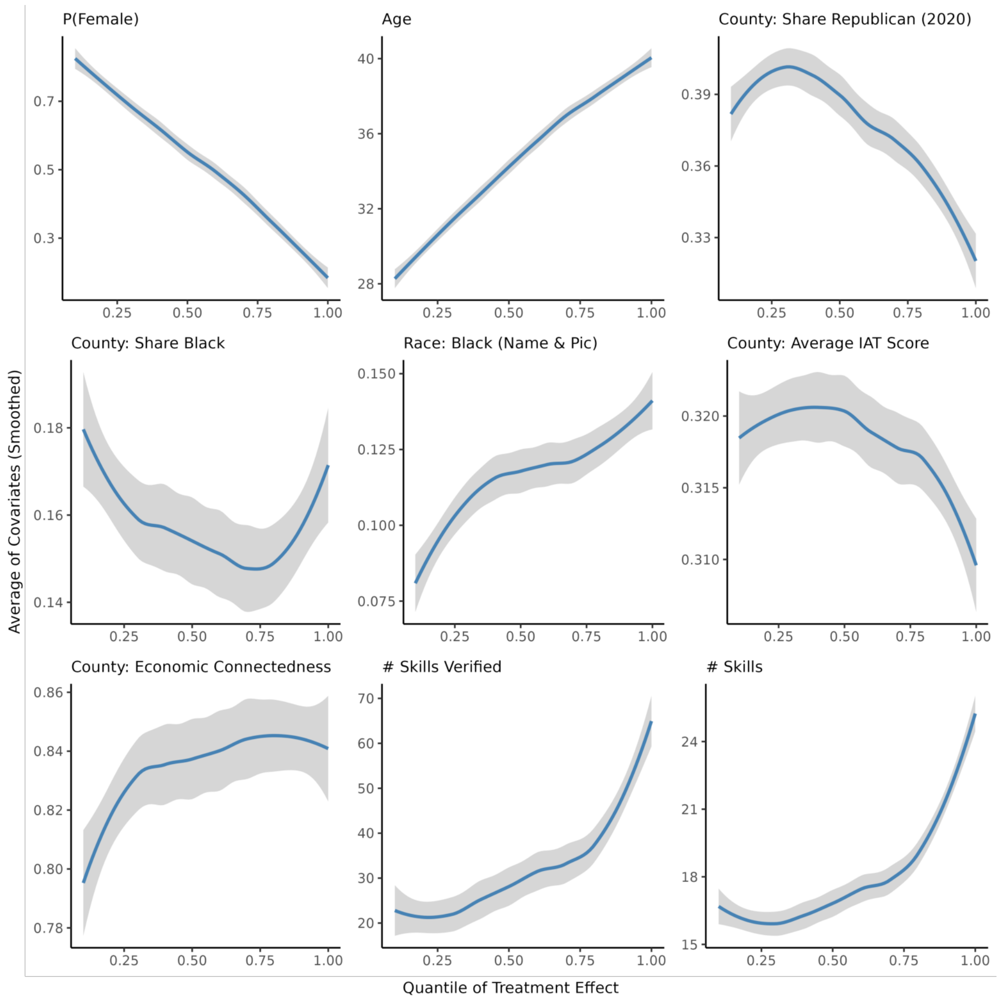

Figure 6: Quantile of estimated CATE and conditional mean of covariate

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). 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, a lower CATE suggests a higher gap and, thus, 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. Estimates are based on a causal forest with 50k trees.

To summarize the findings of the first stage of the experiment: We find a considerable difference in the propensity to accept connection requests from Black and White profiles. This difference emerges instantly and remains over time, resulting in a 13% gap by the end of the first stage. We find that Black and White profiles have a different composition of their resulting networks. For example, the networks of Black profiles encompass more men and users with more contacts. Given detailed information on our targets, we are able to discover multiple correlates of discrimination. For instance, we provide evidence that females and younger users discriminate more. We also document that Black users do discriminate but to a lower extent, mostly driven by Black women discriminating less. This suggests that homophily cannot (fully) explain our results. Finally, in an explorative analysis based on a large set of covariates, we show that discrimination is very widespread and exhibits substantial heterogeneity across individuals with different characteristics and home counties.

4.3 Informational benefits

In this section, we evaluate discrimination in information provision. Our design allows us to differentiate between discrimination stemming from gatekeeping (stage I) and differences in response rates (stage II). Before doing so, we describe the content of the received replies. We then study differences in responses and response rates, i.e., discrimination during stage II. Finally, we estimate the informational benefits a Black and White profile can expect to receive due to discrimination during both the first and second stage.

4.3.1 The value of replies

Overall, acquired connections were very generous in their responses (see Table J.11 in the Appendix). Roughly 21% of all contacts responded to our inquiry. On average, the messages contained roughly 50 words. However, the responses varied widely in their length – while some only included a few words, others spanned over half a page. 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.

To get a better understanding of who provides useful information, we first estimate predictors of a successful reply. Overall, there is little heterogeneity in the propensity to answer a message.[31] The strongest predictors suggest that individuals who are more active or present on LinkedIn, as indicated through more connections, posts, and skills listed, as well as those who have visited the same university as our profiles are more likely to respond. In the second step, we also estimate predictors of highly useful replies, i.e., replies that offer a referral or a meeting.[32] Here, males and targets with more connections are more likely to provide a valuable response.

The results above suggest that the average connection of our Black profiles may be more responsive and likely to provide a useful response. The average connection of our Black profiles is more male and has more connections, both of which are predictive of higher response rates and useful responses. However, Black profiles also contain fewer connections overall and may be discriminated against during stage II.

> **Result 2a:** The newly connected weak ties provide highly useful information.

4.3.2 Swapping profile pictures and response behavior

Before we discuss whether and how the race of our profile affects response behavior, we first address the concern that swapping profile pictures after the first stage of the experiment might induce undesirable side effects. Specifically, the targets could realize that a former White profile is now Black (or vice versa). To alleviate that concern, we provide three pieces of evidence that speak

[31] Appendix G.5 and Figure G.10 provide further details.
[32] Table J.18 reports upon the results.

against it (for more details, see Appendix G.6), by focusing on profile views, link suspensions, and response behavior.

Profile Views If targets were to notice changes in profile pictures, we would expect them to more frequently visit these profiles' pages, increasing profile views of swapped in comparison to non-swapped profiles. This can be studied using a simple diff-in-diff design. We find no difference between swapped and non-swapped profiles prior to the picture change, and, more importantly, we find no difference the weeks after the swapping (see Figure G.11 and Table J.15). Thus, swapping did not change how often our profile's pages have been viewed.

Link suspension An alternative pathway for targets to react is to simply dissolve the connection after observing a change in the profile picture. In general, suspensions were extremely rare and most likely happened due to users leaving LinkedIn rather than actively dissolving our connections. More importantly, suspension rates are virtually identical between swapped and non-swapped profiles (see Figure G.12 and Table J.16).

Response behavior Lastly, we can also directly focus on how targets responded to our message. Notably, we cannot simply compare swapped and non-swapped profiles, as responses could depend on how well the profile 'fits' into the network (independent of whether the targets realize any change in the profile picture). We can, however, leverage the fact that some targets received a connection request at an earlier point in time during stage I. Thus, we have exogenous variation in the time between first seeing our profile with the original picture and seeing it with the new picture. Overall, between 4 and 13 weeks have passed between stage I and II. Using this variation, we, again, find no evidence that the time between accepting a connection request and receiving a message impacts the response probability or message characteristics (see Table J.17).

In summary, these three pieces of evidence indicate that profile picture swapping is highly unlikely to have any effect on targets' behavior.

4.3.3 Discrimination in responses

Prior to stage II, we swapped half of the profile pictures. This 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 a White network, and half of the White profiles have access to Black networks). This, by design, resolves all endogeneity from the first stage of the experiment and allows for a clean investigation of racial preferences in the second stage. As a result, our main analysis in this section simply compares response rates toward requests of Black and White profiles. Independent of stage I, any such difference would suggest that Black and White profiles are treated differentially when asking for advice during the second stage.

Figure 7 compares the response rates. We notice that there is, indeed, some (relatively weak) discrimination in responses. This difference is significant in low-type profiles (left panel), where

White profiles receive more responses than Black profiles. It is almost zero for high-type profiles (middle panel). Aggregating over all profile types (right panel), we find no significant difference in response probabilities (see also Table 2). These results suggest that there is very little discrimination in providing information. Once Black and White profiles are (artificially) equipped with the same networks, response rates are very similar. Table 2 reports upon common regressions confirming the previous observations. TableJ.12 and J.13 in Appendix J reports upon all of these estimations and further studies the length and the usefulness of responses. The main insight remains: using a clean identification, we find only weak evidence of discrimination against Black profiles.[33]

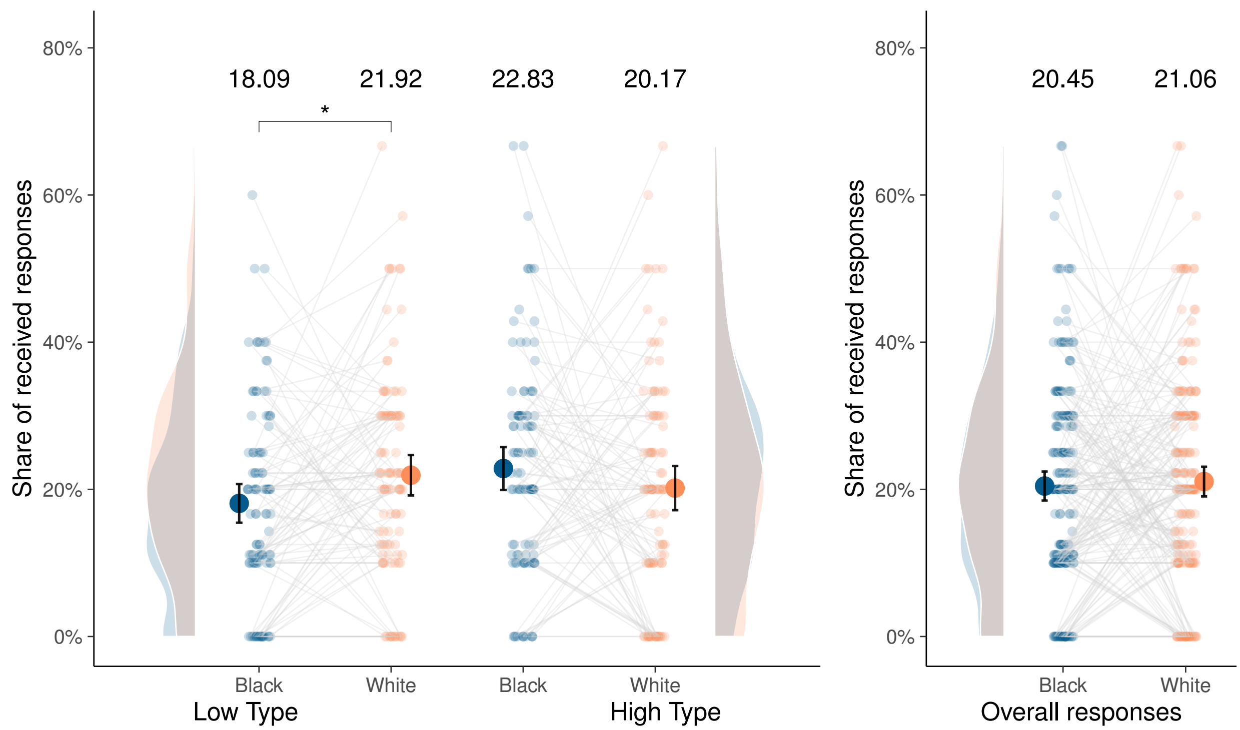

Figure 7: Response rate by race and quality of the profile.

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 middle panel represents profiles indicating attendance at more prestigious universities. The right panel depicts the results aggregating all profiles based on race, i.e. accounting for endogenously grown differences in the network characteristics. Orange objects denote White profiles, while blue objects denote Black profiles. Each dot represents one profile, and twin pairs are connected through grey lines. Whiskers around the mean denote the corresponding 95% confidence intervals. T-tests are used to obtain the following significance levels: $^{\cdot}p<0.10; ^{*}p<0.05; ^{**}p<0.01; ^{***}p<0.001$.

Our experimental setup allows us to study differences in responses in additional detail. Overall differences in information provision can be attributed to three separate effects: first, as shown above, *discrimination* may affect the response rate, i.e., people may react differentially toward Black vs. White messages. Second, the *composition of networks* between those originally built by Black and

[33] In Appendix G.6.1 we zoom in into heterogeneity in this stage and try to find subgroups of targets discriminating more/less. Essentially, find only little heterogeneity. Black and White profiles seem to be treated essentially identically across most target characteristics in this stage of the experiment.

White accounts may differ. As discussed above, some users are more likely to respond than others (e.g., those more active on LinkedIn). If, for example, Black networks contain a higher share of users with a high response rate, the composition would work in favor of their expected informational benefits. Third, responses may be affected by a profile's *'fit'* into the users' networks. Users may have a preference for interacting with individuals with certain characteristics. In a natural setting, the fit might be based on both race as well as other components, such as job interests, skills, etc. Our second stage, however, solely switches the race of half of our profiles, leaving everything else intact. Thus, the 'fit' we measure captures first-stage preferences for interacting with a specific race. It could also be interpreted as a segregation preference. The composition of networks and the 'fit' are both a function of discrimination in stage one. 'Fit' denotes racial preferences in stage one, and the 'composition of networks' results from predictable differences in racial preferences in stage one (e.g., gender).

In studies based on observational data, it is not possible to disentangle these effects. Our experimental design, however, allows us to measure each component explicitly. To operationalize the components, we run a regression with three independent variables: the first describes whether the requesting profile is Black and captures discrimination, as shown in Figure 7. The second indicates whether the profile's network was constructed by a Black profile (as opposed to a White profile) during stage I. It, thus, captures differences in the composition of the networks by comparing response rates between the two types of networks. Finally, regressions include a dummy for whether a given profile's picture was inserted into an alien network. This measures the 'fit' of the network or the segregation preferences. Importantly, by design, all three variables are orthogonal to one another. For instance, the results regarding discrimination in Figure 7 are unaffected by the 'fit' and composition effects: by swapping half of the profile pictures, we provide both Black and White profiles with access to networks with the same composition and 'fit'. This means that, following the swap, half of the Black profiles have a network originally built by a White profile and half those built by a Black profile. The same holds for White profiles. This means that, on average, both Black and White profiles face networks with the same composition. Regarding the 'fit', the swap similarly moves half of the Black profiles into the alien, i.e. White, networks. The same holds for White profiles. Thus, on average, Black and White accounts have the same 'fit'. To summarize, the post-swap race includes no information on whether a given profile was swapped ('fit') or resides in a Black network (composition). Similarly, knowing whether a profile was swapped provides no information on whether the profile resides in a Black network.

Table 1 shows the results of the three effects on response rates. The pooled results in Column 1 show that Black profiles have a slightly reduced response rate. On the other hand, users in Black networks are, overall, more likely to respond to a given message. However, both coefficients are insignificant. Finally, the 'fit' component has the strongest and significant effect.

Separately analyzing low and high types in Columns 2 to 4 reveals some differences in the components' weights. Starting with race, as shown above, we find discrimination among low-type profiles and no such evidence for the high types. At the same time, fit plays a much smaller role for low types, while it is marginally significant among high-type profiles. Finally, the composition variable suggests that individuals in Black networks are more likely to respond, especially among high types. However, the coefficient is insignificant.

In summary, we disentangle three drivers of responses: racial preferences, the composition of the network, and the fit into the network. In the pooled results, the strongest driver is the fit into the network. The weakest driver, once accounted for endogenous differences between Black and White profiles, is race. This suggests that once Black and White accounts are endowed with the same networks, we only find minor differences in users' propensity to respond to their messages. Overall, we can see that in a clean identification, Black profiles are disadvantaged, but this effect is relatively small and insignificant.

	Response Probability			
	Overall	Low quality	High quality	By quality
	(1)	(2)	(3)	(4)
Constant	0.18***	0.17***	0.19***	0.19***
	(0.01)	(0.02)	(0.02)	(0.02)
Fit	0.03·	0.01	0.04*	0.04*
	(0.02)	(0.02)	(0.02)	(0.02)
Composition	0.02	0.01	0.03	0.03
	(0.01)	(0.02)	(0.02)	(0.02)
Discrimination	0.01	0.04*	−0.03	−0.03
	(0.01)	(0.02)	(0.02)	(0.02)
Profile attended worse Uni				−0.02
				(0.03)
Fit x Profile attended worse Uni				−0.03
				(0.03)
Composition x Profile attended worse Uni				−0.02
				(0.03)
Discrimination x Profile attended worse Uni				0.07*
				(0.03)
Picture specific random effects	✓	✓	✓	✓
Log Likelihood	201.83	107.41	89.06	195.17
Observations	400	202	198	400
Notes:	·p<0.10;*p<0.05;**p<0.01;***p<0.001.			

Table 1: Decomposing stage two effects.

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.

Result 2b: We find some evidence of discrimination in the propensity to receive a response, as White profiles are slightly more likely to receive a response, however, only in the low-quality condition.

4.3.4 Expected informational benefits

In this section, we compute 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. On the one hand, it seems likely that White profiles will have an advantage, given their larger networks. On the other hand, as suggested by the composition component above, Black networks are more responsive.

To compute the overall informational benefit of a profile's network, we multiply the expected response rate of each profile's network by the number of connections of the profile's native network. Specifically, we estimate, the probability of responding to a message request for each target.[34] We then aggregate the response probabilities of the acquired connections within a profile's original network to obtain the expected number of responses, i.e., the expected information benefit.

Figure 8 shows the expected information benefit for both high and low-type profiles (see Table 2 and also see Table J.19 in Appendix J for further analysis). For both quality types, White profiles are expected to receive roughly one (22%) more messages than Black profiles. This is the combination of discrimination stemming from both the first and second stages of the experiment, i.e., discrimination originating in the formation of networks and in the response probability to messages. The benefit of Black profiles having a more responsive average network member does not sufficiently improve the response rate, and cannot overcome the disadvantage of fewer contacts. Thus, overall, Black profiles are expected to receive substantially fewer informational benefits through their networks. Given that – on the aggregate – we find no evidence of second-stage discrimination, we conclude that Black profiles' reduced informational benefit is driven by the experiment's first stage, i.e. Black networks being substantially smaller than those of White profiles.

Result 2c: Black and White profiles differ substantially in the overall expected informational benefit of their network.

[34] In more detail, we first use 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, whether the user has obtained a bachelor's degree, the number of contacts the user has, and whether the user lives in a Democratic county. Moreover, we also make use of the most important demographic characteristics like gender, age, race, and whether the user has a senior job to add to the prediction. Thereafter, we estimated the individual response probability of each connected user based on these features, interacted with the race of the profile. Missing values (for example, 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. Therefore, we impute the missing values for all users at this point only with the mean of the respective variable, which just ensures that we have a non-missing prediction for each user's probability of responding to a message. After predicting each user's response probability as a function of their characteristics and the race of our profile, we aggregate the response probability of all connected users by profile.

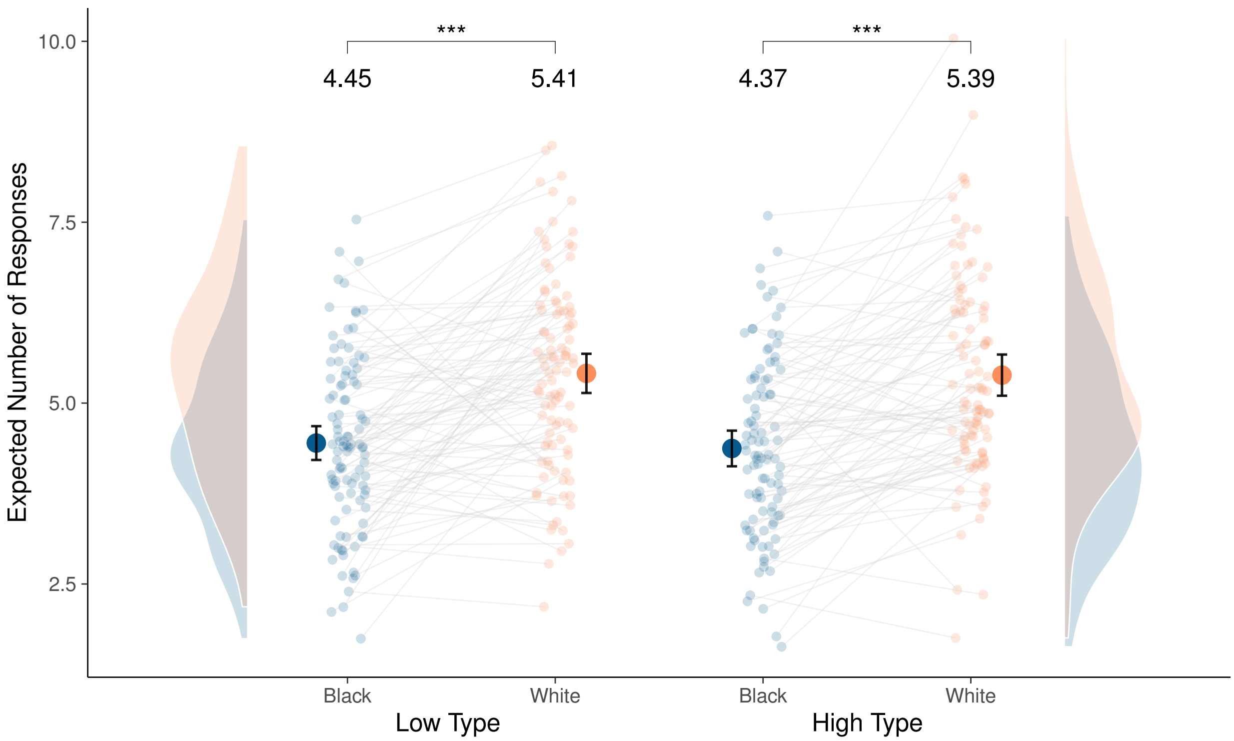

Figure 8: Number of ex-ante expected responses when creating a network.

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 grey lines. Whiskers around the mean denote the corresponding 95% confidence intervals. T-tests are used to obtain the following significance levels: $^{\cdot}p<0.10$; $^{*}p<0.05$; $^{**}p<0.01$; $^{***}p<0.001$.

	Race effect						Race and quality effect					
	Number of Contacts		Response Probability		Informational Benefit		Number of Contacts		Response Probability		Informational Benefit	
	(Stage I)		(Stage II)		(Stage I+II)		(Stage I)		(Stage II)		(Stage I+II)	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Constant	26.13***	35.66***	0.21***	0.28·	5.39***	7.76***	26.06***	35.25***	0.20***	0.27·	5.38***	7.71***
	(0.44)	(5.14)	(0.01)	(0.14)	(0.09)	(1.07)	(0.63)	(5.22)	(0.01)	(0.14)	(0.13)	(1.09)
	[0.45]	[4.45]	[0.01]	[0.13]	[0.10]	[0.93]	[0.65]	[4.58]	[0.02]	[0.14]	[0.14]	[0.96]
Profile is Black	−3.06***	−3.07***	−0.01	−0.003	−0.98***	−0.96***	−3.26***	−3.49***	0.03	0.03	−1.01***	−1.04***
	(0.47)	(0.54)	(0.01)	(0.01)	(0.10)	(0.11)	(0.68)	(0.73)	(0.02)	(0.02)	(0.14)	(0.15)
	[0.48]	[0.58]	[0.01]	[0.01]	[0.10]	[0.12]	[0.66]	[0.77]	[0.02]	[0.02]	[0.13]	[0.16]
Profile attended worse Uni							0.16	−0.09	0.02	0.03	0.02	−0.04
							(0.88)	(0.76)	(0.02)	(0.02)	(0.18)	(0.16)
							[0.91]	[0.76]	[0.02]	[0.02]	[0.20]	[0.16]
Profile is Black and attended worse Uni							0.40	0.80	−0.07*	−0.07*	0.06	0.16
							(0.95)	(0.96)	(0.03)	(0.03)	(0.20)	(0.20)
							[0.95]	[1.02]	[0.03]	[0.03]	[0.20]	[0.21]
State Controls	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓
Job Controls	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓
Firstname Controls	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓
Lastname Controls	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓
Picture trait Controls	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓
Picture specific random effects	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Log Likelihood	-1279.12	-1075.11	206.07	40.68	-653.92	-585.02	-1277.43	-1073.36	203.63	37.82	-655.48	-586.52
Observations	400	400	400	400	400	400	400	400	400	400	400	400

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

Table 2: Main estimates.

The table estimates the number of contacts a profile has by the end of stage one as a function of their race in columns 1, 2, 7, and 8. Columns 3, 4, 9, and 10 estimate the response probability in stage II (after swapping profile pictures). Columns 5, 6, 12, and 12 estimate the expected informational benefit of the profiles. *Profile is Black* denotes a dummy with value one if the profile picture (in the current stage) depicts a Black person, and zero otherwise. *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, use various controls, and all 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. In square brackets, we further display robust standard errors clustered on the twin level.[35]

4.4 Expert Survey

To compare our findings to the priors of experts working on labor economics and/or discrimination, we conducted an expert survey. We reached out to 2,171 labor economists from the Institute of Labor Economics (IZA) network and participants from the NBER's Labor Studies Summer Institute '21 and '22 (see Appendix H for more details and more results). Overall, 269 experts completed our survey. We briefly presented experts with the key features of our experiment and asked them to predict the behavior of targets.

The vast majority of experts correctly predict that White profiles receive more connections than Black profiles. They also accurately forecast educational attainment's positive effect on reducing discrimination. Similarly, experts correctly anticipate that Black users exhibit less discrimination. However, they incorrectly predict that only non-Black users drive discrimination, whereas our data

[35]The estimated coefficients and standard errors between the clustering approach and the mixed effects model are very similar across the board. Subsequently, we restrict our attention to our preferred econometric model by using mixed effect models for reasons of efficiency. However, no result is driven by this estimation choice.

shows that both Black and non-Black users do so.

More strikingly, our findings challenge the priors of experts with regard to how age is associated with discrimination, with experts anticipating older generations to exhibit higher acceptance rate gaps, contrasting our findings that younger generations do so. Furthermore, while experts predict higher discrimination from male users, our results reveal that female users display higher gaps in acceptance rates. Finally, most experts incorrectly forecast gaps in response rates during stage II of our study, whereas our results demonstrate no significant gaps. Interestingly, experts' predictions were extremely homogeneous, with little variation among different expert groups.

In summary, while experts correctly predicted some aspects, such as the overall gap in acceptance rates during stage I, the effect of education, and the directional effect of race, there were surprising deviations concerning discrimination by age, gender, and the persistence of bias in the second stage.

4.5 Back-of-the-Envelope Calculations: Economic Effects of Networks

Our data allow us to roughly estimate annual wage differences between Black and White individuals resulting from discrimination in network formation. These stem from differences in the size of the networks and, hence, the number of informational benefits potentially available through them. We provide two different back-of-the-envelope estimates.

The first estimate relies on the increased probability of finding a job through weak-tie connections on LinkedIn. This is particularly relevant for individuals just entering the job market – similar to our profiles – and unemployed individuals more generally. In the US, Black (youth) unemployment is substantially higher than White (Holzer, 1987; Sorkin, 2023). Further, worker separations are more likely for Black workers and more frequently lead to unemployment (Sorkin, 2023). To obtain an estimate we rely on a recent paper by Rajkumar et al. (2022), who show that each weak tie results in a 0.0047 probability increase of getting a job. A person with a similar occupation as our profiles receives an average annual wage of $45,000 if she obtains a new job (see Table A.6). Our data also suggest that an average White user has 286 connections, which is an underestimate, given that LinkedIn caps the number of reported connections at 500. Further, our first-stage findings suggest that a Black user is expected to have a 13% smaller network due to discrimination. Given these facts, we can calculate the expected wage difference between Black and White users due to differences in acquiring a job through networks on LinkedIn. This results in roughly $2239 being 'lost' by a Black user due to a smaller network in a given year, or an average monthly loss of roughly $200.[36]

Second, to obtain a less narrow estimate, we draw on our own data with respect to targets'

[36]Calculation: The probability of not getting a job offer due to the network is $(1 - 0.0047)$, and the probability of not getting a job from any of the 286 connections of a White user is $(1 - 0.0047)^{286}$. Correspondingly, the probability of not getting a job from any of the connections of a Black user is $(1 - 0.0047)^{286 \cdot 0.87}$. Hence, the difference between a Black user and a White user receiving a job is $(1 - (1 - 0.0047)^{286} - (1 - (1 - 0.0047)^{286 \cdot 0.87}))$. Assuming that a job is obtained only through the network, the corresponding overall difference in yearly income is $(1 - (1 - 0.0047)^{286} - (1 - (1 - 0.0047)^{286 \cdot 0.87})) \cdot 45000 \approx 2239.12$

earnings. The benefits of networks go beyond providing jobs and referrals. For instance, networks can expose individuals to relevant information, such as information on open positions, continued education programs, or other career opportunities. They can also increase their visibility to potential employers (Burt, 1992). To estimate the value of an additional connection in terms of the wage, we run a linear regression of a target's income on her number of connections, controlling for gender, education, race, and a second-order polynomial of age. The result suggests that an additional connection is associated with $70.6 additional yearly income (see Table J.20).[37] For our profiles, the difference between White and Black accounts is three connections. In our data, three connections are associated with a $212 difference in wages. Given that an average White user has 286 connections, we would expect Black profiles, based on our first-stage estimates, to only have 249 contacts. This corresponds to an estimated annual wage loss of $2,612 for Black users due to discrimination, or an average monthly loss of roughly $220.

The calculations above suggest that differences in job networks are likely to translate into substantial economic effects. While they are rather crude, it is additionally worth noting that they only refer to effects at a single point in time. Networks, however, are likely to have long-run effects on labor market outcomes. For instance, a theoretical paper Galenianos (2020) shows that small initial differences in networks between two groups can lead to substantial differences in outcomes. Further, increased job opportunities and referrals at the offset might non-linearly agglomerate over time. Thus, the static estimates above are likely to underestimate total economic effects. Nevertheless, they do provide a first idea of the economic effects of differences in network formation and information provision due to discrimination.

5 Summary and Concluding Remarks

We study the causal role of discrimination in the formation and information provision of job networks. We conduct a large-scale two-stage field experiment on LinkedIn, the largest job networking platform in and outside the U.S. During the first stage, 400+ fictitious LinkedIn profiles develop networks by sending connection requests to 20,000 users. Each user receives requests from two accounts with equivalent CVs, but one profile is Black while the other is White. Race is signaled solely through A.I.-generated profile pictures. In the second stage of the experiment, our fictitious profiles request job-relevant information from their first-stage networks. This allows us to assess how much information Black and White profiles obtain from their job networks. We examine whether Black profiles are discriminated against on the basis of their skin tone and race-specific facial features only. For this, we develop an algorithm that transforms the race of an A.I.-generated picture while keeping other facial features stable. Through survey evidence, we demonstrate that these images are perceived as realistic while keeping variables other than race – such as age, trust, intelligence, looks, and authenticity – constant.

Our experiment yields three main findings. First, we find substantial evidence for discrimina-

[37] Figure I.4 suggests that the functional form is, indeed, well-approximated using a linear function.

tion in the formation of job networks. Black profiles have a 13% lower acceptance rate for their connection requests than the White profiles (23% vs 26%), which is very close to the 2-3 p.p. difference in employers' callbacks found in previous studies (Agan and Starr, 2018; Bertrand and Mullainathan, 2004; Kline et al., 2022; Nunley et al., 2017). Discrimination is widespread, both geographically and across individuals with different characteristics. However, there are some important heterogeneities. For instance, men, older individuals, and those from more Democratic counties show smaller gaps in acceptance rates than women, young users, and individuals from Republican counties respectively. Second, we find that Black users receive substantially fewer informational benefits during the second stage. Overall, Black profiles are expected to receive fewer messages when asking for advice. Third, our novel experimental design allows us to separately identify whether differences in informational benefits originate due to gatekeeping (stage I) or due to discrimination when requesting job-relevant information (stage II). We only find limited evidence of discrimination during the experiment's second stage: when providing Black and White profiles with access to the same networks, we only find marginal effects of race on response rates. We conclude that differences in informational benefits are primarily driven by discrimination during the formation of networks, i.e., gatekeeping.

A survey with more than 250 experts in labor economics further highlights that some of the results go against experts' priors. While experts correctly anticipate discrimination during the experiment's first stage, they expect discrimination to proliferate to the same extent during the second stage. They further predict men and older individuals to discriminate to a higher extent – we find the opposite to be the case.

Overall, our paper presents compelling evidence that discrimination plays a significant role in shaping the informational benefits provided by professional job networks. Around half of all jobs are found through informal networks (Topa, 2011) and they have been shown to have strong effects on individuals' labor market outcomes (Dustmann et al., 2016). Our findings, thus, offer valuable evidence regarding a significant channel that can help explain the disparities in labor market outcomes between minority groups and the White majority in the US labor market. Already, in a 1987 paper in the American Economic Review, Holzer (1987) suggested that *"informal methods of search [...] account for 87-90 percent of the difference in youth employment probabilities between blacks and whites"* (p. 451). Our paper provides the first causal evidence in this regard. Our findings demonstrate that the disparities in networks identified in earlier studies (Fernandez and Fernandez-Mateo, 2006) are at least partially attributable to direct discrimination. They further show that most discrimination is effectively driven by gatekeeping. This accompanies evidence on the importance of creating inclusive institutions and breaking up 'old boys clubs' (Cullen and Perez-Truglia, 2023; Michelman et al., 2021). It further provides justification for affirmative action, given the major role of gatekeeping in explaining outcomes, including through inclusive networking events and workshops. By demonstrating that race affects networking and interactions between individuals, the results further link to recent evidence on economic connectedness and mobility (Chetty et al., 2022a,b). Further, we shed light on the mechanisms through which professional job networking platforms, such as LinkedIn, aid users in advancing their careers. We show that weak-tie networks, such as those of our users, provide substantial informational benefits regarding mentorship advice and job applications. Our study thus complements previous work on the strength of weak-ties (Gee et al., 2017a; Rajkumar et al., 2022), the economic value of professional job networking platforms (Wheeler et al., 2022), and the general literature on online audit studies (e.g. Acquisti and Fong, 2020; Edelman et al., 2017). Finally, our paper provides evidence on networks creating labor market friction through discrimination. This may help bridge the gap between pervasive evidence on discrimination (e.g. Neumark, 2018) and the prediction of its absence based on market-based models. As noted by Arrow (1998), this requires looking for non-market factors that affect economic behavior and "[..] networks seem to be good places to start" (p. 93).

This study opens up numerous avenues for subsequent research. First, we are the first to causally study the effects of discrimination on network formation and information provision on LinkedIn. While, in some aspects, offline job networks may function differently than online networks, the platform provides an ideal setting to cleanly study job networks in general and discrimination more specifically. Ours is the first paper to causally study discrimination in job network formation and information provision. Given that both offline and online networks have been shown to strongly affect labor market outcomes (Dustmann et al., 2016; Wheeler et al., 2022), it is crucial for future research to provide additional evidence on other countries, minorities, and genders. When doing so, our approach to vary race via A.I. generated pictures can offer an alternative to using names as signals, which have been demonstrated to be noisy and potentially biased (Gaddis, 2017; Kreisman and Smith, 2023). The approach is easily adaptable to different contexts, enabling researchers to modify a range of individual attributes, ranging from race to gender or age. While our study specifically focuses on varying one dimension of profile pictures, namely race, our findings might not be directly generalizable to females. Given that women are, e.g., more frequently the subject of sexual harassment on online platforms (Atske, 2021), future papers might want to focus on this question with an adjusted experimental design. Further, the results suggest a potentially important channel for differences in labor market outcomes. Finally, our results highlight heterogeneity in discriminatory behavior, some of which is surprising both to us and to hundreds of experts. These insights emphasize the need for research to further our understanding of who drives discrimination, why, and where it originates. Doing so is important to design effective and well-targeted policies to counter discrimination.

References

Abel, Martin and Rulof Burger, "Unpacking Name-Based Race Discrimination," Working Paper 2023.

Acquisti, Alessandro and Christina Fong, "An Experiment in Hiring Discrimination via Online Social Networks," *Management Science*, 2020, *66* (3), 1005–1024.

Agan, Amanda and Sonja Starr, "Ban the box, criminal records, and racial discrimination: A field experiment," *The Quarterly Journal of Economics*, 2018, *133* (1), 191–235.

Ajzenman, Nicolás, Bruno Ferman, and Pedro C Sant'Anna, "Discrimination in the Formation of Academic Networks: A Field Experiment on# EconTwitter," Technical Report 2023.

Angeli, Deivis, Matt Lowe, and The Village Team, "Virtue Signals," CESifo Working Paper 2023.

Arechar, Antonio A., Simon Gächter, and Lucas Molleman, "Conducting interactive experiments online," *Experimental Economics*, Mar 2018, *21* (1), 99–131.

Arrow, Kenneth and Ron Borzekowski, "Limited Network Connections and the Distribution of Wages," Technical Report 2004.

Arrow, Kenneth J, "What has economics to say about racial discrimination?," *Journal of Economic Perspectives*, 1998, *12* (2), 91–100.

Athey, Susan, "Susan Athey Guest Talk - Estimating Heterogeneous Treatment Effects," October 2020.

_ **and Stefan Wager**, "Estimating treatment effects with causal forests: An application," *Observational Studies*, 2019, *5* (2), 37–51.

_ **, Julie Tibshirani, and Stefan Wager**, "Generalized random forests," *The Annals of Statistics*, 2019.

_ **, Stefan Wager, Vitor Hadad, Sylvia Klosin, Nicolaj Muhelbach, Xinkun Nie, and Matt Schaelling**, "Estimation of Heterogeneous Treatment Effects," May 2020.

Atske, Sara, "The State of Online Harassment — pewresearch.org," https://www.pewresearch.org/internet/2021/01/13/the-state-of-online-harassment/ January 2021. [Accessed 02-Mar-2023].

Auxier, Monica Anderson Brooke, "Social Media Use in 2021," https://www.pewresearch.org/2021/04/07/social-media-use-in-2021-3/ April 2021. [Accessed 02-Mar-2023].

Baker, Edmund and Veronika Grimm, "Monitoring Decreases Discrimination," Working Paper 2021.

Barabási, Albert-László and Réka Albert, "Emergence of scaling in random networks," *Science*, 1999, *286* (5439), 509–512.

Bayer, Patrick and Kerwin Kofi Charles, "Divergent paths: A new perspective on earnings differences between black and white men since 1940," *The Quarterly Journal of Economics*, 2018, *133* (3), 1459–1501.

Beaman, Lori and Jeremy Magruder, "Who Gets the Job Referral? Evidence from a Social Networks Experiment," *American Economic Review*, December 2012, *102* (7), 3574–3593.

_ **, Niall Keleher, and Jeremy Magruder**, "Do job networks disadvantage women? Evidence from a recruitment experiment in Malawi," *Journal of Labor Economics*, 2018, *36* (1), 121–157.

Berinsky, Adam J, Gregory A Huber, and Gabriel S Lenz, "Evaluating online labor markets for experimental research: Amazon.com's Mechanical Turk," *Political Analysis*, 2012, *20* (3), 351–368.

Bertrand, Marianne and Ester Duflo, "Field Experiments on Discrimination," *Handbook of Economic Field Experiments*, 2017, *1*, 309–393.

_ **and Sendhil Mullainathan**, "Are Emily and Greg more employable than Lakisha and Jamal? A field experiment on labor market discrimination," *American Economic Review*, 2004, *94* (4), 991–1013.

Block, Ray, Charles Crabtree, John B. Holbein, and J. Quin Monson, "Are Americans less likely to reply to emails from Black people relative to White people?," *Proceedings of the National Academy of Sciences*, 2021, *118* (52), e2110347118.

Bohren, J. Aislinn, Alex Imas, and Michael Rosenberg, "The Dynamics of Discrimination: Theory and Evidence," *American Economic Review*, October 2019, *109* (10), 3395–3436.

Bohren, J Aislinn, Alex Imas, and Michael Rosenberg, "The dynamics of discrimination: Theory and evidence," *American Economic Review*, 2019, *109* (10), 3395–3436.

_ **, Peter Hull, and Alex Imas**, "Systemic discrimination: Theory and measurement," *Working Paper*, 2022.

Bordalo, Pedro, Katherine Coffman, Nicola Gennaioli, and Andrei Shleifer, "Stereotypes," *The Quarterly Journal of Economics*, July 2016, *131* (4), 1753–1794.

Burt, Ronald S, *Structural holes: The social structure of competition*, Harvard University Press, 1992.

Calvo-Armengol, Antoni and Matthew O Jackson, "The effects of social networks on employment and inequality," *American Economic Review*, 2004, *94* (3), 426–454.

Chernozhukov, Victor, Mert Demirer, Esther Duflo, and Ivan Fernandez-Val, "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 2018.

Chetty, Raj, Matthew O Jackson, Theresa Kuchler, Johannes Stroebel, Nathaniel Hendren, Robert B Fluegge, Sara Gong, Federico Gonzalez, Armelle Grondin, Matthew Jacob et al., "Social capital I: measurement and associations with economic mobility," *Nature*, 2022, *608* (7921), 108–121.

_ , _ , _ , _ , _ , _ , _ , _ , _ , _ , and _ , "Social capital II: determinants of economic connectedness," *Nature*, 2022, *608* (7921), 122–134.

Coenders, Marcel and Peer Scheepers, "Support for Ethnic Discrimination in the Netherlands 1979-1993: Effects of Period, Cohort, and Individual Characteristics," *European Sociological Review*, 1998, *14* (4), 405–422.

Coffman, Katherine B., Christine L. Exley, and Muriel Niederle, "The Role of Beliefs in Driving Gender Discrimination," *Management Science*, June 2021, *67* (6), 3551–3569.

Conzelmann, Johnathan G., Steven W. Hemelt, Brad Hershbein, Shawn M. Martin, Andrew Simon, and Kevin M. Stange, "Grads on the Go: Measuring College-Specific Labor Markets for Graduates," Working Paper 30088, National Bureau of Economic Research May 2022.

Cullen, Zoë B and Ricardo Perez-Truglia, "The old boys' club: Schmoozing and the gender gap," *American Economic Review*, 2023, *113* (7), 1703–40.

Davidov, Eldad and Bart Meuleman, "Explaining attitudes towards immigration policies in European countries: The role of human values," *Journal of Ethnic and Migration Studies*, 2012, *38* (5), 757–775.

Davis, Angela Y, "Rape, racism and the capitalist setting," *The Black Scholar*, 1981, *12* (6), 39–45.

Doleac, Jennifer L and Luke CD Stein, "The visible hand: Race and online market outcomes," *The Economic Journal*, 2013, *123* (572), 469–492.

Dustmann, Christian, Albrecht Glitz, Uta Schönberg, and Herbert Brücker, "Referral-based Job Search Networks," *The Review of Economic Studies*, April 2016, *83* (2), 514–546.

Edelman, Benjamin, Michael Luca, and Dan Svirsky, "Racial discrimination in the sharing economy: Evidence from a field experiment," *American Economic Journal: Applied Economics*, 2017, *9* (2), 1–22.

Edwards, Chad, Brett Stoll, Natalie Faculak, and Sandi Karman, "Social presence on LinkedIn: Perceived credibility and interpersonal attractiveness based on user profile picture," *Online Journal of Communication and Media Technologies*, 2015, *5* (4), 102.

Fernandez, Roberto M. and Isabel Fernandez-Mateo, "Networks, Race, and Hiring," *American Sociological Review*, 2006, *71*, 42–71.

Forbes, "Forbes America's Top Colleges List 2021," Technical Report 2021.

Gaddis, S Michael, "How black are Lakisha and Jamal? Racial perceptions from names used in correspondence audit studies," *Sociological Science*, 2017, *4*, 469–489.

Galenianos, Manolis, "Referral networks and inequality," Working Paper May 2020.

Gallen, Yana and Melanie Wasserman, "Do Male and Female Students Use Networks Differently?," *AEA Papers and Proceedings*, 2021, *111*, 154–158.

_ **and** _ , "Informed choices: Gender gaps in career advice," CEPR Discussion Paper No. DP15728 2021.

Gee, Laura K., "The More You Know: Information Effects on Job Application Rates in a Large Field Experiment," *Management Science*, May 2019, *65* (5), 2077–2094. Publisher: INFORMS.

Gee, Laura K, Jason J Jones, Christopher J Fariss, Moira Burke, and James H Fowler, "The paradox of weak ties in 55 countries," *Journal of Economic Behavior & Organization*, 2017, *133*, 362–372.

_ **, Jason Jones, and Moira Burke**, "Social networks and labor markets: How strong ties relate to job finding on Facebook's social network," *Journal of Labor Economics*, 2017, *35* (2), 485–518.

Georgia Department of Public Health, "Most Popular Baby Names," Technical Report 2022.

Giulietti, Corrado, Mirco Tonin, and Michael Vlassopoulos, "Racial discrimination in local public services: A field experiment in the United States," *Journal of the European Economic Association*, 2019, *17* (1), 165–204.

Goeree, Jacob K, Margaret A McConnell, Tiffany Mitchell, Tracey Tromp, and Leeat Yariv, "The 1/d law of giving," *American Economic Journal: Microeconomics*, 2010, *2* (1), 183–203.

Goncalves, Felipe and Steven Mello, "A few bad apples? Racial bias in policing," *American Economic Review*, 2021, *111* (5), 1406–1441.

Granovetter, Mark, "The Strength of Weak Ties: A Network Theory Revisited," *Sociology Theory*, 1983, *1*, 201–233.

Halevy, Nir, Ori Weisel, and Gary Bornstein, ""In-group love" and "out-group hate" in repeated interaction between groups," *Journal of Behavioral Decision Making*, 2012, *25* (2), 188–195.

Hern, Alex, "FaceApp forced to pull 'racist' filters that allow 'digital blackface'," *The Guardian*, August 2017.

Holzer, Harry J., "Informal Job Search and Black Youth Unemployment," *American Economic Review*, 1987, *77* (3), 446–452.

Hopkins Population Center, "HPC Data Hub Phase 3 Release (11/16/2020)," Technical Report 2020.

Horton, John J, David G Rand, and Richard J Zeckhauser, "The online laboratory: Conducting experiments in a real labor market," *Experimental Economics*, 2011, *14* (3), 399–425.

Hughes, Michael and Steven A Tuch, "Gender differences in whites' racial attitudes: Are women's attitudes really more favorable?," *Social Psychology Quarterly*, 2003, pp. 384–401.

Hum, Noelle J., Perrin E. Chamberlin, Brittany L. Hambright, Anne C. Portwood, Amanda C. Schat, and Jennifer L. Bevan, "A picture is worth a thousand words: A content analysis of Facebook profile photographs," *Computers in Human Behavior*, 2011, *27* (5), 1828–1833.

Ioannides, Yannis M and Linda Datcher Loury, "Job information networks, neighborhood effects, and inequality," *Journal of Economic Literature*, 2004, *42* (4), 1056–1093.

IPEDS, "Integrated Postsecondary Education Data System (IPEDS)," Technical Report 2022.

Jr, Roland G Fryer and Steven D Levitt, "The causes and consequences of distinctively black names," *The Quarterly Journal of Economics*, 2004, *119* (3), 767–805.

Kaas, Leo and Christian Manger, "Ethnic discrimination in Germany's labour market: a field experiment," *German Economic Review*, 2012, *13* (1), 1–20.

Kalmijn, Matthijs, "Intermarriage and homogamy: Causes, patterns, trends," *Annual Review of Sociology*, 1998, *24* (1), 395–421.

Kaplan, Jacob, "jacobkap/predictrace," 2022.

Karras, Tero, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila, "Analyzing and Improving the Image Quality of StyleGAN," in "Proc. CVPR" 2020.

Kline, Patrick, Evan K Rose, and Christopher R Walters, "Systemic Discrimination Among Large U.S. Employers*," *The Quarterly Journal of Economics*, 2022, *137* (4), 1963–2036.

Kreisman, Daniel and Jonathan Smith, "Distinctively Black names and educational outcomes," *Journal of Political Economy*, 2023, *131* (4).

Kroft, Kory, Fabian Lange, and Matthew J Notowidigdo, "Duration dependence and labor market conditions: Evidence from a field experiment," *The Quarterly Journal of Economics*, 2013, *128* (3), 1123–1167.

LinkedIn, "LinkedIn Economic Graph: Career Explorer," 2022.

_ , "About Us," April 2023.

Manant, Matthieu, Serge Pajak, and Nicolas Soulié, "Do Recruiters 'Like' it? Privacy and Online Social Network Profile in Hiring: A Randomized Experiment," Working Paper 2013.

Marmaros, David and Bruce Sacerdote, "How do friendships form?," *The Quarterly Journal of Economics*, 2006, *121* (1), 79–119.

Mayer, Adalbert and Steven L Puller, "The old boy (and girl) network: Social network formation on university campuses," *Journal of Public Economics*, 2008, *92* (1-2), 329–347.

McClintock, Elizabeth Aura, "When does race matter? Race, sex, and dating at an elite university.," *Journal of Marriage and Family*, 2010.

McDonald, Steve, "What's in the "old boys" network? Accessing social capital in gendered and racialized networks," *Social Networks*, 2011, *33*, 317–330.

_ , Nan Lin, and Dan Ao, "Networks of Opportunity: Gender, Race, and Job Leads," *Social Problems*, 2009, *56* (3), 385–402.

McPherson, Miller, Lynn Smith-Lovin, and James M Cook, "Birds of a Feather: Homophily in Social Networks," *Annual Review of Sociology*, 2001, *27* (1), 415–444.

Mejia, Jorge and Chris Parker, "When transparency fails: Bias and financial incentives in ridesharing platforms," *Management Science*, 2021, *67* (1), 166–184.

Michelman, Valerie, Joseph Price, and Seth D Zimmerman, "Old Boys' Clubs and Upward Mobility Among the Educational Elite," *The Quarterly Journal of Economics*, December 2021, p. qjab047.

MIT Election Data and Science Lab, "County Presidential Election Returns 2000-2020," Dataset 2018.

Mortensen, Dale T and Tara Vishwanath, "Personal contacts and earnings: It is who you know!," *Labour Economics*, 1994, *1* (2), 187–201.

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

— , **Roy J Bank, and Kyle D Van Nort**, "Sex discrimination in restaurant hiring: An audit study," *The Quarterly Journal of Economics*, 1996, *111* (3), 915–941.

News, US, "U.S. News Education Rankings," https://www.usnews.com/best-colleges/rankings 2019. Accessed on 29-Jun-2023.

Niche, "The Most Popular College Majors," https://www.niche.com/blog/the-most-popular-college-majors/ 2019. Accessed on 29-Jun-2023.

Nightingale, Sophie J and Hany Farid, "AI-synthesized faces are indistinguishable from real faces and more trustworthy," *Proceedings of the National Academy of Sciences*, 2022, *119* (8), e2120481119.

Nosek, Brian A, Frederick L Smyth, Jeffrey J Hansen, Thierry Devos, Nicole M Lindner, Kate A Ranganath, Colin Tucker Smith, Kristina R Olson, Dolly Chugh, Anthony G Greenwald et al., "Pervasiveness and correlates of implicit attitudes and stereotypes," *European Review of Social Psychology*, 2007, *18* (1), 36–88.

Nunley, John M, Adam Pugh, Nicholas Romero, and R Alan Seals, "The effects of unemployment and underemployment on employment opportunities: Results from a correspondence audit of the labor market for college graduates," *IRL Review*, 2017, *70* (3), 642–669.

Oreopoulus, Philip, "Why Do Skilled Immigrants Struggle in the Labor Market? A Field Experiment with Thirteen Thousand Resumes," *American Economic Journal: Economic Policy*, 2011, *3* (4), 148–171.

Pallais, Amanda and Emily Glassberg Sands, "Why the Referential Treatment? Evidence from Field Experiments on Referrals," *Journal of Political Economy*, December 2016, *124* (6), 1793–1828. Publisher: The University of Chicago Press.

Pebesma, Edzer, "Simple Features for R: Standardized Support for Spatial Vector Data," *The R Journal*, 2018, *10* (1), 439–446.

Peysakhovich, A., M. A. Nowak, and D. G. Rand, "Humans display a 'cooperative phenotype' that is domain general and temporally stable," *Nature Communications*, 9 2014, *5*, 4939.

Quillian, Lincoln, Anthony Heath, Devah Pager, Arnfinn H Midtbøen, Fenella Fleischmann, and Ole Hexel, "Do some countries discriminate more than others? Evidence from 97 field experiments of racial discrimination in hiring," *Sociological Science*, 2019, *6*, 467–496.

Rajkumar, Karthik, Guillaume Saint-Jacques, Iavor Bojinov, Erik Brynjolfsson, and Sinan Aral, "A causal test of the strength of weak ties," *Science*, 2022, *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, "Social heuristics shape intuitive cooperation," *Nature Communications*, 4 2014, *5*, 3677.

Rosenfeld, Michael J., Reuben J. Thomas, and Sonia Hausen, "Disintermediating your friends: How online dating in the United States displaces other ways of meeting," *Proceedings of the National Academy of Sciences*, August 2019, *116* (36), 17753–17758.

Salganik, Matthew J, *Bit by bit: Social research in the digital age*, Princeton University Press, 2019.

Schmutte, Ian M, "Job referral networks and the determination of earnings in local labor markets," *Journal of Labor Economics*, 2015, *33* (1), 1–32.

Silva, Bruno Castanho and Sven-Oliver Proksch, "Fake it 'til you make it: a natural experiment to identify European politicians' benefit from Twitter bots," *American Political Science Review*, 2021, *115* (1), 316–322.

Sommerville, Diane Miller, "The rape myth in the old South reconsidered," *The Journal of Southern History*, 1995, *61* (3), 481–518.

Sorkin, Isaac, "Quantifying Racial Disparities Using Consecutive Employment Spells," Working Paper 2023.

SSA, "Babynames," Dataset 2022.

Suri, Siddharth and Duncan J. Watts, "Cooperation and Contagion in Web-Based, Networked Public Goods Experiments," *PLOS ONE*, 03 2011, *6* (3), 1–18.

Taigman, Yaniv, Ming Yang, Marc'Aurelio Ranzato, and Lior Wolf, "Deepface: Closing the gap to human-level performance in face verification," in "Proceedings of the IEEE conference on computer vision and pattern recognition" 2014, pp. 1701–1708.

Tibshirani, Julie, Susan Athey, Rina Friedberg, Vitor Hadad, David Hirshberg, Luke Miner, Erik Sverdrup, Stefan Wager, and Marvin Wright, "grf: Generalized Random Forests," May 2023.

Topa, Giorgio, "Labor markets and referrals," in "Handbook of Social Economics," Vol. 1, Elsevier, 2011, pp. 1193–1221.

U.S. Bureau of Labor Statistics, "Table F. Distribution of private sector employment by firm size class: 1993/Q1 through 2022/Q1, not seasonally adjusted," Technical Report 2022.

U.S. Census Bureau, "CBSA Shapefiles," Technical Report 2013.

_ , "County Shapefiles," Technical Report 2020.

_ , "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.," 2021.

_ , "Frequently Occurring Surnames from the Census," Technical Report 2022.

Utz, Sonja, "Is LinkedIn making you more successful? The informational benefits derived from public social media," *news media & society*, 2016, *18* (11), 2685–2702.

Vogels, Emily A, "The state of online harassment," *Pew Research Center*, 2021, *13*, 625.

Wager, Stefan and Susan Athey, "Estimation and inference of heterogeneous treatment effects using random forests," *Journal of the American Statistical Association*, 2018, *113* (523), 1228–1242.

Wheeler, Laurel, Robert Garlick, Eric Johnson, Patrick Shaw, and Marissa Gargano, "LinkedIn (to) job opportunities: Experimental evidence from job readiness training," *American Economic Journal: Applied Economics*, 2022, *14* (2), 101–25.

Xu, Frank Kaiyuan, Nicole Lofaro, Brian A. Nosek, Anthony G. Greenwald, Jordan Axt, Lauren Simon, and Nicole Frost, "Race IAT 2002-2022," 2022.

Yadlowsky, Steve, Scott Fleming, Nigam Shah, Emma Brunskill, and Stefan Wager, "Evaluating treatment prioritization rules via rank-weighted average treatment effects," *arXiv preprint arXiv:2111.07966*, 2021.

Zounlome, Nelson OO, Y Joel Wong, Elyssa M Klann, and Jessica L David, ""I'm already seen as a sexual predator from saying hello": Black men's perception of sexual violence," *Journal of Interpersonal Violence*, 2021, *36* (19-20), NP10809–NP10830.

Zschirnt, Eva, "Research ethics in correspondence testing: An update," *Research Ethics*, 2019, *15* (2), 1–21.

A Experimental Design

A.1 Timeline

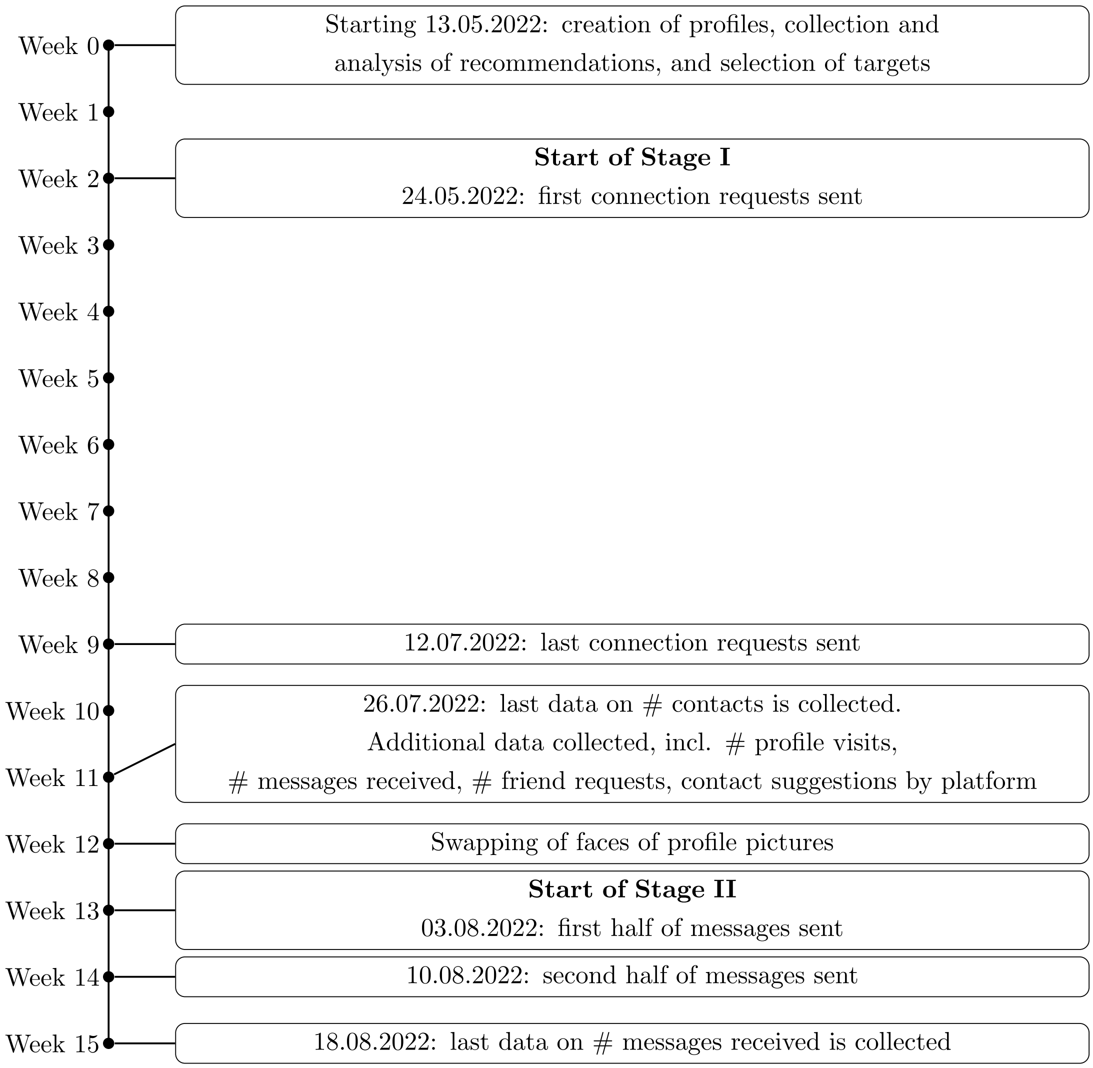

Figure A.1: Timeline of Experiment

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.

A.2 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.1: 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.

A.3 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 Akransas	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.2: 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 proceeded 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.3: 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 "neighbours", we proceeded 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 Forbes 600 ranking and US News Ranking. We also present information on the enrollment at each institution

A.4 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.4: First Names of Profiles

Note: We obtain the most common first names of males born 1997 in Georgia from Georgia Department of Public Health (2022). We then focus on the names which are within the top 50 most common names for *both* White and Black males, 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 SSA (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.5: 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.

A.5 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.6: Job Titles and Average Pay According to Payscale.com (2022)

Job	Description Items
Office Manager	
Description 1	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	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	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	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	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	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	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	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 memorand 4. Greet and proactively assist visitors in a timely manner
Marketing Assistant	
Description 1	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	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.7: Job descriptions

Job descriptions are taken from CV examples on websites like `ideed.com`, `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.8: 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.9: Volunteer Work indicated in Profile

Note: descriptions are taken from CV examples on websites like ideed.com, monster.com, etc.

A.6 Firms

To obtain employers, we first used Statista's Company Data Base to identify the largest employers in each city. If the city was unique in the USA, we used the largest employers as our companies. For cities with too few employers or cities with multiple mentionings, we searched for local information on the largest employers. We used 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 sufficient 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.10.

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	Pharmerica 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	Integris 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	Mountaire 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	Iqor
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	Freeport-Mcmoran 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	Kampgrounds 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	Winco Foods, LLC	118	NH	Easter Seal New Hampshire, Inc.	186	VA	Amerigroup (Anthem)
51	ID	Winco Holdings, Inc.	119	NH	Legacy Echn, 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	Eeco, 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.10: List of employers indicated in profiles

A.7 Process of Profile Creation

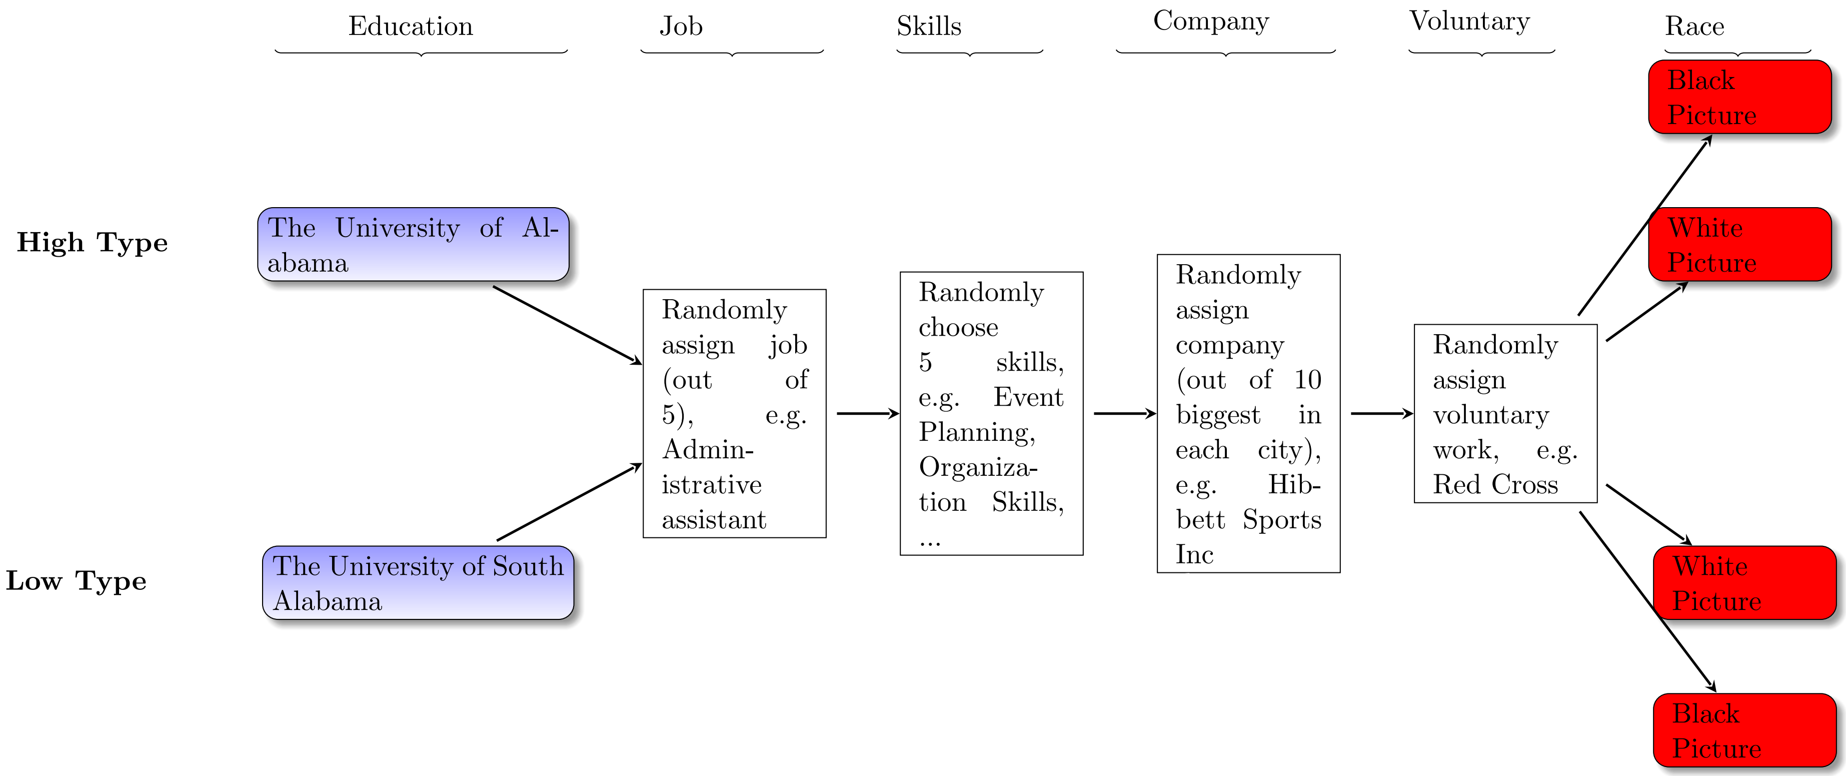

Figure A.2: Profile Creation: Example for Birmingham (Alabama)

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.

A.8 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.

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 A.I. 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 males. 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 females, 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 of each 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 which have the smallest difference between the potential White and Black profile in terms of picture characteristics are used in the final sample.

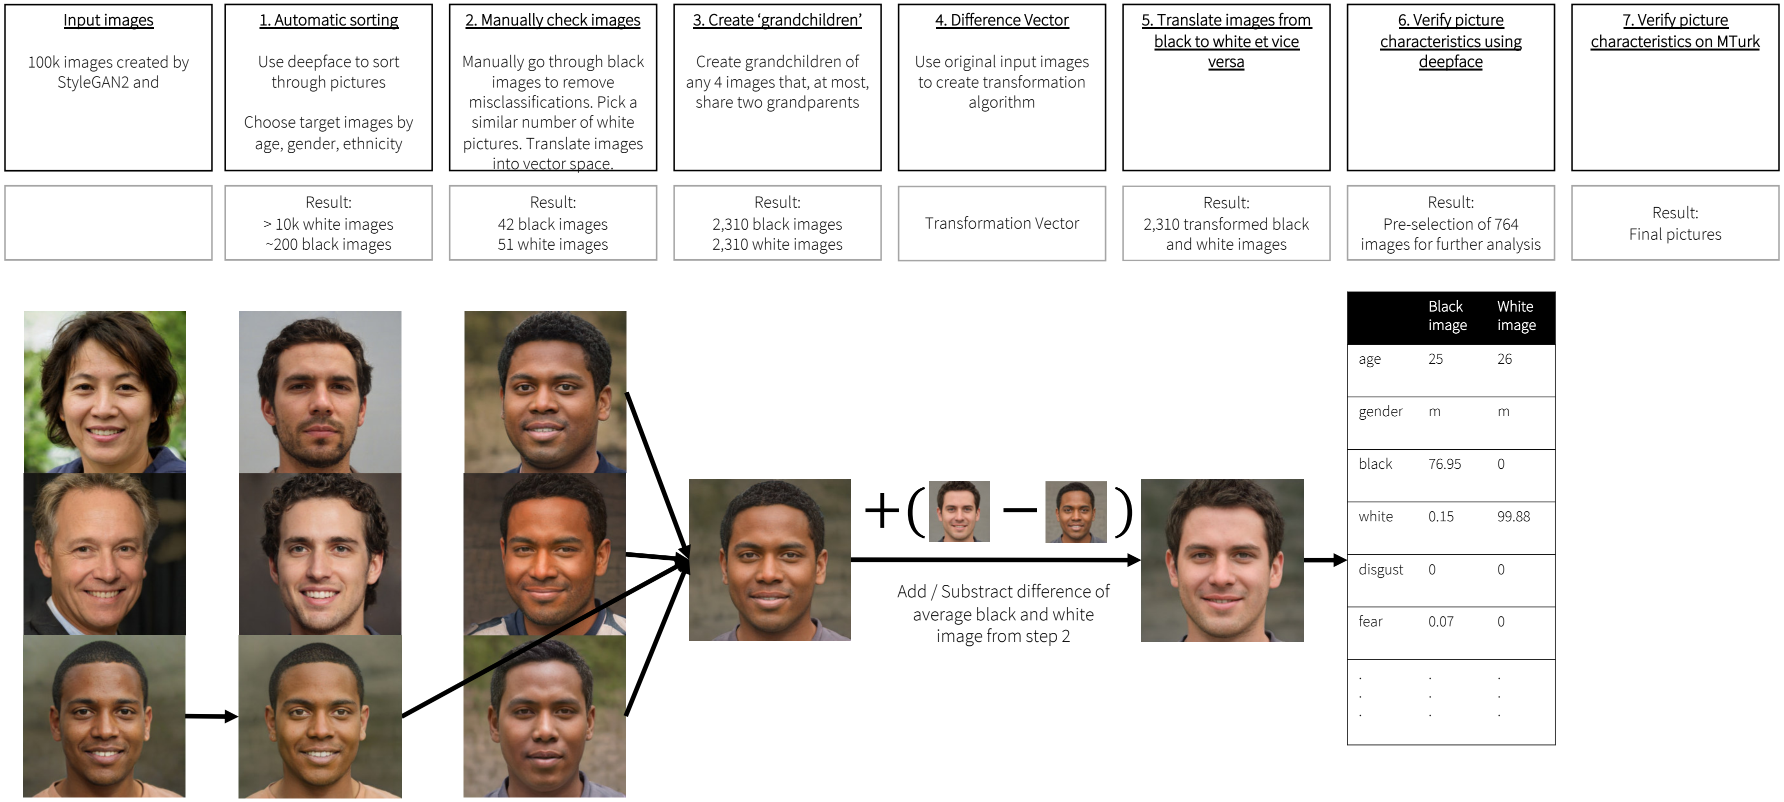

Figure B.1: Picture Creation: Visualization of data processing, selection, and validation of pictures

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.

C.1 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 males and females as potential genders. Given our balancing, around half of the targets are female.

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.

C.2 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 in order 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` (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”.[38] *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.

C.3 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 ≥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.

C.4 Education

Degree The most recently listed degree in CVs is analyzed using a dictionary approach.[39] 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”,

[38] We 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 .

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

"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. (2022) 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. (2022). 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 female, 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.

C.5 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 2). 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 (Name)	18,776	0.52	0.50	1	0	1
Female (Name and Picture)	19,451	0.52	0.50	1	0	1
Black (Name)	17,220	0.05	0.22	0	0	1
White (Name)	17,220	0.72	0.45	1	0	1
Asian (Name)	17,220	0.10	0.30	0	0	1
Hispanic (Name)	17,220	0.13	0.34	0	0	1
American Indian (Name)	17,220	0.01	0.02	0.004	0	0.95
Two Races (Name)	17,220	0.02	0.02	0.02	0	0.38
Black (Name and Picture)	19,465	0.06	0.24	0	0	1
White (Name and Picture)	19,465	0.69	0.46	1	0	1
Age	17,004	34.40	11.00	32	10	82
EMPLOYMENT AND PLATFORM						
Salary	19,362	87,362.68	56,893.03	67,329.50	10,625	951,257
Works in HR	19,481	0.08	0.28	0	0	1
Employed	19,481	0.97	0.18	1	0	1
Retired	19,481	0.003	0.06	0	0	1
Self-Employed	19,481	0.02	0.14	0	0	1
Number of Contacts	18,901	286.18	194.22	294	1	500
Number of Followers	16,639	291.83	268.82	210	0	1,000
Number of Skills	19,481	16.94	14.42	15	0	50
Number of Skill Verifications	19,481	30.50	86.64	8	0	9,090
Number of Posts	19,481	0.84	1.10	0	0	6
Has Profile Picture	19,481	0.69	0.46	1	0	1
EMPLOYER						
Employees	16,775	4,998.75	4,552.15	5,001	0	10,001
Employees on Platform	16,718	29,808.53	76,100.26	3,367	0	962,414
Open Jobs (Platform)	16,945	2,088.69	7,020.23	105	0	107,974
EDUCATION						
None	18,863	0.20	0.40	0	0	1
Some College	18,863	0.12	0.32	0	0	1
Associate	18,863	0.04	0.19	0	0	1
Bachelor	18,863	0.40	0.49	0	0	1
Master	18,863	0.21	0.41	0	0	1
PhD	18,863	0.03	0.18	0	0	1
Undergrads: White	13,985	0.62	0.19	0.66	0.0004	1.00
Undergrads: Black	13,985	0.09	0.12	0.06	0	0.98
Undergrads: Female	13,985	0.54	0.08	0.54	0	1
Forbes Rank	9,743	230.63	164.38	188	1	599
COUNTY						
Distance to Profile (km)	18,121	374.52	873.59	14.14	0	8,068.31
Share Democrat (2020)	17,895	0.60	0.15	0.60	0.09	0.89
Share White	18,119	0.57	0.19	0.57	0.06	0.98
Share Black	18,119	0.16	0.14	0.12	0.002	0.82
Pop. Density	18,119	1,932.86	5,424.99	534.56	0.78	27,755.40
Dissimilarity Index (Black/White)	17,954	54.49	11.70	53	4	85
Dissimilarity Index (Non-White/White)	18,109	40.80	11.86	41	1	69
County: Avg. Race IAT	18,679	0.32	0.05	0.32	-0.33	0.73
EDGES						
Same University	27,736	0.08	0.28	0	0	1

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 SSA (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 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	Distance between profile and reported location of individual	Reported Location in platform CV & Google Maps API
County & CBSA	County & CBSA in which individual lives	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	Vote shares by county in 2020 presidential election	MIT Election Data and Science Lab (2018)
County-Level Demographics	general demographics: population, population density, race shares, and dissimilarity on county-level	from Hopkins Population Center (2020)
Social Capital	Measures of social capital on county and college level	Social Capital Atlas based on Chetty et al. (2022a) and Chetty et al. (2022b)
Implicit Racial Attitudes	Average Race IAT Score by County	Project Implicit (Xu et al., 2022); County-level estimates by Liz Redford height

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 females as males, 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 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.[40] 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.

[40] 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.

Category	Measure	This Study	LinkedIn USA	US Census
Age	Median	32	30-49	38.8
	Share 18-29	40.4%	23.9%	15.7%
	Share 30-49	47.6%	34.8%	19.0%
	Share 50-64	10.4%	31.9%	19.0%
	65+	1.4%	9.3%	16.7%
Gender	Female	52.3%	46.1%	50.5%
	Male	47.6%	53.9%	49.5%
Race	White	69.0%	63.8%	57.8%
	Black	6.0%	12.0%	12.1%
	Hispanic	13.0%	13.7%	18.9%
	Other	12.0%	10.5%	11.2%
Education	Highschool or Less	20.1%	11.5%	36.8%
	Some College	11.5%	13.0%	14.9%
	College +	68.3%	75.5%	48.4%

Table D.1: Caption

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 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.

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 that 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 individuals earning 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	87,363	67,330	56,893	19,362
SALARIES BY CAREER LEVEL				
CEO	167,753	181,804	51,060	123
President	182,513	172,410	70,387	1,113
Director	132,081	128,099	63,346	2,089
Senior	113,919	96,330	56,125	1,402
Assistant	53,969	41,536	35,653	2,254
SALARIES BY EDUCATION				
Degree: None	70,672	56,334	45,364	3,895
Degree: Some College	78,988	60,082	53,162	2,213
Degree: Associate	64,943	52,920	39,278	734
Degree: Bachelor	86,745	67,683	56,463	7,757
Degree: Master	108,272	91,581	62,452	4,154
Degree: PhD	116,812	117,304	62,716	609
Forbes: Top 100	113,650	96,735	67,920	2,329
Forbes: Top 200	101,853	81,120	63,024	5,294
Forbes: Ranked	96,312	76,557	61,098	9,667
Forbes: Not Ranked	78,440	59,392	50,822	9,695
SALARIES BY LINKEDIN VARIABLES				
Num. Skills Verified: >Median	105,445	87,033	61,164	8,028
Num. Skills Verified: <Median	73,019	59,392	47,898	7,898
Num. Skills: >Median	97,639	79,287	58,851	8,170
Num. Skills: <Median	80,649	61,991	54,330	7,756
Num. Contacts: >Median	105,013	85,178	63,263	9,449
Num. Contacts: <Median	70,916	58,408	43,901	9,365
SALARIES BY DEMOGRAPHICS				
Age: <30	72,244	59,021	45,889	6,796
Age: 30-39	90,966	72,414	56,714	5,696
Age: 40-49	108,422	90,992	65,901	2,384
Age: >50	113,635	96,770	68,284	2,012
Female	75,625	59,392	50,046	9,774
Male	100,272	79,817	61,378	8,895
Black	79,165	63,659	49,156	854
White	88,674	67,775	58,613	12,298
Asian	96,391	78,720	62,017	1,692
Hispanic	76,259	59,453	48,169	2,259

Table D.2: Salary Statistics based on Job Titles and Glassdoor.com / Payscale.com

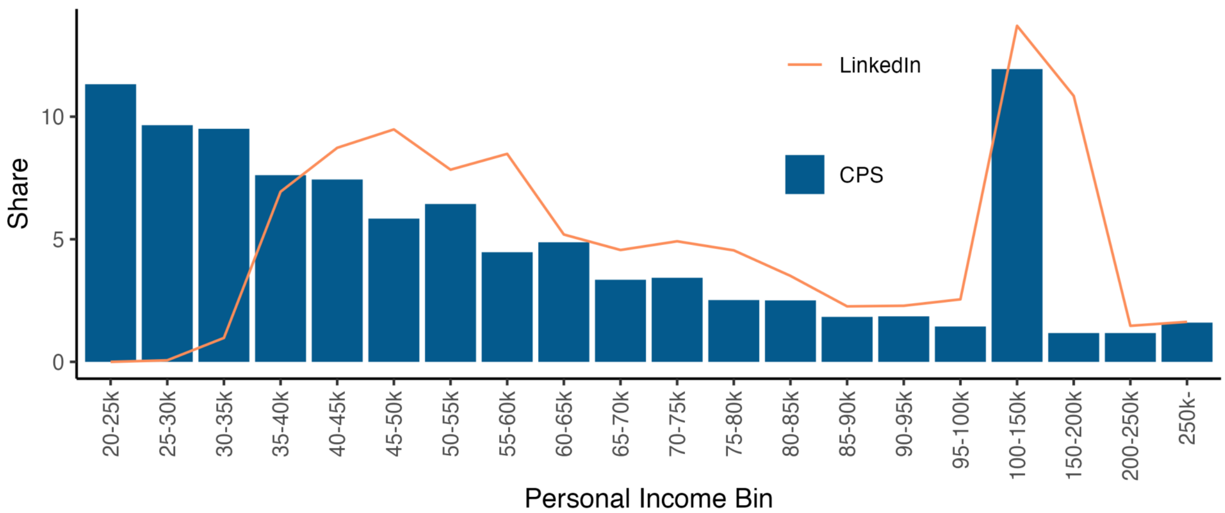

Figure D.1: Income Distribution: LinkedIn Sample vs. Census

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 3.1), 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 to participants before addressing the issues not directly affecting participants.

E.1 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.[41] 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

[41] see LinkedIn's Statistics Page (2023)

April 2023. Thus, the creation of fake profiles does not considerably change the number of 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 researchers[42] and, increasingly, lawmakers[43], 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 making a decision whether to accept our profile's connection request or not. However, the cost of this decision is very minor as it takes just seconds to decide whether to accept a connection request or not. 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

[42] see, e.g., Jeff Hemsley's comment in the Columbia Journalism Review, 2019.
[43] see Center for Democracy and Technology, 2023.

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 only make accessible to contacts. Thus, all the data we obtain is data participants volunteered to be 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.

E.2 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.[44] 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 3.3). 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.

[44]See for example Hern (2017)

E.3 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 better understand 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, minorities 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 to study discrimination in the context of our study. 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 work 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.

E.4 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 a few seconds spent on 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, minorities and the majority 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.

F.1 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 A.I.-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.

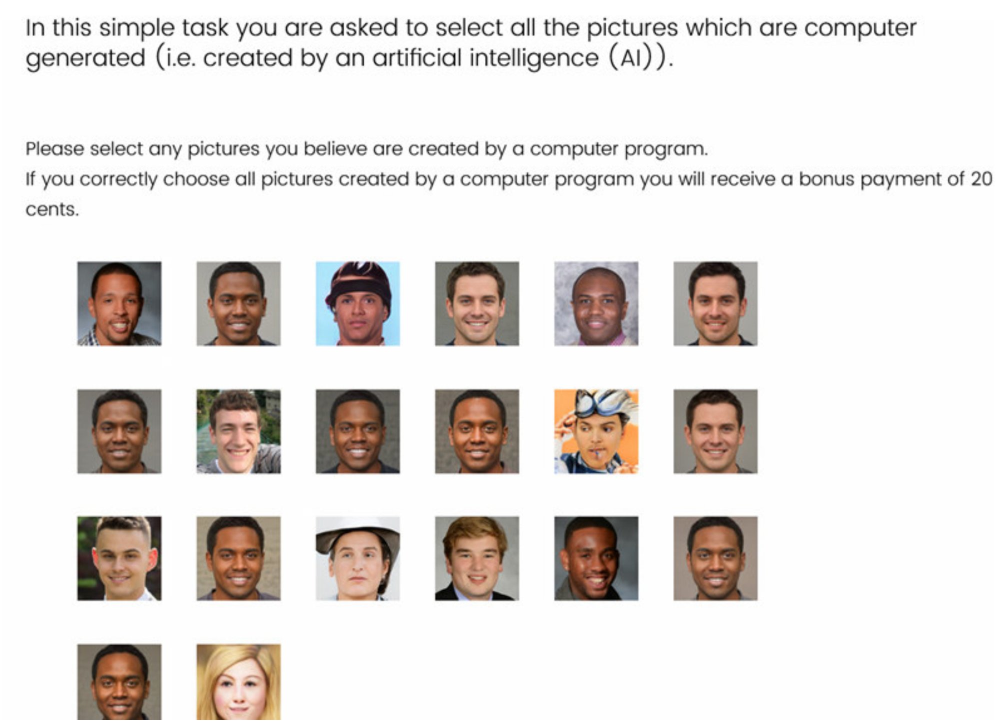

Figure F.1: Screenshot of the Captcha task.

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 female, 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.

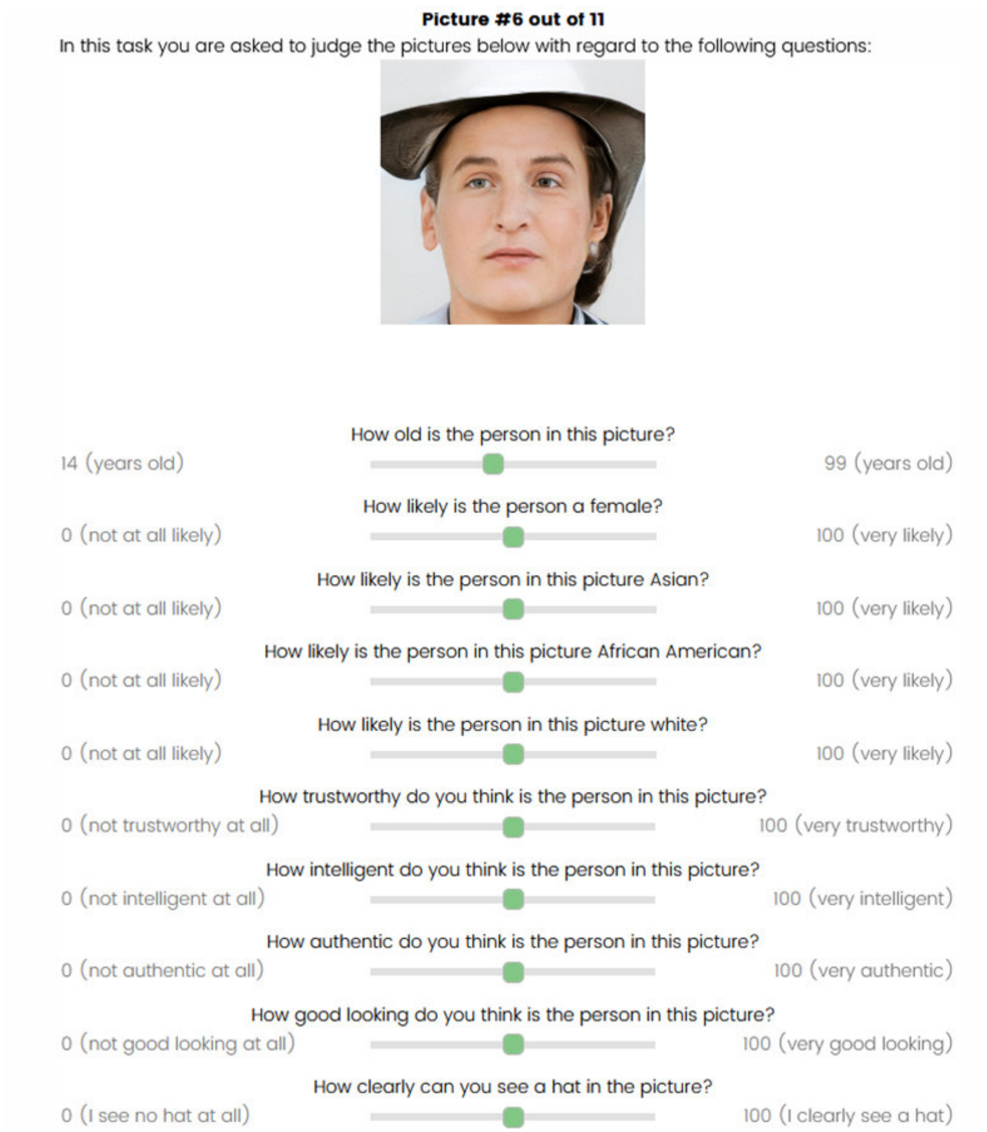

Figure F.2: Screenshot of the individual rating task.
The figure shows a screenshot of the individual rating task for the obviously fake picture.

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.

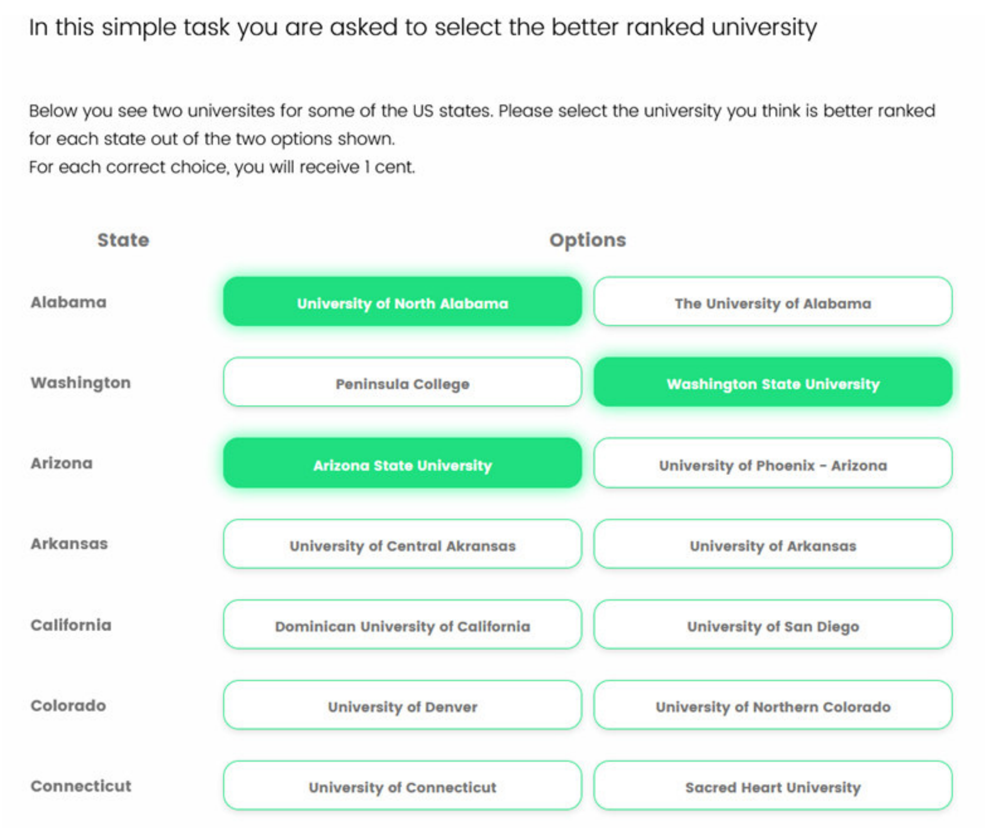

Figure F.3: Screenshot of the university ranking task.

F.2 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.[45] 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.

F.3 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.

F.3.1 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 A.I.-generated pictures have been selected significantly less often as fake compared to obviously computer-generated pictures. More importantly, our A.I.-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 A.I.-generated pictures are sometimes considered less fake than real pictures.

[45]Requesters 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.

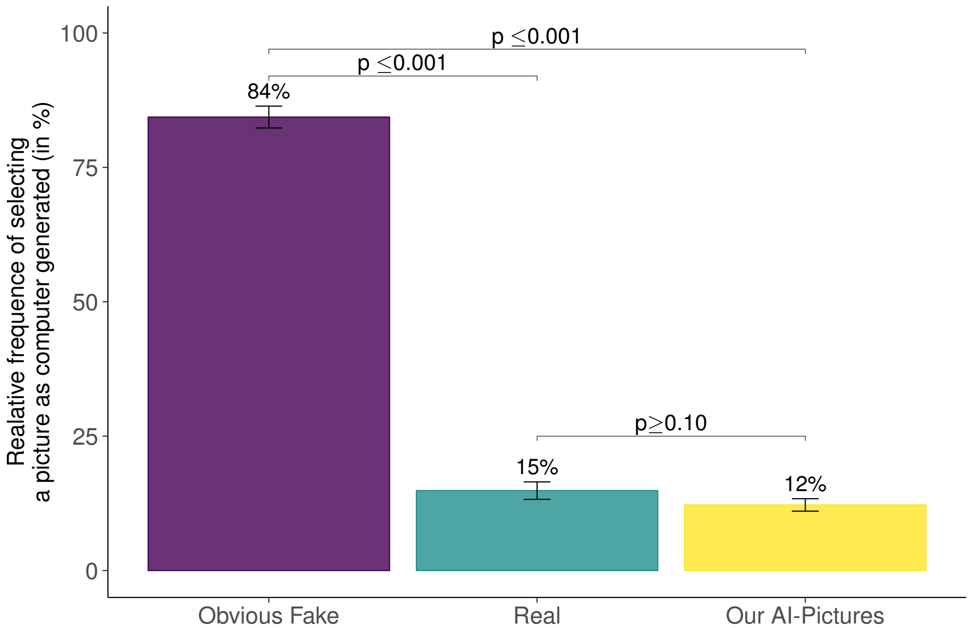

Figure F.4: Captcha-task: Detecting fake pictures

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: $\cdot p<0.10; ^{*}p<0.05; ^{**}p<0.01; ^{***}p<0.001$.

Table J.21 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 A.I.-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 A.I.) are selected as fake at different rates. Table J.22 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.21 and J.22. 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.

F.3.2 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.

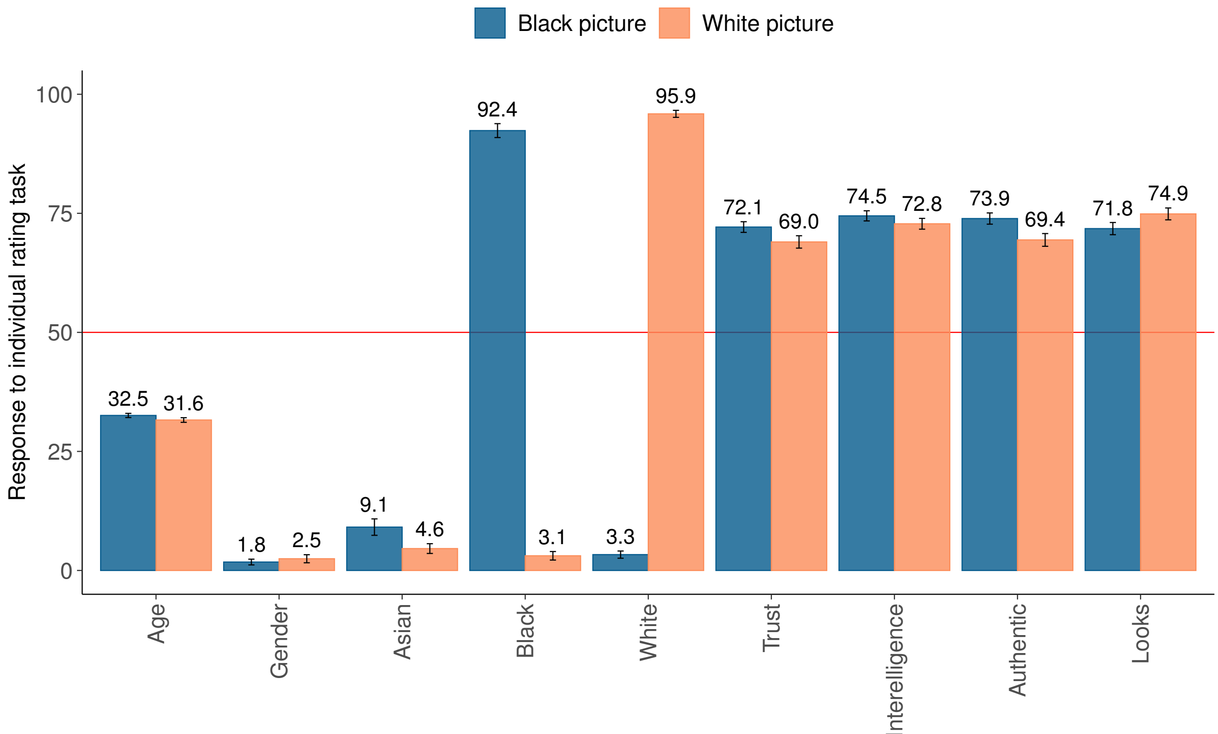

Figure F.5: Average classification of profile pictures

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 female (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.

Table J.23 and Table J.24 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.

F.3.3 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 overage, participants correctly identified the better universities in 37 states. In 8 states, 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.25 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.

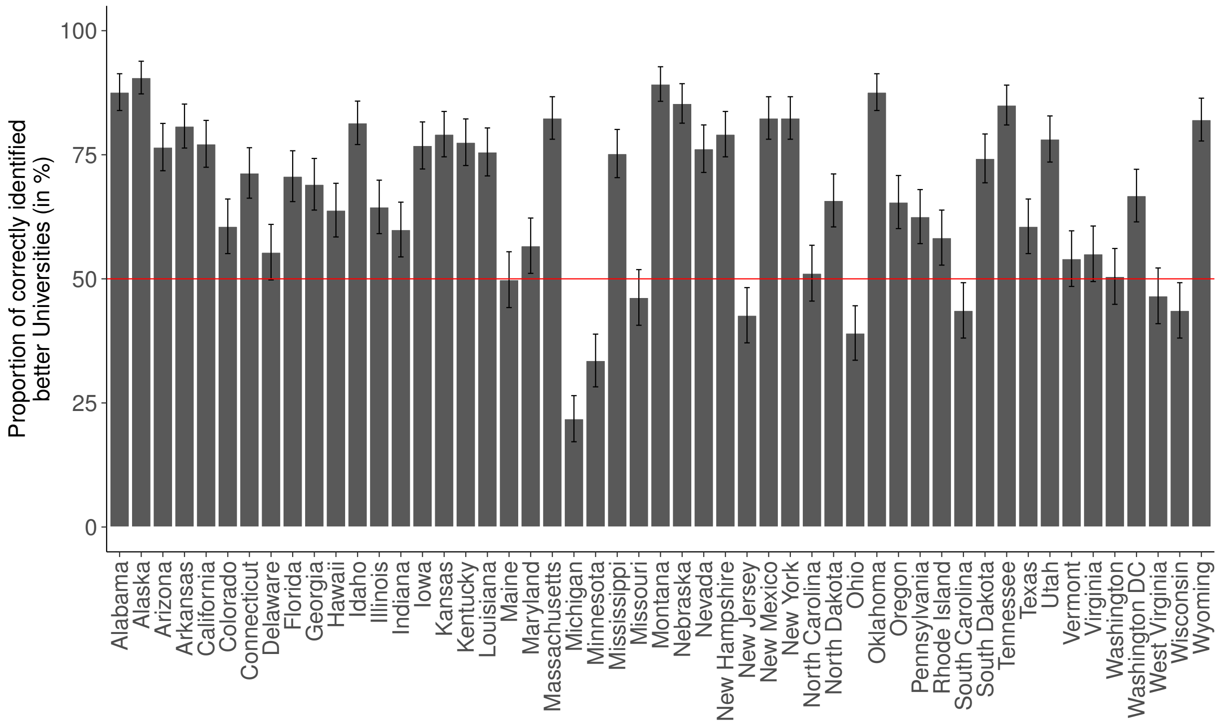

Figure F.6: Propensity of correctly identifying the better ranked university as such.

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

G.1 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.[46] The figure also shows that having attended a better university does not shield one from this experience. The general pattern and the gap in connections are present for both types of profiles.

[46]Obviously, 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 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.

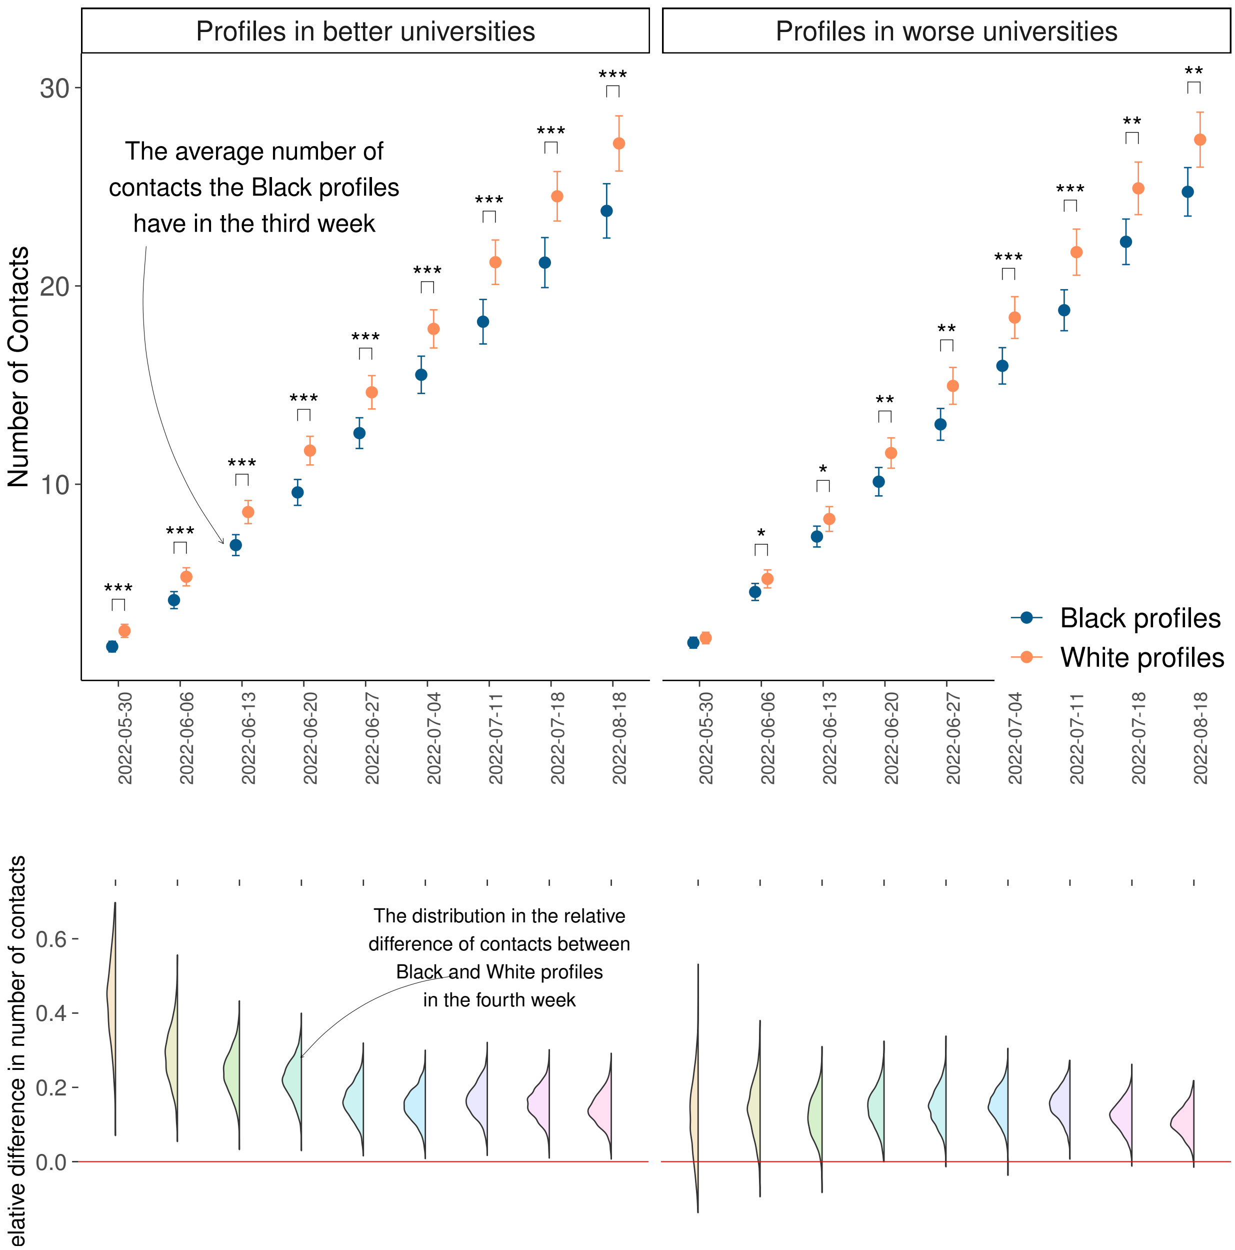

Figure G.1: The evolution of the number of contacts by race and profile quality.
The figure depicts the number of contacts by the week of the experiment for Black and White profiles separately. The left panel denotes results for profiles attending worse universities, while the right panel denotes profiles 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 starts directly at the beginning of the experiment and increases over time (see Table J.2). However, we don't 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.

G.2 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 4.2, we further provide evidence of geographic variation by drawing on targets' home counties.

Figure G.2 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.

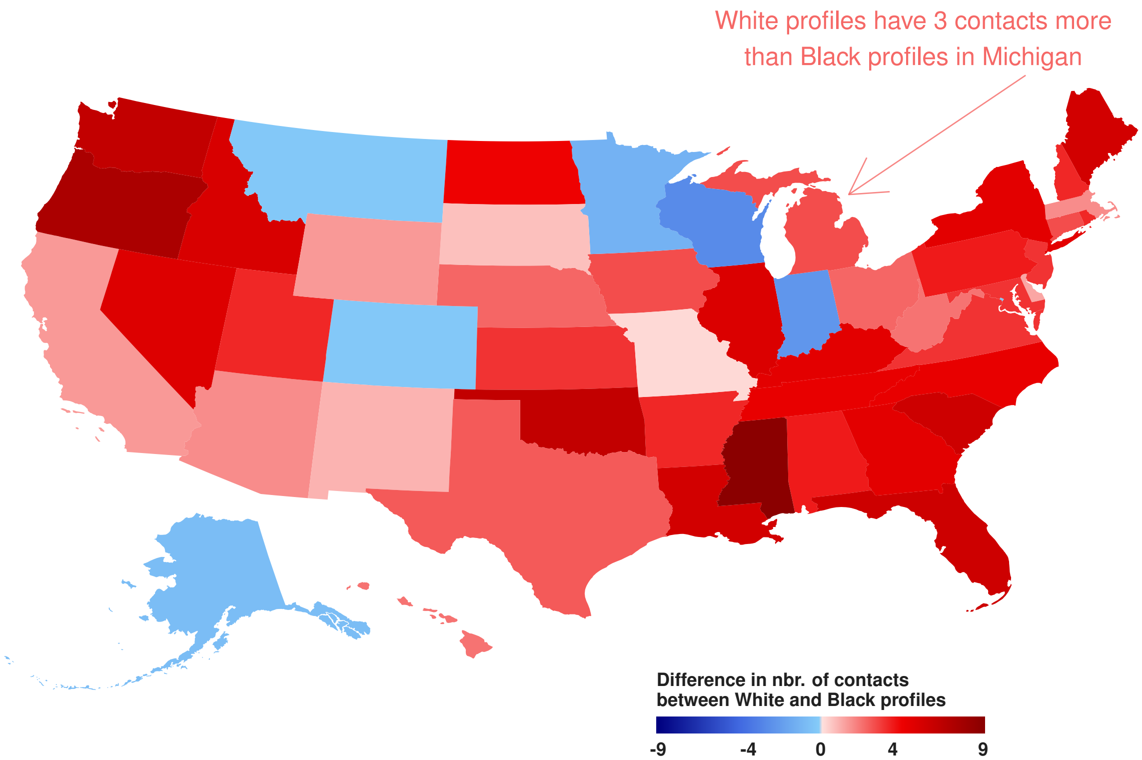

Figure G.2: Difference in the number of contacts by race in each state

G.3 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.3 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 associated degree or a bachelor's degree as their lowest degree) have an almost 5% higher probability of accepting a contact request, and comparably people without a degree have a 6% lower probability of accepting a contact request. Females also seem to have a slightly lower probability of accepting a request (about 1.5%). Further, one standard deviation increase in age reduces the probability of accepting a request by roughly 4%.

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% and 4%, 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% more likely to accept contact requests. In contrast, retired people and users who have a managing position (director, president) are almost 10% and 5% less likely to accept a connection request, respectively.

The area the person lives in 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% higher probability of accepting our connection request. Two other important characteristics of our profile predicting 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% higher, and one standard deviation increase in the likelihood the person on the profile picture is considered Black reduces the acceptance rate by 2%. Notably, the quality of the university our profile attended does not impact the acceptance probability.

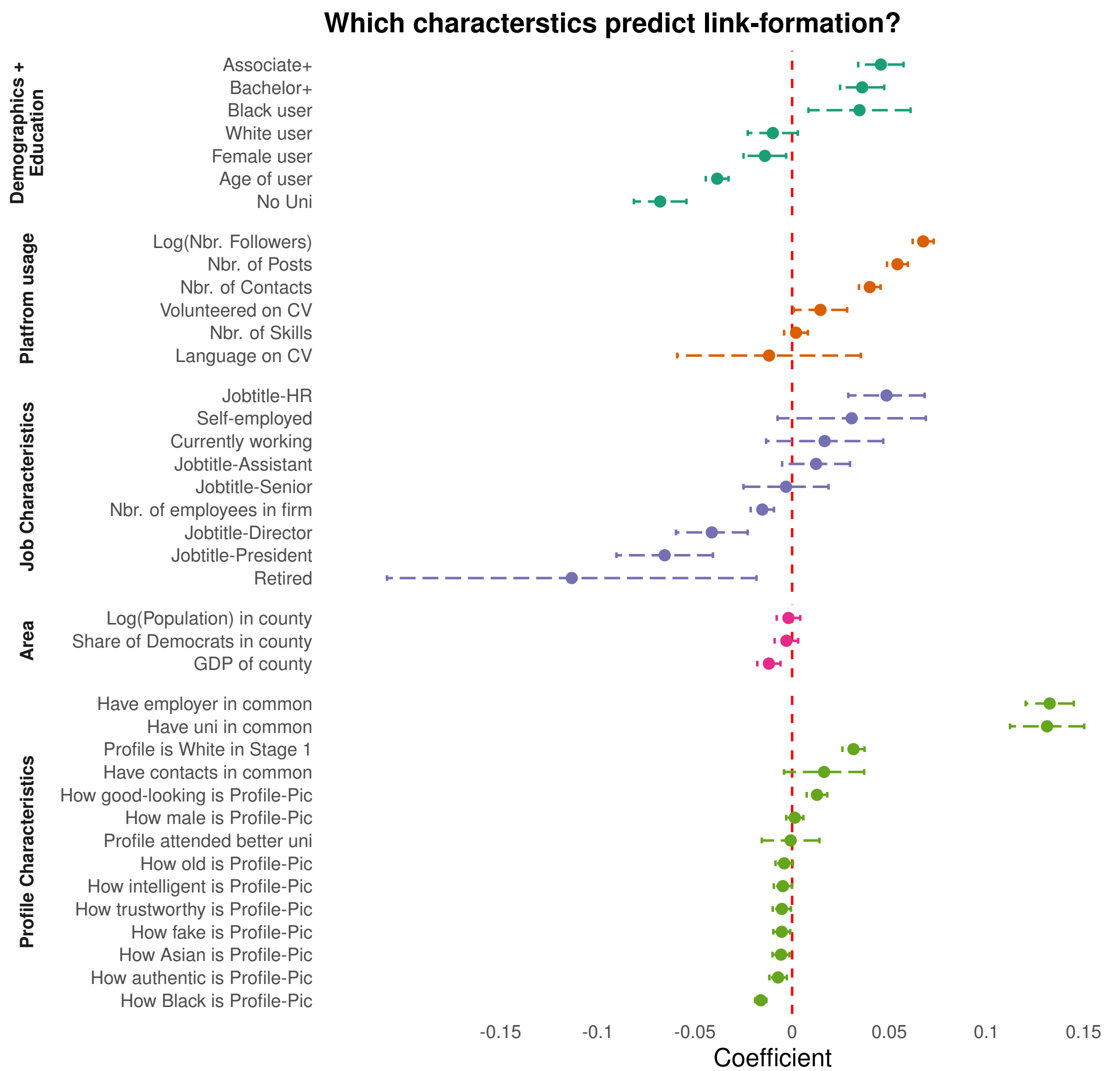

Figure G.3: Predictors of acceptance

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.

G.4 Drivers of discrimination

In the main part of the paper, we have presented multiple relevant correlates of discrimination. In this section, we first use the approach of Wager and Athey (2018) and Athey et al. (2019) to study causal heterogeneous treatment effects using causal forests (Section G.4.1). The results for these exercises are found in Section G.4.1. Next, in Section G.4.2, we zoom in on strong predictors of discrimination: age, gender, and race.

G.4.1 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. Figure G.4 shows the correlation between included covariates. Second and most importantly, many variables are not available for the entire set of observations. In total, 77.5% of observations have full data and are thus included in this section.[47] 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 0.031 in the full and 0.033 in the reduced data set.

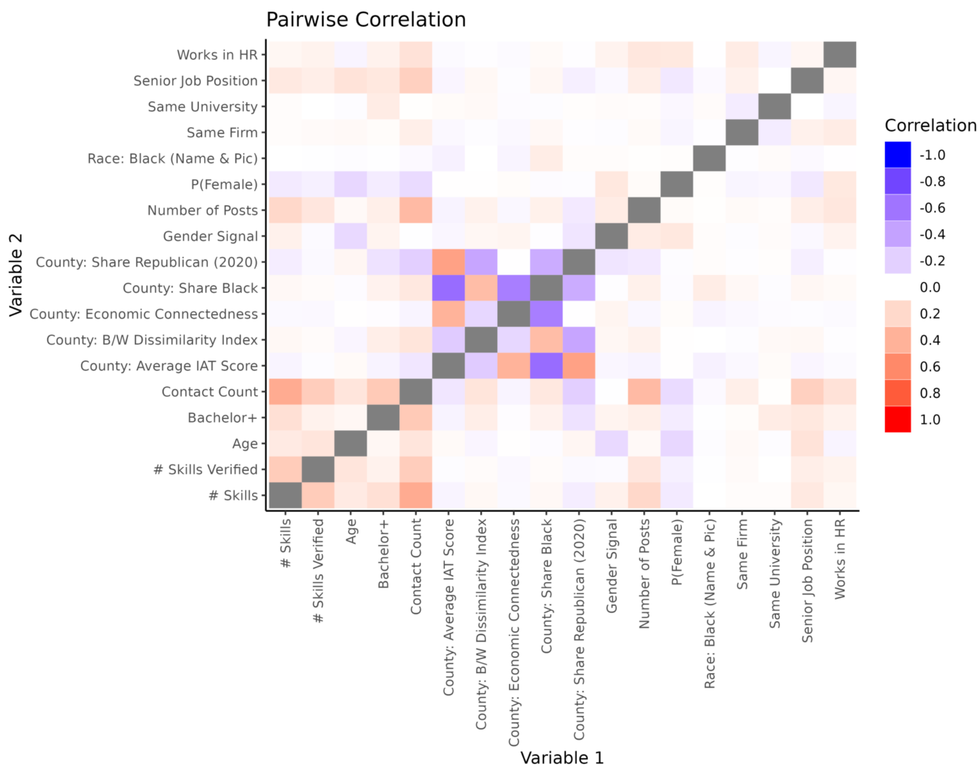

Figure G.4: Correlation between included variables in Causal Forest

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 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.

[47] While 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.

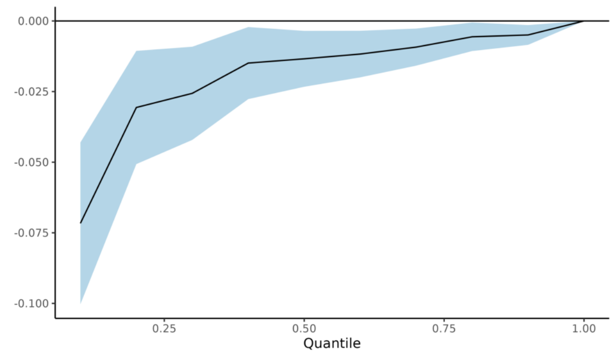

Figure G.5: Target Operating Characteristic (TOC) curve

Note: The above graph compares the average treatment effect on treating everyone with predicted CATE above the qth quintile to the average treatment effect (ATE).

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 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., 2021). 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 qth quintile of CATE. Note that in our context lower CATE mean higher 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 qth quintile of predicted CATE to treating everyone, i.e., the Average Treatment Effect (ATE). The curve is shown in Figure G.5 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: -0.02; $p < 0.01$).

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 0.80 ($p < 0.01$). This suggests that, on average, the ATE is slightly overestimated on the test data. The second coefficient is equal to 0.93 ($p < 0.01$). 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, though it might be slightly overestimated on the test data. 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 use the training forest to obtain out-of-bag CATE estimates. Thus, for each observation, 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 5, suggesting substantial heterogeneity with CATE, i.e. predicted gaps in acceptance rates between White and Black profiles, ranging from around -0.15 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.

Next, we estimate heterogeneity across covariates. For this, we follow Athey et al. (2020) by splitting the data by CATE into four tiles, each containing a similar number of observations. Here, tile 1 includes observations with the lowest CATE, i.e., observations predicted to be most discriminating, while tile 4 includes those with the highest CATE. We then run a regression of a given covariate on tile dummies (without a constant) to obtain the mean and standard error of each covariate in a given tile. The result is shown in Table G.1. As in our main analysis, the table suggests that individuals in tiles with lower predicted CATE tend to be more female, younger, more Republican, to have a lower degree, and job position. The results for females and by age are potentially the most striking showing strong changes from the lowest to the highest tiles. The table provides additional insights though, showing that those in tiles with lower CATE tend to live in areas with lower Economic Connectedness, i.e., levels of cross-class interaction. Regarding targets' race, the results suggest relatively little that Black users discriminate less. It further provides interesting insights regarding variables related to increased engagement with the platform and other users. These include the number of verified skills, listed skills, number of posts, and connections. For all four, the results suggest that users with higher CATE, i.e., those predicted to discriminate less, tend to more strongly engage with the platform and others. Finally, it is worth noting that having visited the same university is not, and working for the same firm is only weakly associated with lower gaps in acceptance rates.

covariates	ntile1	ntile2	ntile3	ntile4
P(Female)	0.734	0.609	0.456	0.261
	(0.006)	(0.006)	(0.006)	(0.006)
Bachelor+	0.694	0.682	0.706	0.795
	(0.006)	(0.006)	(0.006)	(0.006)
Senior Job Position	0.124	0.152	0.195	0.28
	(0.005)	(0.005)	(0.005)	(0.005)
Age	29.81	33.35	36.6	38.91
	(0.143)	(0.143)	(0.143)	(0.143)
Contact Count	266.1	279.1	298.6	366.1
	(2.495)	(2.495)	(2.495)	(2.495)
County: Share Republican (2020)	0.403	0.384	0.364	0.328
	(0.002)	(0.002)	(0.002)	(0.002)
County: Share Black	0.167	0.165	0.152	0.162
	(0.002)	(0.002)	(0.002)	(0.002)
County: B/W Dissimilarity Index	54.97	53.39	54.48	56
	(0.156)	(0.156)	(0.156)	(0.156)
Race: Black (Name & Pic)	0.101	0.108	0.112	0.141
	(0.003)	(0.003)	(0.003)	(0.003)
County: Average IAT Score	0.321	0.317	0.317	0.313
	(0.001)	(0.001)	(0.001)	(0.001)
County: Economic Connectedness	0.814	0.827	0.841	0.85
	(0.002)	(0.002)	(0.002)	(0.002)
# Skills Verified	25.59	25.43	30.68	54.24
	(1.445)	(1.445)	(1.445)	(1.445)
# Skills	17.16	16.53	17.83	22.16
	(0.189)	(0.189)	(0.189)	(0.189)
Works in HR	0.098	0.096	0.087	0.073
	(0.004)	(0.004)	(0.004)	(0.004)
Same University	0.106	0.098	0.098	0.077
	(0.004)	(0.004)	(0.004)	(0.004)
Same Firm	0.076	0.079	0.089	0.101
	(0.004)	(0.004)	(0.004)	(0.004)
Number of Posts	0.836	0.874	0.892	1.074
	(0.015)	(0.015)	(0.015)	(0.015)
Gender Signal	0.137	0.138	0.131	0.143
	(0.005)	(0.005)	(0.005)	(0.005)

Table G.1: Average of covariates by CATE tile

Note: The table describes the average of each covariate by tile. It shows the result of a regression of the respective covariate on tile dummies (without a constant). Each tile contains a fourth of observations. Tile 1 contains observations with the lowest CATE according to the causal forest, and Tile 4 those with the highest CATE. Note that lower CATE mean stronger gaps in acceptance rates. The values in brackets are standard errors of the respective tile dummy from the linear regression (OLS). Colors are assigned based on the position of the subgroup's mean value relative to the standardized empirical distribution of its variable: (x - mean(x))/sd(x).

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.[48] We then restrict our attention to nine variables with above-median variable importance. Based on these variables, we estimate a second causal forest with 50,000 trees on all observations. Figure 6 shows the variables included in the final forest.

G.4.2 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.6 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

[48] Note, that this does not necessarily imply that variables with low importance are not responsible for heterogeneity, which is why we also show Table G.1.

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.6, 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 (β=0.04,t(16963)=4.80, $p \leq 0.001$), and also have more followers than young users (β=0.23,t(16938)=18.75, $p \leq 0.001$). Thus, it seems unlikely that younger users discriminate based on their anticipated interaction with a new contact.

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.

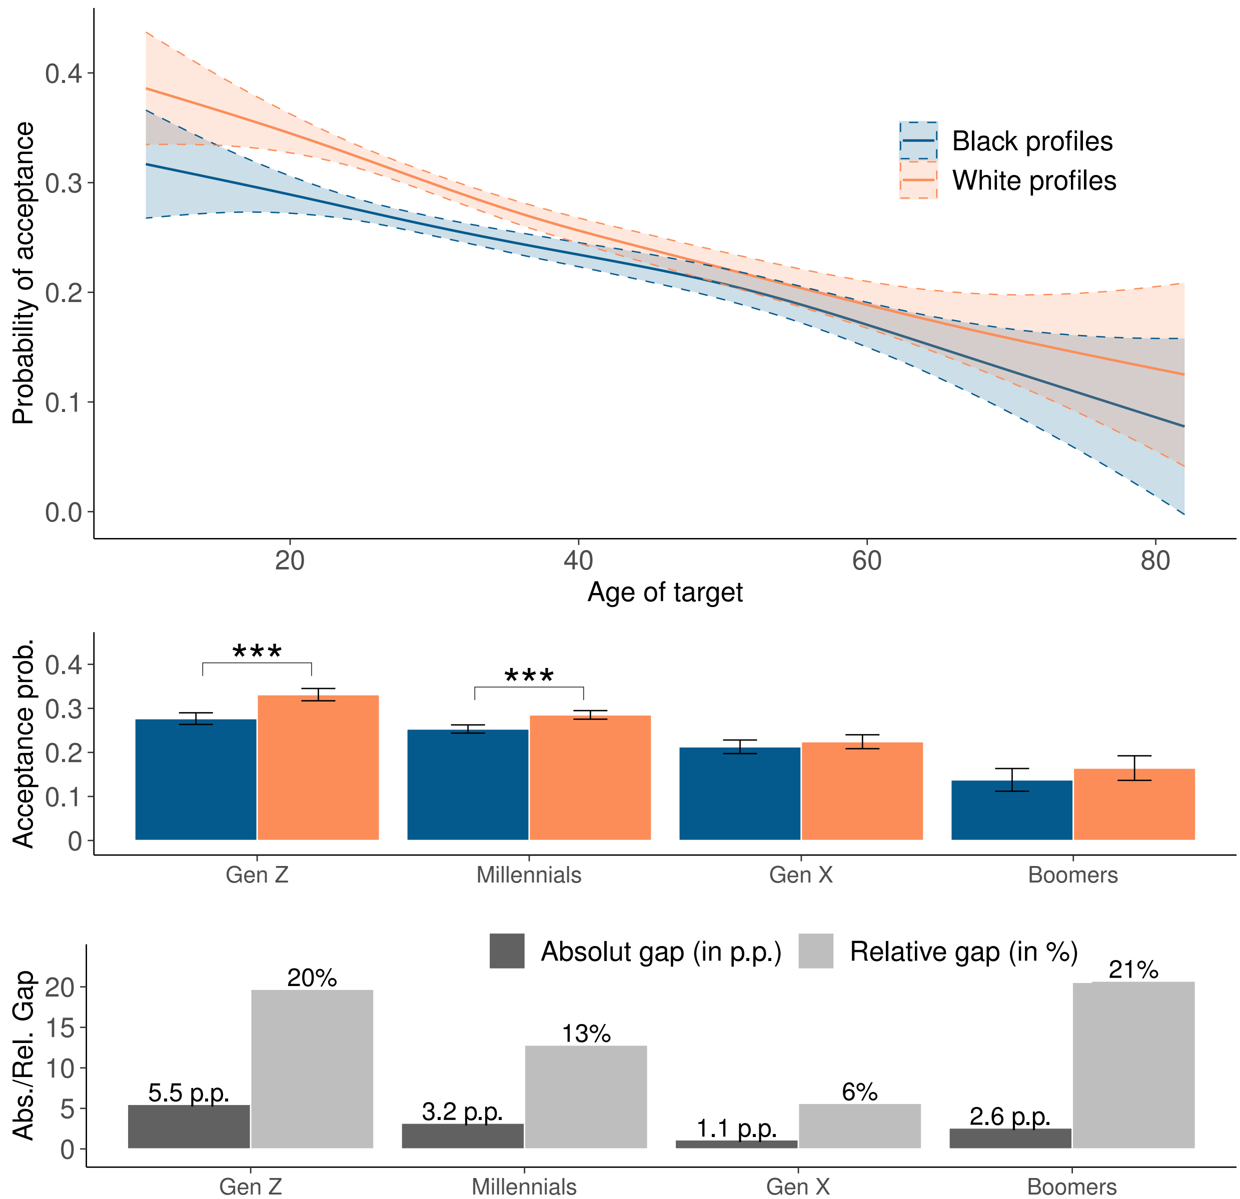

Figure G.6: Acceptance probability of a connection request by a Black and White profile as a function of the target's age.

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: $\cdot p<0.10; ^{*}p<0.05; ^{**}p<0.01; ^{***}p<0.001$.

Gender as a predictor of discrimination The next best predictor of discrimination is gender. As shown in Figure 4, females are discriminating more, not less, than males. Here we take a closer look at how exactly females differ from males in terms of discrimination and study some possible explanations of this phenomenon.

Figure G.7 illustrates the acceptance probability of a connection request of a Black and White profile as a function of the target's probability of being female based on their first name. White profiles have a constant probability of roughly 26% of being accepted independent of the target gender. Thus, males and females seem to accept White profiles to the same extent. For Black profiles, that pattern changes. First, Black profiles have an acceptance rate of 24% if the target is likely male. This acceptance probability, however, reduces monotonically in the target's probability of being female and reaches almost 22% if the target is very likely to be female. Thus, we observe a clear connection gap already for males, but this gap is magnified by females.

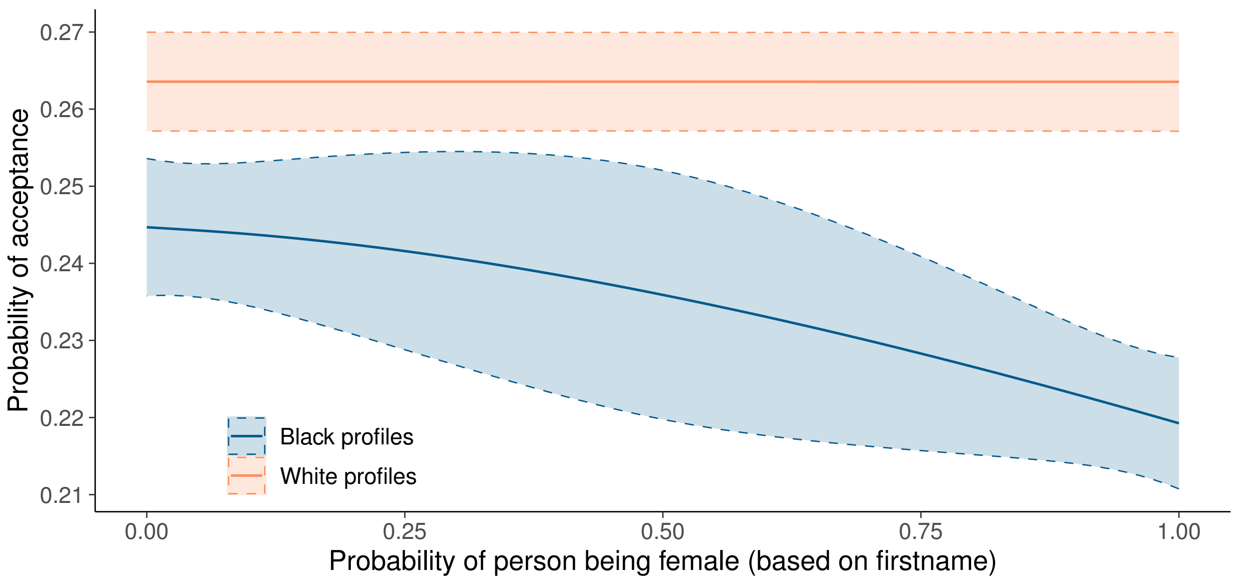

Figure G.7: Acceptance probability of a connection request by a Black and White profile as a function of the target's gender (based on the first name).
The figure illustrates the fitted acceptance probability of a connection request as a function of the target's gender (determined by their first name). 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 females 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 females are less likely to accept a connection request in the first place if 'romantic' advances are anticipated. This possible explanation that females are just generally less likely to accept a connection request can be alleviated by Figure G.7 as it clearly shows that females have the same acceptance probability of a White profile as males. 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.[49] 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 might 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 females than for males. Figure G.8 illustrates the acceptance probability of a connection request of a Black and White profile as a function of the target's probability of being female and the target's probability of being White/Black. First, we see that males 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 females, 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.

[49] See e.g. Insider post or LinkedIn Blog.

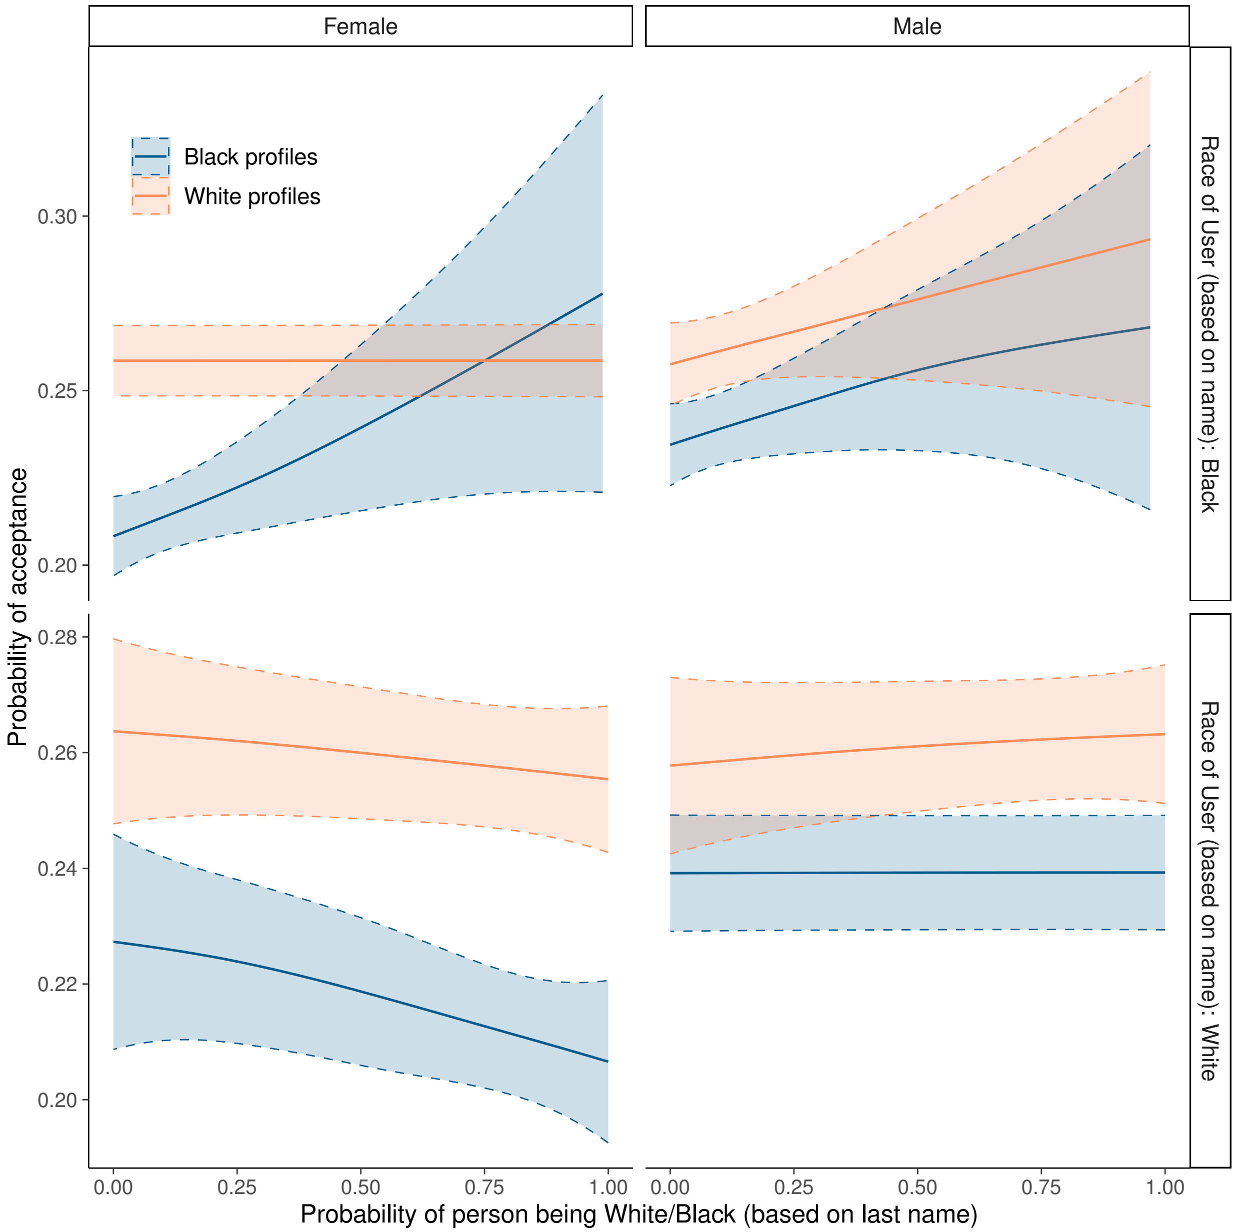

Figure G.8: Acceptance probability of a connection request by a Black and White profile as a function of the target's gender (based on first name) and race (based on last name).
The figure illustrates the fitted acceptance probability of a connection request as a function of the target's race (determined 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 females, while the right panels illustrate the behavior of males. 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 4, race predicts discrimination. Here we focus in more detail on this effect by splitting our measures of race. Figure G.9 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.

Overall, however, we conclude that Black targets are discriminating less. This effect is primarily driven by the behavior toward Black profiles, and we do not find a similar result when focusing on non-Whites. Further, we see from the section above that this race result is mostly driven by females. There, Black male targets do not substantially change their behavior toward Black profiles, while Black female targets do substantially increase their acceptance probability of Black profiles.

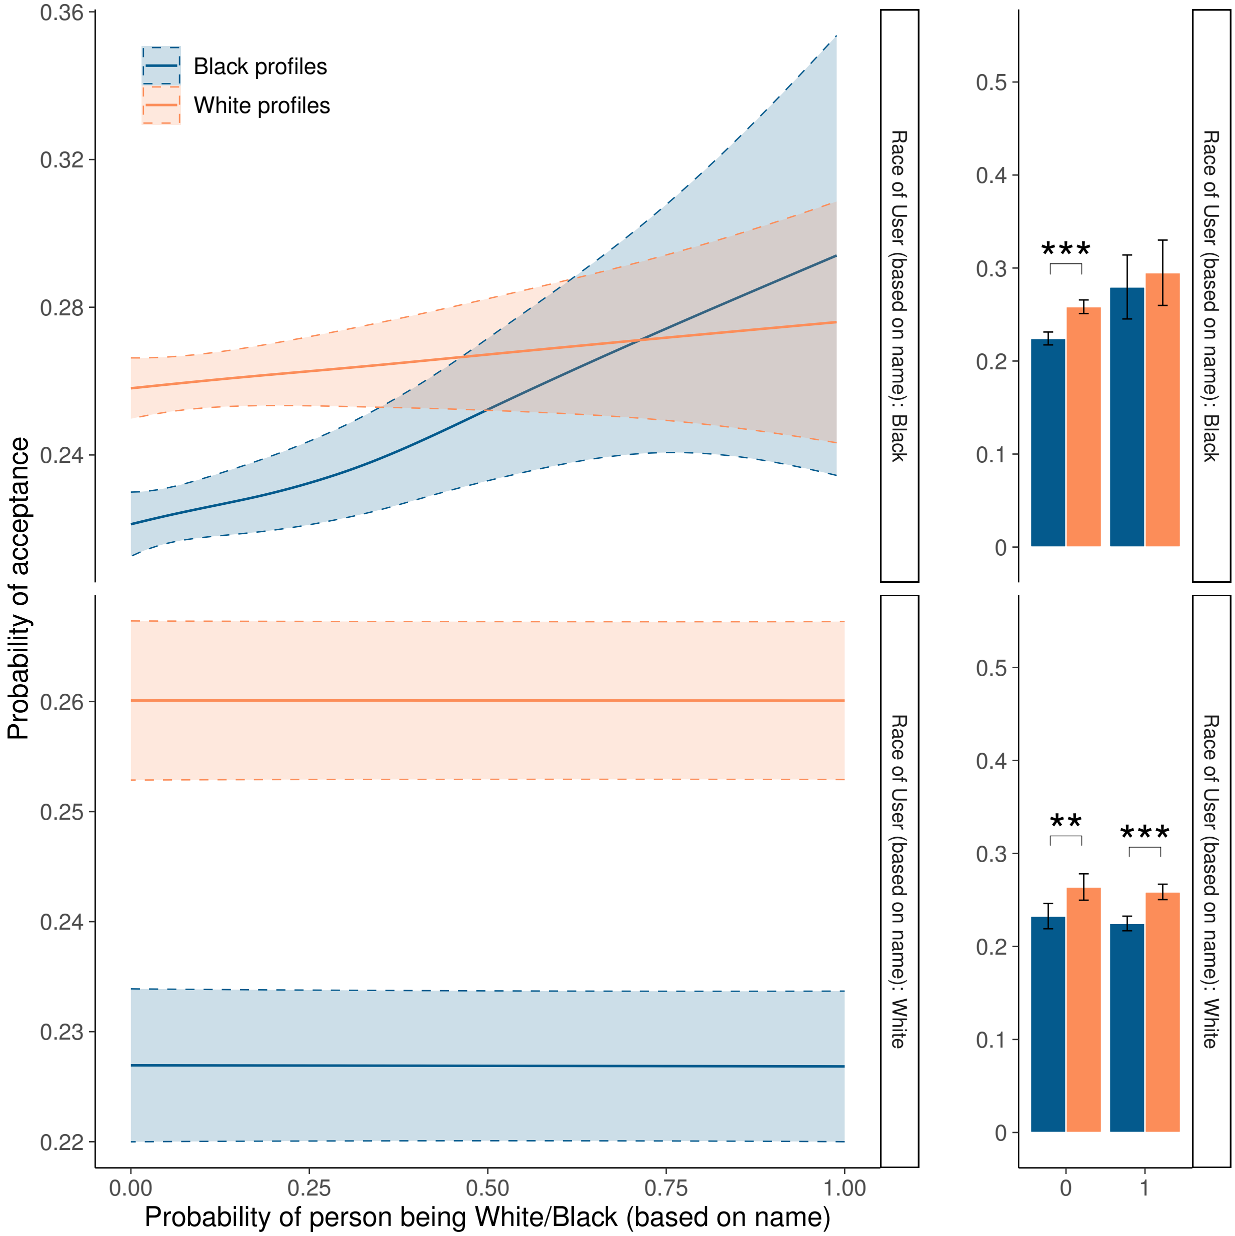

Figure G.9: Acceptance probability of a connection request by a Black and White profile as a function of the target's race (based on last name).

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: $\cdot p<0.10; ^{*}p<0.05; ^{**}p<0.01; ^{***}p<0.001$.

G.5 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.10 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% higher probability of responding to our message, and comparably people without a degree have a 6% lower probability of responding. If our profile and the target attended the same university then they have a 7% 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%. 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% 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 responding to the message.

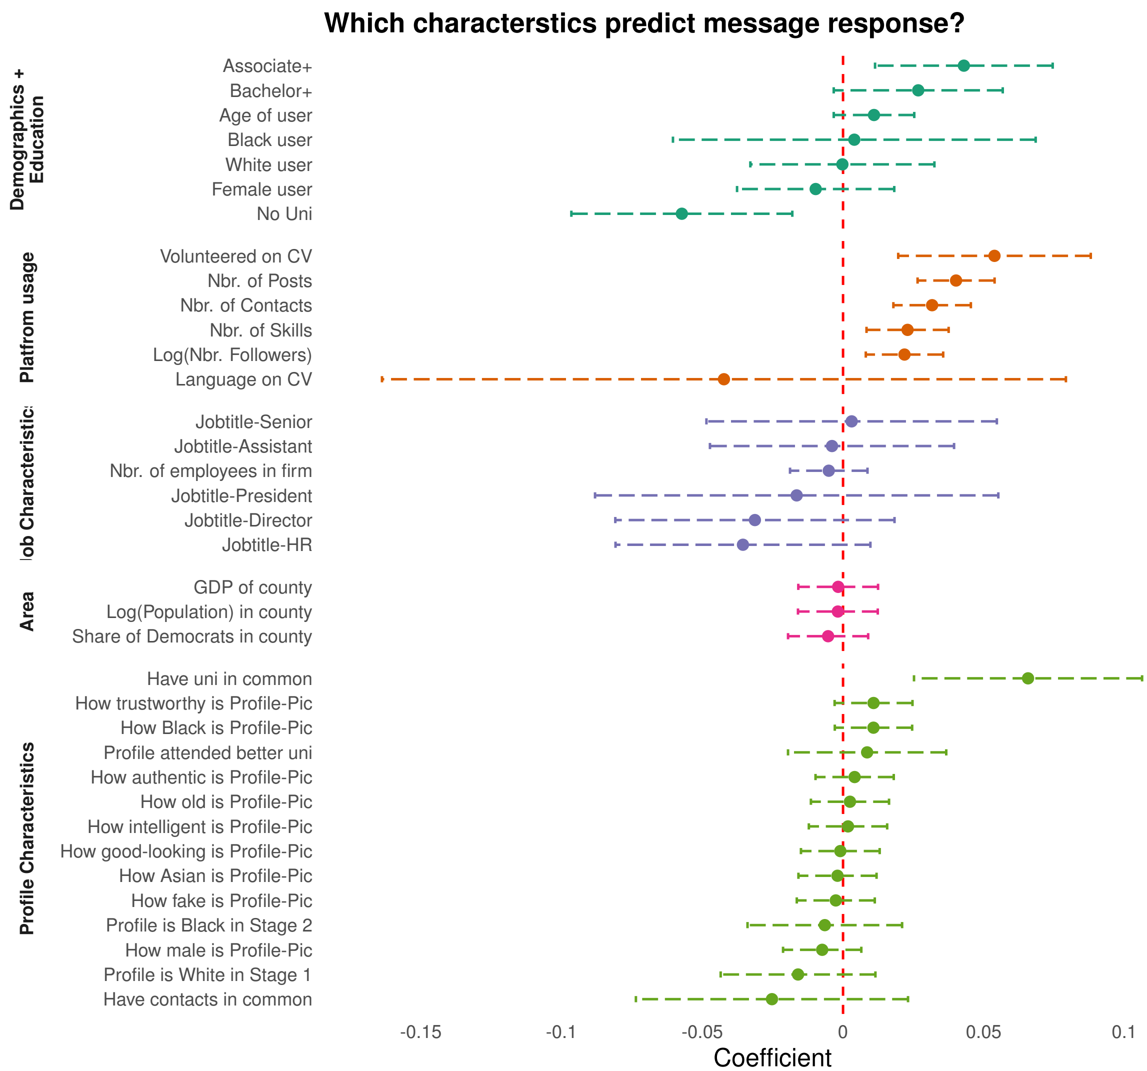

Figure G.10: Characteristics predicting message response.

The figure illustrates the β-coefficients of the following regression: $response_{i,j} = \alpha + \beta \cdot \text{Variable} + \epsilon_i + \epsilon_{i,j}$. 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.

G.6 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.11 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 0.11 (p =0.881) 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: 0.21, 0.42), and remained insignificant even after all messages were dispatched (p =0.779, p =0.579). Regression estimates reported in Table J.15 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.

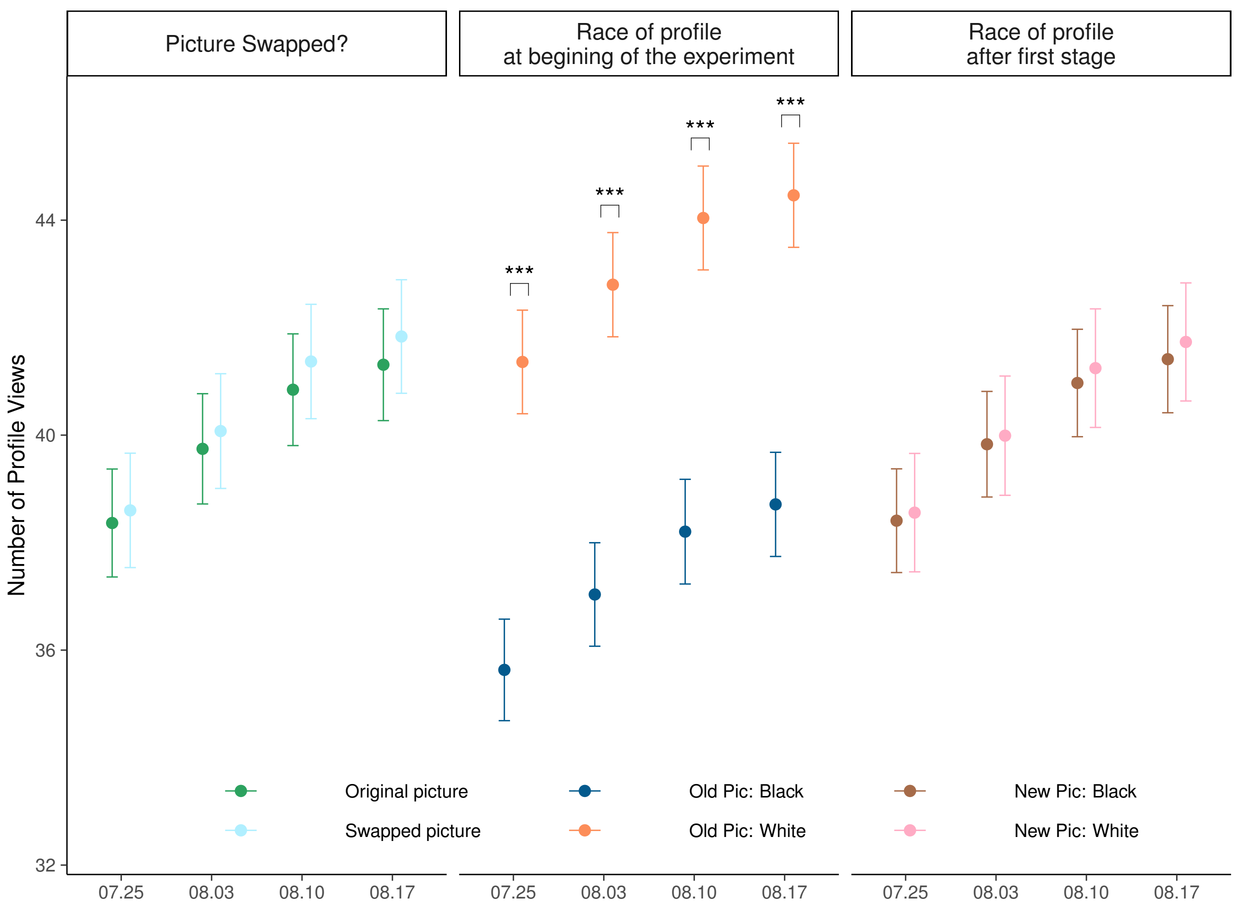

Figure G.11: Number of profile views before and after face-swapping.

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: $\cdot$p<0.10;*p<0.05;**p<0.01;***p<0.001.

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 9486 established links, merely 100 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. 51 suspensions, or 1.02% vs. 1.09% suspensions, p =0.769).

Figure G.12 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.12). 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.16.

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

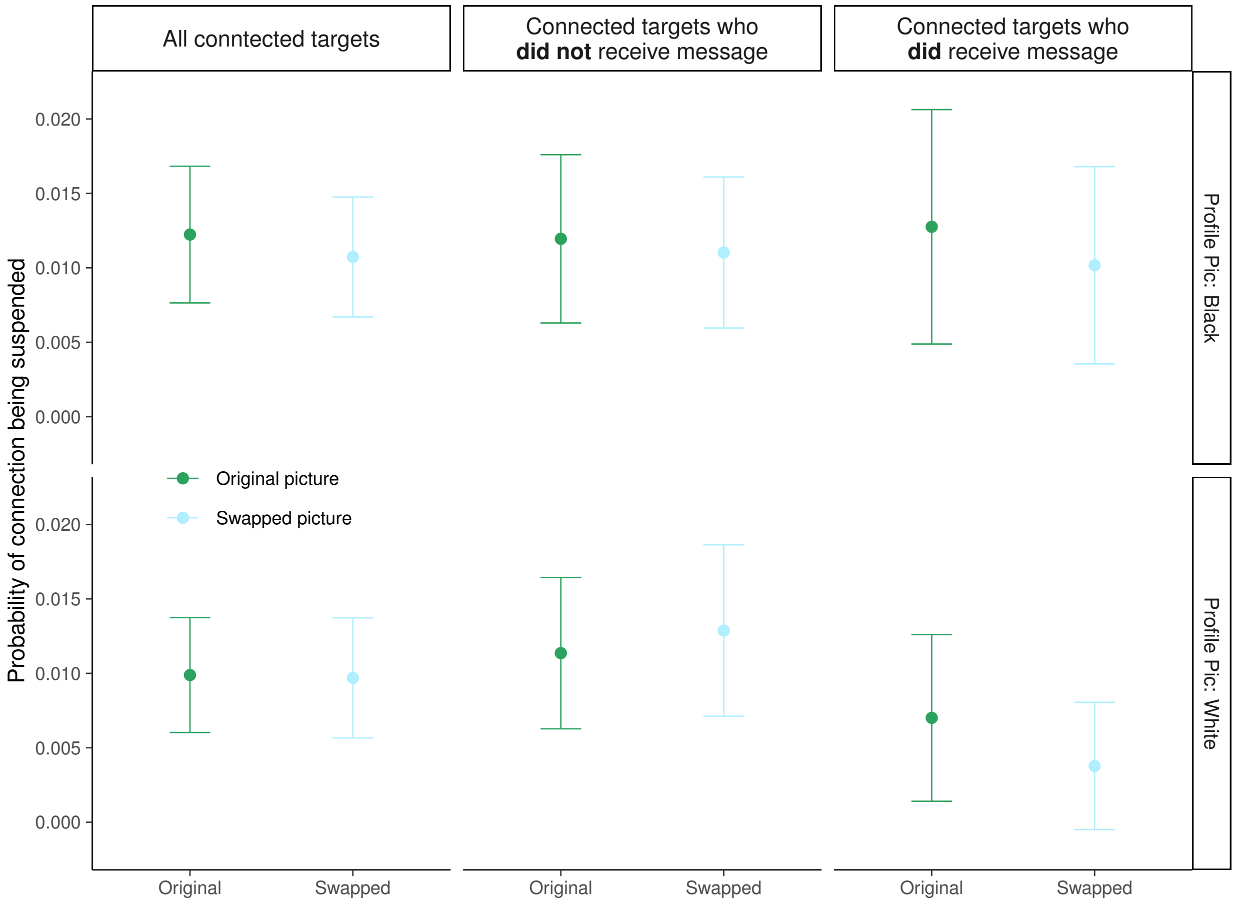

Figure G.12: Suspension probabilities by swapped and original profile pictures.
The figure depicts the probability of a connection suspension for profiles whose picture has not been swapped (i.e., their original picture) in green and profiles whose picture has been swapped to their twin's pictures (i.e., 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: $^{\cdot}$p<0.10;*p<0.05;**p<0.01;***p<0.001.

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.13 presents the response probability, normalized response length (in characters), and the likelihood of the response being highly useful. 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 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.17 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.

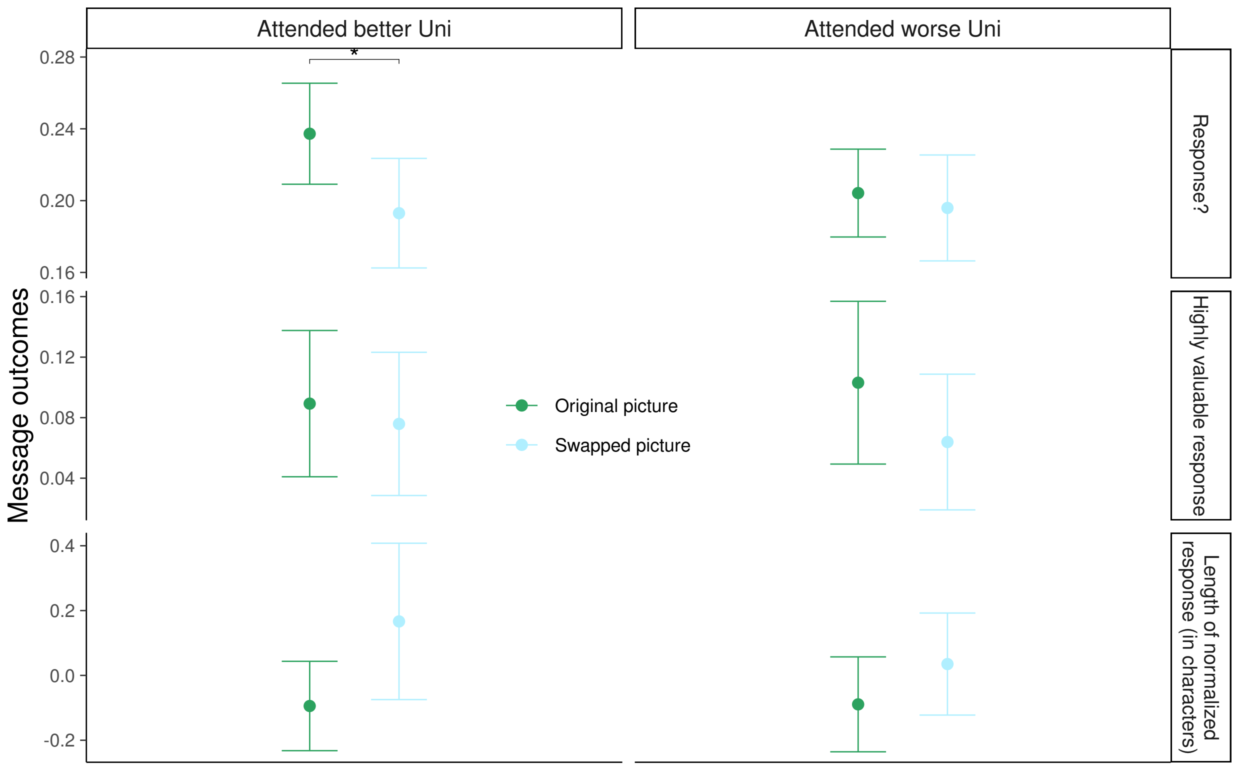

Figure G.13: Response characteristics by swapped and original profile pictures.

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 middle panel illustrates the probability of a response being highly valuable, 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: $\cdot p<0.10$;$^{*}p<0.05$;$^{**}p<0.01$;$^{***}p<0.001$.

G.6.1 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 females who do accept a link are generally more helpful, while males 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.[50] Figure G.14 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.

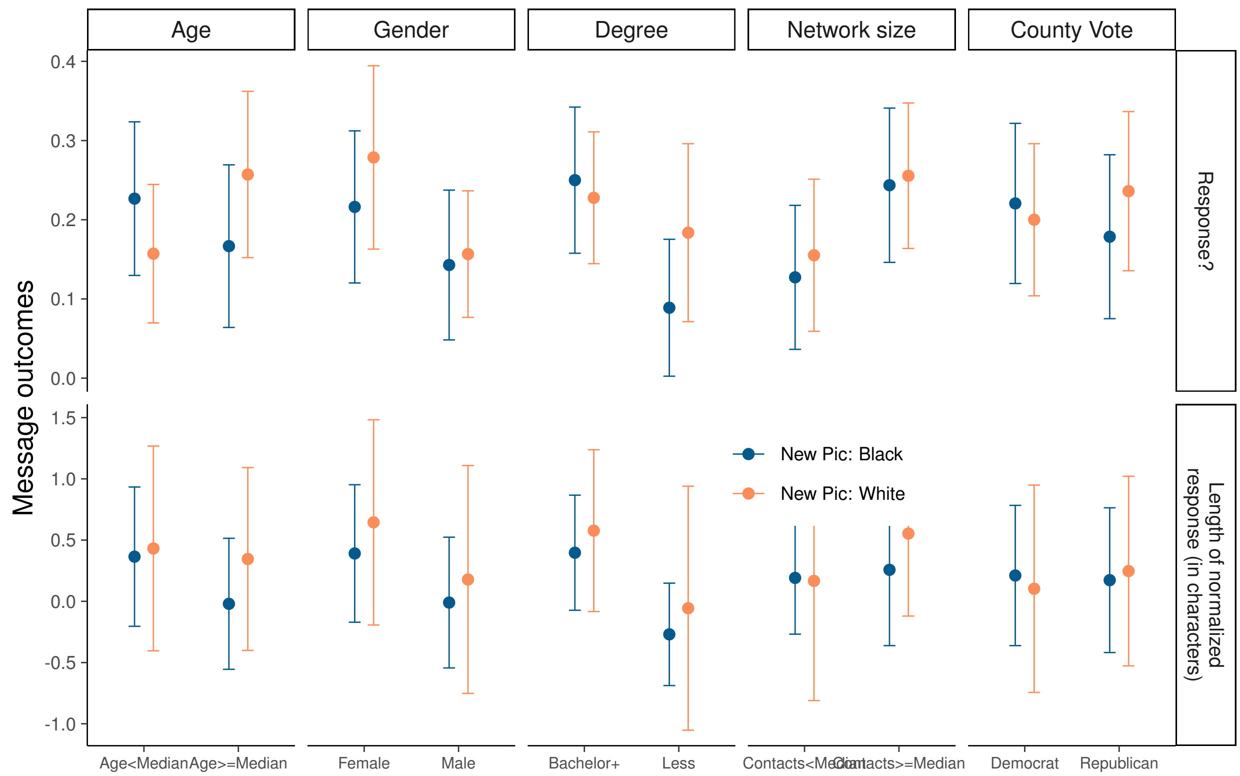

Figure G.14: Response characteristics by race of the profile pictures and basic target characteristics in the second stage of the experiment.

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 profile's requests in stage one, as a broader marker of discrimination. Figure G.15 reveals that there is again not much heterogeneity. Targets who accepted both in stage one 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

[50] We exclude race here due to sample size limitations.

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.14). 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.

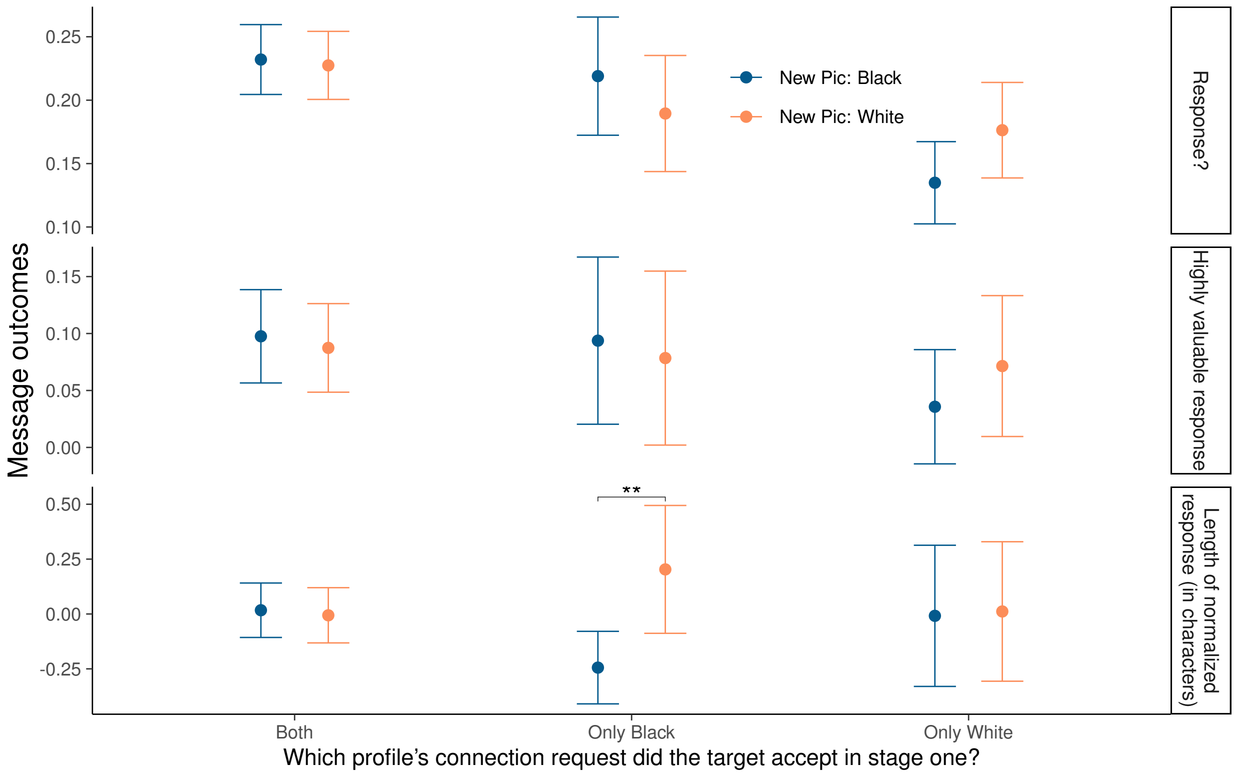

Figure G.15: Response characteristics by race of the profile pictures in the second stage of the experiment.

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, the middle panel illustrates the probability of a response being highly valuable, 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.

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.16 shows the interaction estimates. We find little heterogeneity. Some characteristics interact with the profile's race, but we do not find a clear pattern.

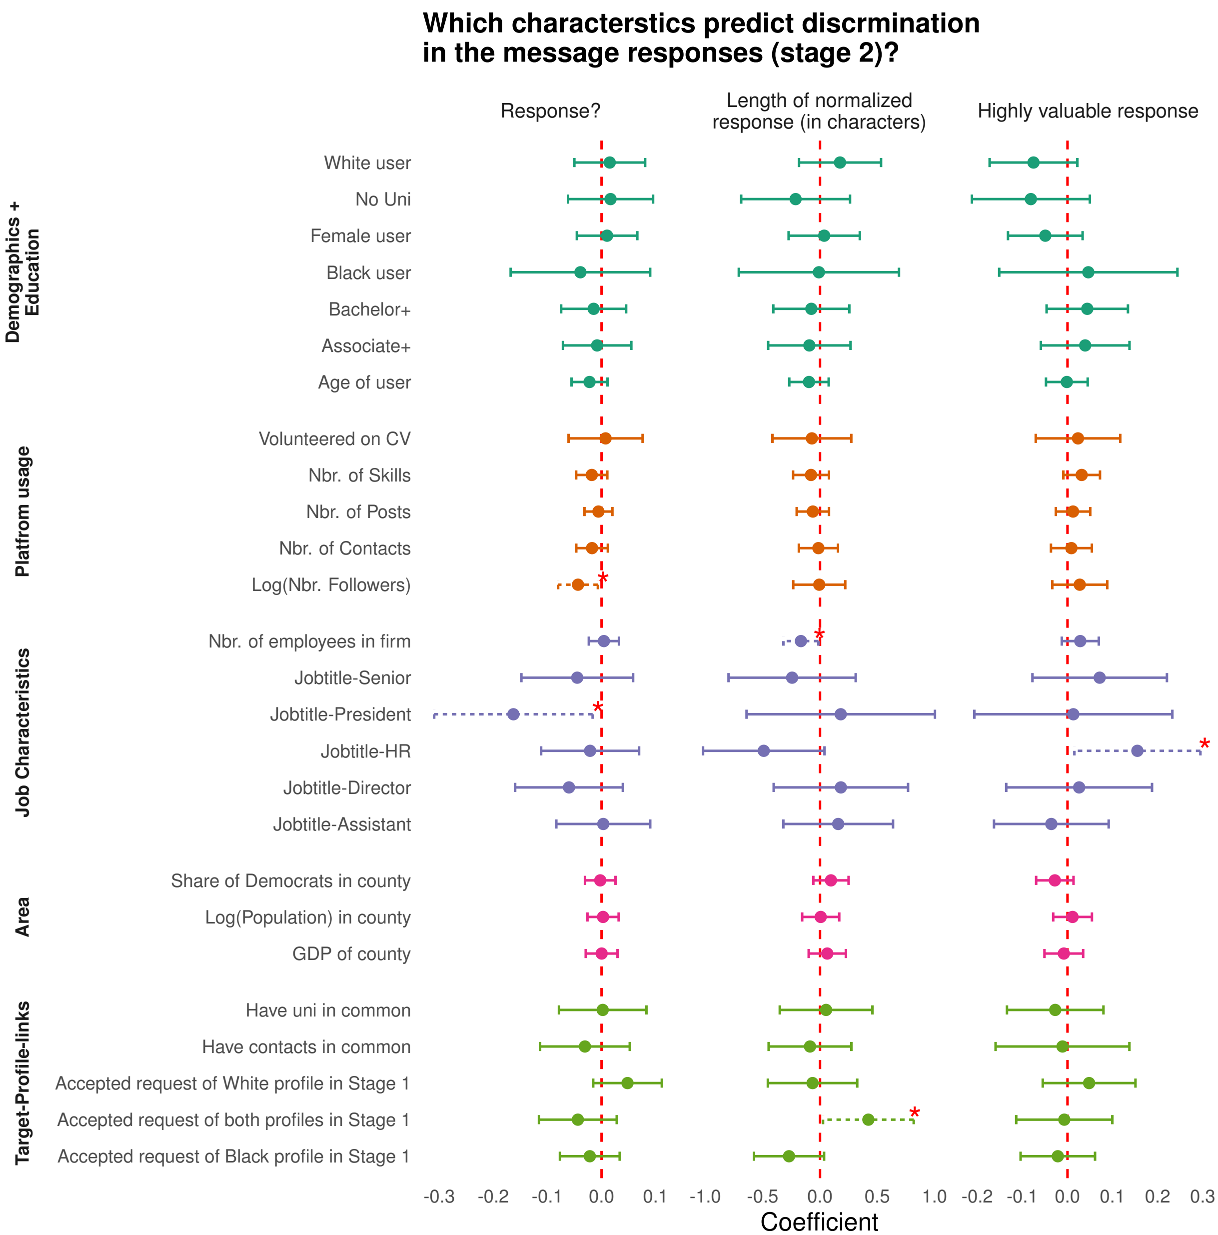

Figure G.16: Correlates of discrimination in message responses.

The figure illustrates the β-coefficients of the following regression: $ResponseChar_{i,j} = \alpha_0 + \alpha_1 \cdot \text{BlackProfile} + \alpha_2 \cdot \text{Variable} + \beta \cdot \text{Variable} \cdot \text{BlackProfile} + \epsilon_i + \epsilon_{i,j}$. ϵ_i is the user random effect with ($\epsilon_i \sim \mathcal{N}(0, \sigma_1^2)$, $\epsilon_{i,j} \sim \mathcal{N}(0, \sigma_3^2)$). *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), and the probability of the response being highly valuable. 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.

G.7 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.10 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

In order 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. In order not to 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.[51] Given some overlap, this results in 2,143 Emails sent.[52] We purposefully designed the survey to be very short in order 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 of 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 female and 71% to be male. 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.

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

[52] of these, 23 Emails could not be delivered.

By the end of the survey, we also asked participants to indicate how confident they were in their 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.[53] 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 with regard 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, 18.1% more contacts relative to Black profiles. This number is relatively close to the actual gap of 13%. In order to have a better understanding with regard to 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 22.3% for boomers to 6.5% 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 17% relative rate, while they expect this relative acceptance gap to be only 8.9% 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.

[53]The 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.

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 (-0.9% relative gap), while they expect non-Black users to highly prefer White profiles with a 14.4% 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 males to discriminate substantially more than females. Specifically, they expect that male users have a relative gap of 15.3 % in favor of White profiles and they expected this gap to reduce to 10.5% ($p \leq 0.001$) for female users. In our data, however, the reverse is true. Male users 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, female users 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. females, males, 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 males to discriminate substantially more than females – 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. 88% of all experts predict a higher response rate towards White profiles relative to Black profiles. On average, experts expect White profiles to receive 12.8% more message responses relative to Black profiles. This prediction is substantially different from what we actually observe, as the actual relative gap is 3% ($p \leq 0.001$). Once again this finding is very robust to a variety of heterogeneity analyses. Professors, experts who work on discrimination, males, and females, 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 males to discriminate substantially more than females. As shown in the main part of the paper, we find the exact opposite: it is females who discriminate substantially more than males. 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.

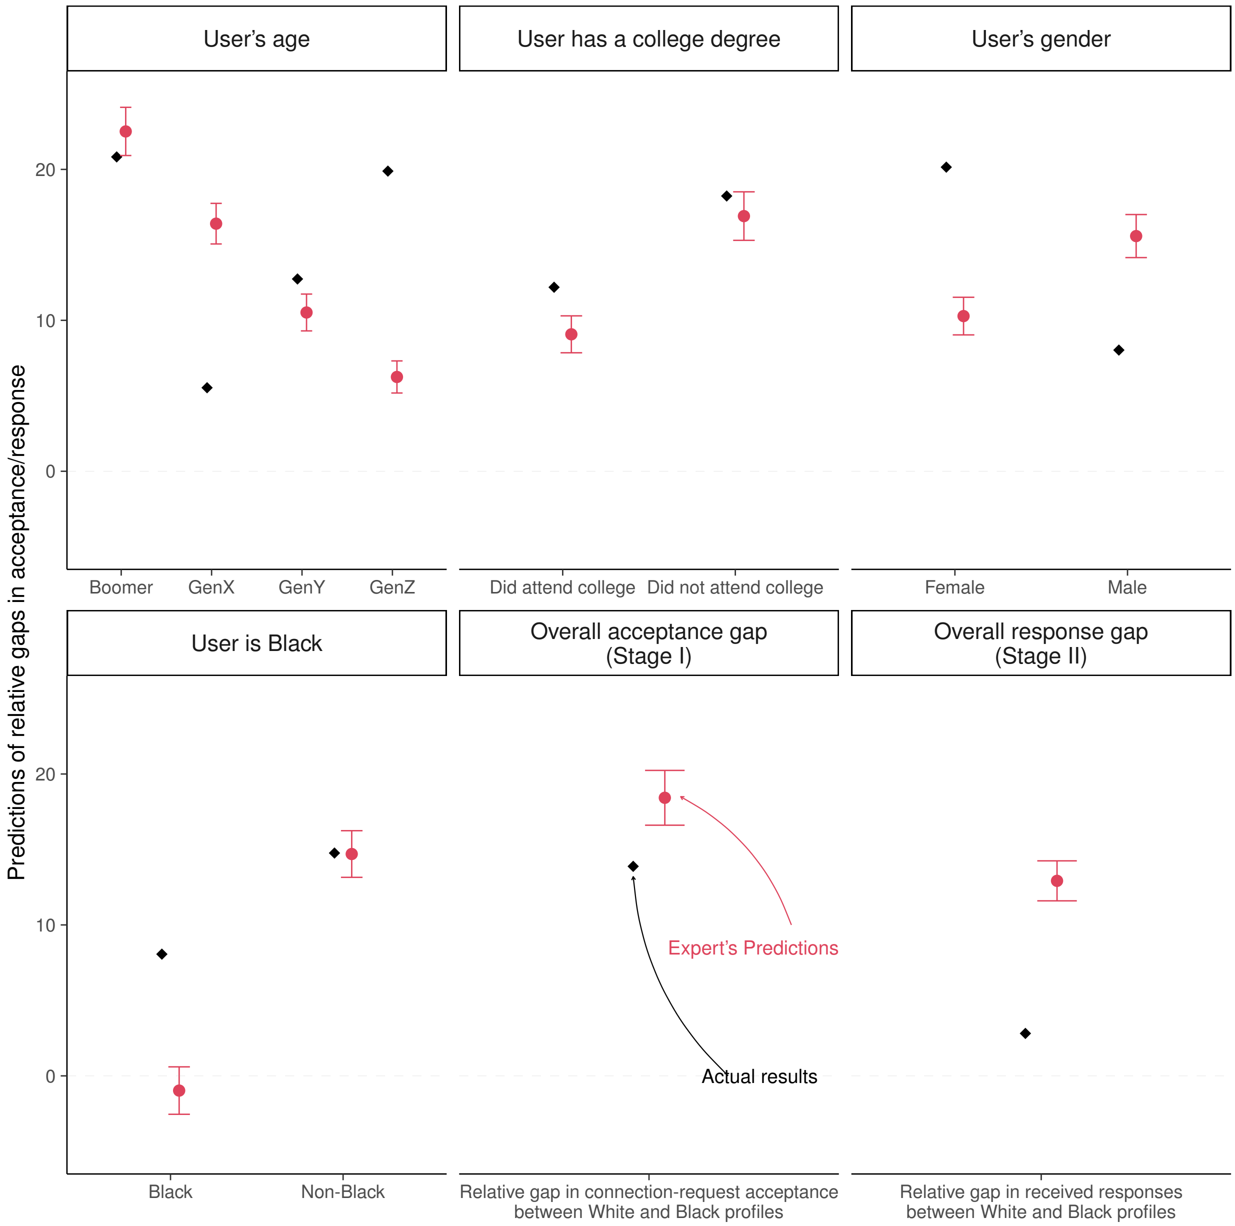

Figure H.1: Experts' predictions of discrimination on LinkedIn

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.

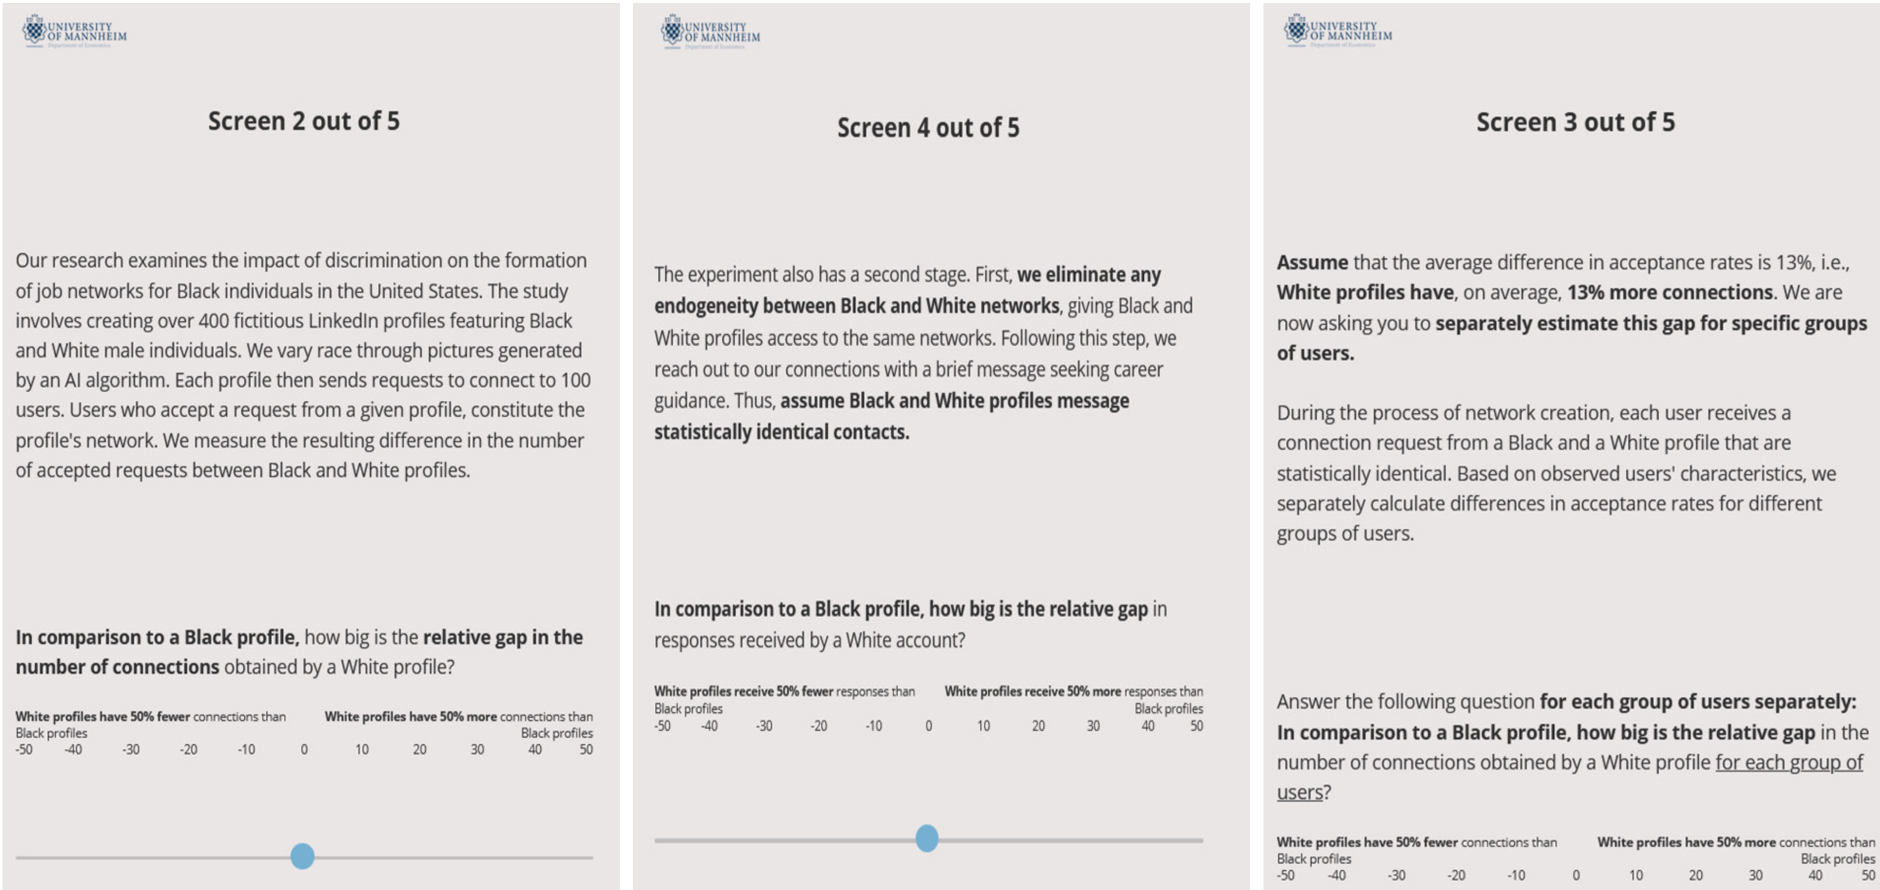

(a) First stage.

(b) Second stage.

(c) Heterogenity.

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

I Figures

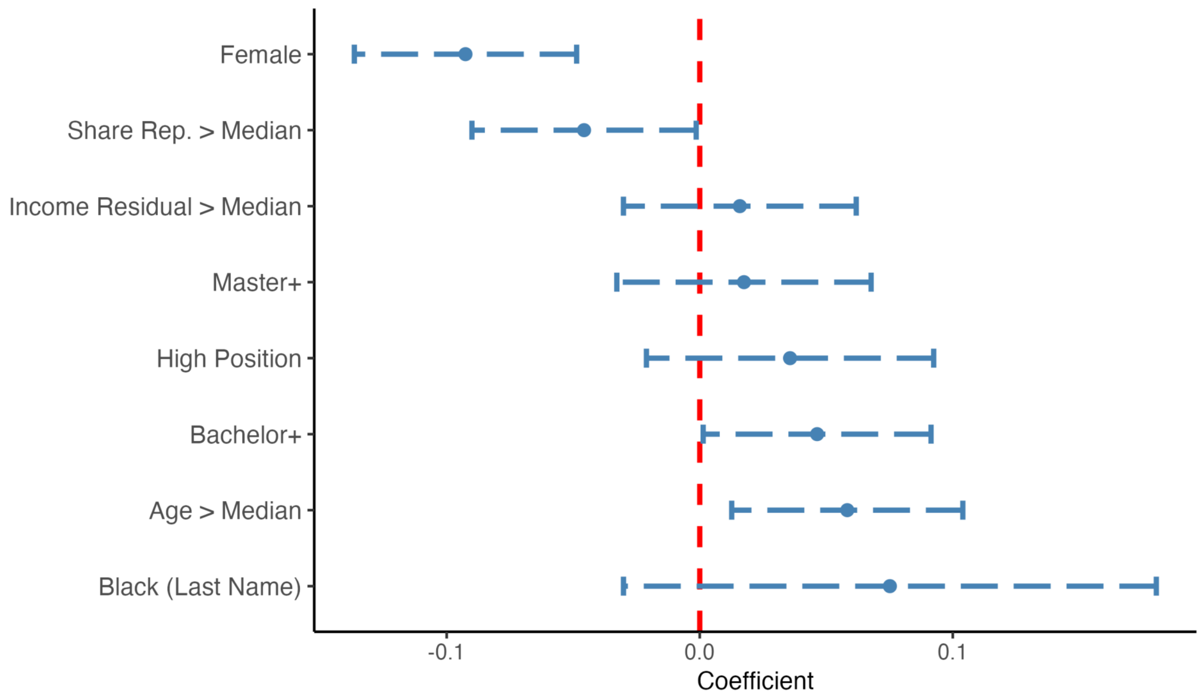

Figure I.1: Correlates of discrimination in relative terms.

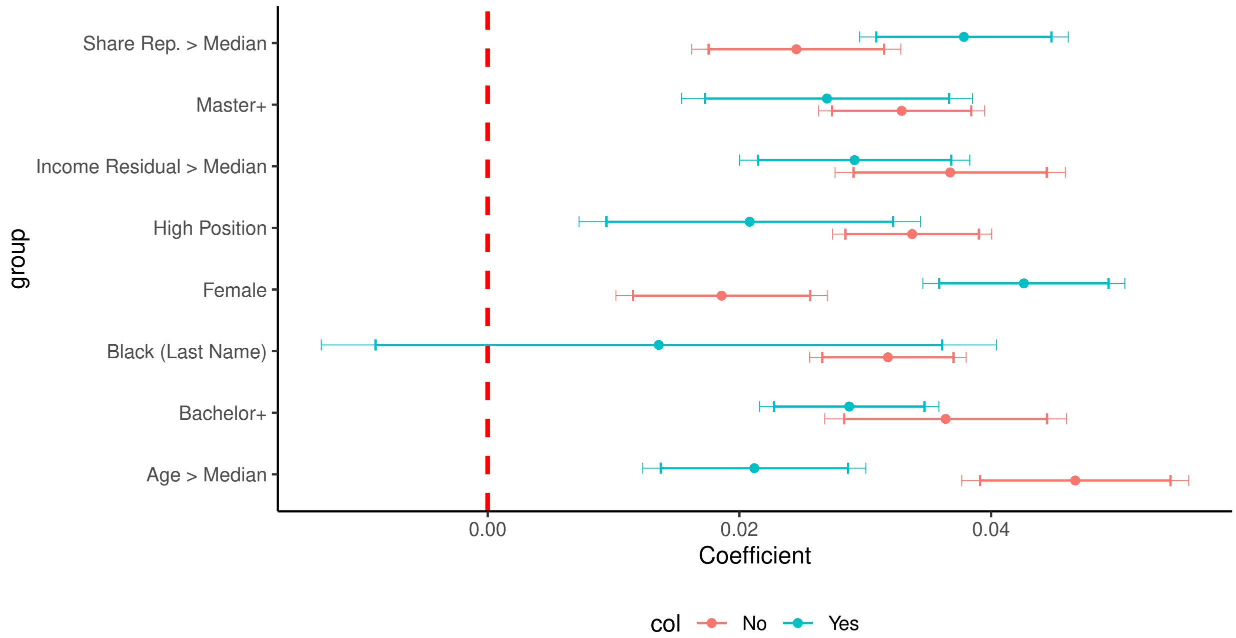

Figure I.2: Heterogeneity: Total gaps in acceptance rates

Note: This Figure plots the total gap in acceptance rates between White and Black individuals across different groups of users. The 90 and 95% confidence intervals are shown. A positive value indicates a higher acceptance rate of White in comparison to Black requests. Estimation is based on Equation 1. However, instead of a separate intercept for individuals with the respective characteristic, the regression includes two interaction terms between the profile being White and the individual holding or not holding the characteristic (both as separate dummy variables): $accepts_{i,j} = \beta_0 + \beta_1 White_j \times has\_characteristic_i + \beta_2 White_j \times has\_not\_characteristic_i + \gamma_i + \omega_j + u_{i,j}$

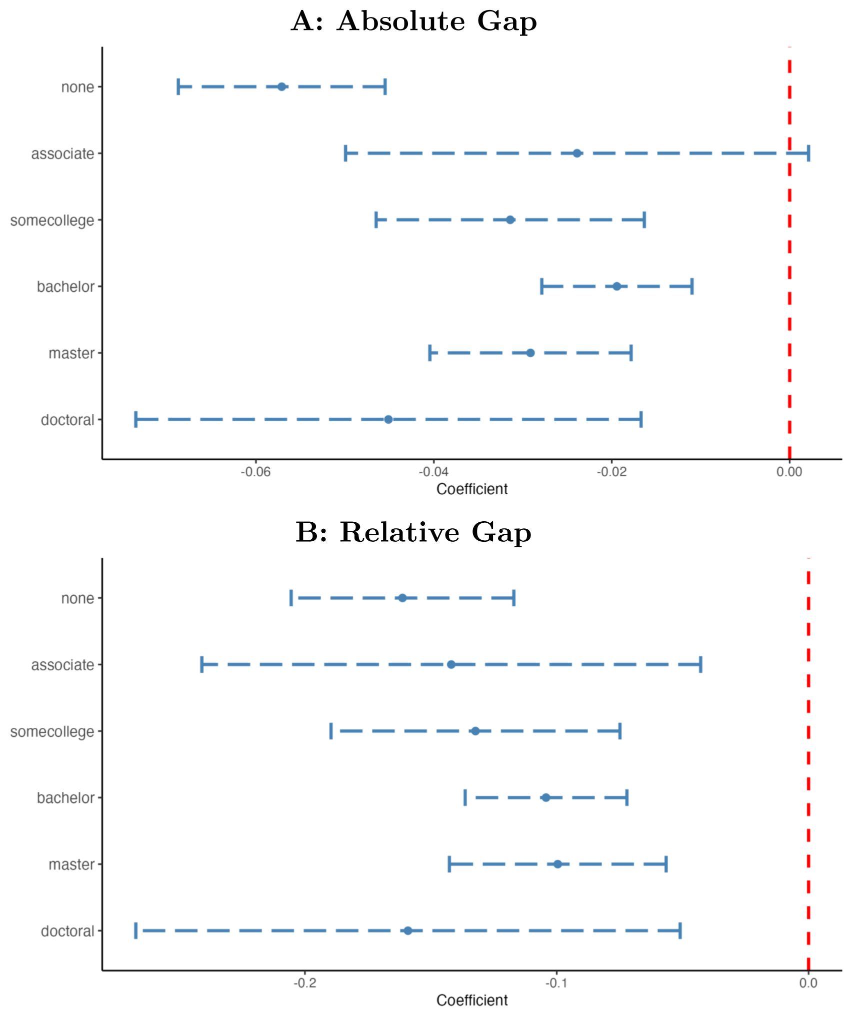

Figure I.3: Absolute and relative gap in acceptance rates across education groups

Note: The figures above are based on the following regression: $accepts_{i,j} = \beta + \sum_k^K \gamma_k White_j \times MaxDegree_i + \gamma_i + \omega_j + u_{i,j}$, where $MaxDegree_{j,k}$ is user j's maximum obtained degree k. γ_i and ω_j are user and profile picture dummies. The regression thus computes a separate gap in acceptance rates for each educational group. In Figure A, the dependent variable is a dummy for whether a given user accepted the request. In Figure B, the dummy is first divided by the acceptance rate of Black requests for education groups. The result then shows the relative gap, i.e. by what percentage the probability of acceptance increases if the request instead stems from a White user. Overall the Figure suggests a slightly decreasing gap with education. However, we find a high gap for users with a Ph.D.

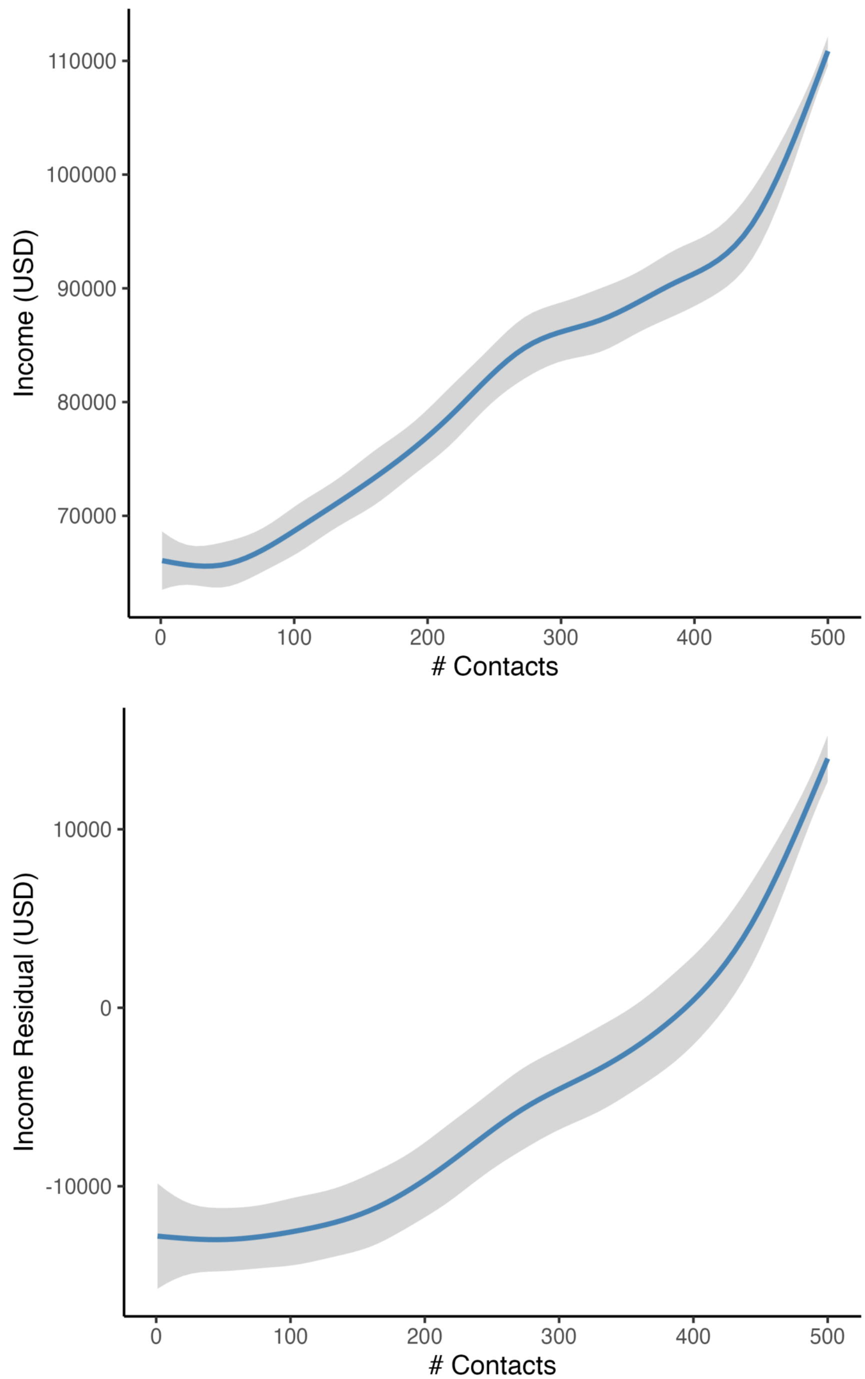

Figure I.4: Income and Income Residual in Relationship to Number of Contacts

Note: The first Figure above shows the smoothed conditional mean income of a given target with respect to her number of connections. Note that the number of connections is capped at 500, as LinkedIn does not report values above 500. The second figure shows the conditional mean of the targets' income residual with respect to the number of connections. Similar to Section 4.2, the income residual is calculated by first regressing income on age, age squared, level of education, race, and gender. Next, the difference between a target's actual income and her predicted income is taken. The residual thus represents a target's income differential relative to what she would have been expected to earn based on these covariates. As the Figure shows, both the income and income residual appear to be close to linear functions of the number of contacts on LinkedIn. Results from a linear regression including the mentioned controls and the number of connections suggest that an additional connection increases the yearly wage by 70.6 USD. While these results are not causal, we use them for a back-of-the-envelope calculation to estimate the economic effects of differences in networks.

J Tables

J.1 Regressions

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

The standard regression we are using is:

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

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., $\gamma = \delta = 0$. $X_k^{\cdot}$ 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 female, Asian.[54]

$Outcome_{i,j}$ 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 characters), and the probability of the response being highly valuable. 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.

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

J.2 Main experiment –First Stage

J.2.1 Aggregate results

Panel A: Aggregate difference in number of contacts

	Number of Contacts						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Constant	26.13***	39.28***	25.96***	26.02***	24.91***	22.70***	35.66***
	(0.44)	(1.91)	(1.06)	(0.87)	(1.15)	(5.06)	(5.14)
Profile is Black	−3.06***	−3.05***	−3.06***	−3.13***	−3.06***	−2.97***	−3.07***
	(0.47)	(0.47)	(0.47)	(0.46)	(0.49)	(0.55)	(0.54)
State Controls	×	✓	×	×	×	×	✓
Job Controls	×	×	✓	×	×	×	✓
Firstname Controls	×	×	×	✓	×	×	✓
Lastname Controls	×	×	×	×	✓	×	✓
Picture trait Controls	×	×	×	×	×	✓	✓
Picture specific random effects	✓	✓	✓	✓	✓	✓	✓
Log Likelihood	-1279.12	-1108.11	-1275.08	-1260.82	-1254.96	-1287.6	-1075.11
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)
Constant	26.06***	39.20***	25.84***	26.04***	24.90***	22.28***	35.25***
	(0.63)	(1.95)	(1.17)	(0.97)	(1.23)	(5.14)	(5.22)
Profile is Black	−3.26***	−3.27***	−3.26***	−3.54***	−3.24***	−3.12***	−3.49***
	(0.68)	(0.67)	(0.68)	(0.67)	(0.70)	(0.74)	(0.73)
Profile attended worse Uni	0.16	0.16	0.17	−0.08	0.02	0.27	−0.09
	(0.88)	(0.72)	(0.89)	(0.89)	(0.90)	(0.89)	(0.76)
Profile is Black and attended worse Uni	0.40	0.42	0.40	0.79	0.36	0.28	0.80
	(0.95)	(0.95)	(0.95)	(0.94)	(0.98)	(0.96)	(0.96)
State Controls	×	✓	×	×	×	×	✓
Job Controls	×	×	✓	×	×	×	✓
Firstname Controls	×	×	×	✓	×	×	✓
Lastname Controls	×	×	×	×	✓	×	✓
Picture trait Controls	×	×	×	×	×	✓	✓
Picture specific random effects	✓	✓	✓	✓	✓	✓	✓
Log Likelihood	-1277.43	-1106.6	-1273.36	-1258.89	-1253.32	-1285.89	-1073.36
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 one 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.

J.2.2 Dynamic effects

Panel A: Aggregate difference in number of contacts over time

	Number of Contacts over time						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Constant	2.06^{***}	9.14^{***}	1.89^{***}	1.94^{***}	$0.66^{\cdot}$	0.20	4.62^{*}
	(0.24)	(1.12)	(0.56)	(0.30)	(0.36)	(1.48)	(1.96)
Profile is Black	-0.51^{**}	-0.51^{**}	-0.51^{**}	-0.57^{**}	-0.47^{*}	-0.49^{*}	-0.56^{**}
	(0.19)	(0.19)	(0.19)	(0.19)	(0.19)	(0.21)	(0.21)
Week	3.24^{***}	3.24^{***}	3.24^{***}	3.24^{***}	3.24^{***}	3.24^{***}	3.24^{***}
	(0.03)	(0.03)	(0.03)	(0.03)	(0.03)	(0.03)	(0.03)
Profile is Black x Week	-0.37^{***}	-0.37^{***}	-0.37^{***}	-0.37^{***}	-0.37^{***}	-0.37^{***}	-0.37^{***}
	(0.05)	(0.05)	(0.05)	(0.05)	(0.05)	(0.05)	(0.04)
State Controls	×	✓	×	×	×	×	✓
Job Controls	×	×	✓	×	×	×	✓
Firstname Controls	×	×	×	✓	×	×	✓
Lastname Controls	×	×	×	×	✓	×	✓
Picture trait Controls	×	×	×	×	×	✓	✓
Picture specific random effects	✓	✓	✓	✓	✓	✓	✓
Log Likelihood	-8340.66	-8214	-8338.81	-8290.36	-8286.96	-8350.3	-8124.43
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)
Constant	2.24^{***}	9.32^{***}	2.04^{**}	2.18^{***}	0.88^{*}	0.35	4.84^{*}
	(0.33)	(1.14)	(0.62)	(0.38)	(0.42)	(1.51)	(1.98)
Profile is Black	-0.84^{**}	-0.84^{**}	-0.84^{**}	-1.06^{***}	-0.76^{**}	-0.78^{**}	-0.94^{***}
	(0.28)	(0.28)	(0.28)	(0.27)	(0.28)	(0.29)	(0.28)
Profile attended worse Uni	-0.36	-0.36	-0.34	-0.51	-0.44	-0.37	-0.57
	(0.47)	(0.40)	(0.48)	(0.47)	(0.47)	(0.47)	(0.41)
Week	3.18^{***}	3.18^{***}	3.18^{***}	3.18^{***}	3.18^{***}	3.18^{***}	3.18^{***}
	(0.05)	(0.05)	(0.05)	(0.05)	(0.05)	(0.05)	(0.05)
Profile is Black and attended worse Uni	$0.66^{\cdot}$	$0.67^{\cdot}$	$0.66^{\cdot}$	0.97^{*}	0.58	0.57	$0.74^{\cdot}$
	(0.39)	(0.39)	(0.39)	(0.38)	(0.39)	(0.39)	(0.38)
Profile is Black x Week	-0.31^{***}	-0.31^{***}	-0.31^{***}	-0.31^{***}	-0.31^{***}	-0.31^{***}	-0.31^{***}
	(0.07)	(0.07)	(0.07)	(0.06)	(0.07)	(0.07)	(0.06)
Profile attended worse Uni x Week	$0.13^{\cdot}$	$0.13^{\cdot}$	$0.13^{\cdot}$	0.13^{*}	0.13^{*}	0.13^{*}	0.13^{*}
	(0.07)	(0.07)	(0.07)	(0.06)	(0.06)	(0.07)	(0.06)
Profile is Black x attended worse Uni x Week	-0.12	-0.12	-0.12	-0.12	-0.12	-0.12	-0.12
	(0.09)	(0.09)	(0.09)	(0.09)	(0.09)	(0.09)	(0.09)
State Controls	×	✓	×	×	×	×	✓
Job Controls	×	×	✓	×	×	×	✓
Firstname Controls	×	×	×	✓	×	×	✓
Lastname Controls	×	×	×	×	✓	×	✓
Picture trait Controls	×	×	×	×	×	✓	✓
Picture specific random effects	✓	✓	✓	✓	✓	✓	✓
Log Likelihood	-8342.17	-8215.63	-8340.29	-8289.2	-8288.93	-8352.28	-8125.7
Observations	3,200	3,200	3,200	3,200	3,200	3,200	3,200

Notes: $^{\cdot}p<0.10$; $^{*}p<0.05$; $^{**}p<0.01$; $^{***}p<0.001$.

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.

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)
Constant	2.72^{***}	4.21^{***}	2.71^{***}	2.73^{***}	2.57^{***}	2.43^{***}	3.87^{***}
	(0.08)	(0.24)	(0.14)	(0.12)	(0.15)	(0.62)	(0.64)
Profile is Black	-0.47^{***}	-0.47^{***}	-0.47^{***}	-0.48^{***}	-0.47^{***}	-0.48^{***}	-0.48^{***}
	(0.11)	(0.11)	(0.11)	(0.11)	(0.11)	(0.11)	(0.11)
Week	0.11^{***}	0.11^{***}	0.11^{***}	0.11^{***}	0.11^{***}	0.11^{***}	0.11^{***}
	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(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)
State Controls	×	✓	×	×	×	×	✓
Job Controls	×	×	✓	×	×	×	✓
Firstname Controls	×	×	×	✓	×	×	✓
Lastname Controls	×	×	×	×	✓	×	✓
Picture trait Controls	×	×	×	×	×	✓	✓
Picture specific random effects	✓	✓	✓	✓	✓	✓	✓
Log Likelihood	-6247.55	-6188.09	-6252	-6247.26	-6255.57	-6272.26	-6229.72
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)
Constant	2.72^{***}	4.21^{***}	2.70^{***}	2.74^{***}	2.58^{***}	2.38^{***}	3.82^{***}
	(0.12)	(0.26)	(0.17)	(0.15)	(0.17)	(0.63)	(0.65)
Profile is Black	-0.50^{**}	-0.50^{***}	-0.50^{**}	-0.53^{***}	-0.50^{**}	-0.50^{**}	-0.54^{***}
	(0.15)	(0.15)	(0.15)	(0.15)	(0.15)	(0.16)	(0.16)
Profile attended worse Uni	0.004	0.002	0.01	-0.03	-0.02	0.02	-0.03
	(0.17)	(0.15)	(0.17)	(0.17)	(0.17)	(0.17)	(0.16)
Week	0.11^{***}	0.11^{***}	0.11^{***}	0.11^{***}	0.11^{***}	0.11^{***}	0.11^{***}
	(0.03)	(0.03)	(0.03)	(0.03)	(0.03)	(0.03)	(0.03)
Profile is Black and attended worse Uni	0.06	0.06	0.06	0.11	0.06	0.04	0.11
	(0.21)	(0.21)	(0.21)	(0.21)	(0.22)	(0.21)	(0.22)
Profile is Black x Week	0.03	0.03	0.03	0.03	0.03	0.03	0.03
	(0.04)	(0.04)	(0.04)	(0.04)	(0.04)	(0.04)	(0.04)
Profile attended worse Uni x Week	0.01	0.01	0.01	0.01	0.01	0.01	0.01
	(0.04)	(0.04)	(0.04)	(0.04)	(0.04)	(0.04)	(0.04)
Profile is Black x attended worse Uni x Week	-0.01	-0.01	-0.01	-0.01	-0.01	-0.01	-0.01
	(0.05)	(0.05)	(0.05)	(0.05)	(0.05)	(0.05)	(0.05)
State Controls	×	✓	×	×	×	×	✓
Job Controls	×	×	✓	×	×	×	✓
Firstname Controls	×	×	×	✓	×	×	✓
Lastname Controls	×	×	×	×	✓	×	✓
Picture trait Controls	×	×	×	×	×	✓	✓
Picture specific random effects	✓	✓	✓	✓	✓	✓	✓
Log Likelihood	-6254.8	-6195.48	-6259.22	-6254.31	-6262.89	-6279.48	-6236.89
Observations	3,200	3,200	3,200	3,200	3,200	3,200	3,200

Notes: $^{\cdot}p<0.10$; $^{*}p<0.05$; $^{**}p<0.01$; $^{***}p<0.001$.

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.

J.2.3 Geographical variation

	Difference in the number of contacts										
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
Constant	2.89^{***}	2.93^{***}	2.82^{***}	4.23^{*}	2.33^{***}	4.79^{**}	4.62^{***}	2.44^{***}	3.19^{***}	2.99^{***}	2.46^{***}
	(0.47)	(0.47)	(0.53)	(1.70)	(0.51)	(1.47)	(1.00)	(0.44)	(0.38)	(0.38)	(0.37)
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					$6.26^{\cdot}$						
					(3.45)						
Share Democratic						−3.64					
						(2.95)					
GDP per Capita (current USD)							$-0.0000^{\cdot}$				
							(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)
Observations	51	51	51	51	51	51	51	51	51	51	51
Notes:							$^{\cdot}p<0.10$; $^{*}p<0.05$; $^{**}p<0.01$; $^{***}p<0.001$.				

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. 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 males 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.[55]

J.2.4 Differences in Networks

	Black profiles (N=4477)	**White profiles (N=5046)**	**p-value**
Male	0.50 (0.50)	0.48 (0.50)	**0.019***
White	0.71 (0.45)	0.71 (0.45)	0.92
Black	0.06 (0.24)	0.06 (0.23)	0.344
Age	32.91 (9.97)	32.67 (9.97)	0.254
Contact Count	319.88 (185.70)	311.97 (185.89)	**0.039***
Skill Count	20.92 (13.77)	20.58 (13.54)	0.24
Num Verific. of Skills	37.36 (57.41)	37.57 (146.00)	0.93
Post/Share on Platform	0.55 (0.50)	0.53 (0.50)	**0.078.**
Log(Nbr.Followers)	5.77 (1.47)	5.70 (1.46)	**0.035***
Profile has volunteering	0.21 (0.40)	0.20 (0.40)	0.215
Profile has language	0.01 (0.11)	0.01 (0.11)	0.791
Profile picture is happy	0.81 (0.39)	0.82 (0.38)	0.377
Folows a philanthropist	0.03 (0.18)	0.03 (0.18)	0.775
Works in HR	0.10 (0.30)	0.10 (0.29)	0.749
Firm: Employees	4739.38 (4529.35)	4699.56 (4521.37)	0.686
Firm: Employees on Platform	25714.65 (69553.80)	25781.80 (70590.17)	0.965
Firm: Jobs on Platform	1903.24 (6516.28)	1939.82 (7209.71)	0.806
Degree: None	0.15 (0.36)	0.16 (0.37)	0.274
Degree: Associate	0.04 (0.20)	0.05 (0.21)	0.666
Degree: BA	0.45 (0.50)	0.44 (0.50)	0.375
Degree: MA	0.21 (0.41)	0.21 (0.40)	0.645
Degree: PhD	0.03 (0.17)	0.03 (0.17)	0.713
County: Share Dem.	0.60 (0.15)	0.60 (0.15)	0.312
County: Share Rep.	0.38 (0.15)	0.38 (0.15)	0.302
County: Share White	0.58 (0.19)	0.58 (0.19)	0.926
County: Share Black	0.17 (0.15)	0.17 (0.15)	0.993
County: Pop. Density	1814.57 (5096.56)	1766.03 (5023.99)	0.647
County: B/W Dissimilarity Index	54.73 (11.62)	54.61 (11.78)	0.636
W/nW Dissimilarity Index	41.33 (11.78)	41.17 (11.94)	0.514

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

The table reports on 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: $^{\cdot}p<0.10$;$^{*}p<0.05$;$^{**}p<0.01$;$^{***}p<0.001$.

[55] States 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.

J.2.5 Individual predictors of discrimination

	Accepted contact request		
	(1)	(2)	(3)
Constant	0.28***	0.24***	0.18***
	(0.004)	(0.01)	(0.03)
Profile is Black	−0.03***	−0.03***	−0.03***
	(0.003)	(0.003)	(0.004)
Z-Scored Age	−0.04***	−0.05***	−0.05***
	(0.003)	(0.003)	(0.004)
Profile is Black x Z-Scored Age	0.01***	0.01***	0.01**
	(0.003)	(0.003)	(0.004)
Picture random effects	✓	✓	✓
Target random effects	✓	✓	✓
Nbr of contacts	×	✓	✓
Other fixed effects	×	×	✓
Log Likelihood	-16876.62	-16693.16	-10529.48
Observations	33,446	32,928	21,089
Notes:	$^{\cdot}p<0.10$; $^{*}p<0.05$; $^{**}p<0.01$; $^{***}p<0.001$.		

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); dummy variables on whether the target has an assistant, senior, or hr-job; a dummy on whether the target is White; a dummy on whether the target is female; the share of democrats in the target's county; whether the target attended the same university and the share of Black and White students in the target's university.

	Accepted contact request								
	All users			White users			Black users		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Constant	0.26***	0.20***	0.20***	0.27***	0.19***	0.16***	0.32***	0.26***	0.13
	(0.01)	(0.01)	(0.03)	(0.01)	(0.01)	(0.03)	(0.02)	(0.03)	(0.13)
Profile is Black	−0.02***	−0.02***	−0.01*	−0.02***	−0.02***	−0.02**	−0.02	−0.02	−0.01
	(0.004)	(0.004)	(0.01)	(0.01)	(0.01)	(0.01)	(0.02)	(0.02)	(0.03)
Female	−0.002	0.01$^{\cdot}$	−0.01	−0.01$^{\cdot}$	0.004	−0.02*	−0.05	−0.03	−0.03
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.03)	(0.03)	(0.04)
Profile is Black x Female	−0.02***	−0.03***	−0.03***	−0.02**	−0.02**	−0.03**	−0.001	−0.001	−0.01
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.03)	(0.03)	(0.04)
Picture random effects	✓	✓	✓	✓	✓	✓	✓	✓	✓
Target random effects	✓	✓	✓	✓	✓	✓	✓	✓	✓
Nbr of contacts	×	✓	✓	×	✓	✓	×	✓	✓
Other fixed effects	×	×	✓	×	×	✓	×	×	✓
All subjects	✓	✓	✓	×	×	×	×	×	×
Only White subjects	×	×	×	✓	✓	✓	×	×	×
Only Non-white subjects	×	×	×	×	×	×	✓	✓	✓
Log Likelihood	-17674.83	-17384.49	-12046.11	-11300.65	-11121.62	-7844	-862.22	-848.04	-587.74
Observations	36,911	35,794	24,035	23,792	23,130	15,640	1,617	1,550	1,027

Notes: $^{\cdot}p<0.10$; $^{*}p<0.05$; $^{**}p<0.01$; $^{***}p<0.001$.

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); dummy variables on whether the target has an assistant, senior, or hr-job; a dummy on whether the target is White; the z-scored age of the target; the share of democrats in the target's county; whether the target attended the same university and the share of Black and White students in the target's university. 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.

	Accepted contact request											
	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)
Constant	0.27^{***}	0.18^{***}	0.25^{***}	0.18^{***}	0.25^{***}	0.17^{***}	0.26^{***}	0.17^{***}	0.26^{***}	0.21^{***}	0.26^{***}	0.19^{***}
	(0.01)	(0.04)	(0.01)	(0.04)	(0.01)	(0.03)	(0.01)	(0.04)	(0.01)	(0.04)	(0.01)	(0.03)
Profile is Black	-0.02^{***}	-0.02^{*}	-0.04^{***}	-0.05^{***}	-0.04^{***}	-0.05^{***}	-0.02^{***}	-0.02^{**}	-0.05^{***}	-0.05^{***}	-0.05^{***}	-0.05^{***}
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
Prop(Target $\neq$ White)	-0.02	-0.02	0.02	$0.03^{\cdot}$	0.02	$0.03^{\cdot}$						
	(0.01)	(0.02)	(0.01)	(0.02)	(0.01)	(0.02)						
Profile is Black x Prop(Target $\neq$ White)	0.01	0.01	0.01	0.01	0.01	0.01						
	(0.01)	(0.02)	(0.01)	(0.02)	(0.01)	(0.02)						
Prop(Target = Black)							0.05	0.02	0.01	-0.03	0.01	-0.03
							(0.03)	(0.04)	(0.03)	(0.04)	(0.03)	(0.04)
Profile is Black x Prop(Target = Black)							0.02	0.03	$0.04^{\cdot}$	0.07^{*}	0.04	$0.06^{\cdot}$
							(0.03)	(0.04)	(0.02)	(0.03)	(0.03)	(0.03)
Male					$0.02^{\cdot}$	0.03^{*}					0.001	0.01
					(0.01)	(0.01)					(0.01)	(0.01)
Profile is Black x Male					0.02^{*}	0.03^{*}					0.02^{**}	0.03^{***}
					(0.01)	(0.01)					(0.01)	(0.01)
Male x Prop(Target $\neq$ White)					$-0.04^{\cdot}$	$-0.05^{\cdot}$						
					(0.02)	(0.03)						
Profile is Black x Male x Prop(Target $\neq$ White)					0.001	0.01						
					(0.02)	(0.02)						
Male x Prop(Target = Black)											0.04	0.05
											(0.04)	(0.05)
Profile is Black x Male x Prop(Target = Black)											-0.02	-0.04
											(0.04)	(0.05)
Picture random effects	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Target random effects	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Other fixed effects	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓
All subjects	×	×	×	×	✓	✓	×	×	×	×	✓	✓
Only Females	×	×	✓	✓	×	×	×	×	✓	✓	×	×
Only Males	✓	✓	×	×	×	×	✓	✓	×	×	×	×
Log Likelihood	-7840.77	-5501.18	-7725.55	-5050.19	-15580.5	-10533.5	-7836.92	-5499.64	-7724.02	-5049.27	-15575.48	-10530.85
Observations	15,771	10,603	17,068	10,486	32,839	21,089	15,771	10,603	17,068	10,486	32,839	21,089

Notes: $^{\cdot}$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); dummy variables on whether the target has an assistant, senior, or hr-job; the z-scored age of the target; the share of democrats in the target's county; whether the target attended the same university and the share of Black and White students in the target's university. 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.

	Accepted contact request by user's race			
	(1)	(2)	(3)	(4)
Constant	0.26^{***}	0.18^{***}	0.26^{***}	0.19^{***}
	(0.01)	(0.03)	(0.005)	(0.03)
Profile is Black	-0.03^{***}	-0.03^{***}	-0.04^{***}	-0.04^{***}
	(0.004)	(0.01)	(0.004)	(0.005)
Prop(Target $\neq$ White)	0.01	0.01		
	(0.01)	(0.01)		
Profile is Black x Prop(Target $\neq$ White)	0.004	0.003		
	(0.01)	(0.01)		
Prop(Target = Black)			0.03	-0.001
			(0.02)	(0.03)
Profile is Black x Prop(Target = Black)			$0.04^{\cdot}$	0.05^{*}
			(0.02)	(0.02)
Picture random effects	✓	✓	✓	✓
Target random effects	✓	✓	✓	✓
Other fixed effects	×	✓	×	✓
Log Likelihood	-16130.47	-10893.55	-16124.57	-10889.32
Observations	33,861	21,739	33,861	21,739
Notes:	$^{\cdot}$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); dummy variables on whether the target has an assistant, senior, or hr-job; the z-scored age of the target; a dummy on whether the target is female; the share of democrats in the target's county; whether the target attended the same university and the share of Black and White students in the target's university. 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.

J.2.6 Ancillary outcomes

	#Views		# unsolicited messages		# unsolicited contact requests	
	(1)	(2)	(3)	(4)	(5)	(6)
Constant	41.38^{***}	40.68^{***}	0.31^{***}	0.29^{***}	0.96^{***}	0.94^{***}
	(0.48)	(0.45)	(0.04)	(0.04)	(0.09)	(0.09)
Profile is Black	-5.72^{***}	-4.30^{***}	-0.09^{*}	-0.06	0.01	0.04
	(0.51)	(0.48)	(0.04)	(0.04)	(0.10)	(0.10)
Nbr. of Contacts		3.00^{***}		0.08^{**}		0.06
		(0.32)		(0.03)		(0.07)
Log Likelihood	-1313.95	-1273.98	-287.39	-285.74	-652.78	-654.13
Observations	400	400	400	400	400	400
Notes:	$^{\cdot}$p<0.10;*p<0.05;**p<0.01;***p<0.001.					

Table J.10: 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.

J.3 Main experiment – Second Stage

J.3.1 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	58.28	39.84	3	300
Nbr. characters	338	279	329.93	227.47	13	1,805
Friendliness	338	3.50	3.51	.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

Table J.11: Summary statistics of the responses received on profile level.

The tables 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.

J.3.2 Message Responses

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

	Response Rate					Message Length (in char)					Highly Useful Message?				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)
Constant	0.21***	0.37***	0.16***	0.20·	0.28·	83.21***	91.65***	76.01***	26.88	14.53	0.08***	0.13	0.01	0.06	0.06
	(0.01)	(0.06)	(0.03)	(0.12)	(0.14)	(4.96)	(25.25)	(18.30)	(58.10)	(73.72)	(0.02)	(0.08)	(0.06)	(0.19)	(0.23)
Profile is Black	−0.01	−0.01	−0.01	−0.01	−0.003	−10.36	−10.05	−8.83	−10.64	−8.84	−0.01	−0.02	−0.02	−0.01	−0.02
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(6.98)	(7.23)	(7.17)	(7.03)	(7.63)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)
State Controls	×	✓	×	×	✓	×	✓	×	×	✓	×	✓	×	×	✓
Job Controls	×	✓	×	×	✓	×	✓	×	×	✓	×	✓	×	×	✓
Firstname Controls	×	×	✓	×	✓	×	×	✓	×	✓	×	×	✓	×	✓
Lastname Controls	×	×	✓	×	✓	×	×	✓	×	✓	×	×	✓	×	✓
Picture trait Controls	×	×	×	✓	✓	×	×	×	✓	✓	×	×	×	✓	✓
Picture specific random effects	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Log Likelihood	206.07	122.4	161.61	166.2	40.68	-1886.09	-1643.32	-1788.97	-1875.26	-1534.91	51.98	-6.71	11.18	14.75	-82.27
Observations	400	400	400	400	400	339	339	339	339	339	338	338	338	338	338

Panel B: Differences in messages accounting for profile quality

	Response Rate					Message Length (in char)					Highly Useful Message?				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)
Constant	0.20***	0.36***	0.15***	0.19	0.27·	88.40***	98.81***	79.94***	39.36	30.48	0.09***	0.15·	0.02	0.09	0.11
	(0.01)	(0.06)	(0.04)	(0.12)	(0.14)	(7.23)	(25.85)	(19.21)	(59.59)	(75.73)	(0.02)	(0.08)	(0.06)	(0.19)	(0.23)
Profile is Black	0.03	0.03	0.03	0.03	0.03	−11.75	−12.27	−10.38	−13.00	−13.17	−0.02	−0.03	−0.03	−0.02	−0.04
	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(9.93)	(10.33)	(10.23)	(10.09)	(10.95)	(0.03)	(0.03)	(0.03)	(0.03)	(0.03)
Profile attended worse Uni	0.02	0.02	0.02	0.02	0.03	−9.80	−12.76	−6.73	−9.14	−10.65	−0.02	−0.04	−0.01	−0.03	−0.04
	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(9.93)	(10.30)	(10.37)	(10.34)	(11.41)	(0.03)	(0.03)	(0.03)	(0.03)	(0.04)
Profile is Black and attended worse Uni	−0.07*	−0.07*	−0.07*	−0.07*	−0.07*	1.87	3.21	2.58	3.88	7.94	0.03	0.03	0.02	0.03	0.04
	(0.03)	(0.03)	(0.03)	(0.03)	(0.03)	(13.98)	(14.50)	(14.66)	(14.23)	(15.69)	(0.04)	(0.04)	(0.05)	(0.04)	(0.05)
State Controls	×	✓	×	×	✓	×	✓	×	×	✓	×	✓	×	×	✓
Job Controls	×	✓	×	×	✓	×	✓	×	×	✓	×	✓	×	×	✓
Firstname Controls	×	×	✓	×	✓	×	×	✓	×	✓	×	×	✓	×	✓
Lastname Controls	×	×	✓	×	✓	×	×	✓	×	✓	×	×	✓	×	✓
Picture trait Controls	×	×	×	✓	✓	×	×	×	✓	✓	×	×	×	✓	✓
Picture specific random effects	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Log Likelihood	203.63	119.95	158.8	163.48	37.82	-1878.86	-1635.63	-1782.17	-1868.26	-1527.78	47.14	-11.06	6.3	10.08	-86.62
Observations	400	400	400	400	400	339	339	339	339	339	338	338	338	338	338

Panel C: Differences in messages accounting for message type

	Response Rate					Message Length (in char)					Highly Useful Message?				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)
Constant	0.29***	0.45***	0.23***	0.24·	0.32*	359.82***	272.38***	358.67***	116.36	7.23	0.08***	0.13·	0.02	0.01	0.03
	(0.01)	(0.06)	(0.04)	(0.12)	(0.14)	(19.39)	(81.35)	(62.20)	(198.13)	(244.22)	(0.02)	(0.07)	(0.06)	(0.20)	(0.23)
Profile is Black	−0.01	−0.01	−0.01	−0.01	−0.01	−51.99·	−42.17	−47.90·	−54.31·	−43.79	−0.01	−0.004	−0.02	−0.001	−0.01
	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(27.75)	(28.23)	(28.38)	(27.97)	(29.77)	(0.03)	(0.03)	(0.03)	(0.03)	(0.03)
Mentor message	−0.15***	−0.15***	−0.15***	−0.15***	−0.15***	−30.72	−27.44	−21.31	−32.31	−18.06	−0.01	0.001	−0.01	−0.01	0.002
	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(30.48)	(31.42)	(30.84)	(30.44)	(31.62)	(0.03)	(0.03)	(0.03)	(0.03)	(0.03)
Profile is Black and Mentor message	0.01	0.01	0.01	0.01	0.01	81.66·	82.12·	71.07	85.01*	75.11·	0.01	−0.02	0.01	0.01	−0.02
	(0.03)	(0.03)	(0.03)	(0.03)	(0.03)	(43.07)	(44.77)	(43.59)	(42.97)	(45.05)	(0.05)	(0.05)	(0.05)	(0.05)	(0.05)
State Controls	×	✓	×	×	✓	×	✓	×	×	✓	×	✓	×	×	✓
Job Controls	×	✓	×	×	✓	×	✓	×	×	✓	×	✓	×	×	✓
Firstname Controls	×	×	✓	×	✓	×	×	✓	×	✓	×	×	✓	×	✓
Lastname Controls	×	×	✓	×	✓	×	×	✓	×	✓	×	×	✓	×	✓
Picture trait Controls	×	×	×	✓	✓	×	×	×	✓	✓	×	×	×	✓	✓
Picture specific random effects	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Profile specific random effects	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Log Likelihood	128.61	45.12	83.69	89.4	-36.88	-3169.94	-2859.47	-3042.62	-3150.06	-2708.79	-2.26	-58.26	-41.79	-39.23	-132.29
Observations	800	800	800	800	800	463	463	463	463	463	462	462	462	462	462

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

Table J.12: 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 send two types of requests (mentor and application), with have a double the sample size in Panel C. The regressions are conducted on the profile level, use various controls, and all follow Equation 1.

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

	Response Rate				Message Length (in char)				Highly Useful Message?			
	Native		Alien		Native		Alien		Native		Alien	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Constant	0.21***	0.34	0.21***	0.36	73.47***	94.17	93.94***	−155.01	0.08***	−0.10	0.07***	0.40
	(0.01)	(0.23)	(0.02)	(0.25)	(6.26)	(121.13)	(7.75)	(140.68)	(0.02)	(0.40)	(0.02)	(0.38)
Profile is Black	0.01	0.01	−0.02	−0.02	0.87	2.20	−22.73*	−16.07	−0.002	0.001	−0.02	−0.06$^{\cdot}$
	(0.02)	(0.02)	(0.02)	(0.02)	(8.83)	(12.39)	(10.85)	(14.19)	(0.03)	(0.04)	(0.03)	(0.04)
State Controls	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓
Job Controls	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓
Firstname Controls	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓
Lastname Controls	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓
Picture trait Controls	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓
Picture specific random effects	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Log Likelihood	113.44	-30.23	90.54	-39.88	-965.65	-586.75	-909.56	-504.49	16.51	-76.91	32.34	-65.85
Observations	200	200	200	200	177	177	162	162	176	176	162	162

Panel B: Differences in messages accounting for profile quality

	Response Rate				Message Length (in char)				Highly Useful Message?			
	Native		Alien		Native		Alien		Native		Alien	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Constant	0.21***	0.35	0.19***	0.33	72.57***	77.84	108.33***	−101.76	0.09**	0.12	0.09**	0.55
	(0.02)	(0.23)	(0.02)	(0.26)	(8.90)	(125.87)	(11.68)	(144.48)	(0.03)	(0.40)	(0.03)	(0.39)
Profile is Black	0.06*	0.05	−0.003	0.003	5.28	2.84	−32.96*	−32.64$^{\cdot}$	−0.01	−0.04	−0.04	−0.08
	(0.02)	(0.03)	(0.03)	(0.03)	(12.45)	(16.98)	(15.69)	(19.64)	(0.04)	(0.05)	(0.04)	(0.05)
Profile attended worse Uni	0.01	0.03	0.02	0.03	1.81	8.22	−25.60	−30.97	−0.01	−0.12*	−0.03	−0.09
	(0.03)	(0.03)	(0.03)	(0.03)	(12.59)	(16.81)	(15.57)	(20.04)	(0.05)	(0.05)	(0.04)	(0.05)
Profile is Black and attended worse Uni	−0.09*	−0.08*	−0.04	−0.04	−9.07	−2.55	16.87	29.34	0.01	0.09	0.04	0.03
	(0.04)	(0.04)	(0.04)	(0.04)	(17.76)	(22.40)	(21.76)	(26.73)	(0.06)	(0.07)	(0.06)	(0.07)
State Controls	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓
Job Controls	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓
Firstname Controls	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓
Lastname Controls	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓
Picture trait Controls	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓
Picture specific random effects	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Log Likelihood	112.47	-33.24	85.96	-44.27	-958.58	-579.1	-900.75	-495.53	12.2	-78.16	28.22	-68.24
Observations	200	200	200	200	177	177	162	162	176	176	162	162

Panel C: Differences in messages accounting for message type

	Response Rate				Message Length (in char)				Highly Useful Message?			
	Native		Alien		Native		Alien		Native		Alien	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Constant	0.29***	0.37	0.28***	0.43$^{\cdot}$	325.47***	35.16	401.98***	−260.85	0.08**	−0.07	0.08***	0.36
	(0.02)	(0.23)	(0.02)	(0.25)	(22.78)	(376.64)	(32.50)	(503.85)	(0.03)	(0.44)	(0.02)	(0.40)
Profile is Black	0.01	0.002	−0.03	−0.03	−36.98	−29.11	−72.09	−56.57	0.004	0.02	−0.02	−0.07
	(0.03)	(0.03)	(0.03)	(0.03)	(32.71)	(42.02)	(46.10)	(59.84)	(0.04)	(0.05)	(0.03)	(0.05)
Mentor message	−0.15***	−0.15***	−0.15***	−0.15***	−23.20	5.62	−47.16	−42.98	0.02	0.02	−0.05	−0.07
	(0.03)	(0.03)	(0.03)	(0.03)	(37.66)	(41.08)	(49.15)	(53.98)	(0.05)	(0.05)	(0.04)	(0.04)
Profile is Black and mentor message	0.01	0.01	0.01	0.01	66.12	22.78	106.35	140.95$^{\cdot}$	0.02	−0.02	0.003	0.03
	(0.04)	(0.04)	(0.04)	(0.04)	(52.78)	(59.01)	(70.01)	(79.40)	(0.07)	(0.08)	(0.06)	(0.06)
State Controls	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓
Job Controls	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓
Firstname Controls	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓
Lastname Controls	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓
Picture trait Controls	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓
Picture specific random effects	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Profile specific random effects	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Log Likelihood	72.86	-70.13	49.07	-83.21	-1648.71	-1168.02	-1491.94	-977.57	-31.9	-118.05	34.26	-73.68
Observations	400	400	400	400	247	247	216	216	247	247	215	215

Notes: $^{\cdot}p<0.10$; $^{*}p<0.05$; $^{**}p<0.01$; $^{***}p<0.001$.

Table J.13: 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

	Response Rate									
	Which profile's connection request did targets accept in the first stage of the experiment?									
	Black only		White only		Both		Not both		All targets	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Constant	0.19***	0.64	0.18***	−0.09	0.23***	0.26	0.19***	0.26	0.23***	0.27$^{\cdot}$
	(0.02)	(0.45)	(0.02)	(0.31)	(0.01)	(0.20)	(0.02)	(0.21)	(0.01)	(0.14)
Profile is Black	0.03	0.03	−0.04	−0.03	0.005	0.01	0.03	0.04	0.005	0.01
	(0.03)	(0.04)	(0.03)	(0.03)	(0.02)	(0.02)	(0.03)	(0.03)	(0.02)	(0.02)
Stage 1: Accepted Only Black									−0.04	−0.03
									(0.03)	(0.03)
Stage 1: Accepted Only White							−0.01	−0.003	−0.05*	−0.05*
							(0.03)	(0.03)	(0.02)	(0.02)
Profile is Black Stage 1: Accepted Only Black									0.02	0.02
									(0.04)	(0.04)
Profile is Black Stage 1: Accepted Only White							−0.07$^{\cdot}$	−0.09*	−0.05	−0.04
							(0.04)	(0.05)	(0.03)	(0.03)
State Controls	×	✓	×	✓	×	✓	×	✓	×	✓
Job Controls	×	✓	×	✓	×	✓	×	✓	×	✓
Firstname Controls	×	✓	×	✓	×	✓	×	✓	×	✓
Lastname Controls	×	✓	×	✓	×	✓	×	✓	×	✓
Picture trait Controls	×	✓	×	✓	×	✓	×	✓	×	✓
Picture specific random effects	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Profile specific random effects	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Log Likelihood	-307.5	-407.73	-337.26	-454	-1027.42	-1162.44	-649.06	-789.87	-1684.84	-1854.15
Observations	591	591	827	827	1,846	1,846	1,418	1,418	3,264	3,264

	Message Length (in char)									
	Black only		White only		Both		Not both		All targets	
	(11)	(12)	(13)	(14)	(15)	(16)	(17)	(18)	(19)	(20)
Constant	0.22$^{\cdot}$	−7.87	−0.002	2.70	0.001	−0.35	0.21	−1.75	0.0001	−1.07
	(0.13)	(6.88)	(0.16)	(4.59)	(0.07)	(1.05)	(0.15)	(1.69)	(0.07)	(0.87)
Profile is Black	−0.45**	−1.61***	−0.01	−0.27	0.02	0.002	−0.45*	−0.21	0.02	0.02
	(0.17)	(0.48)	(0.23)	(0.52)	(0.10)	(0.11)	(0.21)	(0.27)	(0.10)	(0.11)
Stage 1: Accepted Only Black									0.20	0.17
									(0.15)	(0.16)
Stage 1: Accepted Only White							−0.22	0.14	−0.01	−0.02
							(0.20)	(0.28)	(0.14)	(0.15)
Profile is Black Stage 1: Accepted Only Black									−0.45*	−0.46*
									(0.21)	(0.22)
Profile is Black Stage 1: Accepted Only White							0.45	0.13	−0.04	0.005
							(0.29)	(0.39)	(0.20)	(0.22)
State Controls	×	✓	×	✓	×	✓	×	✓	×	✓
Job Controls	×	✓	×	✓	×	✓	×	✓	×	✓
Firstname Controls	×	✓	×	✓	×	✓	×	✓	×	✓
Lastname Controls	×	✓	×	✓	×	✓	×	✓	×	✓
Picture trait Controls	×	✓	×	✓	×	✓	×	✓	×	✓
Picture specific random effects	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Profile specific random effects	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Log Likelihood	-154.47	-84.5	-214.36	-141.8	-567.12	-564.48	-379.47	-343.74	-951.29	-971.1
Observations	121	121	128	128	424	424	249	249	673	673

	Highly Useful Message?									
	Black only		White only		Both		Not both		All targets	
	(21)	(22)	(23)	(24)	(25)	(26)	(27)	(28)	(29)	(30)
Constant	0.06$^{\cdot}$	−5.06	0.07**	−1.29	0.08***	0.24	0.07$^{\cdot}$	0.20	0.08***	0.20
	(0.04)	(8.93)	(0.03)	(0.80)	(0.02)	(0.32)	(0.03)	(0.43)	(0.02)	(0.23)
Profile is Black	0.02	−0.69	−0.04	0.08	0.01	−0.02	0.02	−0.06	0.01	−0.01
	(0.05)	(0.65)	(0.04)	(0.09)	(0.03)	(0.03)	(0.05)	(0.07)	(0.03)	(0.03)
Stage 1: Accepted Only Black									−0.01	−0.0003
									(0.04)	(0.04)
Stage 1: Accepted Only White							0.01	−0.07	−0.01	−0.01
							(0.05)	(0.07)	(0.04)	(0.04)
Profile is Black Stage 1: Accepted Only Black									−0.002	−0.01
									(0.06)	(0.06)
Profile is Black Stage 1: Accepted Only White							−0.06	0.04	−0.05	−0.05
							(0.06)	(0.10)	(0.05)	(0.06)
State Controls	×	✓	×	✓	×	✓	×	✓	×	✓
Job Controls	×	✓	×	✓	×	✓	×	✓	×	✓
Firstname Controls	×	✓	×	✓	×	✓	×	✓	×	✓
Lastname Controls	×	✓	×	✓	×	✓	×	✓	×	✓
Picture trait Controls	×	✓	×	✓	×	✓	×	✓	×	✓
Picture specific random effects	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Profile specific random effects	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Log Likelihood	-7.37	-46.11	2.06	-56.79	-60.33	-176.01	-5.92	-105.75	-69.62	-206.3
Observations	115	115	126	126	411	411	241	241	652	652

Notes: $^{\cdot}p<0.10$; $^{*}p<0.05$; $^{**}p<0.01$; $^{***}p<0.001$.

Table J.14: 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.

J.3.3 Consequences of swapping the profile picture

	Views			
	(1)	(2)	(3)	(4)
Constant	39.22***	38.73***	38.57***	37.63***
	(0.58)	(5.92)	(0.58)	(5.92)
Picture swapped	0.33	0.55	0.26	0.43
	(0.83)	(0.61)	(0.83)	(0.62)
Weeks after swapping			0.99***	0.99***
			(0.04)	(0.04)
Picture swapped x Weeks after swapping			0.11$^{\cdot}$	0.11$^{\cdot}$
			(0.06)	(0.06)
State Controls	×	✓	×	✓
Job Controls	×	✓	×	✓
Firstname Controls	×	✓	×	✓
Lastname Controls	×	✓	×	✓
Picture trait Controls	×	✓	×	✓
Picture specific random effects	✓	✓	✓	✓
Profile specific random effects	✓	✓	✓	✓
Log Likelihood	-3525.63	-3282.82	-3251.06	-3008.25
Observations	1,599	1,599	1,599	1,599

Notes: $^{\cdot}p<0.10$; $^{*}p<0.05$; $^{**}p<0.01$; $^{***}p<0.001$.

Table J.15: 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.

	Probability of connection being suspended																							
	All connected targets												Connected targets who **did not** receive message						Connected targets who **did** receive message					
	All profiles		Black profiles		White profiles		All profiles		Black profiles		White profiles		All profiles		Black profiles		White profiles		All profiles		Black profiles		White 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)
Constant	0.01***	0.02	0.01***	0.03	0.01***	0.03	0.01***	0.02	0.01***	0.03	0.01***	0.02	0.01***	0.01	0.01***	0.02	0.01***	0.03	0.01***	0.03	0.01***	0.08	0.01**	0.03
	(0.001)	(0.01)	(0.002)	(0.02)	(0.002)	(0.02)	(0.002)	(0.01)	(0.002)	(0.02)	(0.002)	(0.02)	(0.002)	(0.02)	(0.002)	(0.02)	(0.002)	(0.03)	(0.002)	(0.03)	(0.004)	(0.06)	(0.003)	(0.04)
Picture Swapped	−0.0004	−0.002	−0.001	−0.004·	−0.0001	−0.004*	0.001	−0.0004	−0.001	−0.004	0.003	−0.001	0.001	−0.0003	−0.001	−0.004·	0.003	−0.002	−0.003	−0.002	−0.003	−0.01	−0.003	−0.01·
	(0.001)	(0.001)	(0.001)	(0.002)	(0.001)	(0.002)	(0.001)	(0.002)	(0.002)	(0.003)	(0.002)	(0.003)	(0.001)	(0.002)	(0.002)	(0.002)	(0.002)	(0.003)	(0.003)	(0.003)	(0.01)	(0.01)	(0.004)	(0.005)
Target received message							0.004·	0.004·	0.003	0.004	0.004	0.004												
							(0.002)	(0.002)	(0.003)	(0.003)	(0.003)	(0.003)												
Target received message x Picture Swapped							−0.005	−0.004	0.001	0.001	−0.01*	−0.01·												
							(0.003)	(0.003)	(0.005)	(0.005)	(0.005)	(0.005)												
State Controls	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓
Job Controls	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓
Firstname Controls	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓
Lastname Controls	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓
Picture trait Controls	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓
Picture specific random effects	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Target specific random effects	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Log Likelihood	9895.09	9550	4751.85	4432.68	5164.35	4850.36	9886.26	9541.17	4744.3	4425.33	5156.65	4842.43	6246.88	5920.96	3300.05	2993.34	3028.1	2732.91	3219.05	2917.87	1366.93	1111.51	1952.89	1659.02
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.16: 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				Message Length (in char)				Highly Useful Message?			
	Swapped Only		All		Swapped Only		All		Swapped Only		All	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Constant	0.19***	0.36	0.19***	0.26·	0.09	−2.52	0.09	−1.47·	0.07***	0.50	0.07***	0.21
	(0.01)	(0.23)	(0.01)	(0.14)	(0.07)	(1.83)	(0.06)	(0.88)	(0.02)	(0.43)	(0.02)	(0.23)
Time passed between 28.07 and acceptance	0.02	0.02·	0.02	0.02·	0.02	0.01	0.02	0.02	−0.01	−0.01	−0.01	−0.01
	(0.01)	(0.01)	(0.01)	(0.01)	(0.07)	(0.08)	(0.07)	(0.07)	(0.02)	(0.02)	(0.02)	(0.02)
Picture not swapped			0.03*	0.03·			−0.17*	−0.24*			0.01	0.02
			(0.01)	(0.02)			(0.09)	(0.10)			(0.02)	(0.02)
Picture not swapped x Time passed between 28.07 and acceptance			−0.01	−0.02			−0.05	−0.07			0.02	0.01
			(0.02)	(0.02)			(0.09)	(0.10)			(0.02)	(0.03)
State Controls	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓
Job Controls	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓
Firstname Controls	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓
Lastname Controls	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓
Picture trait Controls	×	✓	×	✓	×	✓	×	✓	×	✓	×	✓
Picture specific random effects	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Log Likelihood	-818.99	-952.42	-1710.55	-1877.97	-476.99	-448.89	-965.91	-987.22	-30.82	-133.64	-63.98	-202.81
Observations	1,677	1,677	3,317	3,317	320	320	681	681	313	313	660	660

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

Table J.17: 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 one 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), and (10) focus only on those targets that have been messaged by profiles whose picture has been swapped. The remaining columns focus on all targets.

J.3.4 Response probabilities and heterogeneity

	Response Rate		Message Length (in char)		Highly Useful Message?	
	(1)	(2)	(3)	(4)	(5)	(6)
Constant	0.13***	0.11	0.05	−0.10	0.02	−0.02
	(0.03)	(0.08)	(0.19)	(0.44)	(0.05)	(0.13)
White	−0.01	−0.004	0.15	0.23$^{\cdot}$	−0.02	−0.03
	(0.02)	(0.02)	(0.10)	(0.12)	(0.03)	(0.04)
Male	0.01	0.003	−0.03	−0.12	0.07**	0.08**
	(0.02)	(0.02)	(0.09)	(0.10)	(0.03)	(0.03)
Age	0.0005	0.0002	−0.002	−0.001	−0.001	−0.001
	(0.001)	(0.001)	(0.005)	(0.01)	(0.001)	(0.002)
Bachelor+	0.01	−0.01	0.12	−0.02	0.02	0.01
	(0.02)	(0.02)	(0.11)	(0.13)	(0.03)	(0.04)
Contact Count	0.0002***	0.0002**	−0.001$^{\cdot}$	−0.001*	0.0002*	0.0002*
	(0.0000)	(0.0001)	(0.0003)	(0.0003)	(0.0001)	(0.0001)
HR.Job		−0.04		0.04		0.05
		(0.03)		(0.17)		(0.05)
Same Uni		0.04		0.09		−0.004
		(0.03)		(0.13)		(0.04)
UniWhite		0.05		−0.32		−0.10
		(0.09)		(0.45)		(0.13)
UniBlack		−0.01		0.09		0.02
		(0.06)		(0.34)		(0.10)
Share Democrat		−0.09		0.58$^{\cdot}$		0.20$^{\cdot}$
		(0.06)		(0.35)		(0.10)
Nbr. of Profile's friends		0.004*		−0.001		−0.003
		(0.002)		(0.01)		(0.002)
Log Likelihood	-1296.81	-1007.41	-734.38	-577.61	-76.99	-67.6
Observations	2,541	1,952	514	406	511	404
Notes:				$^{\cdot}$p<0.10;*p<0.05;**p<0.01;***p<0.001.		

Table J.18: Response probability and usefulness by target characteristics.
The table reports upon regressions estimating several response characteristics as a function of multiple target characteristics.

J.3.5 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)
Constant	5.39***	8.18***	5.30***	5.41***	5.12***	5.06***	7.76***
	(0.09)	(0.40)	(0.22)	(0.18)	(0.24)	(1.05)	(1.07)
Profile is White	−0.98***	−0.98***	−0.98***	−0.99***	−0.97***	−0.95***	−0.96***
	(0.10)	(0.10)	(0.10)	(0.10)	(0.10)	(0.11)	(0.11)
State Controls	×	✓	×	×	×	×	✓
Job Controls	×	×	✓	×	×	×	✓
Firstname Controls	×	×	×	✓	×	×	✓
Lastname Controls	×	×	×	×	✓	×	✓
Picture trait Controls	×	×	×	×	×	✓	✓
Picture specific random effects	✓	✓	✓	✓	✓	✓	✓
Log Likelihood	-653.92	-561.47	-655.97	-649.47	-653.27	-676.41	-585.02
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)
Constant	5.38***	8.17***	5.29***	5.42***	5.13***	5.00***	7.71***
	(0.13)	(0.41)	(0.25)	(0.20)	(0.26)	(1.07)	(1.09)
Profile is Black	−1.01***	−1.01***	−1.01***	−1.07***	−1.00***	−0.97***	−1.04***
	(0.14)	(0.14)	(0.14)	(0.14)	(0.14)	(0.15)	(0.15)
Profile attended worse Uni	0.02	0.02	0.02	−0.04	−0.01	0.04	−0.04
	(0.18)	(0.15)	(0.19)	(0.18)	(0.19)	(0.19)	(0.16)
Profile is Black and attended worse Uni	0.06	0.07	0.06	0.14	0.07	0.04	0.16
	(0.20)	(0.19)	(0.20)	(0.19)	(0.20)	(0.20)	(0.20)
State Controls	×	✓	×	×	×	×	✓
Job Controls	×	×	✓	×	×	×	✓
Firstname Controls	×	×	×	✓	×	×	✓
Lastname Controls	×	×	×	×	✓	×	✓
Picture trait Controls	×	×	×	×	×	✓	✓
Picture specific random effects	✓	✓	✓	✓	✓	✓	✓
Log Likelihood	-655.48	-563.29	-657.5	-650.83	-654.82	-677.94	-586.52
Observations	400	400	400	400	400	400	400

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

Table J.19: 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.

J.4 Back-of-Envelope Calculation

	Income (USD)	
	(1)	(2)
Number of connections	95.32***	70.53***
	(2.02)	(2.36)
Controls	×	✓
Log Likelihood	-232698.75	-197062.7
Observations	18,898	16,052
Notes:	$^{\cdot}$p<0.10;*p<0.05;**p<0.01;***p<0.001.	

Table J.20: Back-of-Envelope Calculation: The economic benefit of an additional connection

Note: The above table shows the results of a linear regression of income on the number of connections of a given target. Column (2) includes controls for age, age squared, level of education, race, and gender.

J.5 Validation experiment

J.5.1 First stage: Captcha task

	Picture selected as computer generated						
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)
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)		
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
Log Likelihood	−1,571.71	−2,094.90	−2,090.80	−1,578.06	−2,098.44	−2,108.04	−4,279.30
Notes:					$^{\cdot}$p<0.10;*p<0.05;**p<0.01;***p<0.001.		

Table J.21: 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 raters age, gender, education, income, ethnicity, and political preference. In the last column, we reweight our sample to match the sample characteristics of LinkedIn users.

	Picture selected as computer generated						
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)
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)		
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
Log Likelihood	−908.28	−1,606.83	−1,602.05	−920.12	−1,612.53	−1,610.23	−2,825.68
Notes:					·p<0.10;*p<0.05;**p<0.01;***p<0.001.		

Table J.22: 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 raters age, gender, education, income, ethnicity, and political preference. In the last column, we reweight our sample to match the sample characteristics of LinkedIn users.

J.5.2 Second stage: Individual rating task

Panel A: Age of person in the picture								
	How old is the person in this picture?							
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)
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)		
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: Gender of person in the picture								
	How likely is the person in this picture female?							
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)
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)		
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: Person in the picture is black								
	How likely is the person in this picture African American?							
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)
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)		
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: Person in the picture is white								
	How likely is the person in this picture White?							
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)
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)		
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 E: Person in the picture is asian								
	How likely is the person in this picture Asian?							
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)
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)		
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.23: Regression estimates on the differences in rated demographic charactersitcs 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 raters age, gender, education, income, ethnicity, and political preference. In the last column, we reweight our sample to match the sample characteristics of LinkedIn users.

Panel A: Trustworthiness of person in the picture

	How trustworthy do you think is the person in this picture?							
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)
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)		
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?							
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)
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)		
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?							
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)
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)		
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?							
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)
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)		
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.24: 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 raters age, gender, education, income, ethnicity, and political preference. In the last column, we reweight our sample to match the sample characteristics of LinkedIn users.

J.5.3 Third stage: University rating task

	The better universites are correctly identified as such									
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)
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)	
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
Log Likelihood	−9,455.55	−9,624.21	−9,622.93	−9,448.13	−9,458.68	−9,458.28	−9,458.66	−9,620.01	−9,599.60	−17,445.28

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

Table J.25: 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.

J.6 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?							
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)
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)
Observations	269	269	269	269	269	269	269	254	269	269	269	269	269	269	269	254
R^2	0.00	0.005	0.01	0.004	0.0001	0.01	0.001	0.01	0.00	0.01	0.01	0.004	0.0002	0.0000	0.0001	0.01
Adjusted R^2	0.00	0.001	0.01	−0.0000	−0.004	0.01	−0.002	0.01	0.00	0.01	0.002	0.001	−0.004	−0.004	−0.004	0.003

Notes: $\cdot p<0.10$; $^{*}p<0.05$; $^{**}p<0.01$; $^{***}p<0.001$.

Table J.26: 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 one and stage two 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 female, *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.

	How many more connections do White profiles have relative to Black profiles?															
	By user's age								By user's gender							
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)	10.28*** (0.68)	10.73*** (0.71)	10.52*** (0.69)	10.04*** (1.73)	10.21*** (1.03)	2.57 (4.20)	10.63*** (1.59)	14.19*** (1.83)
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									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)	2.39 (2.16)
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.27: 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 male or female. *Is Female* indicates whether the expert is female, *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. As all experts have been asked multiple questions, we account for subject-specific heterogeneity by using a subject-specific random effect.

	How many more connections do White profiles have relative to Black profiles?															
	By user's race								By user's level of education							
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)
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)								
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)
Knows this research		−8.09* (3.66)								−4.43 (3.37)						
NonBlack:Knows this research		11.22* (5.18)														
Knows results			−14.16** (5.28)								−9.98* (4.85)					
NonBlack:Knows results			16.26* (7.47)													
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)
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.28: 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 female, *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. As all experts have been asked multiple questions, we account for subject-specific heterogeneity by using a subject-specific random effect.