

Why Do We Discriminate? The Role of Motivated Reasoning*

Markus Eyting[†]

First version: June, 2022

Last updated in April, 2024

Identifying the cause of discrimination is crucial to design effective policies and to understand discrimination dynamics. This paper links existing models of taste-based and statistical discrimination by studying motivated reasoning as a mechanism for generating discrimination. By selectively acquiring and processing information, individuals form motivated beliefs and consequentially discriminate based on these beliefs. Through a series of experiments in which individuals endogenously acquire and process various signals, I document discrimination stemming from motivated beliefs and demonstrate important differences in how it manifests compared to taste-based and statistical discrimination. Finally, I show that limiting individuals' scope to interpret information reduces such discrimination.

Keywords: discrimination; motivated reasoning; belief formation; information

*I thank Michael Kosfeld, Florian Hett, Aislinn Bohren, Alex Imas, Sendhil Mullainathan, Emir Kamenica, George Loewenstein, Vojtěch Bartoš, Joël J. van der Weele, Theo Offerman, Jeroen van de Ven, Guido Friebel, Navid Sabet, Jan Hausfeld, Simon Heß, Johannes Kasinger, Patrick Schmidt, and seminar and conference audiences at the JGU Mainz, Goethe University Frankfurt, University Chicago, Discrimination & Diversity workshop at University of East Anglia, ESA Bologna, Esa Lyon, EEA Barcelona, Behavioral Measurements conference Frankfurt, Discrimination and Disparities online Seminar, IMEBESS Lisbon, EGPROC at UvA, and the Identity workshop at University of Konstanz for many valuable comments and discussions. I gratefully acknowledge the financial support from the SAFE Experiment Center Frankfurt and the Interdisciplinary Public Policy Mainz.

[†]JGU Mainz, Heidelberg and Stanford University; Email: meyting@stanford.edu

1 Introduction

Discrimination, i.e. the systematically disparate treatment of individuals from different social groups, is widespread and has been documented in various contexts, such as the labor market (Blau and Kahn, 2017; Neumark, 2018), healthcare (Alsan et al., 2019), the justice system (Arnold et al., 2018; Arnold et al., 2021), policing (Goncalves and Mello, 2021), or education (Farkas, 2003). While discrimination from any source can have serious negative consequences for the discriminated, precise identification of the source of discrimination has important implications for policy, welfare analyses, and discrimination dynamics.

Traditionally, economists have categorized discrimination as either taste-based discrimination (Becker, 1957) or statistical discrimination (Phelps, 1972; Arrow, 1973; Aigner and Cain, 1977). Taste-based discrimination arises if an individual experiences an animus towards members of a particular group and therefore discriminates against them to not experience a disutility from interacting with them. Resulting discrimination is hence a direct consequence of the discriminator’s taste and attempts to alleviate this form of discrimination include normative approaches such as blaming or sanctioning (Rathelot and Safi, 2022) or increasing social interactions between groups (Pettigrew and Tropp, 2006). In contrast, statistical discrimination is based on beliefs. When there is uncertainty about focal attributes, evaluators need to rely on their beliefs about group-level statistics to make inference about those attributes. If these beliefs vary systematically across groups, discrimination may arise. Typically, statistical discrimination considers the case where beliefs are accurate and hence beliefs about differences reflect actual differences. However, beliefs that generate discrimination may be inaccurate and resulting discrimination is described as inaccurate statistical discrimination. Correcting these inaccurate group-level beliefs by providing individuals with the true statistics is one potential intervention to reduce this form of discrimination (Bohren et al., 2023).

In this paper, I provide a link between these two long-standing sources of discrimination. I show that individuals hold systematically inaccurate beliefs about the productivity of different social groups, that the formation of these incorrect beliefs is driven by individual motives, and that discrimination is a consequence of these motivated beliefs rather than a direct consequence of underlying tastes.¹ I study

¹Section 4.1 discusses my view on ‘motives’ and its relation to ‘taste’ in more detail.

the behavioral operationalization of the formation of motivated beliefs and demonstrate that individuals exploit ambiguity in information (e.g., if the information is about group-level statistics rather than individual performance, or if the veracity of the information source is unclear) to maintain their prior beliefs. Finally, I show that discrimination based on motivated reasoning can still effectively be mitigated by carefully designed information interventions that limit individuals' wiggle room to systematically interpret information and uphold their motivated beliefs. This contrasts with taste-based discrimination, which is not impacted by information. Consequently, this paper provides a micro-foundation for the formation and persistence of incorrect beliefs that drive a particular form of inaccurate statistical discrimination – *discrimination based on motivated reasoning*.

In a series of four fictitious online hiring experiments, 'employers' repeatedly engage in binary hiring decisions between two 'workers' from different races. Looking at decisions between Hispanic and Asian workers, I find in experiment 1 that employers significantly discriminate against Hispanics after having been provided with actual group-level performance statistics that indicate equal productivity distributions between the two groups.

To identify the underlying mechanism that is driving this disparity, I explore the information acquisition behavior of employers. In experiments 2 and 3, participants are additionally provided with an initial piece of individual-level information about the two workers and have the opportunity to iteratively request more ambiguous (experiment 2) and uncertain (experiment 3) individual-level information signals about the workers during the hiring decisions. Ambiguity and uncertainty are established by varying the veracity of the information source and by varying the information that participants receive about the information structure. In experiment 2, participants do not know the likelihood with which a given information signal is true. In experiment 3, participants know that there is a 60% likelihood that a given signal is true (see Section 2.2.3 for details). Discrimination based on motivated reasoning implies that employers 'fish for good news' Chen and Heese (2021), which means they try to acquire information that supports their motive, e.g., by seeking additional information when a previous piece of information contradicts their motive and stopping to seek additional information once their information set is consistent with their motives. In line with this, I find that if the initial piece of provided information contradicts the employers' motive, they are significantly more likely to search for a second piece of

information and acquire more pieces of information in total, as compared to when the initial piece of provided information confirms their motive.

I also study to what extent employers act consistently with the information, and how this depends on whether or not the information is in line with their motive. Further corroborating the mechanism of discrimination based on motivated reasoning, I find that employers in both experiments are more likely to act consistently with the acquired information if the information confirms their motive and that this effect is larger in experiment 2, where the wiggle room to interpret the information is larger than in experiment 3.

Finally, in experiment 4, I confirm that discrimination can be reduced by limiting the ‘wiggle room’ Dana et al. (2007) of employers to interpret information. After reducing the ambiguity and uncertainty of provided information signals, I find decreasing levels of discrimination and thereby provide evidence of how discrimination based on motivated reasoning can be avoided.

Overall, this form of discrimination links with taste-based discrimination as preferences drive motivated reasoning and thereby contribute to the formation of biased beliefs. However, as mentioned above, in particular the behavior in experiment 4 contrasts with taste-based discrimination, which cannot be changed by the provision of different information. Similarly (inaccurate) statistical discrimination without motivated reasoning cannot fully explain the findings across all experiments, in particular not those in experiments 2 and 3, as in the absence of motivated reasoning, individuals should not selectively engage with presented information depending on its content.

Taken together, this paper makes three main contributions. First, it proposes a new explanation of discrimination that links taste-based discrimination and inaccurate statistical discrimination: discrimination based on motivated reasoning. By selectively acquiring and processing information, individuals form motivated beliefs and consequentially discriminate based on these beliefs. Through a series of experiments, it provides evidence of the existence and the underlying mechanisms of this explanation. By documenting biased information acquisition as well as belief distortions it provides evidence of two potential ways in which motivated reasoning impacts belief formation. Second, it demonstrates that discrimination based on motivated reasoning differs from traditional forms of discrimination by documenting how it can explain discrimination in settings where the basic model of statistical discrimination (without motivated reasoning) as well as taste-based discrimination cannot

fully explain observed behavior. Finally, it shows how these insights can be utilized to design an effective policy intervention to reduce discrimination. By varying the ambiguity of individual-level information, this paper demonstrates that limiting individuals' wiggle room to systematically engage with information may significantly reduce discrimination.

The remainder of this paper is structured as follows. Next, I review related literature on different models of discrimination with a particular focus on belief-based models, and discuss existing work on belief distortions and biased information choices in discrimination. Section 2 provides details on the experimental design and Section 3 demonstrates to what extent group-level beliefs of employers were aligned between treatments. Based on this, Section 4 explains the theoretical background of discrimination based on motivated reasoning and derives testable predictions. Main results are provided in Section 5. Finally, Section 6 concludes.

1.1 Related Literature

This paper connects to the broad economic literature on theories of discrimination and empirical methods to measure it.² In particular, by introducing a particular form of discrimination that links taste-based discrimination and inaccurate statistical discrimination, it contributes to the recently growing literature, that aims to extend the traditional taxonomy of taste-based discrimination (Becker, 1957) versus statistical discrimination (Phelps, 1972; Arrow, 1973; Aigner and Cain, 1977). Other papers that aim to reveal the limits of this long-standing taxonomy by looking more closely at inaccurate beliefs as drivers for discrimination include Coffman et al. (2021), who document that beliefs about average group differences drive gender discrimination in an artificial online hiring experiment, or Barron et al. (2022), who differentiate between explicit and implicit belief-based discrimination between genders.

Setting the stage for this paper, Bohren et al. (2023) argue for the importance of considering incorrect beliefs and document them. They demonstrate that when economic agents have inaccurate beliefs about group attributes, resulting discrimi-

²See, e.g., Charles and Guryan (2011) for a review on challenges to measuring racial discrimination, Bohren et al. (2018) for a discussion about observational vs. experimental data, Bertrand and Duflo (2017), Heckman (1998), and Neumark et al. (2016) for discussions of field experiments, Onuchic (2022) for a recent review on theories of discrimination, Cornell and Welch (1996) for a model of ingroup vs. outgroup screening discrimination, and Bohren et al. (2022) for tools to model and measure systemic discrimination.

nation based on these beliefs can be mistaken for taste-based discrimination. The authors argue that eliciting beliefs or providing multiple information treatments can identify inaccurate beliefs versus animus as potential drivers of discrimination and reason that those who discriminate based on taste are unlikely to change their behavior in response to group-level information. This paper builds upon their findings and takes a significant next step by exploring an important driver of inaccurate beliefs — motivated reasoning — and showing that inaccurate beliefs stemming motivated reasoning may not be corrected by simply providing such information. It considers individuals’ learning behavior and shows that individuals selectively acquire and process information in line with their motives. It demonstrates that after receiving credible group-level and/or ambiguous individual-level information, individuals use the ‘wiggle room’ (Dana et al., 2007) that this information provides for the formation of persistent inaccurate beliefs about unobserved individual-level characteristics. By showing that motivated reasoning leads to inaccurate statistical discrimination, this paper provides a micro-foundation for the emergence of inaccurate beliefs in a discrimination setting.³

Besides this study, to date only a few other studies have argued for the importance of considering incorrect beliefs as a source of discrimination in the economics literature.⁴ However, these studies have either not identified the mechanism for why these belief errors arise or how they can persist in the face of informative signals. As such Bohren et al. (2019) and Sarsons (2017) provide evidence for biased beliefs but do not identify their source, and Coffman et al. (2023a), Bordalo et al. (2016), and Bordalo et al. (2019) argue that incorrect beliefs may be driven by stereotypes but do not directly tie this to discriminatory actions nor does their mechanism explain the persistence of incorrect beliefs over time and in the context of endogenous information acquisition and selective processing.

By introducing an endogenous information acquisition setting in a discrimination

³Other potential mechanisms for non-Bayesian updating or model misspecification resulting in persistent inaccurate beliefs include but are not limited to memory or attention constraints (e.g. Bordalo et al., 2023), correlation neglect (e.g. Enke and Zimmermann, 2019), or failure of contingent reasoning (e.g. Martínez-Marquina et al., 2019). See Bohren and Hauser (2021) for a general framework and more examples of how misinterpreting information impacts learning.

⁴In an in-depth literature review Bohren et al. (2023) show that only 10.5% of 105 papers on discrimination that were published in 10 top economics journals between 1990 and 2018 differentiate between accurate and inaccurate beliefs. These and more recent studies include e.g. Fershtman and Gneezy (2001), Albrecht et al. (2013), Reuben et al. (2014), Bordalo et al. (2016), Bordalo et al. (2019), Bohren et al. (2019), Bursztyn et al. (2020), and Esponda et al. (2022).

setting this paper also adds to the literature on belief distortions and biased information choices through systematic information acquisition and processing in discrimination contexts. Bartoš et al. (2016) pushed the research frontier in this context by modeling and documenting in three field experiments that rational (in)attention can amplify discrimination. They demonstrate that employers pay less attention to a priori less attractive applicants in cherry-picking markets, but more attention to a priori less attractive applicants in lemon-dropping markets. They reason that attention allocation is determined by the likelihood that costly information would change the status quo decision which is to not hire an applicant in the cherry picking market and to hire an applicant in the lemon dropping market.

Focusing more on the extent to which provided information is processed in order to update beliefs, Mengel and Campos-Mercade (2021) attribute disparities in artificial hiring to signal neglect in the belief formation process. They show that employers conservatively update beliefs when confronted with new information and ultimately discriminate against disadvantaged workers.

In closer and very recent work, Coffman et al. (2023b) experimentally study how endogenous information acquisition amplifies discriminatory outcomes between genders in a simulated hiring experiment. While they do not find systematic gender differences in hiring, they observe that relative comparisons – the extent to which a candidate’s gender outperforms the other within the candidate pool – bias decision-making and observe that relative (dis)advantages amplify through endogenous information search. Consistent with the findings of my study, these studies provide evidence that systematic attention allocation and systematic processing of information can impact discrimination. Building on this, I look closer at the mechanisms by considering motivated, rather than cognitive, explanations for behavior and by studying how endogenous information search and processing can lead to persistent and inaccurate beliefs that generate discrimination.

In a non-discrimination context, Chen and Heese (2021) also study biased information choices and show that individuals are more (less) likely to continue acquiring information after they have received information that suggests that acting selfishly is harmful (harmless). Together with Bohren and Hauser (2023) who provide a more general framework for mechanisms of belief distortions and show in an application how self-image concerns can lead to inaccurate beliefs and consequential discrimination, they provide potential theoretical foundations for the observed information acqui-

sition and processing behavior observed in this study. Similarly, this paper builds on theoretical foundations and psychological underpinnings of motivated reasoning (Kunda, 1990; Epley and Gilovich, 2016), which have been documented in various non-discrimination contexts, such as altruistic situations (Di Tella et al., 2015), in response to feedback (Zimmermann, 2020), or in light of uncertainty (Drobner, 2022). More practically, it builds on Thaler (2020), who provides a novel experimental design to identify motivated reasoning and shows that motivated reasoning, even based on uninformative messages, can lead individuals’ beliefs to become more polarized and less accurate. I apply a modified version of the proposed experimental design to identify motivated reasoning in the context of discrimination.

Finally, very recently, two subsequent studies by Rackstraw (2022) and Stötzer and Zimmermann (2022) follow the idea of motivated reasoning as a driver for discrimination and confirm the findings of this paper in complementary experiments. While Rackstraw (2022) finds that experimental employers update their beliefs consistently with their implicit biases in a different hiring experiment, Stötzer and Zimmermann (2022) find that individuals justify selfish actions through ex-post belief updating in a survey experiment.

2 Experiment design

The data collection involves one survey and a series of pre-registered online experiments, programmed in Otree (Chen et al., 2016) and implemented on Prolific, that has shown to provide superior data quality compared to various other platforms (Peer et al., 2022). The following subsections explain the experimental design. For more details and screenshots, see Online Appendix A.

2.1 The pool of workers

In the survey, I collect answers from 96 US participants on a logic quiz, a dictator game, and a real effort task. This is meant to mimic an ‘assessment center’ that provides proxies for workers’ cognitive ability, social competence, and perseverance. Based on their answers in all three tasks, I calculate a score that defines their ‘productivity’ for the subsequent experiments. Additionally, I ask for race and other demographics, past school performance information, and psychological scales for re-

silence, cooperativeness, ambition, and diligence. This information is used to set up profiles of available experimental ‘workers’ for hire in the main experiments. To rule out statistical discrimination based on accurate beliefs as the cause for potential discrimination in the hiring experiment by design, I set up the final pool of workers for the hiring stage by selecting workers so that productivity distributions between race groups are equal. The final pool of workers that was used in the hiring experiments consists of 58 individuals from the US, equally balanced across gender, aged between 18 and 30 with a mean age of 22.83 years.⁵

2.2 The hiring experiments

In four different hiring experiments, subjects act as employers and are asked to hire workers from the constructed pool in a series of binary hiring decisions. In total, the aggregated pool of employers consists of 1138 subjects.

2.2.1 Inducing motives

All experiments involve the same two treatments in a between-subjects design. In both treatments, the employers make hiring decisions between two workers from different races. The two treatments differ in the way the workers are labeled. In treatment ‘Race’, the employers could observe the races of the workers, whereas in treatment ‘Neutral’ the race labels are replaced by neutral shape labels, e.g. ‘Triangle’ workers, or ‘Diamond’ workers. Importantly, the composition of workers in each hiring decision remains constant between treatments; only the labels are changed. This treatment variation induces a motive to hold particular productivity beliefs for employers of treatment Race, but not for those of treatment Neutral. Hence, results from employers in treatment Neutral serve as baseline levels in the subsequent analysis.

2.2.2 Experiment procedure

The experiments consist of a belief stage followed by a hiring stage. In the belief stage, the employers are asked for their subjective belief distribution about the group productivities of the workers. These beliefs indicate the direction of their motives.

⁵Of the 58 workers in the final pool, 17 identified as ‘White’, 16 as ‘African American or Black’, 13 as ‘Asian’, and as 12 ‘Hispanic or Latin’ (hereinafter referred to as ‘Hispanics’).

Subsequently, they are provided with the true productivity distributions for each group of workers and are again asked for their beliefs. All employers were informed that the pool of workers consisted of selected workers from the survey.

As in Coffman et al. (2021), the data on beliefs reveals two things. First, it shows to what extent prior and posterior beliefs reflect the true productivity distributions. Second, and more importantly, it demonstrates the differences in beliefs between the two treatments. Updating beliefs between employers in treatment Race and employers in treatment Neutral is intended to generate identical and correct ex-post beliefs about the productivity distributions among employers in both treatments. This renders the potential motives based on the workers' race the only difference between the two treatments prior to the hiring decisions.⁶

In the hiring stage, I ask the employers to make a sequence incentivized hiring decisions between two workers from the constructed pool of workers. Each employer is repeatedly presented with a pair of workers and asked to hire one of them. After the experiment, one hiring decision is randomly chosen. If they hired the worker with the higher productivity score in this decision the employer gets a bonus payment of \$2. The measurement of productivity has previously been explained to all employers. The workers did not receive an additional payment for being hired. After all hiring decisions the employers answer a few demographic questions before the experiment ends.⁷

2.2.3 Varying wiggle room

Discrimination based on motivated reasoning implies that less wiggle room for the employers to form their beliefs about individual workers can decrease discrimination. Therefore, the experiments differ in the way *individual-level* information is provided in each hiring decision and thus in the extent of wiggle room that the employers have.

In Experiment 1 ('No Information'), employers have the most wiggle room as they do not receive any individual-level productivity information after having received the accurate group-level information during the belief stage. This means, that in each

⁶Until here, the basic structure of the experiments is partly inspired by the design of Coffman et al. (2021) who study gender discrimination in binary hiring decisions, after aligning beliefs between employers in a gender treatment and a birth-month treatment.

⁷The project received IRB approval at the joint ethics board of Goethe University Frankfurt and JGU Mainz. The experiments are pre-registered in the AEA RCT registry under AEARCTR-0008609.

hiring decision in treatment Race employers are simply asked to hire e.g. the Asian worker or the Hispanic worker without receiving any further information about these particular two workers. In each hiring decision in treatment Neutral they are accordingly asked to hire e.g. the Triangle worker or the Diamond worker without receiving any further information about the two workers.⁸ In total 454 subjects successfully completed this experiment, 231 subjects in treatment Race and 223 subjects in treatment Neutral.

In Experiment 2 (‘Ambiguous Information’), employers have a similar level of wiggle room. After the belief stage, they complete the same hiring task as the employers in the ‘No Information’ experiment. However, employers in this experiment receive an initial piece of individual-level information in each hiring decision and are able to request up to 9 additional signals. The signals consist of a binary message from one of two randomly-chosen information sources: True News or Fake News. The message from the True News source is always correct, the message from the Fake News source is never correct. The message reads “The better worker is X” where X is one of the two presented workers.⁹

Note, that if the message comes from the Fake News source it implies that the worker who is not mentioned is indeed the better worker. Importantly, the employers do not know whether or not a message comes from the True news source or from the Fake News source. In fact, in this experiment, the employers are not informed about the likelihood with which a message comes from the True News source or the Fake News source.¹⁰ Since the messages in this experiment are theoretically uninformative, they provide the employer with wiggle room to subjectively interpret each message according to their motive. Apart from this information structure during the hiring phase, this experiment is identical to the ‘No Information’ experiment. In total 470 subjects successfully completed this experiment, 234 subjects in treatment Race and 236 subjects in treatment Neutral.

In Experiment 3 (‘Uncertain Information’) employers are provided with the exact same information structure as in the ‘Ambiguous Information’ experiment. However,

⁸Importantly, the Triangle (Diamond) worker in treatment Neutral is the same worker as the Asian (Hispanic) worker in treatment Race. I vary race-shape assignments between sessions.

⁹This feature is inspired by a novel experimental design to identify motivated reasoning by Thaler (2020). I apply a modified version to identify motivated reasoning in the context of discrimination.

¹⁰Subsequent elicitation of the perceived fraction of messages from the True News source reveals that the vast majority of employers guess that approximately 50% of all messages are true.

in this experiment employers have previously been told that each message has a 60% likelihood to come from the True News source and a 40% likelihood to come from the Fake News source. This adds transparency about the information structure. Following (Dana et al., 2007), who argue that transparency about actions may reduce wiggle room, this additional layer of transparency about the information reduces employers’ wiggle room to subjectively interpret each message compared to the ‘Ambiguous Information’ and ‘No Information’ experiments. Other than that, this experiment is identical to the ‘Ambiguous Information’ experiment. In total 108 subjects successfully completed this experiment, 52 subjects in treatment Race and 56 subjects in treatment Neutral.

In Experiment 4 (‘Tangible Information’) employers do not receive messages that directly displayed the (supposedly) better worker. Instead, employers are given individual-level information about past performances of the two workers.¹¹ Again, for each decision, employers receive one initial piece of information (e.g. their college GPA) and can request up to nine additional random pieces in each hiring decision. This environment still leaves wiggle room for participants but reduces it further as the signals are now always true, but still not conclusively predictive of the better of the two workers. In total 106 subjects successfully completed this experiment, 57 subjects in treatment Race and 49 subjects in treatment Neutral.

On average, the experiments took between 15 minutes (‘No Information’) to 24 minutes (‘Tangible Information’). They were each conducted in December 2021 and with representative samples of the US population.¹²

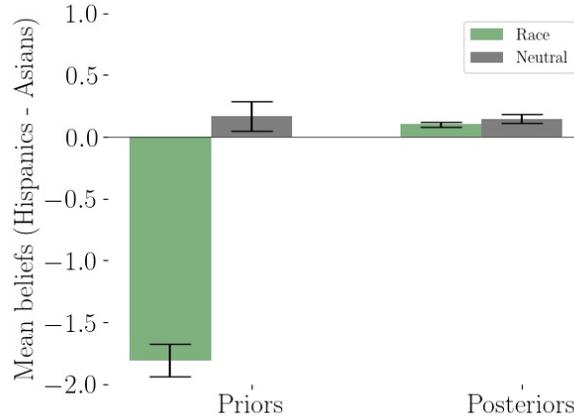
3 Beliefs about group productivities

Before describing the predictions and hypotheses in more detail I first present results of the belief stage that identify employers’ motives. I then demonstrate that employers in treatment Race and treatment Neutral hold identical group-level beliefs about the productivity of workers from different races, once they enter the hiring stage.

¹¹The past performance information includes their college and high school GPA, SAT and ACT score, final high school math and English grade, and psychological measurements of their level of ambition, resilience, diligence, and agreeableness (based on Duckworth et al., 2007; Rammstedt and John, 2005; Sinclair and Wallston, 2004).

¹²Representativeness is ensured in terms of age and gender and is established by Prolific (www.prolific.com). Across all experiments attrition rates were similar and below 3% among participants who were informed about the information structure.

Figure 1: Mean beliefs about group productivities



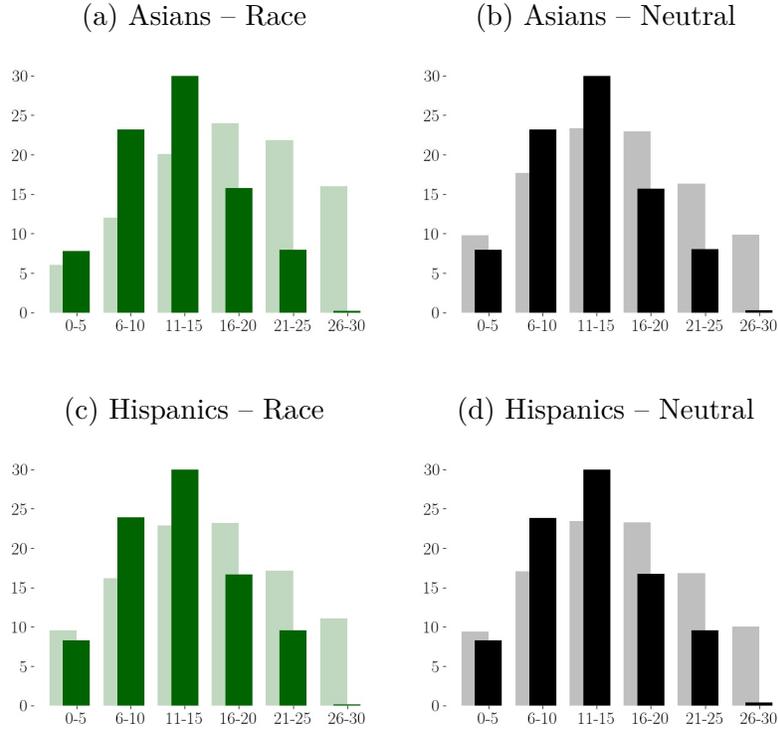
Notes: The first two bars show the mean productivity beliefs *before* the group-level information update of employers in treatment Race and treatment Neutral, respectively, the last two bars show the mean productivity beliefs *after* the group-level information update. Beliefs are plotted as the difference in mean beliefs about Hispanic workers and mean beliefs about Asian workers such that negative values indicate higher beliefs about the productivity of Asian workers than beliefs about the productivity of Hispanic workers, and vice versa. Error bars show the standard errors of the means.

Looking at the prior beliefs about group productivity levels, the first bar of Figure 1 shows a large difference in mean beliefs about Asians and Hispanics in treatment Race. Employers believed Asian workers to be significantly more productive than Hispanic workers. The second bar shows the differences in mean beliefs about Hispanic and Asian workers for employers in treatment Neutral and shows no significant differences in mean productivity levels between the two treatments. Considering that the two groups were presented to them as group Triangle and group Diamond, this result merely constitutes a verification. The difference in differences implies a potential for motivated beliefs that Asian workers are more productive than Hispanic workers.¹³

Next, I check whether the group-level information intervention was successful in aligning beliefs. The two bars on the right of Figure 1 show that the treatment difference in differences in mean beliefs about group-level productivities between Asians

¹³Informed by a pilot study and to attenuate potential experimenter demand effects, I also elicited beliefs and hiring decisions about black and white workers to divert attention from the main decisions. By actually focusing on decisions between Asian and Hispanic workers in the analysis, I consider choices between two minority groups, decreasing the potential for social desirability bias and image concerns, particularly among white employers, e.g., due to recent Black Lives Matter movements. Mean beliefs about black and white workers are shown in Figure B1 in Online Appendix B.

Figure 2: Prior and Posterior belief distributions



Notes: Figures (a) and (b) show employer belief distributions about the productivities of Asian workers before (low opacity) and after (high opacity) the group-level information update in treatment Race and treatment Neutral, respectively. Figures (c) and (d) show employer belief distributions about the productivities of Hispanic workers before (low opacity) and after (high opacity) the group-level information update in treatment Race and treatment Neutral, respectively.

and Hispanics disappeared after the information update. Importantly, this implies, that group-level beliefs about the productivity of Asians and Hispanics were equal across treatments when employers entered the hiring stage. In fact, Figure 2 demonstrates that employers in both treatments learn that the average, variance, and skewness of the productivity distributions of Asian and Hispanic workers are equal. This result holds for all four experiments.¹⁴ In the next section, I present testable predictions for discrimination based on motivated reasoning.

¹⁴Figure B2 in Online Appendix B shows that beliefs were updated homogeneously across employers in each experiment. Theoretically, experimenter demand or confusion could have caused some employers to report similar belief distributions across races, despite not actually holding these beliefs. However, employers knew that the pool of workers was artificially constructed from a larger pool of Asian and Hispanic workers. This implies that there was no reason not to believe the provided group-level information.

4 Discrimination based on motivated reasoning

This section describes the predictions and hypotheses in line with discrimination based on motivated reasoning. Taking the identical group-level belief distributions between employers of treatment Race and employers of treatment Neutral as a point of departure, this section presents various predictions for the hiring stages. The predictions depend on employers' wiggle room to systematically acquire and process information. Before, I start with a short discussion on the definition of a 'motive'.

4.1 What is a motive

Throughout the paper, I remain agnostic about potential rationales of a motive for a particular belief. A motive could be based on a taste against members of a particular group, but it could also be any other reason that leads an individual to systematically and wrongly believe that a member of one group is more/less productive than a member of another group. As an example, consider an employer who has always held the belief that members of one group are more productive than members of another group and who has based previous actions on this belief. If this employer would now come to realize that their long-standing belief was in fact untrue, they might experience a form of cognitive dissonance (Festinger, 1957). In response, to avoid experiencing cognitive dissonance the employer might prefer to stick with their initial belief. In this case, the employer might not hold any taste or animus against the member of a particular group but would still have a motive to hold a particular belief. Importantly, for discrimination based on motivated reasoning to explain disparate treatments of members from certain groups, it does not matter what the motive is based on, as long as it provides any reason to hold a particular belief. Importantly, this form of discrimination still differs from conventional inaccurate statistical discrimination without motivated reasoning, as it is characterized by endogenous information search and processing dependent on the direction of the motive, even when group-level beliefs about the productivity of different groups are equal.

4.2 Predictions under wiggle room

I predict that even though employers believe that there are no group-level productivity differences between Asians and Hispanics, there can still be discrimination against

Hispanics. The reasoning is that without any individual-level information, employers have enough wiggle room to reinstate their initial beliefs of a productive Asians vs. a less productive Hispanic when confronted with a binary decision between two individuals. In the ‘No Information’ experiment, this wiggle allows them to engage in motivated sampling from two similar productivity distributions such that they reinstate the belief of a productive Asian worker and a less productive Hispanic worker. In the ‘Ambiguous Information’ experiment, individuals have the additional wiggle room to misperceive the informativeness of the signal to align with their motives. While employers in both experiments know, that on a group-level there are no differences, they will ultimately reason that between particular two workers from either group, the Asian worker will still be the more productive worker. Note that, discrimination is defined as hiring likelihood disparity regardless of given signal realizations.

Hypothesis 1a *There is significant discrimination against Hispanics in treatment Race, but not in treatment Neutral, among employers in experiments ‘No Information’ and ‘Ambiguous Information’.*

Providing employers with ambiguous individual-level information allows having a first look at the actual mechanism through which beliefs are reinstated, by yielding insights into the information acquisition and processing behavior of employers. Discrimination based on motivated reasoning predicts that individuals will ‘fish for good news’, meaning that they will request additional information signals if the previous signal contradicts their motive, but will not request additional signals if the previous signal confirms their motive. Since all individuals received one initial signal by default, I test for each hiring decision whether or not the content of the initial information signal affects (i) the likelihood to request a second information signal, and (ii) how many signals employers request per decision.

Hypothesis 1b *If the initial information signal suggests to hire the Hispanic worker in the ‘Ambiguous Information’ experiment, employers in treatment Race are more likely to request a 2nd signal and request more signals than employers in treatment Neutral.*

Moreover, discrimination based on motivated beliefs predicts that individuals overweight information signals that confirm their motive and underweight signals that contradict their motive, if their wiggle room allows. To test this, I analyze whether

or not the overall direction of requested signals in a decision affects the likelihood to which employers act according to the signals. More specifically, I test whether individuals are more likely to follow the signals, if the majority of requested signals confirms their motive, and less likely if the majority of requested signals contradict their motive.¹⁵

Hypothesis 1c *If the majority of all considered signals in a decision suggests to hire the Hispanic worker in the ‘Ambiguous Information’ experiment, employers in treatment Race are less likely to act according to the suggestion than employers in treatment Neutral.*

4.3 Predictions under reduced wiggle room

Conceptually, the ‘Uncertain Information’ experiment and the ‘Tangible Information’ experiment are slightly different from each other. In the ‘Uncertain Information’ experiment only 60% of all signals are true and the information structure is rather abstract. In contrast, in the ‘Tangible Information’ experiment all signals are true and the information structure is very simple and includes specific, measurable indications of individuals’ performance that can be directly understood and easily related to. This implies that in the former motivated reasoning may be based on the evaluation of the truthfulness of a signal (as in the ‘Ambiguous Information’ experiment), whereas in the latter, it is based on the predictive power of a signal. Exploratively, I study the extent to which motivated reasoning influences by the evaluation of the truthfulness of constructed information in the ‘Uncertain Information’ experiment differs from motivated reasoning guided by the evaluation of predictive power of the more relatable information in the ‘Tangible Information’ experiment. Importantly, in both experiments, the provided information becomes more meaningful than in the ‘Ambiguous Information’ experiment. This implies that employers have less wiggle room to interpret information according to their motives. In response, motivated belief-based discriminators might still systematically search for news that supports their motive. However, knowing that this kind of individual-level information provides more meaningful suggestions, motivated belief-based discriminators will reduce

¹⁵Even if the previous belief alignment had suffered from confusion or experimenter demand effects and group-level beliefs were hence not fully aligned, systematically different search and processing of this kind of information would still contradict Bayesian behavior as the signals should appear entirely uninformative to a rational Bayesian agent without a motive.

the systematic overweighting of information signals that are in line with their motives (and vice versa).

Hypothesis 2a *If the majority of all considered signals in a decision suggests to hire the Hispanic worker, the difference in the likelihood to follow the suggestion between treatment Race and treatment Neutral is larger in the ‘Ambiguous Information’ experiment than in the ‘Uncertain Information’ and the ‘Tangible Information’ experiments.*

Consequently, limiting employers’ wiggle room should be an effective way to reduce discrimination where discrimination is again defined as hiring likelihood disparity irrespective of given signal realizations.

Hypothesis 2b *There is less discrimination against Hispanics in decisions of employers in the ‘Uncertain Information’ and the ‘Tangible Information’ experiments, than in decisions of employers in the ‘Ambiguous Information’ or the ‘No Information’ experiments.*

By design, accurate statistical discrimination is ruled out as there are no group-level differences between Asians and Hispanics in my constructed pool of workers and hence potential beliefs about differences cannot reflect actual differences. Moreover, taste-based discrimination is unlikely, as it assumes that economic agents experience a disutility from interactions with certain individuals, but my experimental setting does not include any employer-worker interaction and employers are aware of this right from the start.¹⁶ However, it is still possible that individuals consider the mere selection of a worker from the less preferred group as a distasteful interaction with that worker. To argue further that taste-based discrimination is an unlikely explanation for discrimination in this study, I present evidence of behavior that is inconsistent with taste-based discrimination in the analyses of the data from experiments 2 to 4.

5 Results

I first consider to what extent employers discriminate in the ‘No Information’ and the ‘Ambiguous Information’ experiment. Subsequently, the analysis of the informa-

¹⁶In other studies, this design feature is used to rule out taste-based discrimination (e.g. Barron et al., 2022).

tion acquisition and processing behavior in the ‘Ambiguous Information’ experiment provides first insights into the extent of potential motivated reasoning.

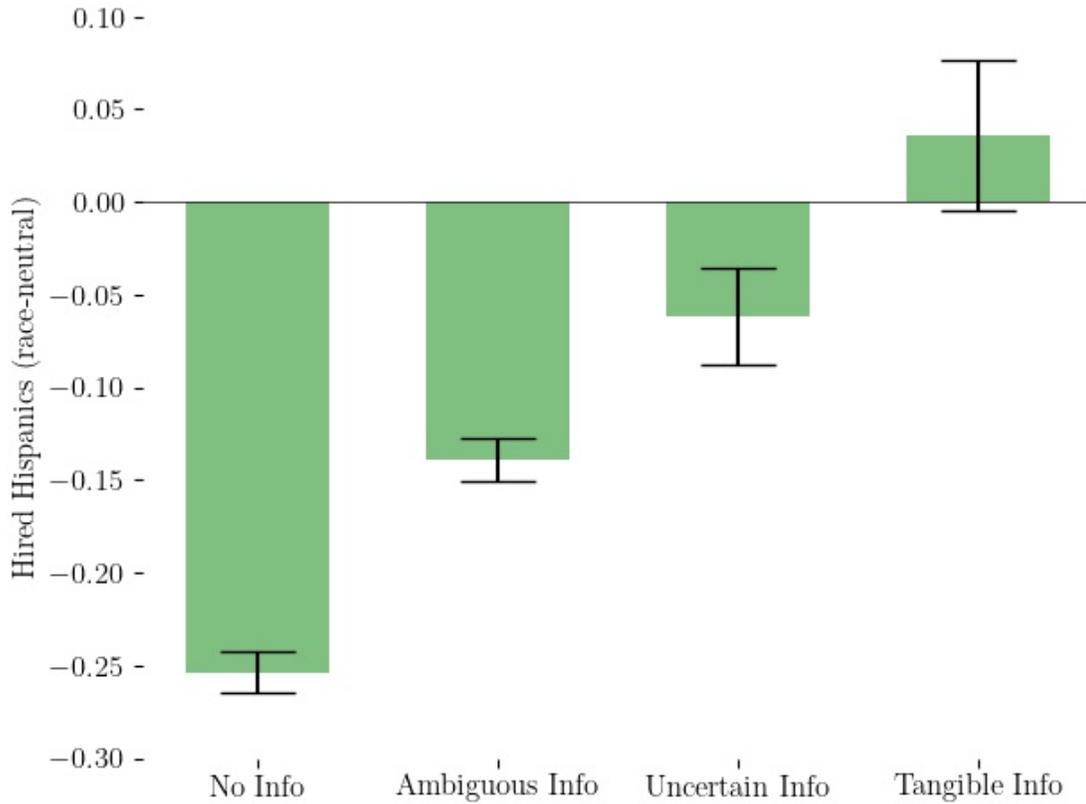
5.1 Wiggle room - No Information & Ambiguous Information

In line with hypothesis 1a, the first two bars in Figure 3 show that discrimination against Hispanics is substantial in these two experiments. Providing employers with the race labels of workers in the ‘No Information’ experiment significantly decreases the hiring rate of Hispanics by 25.36pp. In the ‘Ambiguous Information’ experiment, there is also a significant 13.92pp difference in hiring rates of Hispanics between the two treatments. Albeit this difference is lower than in the case with no additional information, both experiments show significant discrimination against Hispanics, even though employers are aware that group-level productivity scores between the two groups are equal.¹⁷ Usually, this form of discrimination would be described as taste-based discrimination despite the fact that employers and workers do not interact. However, to further investigate the potential cause of this discrimination, I now consider the information acquisition and processing behavior of the employers in the ‘Ambiguous Information’ experiment.

Table 1 provides first evidence that discrimination is based on motivated reasoning. Columns 1 and 2 show that an initial signal that suggests to hire the Hispanic worker has a significantly larger positive effect in treatment Race than in the treatment Neutral on (1) the likelihood to acquire a second signal (coeff=0.1046, p-value=0.004) and on (2) the total number of signals requested per decision (coeff=0.5161, p-value=0.028). These results are consistent with hypothesis 1b and demonstrate that employers were ‘fishing’ for information signals that confirm their motive. Finally, column 3 indicates that employers of treatment Race are significantly less likely to follow the signals if they suggest to hire the Hispanic worker than employers of treatment Neutral (coeff=-0.1881, p-value< 0.001). This lends support to hypothesis 1c. Taken together, the results imply that employers discriminate after having system-

¹⁷The difference can be explained by the decision rule of some individuals to just follow the signals which lowered the overall discrimination rate in the ‘Ambiguous Information’ experiment compared to the one in the ‘No Information’. However, discrimination rates were more extreme (37.33pp.) in the ‘Ambiguous Information’ experiment in decisions in which subjects only requested one additional signal and ended up with two conflicting signals (see also Figure C1 in Online Appendix C.1) which supports the idea of motivated reasoning.

Figure 3: Hiring rates of Hispanic workers across experiments



Notes: The vertical axis displays the fraction of hired Hispanics in treatment Race minus the fraction of hired Hispanics in treatment Neutral. The horizontal axis groups the decisions by experiments ('No Information', 'Ambiguous Information', 'Uncertain Information', and 'Tangible Information'). Error bars indicate standard errors of the means.

atically acquired and processed ambiguous information about the workers. Usually, under the assumption that employers derive disutility from the mere selection of a Hispanic worker, the results so far would still allow classifying observed discrimination as purely taste-based discrimination, in which the discriminators use the available information to justify their taste. In the next section, I provide evidence against this interpretation.

Table 1: Information acquisition and processing behavior under ambiguous information

| | (1) 2nd signal | (2) number of signals | (3) follow signals |
|-------------------------|---------------------|--------------------------|-----------------------|
| race * hispanic | 0.1046 (0.0361) | 0.5161 (0.2345) | -0.1881 (0.0403) |
| race | -0.0667 (0.0427) | -0.3422 (0.2804) | 0.0596 (0.0201) |
| hispanic | -0.0390 (0.0255) | -0.2801 (0.1705) | -0.0695 (0.0268) |
| Observations | 3290 | 3290 | 3290 |
| Baseline mean dep. var. | 0.5483 | 3.6246 | 0.8676 |

Notes: This table shows OLS results of the ‘Ambiguous Information’ experiment. *2nd signal*, is a dummy equal to 1 if an employer requested a second signal. *number of signals*, counts the number of requested signals. *follow signals*, is a dummy equal to 1 if an employer’s hiring decision followed the majority of considered signals. *race* is the treatment dummy equal to 1 if the decision was made in treatment treatment Race. *hispanic* is a dummy equal to 1 if the initial signal (columns 1 and 2) or the direction of all considered signals (column 3) suggests to hire the Hispanic worker. Note that this holds for both treatments even though employers in treatment Neutral did not observe the race of the workers. Units of observation are decision specific. Standard errors are clustered at the individual level and displayed in parentheses.

5.2 Reduced wiggle room - Uncertain Information & Tangible Information

To deepen our understanding of discrimination based on motivated reasoning I now consider the results from the two ‘debiasing’ experiments. In experiments ‘Uncertain Information’ and ‘Tangible Information’, I limit employers’ wiggle room to systematically use the information to form motivated beliefs by providing more meaningful individual-level information. If discrimination in the previous two experiments was based on taste and employers simply used the information to justify their taste, we should not observe any differences between behavior of employers in the ‘Ambiguous Information’ experiment and behavior of employers in the experiments with reduced wiggle room. This is because the hiring behavior of taste-based discriminators should not respond to variations in information. While one could argue that less wiggle room might complicate the ex-post rationalization of animus towards Hispanics, it does not affect the animus itself. The driver of discrimination would still be unaffected by

changes in information and the hiring behavior should therefore not change with reduced wiggle room in experiments ‘Uncertain Information’ and ‘Tangible Information’ if discrimination was based on taste.

First, I consider the case where employers were given similar information as in the ‘Ambiguous Information’ experiment, except that employers are told that signals have a 60% likelihood to come from the True News source and a 40% likelihood to come from the Fake News source. The information acquisition and processing behavior of employers in this experiment is shown in Table 2. Column 1 indicates that employers in treatment Race are insignificantly more likely to acquire a second information signal if the initial signal suggests to hire the Hispanic worker than employers in treatment Neutral and column 2 shows that employers acquire insignificantly more information signals in treatment Race than in treatment Neutral. Note that coefficient sizes are similar to the results from the ‘Ambiguous Information’ experiment and the insignificance is likely to be an artefact of the much smaller sample size in this experiment. Looking at whether or not employers follow the overall direction of signals systematically differently, columns 3 and 4 of Table 2, illustrate that reducing ambiguity and hence the potential wiggle room to interpret the information substantially decreases the treatment difference in the likelihood with which employers act in line with the acquired information. A majority of requested signals that suggests to hire the Hispanic worker, reduces the likelihood to follow the suggestion by 6.08pp more in treatment Race than in treatment Neutral. This is a 12.74pp reduction in the absolute effect size compared to the results from the ‘Ambiguous Information’ experiment. Column 4 shows the results from the triple interaction analysis and provides evidence that this reduction is significant.

Next, I consider the case where employers were given individual-level information about past performances of the two workers in each decision. Across all three outcomes (probability to acquire a second signal, number of signals considered in each decision, likelihood to follow the majority of considered signals) Table 3 displays further reduced effect sizes compared to the ‘Uncertain Information’ experiment. An initial signal that suggests to hire the Hispanic worker has no significant treatment effect on the likelihood to acquire a second signal and on the number of signals requested. Finally, column 4 shows that the effect on the likelihood to follow the signals of (-2.06pp) has also significantly decreased by 16.75pp, compared to the effect in the ‘Ambiguous Information’ experiment. Together with the results from the ‘Uncertain

Table 2: Information acquisition and processing behavior under uncertain information

| | Uncertain Information | | | Uncertain - Ambiguous |
|-------------------------|-----------------------|--------------------------|-----------------------|-----------------------|
| | (1) 2nd signals | (2) number of signals | (3) follow signals | (4) follow signals |
| race * hispanic | 0.0898 (0.0614) | 0.6035 (0.5416) | -0.0608 (0.0494) | 0.1274 (0.0638) |
| race | 0.0083 (0.0800) | 0.2919 (0.7498) | 0.0456 (0.0261) | |
| hispanic | -0.0141 (0.0462) | 0.1543 (0.3632) | -0.0292 (0.0320) | |
| Observations | 756 | 756 | 756 | |
| Baseline mean dep. var. | 0.7313 | 5.4478 | 0.9292 | |

Notes: This table shows OLS results of the ‘Uncertain Information’ experiment. *2nd signal*, is a dummy equal to 1 if an employer requested a second signal. *number of signals*, counts the number of requested signals. *follow signals*, is a dummy equal to 1 if an employer’s hiring decision followed the majority of considered signals. Column (4) refers to the differences in coefficients for *follow signals* in experiments ‘Uncertain Information’ and ‘Ambiguous Information’ and tests hypothesis 2a. *race* is the treatment dummy equal to 1 if the decision was made in treatment Race. *hispanic* is a dummy equal to 1 if the initial signal (columns (1) and (2)) or the direction of all considered signals (column (3)) suggests to hire the Hispanic worker. Note that this holds for both treatments even though employers in treatment Neutral did not observe the race of the workers. Units of observation are decision specific. Standard errors are clustered at the individual level and displayed in parentheses.

Information’ experiment above, this supports hypothesis 2a.¹⁸¹⁹

Summarizing these results, I find that reducing employers’ wiggle room to interpret individual-level information affects the extent to which they follow the signals. While in the ‘No Information’ and the ‘Ambiguous Information’ case, employers ‘fish for good news’ and refuse to follow the signals if this search was unsuccessful, a reduction of the wiggle room to interpret the information reduces the reluctance to follow unwanted signals. Employers may still search for information that confirms

¹⁸Figures C2 and C3 in Online Appendix C.2 provide further evidence in support of motivated reasoning by showing that results flip into the opposite direction for subjects who initially believed that Hispanics had a higher productivity score and that the extent of these results increases in prior belief differences (i.e. employers who initially perceived very strong differences in productivities in favour of Hispanic (Asian) workers are more (less) likely to acquire a second signal when the first signal suggests to hire the Hispanic worker, acquire more (less) signals in that case, and follow the majority signals more (less) often when they suggest to hire the Hispanic worker than employers who initially perceived weaker differences).

¹⁹Hypotheses 1c and 2a compare subsamples who stop requesting signals when the majority of signals suggest to hire the Hispanic worker. Table C1 in Online Appendix C.3 shows results of the “counter-hypothesis” for individuals who stop requesting signals when the majority of signals suggests to hire the Asian worker. Results are similar and confirm motivated information acquisition and processing under wiggle room.

Table 3: Information acquisition and processing behavior under tangible information

| | Tangible Information | | | Tangible - Ambiguous |
|-------------------------|----------------------|--------------------------|-----------------------|-----------------------|
| | (1) 2nd signals | (2) number of signals | (3) follow signals | (4) follow signals |
| race * hispanic | 0.0260 (0.0318) | 0.5709 (0.3937) | -0.0206 (0.0492) | 0.1675 (0.0620) |
| race | 0.0381 (0.0374) | 0.8460 (0.5347) | -0.0328 (0.0287) | |
| hispanic | 0.0051 (0.0235) | 0.1056 (0.3137) | -0.0860 (0.0365) | |
| Observations | 742 | 742 | 742 | |
| Baseline mean dep. var. | 0.9135 | 5.5240 | 0.8894 | |

Notes: This table shows OLS results of the ‘Tangible Information’ experiment. *2nd signal*, is a dummy equal to 1 if an employer requested a second signal. *number of signals*, counts the number of requested signals. *follow signals*, is a dummy equal to 1 if an employer’s hiring decision followed the majority of considered signals. Column (4) refers to the differences in coefficients for *follow signals* in experiments ‘Tangible Information’ and ‘Ambiguous Information’ and tests hypothesis 2a. *race* is the treatment dummy equal to 1 if the decision was made in treatment treatment Race. *hispanic* is a dummy equal to 1 if the initial signal (columns (1) and (2)) or the direction of all considered signals (column (3)) suggests to hire the Hispanic worker. Note that this holds for both treatments even though employers in treatment Neutral did not observe the race of the workers. Units of observation are decision specific. Standard errors are clustered at the individual level and displayed in parentheses.

their motive, but if they cannot find this kind of information, they are now less likely to still act against the information than in the experiments with more wiggle room.

Finally, the last two bars in Figure 3 show how the adapted information processing behavior translates into less discrimination. While there was substantial discrimination in the experiments ‘No Information’ and ‘Ambiguous Information’, absolute treatment differences in hiring rates of Hispanics in the experiments with less wiggle room for the employers decrease significantly to -6.14 pp in the Uncertain Information experiment and to an even positive (but insignificant) rate of 2.81 pp in the Tangible Information experiment (see Table C2 in Online Appendix C.4). This supports hypothesis 2b and illustrates that unambiguous and tangible individual-level information that reduced individuals’ wiggle room for interpretation can effectively decrease discrimination based on motivated reasoning.²⁰

As mentioned above, if discrimination was based on taste, we should not have

²⁰Online Appendix C.5 provides robustness checks of all analyses with more limited employer pools, excluding subjects who correctly answered only 80% or less of training/test questions during the experiment. Results remain unaffected.

observed differences in employers’ behavior in these two experiments compared to employers’ behavior in the first two experiments as the content of information should not affect taste-based discriminators. Since we do observe significant differences in (i) the way employers engage with the information and (ii) hiring rates of Hispanic vs. Asian workers, the results cannot be explained by taste-based discrimination. Similarly (inaccurate) statistical discrimination without motivated reasoning cannot fully explain these findings. In the absence of motivated reasoning, individuals should not systematically differently engage with presented information depending on its content, as prior beliefs had been aligned before the hiring stage and across treatments, leaving the race labels the only difference between treatments. While it is possible, that in the Tangible Information experiment, (inaccurate) statistical discrimination without motivated reasoning might still arise from different beliefs about the informativeness of the signal, this cannot explain observed individual behavior in the Ambiguous Information experiment and the Uncertain Information experiment. Hence, the only explanation that is consistent with behavior in all four experiments is discrimination based on motivated reasoning.

5.3 Further evidence

I pre-registered and conducted two further experiments using the past performance information of the workers. These experiments were identical to the previous experiments, except for the provided individual-level information.

In the first additional experiment (‘One Information’), employers are provided with individual-level information about past performances of the two workers (as in the Tangible Information experiment). However, for each decision, employers receive one piece of information and can not request any additional piece.

In the second additional experiment (‘All Information’), employers are immediately provided with all ten available individual-level pieces of information about past performances of the two workers.

Both experiments reduce employers’ wiggle room at least as much as the Tangible Information experiment, as employers cannot engage in systematic information search and provided information is true and closely related to the respective workers. Hiring rates between treatments in these experiments are almost identical (see Figure D1 of Online Appendix D). This provides further evidence consistent with discrimination

based on motivated reasoning.

5.4 The value of motivated discrimination

To get a first intuition about the extent to which discrimination based on motivated beliefs is costly for employers, I calculate the treatment differences in foregone earnings in each experiment. In the ‘No Information’ and the ‘Ambiguous Information’ experiments, employers forego on average 0.20\$ and 0.14\$ more in treatment ‘Race’ than in treatment ‘Neutral’, respectively. This amounts to a cost of 10% and 7% of the maximum possible bonus of 2\$, respectively. In the experiments with decreased wiggle room, the foregone earnings differences decrease to 0.01\$ (‘Uncertain Information’) and 0.00\$ (‘Tangible Information’).

6 Conclusion

While the existence of discrimination in many contexts has been documented extensively, we still know much less about particular mechanisms of discrimination. This study contributes to filling this gap by taking a closer look at how individuals deal with individual-level information about others. I identify a particular form of inaccurate statistical discrimination that links documented forms of statistical discrimination to the well-known model of taste-based discrimination. When confronted with information that leaves a lot of wiggle room for interpretation, individuals make use of the inconclusiveness of the information and engage with the information systematically differently, depending on the informational content. They search for information that confirms their motive and put less weight on information that contradicts their motive, and ultimately discriminate. When the wiggle room of individuals to interpret information decreases through more conclusive information structures, discrimination decreases. This implies that employers respond to information, but that this response is influenced by their motive.

Distinguishing between various forms of discrimination has important implications for designing targeted policy interventions. This paper complements Bohren et al. (2023) who state that when group statistics are equal between two groups, remaining discrimination looks like taste-based discrimination but might actually be belief-based, as discriminators are not aware of the equal group statistics. I ex-

tend their argument by showing that even if discriminators are aware of equal group statistics, remaining discrimination still does not need to be taste-based. Discriminators might believe that particular individuals of one group outperform individuals of another group, and these beliefs might be driven by motives. In light of discrimination based on motivated reasoning, it seems important that policy interventions take wiggle room of information into account to effectively fight discrimination. Since credible individual-level information is often difficult to provide to decision-makers, the potential for discrimination is large in many contexts. As such, contexts in which individual-level information is sometimes not existent (e.g. about taxi or delivery drivers) or is mainly provided by the potentially discriminated (e.g. labor market) are prone to this form of discrimination. Decision-makers might reason that individual-level information from and about a particular individual is only true if it confirms their motive but false if it contradicts their motive. It is therefore important, that information interventions provide means to signal the validity of information, which reduces decision-makers' wiggle room to engage in motivated reasoning.

Even though this study provides first evidence for the existence of discrimination based on motivated reasoning the exact mechanisms of how information is processed such that decision-makers ultimately discriminate remain yet to be studied. Future research could employ process tracing techniques to look more closely at the information processing behavior (Chen and Fischbacher, 2016; Lahey and Oxley, 2016). Future studies could also study the interplay of individual-level and group-level information in discrimination contexts in more detail. It remains unclear to what extent different information acquisition and processing behavior can be linked to concepts like groupiness (Kranton and Sanders, 2017), or other individual characteristics and attitudes. Finally, studying discrimination based on motivated reasoning in the field could yield first evidence of this form of discrimination in real-world settings.

References

- Aigner, D. J. and Cain, G. G. (1977). “Statistical theories of discrimination in labor markets”. *Ilr Review* 30 (2), pp. 175–187.
- Albrecht, K., Von Essen, E., Parys, J., and Szech, N. (2013). “Updating, self-confidence, and discrimination”. *European Economic Review* 60, pp. 144–169.
- Alsan, M., Garrick, O., and Graziani, G. (2019). “Does diversity matter for health? Experimental evidence from Oakland”. *American Economic Review* 109 (12), pp. 4071–4111.
- Arnold, D., Dobbie, W., and Hull, P. (2021). “Measuring racial discrimination in algorithms”. *AEA Papers and Proceedings*. Vol. 111, pp. 49–54.
- Arnold, D., Dobbie, W., and Yang, C. S. (2018). “Racial bias in bail decisions”. *The Quarterly Journal of Economics* 133 (4), pp. 1885–1932.
- Arrow, K. J. (1973). “The Theory of Discrimination”. *Discrimination in Labor Markets*. Ed. by O. Ashenfelter and A. Rees. Princeton NJ: Princeton University Press.
- Barron, K., Ditzmann, R., Gehrig, S., and Schweighofer-Kodritsch, S. (2022). “Explicit and implicit belief-based gender discrimination: A hiring experiment”.
- Bartoš, V., Bauer, M., Chytilová, J., and Matějka, F. (2016). “Attention discrimination: Theory and field experiments with monitoring information acquisition”. *American Economic Review* 106 (6), pp. 1437–75.
- Becker, G. S. (1957). *The economics of discrimination*. Oregon State monographs: Studies in economics, University Press.
- Bertrand, M. and Duflo, E. (2017). “Field experiments on discrimination”. *Handbook of economic field experiments* 1, pp. 309–393.
- Blau, F. D. and Kahn, L. M. (2017). “The gender wage gap: Extent, trends, and explanations”. *Journal of economic literature* 55 (3), pp. 789–865.
- Bohren, A., Imas, A., and Rosenberg, M. (2018). “The Language of Discrimination: Using Experimental versus Observational Data”. *AEA Papers and Proceedings*. Vol. 108, pp. 169–74.
- Bohren, J. A., Haggag, K., Imas, A., and Pope, D. G. (2023). “Inaccurate statistical discrimination: An identification problem”. *Review of Economics and Statistics*, pp. 1–45.
- Bohren, J. A. and Hauser, D. N. (2021). “Learning with heterogeneous misspecified models: Characterization and robustness”. *Econometrica* 89 (6), pp. 3025–3077.
- (2023). *Behavioral Foundations of Model Misspecification*. Tech. rep. Penn Institute for Economic Research, Department of Economics, University of . . .
- Bohren, J. A., Hull, P., and Imas, A. (2022). *Systemic Discrimination: Theory and Measurement*. Tech. rep. National Bureau of Economic Research.
- Bohren, J. A., Imas, A., and Rosenberg, M. (2019). “The dynamics of discrimination: Theory and evidence”. *American economic review* 109 (10), pp. 3395–3436.
- Bordalo, P., Coffman, K., Gennaioli, N., and Shleifer, A. (2016). “Stereotypes”. *The Quarterly Journal of Economics* 131 (4), pp. 1753–1794.
- (2019). “Beliefs about gender”. *American Economic Review* 109 (3), pp. 739–73.

- Bordalo, P., Conlon, J. J., Gennaioli, N., Kwon, S. Y., and Shleifer, A. (2023). “Memory and probability”. *The Quarterly Journal of Economics* 138 (1), pp. 265–311.
- Bursztyn, L., González, A. L., and Yanagizawa-Drott, D. (2020). “Misperceived social norms: Women working outside the home in Saudi Arabia”. *American Economic Review* 110 (10), pp. 2997–3029.
- Charles, K. K. and Guryan, J. (2011). “Studying discrimination: Fundamental challenges and recent progress”. *Annu. Rev. Econ.* 3 (1), pp. 479–511.
- Chen, D. L., Schonger, M., and Wickens, C. (2016). “oTree – An open-source platform for laboratory, online, and field experiments”. *Journal of Behavioral and Experimental Finance* 9, pp. 88–97.
- Chen, F. and Fischbacher, U. (2016). “Response time and click position: Cheap indicators of preferences”. *Journal of the Economic Science Association* 2 (2), pp. 109–126.
- Chen, S. and Heese, C. (2021). *Fishing for good news: Motivated information acquisition*. Tech. rep. University of Bonn and University of Mannheim, Germany.
- Coffman, K., Collis, M. R., and Kulkarni, L. (2023a). “Stereotypes and belief updating”. *Journal of the European Economic Association*, jvad063.
- Coffman, K., Kostyshak, S., and Saygin, P. (2023b). *Choosing and Using Information in Evaluation Decisions*. Tech. rep.
- Coffman, K. B., Exley, C. L., and Niederle, M. (2021). “The role of beliefs in driving gender discrimination”. *Management Science* 67 (6), pp. 3551–3569.
- Cornell, B. and Welch, I. (1996). “Culture, information, and screening discrimination”. *Journal of Political Economy* 104 (3), pp. 542–571.
- Dana, J., Weber, R. A., and Kuang, J. X. (2007). “Exploiting moral wiggle room: experiments demonstrating an illusory preference for fairness”. *Economic Theory* 33 (1), pp. 67–80.
- Di Tella, R., Perez-Truglia, R., Babino, A., and Sigman, M. (2015). “Conveniently upset: avoiding altruism by distorting beliefs about others’ altruism”. *American Economic Review* 105 (11), pp. 3416–42.
- Drobner, C. (2022). “Motivated beliefs and anticipation of uncertainty resolution”. *American Economic Review: Insights* 4 (1), pp. 89–105.
- Duckworth, A. L., Peterson, C., Matthews, M. D., and Kelly, D. R. (2007). “Grit: perseverance and passion for long-term goals.” *Journal of Personality and Social Psychology* 92 (6), p. 1087.
- Enke, B. and Zimmermann, F. (2019). “Correlation neglect in belief formation”. *The Review of Economic Studies* 86 (1), pp. 313–332.
- Epley, N. and Gilovich, T. (2016). “The mechanics of motivated reasoning”. *Journal of Economic Perspectives* 30 (3), pp. 133–40.
- Esponda, I., Oprea, R., and Yuksel, S. (2022). “Discrimination Without Reason: Biases in Statistical Discrimination”.

- Farkas, G. (2003). “Racial disparities and discrimination in education: What do we know, how do we know it, and what do we need to know?” *Teachers College Record* 105 (6), pp. 1119–1146.
- Fershtman, C. and Gneezy, U. (2001). “Discrimination in a segmented society: An experimental approach”. *The Quarterly Journal of Economics* 116 (1), pp. 351–377.
- Festinger, L. (1957). *A theory of cognitive dissonance*. Vol. 2. Stanford university press.
- Goncalves, F. and Mello, S. (2021). “A few bad apples? Racial bias in policing”. *American Economic Review* 111 (5), pp. 1406–41.
- Heckman, J. J. (1998). “Detecting discrimination”. *Journal of economic perspectives* 12 (2), pp. 101–116.
- Kranton, R. E. and Sanders, S. G. (2017). “Groupy versus non-groupy social preferences: Personality, region, and political party”. *American Economic Review* 107 (5), pp. 65–69.
- Kunda, Z. (1990). “The case for motivated reasoning.” *Psychological bulletin* 108 (3), p. 480.
- Lahey, J. N. and Oxley, D. (2016). “The power of eye tracking in economics experiments”. *American Economic Review* 106 (5), pp. 309–13.
- Martínez-Marquina, A., Niederle, M., and Vespa, E. (2019). “Failures in contingent reasoning: The role of uncertainty”. *American Economic Review* 109 (10), pp. 3437–3474.
- Mengel, F. and Campos-Mercade, P. (2021). “Irrational Statistical Discrimination”. *Available at SSRN 3843579*.
- Neumark, D. (2018). “Experimental research on labor market discrimination”. *Journal of Economic Literature* 56 (3), pp. 799–866.
- Neumark, D., Burn, I., and Button, P. (2016). “Experimental age discrimination evidence and the Heckman critique”. *American Economic Review* 106 (5), pp. 303–08.
- Onuchic, P. (2022). “Recent Contributions to Theories of Discrimination”. *arXiv preprint arXiv:2205.05994*.
- Peer, E., Rothschild, D., Gordon, A., Evernden, Z., and Damer, E. (2022). “Data quality of platforms and panels for online behavioral research”. *Behavior Research Methods* 54 (4), pp. 1643–1662.
- Pettigrew, T. F. and Tropp, L. R. (2006). “A meta-analytic test of intergroup contact theory.” *Journal of personality and social psychology* 90 (5), p. 751.
- Phelps, E. S. (1972). “The statistical theory of racism and sexism”. *American economic review* 62 (4), pp. 659–661.
- Rackstraw, E. (2022). “Bias-Motivated Updating in the Labor Market”. *Available at SSRN 4278076*.
- Rammstedt, B. and John, O. P. (2005). “Kurzfassung des big five inventory (BFI-K)”. *Diagnostica* 51 (4), pp. 195–206.

- Rathelot, R. and Safi, M. (2022). “Taste-Based Discrimination”. *Handbook on Economics of Discrimination and Affirmative Action*. Springer, pp. 1–32.
- Reuben, E., Sapienza, P., and Zingales, L. (2014). “How stereotypes impair women’s careers in science”. *Proceedings of the National Academy of Sciences* 111 (12), pp. 4403–4408.
- Sarsons, H. (2017). “Interpreting signals in the labor market: evidence from medical referrals”. *Job Market Paper*, pp. 141–145.
- Sinclair, V. G. and Wallston, K. A. (2004). “The development and psychometric evaluation of the Brief Resilient Coping Scale”. *Assessment* 11 (1), pp. 94–101.
- Stötzer, L. S. and Zimmermann, F. (2022). “A Note on Motivated Cognition and Discriminatory Beliefs”.
- Thaler, M. (2020). “The Fake News Effect: Experimentally Identifying Motivated Reasoning Using Trust in News”. *Available at SSRN 3717381*.
- Zimmermann, F. (2020). “The dynamics of motivated beliefs”. *American Economic Review* 110 (2), pp. 337–61.

Online Appendix

A Experimental Design

This section provides example screenshots of all relevant screens for both treatments and across experiments.

A.1 Instructions

Figure A1: Background

Background

This study is based on a previous survey. In this previous survey, we set up a pool of potential "workers". These workers are US residents who completed an assessment test that consisted of the following three parts:

1. a **logic/math test**, consisting of 15 questions, including simple math exercises as well as questions in which the workers had to find a pattern in a logical sequence of figures in order to find the next figure of the sequence.
(Click [here](#) for an example screenshot.)
2. a test for their level of **altruism** where we provided each worker with an initial endowment of \$5 and asked him/her to indicate how much of these \$5 he/she would like to anonymously give to another worker.
(Click [here](#) for an example screenshot.)
3. and an **effort test** in which we repeatedly presented each worker with a matrix of zeros and ones and asked to indicate how many times the number "1" appears in that matrix. They had 5 minutes to solve as many matrices as possible. However, they always had the opportunity to abort the task and skip the remainder of it by just clicking an "Abort task & skip part" button.
(Click [here](#) for an example screenshot.)

The workers received a **total score** that consists of the following:

| | |
|--|---|
| the number of logic/math questions (out of 15) that they answered correctly | + |
| the amount of dollars (of the maximum of \$5) that they offered in the altruism test | + |
| the number of matrices that the worker solved correctly in the effort task | + |
| Total Score | |

For example, the total score of a worker who answered 7 out of 15 questions correctly in the math/logic test, who gave \$1.50 out of \$5 to the randomly chosen other worker, and who solved 4 matrices correctly got a score of $7 + 1.50 + 4 = 12.50$.

Figure A2: Decision (group Neutral)

Your decision

In the main part of this study, you will play the role of an employer who is in charge of hiring some of the potential workers. You will repeatedly be presented with two workers and it is your task to hire one of them.

Note that we have set up four subgroups of workers, ◆, ■, ●, and ▲ and we have allocated each worker into one of these four groups.

Out of all hiring decisions that you make in this study, we will randomly select one hiring decision to be the "**decision-that-counts**".

Your bonus payment

If you hired the better of the two workers in the decision-that-counts (i.e. the one with the higher total score in the assessment test), you will get a bonus payment of \$2.00. However, if you hired the worse of the two workers in the decision-that-counts (i.e. the one with the lower score in the assessment test), you will get no bonus payment.

Data

All data that you enter in this survey will be treated anonymously and will solely be used for academic research. Neither the experimenter nor anyone else will be able to identify you from the answers that you give in this survey.

If you have read and understood these explanations, please click "Next" to answer **a few test questions** about the information on this page.

Figure A3: Decision (group Race)

Your decision

In the main part of this study, you will play the role of an employer who is in charge of hiring some of the potential workers. You will repeatedly be presented with two workers and it is your task to hire one of them.

Out of all hiring decisions that you make in this study, we will randomly select one hiring decision to be the "**decision-that-counts**".

Your bonus payment

If you hired the better of the two workers in the decision-that-counts (i.e. the one with the higher total score in the assessment test), you will get a bonus payment of \$2.00. However, if you hired the worse of the two workers in the decision-that-counts (i.e. the one with the lower score in the assessment test), you will get no bonus payment.

Data

All data that you enter in this survey will be treated anonymously and will solely be used for academic research. Neither the experimenter nor anyone else will be able to identify you from the answers that you give in this survey.

If you have read and understood these explanations, please click "Next" to answer **a few test questions** about the information on this page.

Figure A4: Test Questions

True or False: The workers completed a logic/math test, an altruism test and an effort test.

- True
- False

True or False: A very good worker to hire is one that has very good logic/math skills, is very altruistic and puts a lot of effort into tasks.

- True
- False

What is the total score of a worker who solved 5 math/logic questions correctly, who gave \$2.44 to the randomly chosen other worker and who solved 3 matrices correctly in the effort task?

How high would your bonus payoff be, if you hired a worker with a total score of 15.3 instead of a worker with a total score of 10.5 in the decision-that-counts?

 \$

How high would your bonus payoff be, if you instead hired a worker with a total score of 2.60 instead of a worker with a total score of 5 in the decision-that-counts?

 \$

A.2 Belief Stage

This section shows screenshots of the belief stage. The screenshots show the elicitation of beliefs about Asian workers. The elicitation of beliefs about workers from the other groups was conducted identically.

Figure A5: Elicitation of Priors (group Neutral)

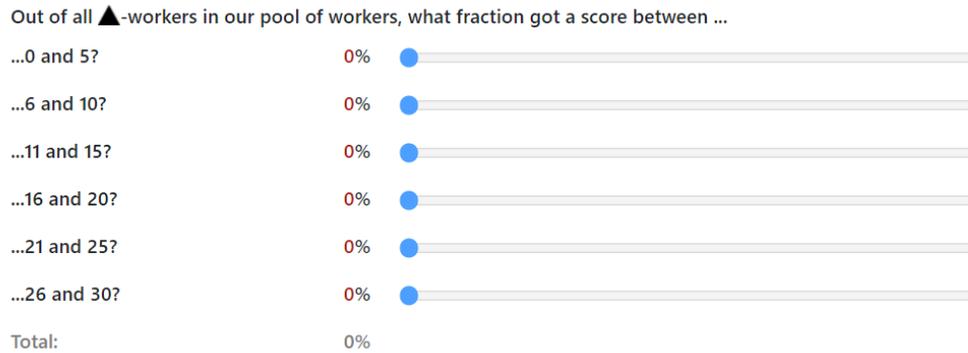


Figure A6: Elicitation of Priors (group Race)

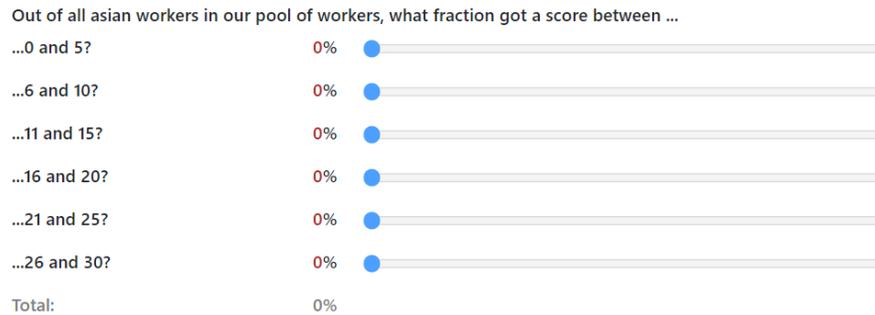
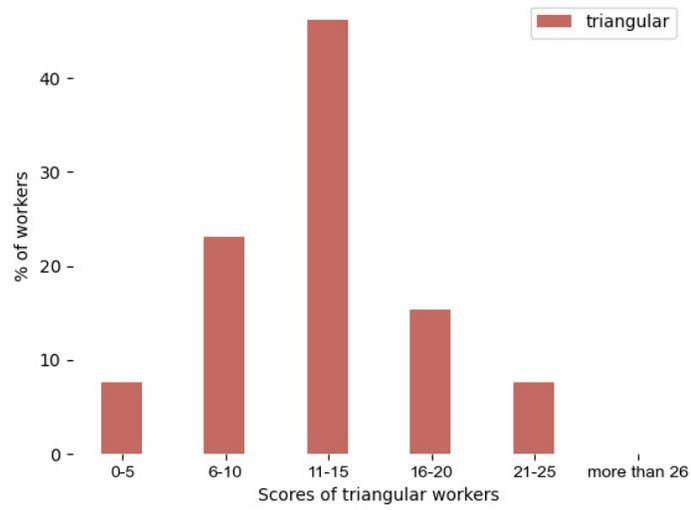


Figure A7: Elicitation of Posteriors (group Neutral)

This graph shows the scores of all ▲-workers in our sample.



Out of all ▲-workers in our pool of workers, what fraction got a score between ...

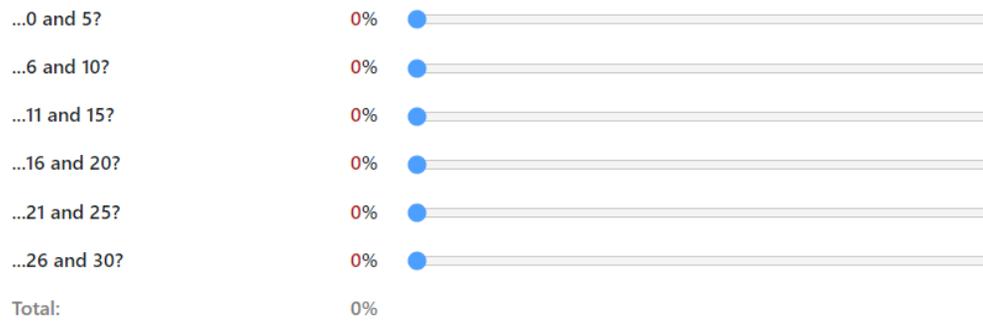
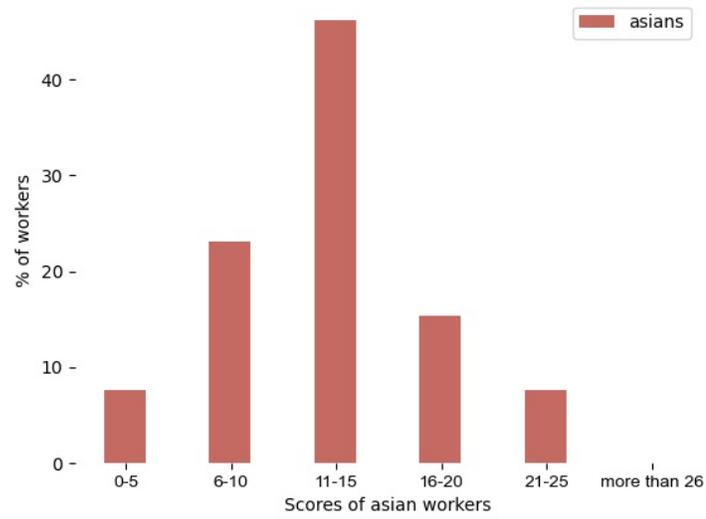
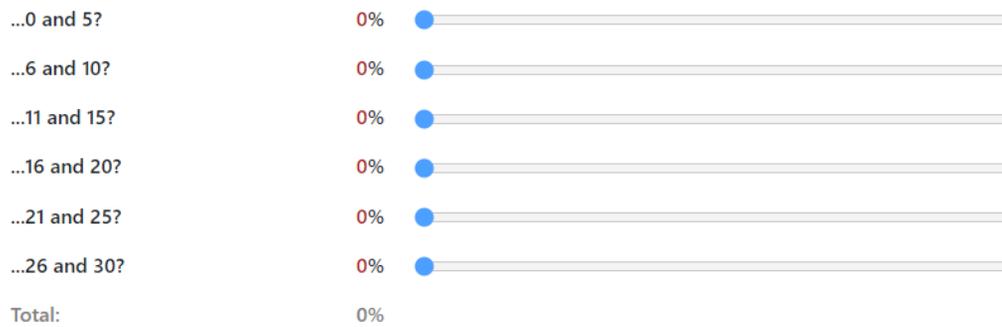


Figure A8: Elicitation of Posteriors (group Race)

This graph shows the scores of all asian workers in our sample.



Out of all asian workers in our pool of workers, what fraction got a score between ...



A.3 Hiring Stage

Figure A9: Hiring Instructions (Experiment ‘No Information’)

You will now play the role of an employer who is in charge of hiring potential workers.

The task: You will repeatedly be presented with a randomly drawn pair of workers from our pool of workers. It is your task to hire one of these two workers.

Your bonus payment: In particular, in the end one of the decisions will be randomly drawn as the **decision-that-counts**. If you hired the better of the two workers in the decision-that-counts (i.e. the one with the **higher total score in the assessment test**), you will get a bonus payment of **\$2.00**. However, if you hired the worse of the two workers in the decision-that-counts (i.e. the one with the **lower score in the assessment test**), you will get **no bonus** payment.

Figure A10: Hiring Instructions (Experiment ‘Ambiguous Information’)

You will now play the role of an employer who is in charge of hiring potential workers.

The task: You will repeatedly be presented with a randomly drawn pair of workers from our pool of workers. It is your task to hire one of the two workers.

Your bonus payment: In the end one of the decisions will be randomly drawn as the **decision-that-counts**. If you hired the **better** of the two workers in the decision-that-counts (i.e. the one with the **higher total score in the assessment test**), you will get a bonus payment of **\$2.00**. However, if you hired the worse of the two workers in the decision-that-counts (i.e. the one with the **lower score in the assessment test**), you will get **no bonus** payment.

Important: On each hiring page, you will be given a message that is meant to tell you who the **better** worker is. However, the message comes from either a **True News** source or a **Fake News** source.

The True News source will always tell you the truth, whereas the Fake News source will never tell you the truth.

If you want, you can request more messages for each decision by clicking on the "**Request Another Message**" button below the two workers.

Whether a message comes from the True News source or the Fake News source is **randomly** determined for each message.

It is up to you how many messages (max. 10) you request and whether or not you follow them.

Figure A11: Hiring Instructions (Experiment ‘Uncertain Information’)

You will now play the role of an employer who is in charge of hiring potential workers.

The task: You will repeatedly be presented with a randomly drawn pair of workers from our pool of workers. It is your task to hire one of the two workers.

Your bonus payment: In the end one of the decisions will be randomly drawn as the **decision-that-counts**. If you hired the **better** of the two workers in the decision-that-counts (i.e. the one with the **higher total score in the assessment test**), you will get a bonus payment of **\$2.00**. However, if you hired the worse of the two workers in the decision-that-counts (i.e. the one with the **lower score in the assessment test**), you will get **no bonus** payment.

Important: On each hiring page, you will be given a message that is meant to tell you who the **better** worker is. However, the message comes from either a **True News** source or a **Fake News** source.

The True News source will always tell you the truth, whereas the Fake News source will never tell you the truth.

If you want, you can request more messages for each decision by clicking on the **"Request Another Message"** button below the two workers.

The likelihood that a message comes from the **True News source is 60%**, the likelihood that a message comes from the **Fake News source is 40%**.

It is up to you how many messages (max. 10) you request and whether or not you follow them.

Figure A12: Hiring Instructions (Experiment ‘Tangible Information’)

You will now play the role of an employer who is in charge of hiring potential workers.

The task: You will repeatedly be presented with a randomly drawn pair of workers from our pool of workers. It is your task to hire one of these two workers.

Your bonus payment: In particular, in the end one of the decisions will be randomly drawn as the **decision-that-counts**. If you hired the better of the two workers in the decision-that-counts (i.e. the one with the **higher total score in the assessment test**), you will get a bonus payment of **\$2.00**. However, if you hired the worse of the two workers in the decision-that-counts (i.e. the one with the **lower score in the assessment test**), you will get **no bonus** payment.

Important: On each hiring page, you will be given a random piece of information about the two workers.

If you want you can request more information for each decision by clicking on the **"Get more information"** button below the two workers.

It is up to you how many pieces of information (max. 10) you request and whether or not you follow them.

Figure A13: Hiring Decision (Experiment 'No Information', group Neutral)

Which of these two workers do you hire?
(Please just click on the shape of the worker who you want to hire.)

 - worker

 - worker

Diamond **Group** **Triangular**

Figure A14: Hiring Decision (Experiment 'No Information', group Race)

Which of these two workers do you hire?
(Please just click on the worker who you want to hire.)

 **Juan**

 **Nansi**

Hispanic or Latin **Race** **Asian**

Figure A15: Hiring Decision (Experiment 'Ambiguous Information' & 'Uncertain Information', group Neutral)

Which of these two workers do you hire?
(Please just click on the shape of the worker who you want to hire.)

| | | |
|--|--------------|---|
|  - worker | |  - worker |
| Diamond | Group | Triangular |

Message 1
The better worker is:

 - worker

[Request Another Message](#)

Figure A16: Hiring Decision (Experiment 'Ambiguous Information' & 'Uncertain Information', group Race)

Which of these two workers do you hire?
(Please just click on the worker who you want to hire.)

| | | | | |
|---|-------------|-------------|--|--------------|
|  | Juan | |  | Nansi |
| Hispanic or Latin | | Race | | Asian |

Message 1
The better worker is:

 **Juan**

[Request Another Message](#)

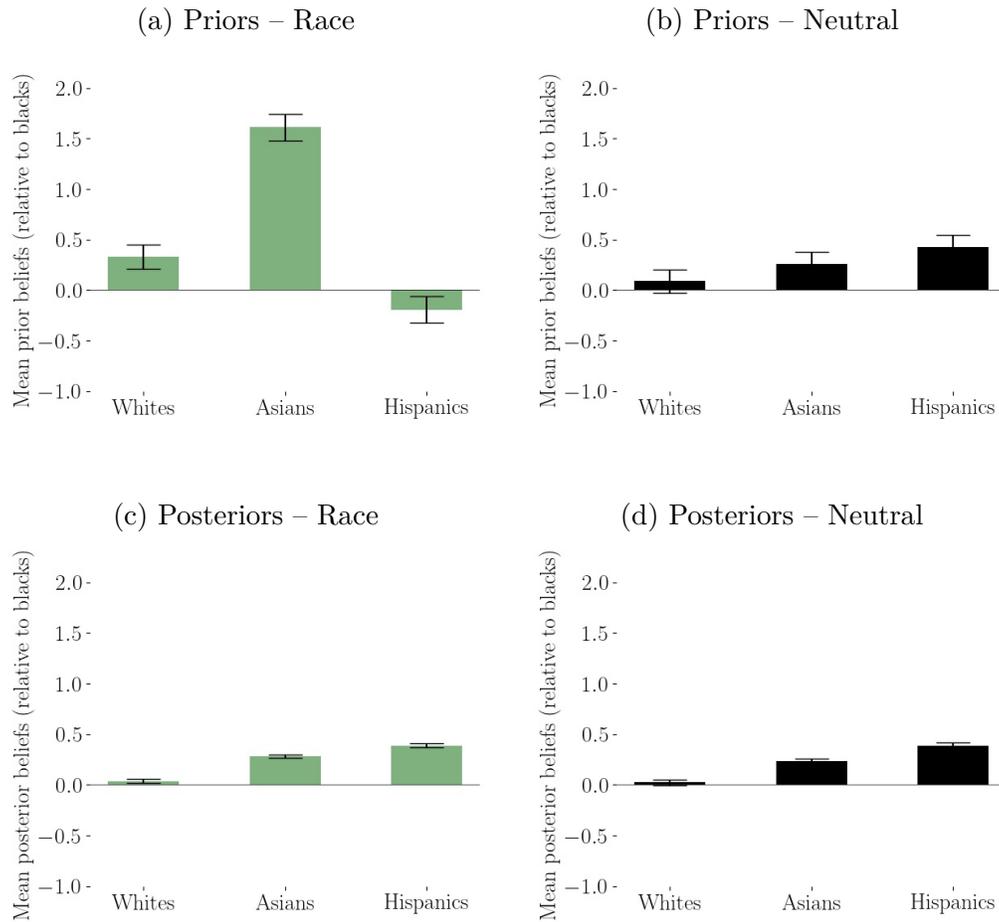
A.4 Deviation from Preregistration

The sample sizes of the different experiments deviate from what has been specified in the preregistration file. While the preregistration aimed at 500 to 600 employers per experiment, especially the experiments with decreased wiggle room (‘Uncertain Information’) and (‘Tangible Information’) include less participants than preregistered. While this is due to resource constraints, the analyses of the hypotheses are well-powered with current sample sizes.

B Group-level Beliefs

Figure B1 shows the relation between mean beliefs about black, white, Asian and Hispanic workers and provides evidence, that the largest difference in prior beliefs, and hence in potential for motivated beliefs, lies in the difference between Asian and Hispanic workers.

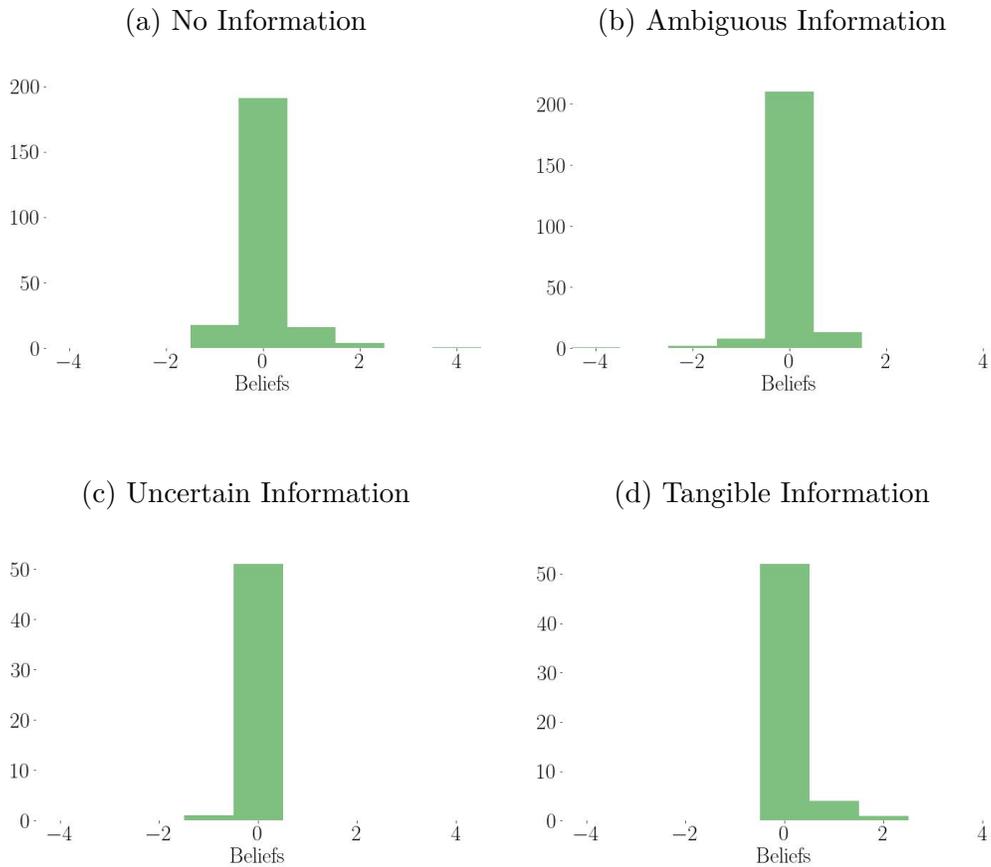
Figure B1: Mean beliefs about group productivities



Notes: Figures (a) and (b) show the mean beliefs before the group-level information update of employers in group Race and group Neutral, respectively, figures (c) and (d) show the mean beliefs after the group-level information update of employers in group Race and group Neutral, respectively. Beliefs are plotted as the difference in means compared to beliefs about blacks. Error bars show the standard errors of the means.

Figure B2 shows differences in differences in posterior beliefs about Asian and Hispanic workers between treatment Race and treatment Neutral for each experiment separately. More specifically, I calculate the differences in mean beliefs about Asian and Hispanic workers for each employer of treatment Race and subtract the difference of the mean of the mean beliefs about Asian and Hispanic workers of employers in treatment Neutral. This is a measure of heterogeneity in posterior beliefs and shows that beliefs were homogeneously updated across experiments.

Figure B2: Heterogeneity in posterior beliefs



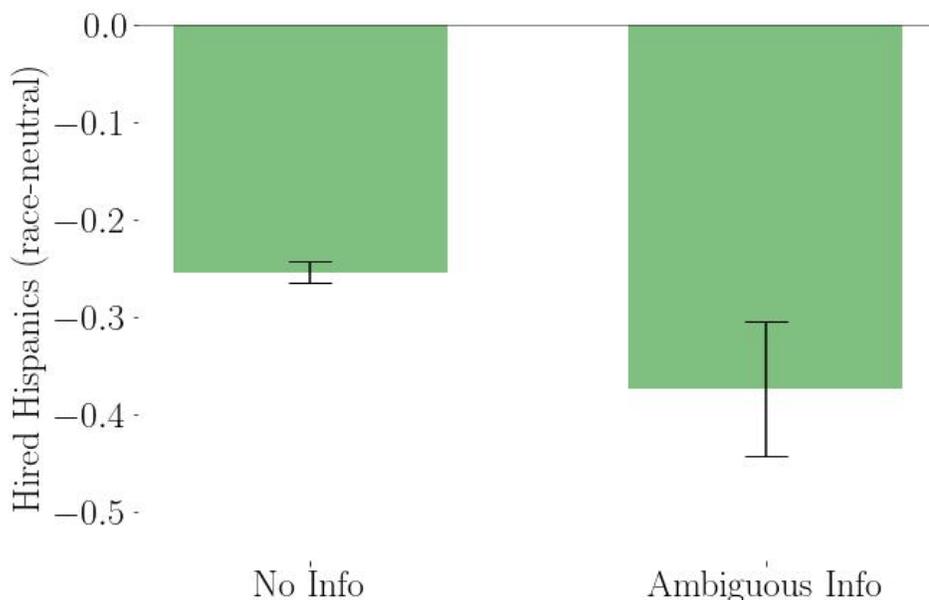
Notes: The figures show histograms of the differences in differences between employer beliefs about Asian and Hispanic workers, and between both treatments. The calculations uses the individual posteriors of treatment Race and the mean posteriors of treatment Neutral, as no employer participated in both treatments. Negative values imply higher beliefs for Asians than for Hispanics in treatment Race than in treatment Neutral.

C Results

This section provides additional results. Section C.1 shows results from decisions in the ‘Ambiguous Information’ experiment, when individuals requested one additional signal and ended up with two conflicting signals. Section C.2 presents results of a heterogeneity analysis for hypotheses 1b, 1c, and 2a with respect to different prior beliefs. Section C.3 provides results of the counter-hypotheses for hypothesis 1c and hypothesis 2a. Section C.4 adds the numbers to Figure 3 of the main text. Finally, Section C.5 provides results of some robustness exercises.

C.1 Discrimination under Conflicting Signals

Figure C1: Hiring rates of Hispanic workers under wiggle room



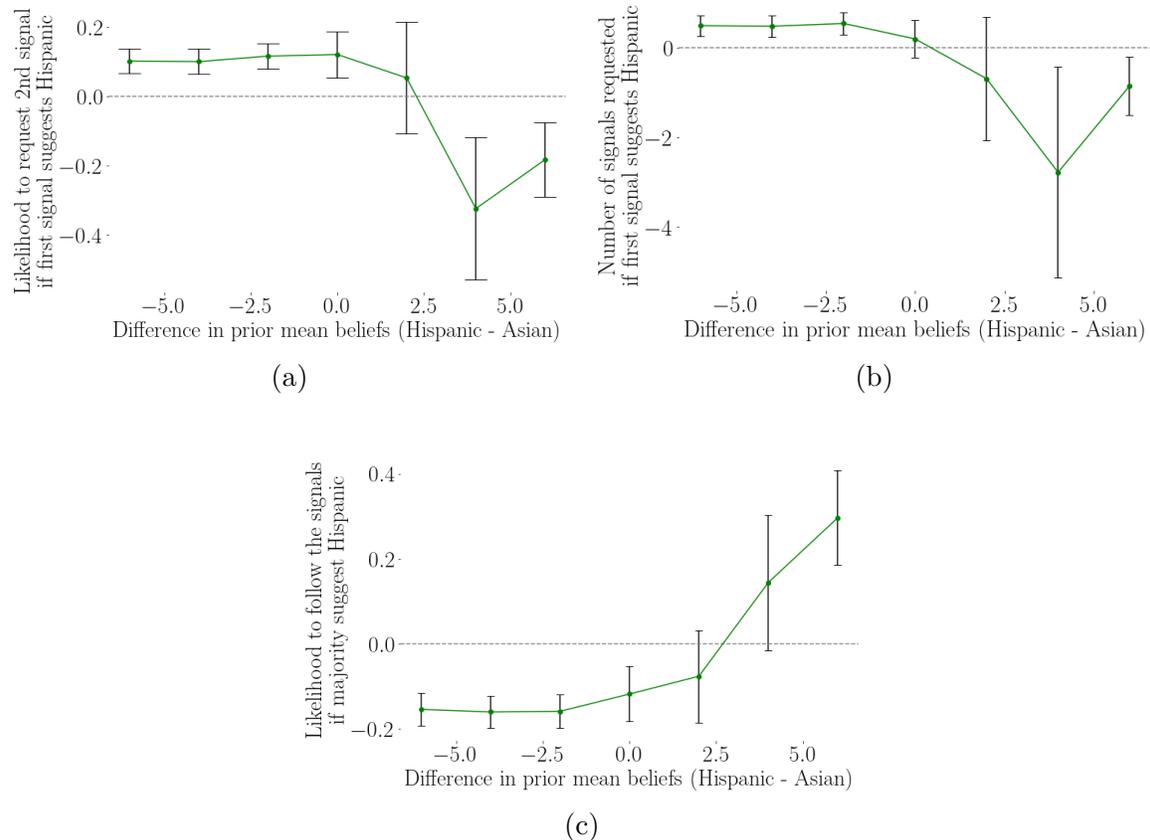
Notes: The vertical axis displays the fraction of hired Hispanics in treatment Race minus the fraction of hired Hispanics in treatment Neutral. The horizontal axis groups the decisions by experiments (‘No Information’, ‘Ambiguous Information’). Decisions in the ‘Ambiguous Information’ experiment only include decisions in which subjects only requested one additional signal and ended up with two conflicting signals. Error bars indicate standard errors of the means.

Figure C1 provides further support to the idea that individuals misperceived the informativeness of the ambiguous signals to align with their motives to form the

belief that the Asian worker is more productive. Given that subjects believe there is a 50% chance of Fake News for each signal, in the absence of motivated reasoning, decisions in the ‘Ambiguous Information’ experiment with two conflicting signals (i.e., one suggesting the Hispanic worker and one suggesting the Asian worker) should be similar to decisions in the No Information experiment. Instead, Figure C1 shows that there is significantly more discrimination in the ‘Ambiguous Information’ experiment in this case.

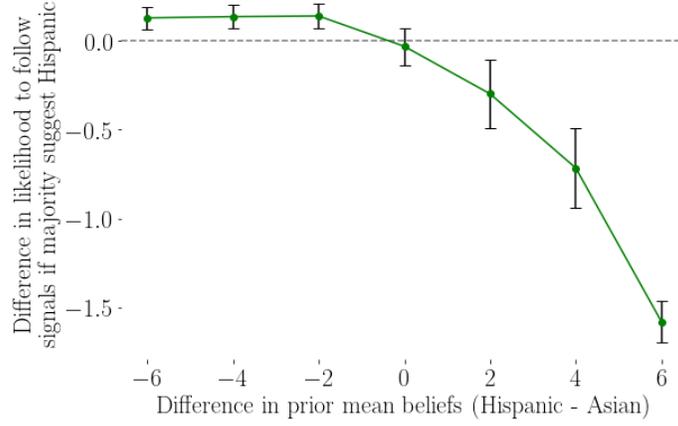
C.2 Heterogeneity of Effects with respect to Prior Beliefs

Figure C2: Ambiguous Information - Heterogeneity with respect to Prior Beliefs (Hypotheses 1b & 1c)

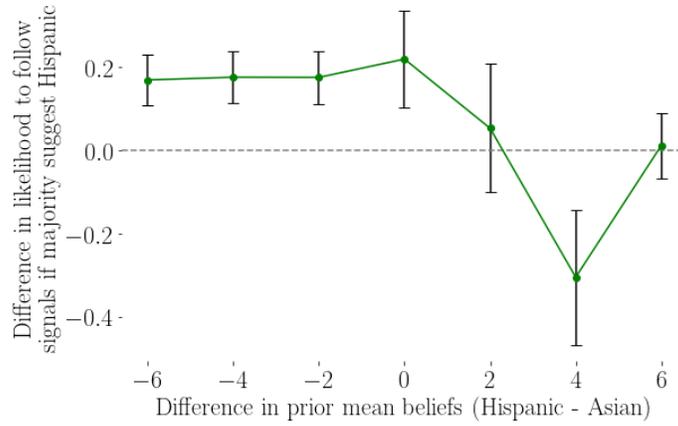


Notes: The vertical axis in panel (a) & (b) displays the regression coefficients as specified in the main analysis for hypotheses 1b, the vertical axis in panel (c) displays those for hypothesis 1c. The horizontal axis shows differences in mean prior beliefs about workers’ productivity scores where positive values indicate a higher prior belief for Hispanic workers than for Asian workers. Error bars indicate standard errors of the coefficient.

Figure C3: Heterogeneity with respect to prior beliefs (Hypothesis 2a)



(a) Uncertain - Ambiguous



(b) Tangible - Ambiguous

Notes: The vertical axis displays the regression coefficients as specified in the main analysis for hypothesis 2a. The horizontal axis shows differences in mean prior beliefs about workers' productivity scores where positive values indicate a higher prior belief for Hispanic workers than for Asian workers. Error bars indicate standard errors of the coefficient.

Figures C2 and C3 show that the results of the main analysis hold for subjects who initially believed that Asian workers were more productive than Hispanic workers and go into the opposite direction for subjects who initially believed that Hispanic workers had a higher productivity. The extent of these results increases with the absolute belief differences, especially for hypotheses 1c (Figure C2c) and 2a (Figure C3a). This is consistent with the idea of motivated reasoning and lends further support to the tested hypotheses.

C.3 Counter-hypotheses

To give more insights into the data structure this section provides results for the following “counter-hypotheses” for hypothesis 1c and hypothesis 2a:

Counter-hypothesis 1c: *If the majority of all considered signals in a decision suggests to hire the Asian worker in the ‘Ambiguous Information’ experiment, employers in treatment Race are more likely to act according to the suggestion than employers in treatment Neutral.*

Counter-hypothesis 2a: *If the majority of all considered signals in a decision suggests to hire the Asian worker, the difference in the likelihood to follow the suggestion between treatment Race and treatment Neutral is larger in the ‘Ambiguous Information’ experiment than in the ‘Uncertain Information’ and the ‘Tangible Information’ experiments.*

Table C1: Results of the Counter-hypotheses for Hypotheses 1c and 2a

| | Ambiguous Information | Uncertain - Ambiguous | Tangible - Ambiguous |
|-------------------------|-----------------------|-----------------------|----------------------|
| | (1) | (2) | (3) |
| | follow signals | follow signals | follow signals |
| race * asian | 0.1861 (0.0381) | -0.1198 (0.0585) | -0.1811 (0.0583) |
| race | -0.1047 (0.0294) | | |
| asian | -0.0135 (0.0265) | | |
| Observations | 3290 | | |
| Baseline mean dep. var. | 0.8423 | | |

Notes: *follow signals* is a dummy equal to 1 if an employer’s hiring decision followed the majority of considered signals. In column (1) employers received ambiguous individual-level information about the two workers. This shows results for the counter-hypothesis of hypothesis 1c. Column (2) and (3) refer to the differences in coefficients for *follow signals* in experiments ‘Uncertain Information’ and ‘Ambiguous Information’ (column 2), and ‘Tangible Information’ and ‘Ambiguous Information’ (column 3). Together they test the counter-hypothesis of hypothesis 2a. *race* is the treatment dummy equal to 1 if the decision was made in treatment treatment Race. *asian* is a dummy equal to 1 if the majority of all considered signals suggests to hire the Asian worker. This holds for both treatments even though employers in treatment Neutral did not observe the race of the workers. Units of observation are decision specific. Standard errors are clustered at the individual level and displayed in parentheses.

These counter-hypotheses compare the subjects that stopped requesting signals when the majority of the received signals suggested to hire the Asian worker (instead

of the Hispanic worker).

The results in Table C1 confirm the “counter-hypotheses” and hence lend further support for the motivated information acquisition behavior described in the main text.

C.4 Discrimination

Table C2 shows discrimination rates across experiments and thereby provides the numbers for Figure 3 of the main text.

Table C2: Discrimination rates across experiments

| | <i>Dep. var: hired hispanic</i> | | | |
|-------------------------|---------------------------------|------------------------|---------------------|--------------------|
| | (1) | (2) | (3) | (4) |
| | no info | ambiguous info | uncertain info | tangible info |
| race | -0.2536*** (0.0301) | -0.1392*** (0.0229) | -0.0614 (0.0408) | 0.0281 (0.0302) |
| Observations | 3178 | 3290 | 756 | 742 |
| Baseline mean dep. var. | 0.5022 | 0.4994 | 0.5230 | 0.4606 |

Notes: In column (1) employers did not receive any individual-level information regarding the two workers, in column (2) they received ambiguous information, in column (3) uncertain information, and in column (4) tangible information. The dependent variable in both models is a dummy equal to 1 if the Hispanic worker was hired and 0 if the Asian worker was hired. ‘race’ is the treatment dummy and equal to 1 if the decision was made in treatment group race (showing the respective races) and 0 otherwise. Standard errors are clustered at the individual level and displayed in parentheses.

C.5 Robustness Checks

During the Instructions, participants in all four experiments had to answer a few training questions in order to make sure that instructions were understood. While the main text includes all participants that answered at least 3 out of 5 of these questions correctly, this section shows results for stricter limits. Section C.5.1 shows results for participants who answered at least 4 out of 5 questions correctly, Section C.5.2 for those who answered all test questions correctly. The results of the main text are robust to these variations.

C.5.1 At least 4 test questions correct

Table C3: Hiring Rates of Hispanic workers across experiments

| | (1) No info | (2) Ambiguous info | (3) Uncertain info | (4) Tangible info |
|-------------------------|---------------------|-----------------------|-----------------------|----------------------|
| race | -0.2491 (0.0330) | -0.1328 (0.0243) | -0.0674 (0.0448) | 0.0102 (0.0319) |
| Observations | 2702 | 2917 | 663 | 586 |
| Baseline mean dep. var. | 0.4921 | 0.4960 | 0.5228 | 0.4728 |

Notes: This table includes decisions of participants who completed at least 4 test questions correctly. In column (1) employers did not receive any individual-level information regarding the two workers, in column (2) they received ambiguous information, in column (3) uncertain information, and in column (4) tangible information. The dependent variable in both models is a dummy equal to 1 if the Hispanic worker was hired and 0 if the Asian worker was hired. 'race' is the treatment dummy and equal to 1 if the decision was made in treatment group race (showing the respective races) and 0 otherwise. Standard errors are clustered at the individual level and displayed in parentheses.

Table C4: Information Behavior across Experiments

| (a) | Ambiguous Information | | |
|-------------------------|-----------------------|---------------------|---------------------|
| | (1) | (2) | (3) |
| | 2nd signals | number of signals | follow signals |
| race * hispanic | 0.0945 (0.0389) | 0.4630 (0.2526) | -0.1736 (0.0430) |
| race | -0.0572 (0.0458) | -0.3650 (0.2988) | 0.0550 (0.0210) |
| hispanic | -0.0268 (0.0271) | -0.2019 (0.1792) | -0.0752 (0.0280) |
| Observations | 2919 | 2919 | 2919 |
| Baseline mean dep. var. | 0.5496 | 3.6684 | 0.8740 |

| (b) | Uncertain Information | | | Uncertain - Ambiguous |
|-------------------------|-----------------------|---------------------|---------------------|-----------------------|
| | (1) | (2) | (3) | (4) |
| | 2nd signals | number of signals | follow signals | follow signals |
| race * hispanic | 0.1437 (0.0634) | 0.6262 (0.5870) | -0.0501 (0.0682) | 0.1235 (0.0620) |
| race | -0.0400 (0.0836) | -0.0779 (0.8118) | 0.0296 (0.0258) | |
| hispanic | -0.0533 (0.0468) | 0.1056 (0.3979) | -0.0478 (0.0329) | |
| Observations | 665 | 665 | 665 | |
| Baseline mean dep. var. | 0.7738 | 5.8571 | 0.9432 | |

| (c) | Tangible Information | | | Tangible - Ambiguous |
|-------------------------|----------------------|--------------------|---------------------|----------------------|
| | (1) | (2) | (3) | (4) |
| | 2nd signals | number of signals | follow signals | follow signals |
| race * hispanic | 0.0260 (0.0318) | 0.5709 (0.3937) | -0.0206 (0.0492) | 0.1675 (0.0620) |
| race | 0.0381 (0.0374) | 0.8460 (0.5347) | -0.0328 (0.0287) | |
| hispanic | 0.0051 (0.0235) | 0.1056 (0.3137) | -0.0860 (0.0365) | |
| Observations | 742 | 742 | 742 | |
| Baseline mean dep. var. | 0.9135 | 5.5240 | 0.8894 | |

Notes: This table includes decisions of participants who completed at least 4 test questions correctly. In panel (a) employers received ambiguous information, in panel (b) uncertain information, and in panel (c) tangible information. *2nd signal*, is a dummy equal to 1 if an employer requested a second signal. *number of signals*, counts the number of requested signals. *follow signals*, is a dummy equal to 1 if an employer's hiring decision followed the majority of considered signals. *race* is the treatment dummy equal to 1 if the decision was made in treatment group *race*. *hispanic* is a dummy equal to 1 if the initial signal (columns 1 and 2) or the direction of all considered signals (column 3) suggests to hire the Hispanic worker. Units of observation are decision specific. Standard errors are clustered at the individual level and displayed in parentheses.

C.5.2 All 5 test questions correct

Table C5: Hiring Rates of Hispanic workers across experiments

| | <i>Dep. var: hired hispanic</i> | | | |
|-------------------------|---------------------------------|---------------------|---------------------|--------------------|
| | (1) | (2) | (3) | (4) |
| | no info | ambiguous info | uncertain info | tangible info |
| race | -0.2599 (0.0413) | -0.1665 (0.0303) | -0.0625 (0.0599) | 0.0131 (0.0388) |
| Observations | 1631 | 1869 | 448 | 364 |
| Baseline mean dep. var. | 0.4820 | 0.5143 | 0.4955 | 0.4845 |

Notes: This table includes decisions of participants who completed all 5 test questions correctly. In column (1) employers did not receive any individual-level information regarding the two workers, in column (2) they received ambiguous information, in column (3) uncertain information, and in column (4) tangible information. The dependent variable in both models is a dummy equal to 1 if the Hispanic worker was hired and 0 if the Asian worker was hired. ‘race’ is the treatment dummy and equal to 1 if the decision was made in treatment group race (showing the respective races) and 0 otherwise. Standard errors are clustered at the individual level and displayed in parentheses.

Table C6: Information Behavior across Experiments

| (a) | Ambiguous Information | | |
|-------------------------|-----------------------|---------------------|---------------------|
| | (1) | (2) | (3) |
| | 2nd signals | number of signals | follow signals |
| race * hispanic | 0.0873 (0.0485) | 0.4469 (0.3259) | -0.1984 (0.0549) |
| race | -0.0111 (0.0577) | -0.0960 (0.3964) | 0.0846 (0.0247) |
| hispanic | 0.0043 (0.0321) | -0.1242 (0.2189) | -0.0608 (0.0343) |
| Observations | 1869 | 1869 | 1869 |
| Baseline mean dep. var. | 0.5210 | 3.6806 | 0.8669 |

| (b) | Uncertain Information | | | Uncertain - Ambiguous |
|-------------------------|-----------------------|--------------------|---------------------|-----------------------|
| | (1) | (2) | (3) | (4) |
| | 2nd signals | number of signals | follow signals | follow signals |
| race * hispanic | 0.0860 (0.0752) | 0.0222 (0.7070) | -0.0385 (0.0731) | 0.1600 (0.0917) |
| race | 0.0154 (0.0958) | 0.6590 (0.9525) | -0.0072 (0.0202) | |
| hispanic | -0.0120 (0.0552) | 0.5375 (0.4779) | -0.0931 (0.0440) | |
| Observations | 448 | 448 | 448 | |
| Baseline mean dep. var. | 0.7826 | 5.7652 | 0.9754 | |

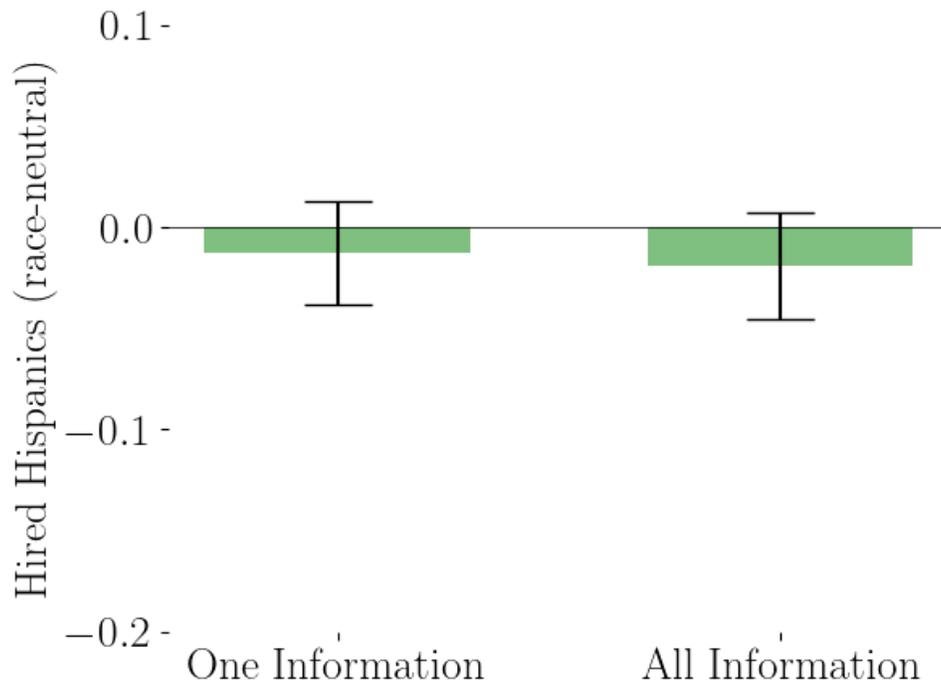
| (c) | Tangible Information | | | Tangible - Ambiguous |
|-------------------------|----------------------|---------------------|---------------------|----------------------|
| | (1) | (2) | (3) | (4) |
| | 2nd signals | number of signals | follow signals | follow signals |
| race * hispanic | 0.0568 (0.0503) | 0.9549 (0.5274) | -0.1000 (0.0670) | 0.0984 (0.0834) |
| race | 0.0486 (0.0664) | 0.7249 (0.7616) | 0.0134 (0.0402) | |
| hispanic | 0.0021 (0.0291) | -0.1436 (0.3887) | -0.0124 (0.0496) | |
| Observations | 364 | 364 | 364 | |
| Baseline mean dep. var. | 0.8812 | 5.6436 | 0.8426 | |

Notes: This table includes decisions of participants who completed all 5 test questions correctly. In panel (a) employers received ambiguous information, in panel (b) uncertain information, and in panel (c) tangible information. *2nd signal*, is a dummy equal to 1 if an employer requested a second signal. *number of signals*, counts the number of requested signals. *follow signals*, is a dummy equal to 1 if an employer's hiring decision followed the majority of considered signals. *race* is the treatment dummy equal to 1 if the decision was made in treatment group *race*. *hispanic* is a dummy equal to 1 if the initial signal (columns 1 and 2) or the direction of all considered signals (column 3) suggests to hire the Hispanic worker. Units of observation are decision specific. Standard errors are clustered at the individual level and displayed in parentheses.

D Additional Experiments

Since neither the One Information experiment nor the All Information experiment allow to study the information acquisition and processing behavior without any process tracing techniques, results of the ‘One Information’ and ‘All Information’ experiments are not included in the main text but mainly serve as additional evidence and are briefly summarized in Figure D1 below. Neither of the experiments shows rates of discrimination, which is consistent with the limited wiggle room in both experiments.

Figure D1: Hiring rates of Hispanic workers under wiggle room



Notes: The vertical axis displays the fraction of hired Hispanics in group Race minus the fraction of hired Hispanics in group Neutral. The horizontal axis groups the decisions by experiments, either the ‘One Information’ experiment or the ‘All Information’ experiment. Error bars indicate standard errors of the means.