

DISCUSSION PAPER SERIES

IZA DP No. 15186

**How to Reduce Discrimination?
Evidence from a Field Experiment in
Amateur Soccer**

Robert Dur
Carlos Gomez-Gonzalez
Cornel Nessler

MARCH 2022

DISCUSSION PAPER SERIES

IZA DP No. 15186

How to Reduce Discrimination? Evidence from a Field Experiment in Amateur Soccer

Robert Dur

Erasmus University Rotterdam and IZA

Carlos Gomez-Gonzalez

University of Zurich

Cornel Nesseler

Norwegian University of Science and Technology

MARCH 2022

Any opinions expressed in this paper are those of the author(s) and not those of IZA. Research published in this series may include views on policy, but IZA takes no institutional policy positions. The IZA research network is committed to the IZA Guiding Principles of Research Integrity.

The IZA Institute of Labor Economics is an independent economic research institute that conducts research in labor economics and offers evidence-based policy advice on labor market issues. Supported by the Deutsche Post Foundation, IZA runs the world's largest network of economists, whose research aims to provide answers to the global labor market challenges of our time. Our key objective is to build bridges between academic research, policymakers and society.

IZA Discussion Papers often represent preliminary work and are circulated to encourage discussion. Citation of such a paper should account for its provisional character. A revised version may be available directly from the author.

ISSN: 2365-9793

IZA – Institute of Labor Economics

Schaumburg-Lippe-Straße 5–9
53113 Bonn, Germany

Phone: +49-228-3894-0
Email: publications@iza.org

www.iza.org

ABSTRACT

How to Reduce Discrimination? Evidence from a Field Experiment in Amateur Soccer

A rich literature shows that ethnic discrimination is an omnipresent and highly persistent phenomenon. Little is known, however, about how to reduce discrimination. This study reports the results of a large-scale field experiment we ran together with the Norwegian Football Federation. The federation sent an email to a random selection of about 500 amateur soccer coaches, pointing towards the important role that soccer can play in promoting inclusivity and reducing racism in society and calling on the coaches to be open to all interested applicants. Two weeks later, we sent fictitious applications to join an amateur club, using either a native-sounding or a foreign-sounding name, to the same coaches and to a random selection of about 500 coaches who form the control group. In line with earlier research, we find that applications from people with a native-sounding name receive significantly more positive responses than applications from people with a foreign-sounding name. Surprisingly and unintentionally, the email from the federation substantially increased rather than decreased this gap. Our study underlines the importance of running field experiments to check whether well-intended initiatives are effective in reducing discrimination.

JEL Classification: C93, J15, Z29

Keywords: ethnic discrimination, intervention, field experiment, correspondence test, amateur soccer

Corresponding author:

Robert Dur
Erasmus University Rotterdam
Department of Economics H9-15
P.O. Box 1738
3000 DR Rotterdam
The Netherlands
E-mail: dur@ese.eur.nl

1. Introduction

Ethnic discrimination is a global and persistent phenomenon. A rich body of experimental research in a variety of contexts demonstrates that the ethnic background of a person matters a great deal for the opportunities one gets in society. In the labor market, for example, numerous correspondence tests have shown that ethnic minorities are less likely to receive callbacks for interviews when applying for jobs (Bertrand and Duflo, 2017; Bertrand and Mullainathan, 2004; Kaas and Manger, 2012; Lancee, 2021; Oreopoulos, 2011; Pager et al., 2009; Quillian and Midtbøen, 2021; Riach and Rich, 2002; Thijsen et al., 2021; Weichselbaumer, 2020; Zschirnt and Ruedin, 2016). Research has also found that ethnic minorities face severe discrimination in housing (Auspurg et al., 2019; Diehl et al., 2013; Sawert, 2020), shopping (Bourabain and Verhaeghe, 2019), transportation (Liebe and Beyer, 2021; Mujcic and Frijters 2021), the sharing economy (Edelman et al., 2017; Ge et al., 2020), online dating (Jakobsson and Lindholm, 2014), and sports (Gomez-Gonzalez et al., 2021; Nessler et al., 2019). Meta-analyses by Quillian et al. (2017) and Heath and Di Stasio (2019) show that the extent to which ethnic minorities are discriminated against has hardly changed during the last three decades.

One reason why ethnic discrimination is so persistent over time is that we still know little about which policies reduce discrimination. In a recent review of more than 300 studies, Paluck et al. (2021) conclude that much research is “ill-suited to provide actionable, evidence-based recommendations for reducing prejudice” (p.533). Similarly, Bertrand and Duflo (2017) write that “While field experiments in the last decade have been instrumental in documenting the prevalence of discrimination, field experiments in the future decade should aim to play as large of a role in isolating effective methods to combat it” (p.383).

This paper contributes to this by analyzing the effectiveness of a low-cost intervention to reduce discrimination. The context we study is amateur soccer. Discrimination is prevalent in amateur soccer. In a correspondence test in 22 European countries, Gomez-Gonzalez et al. (2021) show that when asking to join a training, people with foreign-sounding names are significantly less likely to receive a positive response from clubs than people with native-sounding names. In our study, we measure discrimination in the same way, i.e., through a correspondence test. The correspondence test is preceded by an anti-discrimination intervention, run in collaboration with the Norwegian Football Federation (NFF), and implemented for a random selection of amateur soccer clubs. The experimental design allows us to estimate the causal effect of the intervention on discrimination in the field.

The intervention is an email sent by the NFF encouraging soccer coaches to be open to people interested in membership independent of the ethnic background of the applicant. The email describes the important role that soccer can have in bringing people of diverse backgrounds together. It argues that, in this way, soccer can promote interaction and can be key to social inclusion (cf. Lowe, 2021; Mousa, 2020). It also mentions that the NFF finds it important that soccer is multicultural and diverse, reflecting diversity in society. In addition to encouragement, the email gives some information about current discrimination in amateur soccer, mentioning that studies have shown that “players with foreign-sounding names are less likely to get a

response when contacting a club for the first time”. Providing such information about the prevalence of discrimination can be essential – as shown by Boring and Philippe (2021) – because people are not always aware that they discriminate (Bertrand et al., 2005; Rooth, 2010). The NFF sent the email two weeks preceding the correspondence test.

Our main predictions – which we pre-registered (<https://doi.org/10.1257/rct.8049-1.0>) – were that the email from the NFF would lead to an increase in positive responses to email requests to join a training session from people with a foreign-sounding name, whereas we expected no effects for email requests to join a training session from people with a native-sounding name. We defined a positive response as either an invitation to come to a training or a conditional acceptance (e.g., yes, you’re welcome, but only if you are a defender). This was also pre-registered.

Our results show that coaches are less responsive to applications from people with a foreign-sounding name compared to identical applications from people with a native-sounding name. The gap in the positive response rate for the full sample is 11 percentage points, which is almost the same as the gap found two years ago in amateur soccer in Norway, reported in Gomez-Gonzalez et al. (2021).

Concerning the effect of the intervention, we find a surprising result. Instead of reducing the gap in positive response rates, the gap actually increases as a result of the intervention, implying more discrimination. Underlying this effect is a strong increase in coaches’ positive responses to applications from people with a native-sounding name. Positive responses to applications from people with a foreign-sounding name also go up, but only a little.

One possible explanation for our unexpected finding is that the intervention gave rise to feelings of resentment among some coaches, leading them to be more open towards people with a native-sounding name and less open to people with a foreign-sounding name. Such defiant behavior in response to moral appeals by authorities has been found in earlier studies in other contexts, including tax compliance (Ariel, 2012; Blumenthal et al., 2001), vaccination (Nyhan and Reifler, 2015; Nyhan et al., 2014), and criminal offending (Bouffard and Piquero, 2010). Other coaches, however, may have responded in line with our predictions: they were more open to applications with foreign-sounding names and equally open to applications with native-sounding names. On net, we may then observe an increase in the response to applications with native-sounding names and no, or a very small response, to applications with a foreign-sounding name.

We do some further, more exploratory analysis of our data. Inspired by earlier studies that find more severe discrimination in less populous regions (Huijsmans et al., 2021; Mayda, 2006), we split our data by the population size of the region a club resides in. We show that, in the absence of the intervention, discrimination only occurs in less populous regions. Discrimination is severe there, amounting to a gap in positive response rates of more than 20 percentage points. Strikingly, the treatment effect of our intervention in these regions is close to zero. Clubs located in the most populous regions show no discriminatory responses in the absence of the

intervention, but respond strongly to the intervention in the form of an increase in positive responses to applications with native-sounding names. Hence, if feelings of resentment drive the treatment effects, it should be the non-discriminating coaches who have been affected by this, which seems not very plausible.

A possible alternative interpretation could be that non-discriminating coaches were not aware that discrimination occurs in amateur soccer and that the information contained in the email about the presence of discrimination induced them to conform to the descriptive norm to discriminate. That people tend to conform their behavior to others has been found in many earlier studies in a variety of contexts (Allcott and Rogers, 2014; Bott et al., 2020; Bradler et al., 2016; Chen et al., 2010; Frey and Meier, 2004; Gerber and Rogers, 2009; Hallsworth et al., 2016; Schultz et al., 2007;). Interventions that make harmful behavior of others visible can then backfire, as it makes the harmful behavior seem natural and socially acceptable (Bicchieri and Dimant, 2019; Cialdini et al., 1990; Dur and Vollaard, 2015; Kahan, 1997; Keizer et al., 2008). The descriptive norm theory can, however, not explain that the increased discrimination mainly arose from a strong increase in coaches' positive responses to applications from people with native-sounding names.

Our third and last interpretation – which we owe to one of the reviewers – relates to the part of the email that reminds coaches that soccer can promote inclusivity and inter-ethnic interaction. Informing (or reminding) coaches about this may have encouraged them to invite more natives to expose them to this positive contact. This is also consistent with the finding that the increase in positive responses to applications with native-sounding names occurs amongst (counterfactually) non-discriminating coaches. Thus, our surprising result might be an unintentional consequence of a well-meaning gesture amongst non-discriminatory coaches.

Our paper is inspired by and contributes to a small but growing literature testing interventions to reduce discrimination using field experiments. The study that is closest to ours is Boring and Philippe (2021). They study the effects of a similar intervention – an email that includes information about existing discrimination and an appeal not to discriminate – but in a different context, namely student evaluations in higher education. Their results show that the email is effective in reducing gender bias. Likewise, Alesina et al. (2018) find that informing teachers about their implicit bias against immigrants reduces their bias in grading exams of immigrant students. The field experiments on Twitter by Munger (2017) and Hangartner et al. (2021) find that xenophobic hate speech can be reduced by providing counterspeech.

The rest of the paper is structured as follows. Section 2 describes the experimental design and the data. Section 3 presents the results and provides an interpretation. In Section 4, we give a brief summary and offer some concluding remarks.

2. Experimental setup and data

Our field experiment took place in Norway in the Fall of 2021. Norway has a substantial share of first- and second-generation immigrants. In 2021, 18.5% of inhabitants were immigrants or

Norwegian-born to immigrant parents.¹ The largest groups come from Poland, Lithuania, and Somalia. They are spread out across the country, with some overrepresentation in big cities such as Oslo, Stavanger, and Bergen.

It is widely acknowledged that ethnic discrimination is present in Norway. For example, Midtbøen (2016) and Larsen and Di Stasio (2021) performed correspondence tests in the Norwegian labor market. They find that applicants with a Norwegian-sounding name were significantly more likely to receive a callback than applicants with a foreign-sounding name. Andersson et al. (2012) find similar results for the Norwegian housing market. In the context of amateur soccer, Gomez-Gonzalez et al. (2021) find an 11 percentage points higher response rate to applications from native-sounding names.

The intervention we test was created together with and implemented by the NFF. The NFF's purposes and activities include organizing and promoting soccer in Norway. It provides the framework for regional federations to organize and administer their leagues. Our experiment targets amateur soccer clubs in the lowest soccer divisions. While the clubs participating in these leagues are obligated to comply with the regulations and directions from the NFF, they have autonomy about the level of the membership fees and whom to admit as a member.

The NFF provided us with a dataset of the universe of the NFF's amateur adult soccer clubs in Norway. The dataset contains the email of the coach or the club,² postal code of the club, league in which the club plays, and whether it is a male or female club. If a club had more than one adult team, we randomly selected one team to avoid contamination.

Together with the NFF we created an email message that aims to make coaches more open to admit people with a migration or non-native background to their club. The email points to the important role soccer can play in promoting inclusivity and reducing racism in society. The email also describes findings from studies showing that discrimination is present in amateur soccer. The email calls on the coaches to help keep soccer open to all people that are interested.

The exact text of the email was as follows (translated from Norwegian; the original Norwegian text is included in Appendix A):

Subject: Soccer for everyone - for inclusion and against racism

Dear coach,

¹ For detailed statistics see <https://www.ssb.no/en/befolkning/innvandrere/statistikk/innvandrere-og-norskfodte-med-innvandrerforeldre>.

² 620 of the email addresses we used are the personal email addresses of coaches. The remaining 347 email addresses are the email addresses of administrative personnel of the clubs. The rate and nature of responses to applications sent to the two types of email addresses are very similar and do not differ significantly.

Soccer is the world's most popular sport and it gives us responsibility and the opportunity to unite people from different backgrounds. It is important for all of us that soccer is multicultural and diverse and reflects our entire society - not just for professionals, but for all players. At the amateur level, soccer facilitates integration and promotes interaction on and off the field. We aim to make soccer easily accessible to all members of our society.

Racism and exclusion are a societal problem and, thus, also a problem of soccer.

Scientific studies showed that it is more difficult for people with foreign names to join an amateur soccer club. Regardless of sport or language skills, players with foreign names are less likely to get a response when contacting a club for the first time.

This barrier to participating in sports is not only negative for foreigners, but also for Norwegians with a migration background and the Norwegian society. Creating diverse teams with members from different backgrounds is the key to improving social inclusion.

We have sent this email to all coaches in Norway to applaud you on the work you do and encourage everyone to continue to help keep the door to soccer open to all interested.

2021 will be a strange year, but we hope that at least autumn will be almost normal for most people and that together we can look forward to a glorious 2022.

[Name of NFF's representative]

[Logo of NFF]

The email was sent from the email account of an NFF representative on 15 October 2021 at 1:41 pm.

To be able to estimate the effect of the email on discrimination, it was sent to a randomly selected half of all clubs in Norway.³ The remaining clubs in Norway form the control group. We randomly assigned clubs to the treatment group and the control group using a block design. To achieve balance between treatment and control regarding region and type of club, we created two blocks for each of the federation's 18 regions, one with all male clubs in the region and one with all female clubs in the region. We randomized half of the clubs within each of the 36 blocks either to treatment or to control. Thus, 466 clubs were assigned to treatment and so were sent the email from the NFF, whereas 501 were assigned to control. The slight difference in the number of clubs assigned to treatment and control results from the block design. Figure 1 shows how clubs in treatment and control are distributed geographically.

[Figure 1 near here]

³ This contrasts with the message in the email from the NFF that the email was sent to all coaches. In agreement with the NFF, sending out the rest of the emails was postponed at least until the end of the experiment.

Two weeks and three days after the NFF sent the email to the clubs in the treatment group, we sent fictitious applications to join a training session to all clubs. Sending an email is a common way in Europe to get in touch with a sports club. Soccer clubs in Norway typically provide an email address on their website. For practical reasons, we sent out the emails on two days (on Monday 1 November 2021 and Tuesday 2 November 2021) rather than on a single day. We randomly assigned all clubs (independent of treatment) to one of the two days. We constructed email accounts, using Gmail, with typical Norwegian- and foreign-sounding names. We tested the email accounts to confirm that emails do not end up in a spam folder. Following the approach in Nessler et al. (2019), we created names for the three largest foreign groups in Norway (Polish, Lithuanian, and Somali).^{4,5} Each name was randomly assigned to each club. Block randomization ensured that names were not overrepresented in specific blocks.⁶

Clubs received an email with the following text (translated from Norwegian; the original Norwegian text is included in Appendix A):

Subject: *Training possibility*

Hello,

I am looking for a new soccer club. Do you think I could come and join a training session?

Thanks!

[Name]

We collected responses to fictitious applications for six weeks (from 1 November 2021 until 12 December 2021). Almost all responses arrived soon after we sent the application: 79% of responses arrived within two days and 94% of responses arrived in the first week after sending. We received two responses in the last two weeks; two responses in the fifth week, and no responses in the sixth week. In total, we received 554 responses (57% response rate). From the

⁴ The Norwegian statistics bureau (SSB) categorizes the largest foreign groups. For more information, see <https://www.ssb.no/en/befolkning/innvandrere/statistikk/innvandrere-og-norskfodte-med-innvanderforeldre>.

⁵ We asked students before the start of a lecture (in classes in Ålesund and Trondheim in Spring 2021) to fill out a questionnaire and state if a name sounded Norwegian or foreign. The results of the survey are available in Table S1 in Appendix B and show that the names we use are identified as native or foreign by close to 100% of respondents.

⁶ As a result of this randomization, some clubs received an application with a name associated with a group that is not greatly represented in their area, e.g., an application with a Somali-sounding name to a club in a rural area. However, since we used names for the three largest foreign groups in Norway and immigrants are relatively widely spread geographically in Norway, we do not think that this caused any irregularities. In Norway, young people with an immigrant background historically participate less in sports than those without an immigrant background (Walseth, 2007). The differences are particularly salient among young women. While self-limiting behaviours related to culture or religion could at least partially explain this underrepresentation, our experiment studies whether exclusion by clubs also plays a role.

remaining 413 clubs, we received no response. We coded the responses as positive without additional inquiries (n=291), positive with additional inquiries (n=254), and negative (n=9). Positive responses invited an individual to a training, oftentimes specifying day and time. Positive responses with additional inquiries typically asked applicants for information about their age, previous experience in other clubs or divisions, or their preferred playing position. Negative responses typically denied the opportunity to join a training because clubs were full or had had already too many other requests. The response rates are close to the ones found in 2019 in Norway, reported in Gomez-Gonzalez et al. (2021).

Table 1 gives an overview of the data. The complete dataset is publicly available.⁷

[Table 1 near here]

We received ethical approval from the Human Subjects Committee of the Faculty of Economics, Business Administration, and Information Technology of the University of Zurich (OEC IRB # 2021-031) and registered the experiment before we sent out emails (<https://doi.org/10.1257/rct.8049-1.0>). Our experimental setup had three potential ethical issues. First, the research involves subjects that are uninformed about their participation in the study. This approach has the disadvantage that respondents might participate who do not wish to participate in the research. However, not informing respondents has the advantage that we receive an undistorted response. Second, the subjects are misled. We tell the subjects that a person would like to join their club. However, this person does not exist. The subjects invest time to read the request and to answer it – if they answer. To minimize the effort of the subjects, we write a response email soon after a subject contacts us, informing that the applicant is no longer interested. This assures that subjects do not invest more time in the non-existing individual.⁸ Third, the data is not anonymized at the beginning of the research project. However, the name of the club and the name of any individual have been deleted from the dataset right after processing the data. We published our completely anonymized dataset in HarvardDataverse (the data will be available upon publication). While we acknowledge these potential ethical issues, studying interventions to fight discrimination has the potential to create important benefits for society, which should be traded off against the ethical concerns (Asiedu et al., 2021; Glennerster, 2017).

3. Results

Following the pre-registration, we pooled the responses to the fictitious applications into two categories: “negative” and “positive”. Negative responses include declines and non-responses and positive responses include positive responses with and without additional inquiries. Figure 2 shows our main results. It plots the share of positive responses to applications split by native- and foreign-sounding name of the applicant and treatment and control group of the coach.

⁷ The data (upon publication) is available at <https://doi.org/10.7910/DVN/ZETZDZ> (the page will be activated upon publication of the paper). Individually identifiable data is excluded.

⁸ We considered the possibility of debriefing subjects after the experiment, but decided against it following the advice by the Human Subjects Committee, who argued it may create more harm than good.

Confirming previous research, we find evidence for discrimination: people with a foreign-sounding name receive fewer responses. This holds particularly for coaches in the treatment group – where the gap is 15 percentage points – but also in control, where the gap is 8 percentage points. In the Gomez-Gonzalez et al. (2021) study, the gap found for Norway in 2019 in the same context was 11 percentage points.

Next, let us consider the treatment effect of the intervention. Comparing the positive response rates to applicants with foreign-sounding names for the control and treatment group, we find a tiny increase of slightly more than two percentage points. Surprisingly, there is a large positive treatment effect of 10 percentage points on the positive response rate to people with a native-sounding name. The gap in response rates to applications from people with native- and foreign-sounding names thus widened as a result of the treatment. This runs counter to our prediction and was also not intended by the NFF.

[Figure 2 near here]

Table 2 confirms these results using regression analyses, where the dependent variable *PositiveResponse_i* is a dummy indicating whether we received a positive response to the fictitious application (dependent variable equals one) or not (dependent variable equals zero). The regression equation reads:

$$PositiveResponse_i = \alpha_0 + \beta_1 Foreign-sounding\ name_i + \beta_2 Email\ from\ NFF_i \times Native-sounding\ name_i + \beta_3 Email\ from\ NFF_i \times Foreign-sounding\ name_i + \varepsilon_i$$

The subscript *i* refers to a coach or club. β_1 is an estimate of how much a foreign-sounding name matters for getting a positive response, given that no email was sent to the coach or club by the NFF. β_2 is an estimate of the effect of an email from the NFF on the positive response rate to applications from people with native-sounding names. Likewise, β_3 is an estimate of the effect of an email from the NFF on the positive response rate to applications from people with foreign-sounding names. ε_i is the error term. In some regression models, we also control for several background variables: male teams, the day the application was sent out (Monday or Tuesday), and in which league a team plays.

Table 2 shows that discrimination of people with a foreign-sounding name is statistically significant at the 10% level (and at the 5% level after adding controls, see column 3). The letter from the NFF has a statistically significant positive effect on the response rate to applicants with a native-sounding name and hardly any effect on the response rate to applications with a foreign-sounding name. Surprisingly, we find that emails sent on Monday 1 November 2021 receive fewer responses than those sent on Tuesday 2 November 2021.

[Table 2 near here]

How can we interpret these results? One possible interpretation is that the email led to feelings of resentment among some coaches, resulting in defiant responses. As discussed in the Introduction, defiant responses to moral appeals have been observed before in several other contexts. In our experiment, a defiant response would be to become less open to people with a foreign-sounding name and more open to people with a native-sounding name. The fact that, overall, we do not find a negative effect on the responses to applicants with foreign-sounding names might result from other coaches responding in the predicted and intended way: more open to people with a foreign-sounding name and equally open to people with a native-sounding name. On net, the pattern that we see in the data can result: more positive responses to people with a native-sounding name and no change in responses to people with a foreign-sounding name.

Further exploratory analyses

We performed a few more regression analyses, which were not pre-registered and so are of a more exploratory nature.

First, we studied how much our results depend on the exact definition of a positive response. Table S2 in Appendix B uses positive responses without any additional inquiries as the dependent variable instead of the broader definition used above, which includes positive responses with additional inquiries. Using this stricter definition, we find slightly more discrimination in the control group and a much smaller treatment effect for applicants with native-sounding names, which is statistically indistinguishable from zero. Hence, the sizeable and statistically significant treatment effect reported in Table 2 stems mainly from coaches moving from a negative or no response to a conditional positive response, and not so much to an unconditional positive response. Otherwise, the results in Table S2 are by and large similar to those in Table 2.

Next, we study whether our results differ between more and less populous regions. Earlier studies have shown that ethnic discrimination is particularly prevalent in less populous regions (Huijsmans et al., 2021; Mayda, 2006).⁹ Figure 3 and the accompanying regression results in Table 3 confirm this. Focusing on the control group, we find that discrimination is absent (or even slightly positive) in regions with more than 100,000 citizens. In contrast, in regions with less than 100,000 citizens, there is a sizeable difference of more than 20 percentage points between the response rates to applications from people with native-sounding names and those from people with foreign-sounding names.

Moreover, Figure 3 and Table 3 show a striking difference in the estimated effect of the treatment. Whereas the treatment did not have any effect in regions with less than 100,000 citizens, it had a sizeable effect of 18 percentage points in regions with more than 100,000

⁹ 439 clubs are in regions with less than 100,000 citizens and 528 are in regions with more than 100,000 citizens. A map showing where these clubs are located is available in Figure S1 in Appendix B. In regions with less than 100,000 citizens, 21.0% of the teams are female and 79.0% are male teams. In regions with more than 100,000 citizens, 17.8% of the teams are female and 82.2% are male teams.

citizens on positive responses to applications from people with a native-sounding name. Positive responses to applications from people with a foreign-sounding name also went up in these regions, but to a smaller extent (about 8 percentage points) and statistically insignificant.

Summarizing, for the regions where we find sizeable discrimination in the absence of treatment, we find no effect of the treatment whatsoever. The resulting policy implication is that, if one wishes to reduce discrimination where discrimination is prevalent, the intervention we tested in this paper does not help. For the regions where there is no discrimination in the absence of treatment, we find that the treatment leads to a large increase in positive responses to applicants with native-sounding names, and thus leads to discrimination. Hence, if feelings of resentment drive the treatment effect (as we conjectured above), it should be the non-discriminating coaches who have been affected by this, which seems not very plausible.

The results, however, also suggest an alternative interpretation. As discussed in the Introduction, a rich body of empirical evidence shows that people tend to conform to other people's behavior, i.e., they follow descriptive norms. The email from the NFF mentioned that studies have shown that discrimination is present in amateur soccer. In the less populous regions, this message may contain no news: as our data clearly show, discrimination is prevalent there, and coaches may be well aware of this. In more populous regions, however, discrimination is absent and so the message in the email from the NFF that discrimination occurs may change coaches' beliefs about what is the apparent social norm and they may adjust their behavior accordingly. As a result, in those regions, the intervention may have unintentionally backfired. The descriptive norm theory can, however, not explain that the increased discrimination mainly arose from a strong increase in coaches' positive responses to applications from people with a native-sounding name.

Our third and, we believe, most convincing interpretation – which we owe to one of the reviewers – relates to the part of the email that reminds coaches that soccer can promote inclusivity and inter-ethnic interaction. Informing (or reminding) coaches about the potential of amateur soccer to reduce racism may have encouraged them to invite more natives to expose them to inter-ethnic contact. This would also be consistent with the finding that the increase in positive responses to applications with native-sounding names occurs amongst (counterfactually) non-discriminating coaches. Thus, our surprising result might be an unintentional consequence of a well-meaning gesture amongst non-discriminatory coaches.

[Figure 3 near here]

[Table 3 near here]

4. Concluding remarks

We performed a field experiment together with the NFF to examine if a low-cost intervention decreases discrimination in the context of amateur soccer. The intervention was an email from the NFF which described the important role of soccer for inclusivity and fighting racism and

asked coaches to remain open to all interested applicants. Specifically, the goal was to keep a high response rate for all applicants but to decrease the gap between people with a native- and a foreign-sounding name. The email also mentioned that studies have shown that discrimination occurs in amateur soccer. We measured the effect of the intervention by sending fictitious applications to join a training using a native-sounding or a foreign-sounding name.

Our data show that discrimination is prevalent in Norwegian amateur soccer. The gap in positive responses rates to applications with native- and foreign-sounding names amounts to 11 percentage points in the full sample. These results contribute to the debate regarding discrimination of ethnic minorities.

Surprisingly, the low-cost intervention did not reduce discrimination. On the contrary, we find that the intervention increased discrimination. The increased inequality is driven by coaches being more open to applicants with native-sounding names. We discussed three possible interpretations of our findings. The email might have led some coaches to feel resentful and respond in a defiant way. Alternatively, the information in the email that discrimination occurs in amateur soccer may have changed coaches' perception of the descriptive norm, leading them to discriminate more. A third and, we believe, most convincing interpretation of the findings focuses on the part of the email that stresses the role that soccer can play in reducing racism. This could have encouraged coaches of diverse teams to invite more natives. This interpretation is in line with the fact that the increase in positive responses to applications with native-sounding names occurred among coaches from areas where discrimination in amateur soccer is rare.

To further explore the empirical relevance of possible interpretations of our findings, it would have been great if we could have collected more data – e.g., questionnaire data or interview data – about how the intervention is perceived by different types of coaches. This was outside of the scope of the present study, but we intend to include it in possible follow-up studies. Other interesting avenues for future research – that we were not able to do because of data limitations or lack of statistical power – include: i) studying whether the ethnicity and other background characteristics of the coach and the team matter for the severity and direction of discrimination and for the response to the treatment; ii) studying whether the ethnicity of the applicant matters for these outcomes; iii) studying the effects of adding or taking away particular elements of the intervention email (e.g., regarding the information about current discrimination in amateur soccer); iv) studying the influence on the response rate of different days or weeks on which an application is sent.

Our findings are important for practice and future research as they show that, even though an intervention is well-intended and has been proven to be effective in another context (Boring and Philippe, 2021), it can have no or even a negative effect. Many organizations throughout the world have taken initiatives to fight discrimination. However, interventions are rarely tested through large-scale field experiments. Our results show that refraining from doing so implies running the risk of making things worse rather than better (cf. Behaghel et al., 2015).

Acknowledgements

This study was approved by the Human Subjects Committee of the Faculty of Economics, Business Administration, and Information Technology of the University of Zurich (OEC IRB # 2021-031) and pre-registered at the AEA RCT Registry (AEARCTR-0008049; <https://doi.org/10.1257/rct.8049-1.0>). We are grateful to two reviewers of this journal, Anne Boring, Francesco Capozza, Josse Delfgaauw, Donald Green, Sacha Kapoor, and Arnfinn Midtbøen for helpful comments on an earlier version of this paper. We are also thankful for the support of the Norwegian Football Federation and, especially, the help of Henrik Lunde.

Disclosure statement

The authors report there are no competing interests to declare.

References

- Alesina, A., Carlana, M., La Ferrara, E., & Pinotti, P. (2018). Revealing stereotypes: Evidence from immigrants in schools. National Bureau of Economic Research Working Paper No. 25333.
- Allcott, H., & Rogers, T. (2014). The short-run and long-run effects of behavioral interventions: Experimental evidence from energy conservation. *American Economic Review*, 104(10), 3003-3037.
- Andersson, L., Jakobsson, N., & Kotsadam, A. (2012). A field experiment of discrimination in the Norwegian housing market: Gender, class, and ethnicity. *Land Economics*, 88(2), 233-240.
- Ariel, B. (2012). Deterrence and moral persuasion effects on corporate tax compliance: findings from a randomized controlled trial. *Criminology*, 50(1), 27-69.
- Asiedu, E., Karlan, D., Lambon-Quayefio, M.P., & Udry, C.R. (2021). A call for structured ethics appendices in social science papers. *Proceedings of the National Academy of Sciences*, 118(29), e2024570118.
- Auspurg, K., Schneck, A., & Hinz, T. (2019). Closed doors everywhere? A meta-analysis of field experiments on ethnic discrimination in rental housing markets. *Journal of Ethnic and Migration Studies*, 45(1), 95-114.
- Behaghel, L., Crépon, B., & Le Barbanchon, T. (2015). Unintended effects of anonymous résumés. *American Economic Journal: Applied Economics*, 7(3), 1-27.
- Bertrand, M., Chugh, D., & Mullainathan, S. (2005). Implicit discrimination. *American Economic Review Papers & Proceedings*, 95(2), 94-98.
- Bertrand, M., & Duflo, E. (2017). Field experiments on discrimination. In A.V. Banerjee & E. Duflo (eds.), *Handbook of field experiments, Volume 1* (pp. 309-393). Amsterdam: North-Holland Publishing.
- Bertrand, M., & 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.
- Bicchieri, C., & Dimant, E. (2019). Nudging with care: The risks and benefits of social information. *Public Choice*, forthcoming.

- Blumenthal, M., Christian, C., & Slemrod, J. (2001). Do normative appeals affect tax compliance? Evidence from a controlled experiment in Minnesota. *National Tax Journal*, 54(1), 125-138.
- Boring, A., & Philippe, A. (2021). Reducing discrimination in the field: Evidence from an awareness raising intervention targeting gender biases in student evaluations of teaching. *Journal of Public Economics*, 193, 104323.
- Bott, K.M., Cappelen, A.W., Sorensen, E., & Tungodden, B. (2020). You've got mail: A randomised field experiment on tax evasion. *Management Science*, 66(7), 2801-2819.
- Bouffard, L.A., & Piquero, N.L. (2010). Defiance theory and life course explanations of persistent offending. *Crime & Delinquency*, 56(2), 227-252.
- Bourabain, D., & Verhaeghe, P.P. (2019). Could you help me, please? Intersectional field experiments on everyday discrimination in clothing stores. *Journal of Ethnic and Migration Studies*, 45(11), 2026-2044.
- Bradler, C., Dur, R., Neckermann, S., & Non, A. (2016). Employee recognition and performance: A field experiment. *Management Science*, 62(11), 3085-3099.
- Chen, Y., Harper, F.M., Konstan, J., & Li, S.X. (2010). Social comparisons and contributions to online communities: A field experiment on Movielens. *American Economic Review*, 100(4), 1358-1398.
- Cialdini, R.B., Reno, R.R., & Kallgren, C.A. (1990). A focus theory of normative conduct: Recycling the concept of norms to reduce littering in public places. *Journal of Personality and Social Psychology*, 58(6), 1015-1026.
- Diehl, C., Andorfer, V.A., Khoudja, Y., & Krause, K. (2013). Not in my kitchen? Ethnic discrimination and discrimination intentions in shared housing among university students in Germany. *Journal of Ethnic and Migration Studies*, 39(10), 1679-1697.
- Dur, R., & Vollaard, B. (2015). The power of a bad example: A field experiment in household garbage disposal. *Environment and Behavior*, 47(9), 970-1000.
- Edelman, B., Luca, M., & Svirsky, D. (2017). Racial discrimination in the sharing economy: Evidence from a field experiment. *American Economic Journal: Applied Economics*, 9(2), 1-22.
- Frey, B.S., & Meier, S. (2004). Social comparisons and pro-social behavior: Testing conditional cooperation in a field experiment. *American Economic Review*, 94(5), 1717-1722.
- Ge, Y., Knittel, C. R., MacKenzie, D., & Zoepf, S. (2020). Racial discrimination in transportation network companies. *Journal of Public Economics*, 190, 104205.
- Gerber, A.S., & Rogers, T. (2009). Descriptive social norms and motivation to vote: Everybody's voting and so should you. *Journal of Politics*, 71(1), 178-191.
- Glennerster, R. (2017). The practicalities of running randomized evaluations: Partnerships, measurement, ethics, and transparency. In A.V. Banerjee & E. Duflo (eds.), *Handbook of field experiments*, Volume 1 (pp. 175-243). Amsterdam: North-Holland Publishing.
- Gomez-Gonzalez, C., Nessler, C., & Dietl, H.M. (2021). Mapping discrimination in Europe through a field experiment in amateur sport. *Humanities and Social Sciences Communications*, 8(1), 1-8.
- Hallsworth, M., Chadborn, T., Sallis, A., Sanders, M., Berry, D., Greaves, F., Clements, L., & Davies, S.C. (2016). Provision of social norm feedback to high prescribers of antibiotics in

- general practice: A pragmatic national randomised controlled trial. *The Lancet*, 387(10029), 1743-1752.
- Hangartner, D., Gennaro, G., Alasiri, S., Bahrach, N., Bornhoft, A., Boucher, J., Buse Demirci, B., Derksen, L., Hall, A., Jochum, M., Murias Munoz, M., Richter, M., Vogel, F., Wittwer, S., Wüthrich, F., Gilardi, F., & Donnay, K. (2021). Empathy-based counterspeech can reduce racist hate speech in a social media field experiment. *Proceedings of the National Academy of Sciences*, 118(50), e2116310118.
- Heath, A., & Di Stasio, V. (2019). Racial discrimination in Britain, 1969–2017: A meta-analysis of field experiments on racial discrimination in the British labour market. *British Journal of Sociology*, 70, 1774–1798.
- Huijismans, T., Harteveld, E., Van der Brug, W., & Lancee, B. (2021). Are cities ever more cosmopolitan? Studying trends in urban-rural divergence of cultural attitudes. *Political Geography*, 86, 102353.
- Jakobsson, N., & Lindholm, H. (2014). Ethnic preferences in internet dating: A field experiment. *Marriage & Family Review*, 50(4), 307-317.
- Kaas, L., & Manger, C. (2012). Ethnic discrimination in Germany's labour market: A field experiment. *German Economic Review*, 13(1), 1-20.
- Kahan, D.M. (1997). Social influence, social meaning, and deterrence. *Virginia Law Review*, 83, 349-395.
- Keizer, K., Lindenberg, S., & Steg, L. (2008). The spreading of disorder. *Science*, 322, 1681-1685.
- Lancee, B. (2021). Ethnic discrimination in hiring: Comparing groups across contexts. Results from a cross-national field experiment. *Journal of Ethnic and Migration Studies*, 47(6), 1181-1200.
- Larsen, E.N., & Di Stasio, V. (2021). Pakistani in the UK and Norway: different contexts, similar disadvantage. Results from a comparative field experiment on hiring discrimination. *Journal of Ethnic and Migration Studies*, 47(6), 1201-1221.
- Liebe, U., & Beyer, H. (2021). Examining discrimination in everyday life: A stated choice experiment on racism in the sharing economy. *Journal of Ethnic and Migration Studies*, 47(9), 2065-2088.
- Lowe, M. (2021). Types of contact: A field experiment on collaborative and adversarial caste integration. *American Economic Review*, 111(6), 1807-1844.
- Mayda, A.M. (2006). Who is against immigration? A cross-country investigation of individual attitudes toward immigrants. *Review of Economics and Statistics*, 88(3), 510-530.
- Midtbøen, A. H. (2016). Discrimination of the second generation: Evidence from a field experiment in Norway. *Journal of International Migration and Integration*, 17(1), 253-272.
- Mousa, S. (2020). Building social cohesion between Christians and Muslims through soccer in post-ISIS Iraq. *Science*, 369(6505), 866-870.
- Mujcic, R., & Frijters, P. (2021). The Colour of a Free Ride. *Economic Journal*, 131(634), 970-999.
- Munger, K. (2017). Tweetment effects on the tweeted: Experimentally reducing racist harassment. *Political Behavior*, 39, 629–649.
- Nesseler, C., Gomez-Gonzalez, C., Dietl, H. (2019). What's in a name? Measuring access to social activities with a field experiment. *Palgrave Communications*, 5, 1–12.

- Nyhan, B., & Reifler, J. (2015). Does correcting myths about the flu vaccine work? An experimental evaluation of the effects of corrective information. *Vaccine*, 33(3), 459-464.
- Nyhan, B., Reifler, J., Richey, S., Freed, G.L. (2014). Effective messages in vaccine promotion: A randomized trial. *Pediatrics*, 133(4), e835-e842.
- 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), 148-71.
- Pager, D., Bonikowski, B., & Western, B. (2009). Discrimination in a low-wage labor market: A field experiment. *American Sociological Review*, 74(5), 777-799.
- Paluck, E.L., Porat, R., Clark, C.S., & Green, D.P. (2021). Prejudice reduction: Progress and challenges. *Annual Review of Psychology*, 72, 533-560.
- Quillian, L., & Midtbøen, A.H. (2021). Comparative perspectives on racial discrimination in hiring: The rise of field experiments. *Annual Review of Sociology*, 47, 391-415.
- Quillian, L., Pager, D., Hexel, O., & Midtbøen, A.H. (2017). Meta-analysis of field experiments shows no change in racial discrimination in hiring over time. *Proceedings of the National Academy of Sciences*, 114(41), 10870-10875.
- Riach, P.A., & Rich, J. (2002). Field experiments of discrimination in the market place. *Economic Journal*, 112(483), 480-518.
- Rooth, D.O. (2010). Automatic associations and discrimination in hiring: Real world evidence. *Labour Economics*, 17(3), 523-534.
- Sawert, T. (2020). Understanding the mechanisms of ethnic discrimination: A field experiment on discrimination against Turks, Syrians and Americans in the Berlin shared housing market. *Journal of Ethnic and Migration Studies*, 46(19), 3937-3954.
- Schultz, P.W., Nolan, J.M., Cialdini, R.B., Goldstein, N.J., & Griskevicius, V. (2007). The constructive, destructive, and reconstructive power of social norms. *Psychological Science*, 18(5), 429-434.
- Thijssen, L., Lancee, B., Veit, S., & Yemane, R. (2021). Discrimination against Turkish minorities in Germany and the Netherlands: Field experimental evidence on the effect of diagnostic information on labour market outcomes. *Journal of Ethnic and Migration Studies*, 47(6), 1222-1239.
- Walseth, K. (2007). Bridging and bonding social capital in sport—experiences of young women with an immigrant background. *Sport, Education and Society*, 13(1), 1-17.
- Weichselbaumer, D. (2020). Multiple discrimination against female immigrants wearing headscarves. *ILR Review*, 73(3), 600-627.
- Zschirnt, E., & Ruedin, D. (2016). Ethnic discrimination in hiring decisions: A meta-analysis of correspondence tests 1990–2015. *Journal of Ethnic and Migration Studies*, 42(7), 1115-1134.

Table 1

Descriptive statistics. *

Variable	N	Mean	Std.Dev.	Min	Max
Male or female soccer club (Male=1)	967	.808	.394	0	1
Treatment group (email from federation) or control group (Treatment group=1)	967	.482	.500	0	1
Native- or foreign-sounding name of applicant (Native=1)	967	.481	.500	0	1
Application sent out on Monday or Tuesday (Monday=1)	967	.515	.500	0	1
Responses to applications					
Any response **	554	0.57			
Positive response without further inquiries	291	0.30			
Positive response with further inquiries	254	0.26			
Negative response	9	0.01			
No response	413	0.43			

Notes: * In addition to the variables included in the table, we have data about the clubs' leagues (26 different leagues) and regions (18 regions). In 2017, Norway merged several regions into 11 administrative regions. However, the NFF still follows the previous 2017 version.

** If we only received an automatic response, this is counted as no response.

Table 2
Regression results.

	Dependent variable: positive response = 1, no or negative response = 0		
	Model 1	Model 2	Model 3
Native-sounding name	Omitted	omitted	omitted
Foreign-sounding name	-0.08* (0.04)	-0.08* (0.04)	-0.10** (0.05)
Email from NFF × Native-sounding name	0.10** (0.04)	0.09** (0.04)	0.08* (0.05)
Email from NFF × Foreign-sounding name	0.02 (0.04)	0.02 (0.04)	0.04 (0.05)
Male soccer team		-0.01 (0.04)	0.07 (0.10)
Application sent out Monday or Tuesday (Monday=1)		-0.07** (0.03)	-0.06** (0.03)
League and region control			Yes
Constant	0.58*** (0.04)	0.62*** (0.05)	0.28 (0.22)
Observations	967	967	967
R ²	0.017	0.022	0.053

Notes: Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3

Regression results for subsamples by population size of the region. #

	Dependent variable: positive responses = 1, no or negative response = 0	
	Regions with less than 100,000 citizens	Regions with more than 100,000 citizens
Native-sounding name	Omitted	Omitted
Foreign-sounding name	-0.21*** (0.07)	0.01 (0.06)
Email from NFF × Native-sounding name	-0.04 (0.07)	0.18*** (0.06)
Email from NFF × Foreign-sounding name	-0.01 (0.07)	0.08 (0.06)
Male soccer team	0.09 (0.14)	0.04 (0.15)
Application sent out Monday or Tuesday (Monday=1)	-0.04 (0.05)	-0.08* (0.04)
League and region control	Yes	Yes
Constant	0.21*** (0.07)	0.34 (0.29)
Observations	439	528
R^2	0.090	0.074

Notes: #Excluding/Including only Oslo, Hordaland, Rogaland, Trøndelag, Buskerud, and Østfold.
Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Figure 1

Geographic distribution of clubs by treatment assignment.

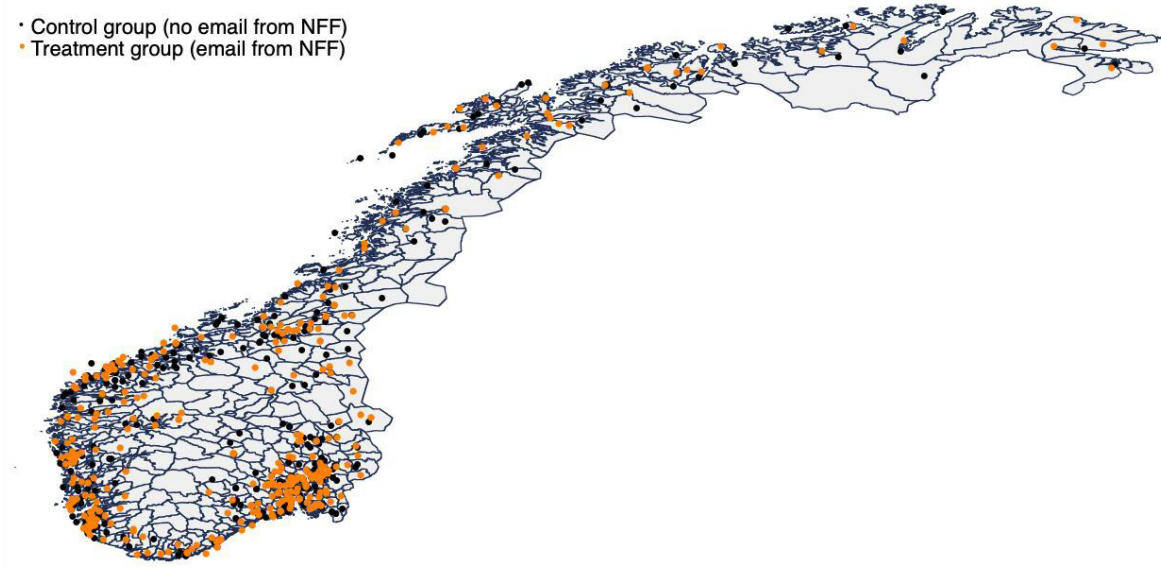


Figure 2

Share of positive responses to applications with native- and foreign-sounding names for treatment and control group coaches.

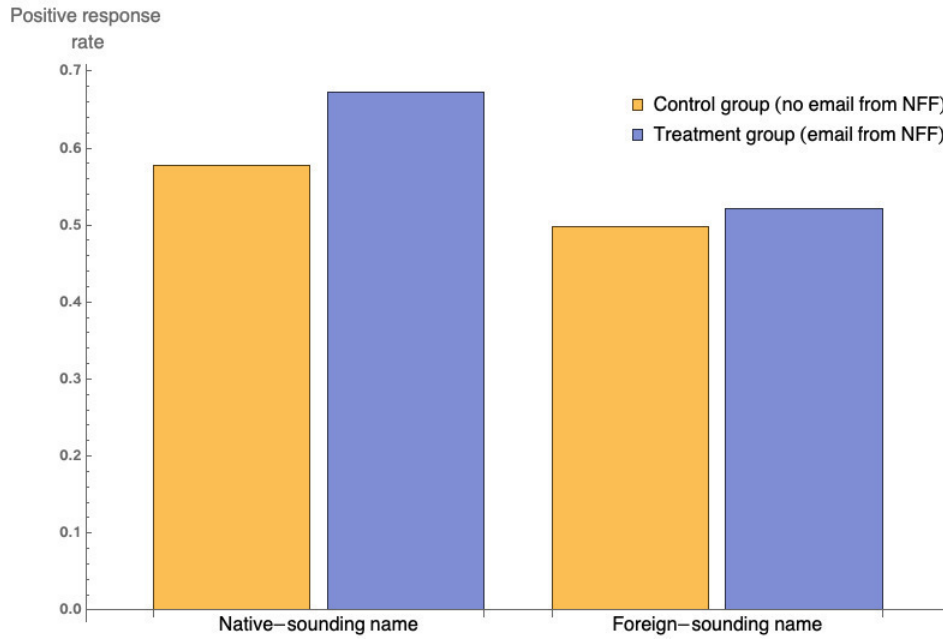
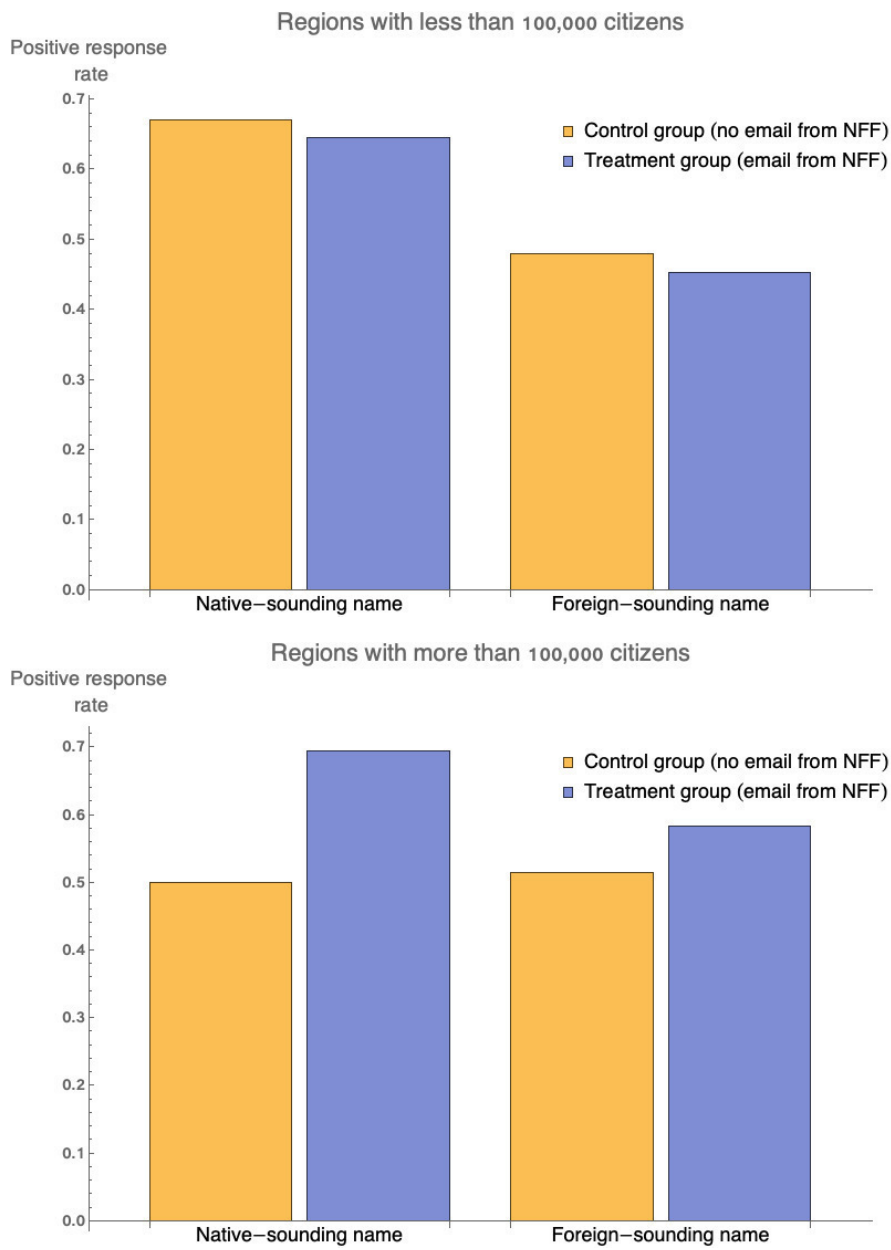


Figure 3

Share of positive responses to applications with native- and foreign-sounding names for treatment and control group coaches by population size of the region.



Appendix

Appendix A: The original emails in Norwegian

Original version in Norwegian of the email sent by the Norwegian Football Federation:

Subject: *Fotball for alle – for inkludering og mot rasisme*

Kjære trener,

Fotball er verdens mest populære sport og det gir oss ansvar og mulighet for å forene mennesker fra en rekke bakgrunner. Det er viktig for oss alle at fotball er flerkulturell og mangfoldig og gjenspeiler hele samfunnet vårt - ikke bare for profesjonelle, men for alle spillere. På amatørnivå letter fotball integrasjon og fremmer interaksjon på og utenfor banen. Vi har som mål å gjøre fotball lett tilgjengelig for alle medlemmer i vårt samfunn.

Rasisme og utestenging er et samfunnsproblem og dermed også et fotballproblem.

Vitenskapelige studier har vist at det er vanskeligere for folk med utenlandske navn å bli med i en amatør fotballklubb. Uavhengig av sport eller språkkunnskap, er det mindre sannsynlig at spillere med utenlandske navn får svar når de kontakter en klubb for første gang.

Denne barrieren for å delta i sport er ikke bare negativ for utlendinger, men også for nordmenn med migrasjonsbakgrunn og det norske samfunnet. Å skape mangfoldige team med medlemmer som har ulik bakgrunn er nøkkelen for å forbedre sosial integrasjon.

Vi har sendt denne eposten til alle trenere i Norge for å heie på den jobben dere gjør og oppfordre alle til å fortsette å bidra til holde døra inn til fotballen åpen for alle interesserte.

2021 blir et snodig år, men vi håper at i alle fall høsten ble tilnærmet normal for de fleste og at vi sammen kan se fram mot et strålende 2022.

[Name of NFFs representative]

[Logo of NFF]

Original version in Norwegian of the email of fictitious applicants:

Subject: *Treningsmulighet*

Hallo,

Jeg ser etter en ny fotballklubb. Tror du jeg kan komme og bli med på en treningsøkt?

Takk!

[Name]

Appendix B: Additional Tables

Table S1

Survey results.

Group	Female names	Sounds native or foreign (in %)	Male names	Sounds native or foreign (in %)
Norwegian-sounding	Marit Andersen	100	Andreas Andersen	94
	Kari Larsen	98	Kristian Larsen	100
	Inger Olsen	96	Kristoffer Olsen	100
	Anne Johansen	100	Joakim Johansen	100
	Ingrid Hansen	100	Martin Hansen	100
Polish-sounding	Julia Kamiński	100	Mateusz Kamiński	100
	Lena Wiśniewski	100	Kacper Wiśniewski	100
	Maja Wójcik	100	Michał Wójcik	100
	Zofia Kowalczyk	100	Dawid Kowalczyk	100
	Zuzanna Nowak	96	Jakub Nowak	100
Lithuanian-sounding	Liepa Vasiliauskas	100	Lukas Vasiliauskas	100
	Emilija Petrauskas	100	Arthur Petrauskas	98
	Austėja Jankauskas	100	Jonas Jankauskas	96
	Viltė Stankevičius	100	Kajus Stankevičius	100
	Gabija Kazlauskas	100	Nojus Kazlauskas	94
Somalian-sounding	Halima Aden	90	Ahmed Aden	96
	Fatima Haji	100	Farah Haji	100
	Sadia Abdullahi	100	Abdi Abdullahi	100
	Khadija Bashir	100	Yusuf Bashir	100
	Sumaya Ali	94	Abdullah Ali	92

Table S2

Regression results for a different definition of the dependent variable.

	Dependent variable: positive responses without further inquiries= 1, otherwise = 0		
	Model 1	Model 2	Model 3
Native-sounding name	Omitted	omitted	omitted
Foreign-sounding name	-0.10** (0.04)	-0.10** (0.04)	-0.11*** (0.04)
Email from NFF × Native-sounding name	0.03 (0.04)	0.02 (0.04)	0.01 (0.04)
Email from NFF × Foreign-sounding name	0.01 (0.04)	0.01 (0.04)	0.03 (0.04)
Male soccer team		-0.06 (0.04)	-0.15* (0.08)
Application sent out Monday or Tuesday (Monday=1)		-0.06* (0.03)	-0.05* (0.03)
League and region control			Yes
Constant	0.34*** (0.03)	0.43*** (0.05)	0.35* (0.20)
Observations	967	967	967
R ²	0.013	0.020	0.055

Notes: Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Figure S1

Geographic distribution of clubs by treatment assignment for regions with more or less than 100,000 citizens.

