The effect of group identity on hiring decisions with incomplete information

FORTUNA CASORIA

Université Bourgogne Franche-Comté, e-mail: fortuna.casoria@bsb-education.com

ERNESTO REUBEN

New York University Abu Dhabi, Center for Behavioral Institutional Design, Luxembourg Institute of Socio-Economic Research, e-mail: ereuben@nyu.edu

CHRISTINA ROTT

Vrije Universiteit Amsterdam, Tinbergen Institute, e-mail: c.e.rott@vu.nl

ABSTRACT

We investigate the effects of group identity on hiring decisions with adverse selection problems. We run a laboratory experiment in which employers cannot observe a worker's ability nor verify the veracity of the ability the worker claims to have. We evaluate whether sharing an identity results in employers discriminating in favor of ingroup workers, and whether it helps workers and employers overcome the adverse selection problem. We induce identities using the minimal group paradigm and study two settings: one where workers cannot change their identity and one where they can. Although sharing a common identity does not make the worker's claims more honest, employers strongly discriminate in favor of ingroup workers when identities are fixed. Discrimination cannot be explained by employers' beliefs and hence seems to be taste-based. When possible, few workers change their identity. However, the mere possibility of changing identities erodes the employers' trust towards ingroup workers and eliminates discrimination.

This version: June 2022

Note: We thank participants of the BeeLab meeting at Maastricht University, and participants of several seminars and conferences for their comments and suggestions. The authors gratefully acknowledge financial support from the Graduate School of Business and Economics of Maastricht University and Tamkeen under the NYU Abu Dhabi Research Institute Award CG005.This is the authors' version of work that was accepted for publication in *Management Science*. Changes resulting from the publishing process may not be reflected in this document. The published version can be found at https://doi.org/10.1287/mnsc.2022.4462.

1 Introduction

A large stream of research has provided consistent evidence of discrimination in hiring decisions.[1] Recently, scholars have shown that employers tend to favor workers similar to themselves in terms of tastes, leisure pursuits, and experiences (e.g., Rivera, 2012). Among other explanations, these findings are consistent with ingroup favoritism—the tendency displayed by individuals to treat members of their identity group more favorably than those from different identities. Whereas ingroup favoritism is commonly associated with unfavorable consequences, it could also help organizations overcome problems arising from individuals' self-interest. After all, many studies show that sharing a common identity increases cooperation (e.g., Eckel and Grossman, 2005), coordination (e.g., Chen and Chen, 2011), and trust (e.g., Falk and Zehnder, 2013).[2]

This study investigates the role of ingroup favoritism in overcoming adverse selection problems in hiring decisions. The hiring process is of interest because employers often do not know the workers' abilities before hiring them and evidence provided by workers is typically unverifiable. If there is a large share of low-ability workers, an adverse selection problem arises as it is sub-optimal for employers to hire at all. The combination of unverifiable communication with visible group identities turns the hiring decision into a strategic game where both *taste-based* and *statistical* discrimination can occur.[3] In this setting, if sharing a common identity decreases the workers' willingness to lie and the employers' belief about the workers' truthfulness, then it can help them overcome the adverse selection problem. Whereas there are numerous studies on ingroup favoritism, to our knowledge, ours is the first to investigate the interaction of group identity and truthfulness in hiring decisions with adverse selection problems.

We run a laboratory experiment using an adverse selection hiring game (Charness and Dufwenberg, 2011). In the game, an employer decides whether to hire a worker or not. The worker can be of low or high ability, but his ability is private information, hence unknown to the employer. Before the hiring decision, the worker sends a message to the employer in which

[1] A lot of the literature studies discrimination against racial and ethnic minorities as well as women in the likelihood of being interviewed or hired. In addition, individuals with lower social status, proxied by their names or region, have been found to experience discrimination in many (high-status) professions (e.g., Riach and Rich, 2002; Bertrand and Mullainathan, 2004; Oreopoulos, 2011; Edo et al., 2019). Some research has even shown that common accents are more hirable, whereas regional accents are discriminated against (Rakić et al., 2011).

[2] We concentrate on situations where identities are observable and therefore a no-identity condition cannot exist. In other words, we focus on differences in the hiring of ingroup and outgroup workers without distinguishing between ingroup favoritism and outgroup hostility. See (Charness and Dufwenberg, 2011) for evidence in situations without identities.

[3] This paper focuses on taste-based discrimination that arises from ingroup favoritism. There can be other types of taste-based discrimination that do not depend on a shared group identity (e.g., animus towards individuals) that we do not study here.

he can claim to be of either ability. The prediction with the standard assumption of own-payoff maximization is for the worker to send the high-ability message and for the employer not to hire. However, as Charness and Dufwenberg (2011) show, the adverse selection problem is alleviated if a substantial fraction of workers are unwilling to lie about their ability, making it profitable for the employer to hire.

We introduce minimal group identities before participants play the game (Tajfel, 1970). Thereafter, we randomly assign participants to roles (employer or worker), workers to abilities (high or low), and employers are matched with either an ingroup or an outgroup worker. We ensured that participants know that abilities are randomly assigned, which rules out an association between ability and group identity by design. Based on insights from the literature, we conjecture that group identities can affect hiring outcomes. First, workers might lie less to ingroup than to outgroup employers. Second, employers might discriminate in favor of ingroup workers because they expect ingroup workers to be more truthful (statistical discrimination) or because they exhibit altruism towards ingroup members (taste-based discrimination). We elicit the employer's expectations about the workers' truthfulness to distinguish between these two forms of discrimination.

Empirical and anecdotal evidence shows that workers respond to discrimination by adapting aspects of their identity. For instance, job seekers might change their name, disguise their accent, or opt for gender-free applications (Arai and Thoursie, 2009; Biavaschi et al., 2017).[4] In fact, Akerlof and Kranton (2000) argue that choosing one's identity may be one of the most important economic decisions people make.[5] Inspired by these papers, we implement two treatments: one with fixed identities and one with flexible identities where workers can choose to keep their initial identity or adopt the employer's identity. We investigate not only whether workers change their identity but also how the option to change identity affects the lying and hiring decisions of ingroup and outgroup employer-worker pairs.

We find that identity does not affect lying because workers lie equally often to ingroup and outgroup employers. However, employers exhibit ingroup favoritism by hiring ingroup workers significantly more often than outgroup workers. Interestingly, employers do not believe messages from ingroup workers more than those of outgroup workers. In other words, discrimination in favor of ingroup workers does not emerge because of statistical discrimination, suggesting that it is driven by tastes. With flexible identities, we find that few workers change their identity.

[4]In these studies, identity and individual traits are potentially correlated. An advantage of our experiment is that we can rule out an association between identity and a worker's ability.

[5]Other research considers identity choice in the context of multidimensional identities. Shayo (2009) analyzes how identification affects support for redistribution. Bernard et al. (2016) study the role of identity choice in shaping social structures. Muñoz-Herrera (2021) studies how changing identity can facilitate the social integration of minority groups.

Nevertheless, the possibility of changing identity eliminates discrimination in favor of ingroup workers but also reduces overall hiring rates.

2 The experiment

2.1 Adverse selection hiring game

We implement a variation of the game used by Charness and Dufwenberg (2011). In the game, an employer is matched with a worker of *low* or *high* ability. The employer knows that the probability of being matched with a high-ability worker is $\frac{1}{2}$, but only the worker knows their realized ability. After learning their ability, the worker sends a cheap-talk message to the employer. Workers choose between the message "I am in the low-ability group" and "I am in the high-ability group". After receiving the message, the employer decides whether to hire the worker or not. The game tree is depicted in Figure 1. If the employer does not hire, the employer and the worker get €7 irrespective of the worker's ability. If the employer does hire, the worker gets €10, and the employer's earnings depend on the worker's ability. If the worker is low ability, the employer gets €0. If the worker is high ability, the employer receives €12 with a probability of $\frac{5}{6}$ and €0 otherwise. As in Charness and Dufwenberg (2011), this feature guarantees that false messages by low-ability workers are contractually nonverifiable.[6]

2.2 Experimental design

The experiment consists of two parts. Participants are informed that they will receive the corresponding instructions at the beginning of each part.

Part 1: Group identity

In part 1, we induce group identity using the minimal group paradigm (Tajfel, 1970). As in many studies, we use the participants' revealed preferences to induce identities.[7] We ask the participants to choose one of two smartphones: an iPhone 6 or a Samsung S6 Edge. Both

[6] Like in Charness and Dufwenberg (2011), there is a final decision, not depicted in Figure 1, where the worker chooses between 'Accept' and 'Reject' after learning the employer's decision. Figure 1 shows the payoffs if the worker accepts. Rejecting is a dominated action because it gives both players a payoff of €5. Not surprisingly, 97.5% (117 out of 120) of the workers accept. Because this decision does not affect the theoretical predictions or results, we omit it from our analysis. This decision is interesting in Charness and Dufwenberg (2011) because it is relevant in their other treatments. In our case, it is not, but we decided to keep it for our results to be comparable to theirs.

[7] A commonly-used approach asks participants for their preference over paintings by Klee and Kandinsky and then assign them to groups according to their stated tastes (e.g., Chen and Li, 2009; Gioia, 2017; Kranton and Sanders, 2017). Others have used preferences over movie genres (Dickinson et al., 2018), colors (Charness et al., 2007), and poetry (Kranton and Sanders, 2017).

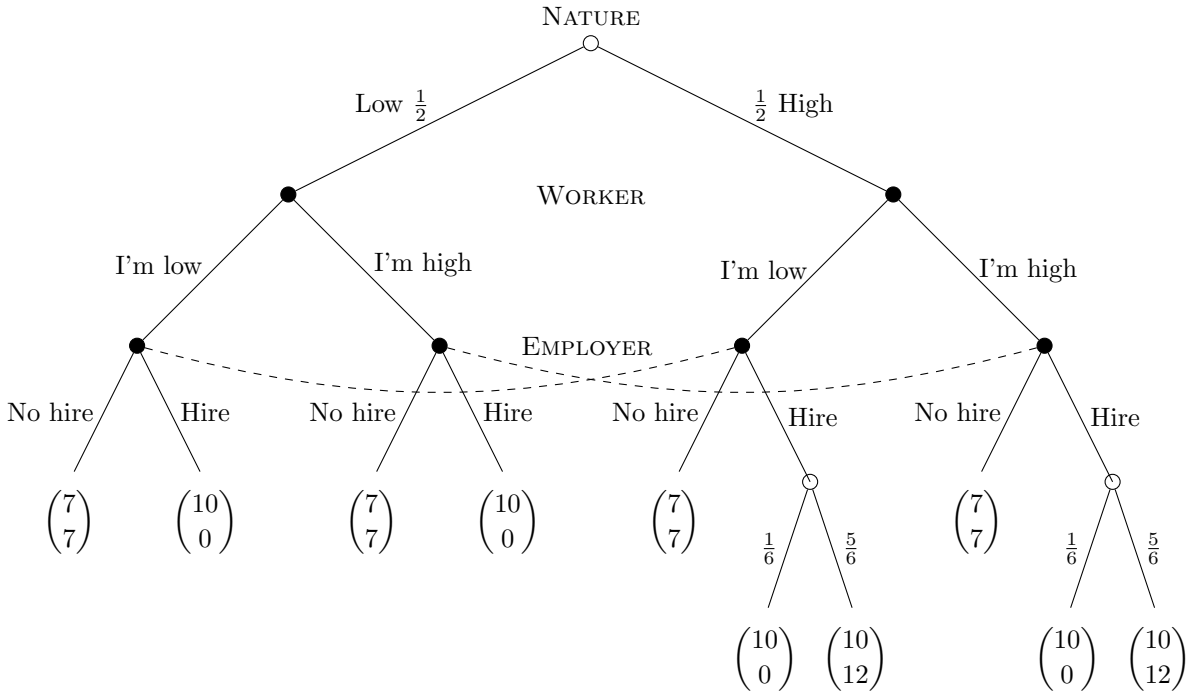

Figure 1. Game tree of the adverse selection hiring game

smartphones have similar functionality, features, and price (around €750 when the study was conducted). To incentivize their decision, we conduct a lottery with a 1-in-750 chance of winning the chosen smartphone.[8]

Part 2: Adverse selection hiring game and belief elicitation

Part 2 consists of two stages. One stage is randomly drawn at the end of the session to determine everyone's payment. In the first stage, participants learn their role (employer or worker), observe each other's identity, and play the one-shot adverse selection hiring game described above. In the second stage, we elicit beliefs. Specifically, immediately after the hiring decision, we ask employers to indicate the probability that they are matched with a low-ability worker. We use the incentive-compatible belief elicitation mechanism proposed by Karni (2009), which is robust to varying levels of risk aversion and deviations from expected utility maximization.[9]

Treatments

We run two treatments. In treatment *Fixed*, the identity chosen in part 1 cannot be changed. In treatment *Flexible*, workers can revise their identity choice after observing their employer's

[8]Participants know that their choice is anonymous and will be used in the second part of the experiment. They also indicate the strength of their preference for the chosen smartphone on a Likert scale.

[9]The mechanism is implemented through a two-step die throw procedure described in the instructions (see the replication materials). We also ask workers to predict their employer's expectation of being matched with a low-ability worker.

identity but before sending their message. Employers know that workers can change identity, but they do not know whether the identity they observe is the worker's initially-chosen identity.

2.3 Conjectures

If all players are rational own-payoff maximizers, the adverse selection hiring game is easily solved. Intuitively, if the employer conditions her hiring on the worker's message, then high-ability and low-ability workers have an incentive to always send the message that results in a higher probability of being hired. However, if workers always send the same message, then messages are uninformative of the worker's ability. In this case, the employer's expected earnings from not hiring (i.e., €7) exceed those from hiring (i.e., $\frac{1}{2} \times \frac{5}{6} \times$ €12 = €5). Hence, in equilibrium, employers do not condition their hiring on the message and never hire, making workers indifferent to what message to send.

Next, we consider how these predictions change if we assume some workers are unwilling to lie. Our goal is to provide a straightforward benchmark describing the conditions under which employers have an incentive to hire. For simplicity, we describe the case where players are risk-neutral, but the general intuition applies to other risk preferences. Over the past decades, substantial evidence has accumulated that some individuals have a preference for truth-telling (e.g., see Abeler et al., 2019). Here, we simply assume that a fraction θ of low-ability workers maximize their monetary earnings whereas the remaining $(1 - \theta)$ are truthful and send the low-ability message. Under this assumption, if we denote the employer's updated belief of being matched with a high-ability worker as b_H, then she prefers to hire as long as her earnings from hiring ($b_H \times \frac{5}{6} \times$ €12) exceed those from not hiring (i.e., €7). In other words, employers hire if their updated belief is above the threshold $b_H^* = \frac{7}{10}$. If employers hire workers who send a high-ability message, then earnings-maximizing high-ability and low-ability workers have a dominant strategy to send the high-ability message. Consequently, because the initial probability of a high-ability worker is $\frac{1}{2}$, the probability that a worker is high-ability conditional on observing a high-ability message boils down to $1/(1 + \theta)$. Combining this with b_H^* gives us the threshold fraction of low-ability workers who lie below which employers are willing to hire: $\theta^* = \frac{3}{7}$.

The role of identity

Here, we discuss the role of group identity, starting with the case where identities are *fixed*. There is considerable evidence that people favor ingroup over outgroup members in numerous domains, from simple allocation decisions to cooperation and trust games (e.g., see Eckel and Grossman, 2005; Chen and Li, 2009; Chen and Chen, 2011; Falk and Zehnder, 2013). A few studies report that ingroup favoritism applies to lying behavior. Rong et al. (2016) find that a shared identity decreases lying in guessing games preceded by a communication stage. Using

natural identities, Maximiano and Chakravarty (2016) find that senders in a sender-receiver game lie less to ingroup (i.e., friends) than to outgroup receivers (i.e., strangers). Butler (2014) finds less lying in ingroup matches in a repeated lemons market game.[10] Given this evidence, we propose the following conjecture.

Conjecture 1 *Fewer workers will lie to ingroup than to outgroup employers.*

One reason for workers to lie less to ingroup employers is altruism towards ingroup members (Chen and Li, 2009). In this case, workers will lie less to ingroup employers to increase their earnings. Alternatively, it is plausible that it is psychologically costlier to lie to an ingroup than an outgroup member. After all, lying is often seen as immoral and moral decisions depend on the closeness between the decision-maker and the potential victim (Bénabou et al., 2020). Although these are distinct reasons, in the simple model above, they boil down to a lower fraction of workers lying in ingroup than in outgroup pairs.[11]

Like workers, there are two straightforward reasons for employers to discriminate in favor of their ingroup. The first reason is discrimination based on taste—i.e., altruism towards ingroup members. Employers increase workers' earnings by hiring them. Hence, altruism towards ingroup members can lead to a higher likelihood of hiring an ingroup worker if the employers' belief b_H is not too extreme.[12] The second reason is statistical discrimination. In other words, employers favor ingroup workers because they believe they are less likely to lie (anticipating Conjecture 1), which is consistent with the evidence showing that individuals expect others to lie less to ingroup members (Benistant and Villeval, 2019). These arguments give us a second conjecture.

Conjecture 2 *Employers are more likely to hire an ingroup than an outgroup worker.*

Because we elicit the employers' beliefs, we can further disentangle empirically taste-based and statistical discrimination. Because ability is randomly assigned, we can rule out differences

[10] Not all studies find evidence of ingroup favoritism in lying. Feldhaus and Mans (2014) find no effect of social identity on lying in a sender-receiver game, whereas Benistant and Villeval (2019) find the same result in a Tullock contest with communication.

[11] One could differentiate between these two reasons with the workers' expectations of each message's impact on the hiring decision. Altruism towards the ingroup predicts a positive association between the relative impact of sending the high-ability message and the likelihood of sending the truthful message. We refrained from measuring these beliefs because incentivized belief elicitation of counterfactual actions is inordinately complicated.

[12] If we define the utility of an employer as $u = \pi + \alpha$, where π is the employer's pecuniary payoff, and α is the utility of increasing the worker's earnings by €2, then the employer hires if her updated belief is above $b_H^* = \frac{7-\alpha}{10}$. If altruism is higher towards ingroup than outgroup members (i.e., $\alpha^I > \alpha^O$), then there is discrimination in favor of the ingroup for beliefs $b_H \in \left(\frac{7-\alpha^I}{10}, \frac{7-\alpha^O}{10}\right)$.

in prior beliefs about the ability of ingroup versus outgroup workers. However, statistical discrimination can arise after employers receive a high-ability message. Namely, if employers expect that ingroup workers are more truthful than outgroup workers, then their updated belief that a worker is a high type conditional on receiving a high-ability message will be higher for ingroup workers than outgroup workers. In other words, assuming that employers correctly anticipate conjecture 1, we can formulate our third conjecture.

Conjecture 3 *Employers' belief of being matched with a high-ability worker after a high-ability message will be higher for ingroup than for outgroup workers, resulting in statistical discrimination.*

To conclude, we discuss the *Flexible* treatment. In this treatment, workers can switch their initial identity before revealing it to employers. Because research on changing minimal identities is rare,[13] it is unclear whether individuals will use initial or final identities to treat others as an ingroup or an outgroup member. If final identities are used and there is ingroup favoritism (Conjectures 1 through 3), then workers have a strong incentive to match the identity of the employer, resulting in overall less lying and more hiring. If initial identities are used, then employers would like to hire workers with whom they share an initial identity, but they cannot tell by observing the final identity whether their initial identities match. This introduces a second adverse-selection problem, which could result in the breakdown of the effect of group identity on hiring. Hence, the effect of flexible identities is ambiguous. Finally, research on natural identities shows that switching one's identity is psychologically costly (Burke, 2006; Chandra, 2006). We use minimal group identities, but even a small psychological cost might be enough to deter workers from switching in the experiment.

2.4 Procedures

We conducted the experiment at BEElab in Maastricht University in 2015. Participants were undergraduate students recruited with ORSEE (Greiner, 2015). We ran ten sessions (five per treatment) with a total of 240 participants (120 per treatment). Each session took one hour, and participants earned, on average, €12, including a €5 show-up fee. Instructions were written in neutral language. The replication materials include samples of the instructions, the z-Tree software programs (Fischbacher, 2007), and the data (see Casoria et al., 2022).

[13] A few researchers have explored settings where individuals can change their affiliation to identity groups. Hargreaves Heap and Zizzo (2009) allow participants to trade group affiliations to play trust games. In Charness and Shmidov (2014), participants playing a public goods game can exit, exclude, and add others to their identity group. Hett et al. (2020) measure group identification preferences and their effect on distributional choices. Robin et al. (2014) find that participants strategically change their opinion to match those of principals, who, in turn, reward like-minded people.

3 Results

We collected 60 independent observations (i.e., employer-worker pairs) per treatment. In *Fixed*, we have 35 outgroup and 25 ingroup pairs, whereas in *Flexible*, we have 28 outgroup and 32 ingroup pairs. Throughout the analysis, we use two-sided tests and power calculations based on $\alpha = 0.10$. We use the worker's initial identity in *Flexible* to determine whether workers and employers form an ingroup or an outgroup pair. The results are unaffected if we use the workers' assumed identity.

Only 10.0% of workers change their identity (6 out of 60). Why is this the case? Intriguingly, the identity change seems unaffected by the employer-worker pairing as workers in ingroup and outgroup pairs change their identity similarly often (3 out of 28 vs. 3 out of 32; χ^2 test $p = 0.863$). Moreover, the strength of preferences for the chosen smartphone does not seem to play a role as it is similar for workers who change identity and those who do not (Mann-Whitney U test, $p = 0.669$). A compelling reason for this behavior is that workers do not expect to benefit from changing identity, which we show later is the correct belief.[14]

3.1 Workers' lying behavior

As expected, low-ability workers lie significantly more often than high-ability workers. Across both treatments, 46.7% of low-ability workers lie (28 out of 60), whereas only 1.7% of high-ability workers do so (1 out of 60; χ^2 test, $p < 0.001$). Henceforth, we focus on the behavior of low-ability workers.

Figure 2 depicts how frequently low-ability workers lie. In *Fixed*, 38.5% of low-ability workers lie in ingroup pairs (5 out of 13), which is slightly less often than the 52.9% who lie in outgroup pairs (9 out of 17). In *Flexible*, it is the other way around: 53.5% of low-ability workers lie in ingroup pairs (8 out of 15), and 40.0% lie in outgroup pairs (6 out of 15). These differences are not statistically significant (χ^2 tests, $p = 0.431$ in *Fixed* and $p = 0.464$ in *Flexible*). Note that these fractions are close to the threshold below which employers are willing to hire ($\theta^* \approx 42.9\%$). We also do not find evidence that being in the *Fixed* or *Flexible* treatment affects lying (χ^2 tests, $p = 0.464$ for outgroup pairs and $p = 0.431$ for ingroup pairs). Overall, we do not find support for Conjecture 1.[15]

[14] Changing identity is costly because it implies giving up the chance to win one's preferred smartphone. Thus, if workers do not expect a monetary benefit from changing, there is no reason to change.

[15] Overall, we can detect differences above 28.9 percentage points with a power of 0.75. In other words, whereas we cannot exclude the possibility of small differences in lying rates, our evidence is consistent with the findings of Feldhaus and Mans (2014) and Benistant and Villeval (2019) that a shared identity does not have a large impact on lying behavior.

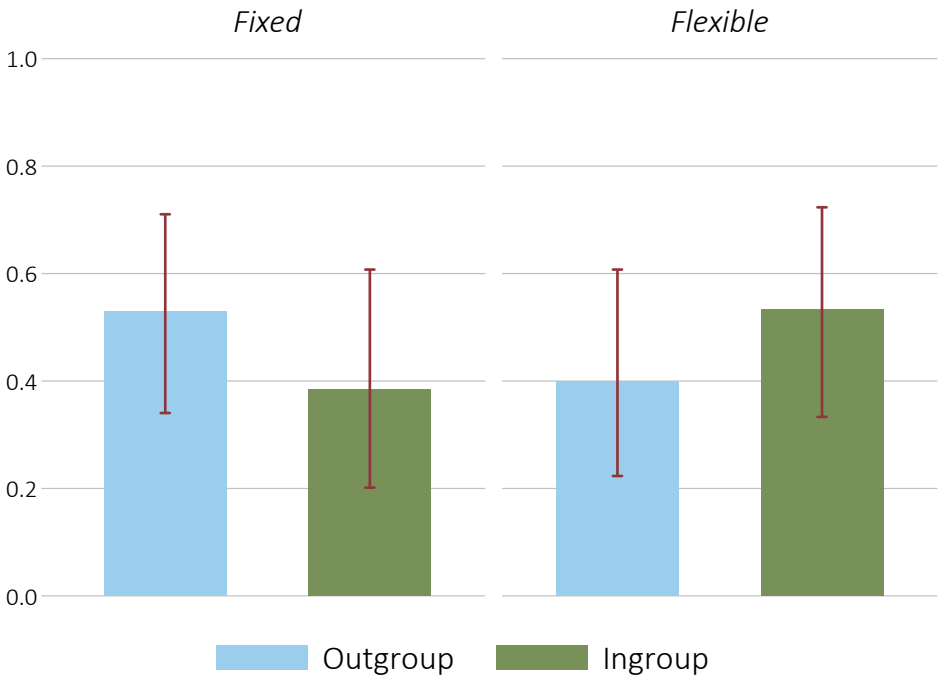

Figure 2. Fraction of low-ability workers who lie to the employer in *Fixed* and *Flexible*

Note: Error bars depict 90% confidence intervals. Ingroup and outgroup pairs are determined by the workers' initial identity.

Result 1 *With both fixed and flexible identities, low-ability workers lie similarly irrespective of whether the employer is an ingroup or an outgroup.*

3.2 Employers' beliefs and hiring behavior

A substantial number of employers hire the worker, and their decision is highly dependent on the worker's message. Overall, 60.9% of the employers who received the high-ability message hire the worker (53 out of 87), whereas only 3.0% of the employers who received the low-ability message do (1 out of 33; χ^2 test, $p < 0.001$). Thus, we focus on the hiring decisions and beliefs of employers who receive the high-ability message from here on.

In *Fixed*, 74.4% of employers hire workers after receiving the high-ability message (32 out of 43). Notably, the employers' average belief of being matched with a high-ability worker equals 69.3%, which is very close to the threshold above which hiring is profitable ($b_H^* = 70\%$). It is also close to the observed fraction of high-ability workers among those who send the high-ability message, namely 67.4%.

Figure 3 depicts the fraction of employers who hire the worker and their mean belief that the worker is of high ability. It shows that 93.8% of employers who receive the high-ability message hire ingroup workers (15 out of 16), but only 63.0% hire outgroup workers (17 out of 27; χ^2 test, $p = 0.025$).[16] Hence, we find evidence of discrimination against outgroup workers,

[16]The statistical power of this test is 0.769.

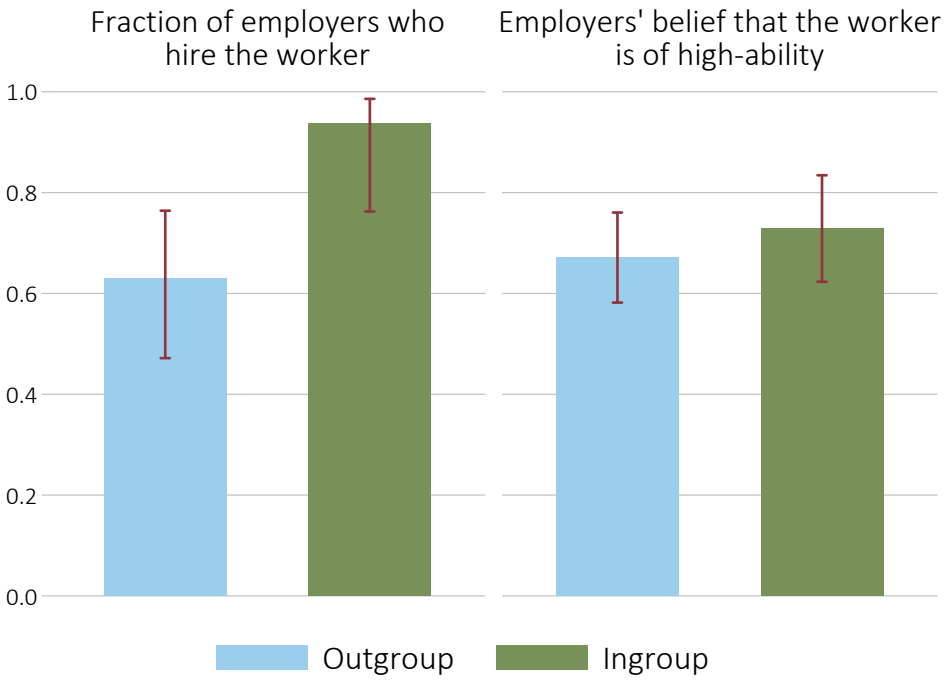

Figure 3. Fraction of employers who hire and their average belief that the worker is of high-ability in *Fixed*

Note: Only for employers who receive the high-ability message. Error bars depict 90% confidence intervals. Ingroup and outgroup pairs are determined by the workers' initial identity.

supporting Conjecture 2. Interestingly, employers' beliefs of being matched with a high-ability worker are not significantly different between ingroup and outgroup pairs (72.9% vs. 67.1%; Mann-Whitney U test, $p = 0.575$), which suggests that employers' discrimination is taste-based rather than statistical (see Conjecture 3).[17] Later, in Table 1, we present further evidence of the importance of taste-based discrimination.

Result 2 *With fixed identities, employers are equally likely to believe the message of ingroup and outgroup workers. However, employers are more likely to hire ingroup than outgroup workers, providing evidence for taste-based rather than statistical discrimination.*

Next, we look at the *Flexible* treatment. In this treatment, only 47.7% of the workers who send the high-ability message are hired (21 out of 44). The fraction of employers who hire is significantly lower in *Flexible* than in *Fixed* (χ^2 test, $p = 0.011$). In line with the lower hiring rate, the employers' belief of being matched with a high-ability worker is significantly lower in *Flexible* than in *Fixed* (54.8% vs. 69.3%; Mann-Whitney U test, $p = 0.020$) and is close to 50%,

[17] As with any null result, we cannot exclude the possibility that there is a small but real difference in beliefs. Given the observed data, the minimal detectable treatment difference with a power of 0.75 is 19.3 percentage points. The observed difference in beliefs of 5.8 percentage points is below this value and is far smaller than the 30.8 percentage point difference in hiring behavior.

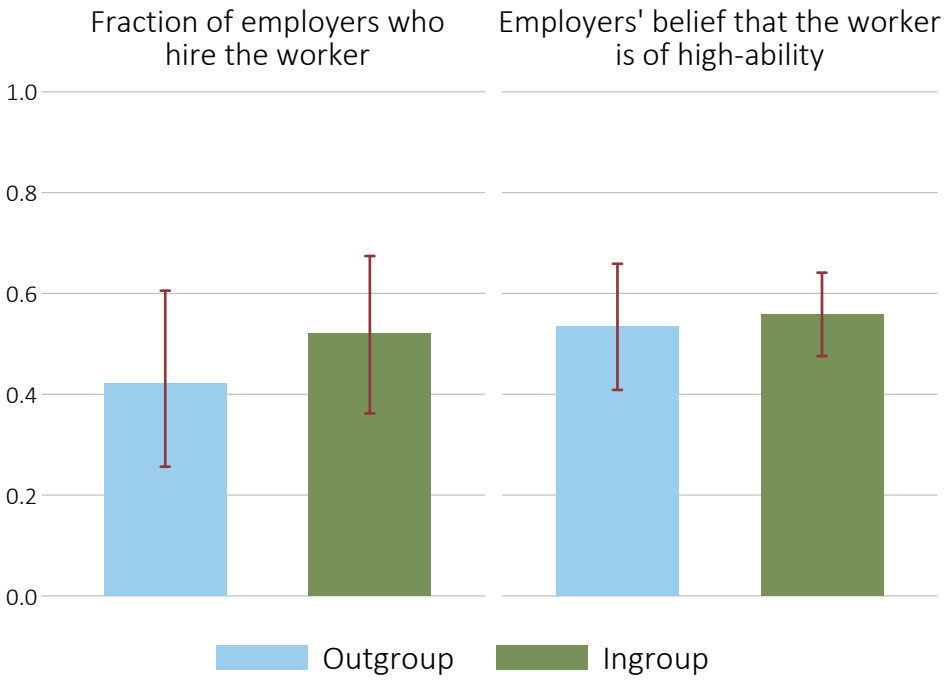

Figure 4. Fraction of employers who hire and their average belief that the worker is of high-ability in *Flexible*

Note: Only for employers who received the high-ability message. Error bars depict 90% confidence intervals. Ingroup and outgroup pairs are determined by the workers' initial identity.

the belief one would hold if the high-ability message is uninformative of the worker's ability.[18]

For employers in *Flexible* who receive the high-ability message, Figure 4 depicts the fraction of them who hire the worker and their mean belief that the worker is of high ability. Overall, 52.0% (13 out of 25) of ingroup workers and 42.1% of outgroup workers are hired (8 out of 19; χ^2 test, $p = 0.515$). Hence, the difference in hiring between *Fixed* and *Flexible* is mainly driven by a difference in the fraction of hired ingroup workers. The hiring rate in ingroup pairs is significantly higher in *Fixed* than in *Flexible* (χ^2 test, $p = 0.005$). At the same time, there is no significant difference for outgroup pairs (χ^2 test, $p = 0.162$). As in *Fixed*, in *Flexible*, we do not find that the employers' belief about the workers' ability differs significantly between ingroup and outgroup pairs (55.8% vs. 53.4%; Mann-Whitney U test, $p = 0.595$). Compared to *Fixed*, employers in *Flexible* are more pessimistic of being matched with a high-ability worker in ingroup pairs (72.9% in *Fixed* vs. 55.8% in *Flexible*; Mann-Whitney U test, $p = 0.049$) and outgroup pairs (67.1% in *Fixed* vs. 53.4% in *Flexible*; Mann-Whitney U test, $p = 0.147$).

In Table 1, we analyze the employers' hiring decisions controlling for their beliefs. Specifically, we run linear probability regressions with the employers' hiring decision as the dependent variable. As above, we restrict the sample to employers who received the high-ability message. In column I, as independent variables, we add indicator variables for the treatment × pair-type

[18]The tests of treatment differences in hiring rates and beliefs have statistical power above 0.80

Table 1. Determinants of the employers' hiring decision

Note: Linear probability regressions. The dependent variable equals one if the employer hires the worker and zero otherwise. Indicator variables for treatment (*Fixed* or *Flexible*) × pair type (ingroup or outgroup), with outgroup pairs in *Fixed* as the reference category. Controls include the employers' self-reported risk aversion, age, nationality, gender, and field of study. Only data from employers who receive a high-ability message. Robust standard errors in parentheses. ***, **, and * indicate statistical significance at 0.01, 0.05, and 0.10.

	I	II	III	IV	V	VI
Ingroup pairs in *Fixed*	0.308***	0.276**	0.301***	0.299**	0.290**	0.327***
	(0.114)	(0.106)	(0.105)	(0.130)	(0.122)	(0.118)
Outgroup pairs in *Flexible*	−0.209	−0.133	−0.130	−0.194	−0.126	−0.124
	(0.150)	(0.145)	(0.145)	(0.147)	(0.143)	(0.144)
Ingroup pairs in *Flexible*	−0.110	−0.048	−0.009	−0.133	−0.073	−0.026
	(0.140)	(0.132)	(0.133)	(0.136)	(0.130)	(0.129)
Belief of high-ability		0.547***			0.507***	
		(0.180)			(0.176)	
Belief of high-ability ≥ 70%			0.359***			0.358***
			(0.096)			(0.096)
Demographic controls	No	No	No	Yes	Yes	Yes
Observations	87	87	87	87	87	87
R^2	0.126	0.212	0.240	0.204	0.272	0.304

combinations (the reference category being outgroup pairs in *Fixed*). This regression simply reproduces the results reported above using non-parametric tests—namely, a significantly higher hiring rate by employers in ingroup pairs in *Fixed*. In columns II and III, we add the employers' belief of being matched with a high-ability worker. In column II, we add beliefs as point predictions. In column III, we add beliefs as a dummy variable that equals 1 if the employer's belief is above or equal to the threshold at which hiring is profitable (i.e., $b^*_H = 70\%$). In both regressions, the coefficient of beliefs is large and statistically significant, confirming the importance of beliefs in the employers' decision.[19]

Interestingly, the introduction of beliefs has two effects on the coefficients of the indicator variables. First, it reduces the magnitude of coefficients of ingroup and outgroup pairs in *Flexible*, making these pairs even more similar to outgroup pairs in *Fixed*. Second, the introduction of beliefs has little effect on the coefficient of ingroup pairs in *Fixed* (post-estimation Wald tests comparing these coefficients across columns I, II, and III, $p > 0.372$). If we interpret this coefficient as the impact of taste-based discrimination, these regressions suggest that employers discriminate because of taste when identities are fixed but not when they are flexible. Columns IV, V, and VI show that these conclusions are robust to controlling for the employers' self-reported tolerance for risk and demographic characteristics (age, gender, nationality, and economics major).

[19]We also ran regressions using dummy variables to divide beliefs into ten equally-spaced categories. The results are qualitatively and quantitatively unchanged in this more flexible specification.

Table 2. Allocation efficiency and average expected earnings by treatment and pair type

Note: Allocation efficiency is the fraction of pairs in which the employer hires a high-ability worker or does not hire a low-ability worker. Statistics are calculated by simulating all possible pairings considering the employers' hiring rate (conditioning on the message received) and the workers' lying rate (conditioning on their ability). Standard deviations in parenthesis.

Condition	Allocation efficiency (in %)	Expected earnings (in €)			
		Overall	Workers	Employers	Employers who hire
Outgroup pairs in *Fixed*	61.9	7.55	8.53	6.57	6.16
	(48.6)	(2.21)	(1.50)	(4.31)	(6.00)
Ingroup pairs in *Fixed*	74.9	7.93	8.83	7.03	7.05
	(43.3)	(2.42)	(1.46)	(4.61)	(5.91)
Outgroup pairs in *Flexible*	62.6	7.46	7.88	7.04	7.14
	(48.4)	(1.75)	(1.37)	(3.20)	(5.89)
Ingroup pairs in *Flexible*	62.1	7.50	8.20	6.81	6.52
	(48.5)	(1.99)	(1.47)	(3.78)	(5.98)

Result 3 *With flexible identities, employers are equally likely to hire and believe the message of ingroup and outgroup workers. Compared to fixed identities, flexible identities reduce the hiring of ingroup members, suggesting that employers no longer discriminate based on taste.*

There are two caveats worth noting concerning Result 3. First, we elicit the employers' average belief, which is a sufficient statistic to measure the impact of beliefs on a binary hiring decision if employers are expected utility maximizers. However, we cannot rule out employers holding more uncertain beliefs for outgroup than for ingroup workers, which combined with non-standard preferences could potentially lead to lower hiring rates. Second, we cannot entirely differentiate whether the ingroup bias differs between majority (iPhone) and minority (Samsung) ingroup pairs or whether there are unobservable differences between these groups that could affect hiring. In Appendix A, we provide a supplementary analysis that suggests that these potential differences do not affect the validity of our results.

3.3 Efficiency

Given the differences in the employers' hiring behavior, it is interesting to investigate the efficiency consequences of ingroup favoritism. We consider two measures of efficiency. The players' earnings and the fraction of pairs in which the employer either hires a high-ability worker or does not hire a low-ability worker, which we call allocation efficiency. To not be constrained by the specific matching of the experiment, we calculated these statistics by considering the employers' mean hiring rate conditional on the message they receive and the workers' lying rate conditional on their ability and then simulating all possible pairings. Table 2 presents the allocation efficiency and average expected earnings by treatment and pair type. It reports the average expected earnings of workers and employers separately.

In all conditions, allocation efficiency and overall earnings are above the no-hiring benchmark obtained with traditional assumptions (i.e., 50% allocation efficiency and €7 in earnings). Comparing across conditions, we see that allocation efficiency and earnings are noticeably higher for ingroup pairs in *Fixed*. This pattern is a direct consequence of there being truthful low-ability workers in all conditions, but a significantly higher hiring rate of workers who send the high-ability message in ingroup pairs in *Fixed*.

Looking at earnings by role, we see that workers earn considerably more than in the no-hiring benchmark (i.e., €7).[20] By contrast, the earnings of employers are close to €7. Looking at the employers' earnings conditional on hiring (last column in Table 2) shows that their expected earnings when they hire are not far from the €7 they earn if they do not hire, especially in ingroup pairs in *Fixed*.[21] This lack of difference might be an important reason we observe taste-based discrimination. Namely, discriminating against outgroup workers is not costly.

4 Conclusion

We examine the effects of group identity on hiring decisions when employers cannot observe the workers' abilities, but workers can communicate their abilities through cheap-talk messages. We ask whether sharing an identity helps workers and employers overcome the adverse selection problem inherent in these decisions and whether the resulting discrimination is statistical or taste-based. We investigate these questions in settings where identities are fixed and flexible.

We find that employers discriminate in favor of ingroup workers when identities are fixed. Notably, employers do not hold differing beliefs about the ability or truthfulness of ingroup and outgroup workers. This leads us to conclude that the observed discrimination is taste-based. In this respect, it is important to note that the workers' truthfulness and the employers' beliefs are such that the average cost of exercising ingroup favoritism is very low, which might be why we observe taste-based discrimination. The literature on identity reports mixed results, from null effects to significant ingroup favoritism (Pechar and Kranton, 2017). The expected cost of discrimination is a plausible explanation for these diverse findings. Another notable result is the effect of group identity on efficiency. Since hiring rates are low because of adverse selection, the increased rate at which employers hire ingroup workers increases overall efficiency. However, because workers are not more honest towards ingroup employers, the benefits of the higher hiring rate are accrued solely by ingroup workers.

[20] Because the workers' earnings depend solely on whether they are hired, their earnings mirror the employers' behavior. Namely, workers who send the high-ability message earn more if they are in an ingroup pair in *Fixed*.

[21] Given the previous results, this is an expected finding. The fraction of low-ability workers who lie is close to $\theta^* \approx 42.9\%$ in all conditions (see Figure 2). In fact, the fraction of lying low-ability workers is not significantly different from this threshold in any treatment $\times$ pair-type combination (Binomial probability tests, $p > 0.443$).

Introducing the possibility to change identity reduces the employers' trust in the workers' truthfulness. Workers are equally likely to lie about their ability, but the employers' belief of being matched with a high-ability worker after seeing a high-ability message is noticeably lower, resulting in a lower hiring rate. This is the case even though the actual number of workers changing identity is extremely low. Flexible identities also eliminate the differential hiring rates between ingroup and outgroup pairs. However, it is unclear why. On the one hand, flexible identities might dampen the taste-based component of ingroup favoritism. On the other hand, the change in the employers' beliefs implies that their expected cost of exercising ingroup favoritism is higher than with fixed identities. Further research would be needed to determine the precise reason for the change.

Overall, our findings in this paper suggest that ingroup favoritism can help alleviate adverse selection problems in hiring decisions. This is a potential explanation for why discrimination in labor markets persists, even if it is taste-based and there is market competition (Becker, 1971). Our findings also suggest that, in hiring decisions where adverse selection is a problem, discrimination ought to be more common for less flexible identities, such as gender and race, than for identities that are easier to change or disguise, such as political and regional identities.

Appendix A Robustness checks

Here, we provide evidence that potential differences between participants who prefer the iPhone or the Samsung smartphones as well as majority/minority group dynamics do not affect the results reported in the main text. In total, 90 (of 120) employers and 83 (of 120) workers chose the iPhone. This stronger preference for the iPhone and the need to have a sufficient number of outgroup pairs resulted in noticeably more iPhone than Samsung ingroup pairs.

For lying behavior of low-ability workers, we did the following robustness checks. First, we compared the lying behavior of ingroup iPhone workers versus that of outgroup iPhone workers. Second, to check for majority/minority group dynamics, we compared the lying behavior of Samsung outgroup workers with iPhone outgroup workers. We do not have enough observations for comparisons with Samsung ingroup pairs. Neither comparisons result in significant or qualitative differences compared to the pooled analysis presented in the main text.

For the hiring behavior in treatment *Fixed*, we checked whether there are differences in the behavior of iPhone versus Samsung employers that could explain the results reported in the main text. To do so, we first ran the regressions I, II, and III from Table 1 with iPhone employers only. Unfortunately, there are too few observations to run these regressions for the Samsung employers only and evaluate the impact of ingroup pairs. The regression results are shown in columns I through III in Table A1. Comparing the coefficient estimates and the significance levels in Table 1 and Table A1 shows that the behavior of iPhone employers is very similar to

Table A1. Robustness checks - determinants of the employer's hiring decision

Note: Linear probability regressions. The dependent variable equals one if the employer hires the worker and zero otherwise, conditional on receiving a high-ability message. Indicator variables for treatment (*Fixed* or *Flexible*) × pair type (ingroup or outgroup) in models I through III, with outgroup pairs in *Fixed* as the reference category. Sample in models I through III are iPhone employers only. Dummy variable for treatment (*Flexible*) and interaction term of treatment dummy (*Flexible*) × employer's identity (iPhone or Samsung) in models IV through VI. Sample in models IV through VI are outgroup employers only. Robust standard errors in parentheses. ***, **, and * indicate statistical significance at 0.01, 0.05, and 0.10.

	I	II	III	IV	V	VI
Ingroup pairs in *Fixed*	0.345**	0.298**	0.322***			
	(0.140)	(0.121)	(0.121)			
Outgroup pairs in *Flexible*	−0.188	−0.058	−0.057			
	(0.202)	(0.189)	(0.184)			
Ingroup pairs in *Flexible*	−0.068	−0.019	0.021			
	(0.161)	(0.142)	(0.144)			
Flexible				−0.256	−0.200	−0.183
				(0.230)	(0.236)	(0.240)
iPhone in *Fixed*				−0.112	−0.044	−0.024
				(0.196)	(0.169)	(0.173)
iPhone in *Flexible*				−0.044	0.139	0.117
				(0.237)	(0.240)	(0.241)
Belief of high-ability		0.606***			0.794***	
		(0.202)			(0.208)	
Belief of high-ability $\geq$ 70%			0.419***			0.467***
			(0.100)			(0.139)
Observations	67	67	67	46	46	46
R^2	0.139	0.238	0.281	0.050	0.238	0.235

the hiring behavior of employers in the pooled analysis. This is particularly important for the (highly) significant coefficient estimates of the indicator variable "Ingroup pairs in *Fixed*". The same ingroup favoritism in treatment *Fixed* can be observed among iPhone employers only, even after controlling for their beliefs. We therefore conclude that differences among participants who prefer iPhone versus Samsung phones do not explain the observed ingroup favoritism.

To check for majority/minority group dynamics, we compare the hiring behavior of Samsung outgroup employers with iPhone outgroup employers. As in the pooled analysis, we condition on employers who receive a high-ability message. Specifically, we ran a series of regressions with the hiring decision as the dependent variable. In all regressions, the independent variables include a dummy variable for treatment *Flexible* and indicator variables for the employer's smartphone choice (iPhone) interacted with treatment dummies (*Fixed* or *Flexible*). The interaction variables allow us to test whether the hiring behavior of iPhone and Samsung employers differs significantly (in each treatment). A significant negative coefficient estimate for "iPhone in *Fixed*" would provide evidence for majority/minority group dynamics in the form of favoritism for majority (iPhone) workers in treatment *Fixed*. The results are shown in columns IV through VI in Table A1. We do not find evidence for this alternative explanation. Thus, we conclude that majority/minority dynamics do not play a role in our experiment.

References

Abeler, J., Nosenzo, D., and Raymond, C. (2019). Preferences for Truth-Telling. *Econometrica*, 87(4):1115–1153.

Akerlof, G. A. and Kranton, R. E. (2000). Economics and Identity. *The Quarterly Journal of Economics*, 115(3):715–753.

Arai, M. and Thoursie, P. S. (2009). Renouncing Personal Names: An Empirical Examination of Surname Change and Earnings. *Journal of Labor Economics*, 27(1):127–147.

Becker, G. S. (1971). *The Economics of Discrimination*. University of Chicago Press, Chicago.

Bénabou, R., Falk, A., and Tirole, J. (2020). Narratives, Imperatives, and Moral Persuasion. Working paper, Princeton University.

Benistant, J. and Villeval, M. C. (2019). Unethical behavior and group identity in contests. *Journal of Economic Psychology*, 72:128–155.

Bernard, M., Hett, F., and Mechtel, M. (2016). Social identity and social free-riding. *European Economic Review*, 90:4–17.

Bertrand, M. and Mullainathan, S. (2004). Are Emily and Greg More Employable Than Lakisha and Jamal? A Field Experiment on Labor Market Discrimination. *American Economic Review*, 94(4):991–1013.

Biavaschi, C., Giulietti, C., and Siddique, Z. (2017). The Economic Payoff of Name Americanization. *Journal of Labor Economics*, 35(4):1089–1116.

Burke, P. J. (2006). Identity Change. *Social Psychology Quarterly*, 69(1):81–96.

Butler, J. V. (2014). Trust, Truth, Status and Identity: An Experimental Inquiry. *The B.E. Journal of Theoretical Economics*, 14(1):293–338.

Casoria, F., Reuben, E., and Rott, C. (2022). Replication materials for: The effect of group identity on hiring decisions with incomplete information. Inter-university Consortium for Political and Social Research, ICPSR-171921-V1.

Chandra, K. (2006). What is ethnic identity and does it matter? *Annual Review of Political Science*, 9(1):397–424.

Charness, G. and Dufwenberg, M. (2011). Participation. *American Economic Review*, 101(4):1211–1237.

Charness, G., Rigotti, L., and Rustichini, A. (2007). Individual Behavior and Group Membership. *American Economic Review*, 97(4):1340–1352.

Charness, G. and Shmidov, V. (2014). Trust and Reciprocity. *Foundations and Trends® in Microeconomics*, 10(3):167–207.

Chen, R. and Chen, Y. (2011). The Potential of Social Identity for Equilibrium Selection. *American Economic Review*, 101(6):2562–2589.

Chen, Y. and Li, S. X. (2009). Group Identity and Social Preferences. *American Economic Review*, 99(1):431–457.

Dickinson, D. L., Masclet, D., and Peterle, E. (2018). Discrimination as favoritism: The private benefits and social costs of in-group favoritism in an experimental labor market. *European Economic Review*, 104:220–236.

Eckel, C. C. and Grossman, P. J. (2005). Managing diversity by creating team identity. *Journal of Economic Behavior & Organization*, 58(3):371–392.

Edo, A., Jacquemet, N., and Yannelis, C. (2019). Language skills and homophilous hiring discrimination: Evidence from gender and racially differentiated applications. *Review of Economics of the Household*, 17(1):349–376.

Falk, A. and Zehnder, C. (2013). A city-wide experiment on trust discrimination. *Journal of Public Economics*, 100:15–27.

Feldhaus, C. and Mans, J. (2014). Who do you lie to? Social identity and the cost of lying. Working Paper Series in Economics 76, University of Cologne.

Fischbacher, U. (2007). z-Tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics*, 10(2):171–178.

Gioia, F. (2017). Peer effects on risk behaviour: the importance of group identity. *Experimental Economics*, 20(1):100–129.

Greiner, B. (2015). Subject pool recruitment procedures: organizing experiments with ORSEE. *Journal of the Economic Science Association*, 1(1):114–125.

Hargreaves Heap, S. P. and Zizzo, D. J. (2009). The Value of Groups. *American Economic Review*, 99(1):295–323.

Hett, F., Mechtel, M., and Kröll, M. (2020). The Structure and Behavioral Effects of Revealed Social Identity Preferences. *The Economic Journal*, 130(632):2569–2595.

Karni, E. (2009). A Mechanism for Eliciting Probabilities. *Econometrica*, 77(2):603–606.

Kranton, R. E. and Sanders, S. G. (2017). Groupy versus Non-Groupy Social Preferences: Personality, Region, and Political Party. *American Economic Review Papers & Proceedings*, 107(5):65–69.

Maximiano, S. and Chakravarty, S. (2016). Deception, Social Preferences, and Friendship. SSRN Working Paper 3589537.

Muñoz-Herrera, M. (2021). Identity change and the cost of social division: Theory and experimental evidence. Working Paper, New York University.

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

Pechar, E. and Kranton, R. (2017). Moderators of Intergroup Discrimination in the Minimal Group Paradigm: A Meta-Analysis. Working paper, Duke University.

Rakić, T., Steffens, M. C., and Mummendey, A. (2011). When it matters how you pronounce it: The influence of regional accents on job interview outcome. *British Journal of Psychology*, 102(4):868–883.

Riach, P. A. and Rich, J. (2002). Field Experiments of Discrimination in the Market Place. *The Economic*

Journal, 112(483):F480–F518.

Rivera, L. A. (2012). Hiring as Cultural Matching. *American Sociological Review*, 77(6):999–1022.

Robin, S., Rusinowska, A., and Villeval, M. C. (2014). Ingratiation: Experimental evidence. *European Economic Review*, 66:16–38.

Rong, R., Houser, D., and Dai, A. Y. (2016). Money or friends: Social identity and deception in networks. *European Economic Review*, 90:56–66.

Shayo, M. (2009). A Model of Social Identity with an Application to Political Economy: Nation, Class, and Redistribution. *American Political Science Review*, 103(2):147–174.

Tajfel, H. (1970). Experiments in Intergroup Discrimination. *Scientific American*, 223(5):96–103.