Exploring Preferences for 'Blinding' One's Own Judgment

by

Sean Fath

Business Administration
Duke University

Date:_____

Approved:

_____
Richard Larrick, Advisor

_____
Aaron Kay

_____
Ashleigh Shelby Rosette

_____
Jane Risen

Dissertation submitted in partial fulfillment of
the requirements for the degree of Doctor
of Philosophy in Business Administration
in the Graduate School
of Duke University

2020

ABSTRACT

Exploring Preferences for 'Blinding' One's Own Judgment

by

Sean Fath

Business Administration
Duke University

Date:_____
Approved:

_____
Richard Larrick, Advisor

_____
Aaron Kay

_____
Ashleigh Shelby Rosette

_____
Jane Risen

An abstract of a dissertation submitted in partial
fulfillment of the requirements for the degree of Doctor
of Philosophy in Business Administration
in the Graduate School
of Duke University

2020

Abstract

I investigate people's degree of preference for "blinding" in decision-making: purposefully restricting the information one sees in order to try to form a more accurate evaluation. For example, when grading her students' papers, a professor might choose to "blind" herself to students' names by anonymizing them, and thus evaluate the papers on content alone. I propose a theoretical framework of individual-level blinding preferences, outlining various factors that may drive evaluators' choices to see or blind themselves to potentially biasing information in an impending evaluation. Next, I discuss 8 studies (N = 5,350) and associated replications (N = 3,720) that (a) explore individuals' preferences for blinding and outline consequences for bias, (b) test the mechanisms driving blinding decisions proposed in my theoretical framework, and (c) explore the efficacy of multiple interventions to encourage a choice to blind one's judgment. I find that people often choose to see potentially biasing information rather than be blind to it, even though they acknowledge they should be blind and that seeing such information will likely bias their evaluations. I also find that interventions that facilitate deliberative reflection before a blinding choice is made can encourage a choice to be blind. I discuss contributions of these studies to research on mental contamination, inequality reduction in organizations, and social perception, as well as implications of these studies for groups concerned with members' decision bias.

Dedication

This dissertation is dedicated to my brother, Graham Fath, who always tried to do the right thing.

Contents

List of Tables

List of Figures

Acknowledgements

I thank Rick Larrick and Aaron Kay, my primary advisors during my time at Fuqua, who taught me how to think and write carefully. Thanks also to Ashleigh Shelby Rosette, Jane Risen, and Jack Soll, who helped me make this work more complete, and to my family, who supported me these past six years without ever really understanding what it is I do every day. Finally, I am eternally thankful for my wife, Devon Proudfoot, who helped me be the best version of myself possible and believed in me when I did not.

1. Introduction

During the first half of the 20th century, the percentage of female members in the foremost symphony orchestras in the United States—like the New York Philharmonic—was under 10%. As documented by Goldin and Rouse (2000), anti-female bias was widespread in the music world—one director notoriously said, "I just don't think women should be in an orchestra." Then, around 1965, the percentage started to increase dramatically. By the mid-1990's, women made up approximately 25% of the membership in each of the top 5 symphony orchestras in the US and represented about 50% of new hires (Goldin & Rouse, 2000). What happened in the 1960s to account for this major change? Throughout that decade, audition policies changed at most top US orchestras: Performers now auditioned from behind a screen. This new policy rendered audition judges and committees "blind" to the gender of people auditioning.

This example illustrates the usefulness of "blinding"—purposefully restricting the information incorporated into a decision or evaluation in order to minimize the risk of being influenced adversely (Robertson & Kesselheim, 2016; Sah et al., 2016; Gilbert, 1993). Blinding has been useful as a structural solution to bias in many domains, from clinical trials of new drugs (Patsy & Prentice, 2010) to peer review in academia (Blank, 1991). In these cases, an external party adopts or imposes a system of blinding to combat bias.

In the absence of systematic policies, however, there are open questions about individuals' preferences for and beliefs about blinding: Do people understand that blinding can improve their own decisions? Will they adopt it for themselves? For instance, a hiring manager has the option to examine the social media profiles of job applicants before evaluating them. An 8th grade teacher can ask a 7th grade teacher about her incoming students. In such situations, will people choose to avoid such information if they believe it could alter their objectivity? In this paper, I[1] explore these questions. First, I discuss the benefits of blinding as an organizational policy. Next, I propose a theoretical framework that organizes several different factors I theorize will inform individuals' own preferences for blinding. Finally, I present eight studies that are, to my knowledge, the first tests of individual-level blinding preferences.

1.1 Background

1.1.1 Mental Contamination

Blinding—sometimes known alternatively as "exposure control" (Gilbert, 1993) or "masking" (Rissing & Castilla, 2016)—helps when it shields a decision from being contaminated by potentially biasing information (Wilson & Brekke, 1994). Information can be contaminative—that is, it can bias or distort evaluations—in multiple ways. Some information, like the gender of an orchestra applicant, is clearly irrelevant and can

[1] Many of the studies in this dissertation are a result of my collaboration with Rick Larrick and Jack Soll. Both Rick Larrick and Jack Soll were instrumental in the development of the theoretical rationale undergirding these studies, their design, and the interpretation of results. This dissertation would not exist without them.

contaminate evaluations if not removed (Goldin & Rouse, 2000). Certain information may be relevant to an evaluation but still carry a risk of biasing that evaluation. Teachers may face this issue when grading an exam that contains multiple-choice and essay sections. Performance on the two tasks is likely to be correlated; the multiple-choice score is a valid cue to essay quality. However, knowledge of the multiple-choice score when grading the essay may bias how the teacher resolves ambiguities in the student's writing. To avoid "double-counting," it is better to grade each section independently before aggregating them.

Generally, the key metacognitive insight is that the sequence of interpreting new information is critical. If an evaluator's goal is to objectively evaluate a new sample of performance—rather than predict a future level of performance—information about past performance (or other background information) can distort the evaluation of that new sample of performance through a process of expectancy-assimilation (Jones, 1986). Initial expectations can lead evaluators to interpret ambiguous evidence in a way that supports the initial belief. For this reason, so long as the new sample of performance can be evaluated (i.e., the new sample of performance is not so noisy as to be impossible to evaluate), the best way to perform an objective evaluation is to first evaluate the new sample of performance independent of any other information. In other words, if an evaluator's goal is to objectively evaluate a new sample of performance (e.g., grade a student's paper), any information that is not that sample of performance (e.g., student's

name, student's grade on a previous paper) can be contaminative until an independent evaluation of the new sample of performance is completed. Hence, being blind to all information that is not the focal sample of performance to be evaluated is an optimal strategy to preserve an objective evaluation of that sample of performance (Axt & Lai, 2019).

1.1.2 Consequences of Contamination in Organizational Contexts

Some studies demonstrate the value of blinding for evaluations of performance. For instance, in the aforementioned study by Goldin & Rouse (2000), the adoption of blinding procedures by major U.S. symphony orchestras facilitated objective evaluations of performers' musical ability by reducing expectancy-assimilation effects consistent with an anti-female bias (Jones, 1986). A large-scale study of review procedures at the academic journal American Economic Review found that "double blind" reviews, where those reviewing article submissions were unaware of (i.e., blind to) the identity of the submissions' authors, often resulted in more favorable outcomes for submissions with low-prestige authors, and less favorable outcomes for submissions with high-prestige authors, than "single blind" reviews, where reviewers were aware of authors' identities (Blank, 1991). In short, the adoption of blinding procedures facilitates objective evaluations of new samples of performance.

Conversely, without blinding policies in place, important evaluations in organizations, such as hiring decisions, can become susceptible to bias. For instance, job

seekers who submit resumes with White-sounding names are perceived as more qualified and are more likely to receive a call-back compared to those with identical resumes who have African American-sounding names (Bertrand & Mullainathan, 2004). Moreover, job applicants with White-sounding names receive a higher return on added years of experience, in terms of call-backs, than applicants with African American-sounding names (Bertrand & Mullainathan, 2004). Similarly, an audit study of science faculty at American universities found that male applicants for a lab manager position were more likely to be hired, offered higher salaries, and received more opportunities for mentorship than female applicants with identical credentials (Moss-Racusin, Dovidio, Brescoll, Graham, & Handelsman, 2012). These audit studies demonstrate the ways in which decision bias—a result of mental contamination—can directly fuel inequality in the workplace.

1.1.3 Strategies for Inequality Reduction in Organizations

Blinding policies are one example of the multitude of approaches organizations have adopted to reduce workplace inequality and encourage diversity among their membership. These initiatives—which were started with the passage of the Civil Rights Act but have become increasingly common in recent years (Tolbert & Castilla, 2016)—generally fall into three broad categories (Kalev, Dobbin, & Kelly, 2006). Some approaches focus on structural changes, such as the creation of positions and hiring opportunities that are explicitly intended to increase workforce diversity (e.g.,

Affirmative Action policies). Other approaches try to use training and accountability programs to reduce stereotyping and bias in organizational evaluations (e.g., hiring decisions). Still other approaches focus on helping members of subordinated social categories expand, and get more out of, their workplace social networks (Kalev et al., 2006).

A longitudinal study of the efficacy of these different approaches found that those of the first type, which focus on structural change in organizations (e.g., the creation of affirmative action hiring plans or diversity focused taskforces), are the most effective in increasing the proportions of women and racial minority employees in a given organization over time (Kalev et al., 2006). As a comparison, diversity initiatives aimed at reducing bias through enhanced training show null to negative effects on diversity propagation in the long run (Kalev et al., 2006), sometimes instead fostering pushback and backlash among employees being "trained" (Dobbin, Schrage, & Kalev, 2015). Indeed, some research suggests that initiatives that aim to encourage meritocracy in organizations, or instill a meritocratic mindset in evaluators (e.g., bias reduction training), may actually, ironically, increase trainees' bias in favor of members of traditionally advantaged social categories (Castilla & Benard, 2010).

Blinding policies represent an inequality reduction mechanism of the structural variety—the scope of information that can possibly be included in an evaluation, like a hiring decision, is altered by a structural change to the evaluation itself (e.g., the

stripping of names from resumes to reduce bias in hiring). Notably, research suggests that structural change initiatives that do not have buy-in from the managers and evaluators whose decision-making they affect are less effective than those that do (Dobbin et al., 2015). For instance, while policy level changes to hiring procedures, such as the mandated use of uniform hiring tests (e.g., vs. wholesale assessments by managers), can help to standardize hiring decisions, some managers balk at the constraints these tests place on their latitude in hiring decisions and only apply them sporadically (Dobbin & Kalev, 2016). In brief, structural inequality reduction mechanisms—like blinding policies—that managers and evaluators personally endorse and prefer to use are the most effective in fostering workforce diversity and combatting the consequences of decision bias.

However, it is not clear exactly how prevalent blinding policies are, in the modern workplace, in the first place. To try to better understand the prevalence of blinding policies in U.S. organizations, I surveyed 828 Human Resources professionals (62.3% Female, M_{age} = 45.33, SD_{age} = 11.52) in partnership with the Society for Human Resource Management (SHRM). These professionals had an average of 13.68 years of work experience in the field of Human Resources. 95.7% of the sample had made at least one hiring decision over their career in HR, 84.5% had made at least 10, and 39.4% had made at least 100—overall, 58.3% answered "yes" to the question, "Is hiring one of the main things you do in your present job?" As such, this represented a sample of

participants—working in the HR field—who could reasonably be expected to have intimate knowledge of "best practices" (e.g., blinding policies) in hiring.

For each participant, I first described blinding policies in hiring: "Sometimes, organizations use what are called 'blinding' policies in the hiring process, where the people making hiring decisions have certain information about job applicants (such as their demographic information) withheld from them until after they make a hiring decision. These policies are intended to reduce the risk of the person making the hiring decision being biased or influenced adversely by irrelevant information." Next, I asked participants the following questions: "Does your organization employ any 'blinding'-style policies in the hiring process, where certain information about job applicants is withheld from the people making the hiring decision until after the decision is made?", "Have you ever worked for an organization that employed 'blinding' policies like the one described above, in a hiring context or otherwise?", "Have you ever received training related to 'blinding' as described above, as a policy or best practice, in any context (hiring or otherwise)?", and "Have you ever heard of 'blinding' policies as described above?" I found that 81% of the sample did *not* work at organizations that used blinding policies in hiring and had never worked for such an organization. Similarly, 80% reported never having received training related to blinding in any evaluative domain. However, 59% of participants did report that they had heard of blinding policies.

These data do not offer a representative survey of U.S. organizations today. With that said, they do suggest that blinding policies may not yet be very common, even though many people performing important evaluations (e.g., hiring decisions) in the workplace may be familiar with blinding as an idea. Given the evidence that blinding strategies can dramatically reduce bias in decision-making (Goldin & Rouse, 2000; Blank, 1991; Sah et al., 2016; Axt & Lai, 2019), these data thus underscore the importance of understanding *individual-level* preferences for blinding in judgments and evaluations. If blinding can help reduce inequality in the workplace—especially if managers and evaluators buy into it themselves (Dobbin & Kalev, 2016; Dobbin et al., 2015)—it is important to understand whether and when individuals will elect to blind their own judgment. I turn to this consideration in the next section, where I propose a theoretical framework of an individual-level blinding decision.

1.2 Theory

In this section, I present a framework that outlines a set of factors that may influence an evaluator's choice to see, or be blind to, a piece or set of information in an impending evaluation (see Figure 1). I do not argue that this framework represents a complete model of an individual-level blinding decision, capturing every antecedent of an evaluator's choice to see (vs. be blind to) potentially contaminative information. Rather, what I present in this section represents a partial framework that may offer a useful starting point to understand individual-level blinding decisions.

In this framework, I assume an evaluator is performing an evaluation of a new sample of performance (e.g., a performance evaluation; a hiring decision) and has the option to see or learn additional information about the target of their evaluation that is not the focal sample of performance to be evaluated (e.g., gossip about an employee to be evaluated; a job candidate's photo). I theorize that at least three different factors could guide an evaluator's choice to see, or be blind to, potentially contaminative information in their evaluation: their *enduring beliefs about information and bias,* their *situational beliefs about information and bias,* and their *motivations.*

First, I propose that at least two *enduring beliefs*—that is, beliefs that carry over across different evaluation situations—will drive evaluators' choices to see (vs. be blind to) potentially contaminative information in an evaluation: (i) a belief that, when evaluating others, it is always better to have more information, and (ii) a belief that, when evaluating others, they are not personally susceptible to bias in their evaluations. Second, I propose that at least one *situational belief*—that is, a belief that is specific to a given evaluation and the information available—will drive evaluators' choices to see (vs. be blind to) potentially contaminative information: a belief that seeing or learning the potentially contaminative information would be helpful for their evaluation (e.g., would improve accuracy). Finally, I propose that at least one *motivation* will drive evaluators' choices to see (vs. be blind to) potentially contaminative information: curiosity to see the information. I elaborate on each of these factors below.

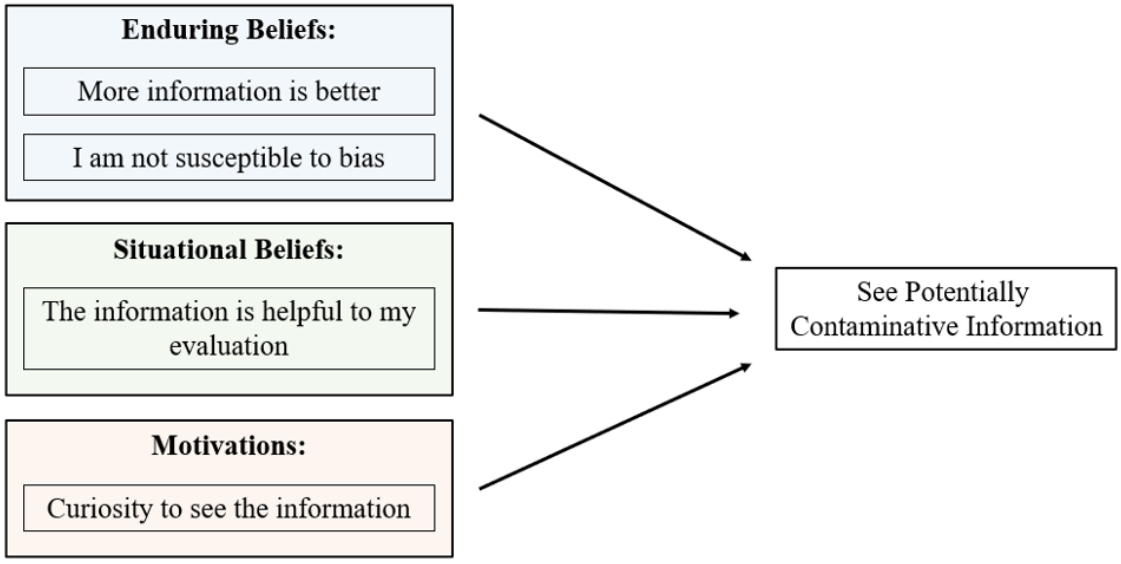

Figure 1: Theoretical Framework

1.2.1 Enduring Beliefs About Information and Bias

First, I propose that evaluators may hold enduring—reliable and consistent across situations—lay beliefs about both information and susceptibility to bias in evaluations. I propose that these enduring beliefs should impact evaluators' preferences for seeing, or being blind to, potentially contaminative information across different evaluative contexts. In what follows, I discuss two such beliefs that I theorize should predict choices to see potentially contaminative information in an evaluation (rather than be blind): a general belief that more information is always better to have and a general belief that one is not susceptible to bias. In both cases, I expect that more (vs. less) endorsement of these beliefs should correspond to more (vs. less) inclination to see potentially contaminative information in an evaluation.

1.2.1.1 A belief that more information is better

I propose that one enduring belief people may hold, shaping their blinding preferences, is that more information is always better to have. One can find instances of this belief in philosophical and scientific discourse as far back as Aristotle, who argued that one should collect as much information as possible when performing analyses, making decisions, etc. (Kerr & Pritchard, 2012). This belief still has adherents today, from celebrity scientist Neil DeGrasse Tyson, who argues that, "to a person engaged in discovery, all information is good, even when it's bad" (Tyson, 2008), to self-help "guru" Simon Sinek, who states, as a general principle in life, that "more information is always better than less" (Sinek, 2009). I argue that, the more evaluators endorse a belief that more information is always better to have when performing evaluations, the more inclined they should be to seek potentially contaminative information, rather than choose to blind themselves to such information.

Notably, some research demonstrates that, under certain conditions, people do tend to prefer having more information to less (Grant, Kajii, & Polak, 1998; Lanzetta & Driscoll, 1966; Lockard, 1963). While people usually seek information (e.g., their odds of winning a lottery) because they view it to be instrumental to a decision (e.g., whether to buy a lottery ticket), they can also prefer information at an intrinsic level, even when it is not instrumental (i.e., no expected utility), merely because it resolves momentary uncertainty (Grant et al., 1998). For instance, when receiving an electric shock is

unavoidable but the time-to-delivery of the shock is variable, people prefer to have a warning signal before the shock is delivered in order to reduce uncertainty (Badia, McBane, & Suter, 1966; Lanzetta & Driscoll, 1966). Similarly, in economic games, research shows that as the degree of uncertainty about the outcome of a choice (i.e., the time-lag between making a gamble and learning the outcome) increases, people become more willing to pay a fee to learn more about the probability of their choice being correct *after* their choice is already made (Eliaz & Schotter, 2007). Even when deciding how much instrumental information to factor into an evaluation, people tend over-estimate how much they will need (Klein & O'Brien, 2018), consistent with a "more information is better" mindset. Overall, this research suggests that people may be especially likely to choose to see or learn non-instrumental information to the extent that doing so relieves a subjective sense of uncertainty, and that decision-makers generally, though perhaps weakly, prefer more information to less (Grant et al., 1998).

On the other hand, certain situations may cue a desire to avoid more information rather than seek it out. People often choose *not* to learn their risk for certain cancers, for instance (Keogh et al., 2004), or whether a completed HIV test returned a positive or negative result (Hightow et al., 2003). Generally, this type of "information avoidance" occurs when people seek to avoid information they perceive as likely to be psychologically threatening (Sweeny, Melnyk, Miller, & Shepperd, 2010), such as information that would threaten an already established, intuitive preference (Woolley &

Risen, 2018). Hence, an evaluator may be more likely to consider excluding information from an evaluation if they believe that avoiding seeing or learning that information would resolve a potential threat to their emotions or established preferences. If information avoidance-related concerns do not apply to a blinding decision, however (e.g., choosing to see or be blind to a job candidate's photograph), an intrinsic preference for having more information over less may be more likely to guide that decision.

In sum, I propose that a person engaging in an evaluation will be less likely to choose to blind their evaluation to potentially contaminative information the more they endorse a belief that more information is always better to have. Some research suggests that people prefer having more information to less in a general sense, and especially to the extent that receiving more information resolves uncertainty. However, other research demonstrates that people do not always desire to have more information, and sometimes avoid receiving more. Therefore, it may be reasonable to conclude that there will be variance between evaluators in the extent to which they believe, when evaluating others, that the more information they can have about the target of their evaluation, the more effective the evaluation will be. I predict that, across different evaluative scenarios and contexts, variance in this belief will inform evaluators' preferences for seeing, vs. being blind to, potentially contaminative information in evaluations.

1.2.1.2 A belief that one is not susceptible to bias

Another enduring belief that evaluators may hold across evaluation situations concerns their self-assessed susceptibility to decision bias. Evaluators who do not believe they are susceptible to decision bias in a general sense may more inclined, across evaluations, to see (vs. be blind to) potentially contaminative information. Those who are less sanguine about their susceptibility to bias may be more inclined to choose blinding.

Some factors might increase or decrease an evaluator's likelihood of believing that they are susceptible to evaluation contamination. Most relevant is research demonstrating that people tend to view themselves as objective in both specific evaluations and their general perceptions of the world (Uhlmann & Cohen, 2007; Armor, 1999). Ironically, research on subjective objectivity demonstrates that increases in belief that one is objective in one's decision-making can lead to increased employment of category based heuristics in evaluations of others (Uhlmann & Cohen, 2007). Relatedly, research on the "bias blindspot" suggests that people (a) specifically believe they are less susceptible to judgment bias than others (Pronin, Lin, & Ross, 2002), and (b) generally believe that they are not personally vulnerable to judgment bias (Pronin, Gilovich, & Ross, 2004). To the extent evaluators view themselves as objective in their reasoning and/or less susceptible to bias than others, they may be relatively less likely to believe

that potentially contaminative information will, in fact, contaminate an impending evaluation.

Even if an evaluator does not believe that they are generally objective (i.e., universally invulnerable to decision bias), they may still believe they are not susceptible to decision bias because of an ability to correct for/remove any such bias from their evaluations. If an evaluator believes that a given piece of information could contaminate their evaluation, but also believes they can correct for that possible contamination to their evaluation, they should be less likely to choose to blind themselves to the potential contaminant than if they do not believe they can successfully correct for it. Some research suggests that people do believe that they can correct for contamination in their evaluations.

Research on flexible correction processes suggests that people employ strategies and/or corrections to their reasoning when they are concerned with susceptibility to bias (i.e., contamination) in that reasoning (Wegener & Petty, 1995; Petty & Wegener, 1993; Ottati & Isbell, 1996). Specifically, given that people are (a) aware of a possibility for bias in an evaluation, (b) motivated to correct for that bias, and (c) have sufficient cognitive resources to deploy a correction process, they will attempt to adjust their evaluation (e.g., moderate its positivity or negativity) in the opposite direction of the considered bias (Wegener & Petty, 1995). Similarly, research on naïve theories of bias correction finds that people often choose to see information in an evaluation (e.g., the gender of a

hypothetical job candidate) even when they acknowledge the potentially biasing nature of that information, ostensibly because they believe they can personally avoid being biased (Wilson & Brekke, 1994; see also Uhlmann & Cohen, 2007). Other research in this area demonstrates that people believe they are more capable at correcting for common decision biases (e.g., the fundamental attribution error) in their own reasoning than others are at correcting for bias in theirs (Van Boven, White, Kamada, & Gilovich, 2003).

This stream of research implies that people may be overconfident in their ability to correct for bias in their evaluations. People are often overconfident about their own abilities and the quality of a given performance or estimate (Soll & Klayman, 2004; Moore & Healy, 2008), and also systematically believe they are "better" (e.g., more talented or accurate) than the average person (Dunning, Meyerowitz, & Holzberg, 1989), especially regarding easy tasks (Larrick, Bunson, & Soll, 2007). It follows that a given evaluator would be (over-)confident in their ability to correct for any bias in their evaluation.

Whether due to a belief that they are objective reasoners or a belief that they can correct for bias in their reasoning, I propose that evaluators' enduring beliefs about their susceptibility to evaluation contamination will guide their blinding choices. Across different evaluations, the more an evaluator believes they are not susceptible to bias in general sense, the more inclined they should be to incorporate potentially biasing information into those evaluations. Conversely, those who believe they are susceptible

to bias (e.g., because they do not believe they are objective or are not confident in their personal bias correction abilities) may be more inclined to consider blinding their evaluations to potentially contaminative information.

1.2.2 Situational Beliefs About Information and Bias

While certain enduring beliefs, across evaluative contexts, may shape evaluators' blinding preferences, I also propose that, *within specific* evaluative contexts—when performing specific types of evaluations or making blinding choices about specific types of information—evaluators' blinding preferences will be guided by situational beliefs. Chiefly, I theorize that, as there is heterogeneity in types of evaluations (e.g., performance evaluations vs. hiring decisions) and in potentially contaminative information (e.g., background performance information vs. a job candidate's demographics), it is likely that there is also heterogeneity in people's beliefs about the usefulness/helpfulness of a given piece of potentially contaminative information for a given evaluation.

Whereas in one specific evaluation situation, an evaluator might deliberate on available, potentially contaminative information, decide it is likely to be helpful to their evaluation (e.g., believing it will improve accuracy), and so choose to see it, in a different evaluation situation, they might judge potentially contaminative information as likely to harm their evaluation (e.g., believing it will lead to bias), and so choose to be blind to it. Put differently, when performing the same evaluation, I theorize that one evaluator

might honestly judge potentially contaminative information to be helpful to their evaluation, and so choose to see it, while another evaluator might judge the same information harmful or contaminative, and so choose to be blind to it. In this way, I propose that, over and above enduring beliefs, people's beliefs about the usefulness or helpfulness of a piece of contaminative information—specific to a given evaluative context—should also inform their blinding preferences, particularly when those beliefs are made salient before a blinding choice is made.

Some research supports the notion that evaluators may deliberate, when performing evaluations, on the extent to which potentially contaminative information is likely to be helpful vs. harmful to their evaluation—and make blinding choices consistent with those situational beliefs. For instance, people are roughly as concerned about judgment contamination as they are about many "serious" issues, such as finding a spouse or getting a job (Wilson & Brekke, 1994), suggesting that people may use situational beliefs about whether information is likely to contaminate (or help) their evaluation to guide blinding preferences. Notably, people are not always accurate in identifying what information is contaminative or biasing, and what information is helpful, in the first place. People dramatically overestimate the expected biasing effect of subliminal messaging, for instance, and misunderstand the value of base-rate information about a product relative to others' subjective experiences with that product (Wilson & Brekke, 1994). Indeed, in some evaluation scenarios, people intentionally

employ, rather than avoid or discount, biasing information, because of situational, naïve theories about its degree of contaminativity vs. usefulness (Dietvorst & Simonsohn, 2018). It is possible that some information is more easily (correctly) identified as potentially contaminative versus other types of information. For instance, one might assume that a teacher is more likely to identify a student's racial category as a possible contaminant in her grading of that student's work than a loud, distracting noise occurring during the grading, or her own fluctuations in mood.

Overall, this research suggests that evaluators are (a) likely to attend to the extent to which available, potentially contaminative information will help or harm their evaluation, (b) not always accurate in these judgments, leading to heterogeneity in beliefs about the helpfulness/harmfulness of different types of information, and (c) likely to use these beliefs to inform their information preferences. In brief, different evaluative situations present different types of potentially contaminative information, and I expect that people will exhibit genuine heterogeneity in beliefs about whether that information is likely to be helpful for their evaluation (e.g., increase accuracy) or harmful (e.g., biasing). While evaluators' enduring beliefs about the value of more information or their susceptibility to bias may factor into their blinding preferences, to the extent these situational beliefs about a given piece of information's degree of helpfulness are made salient before a blinding choice is made, they, too, should guide blinding preferences.

1.2.3 Curiosity Motivation

I propose that one final factor that may shape evaluators' blinding preferences is a motivation to see potentially contaminative information in order to resolve curiosity. Research shows that people are curious by nature and will want to see readily available but hidden information even when there is no advantage to doing so (Berlyne, 1950; 1966; Litman & Jimerson, 2004; Loewenstein, 1994). Indeed, curiosity is a potent driver of decisions (Litman, 2005). In one study, participants were told that a subset of pens in front of them were wired to give an electric shock if clicked. Participants clicked more pens when they were unmarked compared to when they were clearly marked as wired or not (Hsee & Ruan, 2016). That is, participants' curiosity about available, but hidden, information—whether a prank pen would deliver a shock or not—was strong enough to lead them to accept a risk of physical pain in order to satisfy it. I theorize that people are also likely to be curious about targets of evaluation and may seek additional information about those targets in order to satisfy their curiosity. Given the strength of curiosity as a driver of behavior, I propose that piqued curiosity may lead evaluators to elect to see or learn, rather than be blind to, additional information about a target of evaluation *even when accounting for* concerns that it might lead to evaluation contamination.

Some factors might increase or decrease the effect of curiosity motivation on blinding decisions. For instance, it bears noting that curiosity is a "hot" affective state (Loewenstein, 1994)—i.e., experienced as relatively more emotional than rational in

nature (Bazerman, Tenbrunsel, & Wade-Benzoni, 1998). Consequentially, blinding decisions that are framed in such a way as to induce cool-headed reasoning rather than hot-headed impulses might be less likely to be affected by evaluators' curiosity. If an evaluator is making a blinding decision concerning an evaluation set into the distant future (e.g., in one month) versus one set the following day, immediate, affective concerns such as a "hot" state of curiosity should be relatively less likely to drive their blinding choice (Trope & Liberman, 2010; Rogers & Bazerman, 2008).

Similarly, if an evaluator is considering what their blinding preference should be, instead of simply making an on-the-spot blinding choice, they should be more likely to approach the decision in a cool-headed, deliberative mindset and hence less likely to be motivated by their personal curiosity (Bazerman et al., 1998). Blinding decision frames that induce careful, analytic reasoning (i.e., "System 2" Stanovich & West, 2000; Evans & Stanovich, 2013) should be more likely to be driven by motivations such as making a fair or accurate evaluation than frames that induce faster, more intuitive processing strategies (i.e., "System 1"). Indeed, people are more motivated to be accurate and objective in the judgments and decisions they make when using a systematic (e.g., System 2) versus heuristic (e.g., System 1) mode of processing (Chaiken, 1980; Chen, Shechter, & Chaiken, 1996; Chaiken & Stangor, 1987). Overall, I theorize that a "hot" motivation to resolve curiosity will drive choices to see (vs. be blind to) potentially contaminative information. Accordingly, blinding decisions that are framed in such a

way as to generate on-the-spot decisions should be more likely to be motivated by curiosity than blinding decisions that require more deliberate, "cool" reasoning, which may be more motivated by concerns over being fair or accurate in one's evaluation.

1.2.4 Summary

In summary, the theoretical framework I have outlined proposes two "cold"—i.e., thoughtful and deliberative—drivers of choices to see potentially contaminative information: (i) enduring beliefs across evaluative contexts that more information is always better or that one is not susceptible to bias and (ii) situational beliefs within different evaluative contexts that a given piece of information is helpful for one's evaluation. Additionally, the framework proposes one "hot"—i.e., intuitive and impulsive—driver of choices to see potentially contaminative information: a motivation to resolve curiosity. I do not argue that these different factors operate independently of each other, and it is possible that some may influence others. For instance, an evaluator might be so curious to see potentially contaminative information about a job candidate that they convince themselves it is likely to be helpful to them—a motivation influencing a situational belief. Conversely, a unique evaluation situation might offer potentially contaminative information that is so blatantly likely to cause evaluation bias that an evaluator rethinks the validity of a belief that more information is better or that they are invulnerable to bias—a situational belief influencing enduring beliefs. While this potential interplay between factors is interesting, I largely test components of my

theoretical framework—enduring beliefs, situational beliefs, and motivations—individually, in terms of their relationship with blinding preferences, in the studies that follow. I made this choice because of the complexity associated with testing all the factors proposed by my framework simultaneously in one experiment; it would be interesting for future research to endeavor to test these potential interactions. Given the number of factors I propose may lead evaluators to choose to see potentially contaminative information, it is clear that I do not expect everyone to choose blinding as a strategy on their own. Consistent with this appraisal, in the studies I present over the following chapters, a substantial percentage of participants—often approaching 50%—will choose to see, rather than be blind to, potentially contaminative information.

1.3 Overview

Across each of the studies I present in the following chapters, participants took part in an evaluation of some kind and made a choice whether to see, or be blind to, additional and potentially contaminative information about the target of their evaluation. The studies presented in Chapter 2 take an initial look at participants' blinding preferences (i.e., whether to see potentially contaminative information or be blind to it) and the consequences of choosing *not* to be blind to potentially contaminative information for evaluation bias across three different contexts. The studies presented in Chapter 3 test the motivational component of my theoretical framework—exploring whether a "hot" state of curiosity drives choices to see potentially contaminative

information. The study presented in Chapter 4 tests both of the enduring beliefs proposed in the framework—that more information is always better to have and that one is not susceptible to bias—as they relate to blinding preferences. Finally, the studies presented in Chapter 5 explore the relationship between situational beliefs that potentially contaminative information is helpful vs. harmful and blinding preferences, and explore whether interventions that cue such situational beliefs before a blinding choice is made can encourage evaluators to choose to blind their evaluations to potentially contaminative information.

Across these studies, I largely elected to draw convenience samples of lay people from online platforms. I made this choice because the present research represents the first tests of individual-level blinding preferences; hence, I prioritized the ability to have high-powered tests of my predictions (i.e., at least 100 participants per experimental condition; Gervais et al., 2015), using samples more representative of the US population than local convenience samples (e.g., in-person or laboratory based; Berinsky, Huber, & Lenz, 2012). However, there were some exceptions to this rule. For instance, in Study 2, participants took part in my study in the behavioral lab at Duke University. In Study 8, I sampled participants with managerial experience via the same online platform used in the other studies with online samples (i.e., Prolific Academic, an online platform shown to provide data similar or better in quality to data derived from other online sources; Peer et al., 2017). In these cases where I selectively sampled participants, I either did so

for the purposes of the study (as in Study 2) or in order to test the replicability of my effects with a sample of participants who were likely to have made consequential blinding decisions in their careers (as in Study 8). Some studies were pre-registered using AsPredicted.org; I include links to pre-registration documents where applicable. Select study materials are archived at https://osf.io/pzgq2/?view_only=0e05f06ce0fd43dfa1f064a0911fcf21.

2. An Initial Look at Blinding Preferences and Consequences for Evaluation Bias

In the preceding chapter, I argued why individual-level blinding preferences are important to examine: Though research has demonstrated that blinding policies can increase the fairness and accuracy of evaluations (Goldin & Rouse, 2000; Axt & Lai, 2019), such policies do not yet appear to be widespread in organizations, and even when implemented, are most likely to be effective when they are endorsed by those whose evaluations they govern (Dobbin & Kalev, 2016; Brehm, 1966). For these reasons, it is important to understand when evaluators in organizations might choose, on their own, to be blind to potentially contaminative information in evaluations and how a preference for blinding might be encouraged. The first step toward understanding how to encourage evaluators to blind their own judgment is to examine what evaluators' preferences for blinding are at baseline. Thus, in the studies I discuss in this chapter, my goal was to probe evaluators' baseline preferences for seeing, or being blind to, potentially contaminative information about the target of an evaluation across a variety of evaluative contexts. I also sought to demonstrate different ways in which a choice to see (vs. be blind to) potentially contaminative information can lead to bias in evaluations of others, in order to reinforce the importance of understanding individual-level blinding preferences.

The studies in this chapter place participants in three different evaluative contexts: estimating someone's performance in a video-taped task (Study 1), grading the

creativity of someone's response to a brainstorming task (Study 2), and making a mock hiring decision (Study 3). In Study 1, participants chose whether to see or be blind to a profile of background information about a performer. In Studies 2 and 3, participants chose whether to see or be blind to information about the target of their evaluation such as the target's name and photograph. Across these studies, I made no predictions regarding the proportions of participants that would choose to see or be blind to potentially contaminative information about the targets of their evaluation; these studies were exploratory in that regard. However, I did predict, in each study, that choosing to see potentially contaminative information (vs. be blind to it) would make participants' evaluations susceptible to bias (e.g., due to an anchoring effect in Study 1, Wilson et al., 1996; or due to an ageist stereotyping process in Study 3, North & Fiske, 2012).

2.1 Study 1

In Study 1, I sought to explore participants' blinding preferences in a performance evaluation context. Participants estimated someone's performance in a task and had the option to view, or be blind to, potentially contaminative information about the performer before making their evaluation. My main prediction in this study was that participants who chose to view the potentially contaminative information would provide biased estimates. The pre-registration document for this study is available at https://aspredicted.org/wm3qg.pdf.

2.1.1 Method

2.1.1.1 Participants

I recruited 804 participants from Prolific Academic (M_{age} = 32.19, SD_{age} = 11.46, 44.5% female).

2.1.1.2 Materials and procedure

At the outset of the study, participants learned that they would be shown a video of someone's performance in an online, pattern recognition task and that, afterward, they would need to estimate how many times the performer was correct during the task. All participants were provided with basic information about the rules of the pattern recognition task and learned that, for every correct decision in the task, a green check would appear on the screen, whereas for every incorrect decision in the task, an orange X would appear. Before evaluating the target person's performance, some participants were shown background information on the target person, depending on condition. Participants were randomly assigned to one of three conditions: *choice to see profile information (choice), automatically see profile information (automatic-profile),* and *automatically be blind to profile information (automatic-blind).* The two control conditions provided a baseline against which to assess the difference in judgment bias shown by participants who chose to see the profile versus remain blind.

Participants in the *choice* condition were told: "At work, we often have background information about the person we are asked to evaluate. Before you watch

the video of their performance, you will have the choice to view a profile of the person whose performance you are evaluating, which contains the following information: Their name; Their hometown; Where they went to college and their major; Their performance on a different task, which tested math ability." After participants read what would be provided in the profile, they were told: "On the next screen, you will be able to indicate if you'd like to view this person's profile before watching the video of their performance, or not. You do not have to view this profile - it is purely optional." Then, they were asked, "Would you like to view this person's profile information before you watch the video of their performance?" and responded with options "yes" or "no." Participants in the *choice* condition who elected to view the profile were presented with the following profile before viewing the video:

PROFILE

Name: Sean F
Hometown: Chicago, IL
Education: Community college graduate
College Major: English
Math Task Performance: 40% Correct

After viewing this profile, participants answered three attention check items. These items were: "What was his hometown?" (Chicago, IL / Gary, IN / Ithaca, NY / Albuquerque, NM), "What was his college major?" (English / History / Advanced Mathematics / Physics), and "What was his performance on the other, math-related

task?" (50% Correct / 40% Correct / 80% Correct / 90% Correct). Responses to these items ranged from 86-97% correct, indicating that most participants paid attention to the profile when they viewed it. Following the attention check items, participants viewed the video of the performance. Participants in the *choice* condition who chose to be blind to the profile information did not view the above profile or complete attention check items—they viewed the video of the performance alone.

Participants in the *automatic-profile* and *automatic-blind* conditions were not presented with the option to view the profile. Those in the *automatic-profile* condition read: "At work, we often have background information about the person we are asked to evaluate. Before you watch the video of their performance, you will view a profile of the person whose performance you are evaluating, which contains the following information: Their name; Their hometown; Where they went to college and their major; Their performance on a different task, which tested math ability." After viewing the profile, participants completed the same attention check items described above. Responses to the attention check items ranged from 80-99% correct, indicating that, as in the *choice* condition, most participants paid attention to the profile. Following these items, participants viewed the video of the performance. Participants in the *automatic-blind* condition viewed the video alone.

All participants viewed the same performance video, which can be viewed at this link: https://www.youtube.com/watch?v=Gcw2pa7Em0E&feature=youtu.be. I generated

this video by using Camtasia Software (https://www.techsmith.com/video-editor.html) to record my computer screen while performing a pattern recognition task on the website Lumosity.com. The pattern recognition task involved being shown a series of cards and indicating whether the one currently shown matched the one shown previously, using key strokes to indicate "yes" or "no." Each incorrect choice was followed by an orange X, and each correct choice was followed by a green check. The video was 49 seconds long and showed only the performance on the task—no results or summary statistics were shown at the end. Though participants were not made aware of this, the true performance level in the video was 69% correct. The video was sped up to 200% of its original speed in order to make performance in the task more ambiguous.

After viewing the video, participants were asked to estimate the performer's degree of success. Participants read: "Now we would like you to evaluate his performance. Specifically, we would like you to estimate the percentage of times that he made a correct guess during the task, from 0% to 100%. Remember, if you estimate that he was 100% correct, that would mean that he correctly identified each card and received only green checks throughout the entire task. If you estimate that he was 10% correct, that would mean that he correctly identified 1 out of every 10 cards, and that 90% of the marks he received during the task were orange Xs. Please select the option below that is closest to your estimate of the percentage of times he was correct during the task." Participants estimated the quality of the performance using a 21-point scale,

with radio buttons labeled 0%, 5%, 10%, …, 100%. After participants made their estimate, their answer was presented to them. They were given the option of confirming or refining it.

2.1.2 Results and Discussion

Forty-eight percent of participants in the *choice* condition chose to view the profile. Because the profile of the performer contained information about a poor performance on a math task, I expected participants who saw the profile to make performance estimates that were negatively biased relative to the true performance of 69%, whereas I did not expect a bias from those who did not see the profile. To assess this, I subtracted 69 from participants' percentage estimates and analyzed these signed difference scores (the average of which represents bias).

As shown in Table 1, average bias was greater (more negative) among participants who chose to view the profile or who automatically viewed the profile than among those who chose to be blind or who were automatically blinded. In support of this, a 2 (Blinding Decision: ability to choose vs. automatic) x 2 (Blinding Outcome: see profile vs. blind) ANOVA test of bias revealed a main effect of Blinding Outcome, $F(1, 800) = 40.84$, $p < .001$, $\eta_p^2 = .05$, but not a main effect of Blinding Decision, $F(1, 800) = .001$, $p = .978$, $\eta_p^2 = .00$ or an interaction, $F(1, 800) = .32$, $p = .572$, $\eta_p^2 = .00$. Analyses using an absolute error measure (i.e., an unsigned discrepancy) revealed similar findings. These results show that participants gave lower estimates of performance quality when they

chose to see the profile or automatically viewed it than when they chose to be blind or were automatically blinded. The *automatic-profile* and *automatic-blind* control conditions confirm that choosing to see or not see information yields a similar difference in bias as did random assignment to the same states.

Table 1: Means and 95% Confidence Intervals of Variables by Condition

	Chose Profile (n = 195)			Automatic-Profile (n = 195)			Chose Blind (n = 210)			Automatic-Blind (n = 204)		
	M	*CIL*	*CIU*	*M*	*CIL*	*CIU*	*M*	*CIL*	*CIU*	*M*	*CIL*	*CIU*
Est.	64.42	62.44	66.40	64.93	63.00	66.86	70.37	69.01	71.73	69.91	68.46	71.37
Bias	-4.58	-6.56	-2.60	-4.07	-6.00	-2.14	1.37	0.01	2.73	0.91	-0.54	2.37

Note. Est. = estimate. *M* = mean. *CIL* = lower bound of 95% confidence interval around mean. *CIU* = upper bound of 95% confidence interval around mean. *n* = cell size. Bias is calculated as participant's estimate minus true value of 69.

2.2 Study 2

In Study 1, nearly half of participants chose to view potentially contaminative information about the target of their evaluation when given the option. It is possible that the percentage of participants choosing to see the profile was as high as it was because participants did not view their blinding choice to be consequential for the target of their evaluation—there was no discernable link between their estimate of the performance and outcomes for the performer. Conversely, it is possible that participants chose to view the profile of the performer because they thought they *should* see it—perhaps believing it would aid estimate accuracy. My goal in Study 2 was to address these two possibilities.

Participants in a laboratory setting graded the real response of another person to a brainstorming-related question and were told that their grade would determine the respondent's eligibility for a bonus. In other words, in Study 2, participants were led to believe that their evaluations were consequential for the target of their evaluation. Situated in this context, participants made a blinding choice in advance of their evaluation to see or be blind to a profile of the person they were to grade, containing what was ostensibly the respondent's name and picture. In one condition, participants made a straightforward blinding choice—to see the profile or not. In another condition, participants first indicated what their blinding choice should be, and then went on to make a real blinding choice. As such, I was able to determine whether participants' blinding preferences in the moment aligned with their beliefs about what they should do (i.e., see profile or be blind). Finally, participants graded the respondent's creativity in their response to the brainstorming task. Because the profile depicted an obese respondent, I expected participants who chose to see the profile to provide negatively biased grades relative to those who chose to be blind (Levine & Schweitzer, 2015). The pre-registration document for this study is available at http://aspredicted.org/blind.php?x=n4ep28.

2.2.1 Method

2.2.1.1 Participants

I recruited 273 student participants from the behavioral lab at Duke Univeristy (M_{age} = 27.65, SD_{age} = 9.67, 67.5% female).

2.2.1.2 Materials and procedure

In a pretest, I recruited 299 participants from Mechanical Turk (M_{age} = 34.56, SD_{age} = 11.03, 40.5% female). Participants completed the "Alternative Uses" brainstorming task, which measures faculty with divergent thinking (Gilhooly et al., 2007). Participants read the following: "In this task, we would like you to write down as many different uses for a brick as you can think of. One common use for a brick is to build things. What other things could a brick be used for?" Participants were given three minutes to write down as many uses for a brick as they could think of. Responses provided in this pretest were employed in the main study.

In the main study, participants were led to believe that they would be reading and grading a response to the Alternative Uses task written by a local college student. In reality, participants were randomly provided with one of the responses generated in the pretest. Participants read: "In a separate study, we collected responses from local college-age students to the Alternative Uses for a Brick Task. You will be randomly assigned to one of these responses, which we will ask you to grade. Please take this grading seriously -- we will be giving out cash prizes to the people who had the best

responses to this Alternative Uses for a Brick Task, and we need your help in determining who came up with the best responses." In this way, participants learned that (a) they'd be grading someone else's response to a task, and (b) their grade would help to determine the respondent's eligibility for a bonus. Participants learned the rules and structure of the Alternative Uses task and read: "Each response gets a grade from 0-100. Responses with a grade of 100 are the most creative and inventive responses, whereas responses with a grade of 0 are the worst possible responses, not creative or inventive at all. Some things to consider are: How original are the uses for a brick that they came up with? How creative are the uses they came up with?"

Following a false loading screen, used to imply that participants were being "paired" with a response, they read: "You have been paired with a response. Before you see this response, you can see a small profile of the person who wrote it, if you would like to. The profile contains the following information: Their name; Their Facebook profile picture. You do not have to see this profile -- it is completely optional." Participants were randomly assigned to one of two conditions. In the *plain choice* condition, participants made a blinding choice in response to the following question: "Do you want to view the profile of the person you are grading? If you select yes, you will see it on the next screen. If you select no, you will just see their response" (yes/no). In the *should first* condition, participants first responded to the following question (*should choice*): "Do you think you should see the profile of the person you are grading?"

(yes/no) and then went on to make their real blinding choice in the same manner as in the *plain choice* condition (*post-should choice*). Participants across conditions who chose to view the profile read that the respondent's name was "Josh" and saw a picture depicting an obese White man taken from Google Images (see Figure 2).

Figure 2: Depiction of Respondent

After viewing the profile, participants were randomly provided with one of the responses to the Alternative Uses task generated in the pretest. Those who chose to be blind to the profile went directly to the response. Finally, participants graded the response they viewed on a 100-point scale, following the prompt: "Please use the slider below to provide your grade for the response. Remember, 0 indicates that the person whose response you saw gave the worst, least creative response possible, whereas 100 indicates that the person gave the best, most creative response possible."

2.2.2 Results and Discussion

The main point of interest in this study was the percentage of participants choosing to see the profile of the respondent across conditions. In the *plain choice* condition, 46% of participants chose to view the profile. This differed from the percentage who thought they should see the profile in the *should first* condition (*should choice*; 18.4%), $\chi(1)^2 = 23.81$, $p < .001$. Participants who made their blinding choice after first indicating what they thought their choice should be (*post-should choice*) were less likely to choose to see the profile (39%) than those who made a plain blinding choice alone, though this difference was not significant $\chi(1)^2 = 1.37$, $p = .241$.

Thus, the results of Study 2 suggested that, in Study 1, participants were not merely electing to see the profile of the performer because they thought they should see it—in this study, only 18.4% of participants indicated they thought they should see the profile of the respondent when asked, whereas 46% chose to see it when making a plain blinding choice (i.e., no "should" framing). That is, these results suggest that situational beliefs that one should see the profile (e.g., because it would aid evaluation accuracy) may have explained the preferences of a portion of participants who made a plain blinding choice, but another factor (e.g., curiosity motivation) likely also drove choices.

With that said, having participants first indicate what their blinding choice should be and then make a real blinding choice afterward did not significantly depress their inclination to see the profile, relative to those making a blinding choice alone

(though the trend of results was consistent with my predictions). I return to this idea in Study 8. As predicted, participants who chose to view the profile of the respondent gave lower grades on the Alternative Uses task (M = 46.36, SD = 27.26) than did participants who chose not to view the profile (M = 54.92, SD = 25.51), $t(271) = 2.66$, $p = .008$, $d = .34$. This result is consistent with research demonstrating that obese individuals are stereotyped as low in competence (Levine & Schweitzer, 2015).

2.3 Study 3

In Studies 1 and 2, participants who chose to view potentially contaminative information about the target of their evaluation made estimates and judgments that were different from those who chose to be blind to that information. In Study 1, seeing the performer's profile reduced participants' estimates of performance, consistent with an anchoring effect (Wilson et al., 1996). In Study 2, seeing the profile of the respondent led participants to give the respondent lower grades, consistent with stereotypes concerning obese people (Levine & Schweitzer, 2015). I argue that these results represent mental contamination, whereby participants' estimates and judgments were influenced adversely (i.e., biased) by external factors, presumably in a fashion contrary to their intentions (Wilson & Brekke, 1994). However, some research suggests that, in some contexts, people *intentionally use* "to be ignored" information, such as a profile of background information about a performer, because they believe it to be instrumental to their evaluations (e.g., consistent with a situational beliefs account, Dietvorst &

Simonsohn, 2018). Critically, this research suggests that people who do not want, or do not intend to use, "to be ignored" information can successfully correct for its effects in their evaluations (Dietvorst & Simonsohn, 2018). The takeaways from this research would that participants in Studies 1 and 2 who chose to see the potentially contaminative information were not being unintentionally biased in their evaluations—a mental contamination account—but rather actively and purposefully employed the information in the profiles to inform their evaluations. This concern might be mitigated, somewhat, by the fact that a strong majority of participants in Study 2 reported, when asked, that they should not see the profile of the respondent. That result is inconsistent with the notion that all participants were employing "to be ignored" information in that study because they thought that they should do so. Still I sought to test this possibility more directly in Study 3.

I employed a third evaluative paradigm in Study 3, a mock hiring context, and participants chose to see, or be blind to, a job candidate's name and photograph. The job candidate's photograph was randomly varied between participants—some viewed a younger job candidate, and some viewed an older candidate. I expected that participants who saw the photograph would evaluate the candidate in a manner consistent with an ageist bias (North & Fiske, 2012). Aside from the evaluative context, the major change in Study 3, relative to Studies 1 and 2, was that both participants who chose to see the candidate's name and photograph and those who chose not to see were ultimately

provided with the photograph. With this design, I was able to test whether "bias" in participants' estimates was constrained to those who chose to see the name and photograph only—consistent with the notion that people who do not want to use "to be ignored" information can successfully ignore it (Dietvorst & Simonsohn, 2018)—or extended even to those who chose to be blind to the information—consistent with a mental contamination account (Wilson & Brekke, 1994). I predicted that both participants who chose to see the name and photo and those who chose to be blind would demonstrate an ageist bias in evaluations of the job candidate.

2.3.1 Method

2.3.1.1 Participants

I recruited 203 participants from Prolific Academic (M_{age} = 34.41, SD_{age} = 12.13, 48.8% female).

2.3.1.2 Materials and procedure

All participants took part in a mock hiring task. For instance, they read: "Imagine you have been tasked with hiring someone for an open position at your company…You are currently at the stage where you are going through the people who have applied for the job and analyzing whether they fit the job's specific requirements or not." Next, participants were shown a mock job advertisement for a "Social Media Manager and Liaison." In addition to innocuous requirements (e.g., "3 years experience in Marketing field or related"), the job ad communicated that the ideal candidate must have social

media knowledge and expertise (e.g., "Candidate will maintain the company's Social Media presence")[2]. Finally, participants imagined that they were perusing the resumes of the people who had applied for the job and answered the following question: "Do you want the job candidate's name and photo displayed at the top of the resume? We will show you the resume on the following page based on your answer to this question" with options "yes, show name and photo at top of resume" and "no, don't show name and photo at top of resume."

If participants elected to view the name and photo, they were presented with a resume containing name and photo information at the top, preceded by the prompt: "Below is the job candidate's resume. Please review the resume carefully. After you review it, you will decide whether this candidate has a good fit for the job or not." If participants elected to be blind to the name and photo, they were still presented with a resume containing name and photo information at the top, preceded by the prompt: "Sorry, we need more participants to evaluate the candidate after seeing the resume. Below is the job candidate's resume. Please review the resume carefully. After you review it, you will decide whether this candidate has a good fit for the job or not."

---

[2] After viewing the advertisement, participants completed an attention check item: "What did the job description/advertisement call for?" Participants responded with options "Job candidate with strong math skills," "Job candidate who is punctual," "Job candidate with Social Media knowledge and expertise," and "None of the above." 97.5% of participants answered the attention check item correctly. Those who failed to answer correctly were shown the job advertisement and asked to answer the attention check item a second time before advancing.

The candidate was named "Steve Balmann," a made-up name, and participants were randomly assigned to view a professional, business-style headshot of a young vs. old version of Steve Balmann. The "young" vs "old" headshot pictures depicted a man seemingly in his 30's wearing a suit vs. a man seemingly in his 60's wearing a suit; other than the photograph, the resumes were identical across conditions (see Figures 3 and 4). In a pretest, the young and old candidates were not rated differently in fit for a marketing-related job that did not require social media expertise (i.e., "Account Manager and Liaison," $p = .319$). However, in this study, I expected participants to rate the young candidate a better fit for a "Social Media Manager" job than the old candidate, regardless of their blinding choice. Notably, it is not necessarily the case that a young person should be more knowledgeable about social media than an older person, and credentials were identical across the two resumes. Hence, I predicted that participants' estimates would be biased consistent with ageist stereotyping.

STEVE BALMANN

steve.balmann@gmail.com

Education

Duke University
Major: Marketing; GPA: 3.7/4.0

Employment Experience

RGI Marketing Corp (Raleigh, NC)
Brand Management Division

Maxwell Marketing & Design Inc. (Raleigh, NC)
Account Management Division

Skills

Brand management; Client relations; Analytics

Microsoft Office (Powerpoint; Excel; Word): Expert

SPSS Statistics: Knowledgeable

Languages

English: Fluent

Spanish: Conversational

Figure 3: Young Candidate Condition

STEVE BALMANN

steve.balmann@gmail.com

Education

Duke University
Major: Marketing; GPA: 3.7/4.0

Employment Experience

RGI Marketing Corp (Raleigh, NC)
Brand Management Division

Maxwell Marketing & Design Inc. (Raleigh, NC)
Account Management Division

Skills

Brand management; Client relations; Analytics

Microsoft Office (Powerpoint; Excel; Word): Expert

SPSS Statistics: Knowledgeable

Languages

English: Fluent

Spanish: Conversational

Figure 4: Old Candidate Condition

After participants reviewed the candidate's summary, they rated the candidate's fit for the Social Media Manager job. Participants responded to four items (a = .97), adapted from the person-organization fit literature (Lauver & Kristof-Brown, 2001). An example item was, "There is a good match between this job candidate and this job." Participants responded to each item on a scale of 1 (*strongly disagree*) to 7 (*strongly agree*).

2.3.2 Results and Discussion

When making a blinding choice in this mock hiring task, 41.2% of participants chose to see, versus be blind to, the candidate's name and photo information. Consistent with my expectations, a 2 (Blinding Decision: see vs. blind) x 2 (Candidate Age: young vs. old) ANOVA test of job fit revealed a main effect of Candidate Age, $F(1, 200) = 27.76$, $p < .001$, $\eta_p^2 = .10$, but not a main effect of Blinding Decision, $F(1, 200) = .71$, $p = .399$, $\eta_p^2 = .01$ or an interaction, $F(1, 200) = .01$, $p = .919$, $\eta_p^2 = .00$. Participants who chose to see the name and photo rated the young candidate a better fit for the job ($M = 4.95$, $SD = 1.31$) than the old candidate ($M = 3.97$, $SD = 1.64$) $t(82) = 3.02$, $p = .003$, $d = .66$. Moreover, among participants who chose to be blind to the name and photo, the young candidate was also rated a better fit for the job ($M = 4.79$, $SD = 1.47$) than the old candidate ($M = 3.77$, $SD = 1.52$) $t(118) = 3.66$, $p < .001$, $d = .68$.

These results demonstrated, consistent with a mental contamination account, that even participants who did not want to see or use the potentially contaminative information—the candidate's name and photo—made estimates that were biased when they were provided with that information. That is, these results are inconsistent with research suggesting that evaluators who prefer not to use "to be ignored" information can effectively factor it out of their evaluations when it is provided (Dietvorst & Simonsohn, 2018); in this study, even participants who chose to be blind to the candidate's name and photo made evaluations that were affected by that information

when it was provided to them. One issue with the materials used in this study is that, though the candidate's credentials were identical across resumes, they were both relatively scant. This may have made the older candidate less "believable" than the younger candidate. I return to and resolve this issue in Study 5, replicating an ageist bias with more conservative, matched materials.

2.4 Chapter 2 Summary

My goal in the studies presented in this chapter was to take an initial look at people's preferences for seeing, versus being blind to, potentially contaminative information across different evaluative contexts. In Study 1, 48% of participants chose to view a profile of a performer, and seeing the profile biased participants' estimates of performance downward, consistent with an anchoring effect (Wilson et al., 1996). In Study 2, 46% of participants chose to see the name and photograph of the college student they were ostensibly grading—an evaluation that would ostensibly determine the student's eligibility for a bonus—while only 18.4% indicated that they should see the name and photo, when asked. Participants who saw the name and photo information returned grades that were lower than those who did not, consistent with stereotypes about obese people (Levine & Schweitzer, 2015). In Study 3, participants took part in a mock hiring task, and 41.2% elected to see a job candidate's name and photo. Seeing the name and photo information made their judgments of the candidate's fit for the job susceptible to an ageist bias (North & Fiske, 2012), regardless of whether participants

chose to see the name and photo or chose to be blind to that information. In brief, across three different evaluative contexts, a considerable proportion of participants chose to see potentially contaminative information about the target of an evaluation, though it was purely optional to do so, rather than be blind to that information. I demonstrated three ways in which seeing such information can bias evaluations. In Chapter 3, I will begin to explore some of the mechanisms that may drive a choice to see, versus be blind to, potentially contaminative information.

3. Mechanism Studies I: Curiosity Motivation

In Chapter 1, I argued for the importance of understanding individual-level blinding preferences and proposed a theoretical framework, outlining various factors that could be expected to influence evaluators' self-blinding choices. The studies presented in Chapter 2 took the initial step of examining evaluators' baseline blinding preferences by having participants make an on-the-spot blinding choice—to see or be blind to potentially contaminative information about a target of evaluation—across a variety of evaluative contexts. The studies I discuss in this chapter will turn toward the mechanisms that may drive individual-level blinding preferences. Specifically, in the studies presented in this chapter, my goal was to test the motivational factor proposed in my theoretical framework: I explored whether a "hot" state of curiosity drives people to choose to see/learn potentially contaminative information. Moreover, I explored whether decision frames designed to reduce on-the-spot curiosity might encourage choices to be blind to potentially contaminative information.

In Study 4, I employed a performance evaluation context similar to the context of Study 1. In Study 5, I employed a mock hiring evaluation similar to that which was used in Study 3. Across these studies, I predicted that differences in reported curiosity motivation would explain choices to see (vs. be blind to) potentially contaminative information, and that decision frames that reduced participants' in-the-moment

curiosity would encourage a choice to be blind to potentially contaminative information relative to those that did not.

3.1 Study 4

In Study 4, I explored whether a motivational state of curiosity might drive blinding decisions. Using the same performance evaluation paradigm used in Study 1, I explored participants' actual blinding preferences (i.e., to see the performer's profile or be blind to that information), their beliefs about what their preferences would be in an imaginary version of that evaluation, and their beliefs about what their preferences should be, in an imaginary version of that evaluation, in order to make an unbiased evaluation. In this study, participants did not complete the performance evaluation after making their blinding choice. Instead, participants went on to indicate the extent to which their blinding choice was motivated by curiosity, fairness concerns, or accuracy concerns, and their lay beliefs about the effects seeing the profile would have on their evaluations' fairness and accuracy.

I made several predictions in Study 4. First, I expected that participants' self-reported curiosity would explain their choices to see potentially contaminative information—the profile of the performer. Consistent with research demonstrating that curiosity is a "hot" affective state (Loewenstein, 1994), I expected that participants would report greater curiosity when making an on-the-spot, plain blinding choice than when making a more deliberative blinding choice (i.e., what their blinding preferences

would be in a hypothetical sense or what they thought their blinding preferences should be, also hypothetically). When their blinding choice was framed in these more deliberative ways—a plain hypothetical choice or a normative (i.e., "should") hypothetical choice—I expected participants to be less motivated by curiosity and more motivated by concerns over fairness and accuracy (Chaiken, 1980), and hence more likely to favor blinding.

I formally predicted (in the preregistration document) that participants making blinding choices framed in the least deliberative and most intuitive fashion—a plain, on-the-spot choice—would be most inclined to see the profile, most motivated by curiosity, and least motivated by fairness and accuracy concerns, whereas participants making blinding choices framed in the most deliberative and least intuitive fashion—a hypothetical consideration of what their blinding preference should be—would be least inclined to see the profile, least motivated by curiosity, and most motivated by fairness and accuracy concerns (i.e., a "want vs. should" conflict; Bazerman et al., 1998). It follows that participants making a hypothetical, "plain" blinding choice in this study—more deliberative and "cold" a choice than those making a plain blinding choice on-the-spot but less deliberative than those making a hypothetical choice about what their blinding preference should be—would report blinding preferences and degree of curiosity motivation in between those indicating their in-the-moment "want" preferences and their hypothetical "should" preferences. I measured lay beliefs to test

whether participants in different conditions shared a general theory of how contaminative information affected fairness and accuracy or whether it differed by condition (perhaps reflecting motivated reasoning used to justify preferences). The pre-registration document for this study is available at https://aspredicted.org/3qg9x.pdf.

3.1.1 Method

3.1.1.1 Participants

I recruited 459 participants from Prolific Academic (M_{age} = 33.13, SD_{age} = 11.69, 42.6% female).

3.1.1.2 Materials and procedure

Study 4 employed the performance evaluation paradigm used in Study 1. The evaluation revolved around viewing a video of someone's performance in a pattern recognition task and estimating the quality of that performance. The blinding choice participants made was whether to see a profile of background information about the performer or be blind to that information.

Participants were randomly assigned to one of three conditions: *actual preference*, *normative preference*, or *hypothetical preference.* Participants either were told (*actual preference* condition) or imagined (*normative preference* and *hypothetical preference* conditions) that they would view a video of someone's performance on a pattern recognition task and estimate their performance, as in Study 1. Participants either were given (*actual preference* condition) or imagined being given (*normative preference* and

hypothetical preference conditions) the choice to view a profile of the performer before watching the video of the performance. The profile was described as including a "recent picture" of the performer (a change from Study 1).

Participants in the *actual preference* condition indicated their blinding preference in the same way as in Study 1, choosing whether to see or be blind to the profile of the performer in response to the question, "Would you like to view this person's profile information before you watch the video of their performance?" (yes/no). Participants in the *normative preference* condition were asked: "Imagine you are in this scenario. If you want to make an unbiased judgment of the person's performance, do you think you should view their profile information before you watch the video of their performance?" and responded with options "yes" or "no." Conversely, participants in the *hypothetical preference* condition were asked: "Imagine you are in this scenario. Would you like to view this person's profile information before you watch the video of their performance?" and responded with options "yes" or "no." As such, in the *actual preference* condition, participants made a plain, on-the-spot blinding choice. In the *hypothetical preference* condition, participants imagined taking part in the evaluation and indicated what their blinding preference would be, using the same choice wording as in the *actual preference* condition, only in an imaginary sense. In the *normative preference* condition, participants imagined taking part in the evaluation and indicated what they

thought their blinding preference should be in order to be unbiased, were then in such a scenario.

Finally, all participants completed three follow-up questions, in random order, assessing the motivations driving their blinding preferences. They also answered two questions about their lay beliefs, in random order. The order of these two sets of questions was itself randomized. The three motivation questions read: "To what extent was your choice about viewing the profile motivated by [concerns about treating the performer fairly / concerns about making an accurate estimate / curiosity]?" Participants responded to each question on a scale of 1 (*not at all*) to 7 (*completely*). The two lay beliefs questions read: "Please complete this sentence by choosing the option below that corresponds to your beliefs: Compared to not seeing the profile information, seeing the profile information would make my estimate of the performance in the video ...." and participants responded to each on a 7-point scale. The scale for the fairness question was anchored at 1 = "much less fair to the person I am evaluating" and 7 = "much more fair to the person I am evaluating." The scale for the accuracy question was anchored at 1 = "much less accurate" and 7 = "much more accurate." After participants completed these follow-up questions, they were debriefed and the study ended.

3.1.2 Results and Discussion

3.1.2.1 Preferences

I expected there to be a discrepancy between participants' actual and normative blinding preferences—between on-the-spot preferences and beliefs about what one should do, hypothetically, in order to be unbiased. The percentage of participants who chose to view the profile differed across the three conditions, $\chi(2)^2 = 47.52$, $p < .001$. As predicted, whereas only 9.7% of participants indicated they thought they should see the profile (*normative preference*), 45% chose to see it when it was readily available to them in the moment (*actual preference*), $\chi(1)^2 = 47.95$, $p < .001$. In the hypothetical preference condition, 32.5% chose to see the profile; this also differed from the normative preference condition, $\chi(1)^2 = 23.72$, $p < .001$, and the *actual preference* condition, $\chi(1)^2 = 5.04$, $p = .033$.

3.1.2.2 Motivations

Next, I examined the extent to which condition predicted participants' responses to the blinding preference motivation questions (see Table 2 for a breakdown of means by condition and motivation). As I expected, participants who made an actual blinding choice in the moment (*actual preference*) reported that their blinding choice was more motivated by curiosity, $t(450) = 4.84$, $p < .001$, and less by concerns over fairness, $t(450) = -5.98$, $p < .001$, and accuracy, $t(450) = -5.19$, $p < .001$, than participants indicating what they thought their blinding choice should be (*normative preference*). Participants in the

hypothetical preference condition reported motivations in-between those reported by participants in the other two conditions (see Table 2).

Table 2: Reported Motivations by Condition and Blinding Choice

	Chose Profile		Chose Blind		Total	
	M	*SD*	*M*	*SD*	*M*	*SD*
Fairness						
Actual	3.31	2.11	5.76	1.95	4.65	2.36
Hypothetical	3.65	1.90	6.19	1.32	5.37	1.94
Normative	5.47	1.19	6.06	1.58	6.00	1.55
Accuracy						
Actual	3.62	2.21	4.71	2.37	4.21	2.36
Hypothetical	4.35	1.72	5.38	1.81	5.05	1.84
Normative	5.93	0.96	5.37	1.92	5.42	1.86
Curiosity						
Actual	5.94	1.34	1.85	1.49	3.71	2.49
Hypothetical	5.23	1.74	1.94	1.41	3.00	2.17
Normative	5.00	1.73	2.22	1.67	2.49	1.87

In order to test the extent to which participants' reported motivations predicted their blinding choice, I performed separate mediation bootstrapping procedures (Preacher, Rucker, & Hayes, 2007) for each motivation question, in a PROCESS SPSS Macro (Model 4; 5,000 bootstraps; Hayes, 2012). In these analyses, I contrasted participants in the *actual preference* condition (coded as 0) with participants in the *normative preference* condition (coded as 1). Participants chose either to view the information (coded as 0) or to be blinded (coded as 1). Three separate analyses revealed an indirect effect of condition on blinding choice through curiosity motivation, $ab = 1.19$, $SE = .30$, 95% CI [.63, 1.81] (Figure 5), fairness motivation, $ab = .58$, $SE = .14$, 95% CI [.33,

.88] (Figure 6), and accuracy motivation, *ab* = .16, *SE* = .08, 95% CI [.02, .33] (Figure 7). That is, to the extent that participants in the *actual preference* condition reported being relatively more motivated by curiosity, and relatively less motivated by fairness and accuracy concerns, they were relatively less likely to favor blinding.

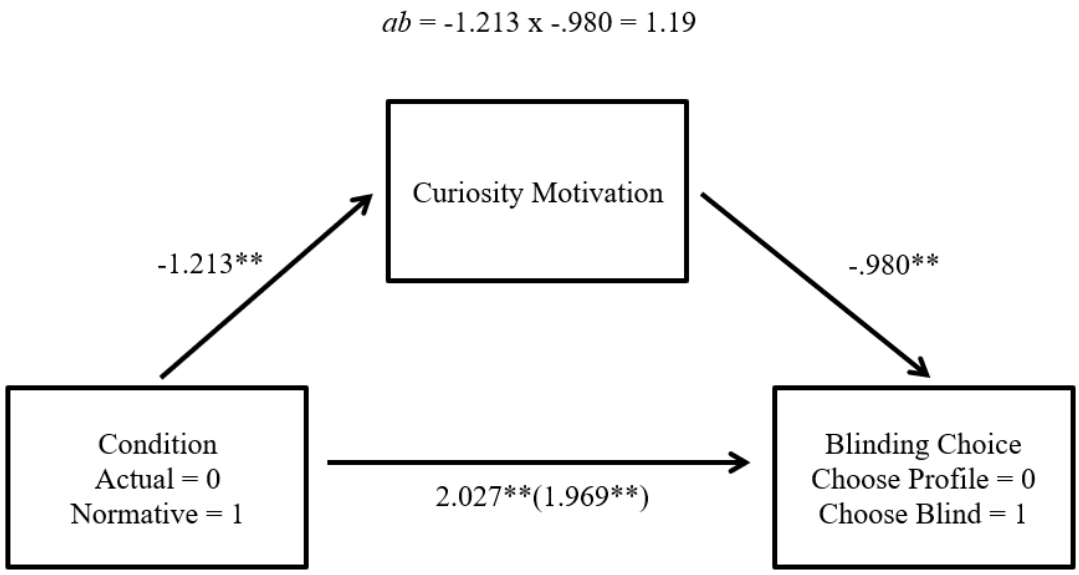

Figure 5: Mediation Model Showing the Indirect Effect of Condition on Blinding Choice Through Curiosity Motivation. Regression Coefficients are Unstandardized. **$p < .001$, *$p < .05$

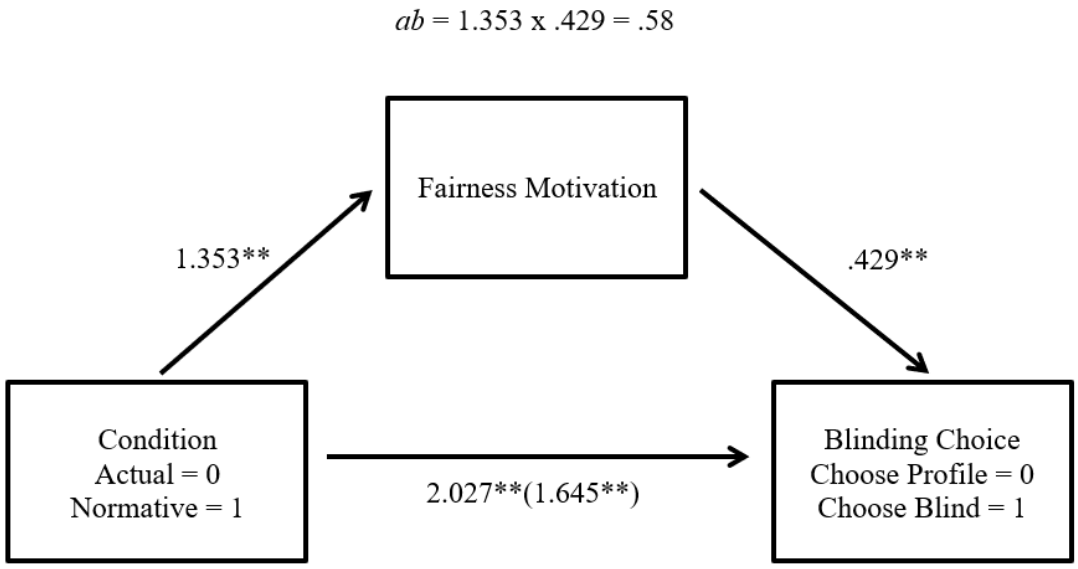

Figure 6: Mediation Model Showing the Indirect Effect of Condition on Blinding Choice Through Fairness Motivation. Regression Coefficients are Unstandardized. **$p < .001$, *$p < .05$

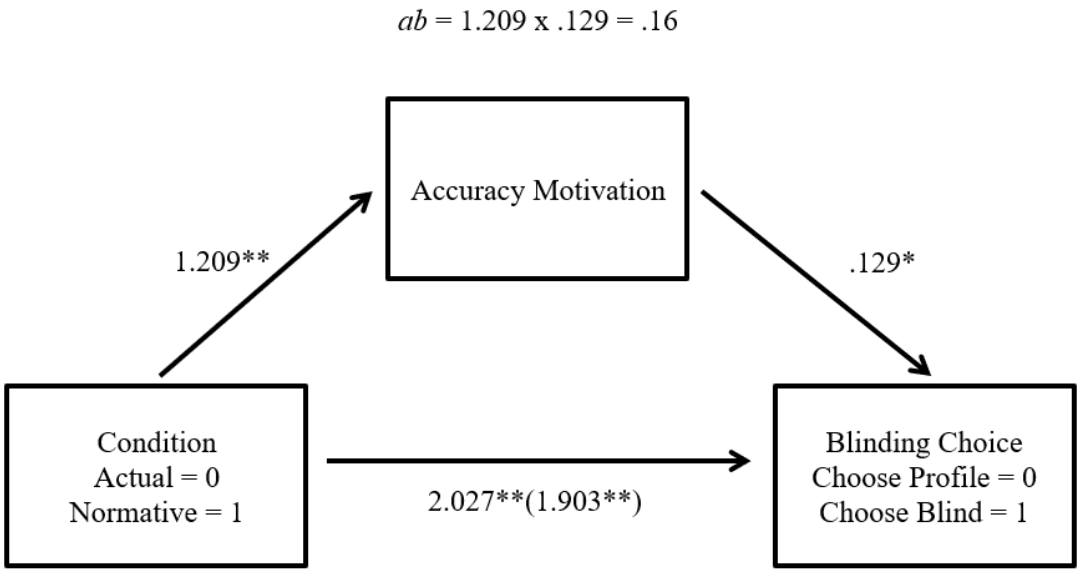

Figure 7: Mediation Model Showing the Indirect Effect of Condition on Blinding Choice Through Accuracy Motivation. Regression Coefficients are Unstandardized. **$p < .001$, *$p < .05$

3.1.2.3 Lay beliefs

Finally, I assessed participants' beliefs about the effects of seeing the profile on the fairness and accuracy of their estimates (see Table 3 for a breakdown of means by condition and lay belief item). This allowed me to test for motivated reasoning, whereby people more likely to choose to see contaminative info might also be relatively more likely to report that seeing contaminative information would increase their estimates' accuracy and fairness. All means were significantly lower than the scale midpoints of 4 (p's < .003), suggesting that all participants understood that choosing blinding would be more accurate and fairer than seeing the profile. With that said, participants in the *actual preference* condition, who were most inclined to see the profile, reported that seeing the profile would be less detrimental to estimate fairness and accuracy than those in the *normative preference* condition (p's < .03), who were least inclined to see it.

Table 3: Reported Lay Beliefs by Condition and Blinding Choice

	Chose Profile		Chose Blind		Total	
	M	*SD*	*M*	*SD*	*M*	*SD*
Effect on fairness						
Actual	4.24	1.07	2.94	1.14	3.53	1.28
Hypothetical	4.29	0.97	2.63	1.17	3.17	1.35
Normative	5.27	1.53	2.63	1.18	2.89	1.44
Effect on accuracy						
Actual	4.37	0.99	3.13	1.05	3.69	1.19
Hypothetical	4.29	1.11	3.17	1.12	3.53	1.23
Normative	5.00	1.41	3.21	1.15	3.38	1.28

These results may be evidence of motivated reasoning in advance of a blinding choice or of a post-hoc justification process. That is, participants may have adjusted their beliefs about the relationship between seeing the profile and estimate fairness/accuracy before their blinding choice, to conform with the choice they knew they wanted to make, or after their blinding choice, to conform with the choice they made. However, these results may also indicate belief-action consistency: The choice to be blind or see the background information was associated with lay beliefs within each condition, such that participants who elected to see the profile thought doing so would increase their estimates' fairness and accuracy, whereas those who elected to be blind thought seeing the profile would decrease their estimates' fairness and accuracy (p's < .001). Indeed, the fact that a subset of people state that seeing background information prior to judging performance increases accuracy and fairness may suggest that there is interesting and genuine heterogeneity in situational beliefs about information usefulness—as proposed in my theoretical framework—and I return to this possibility in Studies 7 and 8.

In sum, the results of Study 4 provided evidence consistent with my prediction that a "hot" state of curiosity would encourage a choice to see or learn potentially contaminative information. The more participants making an actual blinding choice in the moment were motivated by curiosity in their choice, the more likely they were to choose to see the profile of the performer. Decision frames that induced more deliberative thought versus impulsive action—the hypothetical and normative choice

conditions (Trope & Liberman, 2010)—reduced participants' curiosity as well as their inclination to see the profile. That is, a hot state of curiosity was less of a motivator of participants' blinding preferences when those preferences were elicited in a more abstract (i.e., imagined) manner. Moreover, when the blinding choice was framed around a goal to make an unbiased estimate (i.e., the normative choice condition), participants were least likely to choose to see the profile, and this was explained by a greater reported motivation to be fair to the performer as well as make an accurate estimate.

3.2 Study 5

In Study 5, I sought to perform another test of curiosity as a mechanism driving blinding preferences, using a mock hiring context similar to the context of Study 3. I reasoned that, if curiosity encourages people to choose to see potentially contaminative information in advance of an evaluation, they should be less likely to elect to see it before the evaluation if they know they will automatically receive it after the evaluation. That is, I expected that satisfying participants' curiosity to see potentially contaminative information—by guaranteeing their eventual receipt of that information—would reduce their inclination to see it in the moment. Relatedly, in Study 5, I tested whether participants who were encouraged to first perform a "blind" evaluation before learning potentially contaminative information would be less biased by that information, when

they eventually received it, than participants who elected to see it in advance of their evaluation.

As in Study 3, participants evaluated a hypothetical job candidate's fit for a job requiring social media expertise, and their blinding decision revolved around viewing the candidate's credentials alone or viewing the candidate's name and photograph in addition to their credentials. As in Study 3, participants who chose to view the candidate's name and photo were randomly provided with a depiction of a younger versus older candidate, and I tested for an ageist bias in their evaluations. In one condition, participants chose whether to see the candidate's name and photo, or be blind to that information, in advance of their evaluation. In a second condition, participants first evaluated the candidate on credentials alone and then chose whether to see the name and photo with the option, if they chose to see it, to revise their initial, blind evaluation. In a third condition, participants chose whether to see the candidate's name and photo in advance of their evaluation, not at all, or after first performing a blind evaluation (of credentials only). I expected that more participants would choose to see the name and photo information in advance of their evaluation in the first condition than in the third condition, where the option of first performing a blind evaluation and then seeing name and photo afterward was offered. I also predicted that, across conditions, participants who first performed a blind evaluation and then received the name and photo information would show less of an ageist bias in their revised

evaluations than participants who elected to see the name and photo in advance of the evaluation.

3.2.1 Method

3.2.1.1 Participants

I recruited 1199 participants from Prolific Academic (M_{age} = 35.06, SD_{age} = 12.63, 52.4% female).

3.2.1.2 Materials and procedure

Study 5 employed the job candidate evaluation paradigm used in Study 3. The evaluation revolved around deciding the extent to which a job candidate was a good fit for a job requiring social media expertise. The blinding choice participants made was whether to see the job candidate's name and photograph or be blind to that information.

There were two changes to the materials depicting the hypothetical job candidates that are important to note. In Study 3, participants viewed mock resumes of a job candidate, and it was possible—since the credentials were kept constant across conditions—that the older candidate's resume was less "believable" than the younger candidate's resume, or implied something negative about the older candidate, because their credentials were relatively sparse. To resolve this issue, participants in Study 5 evaluated a "summary" of the job candidate, prepared by an assistant, with only certain pieces of information from their overall job application included in the summary (i.e., responses to select questions from the application). In this way, it was more realistic for

the information in the younger and older candidates' summaries to be matched in content. I also updated the information provided in the summary, relative to the information provided in the resume in Study 3, to explicitly indicate that the candidate (in the old, young, and blind versions) had extensive social media knowledge and expertise. As such, these changes to the materials in Study 5, relative to Study 3, presented a more conservative test of an ageist bias: For both the young and old candidates, their credentials were matched in a realistic manner, and both had considerable expertise in social media (negating concerns that the younger candidate might seem implicitly better-suited at social media-related tasks than the older candidate).

Participants in Study 5 were randomly assigned to one of three conditions. In the *choice before* condition, participants read, "As the person making the hiring decision, it is up to you to decide what elements of the job application you want to be included in these summaries of the job candidates that you are reviewing. You are choosing those elements right now." Next, participants made a blinding choice in response to the following question: "Do you want the job candidate's name and photo displayed at the top of their summary? We will show you the summary of the job candidate on the following page based on your answer to this question" (Yes, show name and photo at top of summary vs. No, don't show name and photo at top of summary). Participants who chose to see name and photo information were randomly assigned to a young vs.

old candidate summary (see Figures 8 and 9), whereas those who chose to be blind to that information viewed a summary with name and photo information removed. Finally, participants rated the candidate's fit for the job using the same items used in Study 3 (a = .95).

Applicant photo:

Applicant name:

Steve Balmann

APPLICANT COLLEGE:

Where applicant went to college:

Duke University

Degree applicant received upon graduating:

Bachelor of Arts (With Honors)

Applicant's major:

Marketing

APPLICANT WORK EXP:

Applicant's most recent place of employment:

RGI Marketing Corp. (Raleigh office)

Applicant's position at most recent place of employment:

Social Media Marketing Associate (Online Marketing Division)

Applicant's marketing-related skills and experience:

General: Brand management, Client relations, Analytics

Social Media: Fluency with all major platforms (e.g., Facebook, Instagram), Skilled at boosting followership and growing engagement, Extensive experience writing and deploying Social Media content, Platform analytics knowledge

Figure 8: Young Candidate Summary

Applicant photo:

Applicant name:

Steve Balmann

APPLICANT COLLEGE:

Where applicant went to college:

Duke University

Degree applicant received upon graduating:

Bachelor of Arts (With Honors)

Applicant's major:

Marketing

APPLICANT WORK EXP:

Applicant's most recent place of employment:

RGI Marketing Corp. (Raleigh office)

Applicant's position at most recent place of employment:

Social Media Marketing Associate (Online Marketing Division)

Applicant's marketing-related skills and experience:

General: Brand management, Client relations, Analytics

Social Media: Fluency with all major platforms (e.g., Facebook, Instagram), Skilled at boosting followership and growing engagement, Extensive experience writing and deploying Social Media content, Platform analytics knowledge

Figure 9: Old Candidate Summary

In the *choice after* condition, participants first viewed the blind summary and rated the candidate's fit for the job (a = .92), without any mention of the option to see the candidate's name and photo. After participants made their blind assessment, they read, "Now you have the option to see the name and photo of the job candidate you just evaluated. If you choose to see the candidate's name and photo, you will view an updated version of their summary with their name and photo shown at the top. After viewing the updated summary, you will have the option to revise your assessment of

the job candidate's fit for the job, if you would like to." Participants then responded to the question, "Do you want to see the job candidate's name and photo?" (Yes vs. No). Participants who chose not to see name and photo advanced to the end of the study. Participants who chose to see name and photo were randomly provided with the young vs. old candidate summary and then read, "Now that you have viewed the updated version of the candidate's summary including their name and photo, you may choose to revise your assessment of the job candidate's fit for the job, if you would like to. If you choose to revise your assessment, you will be presented with your initial answers to the four statements assessing the candidate's fit for the job that you completed earlier, and will be prompted to enter new, revised answers." Participants who chose to see the name and photo then responded to the question, "Would you like to revise your assessment of the job candidate's fit for the job?" (Yes vs. No). Participants who elected to revise their assessment completed the same four job fit items as before, with their previous answers provided to them before revised answers were indicated (a = .95).

In the *omnibus choice* condition, participants were able to choose whether to perform a blind evaluation only, view the candidate's name and photo in advance of the evaluation, or first perform a blind evaluation and then see the candidate's name and photo with the option to revise. Participants read, "You have three options regarding the information included in the summaries: If you choose OPTION 1, you will see a summary containing only the candidate's education credentials and work history and

will then assess the candidate's fit for the job. If you choose OPTION 2, you will see a summary containing the candidate's name, photo, education credentials, and work history, and will then assess candidate's fit for the job. If you choose OPTION 3, you will first see a summary containing only the candidate's education credentials and work history, and will then make an initial assessment of the candidate's fit for the job. Following that initial assessment, you will be provided with the candidate's name and photo and will have the option to revise your initial assessment of the candidate's fit for the job." Participants then chose between options 1, 2, and 3. Those who chose options 1 or 2 viewed either a blind summary or a summary including name and photo information (randomized between participants) and evaluated the candidate's fit for the job (*a's* = .94 and .96). Those who chose option 3 first viewed the blind summary and evaluated the candidate's fit for the job (*a* = .90), then were randomly provided with either the young or old summary and chose whether to revise their initial evaluation (*a* = .95).

3.2.2 Results and Discussion

3.2.2.1 Preferences

My central goal in this study was to compare participants' inclination to see the name and photo information (vs. be blind to it) *before* performing their evaluation across the *choice before* condition, where they made that simple blinding choice alone, and the *omnibus choice* condition, where they had a third option to see the information later, after

first performing a blind evaluation. In the *choice before* condition, 43.6% of participants chose to view the name and photo (vs. be blind), and this differed from the 21.3% who chose to view the name and photo in advance of their evaluation in the *omnibus choice* condition (i.e., the percentage who chose option 2), $\chi(1)^2 = 45.55$, $p < .001$. That is, compared to participants making a straightforward blinding choice, participants who had the option to see potentially contaminative information after a blind evaluation were less likely to choose to see before. Of the remaining participants in the *omnibus choice* condition, 38.8% chose to make a blind evaluation only, and 39.8% chose to first make a blind evaluation and then see the candidate's name and photo information.

Within the *choice after* condition, I was not sure how many participants would choose to see the candidate's name and photo when spontaneously offered that information after first making a blind evaluation. Interestingly, 71.5% of participants chose to see the name and photo after making their blind evaluation. Following that choice, however, only 19.1% of participants in the *choice after* condition elected to revise their initial evaluation. In the *omnibus choice* condition, 17% of the participants who chose option 3—to first make a blind evaluation, then see the name and photo with the option to revise—chose to revise their initial, blind evaluation. Next, I explore whether participants who first made a blind evaluation and then saw name and photo information demonstrated less bias in their evaluations than participants who chose to see name and photo in advance of those evaluations.

3.2.2.2 Fit evaluations

I expected that participants who chose to view the name and photo information in advance of their evaluations would make evaluations consistent with an ageist bias. Indeed, in the *choice before* condition, participants who chose to see the candidate's name and photo judged the young candidate a better fit for the job (M = 5.85, SD = .99) than the old candidate (M = 5.04, SD = 1.41), $t(173) = 4.37$, $p < .001$, $d = .66$. Those who elected to be blind to the name and photo also rated the candidate a better fit for the job (M = 6.04, SD = .74), $t(313) = 8.14$, $p < .001$, $d = .90$. Comparing those who chose to be blind in the *omnibus choice* condition (option 1) with those who chose to see the name and photo in advance of their evaluation (option 2), participants who chose to see the candidate's name and photo judged the young candidate a better fit for the job (M = 6.11, SD = .82) than the old candidate (M = 5.13, SD = 1.46), $t(83) = 3.93$, $p < .001$, $d = .83$, and those who elected to be blind to the name and photo also rated the candidate a better fit for the job (M = 5.92, SD = .88), $t(190) = 4.25$, $p < .001$, $d = .66$. That is, even with more conservative materials that carefully matched the young and old candidates in terms of social media expertise, the ageist bias demonstrated in Study 3 replicated in this study.

I also predicted that participants who first made a blind evaluation and then were provided with the name and photo information (as in the *choice after* condition and those who chose option 3 in the *omnibus choice* condition) would demonstrate less ageist bias in their revised estimates than participants who received that information before

their evaluation (Festinger, 1962). To assess this, I contrasted the *choice before* and *choice after* conditions, and participants within the *omnibus choice* condition, in a 2 (Sequence: see photo before vs. after blind evaluation) x 2 (Age: young vs. old candidate) ANOVA (see Table 4). Comparing the evaluations of participants who chose to see the name and photo in the *choice before* condition with the revised evaluations of participants in the *choice after* condition, there was a main effect of sequence, $F(1, 459) = 13.78$, $p < .001$, $\eta_p^2 = .03$, and a main effect of age, $F(1, 459) = 31.12$, $p < .001$, $\eta_p^2 = .06$, qualified by a marginally significant interaction, $F(1, 459) = 3.78$, $p = .053$, $\eta_p^2 = .01$. Contrary to my predictions, in both the *choice before*, $F(1, 459) = 22.76$, $p < .001$ and *choice after* conditions, $F(1, 459) = 8.73$, $p = .003$, the young candidate was rated a better fit for the job than the old candidate. However, while fit evaluations were not different for the young candidate between conditions, $F(1, 459) = 1.54$, $p = .215$, the old candidate was rated a better fit for the job when participants first made a blind evaluation and then saw the candidate's picture (*choice after* condition) than when participants saw the picture in advance of their fit evaluation (*choice before* condition), $F(1, 459) = 16.24$, $p < .001$.

Table 4: Young vs. Old Candidate Fit Evaluations; Choice Before vs. Choice After Conditions

	Name and Photo Seen Before Evaluation		Name and Photo Seen After Blind Evaluation (Revised Evaluations)	
	M	*SD*	*M*	*SD*
Young	5.85	.99	6.04	.75
Old	5.05	1.41	5.65	1.26

I performed the same analysis on participants in the *omnibus choice* condition, comparing those who selected option 2 (see the candidate's name and photo before the evaluation) and option 3 (see the candidate's name and photo after first making a blind evaluation). A 2 (Sequence: see photo before vs. after blind evaluation) x 2 (Age: young vs. old candidate) ANOVA (see Table 5) did not find a main effect of sequence, $F(1, 240) = 2.22$, $p = .138$, $\eta_p^2 = .01$, but did find a main effect of age, $F(1, 240) = 24.12$, $p < .001$, $\eta_p^2 = .09$, qualified by a significant interaction, $F(1, 240) = 4.94$, $p = .027$, $\eta_p^2 = .02$. As before, the young candidate was rated a better fit for the job than the old candidate both among participants who saw the name and photo information before their evaluation, $F(1, 240) = 19.47$, $p < .001$ and participants who saw the name and photo information after their evaluation, $F(1, 240) = 5.22$, $p = .023$. However, while fit evaluations were not different for the young candidate between conditions, $F(1, 240) = .304$, $p = .582$, the old candidate was rated a better fit for the job when participants first made a blind evaluation and then saw the candidate's picture (*option 3*) than when participants saw the picture in advance of their fit evaluation (*option 2*), $F(1, 240) = 6.17$, $p = .014$.

Table 5: Young vs. Old Candidate Fit Evaluations; Omnibus Choice Condition

	Name and Photo Seen Before Evaluation		Name and Photo Seen After Blind Evaluation (Revised Evaluations)	
	M	*SD*	*M*	*SD*
Young	6.11	.81	6.01	.76
Old	5.13	1.46	5.64	1.13

In terms of my predictions, the results of this study were mixed. As I expected, participants were more inclined to see potentially contaminative information when making a plain blinding choice—to see it before an evaluation or not at all (43.6%)—than when the option to see it later was presented (21.3%). This provides more evidence that curiosity is a motivator of blinding preferences: It appeared that satisfying participants' curiosity to see the information *at some point* reduced their inclination to see it in the moment. Moreover, I also replicated the ageist bias demonstrated in Study 3 in this study, using more conservative (both candidates had social media experience) and more realistically matched ("summaries" of the candidates were presented, rather than resumes) materials. However, I found that the ageist bias demonstrated in this study persisted both among participants who saw potentially contaminative information before an evaluation and those who saw it after a blind evaluation. Encouragingly, the interaction results do suggest that having participants first perform a "blind" evaluation limited the extent to which this bias affected the older candidate, such that the gap in fit ratings between the older and younger candidates was closed, somewhat, when participants first evaluated a given candidate on his credentials alone and then saw his photograph. This suggests that, while the ageist bias demonstrated in this study and Study 3 may be difficult to quell entirely, encouraging evaluators to first perform a blind evaluation and then receive potentially contaminative information can at least improve evaluation outcomes for older candidates.

3.2.3 Follow-Up Study

Though participants were more inclined to be blind to potentially contaminative information in the moment when they knew they could learn it after making a blind evaluation, some did still choose to see it in advance of their evaluation, even though they could receive it later. In order to get a better understanding of the motivations of this subset of participants, I ran a follow-up study using 101 different participants on Prolific Academic (M_{age} = 31.65, SD_{age} = 10.74, 49.5% female). The follow-up study closely mirrored the design of Study 5, such that participants made a mock hiring decision for a social media manager job position. Participants read, "You are about to evaluate one job candidate's fit for the job. Imagine that you have in hand their resume, which contains their credentials (e.g., education history, work history, job-related skills). Imagine further that, after you evaluate the job candidate, you will see their name and photograph. That is, after you analyze their fit for the job, you will automatically be shown their name and photo. However, if you would like, you may see their name and photo now (in addition to their resume) before you analyze their fit for the job." Participants then responded to the question, "Do you want to see the job candidate's name and photo now, in addition to their resume, or just see their resume?" with options "See name and photo + resume" and "See resume only."

If participants chose to see the name and photo in addition to the resume, they were presented with one more question: "We are trying to understand decisions like

these better. Please select one of the following options to explain why you chose to see the name and photo now. You may select multiple options." Participants selected between options, "I was curious," "I thought seeing the name and photo now would help me make a better decision," and "Other." If participants selected the third option, they were prompted to explain what they meant by "Other" in an open-response format.

Consistent with the results of Study 5, about a quarter of participants—26.2%—chose to see the candidate's name and photo in advance of their evaluation, even though receiving it after the evaluation was guaranteed. Of those participants, 44.4% chose curiosity as at least one of the reasons why they elected to see the information in advance of their evaluation, 59.3% indicated that they thought it would help them make a better decision, and 11.1% selected "other." This amounted to three participants, who explained their choice in the following ways: "Name doesn't really matter, but the way I saw it, is that a real employer would want to know the name of/how to address their candidates." ; "I believe that how a person represents themselves and the presentation that they make to others says a lot about you [sic] they are. (Ie. If I was hiring a surf board salesman - I might want someone you [sic] looks like they are familiar with surfing and that culture)" ; "I believe attractiveness is always a plus when hiring and vise versa." In other words, these participants articulated a situational belief that the information would be helpful to them in various ways: It would help them categorize candidates, it would provide useful information about the candidate's self-presentation

choices, and it would provide a specific type of useful information: the candidate's degree of attractiveness.

This follow-up study confirmed the results of Study 5—that guaranteeing potentially contaminative information later in an evaluative process can decrease evaluators' desire to see it in the moment, consistent with a curiosity drive to see potentially contaminative information in the moment. Interestingly, some participants seemed to be *so* curious to see the job candidate's name and photo information that they wanted it as soon as possible (i.e., in advance of their evaluation of the resume). However, this follow-up study also helped to shed light on another of the potential mechanisms driving people's blinding choices proposed in the theoretical framework outlined in Chapter 1: a situational belief that seeing or learning potentially contaminative information would be helpful to a given evaluation. I explore these situational beliefs about information helpfulness vs. harmfulness as they relate to blinding preferences in greater depth in Chapter 5.

3.3 Chapter 3 Summary

In Chapter 3, I tested the motivational aspect of the theoretical framework proposed in Chapter 1: that a "hot" state of curiosity encourages choices to see potentially contaminative information. In Study 4, participants making a plain, on-the-spot blinding choice were relatively more motivated in their choice by curiosity than participants making more deliberative choices (e.g., what their blinding preference

should be to satisfy a goal to be unbiased), and this explained their stronger desire to see potentially contaminative information. In Study 5 and the associated follow-up study, guaranteeing participants the ability to receive potentially contaminative information *after* a blind evaluation reduced their desire to have it before the blind evaluation, relative to participants who made a plain choice to see or be blind to potentially contaminative information before their evaluation. Thus, the results of the studies presented in this chapter were supportive of my predictions that a "hot" motivational state of curiosity may drive blinding preferences.

4. Mechanism Studies II: Enduring Beliefs About Information and Bias

In Chapter 1, I proposed a theoretical framework that outlined both "hot"—impulsive and affective—and "cold"—deliberative and rational—drivers of individual-level blinding preferences. The studies presented in Chapter 3 offered evidence for the "hot" factor outlined in my theoretical model—a motivation to relieve momentary curiosity—as a mechanism. In the study I discuss in this chapter, my goal was to test one of the "cold" factors outlined in my theoretical model: general beliefs about information and bias that may have an enduring influence on blinding preferences across evaluative contexts.

Specifically, I explore the connections between blinding preferences and two beliefs: (i) that more information is always better to have when performing evaluations and (ii) that one is not susceptible to bias when performing evaluations. I explore both beliefs in one study, which employs the same job candidate evaluation paradigm as in Studies 3 and 5. In this study, I predicted that each enduring belief would encourage participants to see, rather than be blind to, potentially contaminative information.

4.1 Study 6

In Study 6, I explored whether (i) a belief that more information is better and (ii) a belief that one is not susceptible to bias would inform participants' blinding preferences, using scales I developed for this study. Specifically, I measured belief that more information is better using a 6-item scale. I de-composed my measurement of

belief that one is not susceptible to bias into the two sub-components discussed in Chapter 1: one 3-item scale captured self-assessed objectivity, and another 3-item scale captured self-assessed ability to correct for judgment bias. After participants completed these scales, they took part in a filler task that involved reading and writing about an unrelated article. Finally, participants made a mock hiring decision and chose whether to see or be blind to a set of information about a job candidate. Participants could opt in to seeing up to 10 pieces of information about the candidate, five of which I pre-judged to be diagnostic of the job candidate's value (i.e., "good" information) and five of which I pre-judged to be non-diagnostic and potentially contaminative (i.e., "bad" information). Overall, I predicted that all three enduring belief measures would be positively correlated with the number of pieces of "bad" information about the job candidate that participants chose to see.

I also varied whether participants made a plain blinding choice or a more deliberative choice. In the baseline condition, participants simply selected whichever pieces of information they wanted to see about the candidate. In the normative preference condition, participants selected the pieces of information they thought they should see in order to satisfy a goal to be unbiased. I included this variance in blinding decision framing in order to explore whether participants' enduring beliefs would be more or less predictive of their blinding preferences when approaching their blinding decision in a more deliberative way. For instance, if I am correct in my reasoning that

these enduring beliefs represent a relatively "cold"—i.e., deliberative and thoughtful—process by which one might choose to see potentially contaminative information (i.e., vs. a more "hot" process, like curiosity), it could follow that they would be more predictive of blinding preferences among participants making a relatively more deliberative (i.e., "should") choice.

4.1.1 Method

4.1.1.1 Participants

I recruited 402 participants from Prolific Academic (M_{age} = 32.63, SD_{age} = 12.19, 54.1% female).

4.1.1.2 Materials and procedure

Study 6 employed the job candidate evaluation paradigm used in Studies 3 and 5. The evaluation involved an assessment of a job candidate, as in Studies 3 and 5, but there was no mention of a job advertisement (for a job requiring social media expertise). Participants were simply tasked with hiring "the most competent candidate possible." As in Study 5, participants imagined that there were preparing to view summaries of job candidates, and that they would need to select the information to be included in the summaries. The blinding choice participants made involved selecting different pieces of information out of a list and is described in more detail below.

First, participants completed three different scales measuring three different beliefs, presented in a random order: a belief that more information is always better to

have when evaluating others, a self-assessed objectivity in evaluations of others, and a self-assessed ability to correct bias in evaluations of others. Each scale was preceded by the line, "When you are evaluating the strengths and weaknesses of another person (example: making a hiring decision), to what extent do you agree with the following statements…." The more information is better (MIB) scale contained the following items: "It is always beneficial to have more information about the person you are evaluating," "Sometimes, information about the person you are evaluating is not useful and should even be avoided" (reverse-coded), "The more information you can have about the person you are evaluating, the more accurate your evaluation will be, no matter what," "You can never have too much information about the person you are evaluating," "Sometimes, it's better not to have certain pieces of information about the person you are evaluating" (reverse-coded), and "Some pieces of information can actually make your evaluation of a person less accurate" (reverse-coded). After relevant reverse coding, these items were averaged to form an MIB composite (a = .87). The self-assessed objectivity scale (SAO), which I adapted from Uhlmann & Cohen (2007), contained the following items: "I am reasonable and logical," "My judgments are based on a logical analysis of the facts," and "My decision-making is rational and objective." These items were averaged to form an SAO composite (a = .84). The self-assessed ability to correct bias (SACB) scale contained the following items: "If I encounter information that could make me biased, I do not let that information affect me," "I am good at clearing

distracting information from my reasoning," and "I can recognize any feelings that might bias me and remove them." These items were averaged to form an SACB composite (a = .82).

After participants completed the three scales, they went on to complete a filler task. I included this filler task in order to increase the psychological distance between participants' responses to the scale items and their choices in the subsequent hiring task. The filler task involved reading a short article about plans to colonize Mars. After participants read the article, they responded to the question, "Having read the above article, how likely do you think it is that humans will eventually colonize Mars?" on a scale of 1 (*very unlikely*) to 5 (*very likely*). Next, participants were prompted to write an open-ended response to the following question: "Please use the space below to explain your answer to the above question. Why do you feel the way you do about the likelihood of humans colonizing Mars? Please try to write at least a couple sentences."[1]

Finally, participants completed the job candidate evaluation task. Participants were randomly assigned to one of two conditions: *baseline* vs. *normative*. All participants imagined that they were tasked with hiring someone for an open position at their company, that their goal was to hire the most competent candidate possible, and that their assistant was preparing summaries of job candidates for them to review.

---

[1] For those interested, participants found it likely, on average, that humans would be able to colonize Mars (M = 3.38, SD = 1.09). Their responses to the open-ended question often described their optimism in human ingenuity, which is nice.

Participants read, "As the person making the hiring decision, it is up to you to decide what information you want to be included in these summaries of the job candidates that you are reviewing. You are choosing those elements right now."

In the *baseline* condition, participants chose whether to see (or be blind to) up to 10 items about the job candidate, preceded by the line, "Now, you need to choose what information you want to be included in the summaries of the job candidates. Listed below are the 10 pieces of information available, for each job candidate, to be included in the summaries. Please select the pieces of information you want to be included in the summaries." Five of the items available were items that I pre-judged to be diagnostic of a job candidate's potential value and likely to be realistically incorporated into a hiring decision. Those five "good" items were: Name of candidate's college, Candidate's college GPA, Candidate's previous work experience, Candidate's relevant job-related skills, and Candidate's professional references. The remaining five items were ones that I pre-judged to be potentially contaminative. These five "bad" items were: Candidate's name and photo, Candidate's hometown, Candidate's favorite sport, Candidate's political affiliation, and Candidate's favorite hobbies. Participants selected as many items as they wished to see; items were presented in a random order. In the *normative* condition, participants made the same blinding choice, except that they read the following line in advance of their decision: "Imagine that your goal is to make an

unbiased hiring decision. With that goal in mind, please select the pieces of information you think should be included in the summaries."

I expected that most participants would select all of the "good" items. My chief concern in this study was the number of "bad" items participants chose to see. For each participant, I summed the number of "bad" items selected. I also summed the number of "good" items selected for comparison.

4.1.2 Results and Discussion

4.1.2.1 Preferences

Across conditions, participants chose to see most of the "good" items (M = 3.85, SD = 1.24) but few of the "bad" items (M = .81, SD = .99). This suggests that one way to encourage people to be blind to potentially contaminative information is to (a) ask them to opt in to seeing it and (b) position it next to some (more) clearly diagnostic and useful information. Consistent with the results of Study 4, participants in the *baseline* condition were more inclined to see the bad items (M = .97, SD = 1.05) than participants in the *normative* condition (M = .64, SD = .90), $t(419) = 3.42$, $p = .001$, $d = .34$. Interestingly, participants in the *baseline* condition were also more inclined to see the good items (M = 4.01, SD = 1.16) than participants in the *normative* condition (M = 3.68, SD = 1.30), $t(419) = 2.72$, $p = .007$, $d = .27$, indicating that the more deliberative framing of the blinding decision may have lead participants to be more careful about the items they selected overall.

4.1.2.2 Beliefs

My main prediction in this study was that participants' enduring beliefs—that more information is better (MIB), that their reasoning is objective (SAO), and that they are able to correct for bias (SACB)—would be positively correlated with their inclination to see the bad pieces of information about the job candidate. Across conditions, only MIB beliefs were correlated with the number of bad items selected (see Table 6). As such, these data were supportive of one of the enduring beliefs discussed in my theoretical framework—a belief that more information is always better to have—being a driver of blinding preferences, but not the other belief—a belief that one is not susceptible to bias. To explore the collective effects of the three beliefs on participants' inclination to see the bad items, I regressed the count of the bad items on all three beliefs in one model. In this model, both MIB and SACB emerged as significant predictors, whereas SAO did not. The more participants endorsed a belief that more information is better, the more bad items they chose to see, $b = .20$, $SE = .04$, $\beta = .23$, $t(412) = 4.75$, $p < .001$, whereas the more participants believed they were capable of correcting for bias in their evaluations, the *fewer* bad items they chose to see, $b = -.12$, $SE = .05$, $\beta = -.13$, $t(412) = -2.29$, $p = .022$, while self-assessed objectivity did not predict choices, $p = .864$.

Table 6: Means, Standard Deviations, and Intercorrelations Among Variables

	M	*SD*	1	2	3	4	5
1. MIB	4.16	1.16	-	-	-	-	-
2. SAO	5.69	.86	.18***	-	-	-	-
3. SACB	4.94	1.07	.22***	.52***	-	-	-
4. Bad	.81	.99	.21***	-.03	-.08	-	-
5. Good	3.85	1.24	.06	.04	-.07	.25***	-

Note. MIB = More info better composite, SAO = Self-assessed objectivity composite, SACB = Self-assessed ability to correct for bias composite, Bad = Number of bad items selected, Good = Number of good items selected. *p < .05, **p < .01, ***p < .001.

Next, I explored whether the relationships between participants' enduring beliefs and preferences for the bad information about the candidate differed between conditions (see Table 7 for correlations by condition). To assess this, I mean-centered beliefs and then multiplied them by condition (0 = *baseline,* 1 = *normative*) to form an interaction term. Then, in three separate regression models, I regressed number of bad items selected on the respective belief, condition, and the respective interaction term. No significant interactions emerged for any of the three beliefs, *p's* > .2. As such, it did not appear that the frame of participants' blinding choice (i.e., plain vs. should) impacted the extent to which their enduring beliefs predicted their blinding preferences.

Table 7: Correlations Between Beliefs and Bad Items Selected by Condition

	More Info Better	Self-Assessed Objectivity	Self-Assessed Ability to Correct Bias
Baseline			
Bad item count	.19***	-.09	-.10
Normative			
Bad item count	.25***	.03	-.05

Note. *p < .05, **p < .01, ***p < .001

4.2 Chapter 4 Summary

My goal in Chapter 4 was to explore whether the enduring beliefs proposed in my theoretical framework—that more information is always better to have and that one is not susceptible to bias—would impact evaluators' blinding preferences. Participants completed a job candidate evaluation task similar to the paradigm used in Studies 3 and 5, but made a different type of blinding decision—they chose to see (vs. be blind to) up to ten pieces of information about a job candidate, five of which I judged to be potentially contaminative in nature and five of which I judged to be useful. A belief that more information is better predicted inclination to see the contaminative items but not the useful items, and the relationship between a belief that more information is better and inclination to see the bad items persisted whether participants made a plain, on-the-spot blinding choice or a normative choice. However, neither of the sub-components of a belief that one is not susceptible to bias—self-assessed objectivity or self-assessed ability to correct for bias—explained participants' blinding preferences.

Thus, this study provided evidence for one of the two enduring beliefs discussed in Chapter 1 being a driver of blinding preferences. However, there are reasons to interpret this takeaway cautiously. First, there were tautology concerns in this study, particularly with regard to the relationship between the MIB measure and blinding preferences. I found that participants who indicated a belief that more information is better did, in a subsequent blinding choice, select more information. This concern may

be mitigated by the fact that I inserted a filler task between the MIB scale and the elicitation of participants' blinding preferences, which should have reduced the psychological link between the two. Moreover, I did not find that responses to the MIB scale predicted choices to see the good items; only the bad items. Still, future research could build on this examination by using a time-lagged design to negate tautology concerns.

Another notable result in this study was that the relationship between a belief that more information is better and participants' inclination to see (vs. be blind to) the bad information was not moderated by condition (i.e., baseline vs. normative choice). I argued in Chapter 1 that the enduring beliefs proposed in my theoretical framework represented "cold" beliefs, or processes, which may drive evaluators' blinding preferences. It would follow from this contention that participants in this study who made a normative blinding choice—a relatively more thoughtful choice that required deliberation on what one should do to be unbiased—would demonstrate a stronger connection between MIB beliefs and blinding preferences than those who made a baseline choice, which was relatively less thoughtful and more on-the-spot in nature. I did not find such an interaction, though the correlation between MIB beliefs and blinding preferences was stronger for participants making a normative choice vs. those making a baseline choice. It is possible that the general blinding decision frame employed in this study induced more deliberative reasoning across the board—

participants had to peruse a list of ten possible items and select amongst them—and this may explain why I did not find a difference in the relationship between MIB beliefs and blinding preferences by condition.

It is also noteworthy that this study did not find a relationship between the two measures assessing a sense of invulnerability to bias—self-assessed objectivity and self-assessed ability to correct bias—and participants' blinding preferences. On the one hand, this may be evidence that my theoretical framework is flawed—that an enduring belief that one is invulnerable to bias does not drive individual-level blinding preferences. On the other hand, it is possible that aspects of the study design clouded what is a true relationship in actuality. For instance, neither the SAO scale nor the SACB scale were as reliable as the MIB scale, and both were right-skewed across the sample as a whole (with the mean of SAO approaching ceiling), whereas MIB was normally distributed. It is possible that there was not sufficient variance in these beliefs to capture a relationship to blinding preferences—this would be consistent with research demonstrating that people are overconfident in a general sense (Soll & Klayman, 2004) and specifically regarding their susceptibility to bias relative to others (Pronin et al., 2004; Van Boven et al., 2003). Perhaps future research could strive to manipulate (high vs. low endorsement of) enduring beliefs that one is not susceptible to bias in order to better test their relationship to blinding preferences.

5. Mechanism Studies III: Using Situational Beliefs About Information and Bias to Encourage Blinding

Across Chapters 2-4, participants often chose to see potentially contaminative information when making an on-the-spot decision before an impending evaluation. However, in Studies 2, 4, and 6, a much smaller subset of participants indicated they thought they should see the potentially contaminative information, when asked to indicate normative preferences, than those who made a plain blinding choice. Moreover, in Study 4, all participants indicated, on average, that seeing the contaminative information available to them (i.e., a profile of a performer) would harm their estimates' fairness and accuracy. This pattern of results may indicate two things: First, there does appear to be a subset of people who choose to see potentially contaminative information because they think they should do so—perhaps in accordance with situational beliefs that a given piece of information will actually be helpful to them. Second, this group of people who choose to see potentially contaminative information because they think they should do so—explaining blinding preferences when framed as a normative choice and a subset of preferences when framed as a plain, on-the-spot choice—is comparatively small. Most participants across these studies seemed to hold a situational belief that the available information would not be helpful (e.g., would not increase estimate accuracy or fairness).

None of the previous studies captured participants' situational beliefs about the helpfulness of available, potentially contaminative information in a direct manner.

Studies 2, 4, and 6 assessed these beliefs indirectly when participants were asked to indicate their normative blinding preferences (i.e., a small subset thought they "should" see the information, ostensibly due to situational beliefs about its potentially helpful effects). Study 4 directly asked participants to indicate whether the profile of the performer would increase or decrease their estimates' fairness and accuracy, but participants responded to those questions after a blinding choice had already been made, making it difficult to disentangle their reported beliefs from post-hoc justification pressures.

Hence, in the studies in this chapter, I sought to directly test the final factor in my proposed theoretical framework—evaluators' situational beliefs about information helpfulness—as a predictor of blinding preferences. In both studies, I asked participants to reflect on their situational beliefs about available information—whether it would help or harm accuracy in Study 7 and whether it would increase or decrease bias susceptibility in Study 8—*before* a blinding choice was made. This allowed me to capture their situational beliefs untainted by post-hoc justification pressures. This design choice also presented the opportunity to test a "debiasing" strategy. Results of the previous studies suggested that, in the studies in this chapter, most participants would *not* report situational beliefs that potentially contaminative information would be helpful to their evaluations. Therefore, I explored whether making these beliefs salient to participants

before a blinding choice was made would encourage, on average, a choice to blind their judgment (Kahneman, 2011; Festinger, 1962).

Study 7 employed the performance evaluation paradigm used in some of the previous studies. Study 8 employed the job evaluation paradigm used in others. In both studies, some participants first reflected on the available, potentially contaminative information's likelihood of being helpful vs. harmful before they made a blinding choice. Across both studies, I expected that participants who reflected on their situational beliefs before making a blinding choice, relative to those who did not, would be more inclined to choose to be blind.

5.1 Study 7

Study 7 employed the same performance evaluation paradigm as Studies 1 and 4. My main goal in Study 7 was to test whether making participants' situational beliefs about the helpfulness vs. harmfulness of potentially contaminative information salient to them before a blinding choice would encourage a choice to be blind. I also tested a different intervention for encouraging a choice to be blind: a monetary incentive. Before making their blinding choice (i.e., seeing a profile of the performer or being blind to the profile), participants either reflected or did not reflect on the effect seeing the profile would have on estimate accuracy and were also either given or not given a monetary incentive for the accuracy of their final estimate. That is, some participants articulated their situational beliefs about the effect the available, potentially contaminative

information—the profile of the performer—would have on the accuracy of their estimate before making a blinding choice, while others did not. Some participants also received a monetary incentive for estimate accuracy whereas others did not. Following the choice to be blind or not, I measured the motivations (i.e., curiosity, accuracy, and fairness) underlying that choice in the same manner as in Study 4. I also measured participants' trait Need for Cognition (NFC; Petty, Cacioppo, & Kao, 1984) as a possible individual difference moderator of blinding preferences.

Consistent with research demonstrating that incentives are ineffective in fixing most cognitive biases (Soll, Milkman, & Payne, 2015), I did not expect that providing participants with a monetary incentive for accuracy would affect their blinding choice. However, I did predict that the reflection intervention would encourage a choice to be blind. In Study 4, participants' blinding choices generally aligned with their lay beliefs about the effect seeing the profile would have on estimate accuracy, and most reported a belief that seeing the profile would not help accuracy. Consequentially, in this study, I expected that participants who reflected on the effect of seeing the profile on estimate accuracy—before making a blinding choice—would be relatively more inclined to be blind, thanks to greater self-awareness about the potential harm of the profile information (Kahneman, 2011). The pre-registration document for this study is available at http://aspredicted.org/blind.php?x=n4ep28.

5.1.1 Method

5.1.1.1 Participants

I recruited 1,207 participants from Prolific Academic (M_{age} = 34.35, SD_{age} = 11.67, 48.7% female).

5.1.1.2 Materials and procedure

Study 7 employed the performance evaluation paradigm used Studies 1 and 4. The evaluation revolved around viewing a video of someone's performance in a pattern recognition task and estimating the quality of that performance. The blinding choice participants made was whether to see a profile of background information about the performer or be blind to that information.

Study 7 had a 2 (Reflection: yes vs no) x 2 (Incentive: yes vs no) design. Participants read that they would view a video of someone's performance in a pattern recognition task and estimate that performance and chose whether they would like to view a profile of the performer or not, containing the same information as in Study 4. As in Study 4, participants in Study 7 did not actually view the video or estimate performance.

Participants assigned to reflect on their situational beliefs about accuracy responded to the lay belief question about accuracy used in Study 4 before making their blinding choice: "Please complete this sentence by choosing the option below that corresponds to your beliefs: Compared to not seeing the profile information, seeing the

profile information would make my estimate of the performance in the video ...." (1 = *much less accurate*, 7 = *much more accurate*). I note that this intervention is more subtle than direct instructions that profile information has the potential to be harmful to decision accuracy—rather, participants were nudged to reach this conclusion through their own reflection on their own situational beliefs. Participants assigned to receive an incentive were told, before making their blinding choice, that they should try to be as accurate as possible in their estimate of the performance, and that participants who made estimates within 2% of the exact performance level in the video would receive a bonus equivalent to the base pay of the study[2]. Participants who did not reflect or receive a monetary incentive simply reported their blinding choice as in the *actual preference* condition in Study 4.

After participants made their blinding choice, they completed the three motivation items used in Study 4, in random order, as well as the 18-item Need for Cognition scale (NFC; Petty, Cacioppo, & Kao, 1984), also in random order (a = .93). NFC measures a trait-level preference for engaging in deliberative thought. An example item from this scale was, "I would prefer complex to simple problems." Participants responded to each NFC item on a 7-point scale anchored at 1 (*strongly disagree*) and 7

[2] Participants assigned to receive an incentive were paid the bonus described to them, though they did not actually estimate performance.

(*strongly agree*). The order of the sets of motivations and NFC items was itself randomized.

5.1.2 Results and Discussion

5.1.2.1 Preferences

To examine the effects of the reflection and incentive interventions on blinding choice, I regressed choice on Reflection (yes vs no), Incentive (yes vs no), and the interaction term in a binary logistic regression. Consistent with predictions, there emerged a main effect of Reflection b = .19, Wald $\chi^2(1)$ = 10.57, p = .001, but not of Incentive b = -.04, Wald $\chi^2(1)$ = .45, p = .501, or an interaction b = .03, Wald $\chi^2(1)$ = .24, p = .622. That is, participants who completed the reflection item before making their blinding choice were less likely to choose to see the profile (45.1%) than those who did not complete the reflection (54.5%), as I expected, but providing participants with a monetary incentive for accuracy did not affect blinding choice. Looking specifically at participants who did not complete the reflection item, those who received a monetary incentive for accuracy were no more likely to choose to see the profile (56.1%) than those who did not (i.e., baseline; 52.8%), $\chi(1)^2$ = .68, p = .41. Similarly, among participants who did reflect, the percentage who chose to view the profile did not differ between those who received a monetary incentive (45.4%) and those who did not (44.9%), $\chi(1)^2$ = .016, p = .899.

These results indicate that encouraging reflection on situational beliefs about the effect of seeing contaminative information on one's estimate accuracy, immediately before a blinding choice was made, was an effective strategy to induce a choice to be blind whereas providing a monetary incentive for estimate accuracy was not. I expected that making participants' situational beliefs about the effect-on-accuracy of the information available to them, before their blinding choice, would encourage a choice to be blind because most participants would *not* indicate a belief that seeing the profile would help estimate accuracy. Indeed, in Study 4, only 17% of participants indicated on the lay belief item about accuracy that seeing the profile would help estimate accuracy (i.e., responses above the midpoint of the scale). Similarly, in this study, 26.5% of participants who completed the reflection item (which was the same as the lay belief item about accuracy used in Study 4) responded above the midpoint of the scale, thereby indicating a belief that seeing the profile would help estimate accuracy. 83.8% of participants in this study who indicated on the reflection item that seeing the profile would help accuracy chose to view it. Conversely, of the remaining 73.5% of participants—who ranged from believing the profile would harm accuracy (responses below the midpoint) to believing it would have no effect (responses at the midpoint)—only 31.2% elected to view the profile. Therefore, in line with predictions, it seems the reflection intervention encouraged a choice to be blind to the profile by making people's

situational beliefs about the value of seeing contaminative information salient to them before their blinding choice.

5.1.2.2 Motivations

I also explored the extent to which participants' blinding choices were motivated by concerns over fairness, accuracy, and curiosity, using the same motivation items used in Study 4 (see Table 8 for a breakdown of means by condition and motivation). A 2 (Reflection: yes vs no) x 2 (Incentive: yes vs no) ANOVA test of fairness motivation revealed a main effect of Incentive $F(1, 1201) = 11.99$, $p = .001$, $\eta_p^2 = .01$, but not of Reflection $F(1, 1201) = .72$, $p = .396$, $\eta_p^2 = .001$, or an interaction $F(1, 1201) = .30$, $p = .597$, $\eta_p^2 = .000$, indicating that providing participants with a monetary incentive for accuracy in their estimate decreased the extent to which their blinding choice was motivated by treating the performer fairly ($M = 4.35$, $SD = 2.23$) relative to those who did not receive an incentive ($M = 4.79$, $SD = 2.18$). However, a 2 (Reflection: yes vs no) x 2 (Incentive: yes vs no) ANOVA test of accuracy motivation did not reveal a main effect of Reflection $F(1, 1201) = .001$, $p = .982$, $\eta_p^2 = .000$, Incentive $F(1, 1201) = 2.08$, $p = .150$, $\eta_p^2 = .002$, or an interaction $F(1, 1201) = .24$, $p = .627$, $\eta_p^2 = .000$. Finally, a 2 (Reflection: yes vs no) x 2 (Incentive: yes vs no) ANOVA test of curiosity motivation revealed a main effect of Reflection $F(1, 1201) = 9.81$, $p = .002$, $\eta_p^2 = .01$, but not of Incentive $F(1, 1201) = .16$, $p = .686$, $\eta_p^2 = .000$, or an interaction $F(1, 1201) = .09$, $p = .770$, $\eta_p^2 = .000$, indicating that asking participants to reflect on the effect seeing the profile might have on estimate accuracy

decreased the extent to which their blinding choice was motivated by curiosity (M = 3.68, SD = 2.41) relative to participants who did not reflect (M = 4.12, SD = 2.45).

Table 8: Reported Motivations by Condition

	Did Reflect				Did Not Reflect			
	Incentive Provided		Incentive Not Provided		Incentive Provided		Incentive Not Provided	
	M	*SD*	*M*	*SD*	*M*	*SD*	*M*	*SD*
Fairness	4.33	2.24	4.70	2.22	4.37	2.23	4.87	2.14
Accuracy	4.28	2.25	4.52	2.27	4.34	2.23	4.47	2.20
Curiosity	3.63	2.44	3.73	2.38	4.11	2.48	4.13	2.43

As detailed above, participants who completed the reflection question were relatively less likely to choose to see the profile. I posited that the reflection intervention was effective because it encouraged participants to act on their beliefs about the relationship between the profile and estimate accuracy. These results suggest the reflection intervention may also have been effective because it decreased participants' curiosity to see the profile. As in Study 4, curiosity motivation was strongly related to blinding choice (0 = chose profile, 1 = chose blind), r = -.828, p < .001, such that lower curiosity motivation predicted a choice to be blind to the profile. These explanations are not mutually exclusive: The reflection intervention may have reduced participants' focus on their curiosity to see the profile and shifted it to a consideration of the effects seeing the profile would have on the accuracy of their estimates.

In comparison, providing an accuracy-related incentive may have been ineffective in encouraging blinding in that it affected neither participants' curiosity nor their accuracy motivation, but decreased a motivation to be fair to the performer. As in Study 4, fairness motivation was related to a choice to be blind, $r = .3$, $p < .001$. In sum, results conformed with expectations: Providing participants with a monetary incentive did not affect their blinding choice, but asking participants to articulate their situational beliefs about the helpfulness vs. harmfulness of potentially contaminative information before a blinding choice was made encouraged a choice to be blind[3].

5.2 Study 8

In Study 7, roughly a quarter of participants believed that seeing the profile of the performer would increase the accuracy of their estimates; the rest either believed that the profile was irrelevant or that seeing it would harm their accuracy. Hence, asking participants to articulate these situational beliefs about the helpfulness vs. harmfulness of the profile before they made a blinding choice encouraged them to choose blinding when they did make their choice. In Study 8, I returned to a different type of reflection-based intervention, which I first tested in Study 2, exploring whether participants who first reflected on what their blinding choice should be to be unbiased before making an

[3] I also examined whether participants' Need for Cognition (NFC), a trait-level preference for deliberative thought, impacted their blinding choice. Across conditions, NFC was weakly correlated with blinding choice (0 = chose profile, 1 = chose blind), $r = .062$, $p = .031$, suggesting that the more participants preferred to engage in deliberative thought, the more likely they were to choose to be blind.

actual choice would be more inclined to choose blinding relative to those who made a plain blinding choice alone.

Study 8 employed the same job candidate evaluation paradigm used in Studies 3 and 5, such that participants evaluated a job candidate's fit for a job requiring social media expertise. The blinding choice concerned whether to evaluate job candidates on their credentials alone or to be provided with personally identifying information (i.e., name and photograph) in addition to credentials. To build on Studies 3 and 5, which demonstrated an ageist bias in evaluations of the job candidate, in this study, I tested for an attractiveness bias. Participants who chose to view the job candidate's name and photograph were randomly assigned to an attractive vs. unattractive candidate, using pictures from the Chicago Face Database that were pre-rated to be high vs. low on attractiveness (Ma, Correll, & Wittenbrink, 2015). Consistent with evaluative biases in favor of attractive people (Langlois et al., 2000), I expected that participants would rate the attractive candidate a better fit for the job than the unattractive candidate, even given matched (and impressive) credentials.

Participants either made an actual blinding choice only, or first reflected on what their blinding choice should be in order to perform an unbiased evaluation and then made an actual blinding choice. As in Studies 2, 4, and 6, I expected more participants to prefer blinding when expressing what they should do to be unbiased than when simply making a blinding choice in the moment. As in Study 7, I expected consistency pressures

(Festinger, 1962) to lead participants who first expressed what they should do to be more likely to choose blinding when making a subsequent, actual blinding choice than when making an actual blinding choice alone.

One final change of note in Study 8, relative to the preceding studies, was that I specifically sampled participants with managerial experience in their careers via Prolific Academic. I adopted this new sample criterion in Study 8 in order to rule out the possibility that people with managerial experience will not choose to see potentially contaminative information (perhaps due to advanced training). In this managerial sample, 77.8% of participants had made at least one hiring decision over the course of their careers. Because the blinding decision in Study 8 was embedded in a mock hiring task, the use of this select sample—who had both managerial experience and, for the most part, actual hiring experience—allowed me to better understand the "real-world" generalizability of the effects demonstrated in the preceding studies. The pre-registration document for this study is available at http://aspredicted.org/blind.php?x=hi4nj9.

5.2.1 Method

5.2.1.1 Participants

I recruited 803 managerial-level participants from Prolific Academic (M_{age} = 39.32, SD_{age} = 12.49, 53.2% female). Participants had an average of 18.79 years of work experience (SD = 11.87), had made an average of 27.92 hiring decisions in their careers

(*SD* = 195.50), and had an average number of 4.96 direct reports in their current jobs (*SD* = 21.97). Participants indicated that they were presently employed in 21 different industries, with the most common being Education (13.3%), Information Technology (10.6%), and Health Care (8.2%).

5.2.1.2 Materials and procedure

Study 8 employed the job candidate evaluation paradigm used in Studies 3 and 5. The evaluation revolved around deciding the extent to which a job candidate was a good fit for a job requiring social media expertise. The blinding choice participants made was whether to see the job candidate's name and photograph or be blind to that information.

Study 8 had a 2 (Choice: actual vs. intervention) x 2 (Candidate: attractive vs. unattractive) design. Participants were shown a mock job advertisement for a "Social Media Manager and Liaison," imagined that an assistant was preparing summaries of the job candidates for them based off of the information in their job applications, and imagined that they would review one such summary and assess the candidate's fit for the job. Participants read: "As the person making the hiring decision, it is up to you to decide what elements of the job application you want to be included in these summaries of the job candidates that you are reviewing. You are choosing those elements right now."

Next, participants in the *actual choice* condition made a plain blinding choice in response to the following: "Do you want the job candidate's name and a photo provided by their previous employer (from their employee ID) displayed at the top of their summary? We will show you the summary of the job candidate on the following page based on your answer to this question." Participants responded with options, "Yes, show name and photo at top of summary" and, "No, don't show name and photo at top of summary." Conversely, participants in the *intervention* condition first reflected on what they should do in order to be unbiased ("*intervention reflection*") and subsequently made the same blinding choice as participants in the *actual choice* condition ("*intervention choice*"). That is, participants in the *intervention* condition first responded to the following: "In order to make an unbiased judgment, do you think you should see the job candidate's name and a photo provided by their previous employer (from their employee ID)?" (yes/no), and then made their blinding choice in the same manner as in the *actual choice* condition. As in Study 7, this intervention constituted an independent reflection on the part of participants rather than direct instruction that seeing the contaminative information (i.e., name and photo of candidate) could bias their judgments.

In either condition, participants who chose not to view the job candidate's name and photo when making their blinding choice saw a "blind" summary of the candidate containing only information about their education and their most recent work

experience. As in Study 5, the summary explicitly indicated that the candidate had a high degree of knowledge of social media and experience in social media marketing. Participants who chose to view the job candidate's name and photo viewed summaries containing the same information as in the "blind" summary, plus a name and photo at the top.

The candidate was named "Steve Balmann," a made-up name, and participants were randomly assigned to view a picture of an attractive (Figure 10) vs. unattractive (Figure 11) White male, taken from the Chicago Face Database (CFD; Ma et al., 2015). Pictures in the CFD are standardized, such that the people photographed are matched on their expression and garb. I chose pictures of White men who were rated in the CFD norming data as similar in terms of age ($M_{attractive}$ = 25.82 vs. $M_{unattractive}$ = 23.32) but disparate in terms of attractiveness ($M_{attractive}$ = 4.66 vs. $M_{unattractive}$ = 1.81). I elected to use this specific photo pairing because they were the most disparate possible in terms of attractiveness (i.e., most-attractive vs. second-least-attractive) in the norming data while still rated to be similar in age. For instance, the least attractive White, male photo in the norming data was rated substantially older than the most attractive photo (M = 38.91). As outlined above, the candidate's photo was described to participants as one that was provided from their previous employer (from their employee ID). I framed the photo in this way so that it was realistic for the candidate's photo to depict him wearing the grey tee-shirt all CFD subjects wear, as well as to minimize concerns that either candidate had

made unprofessional self-presentation choices. Even though the two candidates were matched in terms of (strong) credentials, I expected participants to rate the unattractive candidate a worse fit for the job than the attractive candidate and the blind summary. After participants reviewed the candidate's summary, they rated the candidate's fit for the Social Media Manager job using the same items from Studies 3 and 5 (a = .96).

Figure 10: Attractive Candidate Summary

Figure 11: Unattractive Candidate Summary

5.2.2 Results and Discussion

5.2.2.1 Preferences

First, I examined the effect of the intervention on participants' blinding preferences. Consistent with the results the previous studies, more participants chose to view the contaminative information—the job candidate's name and photo—when making an actual choice in the moment (*actual choice*; 37.6%) than when considering what they should do in order to be unbiased (*intervention reflection*; 15.1%), $\chi(1)^2 = 52.45$,

$p < .001$. I expected that participants who first reflected on what they should do to be unbiased (*intervention reflection*) and then made a subsequent, actual blinding choice (*intervention choice*) would be relatively less likely to choose to see contaminative information than those who made an actual blinding choice alone. Confirming the success of this reflection intervention, participants who made a blinding choice after first reflecting on what they should do (*intervention choice*) were less likely to choose to see the candidate's name and photo (21.5%) than those who made an actual blinding choice alone, $\chi(1)^2 = 24.87$, $p < .001$.

5.2.2.2 Job fit ratings

I expected that, across choice conditions, participants who chose to see the job candidate's name and photo would be biased in favor of the attractive (vs. unattractive) candidate. As predicted, the unattractive candidate ($M = 5.30$, $SD = 1.15$) was judged a worse fit for the job than the attractive candidate ($M = 5.79$, $SD = .87$) $t(236) = -3.69$, $p < .001$, $d = .45$, and the "blinded" version (credentials only; $M = 5.91$, $SD = .68$), $t(671) = -6.34$, $p < .001$, $d = .61$. The attractive candidate and the "blinded" version did not receive different fit ratings, $t(695) = -1.49$, $p = .138$, suggesting that the attractive version of the candidate was consistent with what participants expected when assessing the blind summary.

This study provided more evidence, with a managerial sample, that asking participants to reflect on their situational beliefs about the helpfulness vs. harmfulness of

potentially contaminative information—before making a blinding choice—can encourage a choice to be blind. In this case, participants reflected on whether seeing contaminative information would make their judgments susceptible to bias (i.e., vs. the accuracy-based reflection used in Study 7). Most participants who completed the reflection indicated that seeing the candidate's name and photo would bias their judgments. Importantly, they were then less likely, relative to those who did not reflect, to choose to see the candidate's name and photo information when making a subsequent, actual blinding decision.

5.2.3 Replications

I ran six other studies testing the should-first reflection intervention using a job candidate evaluation paradigm, as in Study 8. In each study, participants were less inclined to see the job candidate's name and photograph when they first indicated what their preference should be, relative to when they made a blinding choice alone (see Table 9). I also ran a seventh study, which also used a job candidate evaluation paradigm, where participants either made a blinding choice for themselves, or first reflected on what their blinding prescription would be for others and then made a choice for themselves. While results were consistent, directionally, with the results of these should-first interventions, first reflecting on what one would prescribe for others did not significantly reduce inclination to see the job candidate's name and photo personally, relative to those who made a choice for themselves first. Further details on these

reflection studies are available at

https://osf.io/pzgq2/?view_only=0e05f06ce0fd43dfa1f064a0911fcf21.

Table 9: Six Replications of the Should-First Reflection Intervention

	Actual Choice	Should Choice	Post-Should Choice
Replication 1 (N = 806)	44.1	14.4	23.3***
Replication 2 (N = 804)	36.6	13.2	19.4***
Replication 3 (N = 801)	39.3	9.7	19.2***
Replication 4 (N = 400)	49.0	12.6	22.2***
Replication 5 (N = 405)	45.9	20.5	26.5***
Replication 6 (N = 403)	40.3	13.4	20.3***

Note. Across choice conditions, percentages choosing to see the information are displayed, Post-should Choice is compared to Actual Choice in a chi-square analysis, *p < .05, **p < .01, ***p < .001

5.3 Chapter 5 Summary

The studies presented in Chapter 5 explored the last factor in my theoretical framework—evaluators' situational beliefs about the helpfulness vs. harmfulness of potentially contaminative information—as a driver of blinding preferences. In both studies, I asked some participants to articulate their situational beliefs about the helpfulness vs. harmfulness of available information before choosing whether to see it or be blind to it. In Study 7, participants indicated whether they thought a profile of a performer would increase or decrease their evaluations' accuracy. In Study 8,

participants indicated whether seeing a job candidate's name and photo would help or harm their ability to make an unbiased estimate of the candidate's fit for a job. In both cases, those who thought the information would be helpful for their judgments (in terms of accuracy or bias susceptibility) were more inclined to see it. However, most indicated a belief that the information would be harmful; hence, facilitating a reflection on situational beliefs before a blinding choice was made encouraged participants to choose to be blind. As a comparison, in Study 7, providing participants with a monetary incentive for estimate accuracy did not influence their blinding preferences.

6. General Discussion

In the present research, I performed the first tests of individual-level preferences for "blinding" one's judgment—actively avoiding incorporating potentially contaminative information into an evaluation. In Chapter 1, I discussed research on mental contamination and blinding as a strategy to reduce decision bias and proposed a theoretical framework for a study of individual-level blinding decisions, outlining enduring beliefs, situational beliefs, and motivational factors that may drive evaluators' blinding preferences. In Chapters 2-4, I presented studies that explored people's blinding preferences and underlying motivations across various evaluative contexts, and in Chapter 5, I outlined the results of studies testing multiple interventions to encourage a choice to be blind to potentially contaminative information. Across these studies, I found that people often choose *not* to be blind, and instead to see or learn potentially contaminative information about a target of evaluation. I demonstrated various ways in which a choice to view potentially contaminative information, rather than be blind to it, can foster bias in evaluations of others.

6.1 Contributions of the Present Research

Some studies have demonstrated the value of blinding—alternatively known as "exposure control" (Gilbert, 1993) or "masking" (Rissing & Castilla, 2016)—in reducing decision noisiness and bias (Axt & Lai, 2019). For instance, blinding policies helped to reduce anti-female bias in auditions for roles in major U.S. symphony orchestras (Goldin

& Rouse, 2000), and countered bias in favor of high-prestige authors at a major academic journal (Blank, 1991). Indeed, blinding is so effective a strategy at reducing decision bias that various policy proposals have called for more widespread use of blinding as a policy in domains such as the criminal justice system (e.g., Sah et al., 2016; Robertson & Kesselheim, 2016), where decision bias can have serious consequences.

Studies that have audited hiring decisions provide a stark example of the ways bias can creep into important evaluations—and negatively impact outcomes for members of disadvantaged groups—without blinding policies in place. Applicants with identical credentials are more likely to be hired, and offered higher salaries, when their name communicates that they are White or a male (Bertrand & Mullainathan, 2004; Moss-Racusin et al., 2012). Hiring decisions governed by blinding policies, such as a policy requiring evaluators to be blind to information like a candidate's name, should be less likely to demonstrate such bias.

Yet, even given studies demonstrating the value of blinding for combatting bias and the proliferation of bias in absence of blinding, and policy proposals calling for blinding to be more widespread, blinding policies are not yet mainstream. In the context of hiring, for instance, the opposite is often true, whereby managers prefer to employ unstructured and open interview procedures (Chapman & Zweig, 2005; Dana, Dawes, & Peterson, 2013; van der Zee, Bakker, & Bakker, 2002), increasing the likelihood that irrelevant and potentially biasing information about job candidates will shape their

decisions (Rivera, 2012). Therefore, it is important to understand evaluators' own preferences for blinding: Will people choose blinding for themselves?

As Wilson and Brekke (1994) have noted, however, there are few, if any, studies that have explored individual-level preferences for blinding. Hence, the present research contributes to the literature on inequality reduction mechanisms in organizations (Dobbin & Kalev, 2016), and advances existing research on blinding as a topic (Robertson & Kesselheim, 2016), by exploring people's personal preferences for blinding. The studies presented here also advance the literature on mental contamination processes by exploring an agency oriented view of mental contamination. That is, we build on recent work (Dietvorst & Simonsohn, 2018) demonstrating that mental contamination is not always something that *happens to* people as a function of unconscious responses to outside forces—sometimes, people *actively seek out* information that fosters mental contamination, even when they acknowledge it will do so.

Finally, these studies contribute to research exploring group-level stereotypes and social perception (Cuddy, Fiske, & Glick, 2007; 2008). Recent research in this domain has suggested that biases related to age have persisted in the US population during a time when other category-based biases may be diminishing in prevalence (Charlesworth & Banaji, 2019). Consistent with this perspective, the present research demonstrates a robust, ageism-consistent, evaluative bias across multiple studies, in addition to pro-attractive and anti-obese biases.

6.2 Implications for Organizations Concerned with Evaluator Bias

In Studies 2, 4, 6, and 8, I found that people often choose to see potentially contaminative information even though they acknowledge, when asked, that they should not see it if their goal is to perform an unbiased evaluation. These results underscore the importance of finding feasible interventions that can encourage a choice to be blind to potentially contaminative information. For many decisions, the most efficient means to reduce bias is to prevent access to contaminative information by adopting blinding as policy (Axt & Lai, 2019)—otherwise, people may seek it out and become biased. However, blinding is not always possible to implement at a policy level—such as when evaluations are too variable or complex in nature—and it is likely that potentially contaminative information (e.g., workplace gossip) can still pervade even "blinded" evaluations (e.g., performance evaluations).

Where access to contaminative information cannot be blocked structurally, some strategies may reduce people's temptation to seek it out. Given that the present research demonstrated both "cold"—i.e., deliberative and thoughtful—reasons why people choose to see potentially contaminative information as well as "hot"—i.e., intuitive and impulsive—reasons, it may make sense to tailor the strategy to the cause. For instance, the studies presented in Chapter 3 demonstrated that a hot motivational state of curiosity drives blinding preferences, and that blinding decision frames that reduce evaluators' curiosity can encourage choices to be blind in the moment. As an example,

when participants in Study 5 were guaranteed the ability to learn potentially contaminative information about a job candidate after evaluating their credentials in a "blind" fashion, they were relatively less inclined to see the information before evaluating credentials. This suggests that satisfying evaluators' curiosity to see potentially contaminative information at *some point* can reduce their desire to see it in the moment. It follows that evaluations that offer potentially contaminative information especially likely to pique evaluators' curiosity would benefit from curiosity reduction strategies. One such strategy may be pre-commitment—having people make self-blinding decisions in advance of a given judgment or evaluation. A manager may decide in advance not to look at the social media profiles of future job applicants, for example. Pre-commitment may lead decision-makers to view self-blinding decisions in more abstract terms because they are temporally distant (Trope & Liberman, 2010), thereby avoiding the "hot" state brought on by making a want/should decision in the moment (Bazerman, Tenbrunsel, & Wade-Benzoni, 1998).

Chapter 4 provided evidence consistent with the notion that evaluators use at least one "cold," enduring belief about information and bias to guide their blinding preferences—a belief that more information is always better to have when performing evaluations of others. Training programs or information campaigns that reduce managers' endorsement of this belief could serve to encourage choices to be blind—simple anecdotal examples could be used to invalidate a belief that *all* information is

useful. Chapter 4 also demonstrated the value of blinding decision frames that induce more careful consideration of what information to include vs. exclude from a decision than simple, binary choices about potentially contaminative information (i.e., see it or don't). That is, Study 6 situated potentially contaminative information in a list of other, diagnostic information, and asked participants to opt in to the pieces of information they would prefer to have. Using this frame that required participants to peruse (and likely compare) all available information, participants were generally very likely to opt in to diagnostic information (e.g., a job candidate's references), but unlikely to opt in to biasing information (e.g., a job candidate's name and photo).

Finally, Chapter 5 demonstrated that there is heterogeneity in evaluators' cold, situational beliefs about the helpfulness vs. harmfulness of different types of contaminative information across different types of evaluations. In both Study 7 and Study 8, participants who first reflected on their beliefs about the helpfulness vs. harmfulness of available, potentially contaminative information and then chose whether to see that information vs. be blind to it made blinding choices consistent with their beliefs. Encouragingly, in these studies, most participants did not report a belief that the potentially contaminative information would be helpful to them—would increase evaluation accuracy or decrease evaluation bias susceptibility—hence, making these beliefs salient before blinding choices were made encouraged choices to be blind. These results imply the efficacy of a straightforward intervention: People may be more likely

to choose to be blind if they are asked to articulate their situational beliefs about the likely helpfulness vs. harmfulness of contaminative information before making a blinding choice. However, this intervention is only useful to the extent most people do *not* believe that available, potentially contaminative information will be helpful to their evaluations. Some types, or frames, of information may be more likely to be encoded as harmful or contaminative (e.g., a job candidate's "race and gender information") than others (e.g., a job candidate's "professional headshot picture"). To the extent there is likely to be consensus around the contaminativity of available information, encouraging evaluators to consider their situational beliefs about the usefulness of that information before they choose whether to see it or not should promote better choices.

The above discussion of interventions that may encourage choices to blind one's judgment presupposes that blinding is "good," or something to be encouraged. However, some research suggests that blinding evaluations can sometimes *foster* decision bias. For instance, a study of racial bias in hiring before and after "ban the box" policies (BTB) were instituted—policies that make employers blind to the criminal history of job applicants—found that anti-Black discrimination in hiring *increased* after the institution of BTB (Agan & Starr, 2017). That is, employers who were made blind to the criminal history of job applicants became more likely to discriminate based on racial demographics, presumably due to over-employment of racial demographic information as a cue to criminal history. This research suggests that there can sometimes be a danger

to enacting blinding policies: In some situations, efforts to curb one type of mental contamination may unfortunately increase the likelihood of other types of contamination.

The intentionality of employers in the BTB example is unclear, in that they may not have intentionally engaged in racial discrimination after BTB was implemented because they truly believed that non-Whites are more likely to have a criminal history. Rather, it is possible that merely making evaluators aware that certain pieces of information (e.g., criminal history) have been removed from an evaluation leads them—even unconsciously—to search for other ways to glean the removed information. This concern may be more germane to blinding policies, rather than individual-level blinding decisions, but it does present a valid potential downside to blinding an evaluator's judgment. It is possible that employers would be less likely to exhibit racial discrimination in hiring if they knew they could see a job candidate's "box"—their criminal history—after first evaluating them without that information (e.g., vs. simply removing that information entirely, as is the case with BTB). Indeed, Study 5 demonstrated that participants who first performed a blind evaluation of a job candidate, then received age-communicating information, and then had the option to revise their blind evaluation made evaluations less tainted by an ageist bias than those who saw the age-communicating information before their evaluation (i.e., did not first perform a blind evaluation).

In any case, it is possible that enacting blinding policies may sometimes lead people to seek out the "blinded" information via other means. This may be less likely to occur if people *choose* to be blind, rather than have a blinding policy foisted upon them. In the event that blinding policies are to be instituted, evaluators may be less likely to seek out the "blinded" information via other informational cues (causing other types of contamination) if they know they can receive it after first performing a blind evaluation.

6.3 Limitations of the Present Research

Of course, these studies were not without limitations. One central limitation of these studies is that they explore individuals' blinding preferences in an online setting only (i.e., computer-based surveys). While this setting is aligned with some 'real-world' blinding decisions (e.g., evaluations of job candidates made online), it is not consistent with others. As one example, it is not clear whether these findings would generalize to situations in which contaminative information is offered by a peer or colleague in person (e.g., workplace gossip), where social pressures may make one more inclined to accept contaminative information or may, conversely, make one more inclined to choose not to (e.g., to appear appropriate to others). I elected to employ online experiments in order to leverage the ability to perform high-powered tests of my predictions using rigorous and controlled experimental paradigms, but future research could confirm the generalizability of the effects demonstrated here in non-online, richer settings.

Another limitation of these studies is that they do not explore blinding preferences among people who have expertise in a given evaluative domain, such as those who are expert at making hiring decisions. On the one hand, it is possible that experts in a given evaluative domain may be less likely to elect to see potentially contaminative information in advance of an evaluation, because their expertise grants them a better understanding of their susceptibility to decision bias and/or the scope of what information is useful vs. non-useful. On the other hand, it is possible that expertise in a given evaluative domain may lead evaluators in that domain to be overconfident about their degree of susceptibility to bias, perhaps encouraging a choice to see potentially contaminative information relative to a novice in that domain. While I did try to contend with this issue, to an extent, by sampling managerial-level participants who broadly had hiring experience for some of my hiring-related experiments, addressing the potential effects of domain expertise still suggests a fertile avenue for future research.

These studies were also limited by their focus on individual-level blinding preferences, rather than preferences for the institution of blinding policies in a workplace. That is, the effects demonstrated in these studies were specific to evaluators' inclination to blind their own judgment, and may not necessarily generalize to other questions, such as managers' inclination to institute blinding policies for the employees they supervise or CEOs' inclination to enforce blinding policies as a firm-level change.

Indeed, given the large percentage of participants in these studies who elected to see, rather than be blind to, potentially contaminative information when it was up to them, it may be useful to understand the preferences of people in positions of authority for the institution of broad-scale blinding policies. Some research suggests that, whereas managers may not be very inclined to blind their own judgment, they should be more inclined to blind others'—people generally believe that others are more susceptible to decision bias then they are personally, for instance (Pronin et al., 2002). Indeed, in some of the replication data discussed after Study 8, I found that participants were more inclined to prescribe blinding for other people than they were for themselves. The results of Studies 7 and 8 also suggest that having managers complete a reflection on their situational beliefs about the value of blinding could be a way to encourage them to enact a blinding policy affecting themselves (and/or others) for future evaluations. A reflection-based task—even just an elicitation of "should" beliefs—that leads managers to adopt a "colder" decision-making frame in the moment could be leveraged toward better policy setting for the future. These dynamics should be explored more thoroughly by future research. Given the value of blinding as a solution to bias, it is important to understand not only evaluators' own preferences for blinding but their support for blinding policies more broadly.

There were also some limitations to the theoretical framework proposed in Chapter 1. Chiefly, the framework I proposed was not "complete," in that other factors,

which I did not theorize or test, could also be reasonably expected to influence individual-level blinding preferences. It is best to think of the theoretical framework proposed in Chapter 1 as a partial framework, or starting point, rather than a complete model of a blinding decision.

For instance, one factor that might be expected to influence some blinding decisions is variance in evaluators' preference for the maintenance of social hierarchy (Ho et al., 2015). An evaluator who prefers to maintain social hierarchy and inequality might choose to see a job candidate's name and photo, rather than be blind to that information, in the hopes of using it to differentiate between job candidates of social groups they believe "should" be favored, and those of social groups who should not. Conversely, when available, potentially contaminative information (or the evaluation at hand) is not relevant to the maintenance of social hierarchy (e.g., background performance information in a performance evaluation), this type of enduring motivation—social hierarchy maintenance—might not be expected to drive blinding preferences. These nuances are not captured by the theoretical framework proposed in Chapter 1, and it would be interesting for future work to explore them.

Another motivational state that is not captured by the theoretical framework I proposed but is likely to influence blinding decisions is psychological reactance (Brehm, 1966). Psychological reactance might be expected to operate situationally in relation to blinding preferences, such that, from evaluation situation to situation, evaluators who

feel that their freedom to see potentially contaminative information is threatened might be more likely to react against that constraint and choose to see (vs. be blind to) that information (Brehm, 1966). Some research also suggests that reactance could operate as an enduring motivation across different evaluations (Hong & Faedda, 1996), such that evaluators generally "high" on reactance might be generally more likely to choose not to be blind as a way to assert their freedom to see potentially contaminative information. Future research could also explore these dynamics. I did not intend for the theoretical framework presented in Chapter 1 to represent a complete and full accounting of the factors that drive individual-level blinding preferences. Instead, it was my goal to offer a starting point for a first test of blinding preferences.

7. Conclusions

In many organizational evaluations, such as performance evaluations, internal hiring and promotion decisions, and interviews for new hires, decision bias can have meaningful consequences for both individuals (e.g., members of stereotyped groups) and organizations as a whole. While some studies have explored the value of blinding in reducing decision bias across different evaluative domains, there is still much that is unknown about blinding as a topic. Perhaps most importantly, no existing research has explored evaluators' own preferences for blinding their evaluations. In this dissertation, I proposed a theoretical framework for my study of individual-level blinding decisions, outlining various beliefs and motivations that may lead evaluators to choose to see

potentially contaminative information in an impending evaluation. Eight studies (a) explored individuals' preferences for blinding and outlined consequences for bias, (b) tested the mechanisms driving blinding decisions proposed by my theoretical framework, and (c) explored the efficacy of multiple interventions to encourage a choice to blind one's judgment.

These studies demonstrated that, even though most people agree that one shouldn't view potentially contaminative information before an evaluation, upwards of 50% of people still elect to see such information when it is readily available. This likelihood may be even more pronounced in the "real world," where tempting, contaminative information is broadly available. Thankfully, these studies also provided some insights regarding how to encourage blinding: If people can be nudged to think a bit more deeply than usual about their preferences for information and what those preferences might mean for evaluations of others, they often arrive at, and act in a manner consistent with, good intentions.

References

Agan, A., & Starr, S. (2018). Ban the box, criminal records, and racial discrimination: A field experiment. *The Quarterly Journal of Economics, 133*(1), 191-235.

Armor, D. A. (1999). The illusion of objectivity: Bias in the belief in freedom from bias. *Doctoral dissertation, University of California Los Angeles*.

Axt, J. R., & Lai, C. K. (2019). Reducing discrimination: A bias versus noise perspective. *Journal of Personality and Social Psychology, 117*(1), 26-49.

Badia, P., McBane, B., & Suter, S. (1966). Preference behavior in an immediate versus variably delayed shock situation with and without a warning signal. *Journal of Experimental Psychology, 72*(6), 847-852.

Bazerman, M. H., Tenbrunsel, A. E., & Wade-Benzoni, K. (1998). Negotiating with yourself and losing: Making decisions with competing internal preferences. *Academy of Management Review, 23*(2), 225-241.

Berinsky, A. J., Huber, G. A., & Lenz, G. S. (2012). Evaluating online labor markets for experimental research: Amazon.com's Mechanical Turk. *Political Analysis, 20*(3), 351-368.

Berlyne, D. E. (1950). Novelty and curiosity as determinants of exploratory behavior. *British Journal of Psychology, 41,* 68-80.

Berlyne, D. E. (1966). Curiosity and exploration. *Science, 153*(3731), 25-33.

Bertrand, M., & Mullainathan, S. (2004). Are Emily and Greg more employable than Lakisha and Jamal? A field experiment on labor market discrimination. *The American Economic Review, 94*(4), 991-1013.

Blank, R. M. (1991). The effects of double-blind versus single-blind reviewing: Experimental evidence from the American Economic Review. *The American Economic Review*, 1041-1067.

Brehm, J. W. (1966). *A theory of psychological reactance*. New York: Academic Press.

Castilla, E. J., & Benard, S. (2010). The paradox of meritocracy in organizations. *Administrative Science Quarterly, 55*(4), 543-676.

Chaiken, S. (1980). Heuristic versus systematic information processing and the use of source versus message cues in persuasion. *Journal of Personality and Social Psychology, 39*(5), 752-766.

Chaiken, S., & Stangor, C. (1987). Attitudes and attitude change. *Annual Review of Psychology, 38*(1), 575-630.

Chapman, D. S., & Zweig, D. I. (2005). Developing a nomological network for interview structure: Antecedents and consequences of the structured selection interview. *Personnel Psychology, 58*(3), 673-702.

Charlesworth, T. E., & Banaji, M. R. (2019). Patterns of implicit and explicit attitudes: I. Long-term change and stability from 2007 to 2016. *Psychological Science, 30*(2), 174-192.

Chen, S., Shechter, D., & Chaiken, S. (1996). Getting at the truth or getting along: Accuracy-versus impression-motivated heuristic and systematic processing. *Journal of Personality and Social Psychology, 71*(2), 262-275.

Cuddy, A. J. C., Fiske, S. T., & Glick, P. (2007). The BIAS map: Behaviors from intergroup affect and stereotypes. *Journal of Personality and Social Psychology, 92*(4), 631-648.

Cuddy, A. J., Fiske, S. T., & Glick, P. (2008). Warmth and competence as universal dimensions of social perception: The stereotype content model and the BIAS map. *Advances in Experimental Social Psychology, 40*, 61-149.

Dana, J., Dawes, R., & Peterson, N. (2013). Belief in the unstructured interview: The persistence of an illusion. *Judgment and Decision Making, 8*(5), 512-520.

Dietvorst, B. J., & Simonsohn, U. (2018). Intentionally "biased": People purposely use to-be-ignored information, but can be persuaded not to. *Journal of Experimental Psychology: General, 148*(7), 1228-1238.

Dobbin, F., & Kalev, A. (2016). Why diversity programs fail and what works better. *Harvard Business Review, 94*(7-8), 52-60.

Dobbin, F., Schrage, D., & Kalev, A. (2015). Rage against the iron cage: The varied effects of bureaucratic personnel reforms on diversity. *American Sociological Review, 80*(5), 1014-1044.

Dunning, D., Meyerowitz, J. A., & Holzberg, A. D. (1989). Ambiguity and self-evaluation: The role of idiosyncratic trait definitions in self-serving assessments of ability. *Journal of Personality and Social Psychology, 57*(6), 1082-1090.

Eliaz, K., & Schotter, A. (2007). Experimental testing of intrinsic preferences for noninstrumental information. *American Economic Review, 97*(2), 166-169.

Evans, J. S. B., & Stanovich, K. E. (2013). Dual-process theories of higher cognition: Advancing the debate. *Perspectives on Psychological Science, 8*(3), 223-241.

Festinger, L. (1962). *A theory of cognitive dissonance* (Vol. 2). Stanford university press.

Gervais, W. M., Jewell, J. A., Najle, M. B., & Ng, B. K. (2015). A powerful nudge? Presenting calculable consequences of underpowered research shifts incentives toward adequately powered designs. *Social Psychological and Personality Science, 6*(7), 847-854.

Gilbert, D. T. (1993). The assent of man: Mental representation and the control of belief. In D. M. Wegner, & J. W. Pennebaker (Eds.), *The handbook of mental control* (pp. 57-87). Englewood Cliffs, NJ: Prentice Hall.

Gilhooly, K. J., Fioratou, E., Anthony, S. H., & Wynn, V. (2007). Divergent thinking: Strategies and executive involvement in generating novel uses for familiar objects. *British Journal of Psychology, 98*(4), 611-625.

Goldin, C., & Rouse, C. (2000). Orchestrating impartiality: The impact of "blind" auditions on female musicians. *The American Economic Review, 90*(4), 715-741.

Grant, S., Kajii, A., & Polak, B. (1998). Intrinsic preference for information. *Journal of Economic Theory, 83*(2), 233-259.

Hayes, A. F. (2012). PROCESS: A versatile computational tool for observed variable mediation, moderation, and conditional process modeling. Retrieved from http://is.muni.cz/el/1423/podzim2014/PSY704/50497615/hayes_2012_navod_process.pdf

Hightow, L. B., Miller, W. C., Leone, P. A., Wohl, D., Smurzynski, M., & Kaplan, A. H. (2003). Failure to return for HIV posttest counseling in an STD clinic population. *AIDS Education and Prevention, 15*, 282–290.

Ho, A. K., Sidanius, J., Kteily, N., Sheehy-Skeffington, J., Pratto, F., Henkel, K. E., ... & Stewart, A. L. (2015). The nature of social dominance orientation: Theorizing and measuring preferences for intergroup inequality using the new SDO_7 scale. *Journal of Personality and Social Psychology, 109*(6), 1003-1028.

Hong, S. M., & Faedda, S. (1996). Refinement of the Hong psychological reactance scale. *Educational and Psychological Measurement, 56*(1), 173-182.

Hsee, C. K., & Ruan, B. (2016). The Pandora effect: The power and peril of curiosity. *Psychological Science, 27*(5), 659-666.

Jones, E. E. (1986). Interpreting interpersonal behavior: The effects of expectancies. *Science, 234*(4772), 41-46.

Kalev, A., Dobbin, F., & Kelly, E. (2006). Best practices or best guesses? Assessing the efficacy of corporate affirmative action and diversity policies. *American Sociological Review, 71*(4), 589-617.

Kahneman, D. (2011). *Thinking, fast and slow.* Macmillan.

Keogh, L. A., Southey, M. C., Maskiell, J., Young, M. A., Gaff, C. L., Kirk, J., ... & Hopper, J. L. (2004). Uptake of offer to receive genetic information about BRCA1 and BRCA2 mutations in an Australian population-based study. *Cancer Epidemiology and Prevention Biomarkers, 13*(12), 2258-2263.

Klein, N., & O'Brien, E. (2018). People use less information than they think to make up their minds. *Proceedings of the National Academy of Sciences, 115*(52), 13222-13227.

Langlois, J. H., Kalakanis, L., Rubenstein, A. J., Larson, A., Hallam, M., & Smoot, M. (2000). Maxims or myths of beauty? A meta-analytic and theoretical review. *Psychological Bulletin, 126*(3), 390-423.

Lanzetta, J. T., & Driscoll, J. M. (1966). Preference for information about an uncertain but unavoidable outcome. *Journal of Personality and Social Psychology, 3*(1), 96-102.

Larrick, R. P., Burson, K. A., & Soll, J. B. (2007). Social comparison and confidence: When thinking you're better than average predicts overconfidence (and when it does not). *Organizational Behavior and Human Decision Processes, 102*, 76-94

Lauver, K. J., & Kristof-Brown, A. (2001). Distinguishing between employees' perceptions of person--job and person--organization fit. *Journal of Vocational Behavior, 59*(3), 454-470.

Levine, E. E., & Schweitzer, M. E. (2015). The affective and interpersonal consequences of obesity. *Organizational Behavior and Human Decision Processes, 127*, 66-84.

Litman, J. (2005). Curiosity and the pleasures of learning: Wanting and liking new information. *Cognition & Emotion, 19*, 793-814.

Litman, J. A., & Jimerson, T. L. (2004). The measurement of curiosity as a feeling-of-deprivation. *Journal of Personality Assessment, 82,* 147-157.

Lockard, J. S. (1963). Choice of a warning signal or no warning signal in an unavoidable shock situation. *Journal of Comparative and Physiological Psychology, 56*(3), 526-530.

Loewenstein, G. (1994). The psychology of curiosity: A review and reinterpretation. *Psychological Bulletin, 116*(1), 75-98.

Ma, D. S., Correll, J., & Wittenbrink, B. (2015). The Chicago face database: A free stimulus set of faces and norming data. *Behavior Research Methods, 47*(4), 1122-1135.

Moore, D. A., & Healy, P. J. (2008). The trouble with overconfidence. *Psychological Review, 115*(2), 502-517.

Moss-Racusin, C. A., Dovidio, J. F., Brescoll, V. L., Graham, M. J., & Handelsman, J. (2012). Science faculty's subtle gender biases favor male students. *Proceedings of the National Academy of Sciences, 109*(41), 16474-16479.

North, M. S., & Fiske, S. T. (2012). An inconvenienced youth? Ageism and its potential intergenerational roots. *Psychological Bulletin, 138*(5), 982-997.

Ottati, V. C., & Isbell, L. M. (1996). Effects on mood during exposure to target information on subsequently reported judgments: An on-line model of misattribution and correction. *Journal of Personality and Social Psychology, 71*(1), 39-53.

Patsy, B. M., & Prentice, R. L. (2010). Minimizing bias in randomized trials: The importance of blinding. *JAMA, 304*(7), 793-794.

Peer, E., Brandimarte, L., Samat, S., & Acquisti, A. (2017). Beyond the Turk: Alternative platforms for crowdsourcing behavioral research. *Journal of Experimental Social Psychology, 70*, 153-163.

Petty, R. E., Cacioppo, J. T., & Kao, C. F. (1984). The efficient assessment of need for cognition. *Journal of Personality Assessment, 48*(3), 306-307.

Petty, R. E., & Wegener, D. T. (1993). Flexible correction processes in social judgment: Correcting for context-induced contrast. *Journal of Experimental Social Psychology, 29*(2), 137-165.

Preacher, K. J., Rucker, D. D., & Hayes, A. F. (2007). Addressing moderated mediation hypotheses: Theory, methods, and prescriptions. *Multivariate Behavioral Research, 42*(1), 185-227.

Pronin, E., Gilovich, T., & Ross, L. (2004). Objectivity in the eye of the beholder: divergent perceptions of bias in self versus others. *Psychological Review, 111*(3), 781-799.

Pronin, E., Lin, D. Y., & Ross, L. (2002). The bias blind spot: Perceptions of bias in self versus others. *Personality and Social Psychology Bulletin, 28*(3), 369-381.

Rissing, B. A., & Castilla, E. J. (2014). House of green cards: Statistical or preference-based inequality in the employment of foreign nationals. *American Sociological Review, 79*(6), 1226-1255.

Rivera, L. A. (2012). Hiring as cultural matching: The case of elite professional service firms. *American Sociological Review, 77*(6), 999-1022.

Robertson, C. T., & Kesselheim, A. S. (Eds.). (2016). *Blinding as a solution to bias: Strengthening biomedical science, forensic science, and law*. Academic Press.

Rogers, T., & Bazerman, M. H. (2008). Future lock-in: Future implementation increases selection of 'should' choices. *Organizational Behavior and Human Decision Processes, 106*(1), 1-20.

Sah, S., Tannenbaum, D., Cleary, H., Feldman, Y., Glaser, J., Lerman, A., ... & Spohn, C. (2016). Combating biased decision-making & promoting justice & equal treatment. *Behavioral Science & Policy, 2*(2), 78-87.

Soll, J. B., & Klayman, J. (2004). Overconfidence in interval estimates. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 30*(2), 299-314.

Soll, J. B., Milkman, K. L., & Payne, J. W. (2015). A user's guide to debiasing. *The Wiley Blackwell Handbook of Judgment and Decision Making, 2*, 924-951.

Stanovich, K. E., & West, R. F. (2000). Individual differences in reasoning: Implications for the rationality debate?. *Behavioral and Brain Sciences, 23*(5), 645-665.

Sweeny, K., Melnyk, D., Miller, W., & Shepperd, J. A. (2010). Information avoidance: Who, what, when, and why. *Review of General Psychology, 14*(4), 340-353.

Tolbert, P. S., & Castilla, E. J. (2017). Editorial essay: Introduction to a special issue on inequality in the workplace ("What Works?"). *ILR Review, 70*(1), 3-15.

Trope, Y., & Liberman, N. (2010). Construal-level theory of psychological distance. *Psychological Review, 117*(2), 440-463.

Uhlmann, E. L., & Cohen, G. L. (2007). "I think it, therefore it's true": Effects of self-perceived objectivity on hiring discrimination. *Organizational Behavior and Human Decision Processes, 104*(2), 207-223.

Van Boven, L., White, K., Kamada, A., & Gilovich, T. (2003). Intuitions about situational correction in self and others. *Journal of Personality and Social Psychology, 85*(2), 249-258.

Van der Zee, K. I., Bakker, A. B., & Bakker, P. (2002). Why are structured interviews so rarely used in personnel selection?. *Journal of Applied Psychology, 87*(1), 176-184.

Wegener, D. T., & Petty, R. E. (1995). Flexible correction processes in social judgment: The role of naive theories in corrections for perceived bias. *Journal of Personality and Social Psychology, 68*(1), 36-51.

Wilson, T. D., & Brekke, N. (1994). Mental contamination and mental correction: Unwanted influences on judgments and evaluations. *Psychological Bulletin, 116*(1), 117-142.

Wilson, T. D., Houston, C. E., Etling, K. M., & Brekke, N. (1996). A new look at anchoring effects: Basic anchoring and its antecedents. *Journal of Experimental Psychology: General, 125*(4), 387-402.

Woolley, K., & Risen, J. L. (2018). Closing your eyes to follow your heart: Avoiding information to protect a strong intuitive preference. *Journal of Personality and Social Psychology, 114*(2), 230-245.

Biography

Sean Fath received a B.A. in Psychology from DePaul University in 2011. Sean's work has been published in *Organizational Behavior and Human Decision Processes,* the *Journal of Experimental Social Psychology,* and *Anxiety, Stress, and Coping*.