

Attention Discrimination: Theory and Field Experiments with Monitoring Information Acquisition[†]

By VOJTĚCH BARTOŠ, MICHAL BAUER, JULIE CHYTILOVÁ, AND FILIP MATĚJKA*

We integrate tools to monitor information acquisition in field experiments on discrimination and examine whether gaps arise already when decision makers choose the effort level for reading an application. In both countries we study, negatively stereotyped minority names reduce employers' effort to inspect resumes. In contrast, minority names increase information acquisition in the rental housing market. Both results are consistent with a model of endogenous allocation of costly attention, which magnifies the role of prior beliefs and preferences beyond the one considered in standard models of discrimination. The findings have implications for magnitude of discrimination, returns to human capital and policy. (JEL C93, D83, J15, J16, J24, J71, R31)

Understanding why people discriminate based on ethnicity, gender, or other observable group attributes has been one of the central topics in economics and other social sciences for decades.¹ Since the seminal work of Phelps (1972) and Arrow (1973), it has been widely acknowledged that due to a lack of individual-level information decision makers often rely on a group attribute as a signal of unobserved

*Bartoš: CERGE-EI (a joint workplace of Charles University in Prague and the Economics Institute of the Czech Academy of Sciences), Politických vězňů 7, 111 21 Prague 1, Czech Republic (e-mail: vojtech.bartos@cerge-ei.cz); Bauer: CERGE-EI, Politických vězňů 7, 111 21 Prague 1, Czech Republic, and Institute of Economic Studies, Faculty of Social Sciences, Charles University in Prague (e-mail: bauer@cerge-ei.cz). Chytilová: Institute of Economic Studies, Faculty of Social Sciences, Charles University in Prague, Opletalova 26, Prague 1, 11000, Czech Republic, and CERGE-EI (e-mail: chytilova@fsv.cuni.cz); Matějka: CERGE-EI, Politických vězňů 7, 111 21 Prague 1, Czech Republic (e-mail: filip.matejka@cerge-ei.cz). The authors are grateful for funding from the CERGE-EI Foundation under a program of the Global Development Network (Bartoš), from the Czech Science Foundation (Bauer and Chytilová 13-20217S, Matějka 14-30724S) and from the People Programme (Marie Curie Actions) of the EU's Seventh Framework Programme under REA grant (Bauer and Chytilová 609642). We thank Colin Camerer, Stefano DellaVigna, Randy Filer, Martin Gregor, Christian Hellwig, Štěpán Jurajda, Peter Katusčák, Ulrike Malmendier, Marti Mestieri, Sendhil Mullainathan, Ron Oaxaca, Franck Portier, Chris Sims, Jakub Steiner, Matthias Sutter, five anonymous referees and audiences at Oxford University, NYU, Princeton University, Columbia University, CERGE-EI, Toulouse School of Economics, the University of Gothenburg, the Institute for Advanced Studies in Vienna, UC Berkeley, University of San Francisco, and ASSA meeting in Philadelphia for valuable comments, and Kateřina Boušková, Lydia Hähnel, Vít Hradil, Iva Pejsarová, Lenka Švejdová, and Viktor Zeisel for excellent research assistance. Institutional Review Board approval has not been obtained because the institutions, which the authors are affiliated with, do not have IRBs. The authors declare that they have no relevant or material financial interests that relate to the research described in this paper.

[†]Go to <http://dx.doi.org/10.1257/aer.20140571> to visit the article page for additional materials and author disclosure statement(s).

¹Researchers have produced a vast amount of evidence documenting discriminatory behavior based on ethnicity or gender on labor, housing, and consumer markets. Yinger (1998) and Altonji and Blank (1999) survey regression-based (nonexperimental) evidence, and Riach and Rich (2002) and List and Rasul (2011) summarize related field experiments.

individual characteristics. This may give rise to “statistical discrimination” in selection decisions on various markets.² At the same time, a large body of research in both economics and psychology documents that scarce attention plays an important role in decision making (e.g., Newell, Shaw, and Simon 1958; Kahneman 1973; Gabaix et al. 2006; Chetty, Looney, and Kroft 2009; Fehr and Rangel 2011) and theories assuming costly attention made progress in explaining a range of important economic phenomena (e.g., Sims 2003; Mackowiak and Wiederholt 2009).

While the existing models of discrimination implicitly assume that individuals are fully attentive to available information, we link the literatures on discrimination and scarce attention. We develop a model in which we describe how knowledge of a group attribute impacts the level of attention to information about an individual and how the resulting asymmetry in acquired information across groups—denoted “attention discrimination”—can lead to discrimination in a selection decision. To test the model, we build on the experimental design of Bertrand and Mullainathan (2004) and perform three correspondence tests in two countries. A novel feature of our field experiments are the tools to measure the process of decision making, in addition to selection choices, by monitoring acquisition of information about applicants.

Attention to available information about candidates is crucial input in virtually any selection process: in the recruitment of employees, school admissions, housing market tenant selection, loan provisions, voting in elections, or scientific review processes, to name a few examples. The *Economist* (2012), for instance, describes the process as follows:

They [human resource staff] look at a CV for ten seconds and then decide whether or not to continue reading. If they do, they read for another 20 seconds, before deciding again whether to press on, until there is either enough interest to justify an interview or to toss you into the ‘no’ pile.³

Similarly, qualitative studies of college admissions describe the reading of applications by admission committees as very coarse and quick (Stevens 2009; Deresiewicz 2014). The pioneering field experiment on discrimination in the labor market by Bertrand and Mullainathan (2004) finds that returns to sending higher-quality resumes, in terms of callbacks, are higher for applicants with a white-sounding name compared to applicants with an African American-sounding name in the US labor market. The pattern is consistent with employers not continuing to read once they see an African American name on a resume, thus resulting in greater discrimination among more qualified applicants. These provocative findings motivate the need to find a way to measure the effect of a name on reading effort. For a theory, the findings open the question as to whether choices about inspecting applicants are

²Taste-based discrimination is the second prominent explanation for why people discriminate (Becker 1971). It arises due to preferences, not due to lack of information.

³Also, a recent study found that human resource managers spend on average six seconds reviewing an individual resume (TheLadders 2012). Another study (Dechief and Oreopoulos 2012, p. 21) quotes several recruiters describing the need to have quick routines for selecting resumes: “I’m down to about seven seconds. [The information I’m looking for] needs to pop out so I’m very much onto keyword skimming. I’m almost like a Googlebot, like when you put in a search query. I have to do it really fast. I don’t have time to waste. ... I do realize how unfair the whole process is.”

guided by the expected benefits of reading, as indicated by the qualitative description from practitioners.

To illustrate how the allocation of costly attention affects discrimination, we propose a new model. First, acquiring information is costly and decision makers optimize how much information to acquire based on expected net benefits. This leads to “attention discrimination.” Second, imperfect information affects selection decisions because the less the decision maker knows about an individual, the more he relies on observable group attributes when assessing individual quality. Putting these two key features together, the endogenous attention magnifies (in most types of markets) the impact of animus and prior beliefs about group quality. Discrimination in selection decisions can persist even if perfect information about an individual is readily available, if it is equally difficult to screen individuals from dissimilar groups, and if there are no differences in preferences. The model also implies lower returns to employment qualifications for negatively stereotyped groups in selective markets, and for policy the important role of the timing of when a group attribute is revealed.

The model provides the following testable prediction: In “cherry-picking” markets where only top applicants are selected from a large pool of candidates (e.g., much of the labor market, admissions to top schools, the scientific review process in leading scholarly journals), decision makers should favor acquiring information about individuals from an *a priori* more attractive group. In contrast, in “lemon-dropping” markets where most applicants are selected (e.g., the rental housing market, admissions to nearly open-access schools), decision makers benefit more from acquiring information about individuals from a less attractive group. This is because more information should be acquired when its expected benefits are higher, which is when there is a higher chance that the informed decision differs from the status quo, i.e., when there is a higher chance of accepting the applicant in a market where most applicants are rejected and vice versa.⁴

We test the predictions of the model by monitoring information acquisition in three field experiments: in rental housing and labor markets in the Czech Republic and in the labor market in Germany. We send e-mails responding to apartment rental advertisements and to job openings. In each country we study discrimination against negatively stereotyped ethnic minorities and randomly vary the names of fictitious applicants. In the German labor market we also vary the quality of applicants by signaling recent unemployment in the e-mail. To monitor information acquisition in the labor market, employers receive an e-mail application for a job opening, which contains a hyperlink to a resume. Similarly, in the housing market landlords can click on a hyperlink located in the e-mail and learn more on an applicant’s personal website. We monitor whether employers and landlords open the applicant’s resume (resp. website) as well as the intensity of information acquisition.

While we find strong evidence of discrimination against minorities in a selection decision (invitation to a next stage) on both the housing and labor markets, we also document that systematic discrimination arises even earlier, during the process of information acquisition. The key findings on attention allocation are as follows. In the labor markets in both countries, employers put more effort into opening and

⁴Bose and Lang (2015) use similar logic in a different setting by showing that costly monitoring in the workplace is most beneficial when the employer has neither too low nor too high priors about the quality of a worker.

reading resumes of majority compared to minority candidates, while on the rental housing market landlords acquire more information about minority compared to majority candidates. Signaling unemployment lowers attention to an applicant's resume, similarly as minority name does. The set of results on attention allocation is consistent with the proposed model of discrimination with endogenous attention. The labor markets we study are very selective, as indicated by low invitation rates, and decision makers acquire *less* information about a priori less attractive applicants, whether it be a person with minority ethnic status or unemployed. In contrast, the rental housing market is not selective and decision makers acquire *more* information about applicants who look a priori less attractive. Later, we also discuss alternative explanations.

Methodologically, our paper contributes to efforts to test theory with enhanced measurement tools. In the lab, researchers have fruitfully complemented choice data with measures of the decision-making process to sort through alternative theoretical explanations of observed behavior. These techniques involve eye-tracking (Knoepfle, Wang, and Camerer 2009; Krajbich, Armel, and Rangel 2010; Arieli, Ben-Ami, and Rubinstein 2011; Reutskaja et al. 2011; Devetag, Di Guida, and Polonio 2016), mouse-tracking⁵ (Camerer et al. 1993; Costa-Gomes, Crawford, and Broseta 2001; Costa-Gomes and Crawford 2006; Gabaix, et al. 2006; Brocas et al. 2014), or monitoring brain activity (Bhatt and Camerer 2005; Hare, Malmaud, and Rangel 2011). Camerer and Johnson (2004) and Crawford (2008) summarize how progress in testing theories of human behavior has been facilitated by using information acquisition measures. To the best of our knowledge, ours is the first study that integrates monitoring information acquisition, in addition to selection decisions, into a field experiment.

In order to identify discrimination based on ethnicity, gender, caste, or sexual orientation in the labor and housing markets, previous correspondence tests⁶ estimated the effects of a group-attribute signal (mostly names) in applications (e.g., Neumark, Bank, and Van Nort 1996; Weichselbaumer 2003; Bertrand and Mullainathan 2004; Ahmed and Hammarstedt 2008; Banerjee et al. 2009; Kaas and Manger 2012). These experiments measure the likelihood of a callback (or invitation) as the outcome of interest.⁷ We offer an extension of this widely-used design by measuring effort expended to open and read resumes in the labor market and to acquire information about potential tenants in the rental housing market.⁸ Although

⁵Mouse-tracking, a technique closest to the monitoring tools used in this paper, typically uses Mouselab software, which displays information hidden in boxes on the computer screen and then tracks which and how many pieces of information subjects acquire.

⁶Two types of procedures have been used to test for the extent of discrimination on markets (Riach and Rich 2002). Correspondence tests involve responding to vacancies with written applications. Personal approaches, typically referred to as audit tests, include studies that have trained individuals attending job interviews or applying over the telephone.

⁷An important exception is Milkman, Akinola, and Chugh (2012) who study race and gender discrimination in academia and measure not only callback of faculty members reacting to students' requests to meet but also analyze the speed of their reply. Conditional on receiving a callback, in our experiments we do not find any significant difference in response speed across ethnic groups.

⁸The effort to better inform theories of discrimination by collecting novel types of data and performing experiments across distinct markets relates our work to List (2004), who combines a natural field experiment with artefactual field experiments to distinguish between taste-based and statistical discrimination in a product market, and to Gneezy, List, and Price (2012), who measure discrimination based on disability, gender, race, and sexual orientation across several markets to understand how the controllability of a group attribute affects discrimination.

the interview invitation decision can also be interpreted as a choice about costly information acquisition, the richer data about the decision-making process is useful for at least two reasons. First, since the costs of reading a resume are tiny compared to interviewing an applicant, it is an open question as to whether discrimination manifests itself already at the very outset of the decision-making process. This is potentially important for policy, since very early signals would have a larger impact on outcomes and also because addressing the smallest frictions in the early stage, such as the cost of reading a resume as opposed to the cost of an interview, might be easier. Second, measures of reading effort allow for a richer understanding of how attention discrimination operates on different types of markets.⁹

Our model of attention discrimination contributes to existing theories of discrimination (for a recent survey see Lang and Lehmann 2012). It is related most closely to “screening discrimination” (Cornell and Welch 1996), in which the key assumption is that it is more difficult to understand signals from a culturally dissimilar group (Lang 1986). Also, researchers (e.g., Greenwald, McGhee, and Schwartz 1998; Bertrand, Chugh, and Mullainathan 2005; Stanley, Phelps, and Banaji 2008) have argued that due to negative unconscious attitudes—“implicit discrimination”—people often use simple decision rules biased against negatively stereotyped groups, which may result in little effortful scrutiny of relevant information. In our model, differences in acquired knowledge are an outcome of the agent’s choice and can arise even if the signals provided are equally informative across groups and there are no unconscious biases in attention. This approach relates our model to a growing literature on rational inattention that uses an optimizing framework to study the effects of limited attention to the available information on a range of (mostly macroeconomic) phenomena (e.g., Sims 2003; Mackowiak and Wiederholt 2009; Woodford 2009; Van Nieuwerburgh and Veldkamp 2010; Matějka and Sims 2011; Caplin and Dean 2015; Matějka and McKay 2015).

The rest of the paper is organized as follows. In Section I we develop a model of an agent who decides how much to learn about an applicant and we describe how attention discrimination can arise and its implications for discrimination in selection decisions. We also formulate testable predictions for the field experiments. Sections II–IV detail the experimental designs and present empirical results in the rental housing and labor markets. Section V provides a discussion about how the results map on the proposed model and alternative interpretations. Section VI concludes.

I. The Model of Attention Discrimination

A. Setup of the Model

We model a two-stage decision maker’s (DM) choice about an applicant. A notable difference from existing models is that the level of additional information

⁹Since the invitation decision combines a choice to learn more about an applicant with a preselection decision (narrowing down the pool of applicants), it is difficult to infer the sign of the gap in willingness to acquire information from observed gaps in the likelihood of invitation. This is particularly the case in lemon-dropping markets, where unfavorable stereotypes or preferences are predicted to lead to greater information acquisition but a lower likelihood of invitation.

on the individual-specific quality is endogenous to the group's characteristics. In the first stage, the DM first observes the applicant's group of ethnic origin G , and then decides whether to pay additional attention to the applicant and whether to invite the applicant for an interview. In the case that the applicant is invited to the second stage, then the DM receives additional information about the applicant and chooses whether to accept him or not. The role of the first stage is to preselect applicants to potentially save on costs from inviting unsuitable applicants to the second stage.

For the DM, the applicant is of an inherent unknown payoff π , which consists of two components,

$$\pi = q - d_G,$$

where q is an unknown objective quality of the applicant which can include skill, work ethic, or reliability, and d_G is the DM's known distaste toward the applicant's group G or the distaste of individuals with whom the DM interacts, e.g., customers or neighbors. Quality in group G is distributed according to $N(q_G, \sigma_G^2)$, which is known by the DM and it forms the DM's prior knowledge about q . With respect to information acquisition, the quality q can be expressed as follows:

$$q = q_G + q_1 + q_2,$$

where $q_1 + q_2$ is the deviation of the applicant's quality from the group's mean q_G . We assume that at the beginning of the whole process q_G is observed. Then, in the first stage the DM can acquire information about q_1 only, which is drawn from $N(q_1, \sigma_{G,1}^2)$ and is independent from q_2 . For instance, in the case of a job application, q_1 summarizes all quality that can be inferred from a resume. In the second stage (e.g., during the interview), the quality q is observable.

The DM knows what is the best alternative to the applicant, and thus knows the reservation payoff R from rejecting the applicant in either of the stages.¹⁰ The DM maximizes the expectation of the payoff from accepting or rejecting the applicant less the incurred costs of inspection during the whole process:

$$\text{payoff} = \begin{cases} \pi & \text{if the DM accepts the applicant} \\ R & \text{if the DM rejects the applicant} \end{cases} - \text{inspection costs.}$$

In the first stage, the DM faces two choices. First, he chooses whether to pay the cost of inspection C_1 , e.g., whether to read the applicant's resume. If he pays the cost then he observes q_1 , and his posterior belief about the quality is $N(q_G + q_1, \sigma_G^2 - \sigma_{G,1}^2)$, since what is left to learn is q_2 only, otherwise the belief is $N(q_G, \sigma_G^2)$. Second, based upon the posterior belief he chooses whether to invite the applicant to the second stage, or not. The cost of invitation is C_2 . In the second stage, the DM observes the applicant's quality and decides whether to accept him.

¹⁰The quality q and reservation payoff R also summarize all payoff-relevant implications given by the current market situation, which include the general equilibrium effects or even wage demands by each particular applicant. For instance, if in equilibrium everyone pays more attention to the majority, and filter out good majority candidates while the good minority applicants are still available, then q and R adjust accordingly.

At this stage, when all costs of information acquisition are sunk, the applicant is accepted if and only if $q - d_G > R$.

DEFINITION (The DM's first-stage problem): *Upon observing G , the DM first chooses whether to incur C_1 and receive additional information, or to reject or invite the applicant without it. He chooses the action that maximizes the expected payoff:*

$$\text{payoff}(\text{reject}) = R$$

$$\text{payoff}(\text{invite}) = E[\max(R, q - d_G)] - C_2$$

$$\text{payoff}(\text{info}) = E[\max(R, E[\max(R, q - d_G) | q_1] - C_2)] - C_1.$$

In principle, there can be two types of situations. A cherry-picking market is a selective one in which, without any information except for the group attribute, the DM prefers rejecting the applicant to inviting him to the second stage, $\text{payoff}(\text{reject}) > \text{payoff}(\text{invite})$, and vice versa for the lemon-dropping market. For instance, a cherry-picking market is a labor market with many applicants for one job posting, where a priori a very few applicants are fit for the job, while in many locations the rental housing market is a lemon-dropping market, where an average applicant is acceptable.

Let us emphasize that the qualitative implications of the model would be unchanged if the DM faced a random pool of alternatives to the applicant in the second stage, rather than a given R . This is because the effect of a random pool of alternative applicants is already encompassed in the random component q_2 , which is observed during the second stage. The only thing that matters in the second stage is payoff from the applicant relative to payoff of the alternative; a high draw of q_2 can model a low draw of quality of the alternatives. The model thus also describes the behavior of a DM facing a sequential choice of applicants for a limited number of interview slots. In that case, R and the distribution of q_2 would depend on applicants the DM interacted with previously. For instance, if one applicant is accepted, then the next applicant faces a higher R .

B. Effects of Preferences and Beliefs on Attention Allocation

In this subsection we describe how attention to an applicant is affected by sources of discrimination highlighted by prominent theories of discrimination. In the next subsection we describe how endogenous attention affects the magnitude of discrimination in final selection decisions.

Proposition 1 below describes a new channel through which discrimination can operate: costly attention. It addresses how endogenous attention depends on the DM's choice situation and beliefs.¹¹ We discuss these predictions below, and test them in the empirical part of the paper.

¹¹Note that the model describes the choice between no attention and some attention, only, and not between different positive levels of attention. We do this for the sake of simplicity as well as because most of our empirical

PROPOSITION 1 (Attention Discrimination):

- (i) Applicants from group G that are less attractive a priori (due to lower q_G , σ_G^2 , or higher d_G , C_2) are paid (weakly) less attention in the cherry-picking markets and (weakly) more attention in the lemon-dropping markets.
- (ii) Applicants from a dissimilar group G with higher cost of attention C_1 or lower $\sigma_{G,1}^2$ are paid less attention in either market.

PROOF:

See the online Appendix.

A distaste toward a certain group in our model is captured by parameter d_G (Becker 1971).¹² An increase in d_G decreases the mean payoff from an applicant from group G . The resulting effect on attention differs across markets: a higher distaste implies less attention in the first stage in cherry-picking markets, and more attention in lemon-dropping markets. The reason is that costly information is useful only when there is a chance that it changes the status quo decision (which is to reject in cherry-picking and to invite in lemon-dropping markets). Benefits from information are illustrated by the shaded regions in Figure 1. The figure presents distributions of expected payoffs from inviting the applicant to the second stage. These distributions are determined by distributions of q_1 and by expected payoffs in the second stage for each particular q_1 . In cherry-picking markets, information is useful only when the DM identifies applicants that are better not rejected. When the distaste d_G increases, then the distribution shifts to the left, since there are fewer good candidates, and the benefits from information decrease (panel A). In contrast, in lemon-dropping markets (panel B), benefits from information acquisition are given by the potential of discovery of bad applicants and thus a decrease in the mean payoff increases the DM's attention.¹³

Next, we consider statistical discrimination (Phelps 1972; Arrow 1973), which is driven by differences in beliefs about the applicant's quality. In our model, this channel is represented by a change in the mean q_G or variance σ_G^2 . The implications of a reduction in q_G are the same as for an increase in the distaste parameter d_G : less attention in cherry-picking and more attention in lemon-dropping markets. A decrease in the variance of beliefs σ_G^2 , when holding $\sigma_{G,1}^2$ fixed, has the same effects as a decrease in q_G , since a higher variance

results correspond to such a choice. In an alternative model with a sequential choice of levels of attention, where after some information is acquired the DM could choose to acquire more information, the results of Proposition 1 would hold, too. The only difference would be that the type of market would be conditioned on the information received before the choice of additional information.

¹² An alternative way of introducing differences in taste would be to vary reservation payoff R across groups. Note that an increase in R has the same effect as an increase in d_G of the same size.

¹³ An increase in parameter C_2 has similar effects as an increase in d_G . This makes sense intuitively—higher C_2 may arise due to distaste toward interacting with a certain group during the interview. In our model, C_2 is not a cost of additional information that can be incurred independently of a selection decision, and unlike C_1 it is a cost that needs to be incurred for any accepted applicant and thus deducted from the payoff.

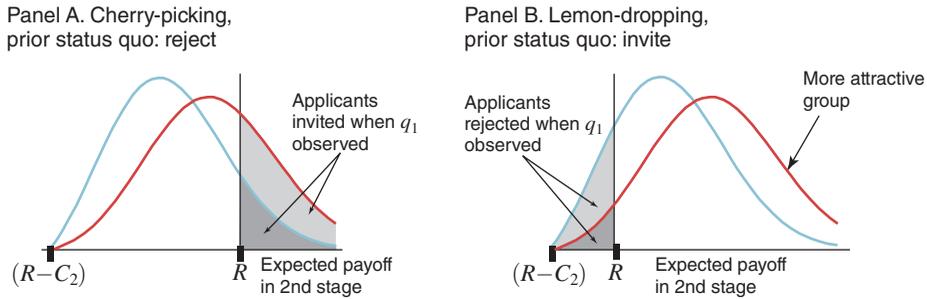


FIGURE 1. EXPECTED BENEFITS FROM INFORMATION ACQUISITION IN THE FIRST STAGE

increases the likelihood of good candidates, while the bad ones are filtered out by the DM in the second stage.¹⁴

Last, we consider the effects on attention of a greater difficulty to understand signals from a culturally dissimilar group (Cornell and Welch 1996). In our model, such a dissimilar group would in the first stage be characterized by either a lower $\sigma_{G,1}^2$ (i.e., less uncertainty can be resolved by reading a resume) or by a higher cost of information C_1 (i.e., reading requires more effort). In both cases, attention in the first stage weakly decreases in either market. This is because both of these characteristics affect $payoff(info)$ only: they determine the level of attention in the first stage, and do not affect the DM’s choices when no information is provided in the first stage.¹⁵

To summarize, there are two types of group characteristics entering the model. Changes in characteristics related to taste and beliefs about groups ($d_G, q_G, \sigma_G^2, C_2$) affect a priori attractiveness of an applicant and have opposite effects on attention across the two types of markets. On the other hand, changes in characteristics related to a process of screening ($C_1, \sigma_{G,1}^2$) have the same effects in both markets.

C. Endogenous Attention and Discrimination in Selection Decisions

In this subsection, we describe how endogenous attention can exacerbate discrimination in selection decisions, discuss cases when this is less likely to happen, and point to potential policy implications.

COROLLARY 1 (Discrimination Exacerbation):

- (i) *If both groups are either in the cherry-picking market or both are in the lemon-dropping market, then endogenous attention further disadvantages the group less attractive a priori (due to lower $q_G, \sigma_{G,2}^2$, or higher C_2) in the selection decision.*

¹⁴If σ_G^2 increases, then $E[\max(R, q - d_G) | q_1] - C_2$ increases for all q_1 , since the bad candidates are filtered out while the good ones are accepted, and thus the distribution in Figure 1 shifts to the right in the sense of first-order stochastic dominance.

¹⁵In contrast, characteristics described above ($d_G, q_G, \sigma_G^2, C_2$) affect not only $payoff(info)$, but also $payoff(invite)$, and thus determine the a priori attractiveness of the group.

- (ii) *Endogenous attention disadvantages a dissimilar group G (with higher C_1 or lower $\sigma_{G,1}^2$) in the cherry-picking market and helps it in the lemon-dropping market.*

PROOF:

See the online Appendix.

The statement that endogenous attention disadvantages the less attractive group means that the difference in acceptance probability between applicants from a majority group and group G is (weakly) higher than if the level of attention to G were exogenously fixed at the level of attention paid to the majority.

The findings above imply that when the general population as well as group G face the same type of market (cherry-picking or lemon-dropping), then endogenous attention magnifies the effects of differences in taste and beliefs about groups. In cherry-picking markets, the applicant is rejected when no additional information about him is acquired and therefore the DM's attention weakly increases an applicant's chances of being invited. It follows that a higher chance of being invited implies a higher chance of being selected in the second stage, since the invitation is a prerequisite for selection and qualities are observed in the second stage. At the same time, the less attractive groups (i.e., groups with higher d_G , lower q_G , or lower σ_G^2) are paid less attention. In lemon-dropping markets, attention decreases the likelihood of an invitation, and the disadvantaged group is paid more attention.

There is one special type of situation when endogenous attention may not magnify discrimination in selection decisions—a “middle-market,” in which preferences for or beliefs about the two groups are very different and the less attractive group faces the selective situation and would be rejected in the absence of additional information, while the other group is in the lemon-dropping market. The disadvantaged group can, in this case, be paid more attention than the majority group, and endogenous attention can work to the minority's advantage.

The endogenous attention also magnifies the role of differences in difficulty to understand signals across groups in cherry-picking markets, where attention is desirable, but culturally dissimilar groups receive less of it. This is not the case in lemon-dropping markets, in which dissimilar groups also receive less attention, but here less attention improves chances of selection.

In the model, for the sake of simplicity, we assume that quality q is perfectly observable in the second stage. Thus, the DM's belief about an applicant's quality in the second stage is independent of the level of attention in the first stage, and hence endogenous attention in the first stage influences a final selection decision only via its effects on invitations to the second stage. If we allowed for the realistic case of imperfect knowledge in the second stage, then the magnifying effect of endogenous attention on discrimination in selection decisions would be further reinforced in cherry-picking markets. However, in lemon-dropping markets, while endogenous attention disadvantages applicants from the group G in terms of invitations to the second stage, the higher attention in the first stage may provide them with an advantage conditional on being invited. The DM would in the second stage possess more precise knowledge about such candidates, which might increase the likelihood of selection.

The findings above suggest the important role of timing of when the group attribute is revealed during the decision-making process in selective markets, an insight that is potentially interesting for policy. The following corollary is an immediate implication of Proposition 1 and Corollary 1.

COROLLARY 2 (Timing of Ethnic Group Revelation): *If both groups are either in the cherry-picking market or both are in the lemon-dropping market, then the probability that an applicant from a less attractive group is accepted is (weakly) lower if he is known to be from G a priori rather than when he is first considered to be from a general population and his membership in G is revealed only before the final selection decision.*

Postponing the revelation helps the disadvantaged group by leveling the attention a DM pays to applicants. The probability that an applicant from a less attractive group is accepted is lower if he is known to be from group G prior to when the DM chooses whether to inspect the applicant in the first stage rather than when the applicant is first considered to be from a general population, and membership in G is revealed only before the final selection decision. This effect is not present in the standard model of statistical discrimination, because there the DM receives signals of exogenously-given precision and forms his posterior knowledge independent of the signals' succession, while in our model the first signal—i.e., the group attribute—affects the choice of whether to acquire an additional signal.

An important question that goes beyond the presented model is what are the dynamic and general equilibrium effects of endogenous attention? Endogenous attention has interesting implications for the persistence of discrimination that is driven by different beliefs across groups. It is known that such discrimination can persist in the long run, for instance if agents can invest in their skills, and if the skills and the investment are not perfectly observed (e.g., Coate and Loury 1993). In this case, a negatively-stereotyped group has less incentive to acquire the skills, which results in a self-fulfilling negative stereotype. With endogenous attention, this disincentive effect is further reinforced in cherry-picking markets, where negatively stereotyped groups face not just lower likelihoods of acceptance but they are also less rewarded for their credentials due to lower attention given to them. The effect is attenuated on the lemon-dropping market. Regarding taste-based discrimination, endogenous attention does not seem to provide novel implications in terms of whether market forces would attenuate or eliminate such discrimination either by growth or endogenous entry of nondiscriminating firms (Becker 1971; Arrow 1973).

Finally, the model also speaks to a dynamic setting, in which the DM is aware of the possibility of having inaccurate beliefs.¹⁶ In our model, the DM pays more attention when the uncertainty $\sigma_{G,1}^2$, which can be resolved in the first stage, is higher. This intuition extends to a dynamic model, too, when the whole selection

¹⁶ Alternatively, the DM can have inaccurate beliefs and be unaware of it. This may arise, for example, when DMs recall only a group's most representative or distinctive types (Bordalo, Gennaioli, and Shleifer forthcoming). Importantly, the effects on attention are driven purely by the form of beliefs, regardless of whether they are accurate or not. In selection decisions true qualities interact with beliefs and therefore inaccurate beliefs may change the predicted effects in either direction. For instance, being deemed as a highly homogeneous group is favorable in lemon-dropping markets, but disadvantageous in cherry-picking markets.

process regarding one candidate plays the role of a first stage for all future candidates. When uncertainty about a group is higher, then the DM would pay more attention and be more likely to invite an applicant from group G simply to learn more about the whole group, and use such information in the future, i.e., to make future beliefs more accurate.

II. Field Experiment in the Rental Housing Market

In the first experiment, we study ethnic discrimination in the rental housing market in the Czech Republic, a market with a low level of selectivity (as we document below, a large fraction of applicants are invited).

We focus on two ethnic minorities: Roma and Asian. The Roma population constitutes the largest ethnic minority in the European Union (estimated at 6 million people, 1.2 percent) as well as in the Czech Republic (1.5–3 percent). Intolerance and social exclusion of Roma is considered one of the most pressing social and human rights issues in the European Union (European Commission 2010). East Asians (mostly Vietnamese but also Chinese or Japanese) are the second-largest ethnic minority group in the Czech Republic (0.6 percent) and migrants from East Asia form large minority groups in many European countries. In the Czech Republic they are mostly self-employed in trade and sales businesses and lack formal employment.

Both minority groups are disadvantaged economically and socially, and face unfavorable stereotypes. The unemployment rate of Roma in the Czech Republic was estimated at 38 percent, compared to 9.4 percent overall unemployment rate in 2012. While 84 percent of the majority population complete a high school or university degree, the proportion is 47 percent and 33 percent for the Vietnamese and Roma adult population, respectively (Czech Statistical Office 2011). An opinion poll revealed that 86 percent and 61 percent of Czechs would not feel comfortable or would find it unacceptable to have Roma and Vietnamese as neighbors, respectively.¹⁷ In an online survey that is discussed below (Survey I), landlords expect individuals with Roma and Asian names to be worse tenants than the majority applicants for apartment rentals.¹⁸

A. Experimental Design

Manipulating Identity of Applicant.—The experiment was based on sending e-mails expressing interest in arranging an apartment viewing. To evoke ethnic minority status we designed three fictitious applicants: representatives of the Asian and Roma ethnic minorities and a control identity of the white majority group. The only real attributes of these identities were a name, an e-mail address, and a

¹⁷For more details about the socioeconomic status of Roma in Central and Eastern European countries see Barany (2002). FRA and UNDP (2012) describe documented inequalities in education, employment, health, and housing outcomes between Roma and majority populations in the Czech Republic and other EU countries. Spaan, Hillmann, and van Naerssen (2005) provide a detailed description of the integration of immigrants from East Asia in Europe.

¹⁸Similarly, another survey (Survey III) documents that university students expect individuals with Roma and Asian names to have a lower socioeconomic status, as measured by education level and quality of housing.

personal website.¹⁹ We selected the names based on name frequency data: Jiří Hájek (white majority-sounding name), Phan Quyet Nguyen (Asian-sounding name), and Gejza Horváth (Roma-sounding name).²⁰ Since the e-mail address contained the applicant's name, the name is arguably the first piece of information that a landlord learns about the applicant. For the sake of brevity, we denote applicants with a white majority-sounding name as "white applicants" or as "majority applicants," applicants with ethnic minority-sounding names (both Asian and Roma) as "minority applicants," and applicants with Asian-sounding and Roma-sounding names as "Asian applicants" and "Roma applicants," respectively. Note that technically the results of our experiments describe the effects of the ethnic sounding-ness of the names rather than the effects of ethnicity itself.

To verify that landlords associated the selected names with respective ethnic groups, we conducted a pre-survey on a sample of 50 respondents. All respondents associated the name Jiří Hájek with the Czech nationality and the name Phan Quyet Nguyen with one of the Asian nationalities (92 percent associated it with the Vietnamese nationality), and the name Gejza Horváth was thought to be a Roma name in 82 percent of cases, indicating a strong link between names and ethnicity. To confirm that the application e-mails from all applicants would be delivered and not identified as spam, prior to the implementation of the experiment we sent each variant of the e-mail message to 40 individuals with e-mail accounts from different providers. In all cases the e-mails were delivered successfully.

Manipulating Access to Information.—In application e-mails, we used three manipulations of access to information about applicants (for an overview of the experimental design see online Appendix Table S1). First, in the No Information Treatment, the e-mail contains a greeting and the applicant's interest in renting an apartment, but does not provide any information about the characteristics of the applicant other than his minority/majority-sounding name.²¹ Invitation rates in this treatment are informative about the type of market, since landlords can make inferences based on the applicant's name (and the short text) only. Recall that we defined the cherry-picking market as one where the status quo, after learning a group attribute (but no other information), is not to invite an applicant—thus with no heterogeneity in the DM's thresholds, the invitation rate should be close to zero. In contrast,

¹⁹There is a difficult trade-off involved in organizing this type of experiment. While informed consent is clearly desirable, it is extremely difficult to measure discrimination with the consent of participants in natural field experiments (List and Rasul 2011). Given their social benefits, audit studies and correspondence tests are considered among the prime candidates for the relaxation of informed consent (Riach and Rich 2002, Pager 2007). Our research has been approved by the Director of the Institute of Economic Studies, Charles University in Prague. We followed the conventional IRB standards for these types of experiments and took the following steps to minimize the landlords' costs and risks. In particular, the information acquisition was designed such that it took little effort and time and we quickly and politely declined invitations for an apartment viewing, within two days at most. We sent only one application to each landlord, and thus are not able to identify discrimination at an individual level, and after the data collection we deleted identifiers of individual landlords. A similar practice was followed in our companion experiments in the labor markets.

²⁰Jiří is the most frequent Czech first name and Hájek is among the top 20 most frequent surnames in the Czech Republic. Nguyen and Horváth are the most frequent surnames for the Asian and Roma minorities, respectively.

²¹Specifically, the text of the e-mail in the No Information Treatment was as follows: "Dear Sir/Madam, I am writing because I am very interested in renting the apartment that you have advertised. When would be a good time to come see the apartment? Best regards, Phan Quyet Nguyen." We avoided syntax or spelling mistakes. For wording of all manipulations of the e-mail, see the online Appendix.

in the lemon-dropping case the status quo is to select all applicants. We consider the invitation rate of 50 percent as an approximate dividing line between the two types of markets. This is the exact dividing line when heterogeneity among DMs is small and symmetric. In this situation a DM hesitates most whether to invite an applicant or not, and thus additional information is predicted to be most beneficial.

Second, in the Monitored Information Treatment, the e-mail uses the same sentence to express interest in viewing an apartment as in the No Information Treatment. The only difference is that it includes the hyperlink to a personal website located in the applicant's electronic signature, which gives landlords an opportunity to acquire more information about an applicant. The link has a hidden unique ID number assigned to each landlord, which allows us to distinguish landlords who decide to acquire information about the applicant. Software similar to Mouselab monitors landlords' information acquisition on the website. Five different boxes are located in the main section of the website, each with a heading representing a type of information that is hidden "behind" the box—age, marital status, smoking habits, occupation, and education. A snapshot is displayed in online Appendix Figure S1. When the boxes were uncovered, landlords learned that the applicant is 30 years old, single, a nonsmoker, and working in trade with a steady income. We randomly varied whether an applicant reported having a high school or college degree.

Since only one box can be opened by a computer mouse at one point in time, the software allows us to identify whether a landlord decides to acquire information on an applicant's website, and how many and which pieces of information receive attention. These monitoring features provide direct insight into the process of information acquisition. In addition to the boxes with personal information, the website also contains tags for a personal blog, pictures and contact information (when accessed, an "under construction" note pops up, to reduce landlord's costs by limiting the time spent on the website). The design of the website is based on a professionally created template, which is freely available on the Internet.²²

In the third manipulation, the e-mail again uses the same introductory sentence as in the No Information Treatment, but instead of providing a hyperlink to a website, the applicant reports the same characteristics directly in the body of the e-mail. Specifically, we added the following text: "I am a thirty-year-old man, I am single, I have a college [a high school] degree, and I do not smoke. I have a steady job (with a regular paycheck) at a company." Again, we randomly varied the education level. This allows us to study the effects of name on how much landlords respond to changes in available information in terms of invitation rates. The motivation for this treatment is to test whether name effects on responsiveness mimic the name effects on attention from the Monitored Information Treatment.

An online survey (Survey I) implemented after the experiment among a different sample of landlords ($N = 60$) shows that the set of applicant characteristics reported in the second and third manipulations are considered attractive, as compared to the typical population of applicants on this market. The landlords were

²²Still, to some landlords the website may appear unusual and this may affect their callback. Nevertheless, it should be noted that the content and the design of the website cannot affect a decision as to whether or not to open it, since the decision happens when the landlord sees only the link. Providing a hyperlink to a personal website is a common feature in an electronic signature.

given two profiles (across subjects), which contained the same set of characteristics as described above, and asked: “Based on your experience with renting an apartment, how would you compare the following applicant to other applicants? 1 = strongly above average, 2 = above average, 3 = average, 4 = below average, 5 = strongly below average.” We find that both profiles, those with a high school and a college degree, were evaluated as substantially above average (2.19 and 1.64, respectively).²³

B. Sample Selection and Data

The experiment was implemented between December 2009 and August 2010 in the Czech Republic, mostly in Prague. Over that period, we monitored four (out of ten) major websites that provide rental advertisements. Placing an ad on these websites requires a small fee, while responding to an advertisement is free. We chose to apply only for small homogeneous apartments, of up to two rooms with a separate kitchen, that looked suitable for a single tenant without a family. We excluded offers mediated by real estate agents and also offers where landlords did not make their e-mail publicly available and relied on a telephone or an online form (66 percent), in order to be able to monitor information acquisition. Overall, we responded to 1,800 rental ads and randomly assigned an applicant name and provided information. We recorded the gender of the landlord, implied by the name, and the characteristics of apartments commonly published as a part of the advertisement such as rental price, the size of the apartment, and whether it is furnished. These characteristics vary little across experimental treatments, indicating that randomization was successful (online Appendix Table S3).

To measure attention in the Monitored Information Treatment, we record whether a landlord visits an applicant’s personal website and how many and which boxes with information he uncovers. To measure responses to the applicant, we distinguish between a positive response, indicating either a direct invitation to an apartment viewing or an interest in further contact, and a negative response, capturing the rejection of an applicant or the absence of response.²⁴ Note that with the correspondence experimental approach a researcher does not measure the ultimate outcomes, i.e., whether an applicant rents the apartment and for what price. Nevertheless, since the invitation is typically a prerequisite for the final positive decision, it is likely that the gaps in the share of positive responses across ethnic groups translate into gaps in final decisions about actual rental.

C. Results

Do Landlords Discriminate Against Minorities?—We start the analysis by looking at whether ethnic minorities are discriminated against when no information about the applicant other than his name is available to a landlord (No Information

²³ These characteristics are likely seen particularly positively (relative to expectations) for minority applicants, given the gaps in education and employment relative to the majority population.

²⁴ As a robustness check, we also estimated the effect of minority-signaling names on callback (online Appendix Table S4), which distinguishes applications that result in contact, regardless of whether it is a positive or negative response. Overall, we find a qualitatively similar impact of names on the callback rate as on the invitation rate.

TABLE 1—CZECH RENTAL HOUSING MARKET: INVITATION RATES AND INFORMATION ACQUISITION BY ETHNICITY, COMPARISON OF MEANS

	White majority name (W) (1)	Pooled Asian and Roma minority name (E) (2)	Percentage point difference: W – E, (<i>p</i> -value) (3)	Asian minority name (A) (4)	Percentage point difference: W – A, (<i>p</i> -value) (5)	Roma minority name (R) (6)	Percentage point difference: W – R, (<i>p</i> -value) (7)	Percentage point difference: R – A, (<i>p</i> -value) (8)
<i>Panel A. Invitation for a flat visit</i>								
No Information Treatment (<i>n</i> = 451)	0.78	0.41	37 (0.00)	0.39	39 (0.00)	0.43	36 (0.00)	3 (0.57)
Monitored Information Treatment (<i>n</i> = 762)	0.72	0.49	23 (0.00)	0.49	23 (0.00)	0.49	23 (0.00)	0 (0.92)
Monitored Information Treatment ^a (<i>n</i> = 293)	0.84	0.66	18 (0.00)	0.71	13 (0.00)	0.62	21 (0.00)	-9 (0.20)
Monitored Information Treatment ^b (<i>n</i> = 469)	0.66	0.37	29 (0.00)	0.35	31 (0.00)	0.39	27 (0.00)	4 (0.51)
Treatment with additional text in the e-mail (<i>n</i> = 587)	0.78	0.52	26 (0.00)	0.49	29 (0.00)	0.55	23 (0.00)	5 (0.29)
<i>Panel B. Information acquisition in the Monitored Information Treatment</i>								
Opening applicant's personal website	0.33	0.41	-8 (0.03)	0.38	-5 (0.24)	0.44	-11 (0.01)	6 (0.15)
Number of pieces of information acquired	1.29	1.75	-0.46 (0.01)	1.61	-0.32 (0.09)	1.88	-0.59 (0.00)	0.27 (0.17)
At least one piece of information acquired	0.30	0.40	-10 (0.01)	0.37	-7 (0.12)	0.44	-13 (0.00)	7 (0.12)
All pieces of information acquired	0.19	0.26	-8 (0.02)	0.24	-6 (0.12)	0.28	-10 (0.01)	4 (0.33)
Number of pieces of information acquired ^a	3.91	4.24	-0.33 (0.06)	4.23	-0.32 (0.15)	4.25	-0.34 (0.09)	0.02 (0.90)
At least one piece of information acquired ^a	0.92	0.98	-6 (0.02)	0.97	-5 (0.15)	0.98	-7 (0.03)	2 (0.47)
All pieces of information acquired ^a	0.56	0.64	-7 (0.23)	0.64	-8 (0.30)	0.64	-7 (0.30)	-0 (0.96)

Notes: Means. Panel A reports how name affects invitation for a flat visit and panel B how it affects information acquisition in the Monitored Information Treatment. Columns 3, 5, 7, and 8 report differences in percentage points, in the parentheses we report *p*-value for a *t*-test testing the null hypothesis that the difference is zero. The differences in the number of pieces of information acquired on the website are reported in absolute terms, not in percentage points.

^a The numbers are reported for the subsample of landlords who opened an applicant's website.

^b The numbers are reported for the subsample of landlords who did not open an applicant's website.

Treatment). In this treatment, the invitation rates reflect the tastes and prior beliefs about the expected characteristics of each group. We find that majority applicants are invited for an apartment viewing in 78 percent of cases, while minority applicants receive invitations in only 41 percent of cases (panel A of Table 1). The gap that arises solely due to name manipulation is large in magnitude (37 percentage points, or 90 percent) and statistically significant at the 1 percent level. Put differently, minority applicants have to respond to almost twice as many advertisements to receive the same number of invitations as majority applicants.

Next, we distinguish between applicants with Asian- and Roma-sounding names. The invitation rates are very similar: 43 percent for the Roma minority and 39 percent for the Asian minority applicants. The difference in invitation rate between the two minority groups is not statistically distinguishable (column 8), while the gap between the majority and each of the two minority groups is large and similar

in magnitude (columns 5 and 7). Table 2 documents the findings in a regression framework, where we control for the landlord's gender and the characteristics of the apartment described in an advertisement (price, size, furnishings).

Observation 1: Applicants with minority-sounding names are discriminated against. If no information about applicants is available, applicants with a majority-sounding name are 90 percent more likely to be invited for an apartment viewing compared to applicants with a minority-sounding name.

In the model, information acquisition is the most valuable and the DM pays the most attention when ex ante expected payoffs from invitation and rejection are equal (the invitation rate without additional information is 50 percent), and thus when the DM hesitates ex ante then any piece of information is useful and can affect the decision. The further the invitation rate is from 50 percent, the less information is acquired. The invitation rate for the majority applicants (78 percent) suggests they are in the lemon-dropping situation and the mean of the prior belief about this group is far above the threshold level of quality necessary for invitation. On the other hand, the invitation rate for minority applicants is 41 percent. Given that some fraction of apartments might have already been rented out by the time we sent the e-mail and thus their owners were unlikely to invite any applicant, it is difficult to say for sure whether the prior mean about a minority group is above (lemon-dropping situation) or below (cherry-picking situation) the threshold. Importantly, though, it's clear the landlords hesitate more about whether to invite applicants from the minority group and thus acquiring more information about minority applicants should be more valuable, compared to acquiring information about majority applicants.

Do Landlords Choose Different Levels of Attention to Information Based on the Ethnicity of an Applicant?—In the Monitored Information Treatment, we find that only less than half of the landlords open the applicant's website even though the cost of acquiring information is very small—literally one click on the hyperlink. Importantly, the applicant's name matters for attention allocation (panel B of Table 1 and Table 3). While 41 percent of landlords opened the website of minority applicants, 33 percent did so for majority applicants. When summing the number of applicant characteristics to which a landlord pays attention (the maximum is five), we find that landlords learn about 1.75 characteristics of a minority applicant and 1.29 for a majority applicant. Similarly, the likelihood of opening at least one of the boxes with information is 40 percent for minority and 30 percent for majority applicants, and the likelihood of opening all the boxes is 26 percent for minority and 19 percent for majority applicants. These differences in information acquisition measures across an applicant's ethnicity are statistically significant and are driven by both a greater likelihood of opening the website as well as more effort to acquire information, conditional on opening the website. Among a subsample of landlords who opened an applicant's website, we still observe that landlords are significantly more likely to open at least one of the boxes with information and to open a higher number of boxes when the applicants have minority names compared to majority name.

Online Appendix Table S5 reports further results about how ethnicity affects the process of information gathering. While the name affects the amount of information

TABLE 2—CZECH RENTAL HOUSING MARKET: INVITATION RATES BY ETHNICITY, REGRESSION ANALYSIS

Dependent variable:	Invitation for an apartment viewing				
	(1)	(2)	(3)	(4)	(5)
<i>Panel A</i>					
Sample:	No Information Treatment		Whole sample		
			All	W majority name	E minority name
Ethnic minority name	-0.39*** (0.04)		-0.37*** (0.04)		
Asian minority name		-0.41*** (0.05)			
Roma minority name		-0.39***	(0.05)		
Monitored Information Treatment			-0.08 (0.06)	-0.06 (0.05)	0.08** (0.04)
Ethnic minority name × Monitored Information Treatment			0.16** (0.06)		
Additional text in the e-mail: with high school			-0.00 (0.07)	-0.00 (0.06)	0.08* (0.05)
Ethnic minority name × Additional text in the e-mail: with high school			0.08 (0.08)		
Additional text in the e-mail: with college			0.00 (0.07)	0.01 (0.06)	0.15*** (0.05)
Ethnic minority name × Additional text in the e-mail: with college			0.14* (0.08)		
Observations	451	451	1,800	599	1,194
<i>Panel B</i>					
Sample:	Monitored Information Treatment				
				Landlords who opened applicant's website	Landlords who did not open applicant's website
	All	All	All		
Ethnic minority name	-0.23*** (0.04)	-0.26*** (0.04)	-0.28*** (0.04)	-0.18*** (0.05)	-0.29*** (0.05)
Opening applicant's website		0.26*** (0.04)	0.21*** (0.07)		
Ethnic minority name × Opening applicant's website			0.07 (0.09)		
Observations	762	762	762	293	469

Notes: Probit, marginal effects (dF/dx) in all columns of both panels. Robust standard errors in parentheses. In columns 1–3 of panel A and all columns of panel B the omitted variable is a white majority name. In columns 4–5 of panel A the omitted variable is a dummy for No Information Treatment. In all columns of both panels, we control for a dummy variable indicating a landlord being a female, a dummy variable indicating an unknown gender of a landlord (the mean of this variable is 0.02), size of an apartment, price of an apartment rental, and a dummy variable indicating an equipped apartment.

***Significant at the 1 percent level.

**Significant at the 5 percent level.

*Significant at the 10 percent level.

TABLE 3—CZECH RENTAL HOUSING MARKET: INFORMATION ACQUISITION BY ETHNICITY, REGRESSION ANALYSIS

Dependent variable:	Opening applicant's personal website		Number of pieces of information acquired		Number of pieces of information acquired about education and occupation		Number of pieces of information acquired about personal characteristics	
	Monitored Information Treatment: all observations				Monitored Information Treatment: subsample of landlords who opened applicant's website			
Sample:	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Ethnic minority name	0.08** (0.04)		0.46*** (0.16)		0.18** (0.09)		0.12 (0.13)	
Asian minority name		0.05 (0.04)		0.31 (0.19)		0.17* (0.10)		0.07 (0.14)
Roma minority name		0.11*** (0.04)		0.60*** (0.19)		0.18* (0.09)		0.15 (0.14)
Observations	762	762	762	762	293	293	293	293

Notes: Probit, marginal effects (dF/dx) in columns 1 and 2. OLS in columns 3–8. Robust standard errors in parentheses. In columns 3 and 4, the dependent variable is number of pieces of information a landlord viewed on applicant's personal website: minimum is zero and maximum is five. In columns 5 and 6 it is a number of pieces of information about education and occupation he/she uncovered: minimum is zero and maximum is two, and in columns 7 and 8 it is a number of pieces of information about personal characteristics (age, smoking habits, marital status) he/she uncovered: minimum is zero and maximum is three. In all columns the omitted variable is a white majority name. In all columns, we control for a dummy variable indicating a landlord being a female, a dummy variable indicating an unknown gender of a landlord, size of an apartment, price of an apartment rental, and a dummy variable indicating an equipped apartment.

*** Significant at the 1 percent level.

** Significant at the 5 percent level.

* Significant at the 10 percent level.

acquired, we do not find a systematic influence on which type of information is acquired as well as on the order in which different pieces of information are acquired. Unconditional on opening the website, the likelihood of opening a box about, for example, education level is 36 percent for minority applicants and 27 percent for majority applicants (panel A). Thus, the difference due to name manipulation is 33 percent (or 9 percentage points). A similar picture arises for other individual characteristics: the likelihood of paying attention to those is 30–46 percent greater for minority applicants compared to majority applicants. Also, the landlords who visit the website are more likely to open each of the boxes for minority applicants compared to majority applicants, but the differences are not statistically significant with the exception of the box with occupation information (panel B).

In terms of the order of uncovering the boxes, we find that conditional on opening the website, the likelihood of uncovering each of the boxes as the first one does not significantly differ across ethnicity (panel C). Similarly, conditional on opening all the boxes, the order of uncovering does not differ across applicants' names (panel D). Together, these results suggest that the observed differences in acquired information are not driven by landlords being worried about a particular single attribute of minority applicants, but rather by a more general effort to screen this group more carefully.

Distinguishing between the two minority groups reveals that, compared to the majority applicant, landlords acquire more information about both Roma and Asian applicants (columns 5 and 7 of Table 1). We also observe that the amount of acquired information is somewhat (although insignificantly) greater for Roma

applicants relative to Asian applicants (column 8). This is interesting given that the landlords appeared to hesitate most on whether to invite Roma applicants, since the invitation rate of this minority was closest to the 50 percent invitation rate. Table 3 documents the findings in a regression framework.

Observation 2: Landlords pay more attention to available information about applicants with a minority-sounding name relative to applicants with a majority-sounding name.

Responsiveness to Available Information.—In order to test whether landlords are more responsive to available information provided by minority applicants compared to majority applicants, as suggested by observed differences in attention, we estimate the effects of three manipulations in the available information on invitation rate: (i) adding a sentence to the e-mail message signaling attractive characteristics of the applicant; (ii) varying the education level between high school degree and college degree in the added sentence; and (iii) having access to an applicant's personal website.

We find that the invitation rate responds to information provided by applicants with minority names,²⁵ a pattern which is consistent for all three manipulations of available information. Column 5 in panel A of Table 2 shows that, relative to the No Information Treatment, the invitation rate increases by 8 percentage points for minority applicants who add the sentence reporting a high school education. The increase is 15 percentage points for minority applicants who add a sentence and report having a college degree. The pure effect of reporting a college degree compared to a high school degree is 8 percentage points, which is marginally significant statistically (column 3 of online Appendix Table S7). In contrast, there is little response in the invitation rate when the same manipulations of available information are performed by the applicant with the majority name. The invitation rate remains at the same level, 78 percent, independent of whether the applicant provides no information, includes a sentence about his characteristics, and also does not respond to changes in his education level (column 4 of Table 2).

Taken together, the decision makers are found to be more sensitive to information provided by minority candidates compared to majority candidates (columns 3 of Table 2). The interaction effect of having a minority name and adding a sentence with a college degree is 14 percentage points. The interaction effect is still positive (8 percentage points) when reporting a high school degree. Similarly, among applicants who provide information in the body of the e-mail, the interaction effect of a minority name and reporting a college degree is 7 percentage points (online Appendix Table S7). The last two interaction effects are not statistically significant. As a consequence, the discrimination in terms of invitation rate is 37 percentage points in the No Information Treatment, and it diminishes to 29 percentage points for applicants who add a sentence and report a high school degree and to 22 percentage points for applicants who add a sentence and report a college degree (online Appendix Table S7).

²⁵ Landlords are responsive to changes in available information about both minority groups (online Appendix Table S6).

Observation 3: The landlords' decision whether to invite an applicant is responsive to manipulations of the available information about applicants with a minority-sounding name, while the decision is not (or only a little) affected by the same changes in the available information about applicants with a majority-sounding name.

Giving access to a personal website leads to an increase of 8 percentage points in the invitation rate for minority applicants, while it causes a moderate but not statistically significant decrease of 6 percentage points for majority applicants. As a result, the gap in the invitation rate between majority and minority applicants decreases from 37 percentage points in the No Information Treatment to 23 percentage points in the Monitored Information Treatment (column 3 of Table 1) and the difference is statistically significant (column 3, panel A of Table 2).

Next, we compare discrimination among landlords who choose to acquire information with those who don't in the Monitored Information Treatment. Note that since attention is not experimentally manipulated, the difference in decisions between these two groups cannot be interpreted causally because we cannot separate the effect of having more information from the self-selection of certain types of landlords. We find a positive relationship between opening a website and the likelihood of an invitation (column 2, panel B of Table 2) and this relationship is slightly higher for the minority candidates (column 3).²⁶ The landlords who did open an applicant's website discriminate less than those who did not—the gaps in the likelihood of invitation are 18 and 29 percentage points, respectively—although the difference is not statistically significant (column 3–5).

Last, among a group of landlords who chose to uncover the box about education, reporting a college degree increases the invitation rate compared to reporting a high school degree. Interestingly, conditional on paying attention, the effect is similar for minority and majority applicants (columns 6–8 of online Appendix Table S7).²⁷

Other Results.—In order to gain some insight about landlords' priors and to better understand possible motivations for observed differences in attention, we conducted an online survey (Survey I) among 60 landlords. The landlords are drawn from the population of landlords who post rental offers online, but are different from those in our experiment.²⁸ Conditional only on name, we directly elicited the mean expected satisfaction with an applicant (in our model, $q_G - d_G$), the variance of expected satisfaction (σ_G^2) and expected informativeness of an applicants' personal website ($\sigma_{G,1}^2, C_1$), since differences in each of these three parameters across groups are predicted to generate asymmetry in attention. Each landlord was given snapshots of a flat rental advertisement and of an e-mail response used in the experiment. The first question was: "How likely it is that the following applicant would be a tenant

²⁶Notice that from the sign of this coefficient we cannot draw inferences about whether attention helps or hurts candidates on this market in general and compare that with the model's predictions, since here we measure a relationship between attention and selection for a candidate with specific (attractive) attributes and not for the whole population of candidates from a given group.

²⁷We don't find any systematic effect of education level on search patterns (online Appendix Table S8).

²⁸We have asked 817 landlords to participate during the months of January and February 2015. The response rate was 7.3 percent. In total we have 89 observations, since some of the 60 landlords answered questions about two or three applicants with different names.

with whom you would be: 1 = highly dissatisfied, 2 = somewhat dissatisfied, 3 = neutral, 4 = somewhat satisfied, 5 = highly satisfied?” The landlords were asked to allocate ten tokens, each representing a 10 percent probability, to the five options. This allows us to measure mean and variance of overall expected satisfaction at the individual level.²⁹ The second question was: “Imagine you have access to the personal website of the applicant. To what extent do you think the website is informative for evaluating him as a prospective tenant? 1 = very uninformative (I will not learn much about an applicant from reading his website), 2 = somewhat uninformative, 3 = somewhat informative, 4 = very informative (I will get a clear idea about the candidate from reading his website).”

Panel A of online Appendix Table S9 shows the results. When compared to the majority name, both Asian and Roma names significantly reduce the mean of expected satisfaction. In contrast, we find virtually no effect of names on the standard deviation of expected satisfaction as well as the expected informativeness of a personal website. Taken at face value, these results support the interpretation that differences in information acquisition across groups observed in experiments are due to unfavorable preferences or prior means, but not due to greater uncertainty about minority candidates or the expected greater informativeness of their website. Nevertheless, these supporting findings need to be taken cautiously, since such direct questions are more vulnerable to social desirability bias and were answered by a sample of landlords different from the decision makers in the experiment.

In principle, the observed greater inspection of personal websites of minority applicants in the housing market could be due to confirmation bias or due to pure curiosity to read about dissimilar individuals,³⁰ both of which would imply that acquired information should affect selection decisions less for minority candidates. Additional results do not provide support for this interpretation: The correlation between opening a website and inviting is similar or greater for minority applicants compared to majority applicants, and the observed differences in attention across groups mimic observed differences in responsiveness to manipulations of available information.

To summarize the main results in the rental housing market, we find that negatively stereotyped names affect both the choices of whether to invite an applicant for an apartment viewing as well as the attention paid to information prior to this decision. Applicants with minority-sounding names are more thoroughly inspected and less likely to be invited for an apartment viewing. Differences in the observed level of inspection across the groups mimic greater responsiveness of the invitation rate to the manipulation of available information about the quality of applicants with minority-sounding names.

III. Field Experiment in the Labor Market: Czech Republic

The second experiment shifts the exploration of discrimination to the labor market. Here, we aim to study discrimination during a selection process in which

²⁹We deliberately focus on measuring priors about *overall* satisfaction instead of priors about specific attributes, since previous work suggests that stereotypes about a group may vary across different dimensions (Fiske et al. 2002).

³⁰Of the landlords in our sample, 93 percent have a white majority-sounding name.

decision makers pick only a few winners out of a large pool of applications, in contrast to the rental housing market.

A. *Experimental Design*

We use the same names as in the rental housing market experiment to evoke Asian, Roma, and white majority ethnic status. The experiment was implemented between August and October 2012 in the Czech Republic. Over that period we monitored the major online job site (www.jobs.cz) and responded to online job advertisements. We implement the treatment with the monitoring of information acquisition, and send an application via e-mail. The e-mail contains a greeting, the applicant's interest in the job opening, his name, and a hyperlink to his professional resume on a website.³¹

We created a conventional resume, following real-life resumes, and we responded to job offers for which the applicant satisfied all the education and qualification criteria. The resume has six parts: education, experience, skills, hobbies, references, and contacts. Applicants are 30-year-old males, have prior work experience as administrative workers, and we randomly determined (orthogonally to name) whether they obtained a high school or a college degree. They report a good knowledge of English, PC skills, and a driver's license. They also list their hobbies and provide two reference contacts. We believe the resume was roughly comparable to that of a standard applicant for these types of jobs.

When employers open the website, they can see a standard version of the resume. Further, they can click on "learn more" buttons placed below each resume category label (contact, education, experience, skills and hobbies). For example, when the website is accessed it reveals basic information about previous employment experience: the name of the firm, the position held and the time period. By clicking on the "learn more" button below the "Experience" label, the website reveals the applicant's responsibilities (document management, administrative support of consultants, work with PC). Thus, in addition to monitoring whether an employer opens the resume, we measure whether an employer decides to acquire more and which type of information. An example of the shorter as well as the expanded form of the resume is in online Appendix Figure S2.

B. *Sample Selection and Data*

We focused on job openings in sales, customer service, and administrative work. We selected these job categories because they have a sufficient flow of new openings and are similar enough not to require subtle adjustments of particular skills in resumes. In addition, we aimed to minimize the costs for the employers of reviewing the resumes and thus we selected job categories that involve less intensive inspection of applicants compared to higher-skill jobs. We also sent only one e-mail to each employer and politely declined all invitations for job interviews within two days.

³¹ The exact wording was as follows: "Dear Sir/Madam, I am writing because I am very interested in job position advertised by your company. You can find my resume in this hyperlink: phanquyetnguyen1982.sweb.cz. Best regards, Phan Quyet Nguyen."

We target the population of employers who use the Internet to advertise job openings. To be able to monitor the opening of an applicant's resume, we had to exclude ads in which employers did not make their e-mail publicly available and required applicants to call or use an online form (59 percent). Overall, we responded to 274 job openings and to each of them we randomly assigned the name of an applicant. We record the type of job, the job requirements and the time when the application was sent. The means of the observable characteristics of job openings are similar across the three groups of applicants (online Appendix Table S10), with the exception of a somewhat higher likelihood of majority applicants applying for openings that required previous job experience, compared to minority applicants (p -value = 0.13) and the lower likelihood of majority applicants applying during the holiday period (p -value = 0.12). In the analysis, we rely on a comparison of means across treatment conditions, as well as a regression analysis in which we control for observable characteristics.

We study how name manipulation affects two types of choices: attention to a resume and the selection decision. First, we measure whether an employer opened an applicant's resume by clicking on the hyperlink to a resume website. Further, we identify which additional information about an applicant an employer uncovered by clicking on the "learn more" buttons. As was the case in the first experiment, we do not measure the ultimate outcome of a selection process (an actual employment offer and wage). The outcome measure is whether the employer e-mailed or called the applicant back with a decision ("callback") and whether the employer decided to invite the applicant for an interview ("invitation"), a more precise outcome of the initial stage of selection process than callback.³²

C. Results

Are Ethnic Minorities Less Likely to be Invited for a Job Interview?—Panel A of Table 4 documents a large amount of discrimination against minority applicants. The callback rate for majority applicants is 43 percent and only 20 percent for minority applicants, making a difference of more than 100 percent, which is highly significant statistically ($p < 0.01$). A similar picture arises when we turn to the invitation rate. While majority applicants are invited in 14 percent of cases, minority applicants receive an invitation only in 6.3 percent of cases. The gap is statistically significant (p -value = 0.03) and is large in magnitude (133 percent).

While we observe the almost identical treatment of applicants with Roma- and Asian-sounding names in the rental housing market, we find some differences in the labor market. Both minority groups are less likely to be invited for a job interview compared to the majority group. The gap, however, is larger and more

³²Since the application was sent via e-mail, the most common response from employers was also via e-mail: 25.9 percent of employers e-mailed back, 9.1 percent invited the applicant for an interview, and 16.8 percent declined the application. Employers could also call the applicant's cell phone number reported in the resume. However, only a few employers called back (5.8 percent). We recorded "missed calls" on each cell phone and then called back to determine the particular employer. Most of the employers who made a phone call also responded via e-mail, and one employer sent a text message with an invitation; in only six cases we cannot directly identify whether the employer who called back meant to invite the applicant or not. In the main estimations we assume they did not, given the large fraction of declines in the e-mail responses and the fact that these employers did not get in touch with the applicant via e-mail.

TABLE 4—CZECH LABOR MARKET: INVITATION RATES AND INFORMATION ACQUISITION BY ETHNICITY, COMPARISON OF MEANS

	White majority name (W) (1)	Pooled Asian and Roma minority name (E) (2)	Percentage point difference: W – E, (<i>p</i> -value) (3)	Asian minority name (A) (4)	Percentage point difference: W – A, (<i>p</i> -value) (5)	Roma minority name (R) (6)	Percentage point difference: W – R, (<i>p</i> -value) (7)	Percentage point difference: R – A, (<i>p</i> -value) (8)
<i>Panel A. Employer's response</i>								
Callback	0.43	0.20	23 (0.00)	0.17	26 (0.00)	0.25	18 (0.01)	8 (0.22)
Invitation for a job interview	0.14	0.06	8 (0.03)	0.05	9 (0.03)	0.08	6 (0.18)	3 (0.46)
Invitation for a job interview ^a	0.19	0.09	10 (0.06)	0.09	10 (0.12)	0.10	9 (0.16)	1 (0.83)
<i>Panel B. Information acquisition</i>								
Opening applicant's resume	0.63	0.56	7 (0.22)	0.47	16 (0.03)	0.66	-3 (0.69)	19 (0.01)
Acquiring more information about qualification ^a	0.16	0.10	6 (0.27)	0.06	10 (0.12)	0.14	2 (0.73)	8 (0.24)
Acquiring more information about other characteristics ^a	0.18	0.18	0 (0.92)	0.19	-1 (0.85)	0.18	0 (0.99)	1 (0.85)

Notes: Means. Panel A reports how name affects callback and invitation for a job interview and panel B how it affects information acquisition. Columns 3, 5, 7, and 8 report differences in percentage points, in the parentheses we report *p*-value for a *t*-test testing the null hypothesis that the difference is zero. Acquiring more information about qualification is a dummy variable indicating whether an employer clicked on “learn more” buttons on a resume to acquire more information about education, experience, and skills. Acquiring more information about other characteristics is a dummy variable indicating whether she/he acquired more information about hobbies and contact information.

^aThe numbers are reported for the subsample of employers who opened applicant's resume.

significant statistically for the Asian minority applicant (5.1 percent invitation rate, *p*-value = 0.03) than the gap for the Roma minority applicant (7.8 percent invitation rate, *p*-value = 0.18). Put differently, Asian applicants need to send 20 applications to receive one invitation, Roma applicants 12.5, and majority applicants 7.5. Columns 1 and 2 of Table 5 demonstrate the evidence in a regression framework. Controlling for observable job characteristics—required high school education, required previous experience, the type of job, and whether the application was sent during the summer holidays—does not affect the size of the observed gaps in the invitation rates and somewhat increases precision. We also find that employers who decide to read a resume are more likely to invite the applicant and those who request previous job experience are less likely to invite the applicant (available upon request).³³

Observation 4: Applicants with minority-sounding names are discriminated against in the labor market. An applicant with a majority-sounding name is 180 percent more likely to be invited for a job interview compared to an applicant with an Asian-sounding name, and 75 percent more likely compared to an applicant with a Roma-sounding name.

³³It is noteworthy that discrimination is not restricted to jobs where language skills and interactions with customers are central (sales and services), and thus employers could presumably discriminate due to a belief about language use or due to the expected taste-based discrimination of their customers (online Appendix Table S11).

TABLE 5—CZECH LABOR MARKET: INVITATION RATE AND INFORMATION ACQUISITION BY ETHNICITY, REGRESSION ANALYSIS

Dependent variable: Sample:	Invitation for a job interview		Opening applicant's resume		Acquiring more information about qualification		Acquiring more information about other characteristics	
	All		All		Employers who open applicant's resume			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Ethnic minority name	-0.09*** (0.04)		-0.08 (0.06)		-0.07 (0.06)		-0.01 (0.064)	
Asian minority name		-0.08** (0.03)		-0.16** (0.07)		-0.10* (0.05)		-0.00 (0.08)
Roma minority name		-0.06* (0.03)		0.03 (0.08)		-0.03 (0.06)		-0.02 (0.07)
Observations	274	274	274	274	160	160	160	160

Notes: Probit, marginal effects (dF/dx), robust standard errors in parentheses. In columns 5 and 6, the dependent variable is a dummy variable indicating whether an employer clicked on “learn more” buttons on a resume to acquire more information about education, experience, and skills; in columns 7 and 8 it indicates whether she/he acquired more information about hobbies and contact information. In all columns the omitted variable is a white majority name and we control for dummy variables indicating required high school education, required previous experience, application being sent during a holiday period (August), and application in the sector of sales and services.

***Significant at the 1 percent level.

**Significant at the 5 percent level.

*Significant at the 10 percent level.

Additional results suggest that human resource managers do not reward applicants for having a higher level of education than requested in the advertisement. While the applicants reported to have either a high school or a college degree, the positions they applied for requested a high school (80 percent) or lower education level. Specifically, conditional on opening the resume, the invitation rate does not differ for applicants who report having a high school degree and for those who report a college degree, for both minority as well as majority applicants (online Appendix Table S12).

The invitation rate (on average 9.1 percent) in the labor market is much lower than 50 percent, despite the fact that the resume signals the relatively high quality of applicants for the selected job types. If we were to link this observation to theory, it would imply that mean prior beliefs about the quality of all groups are below the threshold necessary for an interview invitation and that the labor market is the “cherry-picking” type of market.³⁴ Resumes of applicants with minority names, and with Asian names in particular, are thus predicted to receive less attention compared to resumes provided by majority applicants.

Do Employers Choose Different Levels of Attention to Information Based on the Ethnicity of an Applicant?—We start by looking at the likelihood of opening a resume. Of employers, 58 percent open the resume. Name again matters. We find

³⁴ Low invitation rates seem to be a ubiquitous feature of labor markets. The invitation rates do not get anywhere close to 50 percent in any segment of the market, which does not allow us to test the theoretical prediction about the “switch” in relative attention between minority and majority groups within one type of market (in contrast to testing it across two different markets, as we do in this paper).

that while 63.3 percent of employers visit the web page with the resume of majority applicants, only 47.5 percent of employers do so when they receive an application with the Asian-minority name. The difference is large in magnitude (34 percent) and significant statistically (panel B of Table 4 and column 4 of Table 5) and it demonstrates that ethnicity signaled by name represents a barrier even at the very start of a selection process, before any information about an applicant is acquired. Moreover, in some firms it is common to delegate the printing of all received applications to an assistant, and printed resumes are then screened and evaluated by a different person. In such cases our experimental design fails to measure differences in attention, biasing down the estimated effect of a name on the likelihood of opening a resume. Regarding Roma-minority applicants, i.e., the group with the invitation rate between the majority and the Asian-minority applicants, we find no discrimination in attention: the likelihood of opening the resume is the same as for the majority applicant and higher compared to the Asian-minority applicant (column 4 of Table 5).

Further, we study whether employers differentiate attention after opening the resume. Overall, we find relatively little interest to expand the resume; only 14 percent of employers in our sample clicked on at least one out of five “learn more” buttons and only 1 percent (three employers) clicked on all buttons. Despite relatively little variation, the results (reported in online Appendix Table S13) reveal that employers were somewhat less active in acquiring information about the Asian-minority applicant: they clicked on a lower number of “learn more” buttons and were less likely to click on all the buttons. This is mostly driven by a reduced interest in getting more detailed information about experience (13 percent for the majority and 4 percent for the Asian candidate) and skills (6 percent and 2 percent). Although these differences are large in magnitude, they are not significant statistically at conventional levels. Taken together, employers paid significantly less attention to the qualification of Asian-minority applicants: when considering majority applicants, 16 percent of employers further inspected at least one out of three categories that seem relevant for assessing qualifications (experience, education, and skills), while only 6 percent made that effort when considering Asian-minority applicants (panel B of Table 4). We find no differences in acquiring information about contacts and hobbies.³⁵

In sum, the Asian minority, i.e., the group with the lowest invitation rate, receives the least attention, which is in line with the predictions of our model of attention discrimination. Since the invitation rate of the Roma-minority applicant is lower than the invitation rate of the majority applicant and higher than the invitation rate of the Asian-minority applicant, the model predicts that the amount of acquired information about the Roma applicant should also be somewhere between the Asian applicant and the white applicant. Nevertheless, we do not observe any differences in attention compared to the majority applicant.

Observation 5: Employers are 34 percent more likely to read a resume provided by applicants with majority-sounding names relative to applicants with Asian-sounding names. Conditional on opening a resume, employers more closely inspect the qualifications of applicants with a majority-sounding name relative to

³⁵The fact that only a few employers used the “learn more” buttons prevents us from analyzing differences in the order in which different types of information were acquired.

applicants with an Asian-sounding name. There is little difference in the likelihood of opening a resume as well as in the depth of resume inspection between applicants with majority- and Roma-sounding names.

Other Results.—In order to get a better sense about priors of employers regarding candidates with the names used in our experiment, we conducted an online survey (Survey II) among 39 human resource managers.³⁶ Specifically, employers were shown a snapshot of a job advertisement (similar to one we responded to in the experiment) and of an e-mail response, including the name of the applicant (but not resume). They were asked similar questions as in Survey I, which were adapted to the labor market context. To proxy mean and variance of expected satisfaction with an applicant, we asked: “How likely it is that the following applicant would be an employee with whom you would be: 1 = highly dissatisfied, 2 = somewhat dissatisfied, 3 = neutral, 4 = somewhat satisfied, 5 = highly satisfied?” Employers were asked to allocate ten tokens, each representing a 10 percent probability, to the five options. The second question serves as a measure of the expected informativeness of the applicant’s resume: “Imagine you have access to a professional resume of the applicant. To what extent do you think the resume is informative for evaluating him as a prospective employee? 1 = very uninformative (I will not learn much about an applicant from reading his resume), 2 = somewhat uninformative, 3 = somewhat informative, 4 = very informative (I will get a clear idea about the candidate from reading his resume).”

Panel B of online Appendix Table S9 shows the results. In terms of mean and variance of expected satisfaction, we find a similar pattern for employers as for landlords. On average, employers expect to be significantly less satisfied with minority compared to majority candidates, while the uncertainty about the candidates, in terms of variance, seems to be very similar across names. It is also noteworthy, in light of the lowest attention and invitation rate observed for the Asian applicant in the experiment, that employers have the worst expectations about this group, although the difference between the Asian and Roma applicants is not significant statistically. At the same time, we find that employers expect the resume of minority applicants to be less informative than that of the majority applicant and virtually no differences in this measure across the two minority groups. These results suggest that the observed discriminatory behavior in the experiment is unlikely to be motivated by a lower variance of beliefs about minorities (“you can’t tell them apart”), but rather by negative stereotypes, preferences, or the lower expected informativeness of a resume.

Although the selected names used in the experiments strongly signal ethnicity, they may also signal some other characteristics, social background in particular. In order to assess whether our findings can be explained by discrimination against individuals with low socioeconomic status (SES), we turn to the survey on perceptions, which we conducted among 92 university students (Survey III). For each name used

³⁶The sample of human resource managers in this survey is different from the sample in our correspondence test. The survey was implemented in February 2015. We have directly invited 913 human resource managers via e-mail (response rate 4.3 percent). In total, we have 90 observations, since some of the 39 human resource managers answered questions about two or three applicants with different names.

in the experiment, we measure associations with level of schooling and quality of housing. We find that the majority-sounding name is perceived as having the highest SES, the Roma-sounding name the lowest SES, and that the Asian-sounding name is in the middle (online Appendix Table S2). If the results in the experiment were fully driven by SES, we would expect the Roma minority to be the most discriminated against, and thus this analysis does not provide strong support for the SES interpretation of the name effects.

Motivated by the findings of previous research which documented differences in callbacks between names associated with the same ethnicity (Jacquemet and Yannelis 2012), in Survey III we also included three other names for each ethnicity besides the names used in the experiments. Within each ethnic group, all majority-sounding names and all Asian-sounding names are perceived very similarly (online Appendix Table S14), suggesting our results are likely to be close to the average discrimination of people with Asian-sounding names. The Roma-sounding name used in our experiment is perceived similarly as one of the three names and as signaling a somewhat lower socioeconomic status compared to the two remaining Roma names, indicating that the observed discrimination might be an overestimate of the average discrimination of the Roma minority.

IV. Field Experiment in the Labor Market: Germany

The third field experiment tests two further considerations: (i) the generalizability of attention discrimination against negatively stereotyped ethnic groups to a different country, Germany, and (ii) generalizability to other types of signals of an applicant's quality beyond ethnicity, such as being unemployed.

A. Experimental Design

We study discrimination against the Turkish minority in the German labor market. Individuals with a Turkish background represent the largest minority in the country (2.9 percent of population). Migrants from Turkey came to Germany predominantly in the 1960s and their children and grandchildren, raised in Germany, now represent a significant share of Germany's labor force. The unemployment rate in 2012 among the majority German population was 6.2 percent, while it was 14.4 percent among immigrants. Importantly, Kaas and Manger (2012) found evidence of discrimination against the Turkish minority by employers. In their experiment a white majority-sounding name increased the likelihood of a callback by 14 percent compared to a Turkish-sounding name. We build on these results by focusing on the effects of a minority name on information acquisition prior to a selection decision to invite an applicant for a job interview or not.

We use the same names as Kaas and Manger (2012) to signal white majority and Turkish ethnicity.³⁷ The experiment was carried out between August and September 2013. We used major online job advertisement sites in Germany and sent e-mail

³⁷The first names and surnames of white majority applicants—Denis Langer and Tobias Hartmann—belong to the 30 most common names in Germany. The names of the ethnic minority applicants—Fatih Yildiz and Serkan Sezer—are very common for male descendants of Turkish immigrants.

applications to 745 online job postings in sectors such as information and communication, administration, health and education, manufacturing, and construction. We responded to all job ads that were posted directly by the company and included an e-mail contact in the text of the posting (66 percent of all ads).

The e-mail contains a greeting, the applicant's interest in the job opening, his name, and a hyperlink to his professional resume on a website. Orthogonally to name treatments, we have implemented three additional conditions by randomly varying the text of the application e-mail. In the baseline condition (50 percent of responses) the text was the same as in the Czech Republic. Next, we have implemented two conditions (25 percent of responses each), in which the text of the application e-mail contains a negative signal about an applicant's quality. Specifically, the e-mail includes the following sentence: "I have been searching for a job for two months [a year and a half]." Both unemployment lengths are common in Germany: 45 percent of the unemployed in 2012 were unemployed for at least one year. The negative signals come from revealing the applicant's potential unemployment to the employer, as well as the carelessness demonstrated in releasing this information in the introductory e-mail. Otherwise, the text is identical as in the baseline condition. Observable characteristics of the job openings vary little across experimental conditions (online Appendix Table S15).

As outcomes of interest, we again focus on measures of information acquisition by the employers. We have altered the design of the experiment to aim at more detailed measures of effort to acquire a resume. The provision of a hyperlink allows us to distinguish whether an employer decides to open the resume. In this experiment clicking on the link does not, however, reveal an applicant's resume but instead a browser displays a message indicating a temporary error on the server.³⁸ We measure whether an employer clicked on the link, the number of times the employer attempted to reopen the resume and the likelihood that the employer sent an e-mail requesting the resume to be re-sent, an action that requires nontrivial effort. Since the gap in terms of callback has already been established in previous work (Kaas and Manger 2012), we have not proceeded by responding back with a resume to minimize the costs on the part of employers, and we focus purely on information acquisition.

B. Results

We find that a minority name reduces employer's effort to acquire information about an applicant compared to a majority applicant. This effect holds for all three measures of information acquisition: likelihood of opening an applicant's resume (Table 6 and column 1, panel A of Table 7), number of attempts to open an applicant's resume (column 3) and a likelihood of writing back requesting an applicant to resend the resume (column 5).

³⁸The message is "Database connection error (2): Could not connect to MySQL Server!" Such a message commonly appears when announcing a failure to connect to the database server due to technical problems on the provider's side and thus should not indicate a mistake on the part of the applicant. Nevertheless, we cannot rule out that the failure to open the resume was perceived by some employers as an applicant's mistake and thus a negative signal about his quality.

TABLE 6—GERMAN LABOR MARKET: INFORMATION ACQUISITION BY ETHNICITY, COMPARISON OF MEANS

	(1)	(2)	(3)	(4)	(5)
<i>Panel A. Effect of name on information acquisition</i>					
	White majority name (W) (<i>n</i> = 366)	Ethnic minority name (E) (<i>n</i> = 379)	Percentage point difference: W – E (<i>p</i> -value)		
Opening applicant's resume	0.81	0.75	6 (0.06)		
Number of attempts to open applicant's resume	2.81	2.14	0.67 (0.00)		
E-mail request to re-send resume	0.31	0.19	13 (0.00)		
<i>Panel B. Effect of signal about unemployment on information acquisition</i>					
	No information (N) (<i>n</i> = 372)	2 months unemployed (2M) (<i>n</i> = 187)	18 months unemployed (18M) (<i>n</i> = 186)	Percentage point difference: N – 2M (<i>p</i> -value)	Percentage point difference: N – 18M (<i>p</i> -value)
Opening applicant's resume	0.83	0.73	0.73	10 (0.01)	10 (0.00)
Number of attempts to open applicant's resume	2.66	2.51	2.04	0.15 (0.56)	0.63 (0.01)
E-mail request to re-send resume	0.27	0.26	0.18	2 (0.66)	9 (0.02)

Notes: Means. Panel A reports how information acquisition is affected by name and panel B how it is affected by the signal about recent unemployment. In column 3 of panel A we report differences in means between white majority and ethnic minority group, in the parentheses we report *p*-value for a *t*-test testing the null hypothesis that the difference is zero. The differences in the number of attempts to open applicant's resume are reported in absolute terms, not in percentage points.

Observation 6: Applications with Turkish minority names receive lower attention in all three measures than applications with majority names in the German labor market.

Interestingly, the gap is greater at higher levels of effort. The likelihood of clicking on the resume link at least once is 75 percent for minority applicants and it increases by 8 percent for majority applicants (to 81 percent). For the number of times an employer tried to open the resume the difference is 31 percent (2.1 for minority and 2.8 for majority applicants). Finally, the magnitude of the difference in whether the employer sent an e-mail and requested the resume to be resent is 68 percent (19 percent for minority and 31 percent for majority applicants).

Next, we explore the effect of signaling recent unemployment on attention, i.e., the effect of an unambiguously negative signal about the applicant's quality. We find that employers consistently adjust their attention based on such information. Compared to the baseline condition with no information about unemployment history, indicating an 18-months-unemployment significantly reduces all three measures of information-acquisition effort (Table 6 and panel A of Table 7).

Observation 7: A signal of an applicant's lower quality, observed by an employer prior to reading a resume reduces an employers' attention to the resume.

Similarly as in the case of name effects, we find the magnitude of the gap in attention increases at higher levels of effort to acquire a resume. The difference between the baseline condition and the 18-month unemployment condition is 13 percent for

TABLE 7—GERMAN LABOR MARKET: INFORMATION ACQUISITION BY ETHNICITY, REGRESSION ANALYSIS

Dependent variable:	Opening applicant's resume		Number of attempts to open applicant's resume		E-mail request to re-send resume	
	(1)	(2)	(3)	(4)	(5)	(6)
<i>Panel A. Sample: All</i>						
Ethnic minority name	-0.06*	-0.02	-0.68***	-0.82***	-0.15***	-0.14***
	(0.03)	(0.05)	(0.20)	(0.28)	(0.03)	(0.04)
2M unemployed	-0.10**	-0.04	-0.10	-0.09	-0.01	0.02
	(0.04)	(0.06)	(0.24)	(0.35)	(0.04)	(0.05)
18M unemployed	-0.12***	-0.10*	-0.64***	-0.92***	-0.07**	-0.08*
	(0.04)	(0.06)	(0.24)	(0.34)	(0.03)	(0.04)
Ethnic minority name × 2M unemployed		-0.11		-0.01		-0.06
		(0.09)		(0.48)		(0.06)
Ethnic minority name × 18M unemployed		-0.03		0.57		0.01
		(0.08)		(0.49)		(0.08)
Observations	745	745	745	745	745	745
<i>Panel B. Sample: Two months unemployed and 18 months unemployed</i>						
Ethnic minority name	-0.10**	-0.15**	-0.51*	-0.85**	-0.16***	-0.19***
	(0.05)	(0.07)	(0.29)	(0.40)	(0.04)	(0.05)
18M unemployed	-0.01	-0.06	-0.55*	-0.91**	-0.07*	-0.10**
	(0.05)	(0.07)	(0.28)	(0.41)	(0.04)	(0.05)
Ethnic minority name × 18M unemployed		0.09		0.70		0.09
		(0.08)		(0.57)		(0.09)
Observations	373	373	373	373	373	373

Notes: Probit, marginal effects (dF/dx) in columns 1, 2, 5, and 6, OLS in columns 3 and 4. Robust standard errors in parentheses. In all columns of both panels, the omitted variable is a white majority name and we control for dummy variables indicating required high school education, required previous experience, position in a city with more than a million inhabitants, application being sent in holiday period (August), and a set of four dummy variables indicating the sector (manufacturing and construction, information and communication, administration, and professional, scientific and technical activities). In all columns of panel A, the omitted variable is no information about unemployment, while in all columns of panel B it is two months unemployed.

***Significant at the 1 percent level.

**Significant at the 5 percent level.

*Significant at the 10 percent level.

the likelihood of opening a resume, 23 percent for the number of attempts to open a resume, and 33 percent for the likelihood of requesting the resume to be re-sent. The effect of the two-month unemployment condition is generally smaller in size and less significant statistically (panel A of Table 7). Specifically, the difference between the two-month and 18-month unemployment conditions is statistically significant for the number of attempts to open a resume and for the likelihood of requesting the resume to be re-sent (panel B).

Last, we do not find evidence of a systematic interaction effect of minority names and unemployment conditions on attention (columns 2, 4, and 6, panel A of Table 7). In other words, the 18-month unemployment condition lowers an employer's effort to read a resume for both minority as well as majority applicants. Also, in the baseline condition, which is most comparable to the labor market experiment in the Czech Republic, the minority name significantly lowers the number of clicks on the resume as well as the likelihood of requesting the resume to be re-sent. The negative effect is small and statistically insignificant for the likelihood of opening a resume.

V. Links to Theories

We now consider which models can explain the set of findings from the three correspondence tests. Although it is likely that the observed discrimination in terms of invitation rates arises, at least in part, due to reasons highlighted in standard economic models of discrimination—preference-based and statistical discrimination models—these models cannot explain the complete set of findings, in particular the observed discrimination in attention, an important input for selection decision.

In purely preference-based models of discrimination, individuals do not discriminate due to lack of information and thus imperfect information and attention do not enter the model. There are several models that generate discrimination via imperfect information. Their common feature is that all observable actions prior to selection decisions are the same and discrimination arises at the moment of selection decision when the imperfect information is used. In the first class of statistical discrimination models decision makers take into account observable individual characteristics, while using an observable group attribute, such as ethnicity, to proxy unobservable individual quality (Phelps 1972; Arrow 1973).

The second class of statistical discrimination models emphasizes a lower precision of observable signals as a source of discrimination (Aigner and Cain 1977). Specifically, signals about individuals that economic agents receive are more precise for majority applicants compared to minority applicants, perhaps due to cultural dissimilarity (Cornell and Welch 1996). Thus, the difference in the precision of information about individuals across ethnic groups is assumed, i.e., is exogenous, and is not due to differences in efforts to acquire information. In contrast to these models, the experiments reveal that discrimination begins earlier, already during the information-acquisition stage, creating differences in information imperfection across groups at the moment when agents finally make decisions.

Do decision makers allocate costly attention endogenously? To assess that, we consider whether it can explain the main findings: (i) In the labor market employers pay more attention to majority compared to minority candidates, while in the rental housing market landlords pay more attention to minority compared to majority candidates. (ii) The gap in resume acquisition is greater when acquiring a resume requires writing an e-mail requesting resending a resume compared to simple clicking on a hyperlink. (iii) Signaling recent unemployment—another type of negative signal besides a minority name—lowers attention to an applicant on the labor market, similarly as minority name does.

First, the model predicts a switch in relative attention if markets differ in selectivity. In the labor market, where selectivity is high—since firms select only a few top applicants for an interview—the expected benefits from reading a resume are smallest for the a priori least attractive group, while the benefits of inspection are greater for this group in the housing market, where the overall invitation rate is high. Second, it predicts that at low costs of information acquisition the increased cost increases the motivation of the decision maker to optimize attention and thus the gap in resume acquisition is predicted to increase when an employer needs to write an e-mail compared to clicking on a hyperlink in order to get a resume. Third, it implies that any signal of an applicant's quality should affect attention to subsequent information, independently of whether the signal concerns ethnic status or

some other characteristic relevant for beliefs about quality, such as signaling recent unemployment. Thus, we conclude the predictions of the model are consistent with all three empirical findings.

Note that the switch in relative attention across markets is predicted to arise if decision makers have racist preferences, believe that minority candidates are of lower quality on average, or expect members of a minority group to be more alike. Therefore, based on documenting the switch across markets, we still cannot distinguish between these sources. Indicative evidence on this question is provided by two supplementary online surveys, in which we find that landlords as well as employers expect to be less satisfied with minority candidates, indicating either dislike or belief about a lower objective quality. At the same time, we find virtually no differences in the variance of expected satisfaction across ethnic groups.

It should be noted that the selection of tenants by landlords differs from the hiring decisions of employers in many ways, and thus attributing the switching results to differences in selectivity needs to be taken cautiously. For example, desired applicant's qualities may differ across the markets—landlords may be concerned about a tenant's ability to reliably pay rent and not cause property damage, while employers may focus on the type of education and qualifications relevant for a given job. If minority applicants were considered a priori better tenants and worse employees, then such a combination of beliefs could, in principle, explain the observed switch. However, this interpretation is inconsistent with the observed discrimination in selection decisions as well as results of online surveys, which suggest that minorities are negatively stereotyped in both markets.

Similarly, the expectation of the relative precision of available signals (Aigner and Cain 1977; Cornell and Welch 1996) may also differ across markets. Our model predicts more attention to groups with more precise signals since the benefit of paying a unit of attention to such groups is higher. The switch in relative attention could then be explained if the expected precision of the available signals also switched, i.e., if employers expected signals about minorities to be less informative, while landlords expected the same about majority applicants. We asked employers and landlords about these types of expectations in online surveys. Their responses support the former but not the latter.

Next, it is also possible that knowing more about minority applicants may be more important for a landlord than for an HR manager, perhaps because landlords may be more likely to interact intensively with tenants than HR managers with employees. However, the importance of the decision is predicted to affect overall levels of attention and thus influence the magnitude of attention discrimination, but it is not predicted to lead to a switch in relative attention across groups. Potentially, there might be other differences across markets, which could explain the switch, although the explanation based on differences in selectivity seems to be the most parsimonious.

While we propose a model in which decision makers consciously allocate attention based on expected benefits in each instance, the observed attention choices of employers and landlords in the experiments may be conscious as well as automatic based on simplifying screening heuristics. Our model can help to explain why discrimination heuristics may arise, for example by trial and error or by an initial conscious setting of screening rules that prove to work reasonably well and are later used automatically. Nevertheless, attention allocation can also be affected

by unconscious mental associations against negatively-stereotyped groups, termed implicit discrimination (Bertrand, Chugh, and Mullainathan 2005) and supported by intriguing evidence from implicit association tests (Greenwald, McGhee, and Schwartz 1998; Stanley, Phelps, and Banaji 2008). Such unconscious biases, which may operate in parallel with—and sometimes in contradiction of—one’s conscious intentions, could also explain the lower observed attention to minorities on labor markets.

VI. Concluding Remarks

One of the main insights from information economics is that even very small frictions in information acquisition can have large effects on economic outcomes (Diamond 1971; Sims 2003). At the same time, imperfect information is central to explaining discrimination in markets since the seminal work of Phelps (1972) and Arrow (1973). Yet, there is no theory or direct evidence studying how the small costs of information acquisition may create differences in the form of imperfect information about individuals based on their observable group attributes. This is what we provide.

We first describe how choices of attention affect discrimination in theory. We show that if attention is costly, prior beliefs about ethnic groups enter the final decision not only through Bayesian updating, as in the standard model of statistical discrimination, but also earlier through the choice of attention to available information (“attention discrimination”). As a result, prior beliefs have the potential for a larger impact on discrimination (in most types of markets) and discrimination in the selection of applicants can arise even when the decision makers have the same preferences across different groups, when all the relevant information is available, and when obtaining information about different groups is equally difficult. Costly attention can similarly magnify the role of animus.

In the empirical part, we identify attention discrimination in practice. We develop new tools for field experiments using the Internet, and monitor information acquisition by employers and landlords about applicants prior to a selection decision for a job interview and an apartment viewing. A set of three experiments in two countries reveals that signals of an applicant’s minority status systematically affect attention to easily available information about the applicant (e.g., resume). In line with the model, the observed patterns of attention allocation are consistent with economic agents considering reading applications for a job or apartment as a costly activity and choosing the level of attention with an eye for expected benefits of reading, taking into account an applicant’s observable group attribute and desired level of quality in a given market. We also discuss alternative interpretations, in particular the potential role of an unconscious bias in attention.

The key insight that willingness to process information at hand represents an additional barrier for applicants with unfavorable group attributes points toward several promising directions for future research as well as thoughts about policy, and we mention a few. The model implies the important role of the timing of when a group attribute is revealed—the later a decision maker learns a group attribute, such as name, the smaller the asymmetry in attention to subsequent information such as education or qualification. It is intriguing that employers in the public as well

as private sector have recently started to introduce name-blind resumes,³⁹ in part because researchers produced evidence indicating that blind auditioning (Goldin and Rouse 2000) and name-blind resumes (Åslund and Skans 2012) can reduce discrimination. Understanding practical implications of attention discrimination and which policies may be the most appropriate to attenuate it, without imposing too many restrictions on a firm's choices, is an important area to explore. The idea that early signals have disproportionately large effects on outcomes also has implications for members of negatively stereotyped groups who cannot take for granted that employers will learn their qualities when reading a resume. It might help to provide positive signals early on, for example, by mentioning previous relevant job experience already in the introductory e-mail. Although such information does not reveal anything new since it is fully contained in the resume, it may prevent the employers from putting the resume aside.

Next, the lower predicted and observed attention to negatively stereotyped groups in selective markets can help explain why African Americans and minorities were found to face lower returns to higher quality resumes in the labor markets in the United States and Sweden, respectively (Bertrand and Mullainathan 2004, Bursell 2007). Based on this, we speculate that in the long-term endogenous attention lowers incentives of negatively stereotyped groups to acquire human capital in the first place, and could make beliefs about differences in quality potentially self-fulfilling. Last, if the effect of recent unemployment on the attention of employers is similar when unemployment is signaled in the introductory e-mail, as in our experiment, just as when it is reported in the resume, then endogenous attention may also contribute to greater long-term unemployment.

In addition to presenting novel empirical findings, the experimental design distinguishes itself by offering a methodological contribution. Our analysis joins efforts in laboratory settings to test decision-making processes with enhanced measurement tools, in particular by monitoring information acquisition (Camerer et al. 1993; Costa-Gomes, Crawford, and Broseta 2001; Gabaix et al. 2006). We show that the widespread use of the Internet by economic decision makers opens the possibility of collecting "process data" as a part of a natural field experiment as well. By this, researchers can study in greater detail the processes taking place inside the "black box" and can better inform theories and policy-makers on issues, including those that are sensitive and hard to study in the laboratory (Levitt and List 2007), of which discrimination is one important example.

REFERENCES

- Ahmed, Ali M., and Mats Hammarstedt. 2008. "Discrimination in the Rental Housing Market: A Field Experiment on the Internet." *Journal of Urban Economics* 64 (2): 362–72.
- Aigner, Dennis J., and Glen G. Cain. 1977. "Statistical Theories of Discrimination in Labor Markets." *Industrial and Labor Relations Review* 30 (2): 175–87.
- Altonji, Joseph G., and Rebecca M. Blank. 1999. "Race and Gender in the Labor Market." In *Handbook of Labor Economics*, Vol. 3C, edited by Orley C. Ashenfelter and David Card, 3143–3259. Amsterdam: Elsevier North Holland.

³⁹Name-blind resumes have recently been implemented for hiring workers in the public sector in Belgium, the Netherlands, and Sweden. The policy is being piloted in Germany among several major companies, including Deutsche Post, Deutsche Telekom, L'Oréal, and Procter and Gamble.

- Arieli, Amos, Yaniv Ben-Ami, and Ariel Rubinstein. 2011. "Tracking Decision Makers under Uncertainty." *American Economic Journal: Microeconomics* 3 (4): 68–76.
- Arrow, Kenneth J. 1973. "The Theory of Discrimination." In *Discrimination in Labor Markets*, edited by Orley Ashenfelter and Albert Rees, 3–33. Princeton, NJ: Princeton University Press.
- Åslund, Olof, and Oskar Nordström Skans. 2012. "Do Anonymous Job Application Procedures Level the Playing Field?" *Industrial and Labor Relations Review* 65 (1): 82–107.
- Banerjee, Abhijit, Marianne Bertrand, Saugato Datta, and Sendhil Mullainathan. 2009. "Labor Market Discrimination in Delhi: Evidence from a Field Experiment." *Journal of Comparative Economics* 37 (1): 14–27.
- Barany, Zoltan. 2002. *The East European Gypsies: Regime Change, Marginality, and Ethnopolitics*. Cambridge, UK: Cambridge University Press.
- Bartoš, Vojtěch, Michal Bauer, Julie Chytilová, and Filip Matějka. 2016. "Attention Discrimination: Theory and Field Experiments with Monitoring Information Acquisition: Dataset." *American Economic Review*. <http://dx.doi.org/10.1257/aer.20140571>.
- Becker, Gary S. 1971. *The Economics of Discrimination*. 2nd ed. Chicago: University of Chicago Press.
- Bertrand, Marianne, Dolly Chugh, and Sendhil Mullainathan. 2005. "Implicit Discrimination." *American Economic Review* 95 (2): 94–98.
- Bertrand, Marianne, and Sendhil Mullainathan. 2004. "Are Emily and Greg More Employable than Lakisha and Jamal? A Field Experiment on Labor Market Discrimination." *American Economic Review* 94 (4): 991–1013.
- Bhatt, Meghana, and Colin F. Camerer. 2005. "Self-Referential Thinking and Equilibrium as States of Mind in Games: MRI Evidence." *Games and Economic Behavior* 52 (2): 424–59.
- Bordalo, Pedro, Katherine Coffman, Nicola Gennaioli, and Andrei Shleifer. Forthcoming. "Stereotypes." *Quarterly Journal of Economics*.
- Bose, Gautam, and Kevin Lang. 2015. "Monitoring for Worker Quality: Task Assignment, Job Ladders, Wages and Mobility in Nonmanagerial Internal Labor Markets." <http://people.bu.edu/lang/monitor.pdf> (accessed July 13, 2015).
- Brocas, Isabelle, Juan D. Carrillo, Stephanie W. Wang, and Colin F. Camerer. 2014. "Imperfect Choice or Imperfect Attention? Understanding Strategic Thinking in Private Information Games." *Review of Economic Studies* 81 (3): 944–70.
- Bursell, Moa. 2007. "What's in a Name? A Field Experiment Test for the Existence of Ethnic Discrimination in the Hiring Process." Stockholm University Linnaeus Center for Integration Studies Working Paper 2007:7.
- Camerer, Colin F., and Eric J. Johnson. 2004. "Thinking about Attention in Games: Backward and Forward Induction." In *The Psychology of Economic Decisions*, Vol. 2, edited by Isabelle Brocas and Juan D. Carrillo, 111–30. Oxford, UK: Oxford University Press.
- Camerer, Colin F., Eric J. Johnson, Talia Rymon, and Sankar Sen. 1993. "Cognition and Framing in Sequential Bargaining for Gains and Losses." In *Frontiers of Game Theory*, edited by Ken Binmore, Alan Kirman, and Piero Tani, 27–48. Cambridge, MA: MIT Press.
- Caplin, Andrew, and Mark Dean. 2015. "Revealed Preference, Rational Inattention, and Costly Information Acquisition." *American Economic Review* 105 (7): 2183–2203.
- Chetty, Raj, Adam Looney, and Kory Kroft. 2009. "Salience and Taxation: Theory and Evidence." *American Economic Review* 99 (4): 1145–77.
- Coate, Stephen, and Glenn C. Loury. 1993. "Will Affirmative-Action Policies Eliminate Negative Stereotypes?" *American Economic Review* 83 (5): 1220–40.
- Cornell, Bradford, and Ivo Welch. 1996. "Culture, Information, and Screening Discrimination." *Journal of Political Economy* 104 (3): 542–71.
- Costa-Gomes, Miguel A., and Vincent P. Crawford. 2006. "Cognition and Behavior in Two-Person Guessing Games: An Experimental Study." *American Economic Review* 96 (5): 1737–68.
- Costa-Gomes, Miguel, Vincent P. Crawford, and Bruno Broseta. 2001. "Cognition and Behavior in Normal-Form Games: An Experimental Study." *Econometrica* 69 (5): 1193–1235.
- Crawford, Vincent P. 2008. "Look-Ups as the Windows of the Strategic Soul." In *The Foundations of Positive and Normative Economics: A Handbook*, edited by Andrew Caplin and Andrew Schotter, 249–280. Oxford, UK: Oxford University Press.
- Czech Statistical Office. 2011. *Census of the Czech Republic*. <https://www.czso.cz/csu/sldb/home> (accessed July 13, 2015).
- Dechief, Diane, and Philip Oreopoulos. 2012. "Why Do Some Employers Prefer to Interview Matthew but Not Samir? New Evidence from Toronto, Montreal and Vancouver." Canadian Labour Market and Skills Researcher Network Working Paper 95.
- Deresiewicz, William. 2014. *Excellent Sheep: The Miseducation of the American Elite and the Way to a Meaningful Life*. New York: Free Press.

- Devetag, Giovanna, Sibilla Di Guida, and Luca Polonio.** 2016. "An Eye-Tracking Study of Feature-Based Choice in One-Shot Games." *Experimental Economics* 19 (1): 177–201.
- Diamond, Peter A.** 1971. "A Model of Price Adjustment." *Journal of Economic Theory* 3 (2): 156–68.
- Economist.** 2012. "Ask the Expert: How to Write a CV." *The Economist*. <http://www.economist.com/node/21559508> (accessed July 13, 2015).
- European Commission.** 2010. *The Social and Economic Integration of the Roma in Europe*. <http://eur-lex.eu/legal-content/EN/TXT/PDF/?uri=CELEX:52010DC0133&from=EN> (accessed July 13, 2015).
- Fehr, Ernst, and Antonio Rangel.** 2011. "Neuroeconomic Foundations of Economic Choice: Recent Advances." *Journal of Economic Perspectives* 25 (4): 3–30.
- Fiske, Susan T., Amy J. C. Cuddy, Peter Glick, and Jun Xu.** 2002. "A Model of (often Mixed) Stereotype Content: Competence and Warmth Respectively Follow From Perceived Status and Competition." *Journal of Personality and Social Psychology* 82 (6): 878–902.
- FRA & UNDP.** 2012. *The Situation of Roma in 11 EU Member States: Survey Results at a Glance*. http://fra.europa.eu/sites/default/files/fra_uploads/2099-FRA-2012-Roma-at-a-glance_EN.pdf (accessed July 13, 2015).
- Gabaix, Xavier, David Laibson, Guillermo Moloche, and Stephen Weinberg.** 2006. "Costly Information Acquisition: Experimental Analysis of a Boundedly Rational Model." *American Economic Review* 96 (4): 1043–68.
- Gneezy, Uri, John List, and Michael K. Price.** 2012. "Toward an Understanding of Why People Discriminate: Evidence from a Series of Natural Field Experiments." National Bureau of Economic Research Working Paper 17855.
- Goldin, Claudia, and Cecilia Rouse.** 2000. "Orchestrating Impartiality: The Impact of 'Blind' Auditions on Female Musicians." *American Economic Review* 90 (4): 715–41.
- Greenwald, Anthony G., Debbie E. McGhee, and Jordan L. K. Schwartz.** 1998. "Measuring Individual Differences in Implicit Cognition: The Implicit Association Test." *Journal of Personality and Social Psychology* 74 (6): 1464–80.
- Hare, Todd A., Jonathan Malmaud, and Antonio Rangel.** 2011. "Focusing Attention on the Health Aspects of Food Changes Value Signals in vmPFC and Improves Dietary Choice." *Journal of Neuroscience* 31 (30): 11077–87.
- Jacquemet, Nicolas, and Constantine Yannelis.** 2012. "Indiscriminate Discrimination: A Correspondence Test for Ethnic Homophily in the Chicago Labor Market." *Labour Economics* 19 (6): 824–32.
- Kaas, Leo, and Christian Manger.** 2012. "Ethnic Discrimination in Germany's Labour Market: A Field Experiment." *German Economic Review* 13 (1): 1–20.
- Kahneman, Daniel.** 1973. *Attention and Effort*. Englewood Cliffs, NJ: Prentice Hall.
- Knoepfle, Daniel T., Joseph Tao-yi Wang, and Colin F. Camerer.** 2009. "Studying Learning in Games Using Eye-Tracking." *Journal of the European Economic Association* 7 (2-3): 388–98.
- Krajbich, Ian, Carrie Armel, and Antonio Rangel.** 2010. "Visual Fixations and the Computation and Comparison of Value in Simple Choice." *Nature Neuroscience* 13 (10): 1292–98.
- Lang, Kevin, and Jee-Yeon K. Lehmann.** 2012. "Racial Discrimination in the Labor Market: Theory and Empirics." *Journal of Economic Literature* 50 (4): 959–1006.
- Lang, Kevin.** 1986. "A Language Theory of Discrimination." *Quarterly Journal of Economics* 101 (2): 363–82.
- Levitt, Steven D., and John A. List.** 2007. "What Do Laboratory Experiments Measuring Social Preferences Reveal about the Real World?" *Journal of Economic Perspectives* 21 (2): 153–74.
- List, John A.** 2004. "The Nature and Extent of Discrimination in the Marketplace: Evidence from the Field." *Quarterly Journal of Economics* 119 (1): 49–89.
- List, John A., and Imran Rasul.** 2011. "Field Experiments in Labor Economics." In *Handbook of Labor Economics*, Vol. 4A, edited by Orley Ashenfelter and David Card, 103–228. Amsterdam: Elsevier North Holland.
- Mackowiak, Bartosz, and Mirko Wiederholt.** 2009. "Optimal Sticky Prices under Rational Inattention." *American Economic Review* 99 (3): 769–803.
- Matějka, Filip, and Alisdair McKay.** 2015. "Rational Inattention to Discrete Choices: A New Foundation for the Multinomial Logit Model." *American Economic Review* 105 (1): 272–98.
- Matějka, Filip, and Christopher A. Sims.** 2011. "Discrete Actions in Information-Constrained Tracking Problems." CERGE-EI Working Paper 441.
- Milkman, Katherine L., Modupe Akinola, and Dolly Chugh.** 2012. "Temporal Distance and Discrimination: An Audit Study in Academia." *Psychological Science* 23 (7): 710–17.
- Neumark, David, Roy J. Bank, and Kyle D. van Nort.** 1996. "Sex Discrimination in Restaurant Hiring: An Audit Study." *Quarterly Journal of Economics* 111 (3): 915–41.

- Newell, Alan, J. C. Shaw, and Herbert A. Simon.** 1958. "Elements of a Theory of Human Problem Solving." *Psychological Review* 65 (3): 151–66.
- Pager, Devah.** 2007. "The Use of Field Experiments for Studies of Employment Discrimination: Contributions, Critiques, and Directions for the Future." *The Annals of the American Academy of Political and Social Science* 609 (1): 104–33.
- Phelps, Edmund S.** 1972. "The Statistical Theory of Racism and Sexism." *American Economic Review* 62 (4): 659–61.
- Reutskaja, Elena, Rosemarie Nagel, Colin F. Camerer, and Antonio Rangel.** 2011. "Search Dynamics in Consumer Choice under Time Pressure: An Eye-Tracking Study." *American Economic Review* 101 (2): 900–26.
- Riach, P. A., and J. Rich.** 2002. "Field Experiments of Discrimination in the Market Place." *Economic Journal* 112 (483): F480–518.
- Sims, Christopher A.** 2003. "Implications of Rational Inattention." *Journal of Monetary Economics* 50 (3): 665–90.
- Spaan, Ernst, Felicitas Hillmann, and Ton van Naerssen.** 2005. *Asian Migrants and European Labour Markets: Patterns and Processes of Immigrant Labour Market Insertion in Europe*. New York: Routledge.
- Stanley, Damian, Elizabeth Phelps, and Mahzarin Banaji.** 2008. "The Neural Basis of Implicit Attitudes." *Current Directions in Psychological Science* 17 (2): 164–70.
- Stevens, Mitchell L.** 2009. *Creating a Class: College Admissions and the Education of Elites*. Cambridge, MA: Harvard University Press.
- TheLadders.** 2012. "Eye Tracking Online Metacognition: Cognitive Complexity and Recruiter Decision Making." <http://cdn.theladders.net/static/images/basicSite/pdfs/TheLadders-EyeTracking-StudyC2.pdf> (accessed July 13, 2015).
- Van Nieuwerburgh, Stijn, and Laura Veldkamp.** 2010. "Information Acquisition and Under-Diversification." *Review of Economic Studies* 77 (2): 779–805.
- Weichselbaumer, Doris.** 2003. "Sexual Orientation Discrimination in Hiring." *Labour Economics* 10 (6): 629–42.
- Woodford, Michael.** 2009. "Information-Constrained State-Dependent Pricing." *Journal of Monetary Economics* 56 (15): S100–24.
- Yinger, John.** 1998. "Evidence on Discrimination in Consumer Markets." *Journal of Economic Perspectives* 12 (2): 23–40.