Working Paper Series **689**
(ISSN 1211-3298)

Learning about Ethnic Discrimination from Different Information Sources

Darya Korlyakova

CERGE-EI
Prague, March 2021

ISBN 978-80-7343-496-0 (Univerzita Karlova, Centrum pro ekonomický výzkum a doktorské studium)
ISBN 978-80-7344-585-0 (Národohospodářský ústav AV ČR, v. v. i.)

Learning about Ethnic Discrimination from Different Information Sources*

Darya Korlyakova †

March 3, 2021

Abstract

We experimentally study whether public beliefs about ethnic discrimination, an emotionally loaded issue, are shifted more by information from experts or from ordinary people. We also examine whether people are inclined to choose the most influential sources. For this purpose, we combine, in a novel design, the random provision of information from different sources with endogenous information acquisition from the same sources. We find that individuals update their beliefs most in response to information from experts, namely researchers studying ethnic minorities and human resource managers. Exogenous adjustments in beliefs do not induce changes in attitudes to ethnic minorities. Consistent with the strength of belief updating, more individuals choose information from experts over information from ordinary people. This result suggests that, in the aggregate, people behave rationally as they favor a source that is perceived to be relatively accurate. The findings have implications for information dissemination policies.

Key words: ethnic discrimination, beliefs, information sources, experts

JEL Classification: C90, D83, J71

*This study was supported by Charles University in Prague, GAUK project No. 132218, ERC-CZ/AV-B research project No. 300851901 and by the H2020-MSCA-RISE project GEMCLIME-2020 GA No. 681228.

†CERGE-EI, a joint workplace of Center for Economic Research and Graduate Education, Charles University and the Economics Institute of the Czech Academy of Sciences, Politickych veznu 7, P.O. Box 882, 111 21 Prague 1, Czech Republic. E-mail: darya.korlyakova@cerge-ei.cz. I am indebted to Michal Bauer, Julie Chytilová, Štěpán Jurajda, Andreas Menzel, and Filip Matějka for valuable advice and encouragement. I thank Mark Dean, Evan Friedman, Maria Kogelnik, Maximillian Mueller, Pietro Ortoleva, Christian Ochsner, Silvio Ravaioli, Rastislav Rehák, Jan Švejnar, and Sevgi Yuksel for helpful comments. I am also grateful to ME-DIAN for outstanding cooperation on data collection. Jaroslav Groero, Michal Hakala, and Veronika Vraná provided excellent research assistance. The study was pre-registered in the AEA RCT Registry; AEARCTR-0003500. The experimental design was approved by CERGE Ethical Committee.

1 Introduction

Individuals often have to choose among multiple information sources and they may apply different criteria to identify their preferred source. Consistent with the conventional economic view, people select sources whose information provides them with an accurate representation of reality. However, in the context of controversial social issues, truthful information about ingroup members' unethical behavior can induce an unpleasant emotional response and thus the truth may be avoided or distorted. This idea is in line with a growing body of behavioral research which argues that people have a taste for *like-minded* information sources (e.g. Martin and Yurukoglu, 2017; Gentzkow and Shapiro, 2010; Mullainathan and Shleifer, 2005). Competing theoretical predictions raise the following questions: What information sources are the most influential? Will people choose these sources to acquire information about an emotionally loaded issue?

This paper aims to identify the most influential source(s) by investigating the causal responses of beliefs to information about ethnic discrimination from three distinct sources. Specifically, we gauge willingness to learn from the most influential source by simultaneously studying information acquisition from the same sources, namely ordinary people, human resource managers, and researchers studying ethnic minorities. We chose ordinary people and experts because each group is likely to possess one positive characteristic - low social distance or high accuracy - but never both of them. Therefore, subjects will face a trade-off between listening to potentially uninformed individuals similar to them and learning from HR managers or researchers who have relevant expertise and knowledge but engage less frequently in social interactions with subjects.

More broadly, our choice of information sources helps to add to the ongoing discussion on whether the public distrust experts and science. On one side of the spectrum, studies find that people learn less effectively from more knowledgeable peers (Ambuehl, Bernheim, Ersoy, and Harris, 2018), take less advice from experts than from peers (Läpple and Barham, 2019), and exhibit a *modest* level of trust in experts (Johnston and Ballard, 2016). However, recent results of surveys with representative samples from the US (Pew Research Center, 2019, 2020) and the Czech Republic (Public Opinion Research Center, 2019a) as well as the analysis of EU and UK survey data by Dommett and Pearce (2019) do not reveal growing anti-expert sentiment. Furthermore, US and Czech survey data,

along with Ipsos global poll results (2019), illustrate a generally high level of confidence in the scientific profession. This evidence suggests that widely discussed negative perception of experts may be exaggerated, and hence further scientific investigation is warranted.

Studying differences in belief updating among individuals who *selected* different sources has a number of problems. We would not be able to reliably conclude that a source with the largest belief responses is the most influential, because it may simply be chosen most frequently by people with a high propensity to update. To rule out selection-driven interpretation of results, we conduct a pre-registered survey experiment with a large representative sample of Czech individuals. We consider the Czech Republic to be particularly suitable for our information intervention because Czech society is presumably unaware of *widespread* local ethnic discrimination. According to Special Eurobarometer 437 (2015), 44 percent of Czech people believe that ethnic discrimination is a rare phenomenon locally, even though research and survey evidence documents prevalent unfavorable attitudes to Roma people, Asians, and other ethnic and national minorities in the Czech Republic (e.g. Bartoš, Bauer, Chytilová, and Matejka, 2016; Public Opinion Research Center, 2019b; Bartoš, Bauer, Cahlíková, and Chytilová, 2020). East Asians, whose discrimination our information intervention targets, are one of the two largest ethnic minority groups in the Czech Republic.

In the experiment, we randomly assign subjects to five groups. After eliciting prior beliefs about discrimination, we inform three groups of participants about how many applications a job seeker with an Asian-sounding name has to send to receive one interview invitation. This information is provided from a specific source; either it is a group of 9 ordinary people, 9 HR managers, or 9 researchers who mainly study issues that ethnic minorities face in the Czech Republic.

It would be problematic to isolate the role of an information source if we simultaneously vary the information content. To mitigate this confound, we oversample experts and ordinary people in a supplementary survey and randomly select from these groups the subsets of 9 individuals who, on average, agree on the extent of discrimination. As a result, across three treatments, experimental participants see the *same* number of applications that an Asian has to send, regardless of the source that provides the information. In the Control group, subjects receive neutral information unrelated to Asians. To inves-

tigate whether participants select sources that induce the strongest shifts of beliefs, we allow the last group to choose a preferred source and subsequently acquire information from it. After the information stage, we collect data on posterior beliefs, self-reported attitudes to Asians, and donations to a pro-Vietnamese[1] charity. In a follow-up study one week later, we examine the persistence of treatment effects on these three outcomes.

A complementary goal of this paper is to present a theory-based interpretation of the experimental findings. Traditional theory would predict that signal precision underlies people's decisions related to information acquisition and belief updating. To test whether individuals update more in response to more accurate sources and choose these sources more frequently, we elicit subjects' perceptions regarding the accuracy of three information sources. The second alternative, Akerlof's theory of social distance (1997), highlights the dependence of individual utility on the actions and opinions of others. In accordance with this theory, subjects are expected to learn most effectively from their social network members. To verify this hypothesis, we compare patterns in belief updating and information choices to differences in a self-reported likelihood of befriending an ordinary person, researcher, and HR manager. Finally, we examine a link between our results and the predictions of the confirmation-bias theory (Rabin and Schrag, 1999). In our setting, confirmation bias would manifest itself as a tendency to select a source whose opinion subjects expect to be the closest to their prior belief[2]. We measure subjects' expectation about a source opinion by asking them to predict the estimate of ethnic discrimination for each of the three sources.

We find that individuals' beliefs respond to information from every source, but the strongest response is to the message from experts, i.e. researchers and HR managers. These differences in belief responses persist over a one-week period, although they become smaller in magnitude. Despite revising their beliefs about discrimination, subjects do not change their attitudes to Asians and donate at similar rates to a local pro-Vietnamese

[1] Vietnamese people constitute the largest Asian minority in the Czech Republic.

[2] In the existing literature, there is no unique commonly used test aimed at identifying the presence of the confirmation bias. Charness and Dave (2017) specify a regression that allows them to estimate an extra weight that one places on confirming information. Charness, Oprea, and Yuksel (2020) define confirmatory-seeking types as those who, under different conditions, choose an information structure biased in the direction of their prior beliefs. Jones and Sugden (2001) study several forms of the confirmation bias. First, they explore in an incentive-compatible design whether subjects tend to choose an uninformative but potentially confirming card. Second, they examine whether subjects become more confident in the truth of the hypothesis if they receive a statement confirming it.

charity. Consistent with the strength of belief updating, more individuals prefer to acquire information from experts over information from ordinary people. The auxiliary analysis suggests that revealed patterns in belief updating and preferences for sources can be explained by the relatively high (perceived) accuracy of experts. At the aggregate level, our results accord with economic theory highlighting rationality in information acquisition and belief updating. The counterbalancing effects of social distance do exist, but they are not strong enough to reverse the gap between information sources.

Existing literature studies exogenous information provision separately from endogenous information acquisition, which does not inform policymakers about whether providing the most influential information is effective, measured in terms of people's willingness to acquire it. From an ex-ante perspective, it is conceivable that individuals will listen most to those whose opinions they perceive to be relatively accurate. However, they may *choose* to learn primarily from like-minded people, thereby leaving the dissemination of information from experts unattended. By investigating, in the same setting, the choices of information sources and the effects of exogenous information from these sources on beliefs, our study uncovers a general consistency between information preferences and causal updating.

Random assignment of information relates our paper to numerous studies that explore information treatment effects on individuals' beliefs, policy preferences (Alesina, Stantcheva, and Teso, 2018; Lergetporer, Schwerdt, Werner, West, and Woessmann, 2018; Kuziemko, Norton, Saez, and Stantcheva, 2015; Cruces, Perez-Truglia, and Tetaz, 2013) and behavior (Haaland and Roth, 2020; Alesina, Miano, and Stantcheva, 2018; Grigorieff, Roth and Ubfal, 2020). The closest work to ours is the study by Haaland and Roth (2019), who investigate whether exposure to research evidence about the extent of racial discrimination in the US labor market reduces political polarization in preferences for pro-black policies. In contrast to the above studies, we additionally create exogenous variation in a source that disseminates information. Several studies have already randomly manipulated an information source (Coibion, Gorodnichenko, and Kumar, 2018; Jacobsen, 2019; Cavallo, Cruces, and Perez-Truglia, 2017; Alt, Lassen, and Marshall, 2016). However, we differ from these by fixing the information content across treatments to cleanly identify the effects of a source alone and by varying a source expertise rather than its independence.

Simultaneously, our findings inform the literature on information acquisition (Bartoš et al., 2016; Hoffman, 2016; Hoopes, Reck, and Slemrod, 2015; Charness, Oprea, and Yuksel, 2020). Similar to our work, Fuster, Perez-Truglia, Wiederholt, and Zafar (2018) give survey respondents an opportunity to choose among several pieces of information that differ in their informativeness. However, the authors primarily examine what information about *home price changes* subjects acquire and whether they incorporate favorite information into their forecasts, which is very different from the focus of our study. In addition, the authors concentrate on the ex-ante predictive power of a source as the criterion that respondents use while ranking information sources. In contrast, we differentiate between perceived source accuracy, social distance and the potential to confirm one's initial belief.

The rest of the paper is organized as follows. Section 2 describes the experimental design, while Section 3 discusses our sample and follow-up attrition. Section 4 presents the experimental results, and Section 5 links our findings to theories. Section 6 concludes.

2 Experimental Design

Figure 1 outlines the main elements of our two-wave experiment, which randomly assigns participants to five different groups. In the first wave (main experiment), we elicited beliefs about the prevalence of local labor market discrimination against Asians. Subsequently, three subsets of subjects received exogenous information about discrimination from three different sources. The Control group was exposed to neutral irrelevant information, while the last subset of subjects was offered to choose from a list of three information sources and no-information option. Next, we measured subjects' posterior beliefs, attitudes to Asians, and donation behavior. In the second wave (obfuscated follow-up), which took place a week after the main experiment[3], we tested whether potential treatment effects on the participants' beliefs and attitudes persist. In both waves, we asked several questions on potential explanations for (i) belief responsiveness to information from specific sources and (ii) information choices.

[3]The minimum (maximum) number of days between the two waves was 5.43 (19.17), while the same number for an average participant was 6.95 days.

2.1 First Wave

2.1.1 Prior beliefs about the prevalence of discrimination against Asians

To put subsequent belief elicitation into the context, we first provided all subjects with a brief description of the labor market experiment from the correspondence study by Bartoš, Bauer, Chytilová and Matějka (2016). In particular, participants were told that researchers from CERGE-EI studied the extent of discrimination against Asians in the Czech labor market by sending job applications that were identical except for a job seeker's name (i.e. a Czech- vs. Asian-sounding name) to signal ethnicity.

Next, we informed all subjects that, according to the researchers, a job seeker with a Czech-sounding name has to send on average 7.5 applications in order to receive one interview invitation. Subsequently, we elicited participants' beliefs about the prevalence of local discrimination against Asians[4] by asking them to estimate the number of applications a job seeker with an Asian-sounding name has to send to receive one interview invitation. We incentivized correct answers (i.e. estimates that were the same as the CERGE-EI researchers' finding) with a 22-cent bonus[5] paid in addition to the participation fee. Research evidence indicates that small bonuses could serve as effective incentives (e.g. DellaVigna and Pope, 2017; Bullock, Gerber, Hill and Huber, 2015). In addition, large rewards for accuracy could have motivated our subjects to search for Bartoš et al.'s findings (2016), which are easily accessible online, and state the number they found instead of a true belief. Furthermore, a small pilot prior to the main experiment indicated that individuals who were randomly provided with a 22-cent bonus incentive (in the follow-up) took the posterior elicitation task more seriously. Specifically, these subjects spent more time on a page with the posterior-belief question and were more likely to update their beliefs (see Table B.1).

2.1.2 Information provision or acquisition

After we measured subjects' confidence in their prior beliefs about the extent of local discrimination against Asians, participants moved to an information provision or infor-

[4] In both waves, we asked subjects to consider among Asians mainly Vietnamese, Chinese and Japanese people.

[5] The bonus was paid in Czech crowns. Before giving consent to participate, potential subjects were told that they would receive information about the bonus at the end of November 2018 (instead of learning about it immediately after their completion of the survey). We chose to postpone this information to avoid its influence on subjects' posterior beliefs that were repeatedly collected in the follow-up.

mation acquisition stage. Three subsets of participants (hereafter Exogenous-Info groups) obtained the *same* piece of information about discrimination, namely the estimated number of applications that a job seeker with an Asian-sounding name has to send to receive one interview invitation. However, the source whose estimate participants saw was different across three treatment arms. Subjects in Laymen-Info group obtained the average estimate of 9 ordinary people. Participants in Practitioners-Info group received the average estimate of 9 HR managers, while participants in Researchers-Info group received the average estimate of 9 researchers who primarily study issues that ethnic minorities face in the Czech Republic. In the context of our study, we consider HR managers and researchers to be experts because the former are likely to have practical experience relevant for estimating the prevalence of discrimination and the latter can back up their perception of discrimination by theoretical knowledge and research findings.

To avoid deceiving subjects by presenting *fictional* individuals' estimates, we truthfully surveyed passers-by (whom we refer to as "ordinary people"), HR managers, and researchers in June-July 2018, before running the experiment. To elicit beliefs of passers-by, research assistants approached people in Prague parks and squares and asked whether they were willing to take part in a brief survey[6]. If the passer-by agreed, research assistants briefly described the labor market experiment by Bartoš et al. (2016) and stated the number of applications a job seeker with a Czech sounding name has to send to receive one interview invitation. Next, they asked the passer-by to estimate the number of applications that a job seeker with an Asian-sounding name has to send to obtain an invitation for one interview[7].

HR managers and researchers were contacted via an email that explained that a survey is a part of a PhD student's dissertation, mentioned the topic and length of the survey, and included a link to the online questionnaire. After clicking the link (which differed between the two groups), HR managers and researchers saw the survey consent page[8] followed by the same text and questions that passers-by had obtained. We retrieved the

[6]Research assistants also asked whether a person is over 18 years old when it was not obvious.

[7]We also elicited the corresponding beliefs of passers-by, HR managers, and researchers about an applicant with a Roma-sounding name. However, later on we did not find a match, i.e. an average estimate that was the same across three groups, and thus we did not use information about a Roma job seeker in the experiment.

[8]We decided to collect email addresses (which were deleted within 24 hours) in an attempt to avoid multiple responses. In addition, this allowed us to send reminders without disturbing individuals who had already filled in the questionnaire.

emails of HR managers working in various parts of the Czech Republic by mainly using the Czech Labor office database. In the case of researchers, we contacted individuals from different Czech universities and research institutes (mostly from Charles University and Masaryk University) who hold a PhD degree and whose research is focused on inter-ethnic relations, integration of foreigners, migration, attitudes to ethnic minorities, etc.

Overall, we collected 53 responses from passers-by, 36 responses from HR managers, and 20 responses from researchers[9]. Subsequently, we randomly divided each group into the subsets of 9 people, searched for an estimate that was the same across three subsets (belonging to different groups) after being rounded off to the nearest integer, and found that 14 applications was the match. This number is different from the actual result of the research study by Bartoš et al. (2016), 20 applications. In the experiment, we truthfully informed participants in Exogenous-Info groups that the average estimate of a group of 9 individuals was 14 applications, but we did not mention that the group whose estimate they were given was random. If this information had been provided, it would likely have led to subjects' questions[10] and confusion, and eventually might have decreased their trust in our questionnaire. Some other studies (see, for instance, Falk and Zimmerman, 2017) withheld information from their experimental participants in a similar fashion.

We did not use the averages of the whole groups for the following reasons. First, the averages of all surveyed passers-by, HR specialists and researchers were equal to 20, 14, and 13 applications[11], respectively. Thus, if we had provided different numbers from different sources, we would not have been able to disentangle the source effect from differences in the message content on belief updating.

The supplementary survey results were communicated to subjects in Exogenous-Info groups in the following manner:

> *We asked 9 passers-by/HR managers/researchers who primarily study ... to estimate the number of applications a job seeker with an Asian-sounding name has to send to receive one interview invitation. The average estimate of 9*

[9]The response rate for HR managers and researchers was about 6.5 percent and 37 percent, respectively.

[10]For example, participants might have started wondering why we did not provide the estimates of remaining individuals we had surveyed.

[11]We mention the values after rounding off. After excluding two passers-by with extreme beliefs (which we define in Section 4 as estimates larger than 50 or smaller than 1), the average for this group drops to 16 applications.

passers-by/HR managers/researchers was 14 applications.

Participants assigned to Exogenous-Info treatments also saw a bar chart comparing their prior belief to 14 applications, i.e. the average estimate of the respective source. In Researchers-Info group, we additionally informed participants that 9 researchers whose estimate they saw are not related in any way to the authors of the earlier discussed study. This information is truthful, and it was conveyed to prevent subjects from listening to researchers only because the latter are assumed to be familiar with the study results.

The Control group saw a placebo message that, in order to account for anchoring (Tversky and Kahneman, 1974), contained the same number (14) that Exogenous-Info groups saw. Specifically, untreated participants were exposed to the following text:

We compared the prices of granulated sugar in 9 Czech regions in August 2018.
The average price of sugar in these 9 regions was 14 crowns/kg.

We chose to provide a neutral message (rather than no information) to the Control group to mitigate potential experimental demand effects in the subsequent (main-experiment) collection of posterior beliefs. Participants could think that it is better to state a different estimate of discrimination if they were asked to reconsider their belief immediately after it had been elicited. In addition to the above text, we showed to untreated subjects a graph that compares August 2018 prices of sugar across 9 Czech regions. The average price across regions (14) was highlighted to draw subjects' attention to the number of interest.

The last subset of participants (Info-Choice group) was presented with a list of three groups of 9 individuals to receive discrimination-related information from, and an alternative of no information. Subjects in this group were told that they have a chance to obtain the average estimate of one source and they were asked to rank the options according to their preference[12]. Subsequently, the most preferred alternative was implemented.

[12]We did not mention that information provided by a source could be useful for a subject's subsequent estimation of discrimination. Other studies (e.g. Haaland and Roth, 2019; Chopra, Haaland and Roth, 2019) also do not emphasize the instrumental value of information that participants could acquire. It is probable that participants anticipated that we would ask them again to estimate the extent of discrimination against Asians. Subjects knew that they had completed a small part of the survey and thus they were likely to expect further related questions. Even if subjects were not sophisticated enough to anticipate the repeated estimation of discrimination, they could value information beyond the survey context. Consistent with this idea, Fuster et al. (2018) find that individuals are willing to pay for their favorite information much more than they will gain if their subsequent forecast of house prices is perfectly accurate.

If a participant chose to obtain a source estimate, he or she saw the following message (depending on the first-ranked source):

> *The average estimate of 9 passers-by/9 HR managers/9 researchers who primarily study ... was 14 applications.*

Similar to Exogenous-Info groups, subjects who selected one of the three sources additionally saw a graphical comparison of their prior belief and the average estimate of the respective source. If a participant preferred to see no information over all other options, the information-provision stage was omitted for him or her.

2.1.3 Collection of posterior beliefs, attitudes, and donations

Next, we asked all subjects whether they would like to revise their initial estimate of the number of applications a job seeker with an Asian-sounding name has to send to receive one interview invitation. Answers that coincided with the CERGE-EI researchers' finding were again incentivized with a 22-cent bonus.

We hypothesized that, if people are broadly unaware of existing discrimination against Asians by employers, exogenously shifting their beliefs may lead individuals to feel sympathetic toward Asians. Hence, we asked subjects (in a randomized order) whether they agree or disagree that Asian workers (i) take Czech people's jobs and (ii) produce more disadvantages than advantages for the Czech labor market. To test whether revised beliefs shift *overall* attitudes to the minority of interest, we included a question on how comfortable or uncomfortable a person would feel if his/her neighbor was Asian.

Since self-reported outcomes are subject to social desirability bias, we additionally employed a measure that makes concealing true attitudes to Asians costly. At the very end of the first wave, we offered participants three alternatives in relation to their reward: (i) sending the reward to their personal bank account; (ii) donating the reward to a specific pro-Vietnamese charity[13]; (iii) declining the reward. A decision regarding one's reward is a common part of MEDIAN's[14] surveys, but we modified the donation option by including solely a pro-Vietnamese charity.

[13]We did not find a charity that helps various Asian minorities in the Czech Republic; therefore, we selected a non-profit organization that supports the integration of Vietnamese people into Czech society.

[14]MEDIAN is the company we cooperated with on data collection (see section 3).

2.1.4 Predicted estimates of three information sources

Before facing the choice of an information source, a random half of Info-Choice group was asked to predict for each of the three groups (ordinary people, HR managers, and researchers) how many applications, according to a group, a job seeker with an Asian-sounding name has to send to receive one interview invitation[15]. The inclusion of this task enables us to assess the presence of *confirmation bias* which could imply that individuals tend to select information sources whose predicted beliefs are close to the individuals' prior beliefs.

2.2 Second Wave

2.2.1 Hiding the connection between two waves

Following Haaland and Roth (2019, 2020), we performed an obfuscated follow-up survey. Specifically, to mitigate potential experimental demand effects, we did not tell subjects that the two waves were connected. Moreover, the topics of both surveys were somewhat different[16] as was the text that participants saw before agreeing to participate in the surveys. In addition, at the beginning of the follow-up, we asked subjects several demographic questions to present the follow-up as an independent survey. Next, subjects answered a series of questions that concerned the Czech labor market but were not related to ethnic discrimination in order to further obfuscate the connection with the first wave.

2.2.2 Posterior beliefs, attitudes, donations and willingness to share information

After the obfuscation part, subjects received a series of attitudinal questions that were mostly reformulated[17]. We additionally measured participants' attitudes toward Ukrainians, a large national minority in the Czech Republic, to make the relationship between

[15]These predictions were not incentivized. Several patterns indicate that subjects took this belief elicitation task seriously. We find that an average (median) person spent about 74 (56) seconds on a page with the respective questions. Moreover, no subject wrote down the same number for three groups, and no more than 8 percent of subjects stated the same prediction for two groups. Finally, no more than 5 percent of individuals recorded an "extreme" belief for a source, i.e. higher than 50 applications that an Asian has to send for one interview invitation. For comparison, about 4 percent of subjects in the Info-Choice group stated an "extreme" prior belief which was elicited using an incentivized procedure.

[16]While the topic of the main experiment was "Attitudes toward social issues", the topic of the follow-up was "Economic and social issues".

[17]We used the same formulation of the neighbor-related question in both waves as MEDIAN uses it in other surveys and thus connection to the first wave should not be evident.

two waves less evident. Simultaneously, we aimed to explore whether there are spillovers from shifted attitudes to Asians (as a result of the information treatment) to the perception of other local minorities.

Later, subjects saw the text about Bartoš et al.'s experiment (2016) from the first wave and faced familiar incentivized elicitation of beliefs about the number of applications a job seeker with an Asian-sounding name has to send to receive one interview invitation. Participants were not reminded of their prior beliefs collected in the first wave.

At the end of the second wave, subjects were again offered to donate their experimental earnings to a pro-Vietnamese charity. This time, we used a different non-profit organization and, in both waves, we explicitly informed subjects that the charity is not related to the client who ordered the survey.

In addition to repeatedly collecting the above outcomes, we intended to access the value that individuals attach to information they had received the week before. To this end, we measured in Exogenous-Info groups the willingness to share with one's friends the source estimate of the number of applications a job seeker with an Asian-sounding name has to send to receive one interview invitation[18].

2.2.3 Perceived accuracy of information sources and social distance from them

While deciding how much weight to attach to a signal relative to the prior belief and what signal to choose, subjects may be guided by the signal precision. To test this hypothesis, we asked (for each of the three sources) how accurate, in one's opinion, an average source estimate of ethnic discrimination would be[19]. Alternatively, individuals may want to form the same opinion on the prevalence of local ethnic discrimination as their social network has. In this regard, we asked how likely it is that a subject would become friends with an ordinary person, HR manager, and a researcher who primarily studies issues that ethnic minorities face in the Czech Republic[20]. The order of the information sources within each

[18]This question came *after* the posterior belief elicitation. The type of information source that was mentioned depended on one's initial treatment assignment.

[19]We exposed a random half of follow-up participants to the questions about accuracy and social distance to prevent some participants from thinking too thoroughly about source characteristics before their posterior beliefs were collected.

[20]Both accuracy and social distance were measured using a 5-point Likert scale.

series of questions was randomly determined.

3 Sample Characteristics and Follow-up Attrition

We recruited 3216 subjects in cooperation with MEDIAN, which is a leading survey agency in the Czech Republic. Of these 3216 participants, 2233 subjects completed the follow-up survey, which took place, on average, a week after the main experiment. Subjects randomly assigned to the Info-Choice group were not invited to the follow-up due to financial constraints. Taking this into account, the response rate in the follow-up made up about 87 percent.

Table B.2 presents summary statistics for our sample. In Table B.3, we compare the demographics of our subjects (that we pre-specified to target) to the corresponding characteristics of the Czech population (mostly using data from the Czech Statistical Office). In both the main experiment and follow-up, our sample is fairly representative of the Czech population in terms of gender, age, education, and geography. Tables B.4 and B.5 illustrate that most covariates are balanced across five (four) treatment arms in the main experiment (follow-up survey).

As Table B.6 shows, overall attrition is unrelated to the majority of observables, although it is not entirely random[21]. Crucially, subjects with different prior beliefs about the extent of discrimination against Asians are equally likely to attrit from the follow-up [22]. A key concern, however, is that subjects in Practitioners-Info group are 4 percentage points less likely to complete the follow-up compared to untreated participants ($p<0.05$). Despite this, attrition does not differ significantly among three treatment arms with different information sources ($p>0.10$ in all cases) which is important for our comparison of belief updating across Exogenous-Info groups. In addition, we do not observe more covariate imbalances in the follow-up relative to the main experiment (Tables B.4 and B.5). Furthermore, while focusing on individuals who would appreciate particularly a practitioner's advice on an important issue (N = 1146), we find that subjects from the Practitioners-

[21] While several covariates predict a likelihood of participation in the follow-up (or its completion), this does not significantly affect the sample representativeness in the second wave.

[22] We ran two alternative regression specifications in which, instead of the underestimator dummy, we included (i) the continuous measure of pre-treatment beliefs or (ii) the dummy for above-median initial misperceptions regarding discrimination. Both predictors are insignificant ($p=0.85$ and $p=0.54$, respectively).

Info group are 6.3 percentage points *less* likely to participate in the follow-up relative to untreated participants ($p < 0.05$) [23]. The somewhat higher attrition rate for this subsample would, if anything, attenuate the effects of information from HR managers on subjects' beliefs collected in the follow-up. Finally, only 9.8 percent of non-participants (in the follow-up) quit the survey *after* opening our questionnaire, including those who quit before opening the page with the first (obfuscation) question, and this percentage does not differ significantly across treatment arms ($p>0.10$). Note that individuals who did not respond to the invitation to participate in a new survey were unlikely to know that it was related to the first questionnaire (as the topics of both surveys differed). While looking at the potential reasons for non-participation, we observe that untreated subjects were *more* likely to receive a reminder to fill in a questionnaire compared to their treated counterparts ($p<0.05$ when controls are included). The difference is (insignificantly) higher in Practitioners-Info group relative to the two other Exogenous-Info groups, which could at least partially account for the higher attrition rate in this group[24].

4 Experimental Results

We present four sets of results. First, we discuss the distribution and correlates of prior beliefs about the extent of labor market discrimination against Asians. Second, we analyze how individuals' beliefs respond to the same information about discrimination from different sources. Third, we discuss whether causal shifts in beliefs are accompanied by changes in self-reported attitudes to Asians and donation behavior. Finally, we analyze which information sources individuals select to acquire information about discrimination against Asians.

4.1 Prior beliefs about the extent of discrimination against Asians

As Table B.2 indicates, an average subject from the main experiment believes that a person with an Asian-sounding name has to send 15.75 applications to receive one interview

[23]The other categories of potential advisors included: a person like me, academician, family member, and colleague. Although we asked subjects the respective question at the end of the first wave, after some of them were treated, the sample is balanced on the fraction of participants who would value practitioners' advice most.

[24]Regressions supporting the above findings about subjects who quit the follow-up survey or chose not to participate in it are available upon request.

invitation[25]. Hence, it is reasonable to expect that the source estimate (14 applications) will not significantly shock the mean prior belief. Nevertheless, the provided information may be surprising if there is variation in pre-existing misperceptions at the *individual* level, which the average belief is not informative about. Figure A.1 plots the cumulative distribution function of prior beliefs about discrimination against Asians. Compared to the finding of Bartoš et al. (2016), which indicates that a job seeker with an Asian-sounding name has to send on average 20 applications to receive one interview invitation, almost 70.5 percent of participants underestimate ethnic discrimination, while 10.2 percent are correct about it. This finding contrasts markedly with the result by Haaland and Roth (2019), who compare their American subjects' priors against Bertrand and Mullainathan's finding (2004) and find that only 35 percent of individuals underestimate racial discrimination in the U.S. labor market.

For our further analysis, it is important to compare people's beliefs to the source estimate because this is the number that our subjects subsequently see. From now on, we refer to subjects as ex-ante underestimators (overestimators) if their initial belief is lower (higher) than *14* applications. We find that 51.2 percent of individuals underestimate discrimination against Asians, while 46.8 percent of individuals overestimate it. A median underestimator and overestimator believe that an Asian applicant has to send 10 and 20 applications, respectively.

Our classification based on a prior belief does not set apart respondents who may believe in more favorable or equal treatment of Asians in the Czech labor market. It may be more difficult to persuade such individuals because they are not wrong about the extent of discrimination but about the existence of discrimination per se. We find that about 15.6 percent of participants believe that an Asian has to send *fewer* applications than a Czech person with the same qualifications, whereas 2.9 percent believe that Czech and Asian job seekers are treated equally[26]. Low-educated individuals are more likely to believe that employers treat Asians more positively or at least equally, with 22.7 percent stating a number below or equal to 7.5 applications. This fraction constitutes a 49.3 percent increase relative to the fraction of higher-educated individuals with similar beliefs.

[25]The respective number for a median subject is 13 applications.

[26]The number of applications (7.5) that a job seeker with a Czech-sounding name has to send on average for one interview invitation was taken from the study by Bartoš et al. (2016).

In addition to discussing the overall heterogeneity of prior beliefs, it is worth mentioning how prior beliefs differ depending on subjects' background characteristics. Figure A.2 illustrates that, compared to male subjects, females believe that a job seeker with an Asian sounding name has to an additional 1.5 applications (on average) in order to be invited for one interview ($p<0.01$). Participants from Prague and those with a university degree also believe in more discrimination against local Asians ($p<0.10$). Relative to their counterparts, employed subjects[27] and participants with above-median income believe that an Asian person has to send 1.20 and 1.25 *additional* applications on average for one interview invitation ($p<0.05$ and $p<0.01$, respectively). In contrast to Haaland and Roth (2019), we do not find differences in beliefs about discrimination based on political orientation, which may suggest lower political polarization in Czech people's views on ethnic/racial inequalities compared to Americans. Multiple research evidence (e.g. Draca and Schwarz, 2018; Boxell, Gentzkow, and Shapiro, 2020) indicates a more pronounced political divide in the US relative to other countries. Finally, we do not find that subjects who had a direct contact with Asians in the past perceive the extent of discrimination differently than those with no previous exposure ($p=0.45$).

Result 1: *An average person is unlikely to be surprised by provided information because his or her prior belief about the extent of discrimination against Asians is fairly close to the source estimate. However, the average belief masks a substantial amount of heterogeneity. Compared to the source estimate, about half of people underestimate to different degrees the extent of local labor market discrimination against Asians.*

4.2 Treatment effects on beliefs

We start by presenting graphical evidence indicative of existing information treatment effects and, more importantly, of differences in updating depending on a source whose estimate subjects see. Figure 2 shows the kernel densities of posterior beliefs collected in the main experiment. The modes of densities for Exogenous-Info treatments shift[28] markedly in the direction of the number that participants saw. The spikes around 14 applications and reduction in belief uncertainty are most noticeable in groups that were

[27] The comparison group is composed of the unemployed, retired, subjects on parental leave or engaged in housework, students and others.

[28] For comparison, Figure A.3 illustrates the densities of prior beliefs that look very similar across treatment arms due to the virtue of randomization.

given information from experts [29]. Figure A.4 illustrates persistent but smaller differences across information sources in posterior beliefs elicited one week later.

[Figure 2 here]

To establish the information source effects more formally, we present the results in a regression framework, where, in the even-numbered columns of Table 1, we include the pre-specified covariates[30]. Panel A shows that mean posterior beliefs are not in general statistically distinguishable between the control and treatment groups. As was anticipated, an average person whose prior belief closely resembled the source estimate was not greatly surprised by information from a source.

Splitting the sample into underestimators and overestimators[31] based on comparison with the *provided* information (Panel B of Table 1) reveals significant adjustments in treated participants' beliefs in expected directions[32]. The effect of exogenous information about the extent of ethnic discrimination exists in all treatments, which is in line with Haaland and Roth's findings (2019). Our novel result is that a source that provides information matters for the strength of the information effect. Specifically, in the main experiment, underestimators who receive information from ordinary people raise their estimates on average by 0.8 applications, which represents a 9 percent increase compared to the control mean ($p<0.05$). At the same time, underestimators who learn from HR managers (researchers) increase their beliefs by 2 (1.8) applications or equivalently by 22 (20) percent compared to the control mean ($p<0.01$). Both changes are significantly different from the Laymen-Info group increase. Overestimators who obtain information from ordinary people reduce their estimates on average by 2 applications or by 8.3 percent compared to the control mean ($p<0.05$). Overestimators from Practitioners- and Researchers-Info

[29] Differences *between* densities for Experts-Info groups are less apparent (but significant from Laymen-Info group) if we limit the analysis to individuals who participated in both parts of the experiment. The results are available upon request.

[30] Eventually, we found the inclusion of income dummies to be more reasonable instead of coding household income as the log of the interval chosen by a respondent (which was pre-specified). We also slightly deviated from the pre-analysis plan by recoding beliefs lower than 1 application. However, the fraction of subjects with such estimates never exceeds 1 percent, and we describe below in the robustness checks the exclusion of participants with seemingly unreasonable beliefs.

[31] This category includes both overestimators, i.e. subjects with prior belief above 14 applications, and subjects whose prior coincided with a source estimate. However, the latter subgroup makes up only about 2 percent of our sample, and the results remain virtually the same if we exclude these subjects.

[32] Heterogeneity analysis by a prior belief was a part of the pre-analysis plan. See Tables B.7-B.9 for the results of exploring differences in belief updating along other pre-specified dimensions, such as confidence in a prior belief, previous exposure to Asians, and political affiliation.

group decrease their beliefs by 3.5 and 2.7 applications or by 14.5 and 11 percent relative to the control mean ($p<0.01$). The former change is marginally significantly different from the Laymen-Info group decrease.

In the follow-up, we continue to observe highly significant information treatment effects on both underestimators and overestimators. Differences in belief responses to information from experts relative to information from ordinary people become less pronounced and fall short of statistical significance. Nevertheless, it is likely that in the case of overestimators, the results are influenced by the presence of outliers. In Table B.10, we exclude these subjects instead of top- and bottom-coding their beliefs. Focusing on this subsample (N=2136) reveals a 6.7 and 5.9 percentage point reduction in beliefs of overestimators who learn from HR managers and researchers *in addition to* a downward adjustment that subjects in the Laymen-Info group make ($p<0.05$ and $p<0.10$, respectively).

[Table 1 here]

Next, we investigate individual-level patterns in belief updating depending on a source that provides information. For this exercise, we exploit our experimental design that allows us to measure posterior beliefs both instantly and with a delay. Although the outcomes we analyze were not pre-specified, they are closely connected to the hypotheses stated in our pre-analysis plan. Furthermore, studying *which* individuals are affected more by experts' opinions on discrimination relative to ordinary people's views may provide useful insights for policymakers.

In this respect, Table 2 classifies subjects from the Exogenous-Info groups into four categories. We omit comparisons to the Control group from the subsequent analysis because our current focus is on differences between individual information sources. First, we consider participants who update their beliefs immediately and retain information, i.e. state the same posterior belief or partly move to their prior belief in the follow-up. The message from HR managers and researchers increases the share of such subjects by 7 and 10 percentage points, respectively, compared to the group that saw the message from ordinary people ($p < 0.01$, Laymen-Info group mean = 10%). Second, we focus on the other extreme - participants who stick to their prior belief both in the main experiment and in the follow-up survey. Subjects from Laymen-Info group never update their beliefs in 21 percent of cases. The percentage of non-updaters decreases to 13% and 14%

in Practitioners- and Researchers-Info treatment arm, respectively ($p < 0.01$). The remaining categories include subjects (i) who update beliefs with a lag or (ii) who update beliefs initially but forget information with time, i.e. return to their preconceptions in the follow-up. We do not find systematic differences in the shares of such subjects across Exogenous-Info treatment arms[33]. Similarly, we do not find that subjects randomly assigned to different sources are more or less likely to be unclassified, e.g. to shift their beliefs away from the signal value both in the main experiment and follow-up[34].

[Table 2 here]

Result 2: *Individuals update their beliefs about ethnic discrimination more strongly when they obtain an expert's estimate relative to ordinary people's estimate. Larger responses of people's beliefs to information from HR managers and researchers persist over a one-week period.*

In the remainder of this subsection, we discuss the robustness of the above result. First, note that the effects of the information sources on updating outcomes remain virtually the same if we include the pre-specified covariates in the regressions (Tables 1 and 2). Furthermore, our findings are robust to excluding (i) subjects with extreme beliefs (i.e. those who state prior and/or posterior beliefs below 1 or above 50), (ii) subjects who report searching for an answer after learning about the CERGE-EI researchers' study[35], and (iii) the top and bottom 2 percent of the survey time distribution (see Appendix Tables B.10, B.11 and B.12). The results are qualitatively similar if we solely focus on attentive subjects (Table B.13), but we do not use this specification as the preferred one for the following reasons[36]. First, Table B.14 shows that passing the attention check is

[33] An exception is a *lower* fraction of subjects with delayed updating in the Researchers-Info group relative to the Laymen-Info group ($p<0.10$).

[34] The percentages of immediate updaters and non-updaters in the Control group are 4% and 27%, respectively. About 43% (2%) of untreated participants shifted their beliefs with a lag (updated initially but returned to their priors later). One fourth of the Control Group are unclassified, which is significantly higher than in any of the Exogenous-Info treatments. Untreated respondents may have changed their beliefs for several reasons, e.g. eventually understood the elicitation task better or forgot their prior belief recorded one week earlier. Fuster et al. (2018) discuss further possibilities of belief revision when no signal is sent. Note, however, that the information treatment effects we observe are *above* those that may arise from pure forgetting, anchoring or lack of initial comprehension. This could be seen from the Laymen-Info group means presented above.

[35] Only 4.39 percent of follow-up participants indicate that they were looking for the study results.

[36] 29 (27.9) percent of our main-experiment (follow-up) sample did not pass the attention check. The corresponding percentages documented by other researchers who use diverse national samples are below 30 percent, e.g. 19.9 percent and 27.7 percent in Berinsky, Margolis, and Sances (2014), and 22.4 percent in Haaland and Roth (2020).

correlated with a number of observables. Thus, excluding inattentive respondents would make our sample less representative. More importantly, selecting subjects on the basis of a post-treatment attention check could affect the balance across groups[37]. In addition to the above robustness checks, we run probit regressions (where appropriate) instead of OLS and find no qualitative differences (the results are available upon request).

4.3 Treatment effects on attitudes, donations and willingness to share information

Table 3 shows the effects of information from different sources on self-reported attitudes to Asians collected in the main experiment (Columns 1-4) and in the follow-up survey (Columns 5-8). All outcomes are standardized and recoded such that higher values imply better attitudes to Asians. Overall, the attitudes do not respond to information about discrimination regardless of the source that provides it[38]. This result is consistent with some other studies that find no or ambiguous effects of information on self-reported measures (Barrera, Guriev, Henry, and Zhuravskaya, 2020; Hopkins, Sides, and Citrin, 2019; Haaland and Roth, 2019; Lergetporer, Piopiunik, and Simon, 2018; Barnes, Feller, Haselswerdt, and Porter, 2018; Kuziemko et al., 2015).

[Table 3 here]

A plausible explanation for no treatment effects on mean attitudes could be that ex-ante underestimators and overestimators shift attitudes to Asians in opposite directions consistent with their changes in beliefs. However, we do not observe counteracting information effects when the sample is split by a prior belief (Panel B of Table 3). Another reason for insignificant effects could be that shifts in attitudes require *additionally* correcting people's misperceptions about the level of support that Asians receive from the Czech government, an Asian-specific criminality rate, the share of Asians who could speak Czech, etc. Alternatively, it is possible that people's attitudes to minorities are formed

[37] We chose to administer it after most of our main outcomes were collected to keep participants continuously focused. Subjects were informed at the very beginning of the main experiment that their attention would be checked at some point of the survey.

[38] Although subjects who receive information from ordinary people become more likely to disagree that Asians bring more disadvantages than advantages to the local labor market ($p<0.01$), we do not focus on this finding for two reasons. First, other attitudinal questions do not support this conclusion, and second, the positive effect drops in magnitude and becomes insignificant in the case of the reformulated question in the follow-up survey (Column 6 of Table 3).

by emotional experiences and are largely independent of beliefs about minorities' characteristics and their treatment.

Columns 1 and 2 of Table 4 show the treatment effects on charity donations (in the main experiment and follow-up survey, respectively). We do not observe that information from any source significantly affects the share of subjects (22.5% in the main experiment and 20% in the follow-up) who decide to donate their experimental earnings to a pro-Vietnamese charity[39]. Column 3 of Table 4 illustrates whether individuals' willingness to share information with friends that they obtained in the first wave differs across the three Exogenous-Information groups. Coefficients on Experts-Info treatment indicators are positive but small in magnitude and insignificant.

[Table 4 here]

Result 3: *Self-reported attitudes to Asians and donations of one's own experimental earnings to a pro-Vietnamese charity are generally unresponsive to information about the extent of local labor market discrimination against Asians. Willingness to share information with friends is not significantly affected by a source that provides this information.*

4.4 Information choices

Figure 3 illustrates the distribution of information options that were ranked first by experimental participants[40]. About 38 (32) percent of subjects favored the HR managers'

[39] In contrast to previous studies where participants were asked to (partially) donate a windfall income (Roth and Haaland, 2019; Alesina, Miano, and Stantcheva, 2018; Grigorieff, Roth, and Ubfal, 2020), we asked subjects to donate their *earnings* from the experiment. Although some research (e.g. Bekkers, 2007) has found that most individuals are reluctant to fully sacrifice their earnings, we still decided in favor of this donation measure due to its higher external validity.

[40] Recall that a random half of the Info-Choice group had to predict the beliefs of three sources about the extent of local discrimination against Asians *before* they selected a preferred piece of information. Figure A.5 illustrates that this additional belief elicitation task does not significantly affect the distribution of subjects' information choices relative to the subgroup for whom this task was omitted (p=0.6). Hence, we use data from the whole Info-Choice group in our analysis. During the experiment, some participants experienced difficulties with ranking information choices that was presented as a drag-and-drop task. Therefore, we had to change the question format, which seems to influence significantly the distribution of information choices (Figure A.5, p<0.01). However, the differences between distributions may have resulted from a programming error that was fixed *after* the info-choice question was modified. To check this intuition, we examine data from the Control group in which no programming error occurred. A random half of this group was asked to rank the same four information options but was not given an opportunity to see any information. Restricting attention to this subsample, we do not find significant differences between distributions of the information options before and after the respective question change (A.5, p=0.13).

(researchers') average estimate of the number of applications an Asian job seeker has to send to receive one interview invitation. Almost 23 percent of participants ranked the corresponding estimate of ordinary people as the top alternative. The frequencies of choosing three information sources significantly differ at the 1% level from frequencies that would be observed if subjects were randomizing uniformly across three options. Apparent preference for the experts' opinion suggests consistency, at least at the aggregate level, between people's information choices and their updating behavior when an information source is randomly manipulated. We present further evidence reinforcing this finding later in the subsection.

[Figure 3 here]

It is worth mentioning that only a *small* fraction of individuals (7 percent) do not want to see any estimate of discrimination against local Asians. This observation could be interesting given that information available for choice concerns an emotionally-charged issue[41]. Table B.15 presents the relationships between subjects' observable characteristics and choosing no information. Table 5 shows the determinants of preferences for information sources[42]. Similar to Fuster et al. (2018), we find that only a handful of observable characteristics predict information choices. Intuitively, less-educated subjects are more likely to favor ordinary people's estimate of discrimination against Asians over experts' estimate ($p<0.10$ and $p<0.05$ relative to participants with middle and high education, respectively). Specifically, these individuals decide to learn from ordinary people in 29 percent of cases, while higher-educated individuals choose similarly in 20 percent of cases. Subjects with low income are more likely to decide in favor of acquiring information from a source other than HR managers. Finally, focusing on subjects who prefer one type of experts over the other, we find that older respondents, subjects with right-wing orientation and those with more confidence in their prior belief are *less* likely to choose researchers ($p<0.10$)[43].

[Table 5 here]

[41] In paper by Fuster et al. (2018), 4.3 percent of subjects preferred no information about *home prices*.

[42] This table uses data from a random half of Info-Choice group to include a gap between one's prior and a predicted belief of a source. We return to this variable in Section 5. The regression was pre-specified.

[43] Higher age and right-wing orientation do not significantly predict the choice of researchers if we run a multinomial logit instead of linear probability models. Other correlations established in Table 5 seem to be robust.

Education was identified as the only observable that determines the strength of subjects' preference for information from experts. Table B.16 provides complementary evidence by further exploiting participants' ranking of information options. Conditional on choosing either type of experts (i.e. researchers or HR managers) as the preferred information option, low-educated subjects are 10 percentage points *more* likely to rank ordinary people – but not the other type of experts - as the second-best alternative (mean = 33%, $p < 0.05$). The coefficient of interest barely changes if multiple covariates are included.

The main result of this subsection is that people tend to choose information sources that will shift beliefs about discrimination most. Consistent with this finding, low-educated individuals who exhibit a stronger preference for information from ordinary people should also update more in response to a lay opinion on discrimination compared to higher-educated individuals. Table 6 splits the sample (excluding the Info-Choice group) by education. In the main experiment, low-educated subjects who see ordinary people's estimate move their beliefs on average by 1.69 applications, while higher-educated subjects shift their beliefs by 0.95 applications ($p < 0.05$ from a raw comparison of means and $p < 0.10$ from a model with the interaction term). This difference arises not because low-educated individuals are more easily swayed by others' opinions. In fact, updating[44] by subjects who see either experts' estimate does not differ significantly depending on their education.

Choosing ordinary people's estimate at a higher rate implies, in our context, that low-educated participants will differentiate less among the three sources relative to higher-educated participants. The follow-up patterns in updating (Columns 4-6 of Table 6) support this intuition. One week later, significant differences in belief shifts across information sources persist only for higher-educated participants. In contrast, low-educated individuals update similarly regardless of a randomly assigned information source.

[Table 6 here]

Result 4: *Individuals predominantly choose to consult experts, i.e. the more influential information source, while learning about the prevalence of local discrimination. Only*

[44]Even though this analysis is exploratory, it naturally follows our finding on the role of education in determining people's information preferences. Updating is defined as an absolute difference between a person's posterior belief and his/her prior estimate of discrimination. 93.7 (87.3) percent of main-experiment (follow-up) participants who shifted their beliefs in response to exogenous information from a source updated in a logical direction.

a few observable characteristics, particularly education, are correlated with information choices.

5 Connection to theories

We now discuss the agreement of our experimental results with theoretical predictions outlined below. The experiment was designed to allow differentiation between three theories that highlight the role of source accuracy, social distance, and confirmation bias, mostly with the help of descriptive and correlational evidence.

5.1 Perceived accuracy

One of our hypotheses was that source accuracy considerations underlie the strength of causal updating and information preferences. In light of our findings, this should imply that individuals perceive experts' opinion on discrimination to be more accurate relative to ordinary people's opinion. Figure A.6 compares subjects' judgements regarding each source accuracy. The average estimates of researchers and HR managers are perceived to be significantly more accurate than the average estimate of ordinary people ($p < 0.01$)[45,46].

Another way to explore the role of accuracy in explaining our results is to look at the reduction of uncertainty in subjects' beliefs about discrimination that should follow from seeing a source opinion. After eliciting prior beliefs in the main experiment and posterior beliefs in the follow-up, we asked subjects to state on a 5-point Likert scale how confident they are in their estimate. Although this task provided us with a less granular measure

[45]It should be noted here that additional randomization (which divided follow-up participants into those who answered questions on accuracy and social distance and those for whom these questions were omitted) was not successful. The omission was needed to ensure that *subsequent* posterior belief elicitation was not influenced by subjects' judgements in relation to source characteristics. Table B.17 shows more covariate imbalances across two large subsets relative to main randomization. Nevertheless, data on the perceived accuracy of three sources collected in the pilot (where *all* follow-up subjects answered the question of interest) confirm the finding about higher accuracy of experts (see Figure A.9).

[46]In the pre-analysis plan, we specified to test whether subjects who perceive an information source to be very accurate or accurate (very inaccurate or inaccurate) respond more strongly (weakly) to information from this source. However, we do not perform the pre-specified heterogeneity analysis because information treatment seems to affect the perception of the source accuracy in some cases (see Table B.18 for details). We refrained from collecting the data on accuracy and social distance in the main experiment to avoid priming subjects to think thoroughly about the information sources *before* they were treated. We planned to use information from the Control group only if the treatment contaminated perceptions of accuracy and social distance. However, in that case, we were able to investigate only general patterns (such as Figure A.6 presents) but not heterogeneity.

of belief uncertainty compared to asking subjects to assign percent chances to different alternatives (i.e. ranges of estimates), we preferred the Likert-scale question due to its simplicity and efficiency. Based on our data, we construct *several* measures of uncertainty reduction to validate the conclusions of this exploratory analysis.

In line with earlier evidence on experts' higher perceived accuracy, individuals who receive information from experts become (insignificantly) more sure about their estimate of local discrimination relative to those who receive information from ordinary people (Table 7, p=0.12)[47]. The effect is somewhat larger and marginally significant for the Researchers-Info group, in which subjects' posterior uncertainty decreases by almost 0.10 of a standard deviation (p=0.10). We also observe that a fraction of subjects who become more confident in their belief about discrimination grows by 4.4 percentage points after subjects see information from experts (Laymen-Info group mean = 29%, p = 0.07).

The final piece of evidence in favor of the accuracy explanation concerns posterior uncertainty among different educational groups. Recall that low-educated subjects do not differentiate much between experts' estimate of discrimination and ordinary people's estimate, which is noticeable in their belief updating and information choices. Conversely, higher-educated subjects seem to place higher importance on experts' opinion and thus this source should reduce their posterior uncertainty more strongly. We do not find that a *gap* between average confidence of subjects who receive ordinary people's estimate and those who receive experts' estimate is smaller if we restrict the analysis to the low-educated subsample (Table B.20). However, we do find that the fraction of low-educated subjects who become more confident than initially is almost the same across the Exogenous-Info groups (Laymen-Info group mean = 32%, p = 0.98). In contrast, the fraction of higher-educated subjects who report belief uncertainty reduction is larger by a significant 7.8 percentage points if they were exposed to information from experts one week earlier (Laymen-Info group mean = 27%, p = 0.01).

[Table 7 here]

Overall, perceived source accuracy seems to be a plausible explanation for our findings[48]. Individuals tend to consider experts who estimate the extent of ethnic discrimination on

[47] Table B.5 demonstrates that confidence in a prior belief is balanced across treatment arms.

[48] Appendix C discusses why a non-trivial portion of individuals may have chosen information from ordinary people, the source with the lowest *average* accuracy.

the Czech labor market to be more accurate relative to ordinary people. The complementary result is that posterior uncertainty decreases more if individuals receive experts' estimate. This decrease is driven by the higher-educated subsample whose beliefs respond more to a message from experts and who choose this source more frequently.

5.2 Social distance

Our second hypothesis was that people's beliefs will be moved most by a source with the lowest social distance, which will also be their modal information choice. In light of our findings, this should imply that individuals perceive experts to be socially closer relative to ordinary people. Figure A.7 compares self-reported likelihood of friendship with an ordinary person to likelihood of friendship with an HR manager or researcher who primarily studies issues that ethnic minorities face in the Czech Republic. The graphical comparison clearly indicates that subjects consider themselves to be much more socially *distant* from experts than from ordinary people ($p < 0.01$ from the raw comparison of means).

Next, we examine the relevance of social distance theory for explaining our findings using a regression in which we interact a social distance dummy with each treatment indicator[49]. We exclude the Practitioners-Info group from the interpretation of results because this treatment arm seems especially affected by imperfect *additional* randomization (see footnote 45) implemented in the follow-up. Hence, patterns in belief updating of subjects who were assigned to this group and did not see the questions on accuracy and social distance do not replicate findings discussed earlier. Table B.21 illustrates that subjects who report above-median social distance from researchers are less likely to shift their beliefs in the main experiment in response to information from this source (p=0.06). A coefficient on the interaction term is similar in magnitude but more noisy in the follow-up (p=0.28). At the same time, social distance does not seem to significantly influence the updating behavior of subjects in the Laymen-Info group.

[49] Table B.19 shows that information treatment does not affect social distance from any source. We pre-specified the heterogeneity analysis by social distance. We planned to use a posterior belief as an outcome, but eventually we used updating to increase statistical power by *not* dividing the sample into over- and underestimators. We also planned to use three levels of social distance, which leads to almost no variation in a dummy variable that is equal to 1 if a subject is unlikely or very unlikely to befriend an ordinary person. Hence, we split the sample by *median* social distance.

The regression results suggest that social distance, if anything, *reduces* the gap between belief responses to information from researchers compared to information from ordinary people. At the same time, social distance theory does not provide insight into the existence of the gap per se. Our primary finding does not confirm this theory prediction: information from experts, a *more* socially distant group relative to ordinary people, causes stronger belief responses and is chosen more frequently.

5.3 Confirmation bias

Finally, we hypothesized that people will acquire information from a source that is most likely confirm their original perception of local discrimination. In the light of our findings, this should imply that individuals expect both experts' estimates of discrimination to be closer to their prior belief compared to ordinary people's estimate[50]. Figure A.8 depicts prior beliefs about the number of applications a job seeker with an Asian-sounding name has to send to receive one interview invitation along with the corresponding *predicted* beliefs of the three information sources[51]. According to subjects' predictions, experts disagree on the extent of local labor market discrimination against Asians. In particular, an average subject expects HR managers to state 14.06 applications, while the corresponding number for researchers is 15.98 applications ($p<0.01$). Furthermore, an average participant believes that his or her prior belief is the closest to *researchers'* respective belief ($p=0.74$) and is the furthest from *HR managers'* belief ($p<0.01$)[52].

When further examining the presence of confirmation bias in the regression framework, we do not find that lower (absolute) distance between a subject's prior belief and a source predicted belief is associated with stronger preference for the source (see Table 5). The results are similar if we additionally include a gap between a subject's prior belief and predicted belief of the *alternative* source which may also play a role when a person decides between two sources (see Table B.22 that additionally excludes individuals with extreme beliefs). In an alternative specification (Table B.23), we explore the role of a *relative* gap

[50] Confirmation bias may also manifest itself in a decision to acquire no information to avoid challenging one's prior belief. Recall that a *handful of* our experimental participants go for the no-information option.

[51] The comparison is based on the data from a random half of the Info-Choice group. We also asked untreated participants in the follow-up (N = 570) to predict the beliefs about discrimination against Asians for each of three sources. The comparison of prior beliefs and predicted source estimates looks very similar for this group, and can be provided upon request.

[52] The difference between the mean prior belief and the mean predicted belief of ordinary people, 14.57 applications, is significant at the 5 percent level.

between one's prior belief and predicted belief of a source. We find that subjects whose prior belief is closer to the researchers' predicted estimate of discrimination relative to the predicted estimate of another source are 10 percentage points more likely to favor information from researchers. However, the respective coefficients fall short of statistical significance (p=0.15 and 0.13 when ordinary people and HR managers, respectively, serve as a comparison group)[53].

In general, confirmation bias does not appear to be a leading explanation for our findings. We find, at best, weak evidence indicating the presence of this bias in the choices of participants who decided to learn from researchers. In addition, this theory does not explain why subjects frequently select the HR managers' estimate of discrimination if they do *not* expect HR managers' beliefs to be consonant with their prior.

6 Conclusion

Recent discussions on anti-intellectualism provoked by Michael Gove's famous quote[54] have given rise to a commonly held belief that "the death of expertise" is real. However, there seems to be no clear evidence from economic experiments, at least with representative samples, that the public is receptive to information from ordinary people rather than from intellectual elites. In this paper, we cast doubt on whether experts are no longer influential. Our design distinguishes itself from previous research by capitalizing on synergy between exogenous information provision from distinct sources and endogenous information acquisition. This enables us to address an important policy relevant question: Will individuals become *voluntarily* exposed to information from a source whose influence is the strongest when the message is unavoidable? Answering this question is important to prevent unnoticed information dissemination and thus wasteful governmental spending targeted at raising societal awareness on sensitive topics.

Our findings give reasons for both optimism and pessimism. The good news is that most individuals choose experts as their preferred source of information about ethnic discrimination, a topic that often evokes strong emotions. Moreover, experts, whose information

[53] We find some significant correlations between the distance of researchers' predicted estimate and the choice of that information source when we use a multinomial logistic regression. There is no evidence, though, in favor of confirmation bias in the case of other sources. The results are available upon request.

[54] During his interview with Faisal Islam at Sky News (June 3, 2016), Michael Gove said in relation to Britain's exit from the European Union that "people in this country have had enough of experts".

individuals perceive to be relatively accurate, are more influential relative to ordinary people. Hence, there is consistency, at least at the aggregate level, between individuals' information choices and their causal belief updating. In this regard, our findings support the standard theory predictions. At the same time, information from experts, as well as information from ordinary people, does not affect individuals' self-reported attitudes to ethnic minorities and their donations to an ethnic minority charity. Null effects of information on self-reported measures corroborate some earlier research results (e.g. Haaland and Roth, 2019; Kuziemko et al., 2015).

This paper opens several potentially interesting avenues for future research. We study responses to expert and non-expert information and its acquisition in the context of ethnic discrimination. It is an open question as to how our results would extend to settings in which individuals are exposed to other sensitive issues, such as climate change, gun control or vaccination. Another possible extension is to focus messages on an ethnic minority whom society perceives *very* negatively. We have shown that, when there are merely unfavorable attitudes to a minority, individuals are generally willing to correct their initial misperceptions regarding the treatment that this minority faces. Clearly, this finding does not guarantee that information will not cause "backfiring effects", i.e. an increase in original misperceptions, in the case of a minority that is a common target of societal hatred and violence. Finally, it would be useful to vary the nature of provided information. It is plausible that the public learns from experts when they describe facts but not when they give advice, due to the dislike of being preached to.

Our findings have important implications for information dissemination policies. If the goal of an information campaign is solely to raise awareness about the prevalence of local ethnic discrimination, sharing the opinion of practitioners or academic experts, who are perceived to be socially distant but relatively accurate, could contribute to higher effectiveness of the campaign. If a campaign aims instead to improve attitudes to ethnic minorities, correcting people's beliefs about the extent of discrimination against these minorities could be insufficient or even irrelevant.

References

Akerlof, G. A. (1997). Social distance and social decisions. *Econometrica: Journal of the Econometric Society*, 1005-1027.

Alesina, A., Miano, A., & Stantcheva, S. (2018). Immigration and redistribution. NBER Working Paper No. w24733.

Alesina, A., Stantcheva, S., & Teso, E. (2018). Intergenerational mobility and preferences for redistribution. *American Economic Review*, 108(2), 521-54.

Alt, J. E., Marshall, J., & Lassen, D. D. (2016). Credible Sources and Sophisticated Voters: When Does New Information Induce Economic Voting?. *The Journal of Politics*, 78(2), 327-342.

Ambuehl, Bernheim, Ersoy, & Harris (2018). Peer Advice on Financial Decisions: A case of the blind leading the blind? NBER Working Paper No. w25034.

Barnes, L., Feller, A., Haselswerdt, J., & Porter, E. (2018). Information, knowledge, and attitudes: An evaluation of the taxpayer receipt. *The Journal of Politics*, 80(2), 701-706.

Barrera, O., Guriev, S., Henry, E., & Zhuravskaya, E. (2020). Facts, alternative facts, and fact checking in times of post-truth politics. *Journal of Public Economics*, 182, 104-123.

Bartoš, V., Bauer, M., Chytilová, J., & Matějka, F. (2016). Attention discrimination: Theory and field experiments with monitoring information acquisition. *American Economic Review*, 106(6), 1437-75.

Bartoš, V., Bauer, M., Cahlíková, J., & Chytilová, J. (2020). Covid-19 Crisis Fuels Hostility against Foreigners. Available at SSRN: https://ssrn.com/abstract=3618833.

Bekkers, R. (2007, December). Measuring Altruistic Behavior in Surveys: The All-or-Nothing Dictator Game. In *Survey Research Methods* (Vol. 1, No. 3, pp. 139-144).

Berinsky, A. J., Margolis, M. F., & Sances, M. W. (2014). Separating the shirkers from the workers? Making sure respondents pay attention on self-administered surveys. *American Journal of Political Science*, 58(3), 739-753.

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.

Boxell, L., Gentzkow, M., & Shapiro, J. M. (2020). Cross-Country Trends in Affective Polarization. NBER Working Paper No. 26669.

Bullock, J. G., Gerber, A. S., Hill, S. J., & Huber, G. A. (2015). Partisan Bias in Factual Beliefs about Politics. *Quarterly Journal of Political Science*, 10, 519-578.

Cavallo, A., Cruces, G., & Perez-Truglia, R. (2017). Inflation expectations, learning, and supermarket prices: Evidence from survey experiments. *American Economic Journal: Macroeconomics*, 9(3), 1-35.

Charness, G., & Dave, C. (2017). Confirmation bias with motivated beliefs. *Games and Economic Behavior*, 104, 1-23.

Charness, G., Oprea, R., & Yuksel, S. (2020). How do people choose between biased information sources? Evidence from a laboratory experiment. Working Paper.

Chopra, F., Haaland, I., & Roth, C. (2019). Do People Value More Informative News?. Available at SSRN: https://ssrn.com/abstract=3342595.

Coibion, O., & Gorodnichenko, Y. (2015). Is the Phillips curve alive and well after all? Inflation expectations and the missing disinflation. *American Economic Journal: Macroeconomics*, 7(1), 197-232.

Cruces, G., Perez-Truglia, R., & Tetaz, M. (2013). Biased perceptions of income distribution and preferences for redistribution: Evidence from a survey experiment. *Journal of Public Economics*, 98, 100-112.

DellaVigna, S., & Pope, D. (2017). What motivates effort? Evidence and expert forecasts. *The Review of Economic Studies*, 85(2), 1029-1069.

Dommett, K., & Pearce, W. (2019). What do we know about public attitudes towards experts? Reviewing survey data in the United Kingdom and European Union. *Public Understanding of Science*, 28(6), 669-678.

Draca, M., & Schwarz, C. (2018). How polarized are citizens? Measuring ideology from the ground-up. Measuring Ideology from the Ground-Up.Available at SSRN: https://ssrn.com/abstract=3154431.

Falk, A., & Zimmermann, F. (2017). Information processing and commitment. *The Economic Journal*, 128(613), 1983-2002.

Fuster, A., Perez-Truglia, R., Wiederholt, M., & Zafar, B. (2018). Expectations with Endogenous Information Acquisition: An Experimental Investigation. NBER Working Paper No. w24767.

Gentzkow, M., & Shapiro, J. M. (2010). What drives media slant? Evidence from US daily newspapers. *Econometrica*, 78(1), 35-71.

Grigorieff, A., Roth, C., & Ubfal, D. (2020). Does information change attitudes towards immigrants? *Demography*, 57(3), 1117-1143.

Haaland, I., & Roth, C. (2020). Labor market concerns and support for immigration. *Journal of Public Economics*, 191, 104256.

Haaland, I., & Roth, C. (2019). Beliefs about Racial Discrimination and Support for Pro-black Policies. Available at SSRN: https://ssrn.com/abstract=3039629.

Hoffman, M. (2016). How is information valued? Evidence from framed field experiments. *The Economic Journal*, 126(595), 1884-1911.

Hopkins, D. J., Sides, J., & Citrin, J. (2019). The muted consequences of correct information about immigration. *The Journal of Politics*, 81(1), 315-320.

Hoopes, J. L., Reck, D. H., & Slemrod, J. (2015). Taxpayer search for information: Implications for rational attention. *American Economic Journal: Economic Policy*, 7(3), 177-208.

Hopkins, D. J., Sides, J., & Citrin, J. (2019). The muted consequences of correct information about immigration. *The Journal of Politics*, 81(1), 315-320.

Ipsos (2019). It's a fact, scientists are the most trusted people in world.

Jacobsen, G. D. (2019). How do different sources of policy analysis affect policy preferences? Experimental evidence from the United States. *Policy Sciences*, 1-28.

Johnston, C. D., & Ballard, A. O. (2016). Economists and public opinion: Expert consensus and economic policy judgments. *The Journal of Politics*, 78(2), 443-456.

Jones, M., & Sugden, R. (2001). Positive confirmation bias in the acquisition of information. *Theory and Decision*, 50(1), 59-99.

Kuziemko, I., Norton, M. I., Saez, E., & Stantcheva, S. (2015). How elastic are preferences for redistribution? Evidence from randomized survey experiments. *American Economic Review*, 105(4), 1478-1508.

Läpple, D., & Barham, B. L. (2019). How do learning ability, advice from experts and peers shape decision making?. *Journal of Behavioral and Experimental Economics*, 80, 92-107.

Lergetporer, P., Piopiunik, M., & Simon, L. (2018). Does the Education Level of Refugees Affect Natives' Attitudes?. CESifo Working Paper No. 6832.

Lergetporer, P., Schwerdt, G., Werner, K., West, M. R., & Woessmann, L. (2018). How information affects support for education spending: Evidence from survey experiments in Germany and the United States. *Journal of Public Economics*, 167, 138-157.

Martin, G. J., & Yurukoglu, A. (2017). Bias in cable news: Persuasion and polarization. *American Economic Review*, 107(9), 2565-99.

Mullainathan, S., & Shleifer, A. (2005). The market for news. *American Economic Review*, 95(4), 1031-1053.

Pew Research Center (2019). Trust and Mistrust in Americans' Views of Scientific Experts.

Pew Research Center (2020). Trust in Medical Scientists Has Grown in U.S., but Mainly Among Democrats.

Prior, M., Sood, G., & Khanna, K. (2015). You cannot be serious: The impact of accuracy incentives on partisan bias in reports of economic perceptions. *Quarterly Journal of Political Science*, 10(4), 489-518.

Public Opinion Research Center (2019a). Prestiž povolání – červen 2019 [Occupational prestige - June 2019].

Public Opinion Research Center (2019b). Vztah české veřejnosti k národnostním skupinám žijícím v ČR – březen 2019 [Public attitudes toward national groups living in the Czech Republic - March 2019].

Rabin, M., & Schrag, J. L. (1999). First impressions matter: A model of confirmatory bias. *The Quarterly Journal of Economics*, 114(1), 37-82.

Special Eurobarometer 437 (2015). Discrimination in the EU in 2015.

Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, 185(4157), 1124-1131.

Main figures

Figure 1: Main elements of the experimental design

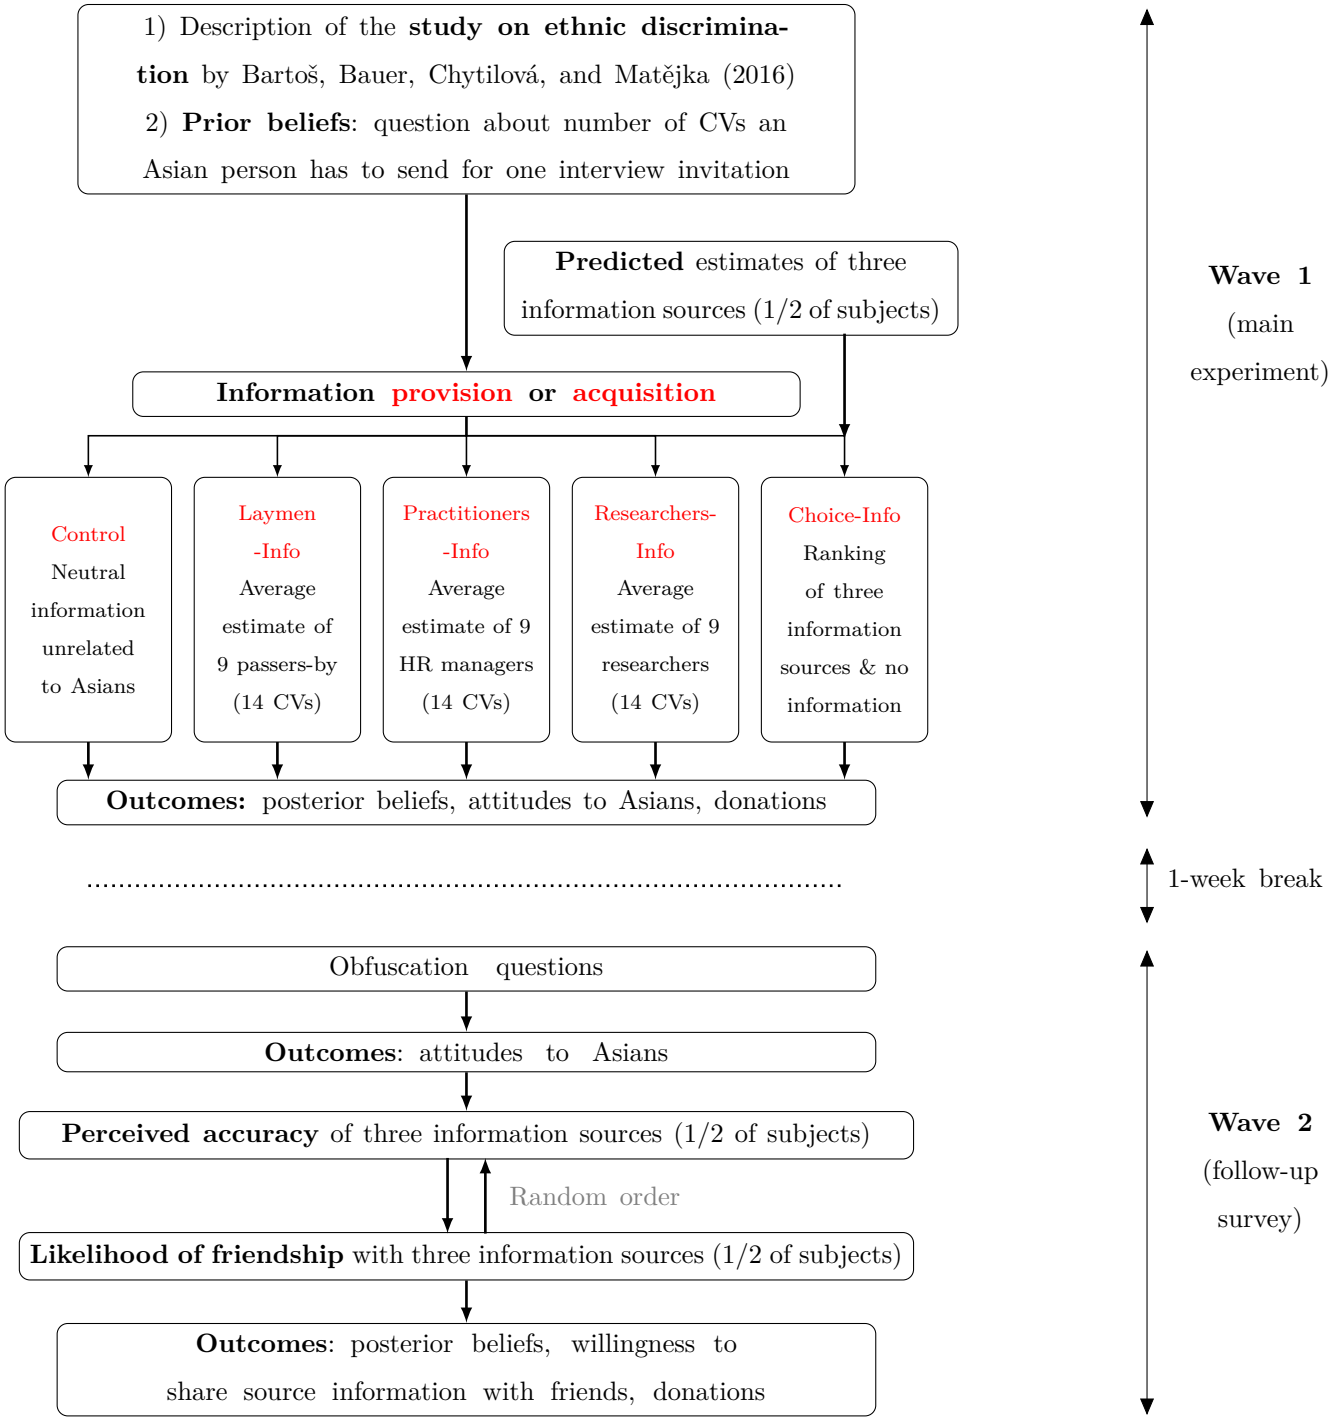

Figure 2: Posterior beliefs about discrimination against Asians (main experiment)

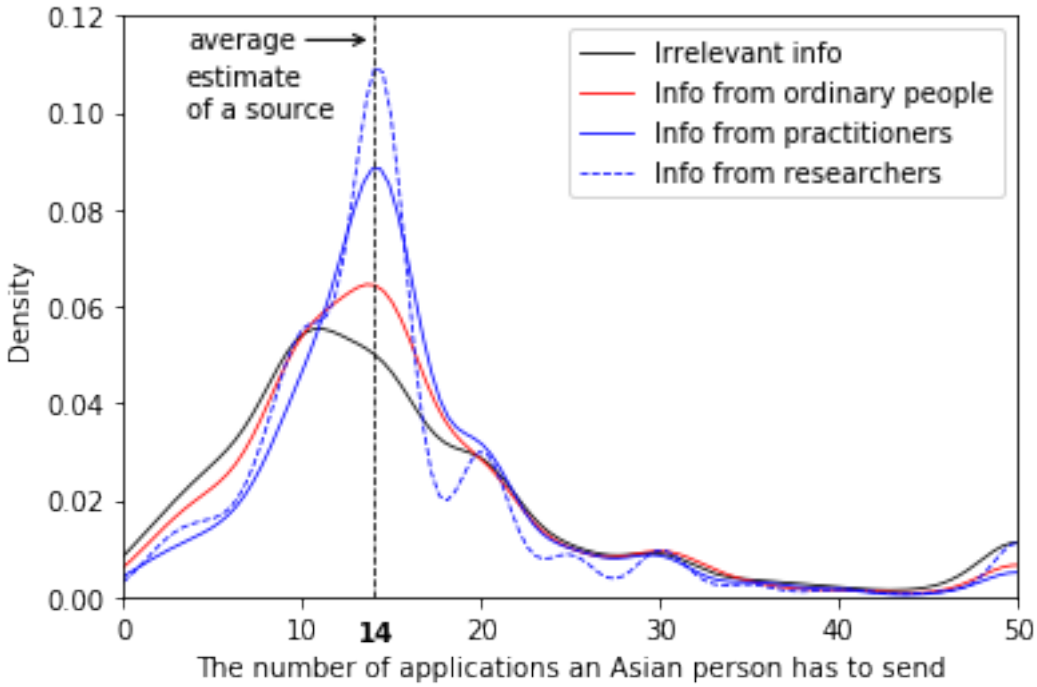

Notes: The figure plots the kernel densities of posterior beliefs using main-experiment data from the Control group and Exogenous-Info groups (N = 2571). The dotted vertical line indicates the number of applications estimated by a source. Extreme beliefs (i.e. estimates higher than 50 and lower than 1) were re-coded accordingly.

Figure 3: Chosen information options

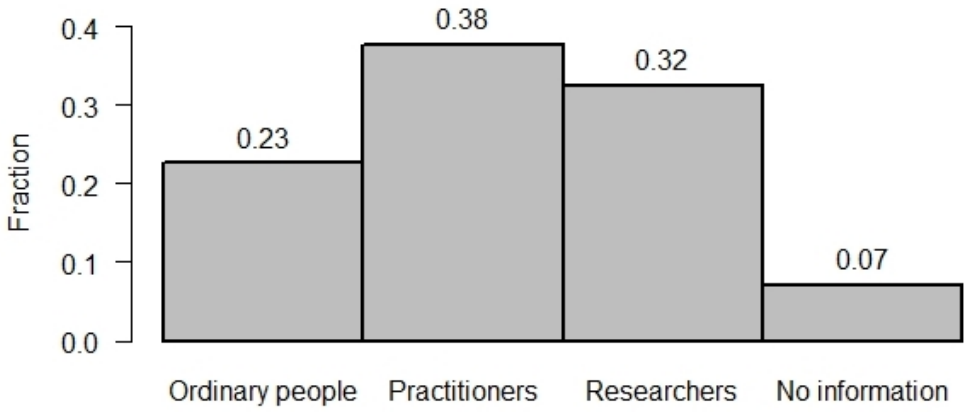

Notes: The figure shows the distribution of alternatives that participants in the Info-Choice group (N = 645) ranked as the most preferred. We often use interchangeably "practitioners" and "HR managers" in the figures.

Main tables

Table 1: Posterior beliefs about discrimination: main experiment and follow-up survey

	Posterior: main		Posterior: follow-up	
	(1)	(2)	(3)	(4)
Panel A: Main specification				
Laymen-Info	-0.68	-0.60	-0.45	-0.40
	(0.60)	(0.46)	(0.58)	(0.52)
Practitioners-Info	-0.51	-0.82*	-1.00*	-1.13**
	(0.57)	(0.46)	(0.56)	(0.52)
Researchers-Info	-0.71	-0.42	-0.77	-0.82
	(0.60)	(0.48)	(0.58)	(0.53)
Observations	2,571	2,571	2,233	2,233
Control mean	16.44	16.44	15.92	15.92
Covariates	No	Yes	No	Yes
Panel B: Prior heterogeneity				
Laymen-Info (a)	-2.08**	-2.03**	-3.04***	-3.15***
	(0.85)	(0.85)	(0.88)	(0.88)
Practitioners-Info (b)	-3.47***	-3.52***	-4.26***	-4.39***
	(0.81)	(0.80)	(0.88)	(0.88)
Researchers-Info (c)	-2.74***	-2.66***	-4.13***	-4.30***
	(0.89)	(0.89)	(0.89)	(0.88)
Underestimator	-15.08***	-14.91***	-10.91***	-10.84***
	(0.71)	(0.71)	(0.79)	(0.80)
Underestimator × Laymen-Info (d)	2.92***	2.83***	5.33***	5.42***
	(0.94)	(0.95)	(1.04)	(1.04)
Underestimator × Practitioners-Info (e)	5.44***	5.42***	6.35***	6.54***
	(0.90)	(0.91)	(1.01)	(1.02)
Underestimator × Researchers-Info (f)	4.51***	4.35***	6.76***	6.86***
	(0.97)	(0.97)	(1.05)	(1.05)
Observations	2,571	2,571	2,233	2,233
Intercept	24.11		21.43	
Covariates	No	Yes	No	Yes
p-value: $a = b$	0.06	0.05	0.10	0.10
p-value: $a = c$	0.42	0.45	0.15	0.13
p-value: $b = c$	0.34	0.27	0.86	0.91
p-value: $a + d = b + e$	0.00	0.01	0.70	0.83
p-value: $a + d = c + f$	0.02	0.02	0.56	0.62
p-value: $b + e = c + f$	0.62	0.59	0.31	0.46
p-value: $a + d = 0$	0.04	0.05	0.00	0.00
p-value: $b + e = 0$	0.00	0.00	0.00	0.00
p-value: $c + f = 0$	0.00	0.00	0.00	0.00

Notes: Underestimator equals to 1 if the value of an individual's prior belief is lower than 14. In columns (2) and (4), the following pre-specified covariates are included: gender, age, household size, regional, educational and income dummies, confidence in a prior belief, municipality size, employment status, exposure to Asians, and political orientation. Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table 2: Fractions of subjects with different updating patterns

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	Share of those who									
	updated immediately & retained information		updated beliefs with delay		updated immediately but forgot later		never updated beliefs		are unclassified	
Practitioners-Info	0.08***	0.07***	-0.02	-0.01	0.00	0.00	-0.08***	-0.08***	0.02	0.02
	(0.02)	(0.02)	(0.03)	(0.03)	(0.01)	(0.01)	(0.02)	(0.02)	(0.02)	(0.02)
Researchers-Info	0.10***	0.10***	-0.06**	-0.05*	0.00	0.00	-0.07***	-0.07***	0.04*	0.03
	(0.02)	(0.02)	(0.03)	(0.03)	(0.01)	(0.01)	(0.02)	(0.02)	(0.02)	(0.02)
Observations	1,663	1,663	1,663	1,663	1,663	1,663	1,663	1,663	1,663	1,663
Covariates	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes
Laymen-Info group mean	0.10	0.10	0.54	0.54	0.02	0.02	0.21	0.21	0.13	0.13

Notes: The regression uses data from three Exogenous-Info groups. Individuals who participated in both waves are included. We consider among those who *updated immediately & retained information* subjects who shifted their beliefs and (i) kept them at the same level over one-week period or (ii) moved to some extent back to their prior belief. Subjects who *updated their beliefs with delay* include those who initially kept their prior or did not update fully but moved their beliefs (more) in the follow-up. Subjects who *updated immediately but forgot later* returned to their prior belief over one-week period. Participants are considered *unclassified* if they update in a non-standard manner, for example move further from the signal in the main experiment but choose the opposite direction in the follow-up. Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table 3: Self-reported attitudes to Asians

	Main experiment				Follow-up			
	(1) Asians take jobs	(2) Asians bring disadvantages	(3) Asian neighbor	(4) Pro-Asian index	(5) harder to find job due to Asians	(6) Asians bring advantages	(7) Asian neighbor	(8) Pro-Asian index
Panel A: Main specification								
Laymen-Info	0.055	0.153***	0.021	0.077*	0.060	0.038	-0.029	0.023
	(0.055)	(0.055)	(0.052)	(0.042)	(0.056)	(0.057)	(0.057)	(0.041)
Practitioners-Info	-0.018	0.032	-0.031	-0.006	0.033	-0.048	-0.036	-0.017
	(0.056)	(0.056)	(0.054)	(0.044)	(0.058)	(0.058)	(0.058)	(0.042)
Researchers-Info	0.034	0.067	-0.051	0.017	0.025	-0.013	-0.025	-0.004
	(0.054)	(0.054)	(0.053)	(0.042)	(0.057)	(0.058)	(0.056)	(0.040)
Observations	2,571	2,571	2,571	2,571	2,233	2,233	2,233	2,233
Covariates	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Panel B: Prior heterogeneity								
Laymen-Info	0.098	0.144*	0.025	0.089	-0.005	0.021	-0.062	-0.015
	(0.080)	(0.080)	(0.073)	(0.060)	(0.075)	(0.080)	(0.081)	(0.055)
Practitioners-Info	-0.019	-0.008	-0.050	-0.026	0.042	-0.009	-0.095	-0.021
	(0.081)	(0.083)	(0.075)	(0.063)	(0.079)	(0.080)	(0.082)	(0.058)
Researchers-Info	0.122	0.058	-0.023	0.052	-0.069	-0.011	0.013	-0.023
	(0.078)	(0.083)	(0.076)	(0.061)	(0.078)	(0.082)	(0.081)	(0.057)
Underestimator	-0.023	-0.032	0.011	-0.015	-0.203**	0.082	-0.078	-0.066
	(0.076)	(0.076)	(0.075)	(0.058)	(0.081)	(0.082)	(0.080)	(0.057)
Underestimator × Laymen-Info	-0.084	0.017	-0.008	-0.025	0.127	0.031	0.063	0.074
	(0.110)	(0.110)	(0.105)	(0.085)	(0.112)	(0.114)	(0.114)	(0.081)
Underestimator × Practitioners-Info	0.005	0.079	0.038	0.041	-0.019	-0.079	0.118	0.007
	(0.111)	(0.112)	(0.106)	(0.087)	(0.116)	(0.116)	(0.115)	(0.084)
Underestimator × Researchers-Info	-0.168	0.018	-0.053	-0.068	0.185	-0.003	-0.072	0.037
	(0.108)	(0.109)	(0.106)	(0.084)	(0.113)	(0.115)	(0.112)	(0.080)
Observations	2,571	2,571	2,571	2,571	2,233	2,233	2,233	2,233
Covariates	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: OLS in all columns in both Panels. The outcomes mentioned in columns (1)-(3) and (5)-(7) were measured on a scale from 1: "Strongly agree" to 5: "Strongly disagree", and *Asians bring advantages* was re-coded so that higher values mean more positive attitudes to Asians. These outcomes are z-scored using respective means and standard deviations in the Control group. *Pro-Asian index* is an unweighted average of the outcomes mentioned in the previous three columns. Both indices and covariates included in all regressions were pre-specified. Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table 4: A decision to donate earnings and willingness to share information with friends

	Main experiment	Follow-up	
	(1) Donation	(2) Donation	(3) Share with friends
Panel A: Main specification			
Laymen-Info	-0.005	-0.007	
	(0.022)	(0.023)	
Practitioners-Info	-0.003	-0.006	0.027
	(0.022)	(0.023)	(0.061)
Researchers-Info	-0.034	-0.019	0.057
	(0.021)	(0.022)	(0.060)
Observations	2,571	2,233	1,663
Control mean	0.225	0.202	
Covariates	Yes	Yes	Yes
Panel B: Prior heterogeneity			
Laymen-Info	-0.002	-0.020	
	(0.032)	(0.031)	
Practitioners-Info	0.029	0.003	0.010
	(0.033)	(0.032)	(0.084)
Researchers-Info	-0.021	-0.035	0.022
	(0.032)	(0.030)	(0.086)
Underestimator	0.029	-0.024	-0.106
	(0.031)	(0.032)	(0.086)
Underestimator × Laymen-Info	-0.008	0.029	
	(0.044)	(0.049)	
Underestimator × Practitioners-Info	-0.057	-0.020	0.033
	(0.038)	(0.043)	(0.120)
Underestimator × Researchers-Info	-0.026	0.035	0.066
	(0.042)	(0.05)	(0.121)
Observations	2,571	2,233	1,663
Covariates	Yes	Yes	Yes

Notes: In columns (1) and (2) probit, marginal effects, standard errors in parentheses. The outcome is a binary variable indicating whether a respondent decided to donate his/her earnings from the experiment to a pro-Vietnamese charity. In column (3) OLS, robust standard errors in parentheses. The outcome was measured on a scale from 1: "Very willing" to 5: "Very unwilling" and re-coded so that higher values mean higher willingness to share information with friends. This outcome is z-scored using the mean and standard deviation in the Laymen-Info group. All regressions include pre-specified controls. *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

Table 5: Correlates of preferences for information sources

	Equals to 1 if chose ...			
	(1) ordinary people over experts	(2) HR managers over ordinary people	(3) researchers over ordinary people	(4) researchers over HR managers
Ordinary-prior belief gap	0.001			
	(0.003)			
HR-prior belief gap		0.002		
		(0.005)		
Research-prior belief gap			-0.001	-0.002
			(0.005)	(0.004)
Male	0.034	-0.048	-0.059	-0.034
	(0.054)	(0.076)	(0.081)	(0.072)
Age	-0.002	0.004*	0.000	-0.004*
	(0.002)	(0.002)	(0.003)	(0.002)
Middle education	-0.098*	0.114	0.121	0.023
	(0.059)	(0.081)	(0.085)	(0.077)
High education	-0.153**	0.201*	0.190*	0.058
	(0.073)	(0.104)	(0.106)	(0.098)
Prague	-0.005	-0.041	0.019	0.133
	(0.107)	(0.160)	(0.144)	(0.146)
2nd income quartile	-0.075	0.227*	-0.058	-0.265**
	(0.100)	(0.135)	(0.136)	(0.125)
3rd income quartile	-0.150*	0.305***	0.028	-0.314***
	(0.089)	(0.114)	(0.130)	(0.101)
4th income quartile	-0.110	0.290**	-0.069	-0.407***
	(0.097)	(0.128)	(0.137)	(0.119)
Income missing	0.075	0.026	-0.198	-0.249
	(0.130)	(0.161)	(0.169)	(0.165)
Employed	-0.012	0.005	0.033	0.023
	(0.056)	(0.077)	(0.087)	(0.072)
Right-wing oriented	0.036	0.021	-0.133	-0.139*
	(0.063)	(0.077)	(0.096)	(0.080)
Household size	-0.012	-0.006	0.049	0.049
	(0.024)	(0.034)	(0.037)	(0.031)
Above-median municipality size	0.051	-0.104	0.036	0.085
	(0.086)	(0.128)	(0.116)	(0.121)
Underestimates discrimination	0.037	-0.033	-0.043	-0.013
	(0.056)	(0.072)	(0.081)	(0.069)
Sure about a prior belief	-0.005	0.094	-0.195	-0.209*
	(0.095)	(0.105)	(0.181)	(0.107)
Unsure about a prior belief	0.009	-0.048	0.043	0.060
	(0.056)	(0.078)	(0.082)	(0.074)
Exposure to Asians	-0.011	0.058	-0.068	-0.081
	(0.052)	(0.070)	(0.086)	(0.069)
Mean	0.26	0.61	0.56	0.45
Observations	303	202	179	225

Note: This table uses data from a random half of the whole Info-choice group. *Source-prior* belief gap is defined as the absolute difference between a predicted belief of a source (top-coded at 50 and bottom-coded at 1) and the prior belief. Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table 6: Causal shifts of beliefs by education

	(1)	(2)	(3)	(4)	(5)	(6)
	Updating: main			**Updating: follow-up**		
	Low educated	Higher educated	All	Low educated	Higher educated	All
Practitioners-info	1.12***	1.24***	1.24***	0.90	1.28***	1.15**
	(0.43)	(0.29)	(0.29)	(0.60)	(0.50)	(0.50)
Researchers-Info	1.30***	0.93***	0.95***	0.84	1.32***	1.22**
	(0.41)	(0.26)	(0.27)	(0.62)	(0.51)	(0.51)
Low educated			0.61*			0.67
			(0.32)			(0.55)
Practitioners-Info × Low educ			-0.08			-0.28
			(0.53)			(0.78)
Researchers-Info × Low educ			0.38			-0.48
			(0.49)			(0.80)
Intercept	1.69***	0.95***	1.66*	5.11***	4.40***	2.33
	(0.29)	(0.13)	(0.95)	(0.45)	(0.29)	(1.62)
Observations	819	1,042	1,861	688	904	1,592
Covariates	No	No	Yes	No	No	Yes

Notes: The regression uses data from three Exogenous-Info groups. Individuals with extreme beliefs, i.e. outliers, are excluded. *Updating* is defined as the absolute difference between a person's posterior belief and his/her prior estimate of discrimination. In Columns 3 and 5 pre-specified covariates are included. Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table 7: Reduction of uncertainty in beliefs about the extent of discrimination

	Confidence in posterior		Share of more confident subjects than initially	
	(1)	(2)	(3)	(4)
Laymen-Info (a)	0.142**	0.142**	0.050*	0.050*
	(0.054)	(0.054)	(0.027)	(0.027)
Practitioners-Info (b)	0.200***		0.085***	
	(0.055)		(0.027)	
Researchers-Info (c)	0.239***		0.102***	
	(0.056)		(0.027)	
Experts-Info (d)		0.220***		0.094***
		(0.047)		(0.023)
Observations	2,233	2,233	2,233	2,233
Control mean			0.24	0.24
Covariates	Yes	Yes	Yes	Yes
p-value: $a = d$		0.12		0.07
p-value: $a = b$	0.32		0.21	
p-value: $a = c$	0.10		0.06	
p-value: $b = c$	0.51		0.53	

Notes: *Confidence in posterior* was measured on a scale from 1: "Very sure" to 5: "Very unsure", and was re-coded so that higher values mean higher confidence in a posterior belief (collected in the follow-up). This outcome is z-scored using the mean and standard deviation in the Control group. *Share of more confident subjects* is a binary variable indicating whether a respondent reported higher confidence in his/her own posterior belief about discrimination compared to his/her confidence in the respective prior belief. Controls included in all regressions were pre-specified. Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Abstrakt

Ve své práci pomocí experimentu zkoumáme velmi citlivou problematiku etnické diskriminace. Zkoumáme, zda se veřejné mínění ohledně diskriminace změní více, když subjektům poskytneme informace od expertů nebo od obyčejných lidí. Dále se zabýváme tím, zda si lidé vybírají spíše vlivnější zdroje. Námi vytyčené otázky zodpovídáme pomocí experimentálního designu, v němž nově kombinujeme dvě možnosti výběru informačních zdrojů, a to zcela náhodný výběr informačního zdroje a endogenní volbu informačního zdroje. Docházíme ke zjištění, že lidé nejvíce pozměňují svá apriorní mínění v reakci na informace od expertů, konkrétně od vědeckých pracovníků zabývajících se tématikou etnických minorit a HR manažerů. V souladu s tímto zjištěním se také ukazuje, že si lidé často jako zdroj svých informací vybírají spíše odborníky. Exogenní manipulace apriorního mínění nijak nepozměňuje postoje k etnickým menšinám. Obecně se dá tedy říci, že se lidé chovají racionálně, jelikož si vybírají informace z relativně přesnějších zdrojů. Výsledky této studie se dají využít při utváření politik informovanosti.

The appendix to this working paper is available at https://www.cerge-ei.cz/working-papers/.

Working Paper Series
ISSN 1211-3298
Registration No. (Ministry of Culture): E 19443

Individual researchers, as well as the on-line and printed versions of the CERGE-EI Working Papers (including their dissemination) were supported from institutional support RVO 67985998 from Economics Institute of the CAS, v. v. i.

Specific research support and/or other grants the researchers/publications benefited from are acknowledged at the beginning of the Paper.

Published by
Charles University, Center for Economic Research and Graduate Education (CERGE)
and
Economics Institute of the CAS, v. v. i. (EI)
CERGE-EI, Politických vězňů 7, 111 21 Prague 1, tel.: +420 224 005 153, Czech Republic.
Printed by CERGE-EI, Prague
Subscription: CERGE-EI homepage: http://www.cerge-ei.cz

Phone: + 420 224 005 153
Email: office@cerge-ei.cz
Web: http://www.cerge-ei.cz

Editor: Byeongju Jeong

The paper is available online at http://www.cerge-ei.cz/publications/working_papers/.

ISBN 978-80-7343-496-0 (Univerzita Karlova, Centrum pro ekonomický výzkum
a doktorské studium)
ISBN 978-80-7344-585-0 (Národohospodářský ústav AV ČR, v. v. i.)