THINKING ABOUT SELF AND OTHERS IN THE CONTEXT OF KNOWLEDGE AND EXPERTISE

A Dissertation

Presented to the Faculty of the Graduate School

of Cornell University

In Partial Fulfillment of the Requirements for the Degree of

Doctor of Philosophy

by

Stav Stav Atir

August 2018

THINKING ABOUT SELF AND OTHERS IN THE CONTEXT OF KNOWLEDGE AND EXPERTISE

Stav Stav Atir, Ph.D.

Cornell University 2018

The work presented here encompasses two lines of research broadly concerned with the way people think of themselves and of others in the context of knowledge and expertise. The first line of research focuses on self-judgments of knowledge. People are often required to evaluate their own knowledge in the service of various decisions; how is this evaluation made, and when does it go awry? One way to study overestimation of knowledge is by asking people to assess their knowledge of terms, places, people, etc. that were invented for the sake of research and do not exist. This approach helps reveal when and why people mistakenly believe themselves to have knowledge that they could not possibly have, a phenomenon called overclaiming. Across two papers, I demonstrate that people are more likely to overclaim knowledge about nonexistent terms when they perceive themselves to be overall knowledgeable about the topic from which the terms are drawn. However, self-perceptions of overall knowledge can be dissociated from true knowledge. I also find that the more people are genuinely knowledgeable about a topic, they less likely they are to overclaim – i.e., to mistakenly believe they have knowledge of nonexistent terms related to their topic of expertise.

The second line of research investigates the role of gender in the way people think and speak about experts, i.e., about professionals. In many professions, people commonly refer to others by their surname only – e.g., *Smith* – when discussing them or their work. Across four

studies that combine archival and experimental data in multiple domains, I demonstrate that gender influences use of this type of reference; specifically, people are more likely to refer to male vs. female professionals by surname only. In five additional studies I find evidence that use of surname affects career-related judgments: researchers who are referred to by their surname only (vs. their full name) are judged to be more famous and eminent, and consequently are seen as higher status and more deserving of a career award.

BIOGRAPHICAL SKETCH

Stav Atir grew up in sunny Tel Aviv. After high school she began pursuing her dream of becoming a doctor by attending the medical school at Tel Aviv University. Midway through the first year, lured by the charming idea of "college", she decided to take a quick detour on the road to medical doctordom and get a bachelor's degree at Yale University in far away New England. One semester later she enrolled in Introduction to Psychology and fell head over heels for the subject. It had taken a few more semesters of her opening up the course catalogue and immediately flipping to the psychology section, but eventually it became clear that the detour had turned into a trip towards an entirely new destination. After graduating, Stav worked as a research specialist for two years at the University of Pennsylvania, and then entered into a psychology PhD program at Cornell University. Her next adventure will be at the University of Chicago Booth School of Business, where she will be a postdoctoral researcher.

ACKNOWLEDGMENTS

I have a superb group of people to thank for making the writing of this dissertation possible. First, Melissa Ferguson. I feel fortunate to have had an advisor who possesses the rare combination of formidable intellect and exceptional kindness. It's been a pleasure developing ideas with you, discussing research (and life) over *Gimme!* coffee at Gates, drinking celebratory champagne at 11 in the morning, and even trying to nap on uncomfortable chairs at the Rochester airport on the way to a conference. Thank you for not only being there but truly sharing with me the highs and lows of research – the dread of waiting for a result on which much hinges, the frustration of mixed results, the joy of new discovery, and the thrill of learning that a manuscript has found a good home. I am more grateful than I can ever express for your guidance, your encouragement, your support, your generosity, and your friendship.

Thank you to the other members of my committee and my other collaborators, Dave Dunning, Thomas Gilovich, Khena Swallow, and Emily Rosenzweig. I am thankful to have gotten the opportunity to learn from each of you. Getting to work with you, to discuss ideas with you, and to see your individual styles and approaches to research, each clever and thoughtful, was inspiring and has made me a better, more well-rounded researcher.

I would also like to thank my family for their unwavering support. They were kind not to ask too often when I thought I might be graduating. Thank you for always encouraging me and standing by me no matter what I chose to do – or where I chose to do it, even if it meant being far away from you. Thank you to my friends, and especially my best friend, who has been there for me through everything since the fourth grade. I appreciate each of you and feel indescribably lucky to have you.

Finally, to my partner, Ido Potter, who has been with me every step of the way and who has shown me more love and understanding than anyone has any right to expect. Thank you for supporting me in more ways than I can count, for giving me perspective, and for never showing any doubt that I could do whatever I set out to do. Doing a PhD is a rollercoaster of excitement, frustration, anxiety, and joy, and that's just in the course of one statistical analysis. Thank you for making the bad parts far more bearable and the good parts even better. Ani ohevet otcha.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

CHAPTER I

INTRODUCTION

The work presented here is broadly concerned with how people think of themselves and of others with regards to knowledge and expertise. The first section focuses on how people think of their own knowledge; specifically, how they assess it and when and why they overestimate it. The second section explores how people think and speak about others who possess a high level of expertise; specifically, how people refer to and make judgments about male and female professionals.

Self-judgments of knowledge

Consider the process of studying for an exam. Periodically, the learner must query herself: Do I understand this material? Do all of these terms make sense? Am I ready to move on to the next section? Put simply, she has to assess her own knowledge. The outcome will determine whether she will spend additional time reviewing the same material, will decide her studying time is better spent covering new material, or will even set aside her books and notes and proclaim herself sufficiently prepared. If she misestimates her knowledge, she either needlessly wastes her time going over material she already knows well, or she moves on prematurely and may discover, when asked to give evidence of her knowledge during the exam, that her mistake has cost her precious points.

The question of how much one knows isn't relevant only when preparing for an exam. We are required to make similar assessments of our knowledge in many different contexts, some of which have high stakes. Should I make this investment or business decision or do I need to learn more about it first? Does my knowledge qualify me to take on this project or to do well at this job? Am I ready for this meeting presentation? Am I sufficiently informed to vote? Should I

agree to teach or tutor on this subject? Should I dispense advice on this topic? The consequences can be dire if the wrong assessment is made. It is therefore crucial to understand how we make knowledge assessments, and when the process leads us to the wrong conclusion.

It may seem implausible that the process should lead to the wrong conclusion all that often. Surely we know our own minds well enough to tell what we know and what we don't know – don't we? More broadly, we tend to feel that if there is one thing on which everyone is an expert, it is their own self (Pronin, 2009). And in a way, this is correct. A person generally has much more information about herself than any observer could because she has access to her own feelings, preferences, thoughts, memories, knowledge – in short, to her own mind. Yet the notion that we make meta-cognitive judgments by simply consulting some orderly and complete store of mental content is not quite accurate.

Previous research has demonstrated that all sorts of meta-cognitive judgments can be pushed around in ways that suggest that people do not truly "know themselves" in this direct manner. For example, through the use of bogus heartbeat feedback, people can be fooled into thinking they feel certain emotions, like attraction (Valins, 1966) and distress (Gu, Zhong, & Page-Gould, 2013). They can also be made to think they prefer one option when they in fact expressed a preference for the other option (Johansson, Hall, Sikstrom, & Olsson, 2005), and they can become convinced that they have memories of things that they have never seen (e.g., a word that did not appear on a list they had memorized; Roediger & McDermott, 1995) or that never happened [e.g., a false childhood memory of getting lost in the mall (Loftus & Pickrell, 1995) or of taking a trip in a hot air balloon (Wade, Garry, Read, & Lindsay, 2002)]. This and much additional research teaches us that in many instances, people do not rely on direct introspective access to assess what they feel, prefer, and remember.

What about knowledge? Again, our intuition may insist that reporting on one's knowledge is easy – that it is a simple matter of directly consulting one's store of mental content. We may concede that *broad* judgments of knowledge can be tricky; it can be difficult to know how to evaluate one's general knowledge or one's overall expertise within a domain. But judgments about knowing a specific piece of information seem like they should be straightforward. However, the evidence suggests that this may not always be the case. When asked to make judgments about their own knowledge, people often rely on a *feeling of knowing* (Koriat, 1993). The tip-of-the-tongue phenomenon provides a nice demonstration of the dissociation between a feeling of knowing and actually knowing; this phenomenon refers to the sensation that one knows a piece of information yet cannot report the information itself (Brown, 1991). Some researchers have explicitly argued or implicitly assumed that this feeling of knowing is based purely on access to real knowledge (the *trace access* account; Hart, 1967; Yaniv & Meyer, 1987), which predicts that a feeling of knowing should be reliable and accurate. However, the feeling of knowing sometimes fails to align with real knowledge (Koriat & Leiblich, 1977; Koriat, 1995). It turns out that even when people feel strongly that they know a piece of information – e.g., a fact or a name – sometimes they do not recognize that piece of information when it is presented to them. In other words, sometimes people feel that they know something even though they do not.

If feeling of knowing does not reflect direct access to knowledge, what does it reflect? In the same way that people sometimes infer their emotions based on heuristics - e.g., that a quick heart rate generally signals a strong, high-arousal emotion – perhaps feeling of knowing is likewise informed by heuristics. Previous work has tied judgments of knowledge to various aspects of the learning experience. For example, people are more likely to judge that they know

information the less time it took them to initially learn it (Koriat, 2008; Koriat, Ackerman, Lockl, & Schneider, 2009a), and the less effort they believe they expended to learn it (Koriat, Ackerman, Lockl, & Schneider, 2009b). Presumably, these are cues that the information is easy to learn and therefore that it has likely been successfully learned. The fluency with which information is encoded or retrieved during learning similarly predicts these judgments (Benjamin, Bjork, & Schwartz, 1998; Matvey, Dunlosky & Guttentag, 2001).

In the work presented here, I explore other heuristics that shape meta-knowledge judgments. First, I posit and find that judgments about having specific knowledge are informed not only by features of the stimuli at hand or the learning process, but also by broad self-perceptions. Specifically, using correlational and experimental studies I demonstrate that people who see themselves as generally knowledgeable about a topic are more likely to believe they have specific knowledge related to that topic, even when they do not. Second, in additional work I find that self-perceptions of knowledge are related to overestimation of knowledge not because self-perceptions of knowledge are particularly accurate. Rather, genuine knowledge and self-perceived knowledge have diverging effects on meta-knowledge judgments. The description of these two lines of work makes up the first two body chapters of this dissertation.

Chapter II is a reproduction of work I published with co-authors Emily Rosenzweig and David Dunning in 2015 in the journal *Psychological Science*, titled "When Knowledge Knows No Bounds: Self-Perceived Expertise Predicts Claims of Impossible Knowledge." The paper examines how people judge their own knowledge and why they sometimes overestimate it. This is accomplished through use of the overclaiming paradigm, in which people are asked to assess their knowledge of real and bogus terms, places, people, etc. that are all related to a particular domain. For example, in the political domain, people may be asked about their knowledge of real

political figures like *John Kerry*, as well as fictitious ones like *Thomas Langer.* If people indicate they have knowledge of the bogus items – called *foils* – they are, in effect, overclaiming, i.e., they mistakenly believe they have knowledge that they cannot have.

The paper tests what underlies such impossible claims of knowledge. The results of five studies consistently find that the way people view their own knowledge predicts their overclaiming; the more people perceive themselves to be knowledgeable, the more likely they are to overclaim knowledge within their domain of self-perceived expertise. In Studies 1a and 1b, the more favorable participants' perceptions of their knowledge in personal finance, the more likely they were to believe they were familiar with real financial terms like *home equity* and *revolving credit*, but also to mistakenly believe they were familiar with fictitious financial terms like *fixed rate deduction.* The same pattern emerged in other domains, including biology, history, literature, and philosophy. Study 2 demonstrated that people's mistakes show specificity; the more people perceived themselves as knowledgeable within a domain, the more they overclaimed knowledge related to that particular domain, even controlling for their perceptions of knowledge in other domains. People who saw themselves as experts in biology, for example, were more likely to overclaim knowledge of bogus biology terms but not of bogus history terms.

Even when people received an explicit warning that some of the terms in the task will be fictitious (Study 3), they overclaimed only slightly less, suggesting that their judgments were honest. If people were merely pretending knowledge of the foils for impression management motives, presumably they would be hyper-careful in their judgments to avoid the embarrassment of asserting knowledge of made-up terms. Importantly, even when a warning was included, self-perceptions of knowledge still predicted overclaiming mistakes. When people's perceptions of knowledge were manipulated (Study 4), their mistaken knowledge assessments changed as well:

those who were experimentally induced to feel more knowledgeable – even though no knowledge was actually gained – became more likely to overclaim, i.e., to mistakenly perceive familiarity with bogus terms. Together, the results suggest that self-judgments of knowledge are indeed not necessarily based on direct access to one's knowledge. Rather, knowledge judgments are influenced by one's overall self-perception of knowledge within the relevant domain.

In Chapter III I present a reproduction of a manuscript with the same co-authors – Emily Rosenzweig and David Dunning – that will be submitted for publication at *Psychological Science* in 2018. The paper builds on the findings reported in Chapter II linking overclaiming to self-perceptions of knowledge by noting that the way people perceive their knowledge is not a perfect index of how knowledgeable they genuinely are, and then unpacking the effects of genuine knowledge and self-perceived knowledge on overclaiming. The first study documented a negative association between genuine knowledge and overclaiming in an ecologically valid context; three groups of people with varying levels of expertise in the medical domain – medical doctors, medical students, and premedical students – displayed different levels of overclaiming. Specifically, those who had more genuine knowledge in medicine – doctors and medical students – were less likely to mistakenly believe they had knowledge of bogus medical terms compared with premedical students.

Across three additional studies, the paper disentangles the contributions of genuine knowledge and self-perceived knowledge and concludes that each one predicts overclaiming independently and in opposite directions. Study 2 shows that participants who viewed themselves as knowledgeable in science were more likely to overclaim knowledge of bogus science terms, whereas those who genuinely knew more about science – as indexed by a science quiz – were less likely to overclaim.

The results of Study 3 suggest that the protective effect of genuine knowledge is related to the associations people have with the real items and the foils. Specifically, the effect of genuine knowledge on overclaiming was mediated by the gap between the number of associations that people generated about the real items and the foils, such that those with greater scientific expertise had a larger gap. Finally, Study 4 provides evidence that more genuinely knowledgeable people make knowledge assessments in a more automatic way; they "just know" how much they know – or don't know – about the object of assessment. In contrast, people who only viewed themselves as knowledgeable tended to report using a more deliberative thinking process to judge their knowledge about the terms. Perhaps deliberating can lead to thinking oneself into a false sense familiarity with the term. The results of the four studies combined suggest that self-perceived knowledge and genuine knowledge are both associated with self-judgments of knowledge through distinct mechanisms and in opposing directions.

The overarching aim of the work presented here is to understand how people think about both themselves and others in the context of knowledge and expertise. My second research project focuses on the way people think, form judgments, and speak about people who are experts: professionals.

Speaking About Others in Professional Contexts

The way we think about others can have profound implications for how we treat, judge, and evaluate them. And the way we think about others is tightly bound up with the way we speak about them. Our thoughts shape not just *what* we say, but *how* we say it (Hogg & Reid, 2006; Maass, Salvi, Arcuri, & Semin, 1989) and, more surprisingly, the converse is true as well: the way things are said, phrased, or named shapes what we think about them (Boroditsky, Schmidt, & Phillips, 2003; Lofthus & Palmer, 1974). An immense literature supports the notion that the

way we think and speak about others is influenced not only by their individual attributes but also by their group memberships, be they relevant (e.g., Nelson, Biernat, & Manis, 1990) or not (e.g. Gorham, 2006; Duncan, 1976; Swim & Sanna, 1996; Taylor & Jaggi, 1974; Hodson, Dovidio, & Gaertner, 2002). Such group-based biases may help explain otherwise puzzling inequalities between groups of people. I will focus here on gender.

Women make up roughly half of the population, and yet remain astonishingly underrepresented in many professions, especially prestigious and lucrative ones, including politics, business, and science, technology, engineering, and math (STEM; National Conference of State Legislatures, 2018; Catalyst, 2018, National Science Board, 2016). For example, women make up only 20% of Google's tech workforce (Google, 2018), and the numbers are similar in other major tech companies like Apple and Facebook (Apple, 2017; Facebook, 2018). The number of female CEOs among Fortune 500 companies has been inching upwards very slowly over the years, but has actually declined since 2017 and stands at an unimpressive 5% (Fortune, 108). Across the board, women continue to earn less than men (Blau & Kahn, 2017) and hold fewer positions of power and influence (The American Association of University Women, 2016).

Why is this still the case? Though there is evidence for some psychological differences between men and women (e.g., Hegarty, & Sims, 1994; cf Hoffman, Gneezy, & List, 2011), the genders appear to be fairly equally matched in cognitive abilities overall; women perform similarly or better than men in high school and college (Voyer & Voyer, 2014), and score similarly to men on general intelligence tests (Colom, Garcia, Juan-Espinosa, & Abad, 2002). The professional representation of women varies widely from place to place; for example, though women account for only 29% of scientific researchers in a country on average, they make up around half of researchers in certain countries (e.g. 51% in Latvia, 56% in Thailand), and

even more than half in others – for example, 63% in Bolivia (United Nations Educational, Scientific, and Cultural Organization, 2018). Thus, innate differences in aptitude seem a poor explanation.

Though there are doubtless multiple causes that interact to explain this unequal state, it is likely that gender biases contribute to it. Blatant prejudice is becoming less and less accepted (Swim, Aikin, Hall, & Hunter, 1995), but men and women are still treated differently in the workplace. Subtle biases may persist, possibly without people's awareness and intent, and may thus be challenging to identify and eradicate, though they could be just as detrimental to women's professional advancement. One such bias may lie in the way people speak about male and female professionals. When speaking about professionals or their work, it is common in many countries and in many professions to refer to them by their surname only. For example, people often talk about *Darwin's theory of evolution*, *Shakespeare's plays, Trump's policies,* etc. It is more difficult to come up with prominent *female* figures who are referred to in this way. Anecdotally, the famous Marie Curie, Jane Goodall, Jane Austen, and Angela Merkel seem usually to be fully named. Are men and women equally likely to be referred to by their surname alone? And if so, is this gender bias innocuous or is there any reason to suspect that it may affect important professional outcomes? Broadly, these are the questions I explore in the final body chapter of this work.

Chapter IV is a reproduction of work I published with co-author Melissa Ferguson in 2018 in the journal *Proceedings of the National Academy of Sciences of the United States of America*, titled "How Gender Determines the Way We Speak About Professionals." The paper describes the results of nine studies. The first four provide archival and experimental evidence that people are more likely to use a surname reference when discussing a male professional than

a female professional. Study 1 reports an analysis of students' reviews of their professors on the popular website RateMyProfessors.com, which revealed that students were more likely to refer to their male than their female professors by surname. A similar pattern emerged for the political domain in Study 2, in which the transcripts of politically diverse radio shows were coded for reference type. As predicted, speakers on the shows were more likely to use a surname reference when talking about a male vs. female target. Study 3 replicated and extended the effect to other domains. Specifically, participants were asked how they normally referred to each one of a list of well known figures, some male and some female, in various domains (e.g., literature, sports, politics, science) when speaking about them in every day life. An analysis of their responses found that people were more likely to refer to the male figures than to the female figures by surname. Additional data was collected to test whether other factors might account for this bias, including gender differences in fame, influence, frequency of first and last name, and the sharing of a surname with a famous family member. These factors failed to explain the effect of gender. Finally, the results of Study 4 suggest that even when gender is the only factor that differs between two individuals, people may still refer to them differently. Participants read information about a fictitious scientist and were asked to rewrite the information in their own words. The group of participants who thought the scientist was male was more likely to refer to him by surname than the group who thought the scientist was female.

The final five studies in the paper explore the possibility that surname use is associated with professional benefits. Specifically, they found that fictional scientists were judged to be more famous and eminent when they were referred to by surname than when they were referred to by full name. This was the case in studies that used a within-subjects design or a between-subjects design. People may infer fame from a surname reference because the listener or reader is

expected to recognize the target using only one identifying piece of information (surname) as opposed to two (first and last name). Judgments of fame and eminence were in turn associated with additional professional benefits: fictional surnamed (vs. full named) scientists were judged to be more high status, more likely to win a prize for their work, and more worthy of a career award. Together, the studies provide evidence of a gender bias in use of surname, and find that a surname reference boosts judgments of fame, eminence, status, and deservingness of awards.

CHAPTER II

WHEN KNOWLEDGE KNOWS NO BOUNDS: SELF-PERCEIVED EXPERTISE PREDICTS CLAIMS OF IMPOSSIBLE KNOWLEDGE[1]

Abstract

People overestimate their knowledge, at times claiming knowledge of concepts, events, and people that do not exist and cannot be known, a phenomenon called overclaiming. What underlies assertions of such impossible knowledge? We found that people overclaim to the extent that they perceive their personal expertise favorably. Studies 1a and 1b showed that self-perceived financial knowledge positively predicts claiming knowledge of nonexistent financial concepts, independent of actual knowledge. Study 2 demonstrated that self-perceived knowledge within specific domains (e.g., biology) is associated specifically with overclaiming within those domains. In Study 3, warning participants that some of the concepts they saw were fictitious did not reduce the relationship between self-perceived knowledge and overclaiming, which suggests that this relationship is not driven by impression management. In Study 4, boosting self-perceived expertise in geography prompted assertions of familiarity with nonexistent places, which supports a causal role for self-perceived expertise in claiming impossible knowledge.

Keywords: knowledge level, judgment, inference, thinking

---

[1] Published as: Atir, S., Rosenzweig, E., & Dunning, D. (2015). When knowledge knows no bounds: Self-perceived expertise predicts claims of impossible knowledge. *Psychological Science*, *26*, 1295-1303.

The menace to understanding [is] not so much ignorance as the illusion of knowledge.

—Daniel J. Boorstin (1987, p. 53), American historian, 1914–2004

Can people differentiate what they know from what they do not? Several lines of research suggest that people are not always accurate judges of their knowledge and often overestimate how much they know (Dunning, 2011; Kruger & Dunning, 1999). Research on overconfidence finds that people commonly judge the accuracy of their judgments too favorably (Fischhoff, Slovic, & Lichtenstein, 1977; Lichtenstein, Fischhoff, & Phillips, 1982; Moore & Healy, 2008) and typically overestimate how well they perform everyday tasks relative to other people (Alicke & Govorun, 2005; Dunning, Heath, & Suls, 2004). Work on the *illusion of explanatory depth* demonstrates that participants tend to think they have a better understanding of how objects work (e.g., a ballpoint pen) than they can demonstrate when that understanding is put to the test (Rozenblit & Keil, 2002).

At times, people even claim knowledge they cannot possibly have, because the object of their knowledge does not exist, a phenomenon known as *overclaiming*. For example, in the late 1970s, nearly a third of American respondents expressed an opinion about the "1975 Public Affairs Act" when asked about it directly, even though the act was a complete fiction (Bishop, Oldendick, Tuchfarber, & Bennet, 1980). Approximately a fifth of consumers report having used products that are actually nonexistent (Phillips & Clancy, 1972). More recent research has asked participants to rate their familiarity with a mix of real and nonexistent concepts, names, and events in domains such as philosophy, life sciences, physical sciences, and literature. Participants reported being familiar with the real items but also, to a lesser degree, with the nonexistent ones. (e.g., Paulhus, Harms, Bruce, & Lysy, 2003).

Although these and other studies document a tendency to claim nonexistent knowledge, little work has explored when or why people are likely to exhibit this tendency. Herein, we focus on the role of self-perceived domain knowledge. For example, if Janet believes her biology knowledge is excellent and Brad believes his is shaky, we suspect that Janet will be more likely than Brad to overclaim knowledge about biology terms. This should also apply within subjects: If Janet considers herself highly knowledgeable in biology but thinks her philosophy knowledge is poor, she will be more likely to overclaim knowledge of biological concepts than of philosophical ones.

A sizable body of work on how people evaluate their own knowledge suggests that they rely not only on a direct examination of their mental contents but also on a *feeling of knowing* (for a review, see Nelson & Narens, 1990; Reder & Ritter, 1992). Notably, a feeling of knowing is often only weakly predictive of actual knowledge (Nelson, 1984) and appears to be informed, at least in part, by top-down inferences about what should be or probably is known (e.g., Costermans, Lories, & Ansay, 1992; Koriat, 1995; but see Hart, 1965, and see Yaniv & Meyer, 1987 for a noninferential account). We theorized that such inferences are drawn from people's preconceived notions about their expertise, inducing a feeling of knowing that then prompts overclaiming.

Several findings suggest that preformed impressions of expertise might influence overclaiming. People judge their quiz performance more favorably when it is framed as testing an ability they think they have (e.g., abstract reasoning) rather than one they think they lack (e.g., computer programming; Ehrlinger & Dunning, 2003), at least partially because their self-perceptions alter the way they experience the task (e.g., whether they answer questions quickly or slowly; Critcher & Dunning, 2009). In addition, level of self-perceived expertise is positively

correlated with providing answers to exceedingly difficult questions and with feelings of certainty but not with answering such questions correctly (Bradley, 1981).

The current investigation tested the relationship between self-perception of domain knowledge and overclaiming knowledge of nonexistent concepts within that domain. We measured overclaiming by asking participants about their familiarity with and knowledge about both real and nonexistent concepts, names, and places (Paulhus et al., 2003), which allowed us to make a clear inference of inappropriate claims of knowledge and to control for claimed knowledge of real items.

After an initial examination of the relationship between self-perceived knowledge and overclaiming in the domain of personal finance (Studies 1a and 1b), we tested the domain specificity of this effect. For example, does self-perceived knowledge in one domain (e.g., biology) predict overclaiming in that domain over and above self-perceived knowledge in other domains (Study 2)? Next, we tested whether overclaiming prompted by self-perceived knowledge was "honest" or was driven by impression-management concerns (Study 3). Finally, to assess whether self-perceived knowledge plays a causal role in overclaiming, we manipulated self-perceived knowledge in geography and measured reported familiarity with nonexistent places (Study 4). In all studies, we assessed overclaiming using a modified version of the signal-detection method recommended by Paulhus et al. (2003).[2]

Studies 1a and 1b

[2] Paulhus et al. (2003) recommended using bias (false alarm rate + hit rate) controlling for accuracy (false alarm rate – hit rate). We used the false alarm rate in place of bias, because hit rate is related to self-perceived knowledge and its inclusion in the dependent variable might inflate our results. We performed the analyses using the recommended method, and all results were essentially unchanged.

In Study 1a, we tested whether individuals who perceive themselves as more knowledgeable in personal finance would be more likely to claim knowledge of nonexistent financial terms. In the realm of finance, failure to recognize or admit one's knowledge gaps could lead to uninformed financial decisions with devastating consequences (Lusardi & Mitchell, 2014). Do self-proclaimed financial experts claim more financial knowledge than they can possibly have?

However, note that self-perceived knowledge may predict overclaiming because it is confounded with genuine knowledge. To address this possibility, we also asked participants in Study 1b to complete a standard financial-literacy quiz, which allowed us to test whether self-perceived knowledge predicted overclaiming over and above any potential relationship between genuine knowledge and overclaiming.

Method

Participants. Study 1a had 100 participants (33 women, 66 men; mean age = 31 years, *SD* = 9.7; 1 participant did not report demographic information). Two additional participants failed to complete the entire study and were excluded from all analyses. Study 1b had 202 participants (85 women, 115 men, 2 whose gender was not reported; mean age = 33.5 years, *SD* = 10.1). Twelve additional participants failed to complete the entire study and were excluded from all analyses. Both samples were recruited through Amazon's Mechanical Turk and were restricted to respondents within the United States. The sample size for Study 1a allowed an 80% probability of identifying a significant effect if the true correlation was .30; the sample was doubled in Study 1b to account for the possibly smaller correlations (.20) after genuine knowledge was controlled for.

Procedure. After providing informed consent, participants rated their general knowledge of personal finance and completed the overclaiming task in counterbalanced order. The questions on personal finance were "In general, how knowledgeable would you say you are about personal finance?" (1 = *not knowledgeable at all*, 7 = *extremely knowledgeable*) and "How would you rate your general knowledge of personal finance compared to the average American?" (1 = *much less knowledgeable*, 7 = *much more knowledgeable*).

The overclaiming task was modeled after the Overclaiming Questionnaire (Paulhus et al., 2003). Participants were asked to rate their knowledge of various personal-finance-related terms:

> We are interested in common knowledge about personal finance. You will see 15 terms related to personal finance. Please rate your knowledge about each term by choosing the appropriate number from 1 (*never heard of it*) to 7 (*very knowledgeable*).

The 15 items were presented one at a time in random order for each participant. Twelve of the 15 were real terms (e.g., *tax bracket*) collected from various finance Web sites, and 3 were nonexistent foils invented by the researchers (e.g., *pre-rated stocks*; full list of items available in Appendix A). Finally, participants filled out a demographic questionnaire and provided information for payment.

The procedure in Study 1b was identical to that in Study 1a, except that participants also completed a widely used financial-literacy quiz (FINRA Investor Education Foundation, 2009; Lusardi & Mitchell, 2011) after completing the other two tasks. The quiz included five questions assessing financial capability (e.g., "A 15-year mortgage typically requires higher monthly payments than a 30-year mortgage, but the total interest paid over the life of the loan will be less"; response options: "true," "false," and "don't know").

Results

Overclaiming was measured by calculating the false alarm rate, which is the proportion of nonexistent foils about which a participant claimed knowledge. We averaged the false alarm

rates for each of the six potential knowledge cutoff points (i.e., we computed the proportion of foils rated as 2 or higher, the proportion of foils rated as 3 or higher, and so on for 4, 5, 6, and 7, and then averaged these proportions), which resulted in an overclaiming value ranging from 0 to 1. In Study 1a, mean overclaiming was 0.29 (*SD* = 0.20), and 93% of participants claimed at least some knowledge of at least one foil. In Study 1b, mean overclaiming was 0.31 (*SD* = 0.23), and 91% of participants claimed some knowledge of at least one foil.

Following the operationalization in previous work on overclaiming (e.g., Paulhus et al., 2003), we looked at overclaiming while controlling for accuracy. Accuracy was obtained by subtracting the averaged false alarm rate from the averaged hit rate (i.e., the proportion of real items about which each participant claimed knowledge, averaged across all six potential cutoff points).

In Study 1a, to test whether self-perceived knowledge predicted overclaiming, we averaged the responses to the two questions measuring self-perceived knowledge in personal finance (α = .91). We next entered self-perceived knowledge of personal finance (*M* = 4.23, *SD* = 1.22) and accuracy into a regression model predicting overclaiming. Self-perceived knowledge positively predicted overclaiming, $b = 0.09$, $t(97) = 9.17$, $p < .001$. The more participants viewed themselves as knowledgeable about personal finance, the more they claimed knowledge of nonexistent personal finance terms.

Likewise, in Study 1b, self-perceived knowledge of personal finance (*M* = 4.43, *SD* = 1.17) positively predicted overclaiming, $b = 0.10$, $t(199) = 13.07$, $p < .001$. In addition, an unanticipated order effect emerged in Study 1b (but not in Study 1a), such that overclaiming was higher when self-perceived knowledge was assessed first (*M* = 0.34, *SD* = 0.24) rather than second (*M* = 0.27, *SD* = 0.21), $t(200) = 2.21$, $p < .05$. However, self-perceived knowledge

significantly predicted overclaiming regardless of whether it was assessed before the overclaiming task, $b = 0.12$, $t(98) = 9.83$, $p < .001$, or after it, $b = 0.09$, $t(98) = 8.47$, $p < .001$. Thus, the order of the tasks was fixed for Studies 2 to 4. Genuine knowledge, as assessed by the financial literacy quiz, also positively predicted overclaiming, $b = 0.05$, $t(199) = 4.92$, $p < .001$. Self-perceived knowledge was positively correlated with genuine knowledge, $r(200) = .32$, $p < .001$.

Finally, we tested whether self-perceived knowledge predicted overclaiming while controlling for genuine knowledge. Scores on the financial literacy quiz could range from 0 to 5. On average, participants answered 3.7 ($SD = 1.9$) questions correctly. By comparison, a nationally representative U.S. sample had average scores of 3.0 and 2.9 in 2009 and 2012, respectively (FINRA Investor Education Foundation, 2013). When genuine knowledge was entered into the model with self-perceived knowledge, self-perceived knowledge remained a highly significant predictor of overclaiming, $b = 0.09$, $t(198) = 11.73$, $p < .001$. Interestingly, genuine knowledge also remained a positive predictor of overclaiming, $b = 0.02$, $t(198) = 2.14$, $p = .033$, which provides preliminary evidence of an independent effect of genuine knowledge on overclaiming.

Study 2

Studies 1a and 1b provided initial evidence that self-perceived knowledge in a particular domain is positively associated with overclaiming within that domain. An alternative interpretation of the result is that it captures only an association between two more general individual differences; people who generally perceive themselves as more knowledgeable are also generally more likely to overclaim in any domain. If that is the case, self-perceived knowledge in a particular domain should predict overclaiming equally well within that domain

and within unrelated ones. In contrast, we hypothesized that self-perceived knowledge has a domain-specific effect on overclaiming.

In Study 2, we explored the question of generality versus specificity by measuring self-perceived knowledge and overclaiming in several domains. We also varied whether the overclaiming questionnaire asked participants about their familiarity with items (as in the original overclaiming questionnaire; Paulhus et al., 2003) or their knowledge of them (as in Studies 1a and 1b).

Method

Participants. One hundred twenty-four people[3] (52 women, 71 men; mean age = 33 years, *SD* = 12.7; 1 participant did not report age or gender) participated online through Amazon's Mechanical Turk. Participation was restricted to respondents within the United States. Four additional participants failed to complete the entire study and were excluded from all analyses. Pretesting showed that a sample size of 100 participants would provide sufficient power to detect the relevant relationship (a correlation of roughly .30).

Procedure. After providing informed consent, participants were asked to rate their general knowledge in various domains: "Please rate your knowledge of the following topics using the following scale: 0 = *no knowledge*; 1 = *limited knowledge*; 2 = *moderate knowledge*; 3 = *substantial knowledge*; 4 = *extensive knowledge*." Participants were asked about three domains of interest (biology, philosophy, and literature) as well as four filler domains (mathematics, architecture, computer programming, and 20th-century art), presented in random order for each participant.

[3] An additional 27 participants completed a version of the study that did not include the perceived knowledge questions. They were not included in the analyses.

Participants then completed an overclaiming questionnaire for the domains of interest (items borrowed from Paulhus et al., 2003). For each of these domains, they saw 15 items, presented in a random order (randomly determined for each participant) on the same page. Twelve items were real (e.g., in biology: *mammal*, *adrenal gland*, *sciatica*) and three were foils (e.g., in biology: *meta-toxins*; full lists available in Appendix B). Approximately half of the participants (n = 61) were asked to rate their knowledge of each item, and the rest (n = 63) were asked to rate their familiarity with each item (1 = *never heard of it*, 7 = *very knowledgeable* or *very familiar*). Finally, participants filled out a demographic questionnaire and provided information for payment.

Results

The majority of participants claimed at least some familiarity with (92%) or knowledge of (87%) at least one foil. Our results replicated the positive relationship between self-perceived knowledge (biology: M = 2.70, SD = 0.92; philosophy: M = 2.34, SD = 0.83; literature: M = 2.80, SD = 0.99) and overclaiming (familiarity—biology: M = 0.26, SD = 0.24; philosophy: M = 0.20, SD = 0.20; literature: M = 0.09, SD = 0.18; knowledge—biology: M = 0.21, SD = 0.22; philosophy: M = 0.19, SD = 0.22; literature: M = 0.09, SD = 0.19). In each domain, self-perceived knowledge positively predicted overclaiming when we controlled for accuracy. Table 1 shows that this relationship emerged both for participants rating their familiarity with items and for participants rating their knowledge of the same items.

Domain	Familiarity		Knowledge	
	Domain SPK only	Controlling for SPK in other domains	Domain SPK only	Controlling for SPK in other domains
Biology	0.10***	0.10***	0.15***	0.12***
Philosophy	0.11***	0.09*	0.16***	0.10***
Literature	0.08***	0.05*	0.10***	0.06**

Table 1. Results of regression analyses: Domain-related self-perceived knowledge (SPK) as a predictor of overclaiming before and after controlling for SPK in other domains.

Note: The first two columns use overclaiming familiarity as the dependent variable; the last two columns use overclaiming knowledge as the dependent variable. The table presents unstandardized coefficients. We controlled for accuracy in all analyses.

*$p < .05$. **$p < .01$. ***$p < .001$.

Next, we tested whether domain-specific self-perceived knowledge remained a significant positive predictor of overclaiming within that domain after controlling for self-perceived knowledge in other domains. As Table 1 shows, when we predicted overclaiming in a specific domain (e.g., biology) from self-perceived knowledge in all three domains simultaneously, we found that self-perceived knowledge in the relevant domain continued to significantly predict overclaiming (within-domain *b*s ranged from 0.05 to 0.12), all *p*s < .05. Of the 12 cross-domain correlations created by this analysis, only 2 proved to be significant at $p < .05$ (*b*s ranged from –0.04 to 0.065). Thus, general individual differences may account for some of the association between self-perceived knowledge and overclaiming. However, these results suggest that, beyond the effect of individual differences, there is a distinct positive association between self-perceived knowledge in a particular domain and the likelihood of overclaiming

within that domain. Note that these results hold whether knowledge or familiarity is used as the measure of overclaiming.

Study 3

We hypothesized earlier that self-perceived knowledge prompts a top-down inference of familiarity that arises when reading and processing the items (Critcher & Dunning, 2009). An alternative explanation is that self-perceived knowledge increases the pretense of knowledge—a phenomenon driven by impression-management goals. Individuals with higher self-perceived knowledge might not experience bogus items as more familiar but may instead simply alter their ratings to portray themselves as knowledgeable.

To test this possibility, we modified our procedure by adding a warning manipulation, which in previous research has been shown to decrease overclaiming overall (Paulhus et al., 2003). Half of participants were warned that some of the items they would be shown did not exist. If individuals with high self-perceived knowledge are only feigning, this warning should serve as a counterincentive, as claiming nonexistent knowledge would be detrimental to the impression of expertise they might wish to put forth. Thus, the warning should reduce overclaiming and diminish the relationship between self-perceived knowledge and overclaiming. However, if people with greater self-perceived expertise truly experience the foils as more familiar, they should still be more likely to overclaim. We therefore predicted that warning participants that some items do not exist would reduce overclaiming overall but would not alter the relationship between self-perceived knowledge and overclaiming.

Method

Participants. Ninety-seven individuals (47 women, 50 men; mean age = 34 years, *SD* = 11) participated online through Amazon's Mechanical Turk. Participation was restricted to

respondents within the United States. Two additional participants failed the attention check and were excluded from all analyses. We had determined that a sample size of 100 participants would provide an 80% probability of identifying a difference between conditions if the true effect size (*d*) was 0.5.

Procedure. After providing informed consent, participants were asked to rate their general knowledge in various domains (1 = *not knowledgeable at all*, 7 = *extremely knowledgeable*), including three domains of interest (biology, philosophy, and history) and four filler domains (American literature, mathematics, computer programming, and 20th-century art), presented in random order (randomly determined for each participant).

Participants were then randomly assigned to one of two conditions: warning ($n = 49$) and no warning ($n = 48$). All participants read the same instructions explaining that they would see items in three different categories and would be asked to rate their familiarity with each item. These instructions were followed by either a warning or a control sentence, bolded and underlined: "Note that some of the items in this inventory do not exist [are very difficult]" (Paulhus et al., 2003). To check that participants read the warning, we presented them with an instructions comprehension check, which constituted the statement "Some of these items do not exist" (response options were "true," "false," and "I'm not sure") and two general-attention filler questions: "In this part of the study, you will see items from how many categories?" (response options were "1," "2," "3," "4," and "I'm not sure") and "You will be asked to rate your familiarity with different items" (response options were "true," "false," and "I'm not sure").

Participants then completed an overclaiming questionnaire for the domains of interest (i.e., biology, philosophy, and history; items borrowed from Paulhus et al., 2003; full lists available in Appendix B). For each of these, they used a 7-point scale to rate their familiarity

with 15 domain-related items, presented in random order for each participant, with each item on a separate page. Twelve of the 15 items were real, and 3 were foils. Finally, participants filled out a demographic questionnaire and provided information for payment.

Results

Two participants whose overclaiming rate was more than 3 *SD*s above the mean were excluded from all analyses. Participants in the warning condition were more likely to indicate that the statement "Some of these items do not exist" was true than those in the no-warning condition ($M = 81.3\%$ vs. $M = 6.4\%$), $\chi^2(1, N = 95) = 61.8$, $p < .001$, which confirms that participants read the warning. Eighty-five percent of participants claimed at least some familiarity with at least one foil. To increase our power to detect any relationship between the warning condition and perceived knowledge, we used a linear mixed model, which included a fixed effect for the warning condition and fixed indicator variables for domain. We also included a random intercept for participant to control for within-subject variance in overclaiming and for the nonindependence of each participant's responses. As found previously (Paulhus et al., 2003), participants who were warned overclaimed less (history: $M = 0.07$, $SD = 0.12$; philosophy: $M = 0.21$, $SD = 0.22$; biology: $M = 0.17$, $SD = 0.12$) than those who were not warned (history: $M = 0.11$, $SD = 0.13$; philosophy: $M = 0.30$, $SD = 0.24$; biology: $M = 0.26$, $SD = 0.21$), $t(93) = -2.20$, $p = .030$, $d = 0.45$. This relationship was in the same direction but nonsignificant when accuracy was entered into the model, $t(93.55) = -1.5$, $p = .128$.

We then tested whether the relationship between self-perceived knowledge and overclaiming interacted with the warning condition. We added accuracy, self-perceived knowledge, and the interaction between self-perceived knowledge and warning to the model. As we found previously, self-perceived knowledge (history: $M = 3.98$, $SD = 1.47$; philosophy: $M =$

3.08, *SD* = 1.54; biology: *M* = 3.34, *SD* = 1.46) positively predicted overclaiming, $b = 0.05$, $t(262.57) = 7.44$, $p < .001$. Moreover, the effect of self-perceived knowledge did not interact with the warning condition, $b = -0.005$, $t(260.89) = -0.52$. Thus, warning people that some of the items do not exist reduced overclaiming as a whole but neither eliminated nor attenuated the positive relationship between self-perceived knowledge and overclaiming.

Study 4

In Study 4, we tested the causal role of self-perception by manipulating self-perceived knowledge to determine whether it influenced overclaiming. Manipulating self-perceived knowledge also allowed us to assess the effect of self-perceived knowledge independently of genuine knowledge, two variables that could not be decoupled in Studies 1 to 3. We shifted participants' perceptions of their North American geography knowledge by giving them an easy or difficult North American geography quiz before the overclaiming questionnaire (method taken from Ehrlinger & Dunning, 2003).

Method

Participants. One hundred forty-eight individuals (55 women, 94 men; mean age = 28 years, *SD* = 9) participated online through Amazon's Mechanical Turk. Participation was restricted to respondents within the United States. One additional participant who participated twice was excluded from all analyses. Pretesting showed that a sample size of 150 participants would provide an 80% probability of identifying a difference between conditions if the true effect size (*d*) was 0.5.

Procedure. After providing informed consent, participants were randomly assigned to one of three conditions: easy quiz ($n = 49$), difficult quiz ($n = 50$), or no quiz ($n = 49$). The quiz included questions about North American travel and geography. The questions in the easy

condition were meant to give participants the sense that they were relatively well traveled and well versed in North American geography because these participants were likely to answer "yes" (e.g., "Have you ever been to New York? Yes/No"; full list available in Appendix C) or to choose a high-numbered answer (e.g., "How many state capitals can you name? (a) 1–2, (b) 3–4, (c) 5 or more). Questions in the difficult condition were similar but meant to induce the opposite feeling because participants were likely to answer "no" (e.g., "Have you ever been to North Dakota? Y/N") or to choose a low-numbered answer (e.g., "How many state capitals can you name? (a) 1–10, (b) 11–30, (c) 31 or more"). As a manipulation check, all participants then rated their knowledge of North American geography (1 = *my geography knowledge is very weak*, 10 = *my geography knowledge is very strong*).

Participants then completed an overclaiming questionnaire in which they were presented with 15 randomly ordered places in the U.S. Twelve of these places were real (e.g., *Philadelphia, Pennsylvania; The National Mall; Acadia National Park*; full list available in Appendix D) and 3 were foils (e.g., *Monroe, Montana*). Participants were asked to rate their familiarity with each (0 = *never heard of it*, 6 = *very familiar*). Finally, participants filled out a demographic questionnaire and provided information for payment.[4]

Results

Compared with participants in the difficult-quiz condition, participants in the easy-quiz condition reported having visited more of the places about which they were asked ($M = 2.61$, $SD = 1.56$, vs. $M = 0.94$, $SD = 1.17$), $t(97) = 6.0$, $p < .001$, and they reported better geography knowledge on the manipulation check ($M = 2.32$, $SD = 0.38$, vs. $M = 1.51$, $SD = 0.37$), $t(97) = 10.8$, $p < .001$. Thus, as expected, participants who completed the easy quiz rated their

[4] Participants then completed an unrelated study not discussed here.

knowledge of North American geography higher (M = 6.37, SD = 2.28) than did those who completed the difficult quiz (M = 5.40, SD = 2.30), $t(97) = 2.10$, $p = .038$, $d = 0.42$. The ratings of the participants who completed no quiz fell in between the ratings from participants in the difficult- and easy-quiz conditions (M =5.9, SD = 2.3) but did not differ significantly from either, ts < 1.07.

Forty-three percent of participants claimed at least some familiarity with at least one foil. To test whether manipulating self-perceived knowledge would influence overclaiming, we entered accuracy and quiz condition (easy, difficult, none) into a regression model to predict overclaiming in North American geography. We found a significant effect of condition, $F(2, 144) = 6.73$, $p = .002$, $\eta^2 = .09$. Participants in the easy-quiz condition overclaimed more (M = 0.16, SD = 0.20) than did those in the difficult-quiz condition (M = 0.05, SD = 0.10), $t(96) = 2.78$, $p = .007$, $d = 0.57$, and those in the no-quiz condition (M = 0.07, SD = 0.13), $t(95) = 2.92$, $p = .004$, $d = 0.60$. Participants in the difficult- and no-quiz conditions did not differ significantly, $t(96) = 0.10$. Thus, participants induced to feel more knowledgeable about North American geography were more likely to claim familiarity with nonexistent places in the United States, which is consistent with a causal account of the role of self-perceived knowledge in overclaiming.

General Discussion

Our work suggests that the seemingly straightforward task of judging one's knowledge may not be so simple, particularly for individuals who believe they have a relatively high level of knowledge to begin with. In Study 1a, we found that self-perceived knowledge of personal finance positively predicted claiming knowledge of nonexistent domain-related terms. The results of Study 1b indicated that this effect was not driven by genuine domain knowledge; self-

perceived knowledge remained a significant predictor of overclaiming when we controlled for genuine knowledge of personal finance. In Study 2, we found that self-perceived knowledge had domain-specific effects on overclaiming. Study 3 revealed that warning participants that some of the items they would encounter were bogus did not alter the relationship between self-perceived knowledge and overclaiming, suggesting that self-perceptions were prompting mistaken but honest claims of knowledge. Finally, Study 4 demonstrated a causal influence of self-perceived knowledge on overclaiming. Experimentally enhancing self-perceived knowledge in geography increased overclaiming knowledge of nonexistent places. These results converge to demonstrate that the more individuals believe they know about a domain, the more likely they are to claim knowledge in that domain that they cannot possibly possess.

These findings add to the body of work on how individuals assess their own knowledge. Our results suggest that people do not simply consult a "mental index" that catalogues their knowledge but instead draw on preexisting self-perceptions of knowledge to make inferences about what they should or probably do know (e.g., Koriat, 1995). For domains of high self-perceived expertise, these inferences may induce a sense of familiarity with terms that sound plausibly real but are not.

An alternative explanation, which does not exclude the first, is that greater self-perceived knowledge leads people to be more motivated to search their memories for relevant knowledge. Individuals who perceive themselves as more knowledgeable in biology, for example, may be more motivated to construct a plausible notion of what "bio-sexual" means. Independent of differences in people's initial sense of familiarity, high self-perceived domain expertise may lead to a confirmation-biased memory search (e.g., Kunda, 1990) for some way that the nonexistent term might indeed be familiar.

It is easy to imagine how a tendency to overclaim, especially in self-perceived experts, could have adverse consequences. Self-perceived experts may give bad counsel when they should give none. For instance, an individual considering a financial decision may consult a friend who expresses confidence in her financial knowledge (Zarnoth & Sniezek, 1997). That friend may provide inappropriate advice because she fails to recognize her insufficient familiarity with the question. Further, a tendency to overclaim may discourage individuals from educating themselves in precisely those areas in which they consider themselves knowledgeable and that may be important to them (Metcalfe, 2009). In other words, overclaiming may hinder people from truly achieving a valuable level of genuine knowledge.

Future research should investigate these and other potential consequences of overclaiming. Another area to explore is the relationship between overclaiming and genuine expertise. In Study 1b, self-perceived knowledge predicted overclaiming over and above genuine knowledge, which suggests that the observed relationship between self-perceived knowledge and overclaiming is not the result of a confound with genuine knowledge. However, genuine knowledge also emerged as an independent predictor of overclaiming, albeit a weaker one than self-perceived knowledge. Future research should explore this relationship. Potentially, individuals with more domain knowledge overclaim more because they have a larger knowledge base from which to draw when assessing familiarity with plausible-sounding foils in that domain. Education may unwittingly aid and abet that phenomenon. Continuing to explore when and why individuals overclaim may prove important in battling that great menace—not ignorance, but the illusion of knowledge.

References

Alicke, M. D., & Govorun, O. (2005). The better-than-average effect. In M. D. Alicke, D. Dunning, & J. Kruger (Eds.), *The self in social judgment* (pp. 85–106). New York, NY: Psychology Press.

Bishop, G. F., Oldendick, R. W., Tuchfarber, A. J., & Bennett, S. E. (1980). Pseudo-opinions on public affairs. *Public Opinion Quarterly*, *44*, 198–209.

Boorstin, D. J. (1987). *Hidden history: Exploring our secret past.* New York: Harper and Row.

Bradley, J. V. (1981). Overconfidence in ignorant experts. *Bulletin of the Psychonomic Society*, *17*, 82–84. doi:10.3758/BF03333674

Costermans, J., Lories, G., & Ansay, C. (1992). Confidence level and feeling of knowing in question answering: The weight of inferential processes. *Journal of Experimental Psychology*: *Learning*, *Memory*, *and Cognition*, *18*, 142–150. doi:10.1037/0278-7393.18.1.142

Critcher, C. R., & Dunning, D. (2009). How chronic self-views influence (and mislead) self-assessments of task performance: Self-views shape bottom-up experiences with the task. *Journal of Personality and Social Psychology*, *97*, 931–945. doi:10.1037/a0017452

Dunning, D. (2011). The Dunning-Kruger effect: On being ignorant of one's own ignorance. In J. Olson & M. P. Zanna (Eds.), *Advances in experimental social psychology* (Vol. 44, pp. 247–296). New York, NY: Elsevier.

Dunning, D., Heath, C., & Suls, J. M. (2004). Flawed self-assessment: Implications for health, education, and the workplace. *Psychological Science in the Public Interest*, *5*, 69–106. doi:10.1111/j.1529-1006.2004.00018.x

Ehrlinger, J., & Dunning, D. (2003). How chronic self-views influence (and potentially mislead) estimates of performance. *Journal of Personality and Social Psychology*, *84*, 5–17. doi:10.1037/0022-3514.84.1.5

FINRA Investor Education Foundation. (2009). *National Financial Capability Study*: *2009 national survey questionnaire*. Retrieved from http://www.usfinancialcapability.org/downloads/NFCS_2009_Natl_Qre_Eng.pdf

FINRA Investor Education Foundation. (2013). *Financial capability in the United States*: *Report of findings from the 2012 National Financial Capability Study*. Retrieved from http://www.usfinancialcapability.org/downloads/NFCS_2012_Report_Natl_Findings.pdf

Fischhoff, B., Slovic, P., & Lichtenstein, S. (1977). Knowing with certainty: The appropriateness of extreme confidence. *Journal of Experimental Psychology*: *Human Perception and Performance*, *3*, 552–564. doi:10.1037/0096-1523.3.4.552

Hart, J. T. (1965). Memory and the feeling-of-knowing experience. *Journal of Educational Psychology*, *56*, 208–216. doi:10.1037/h0022263

Koriat, A. (1995). Dissociating knowing and the feeling of knowing: Further evidence for the accessibility model. *Journal of Experimental Psychology*: *General*, *124*, 311–333. doi:10.1037/0096-3445.124.3.311

Kruger, J., & Dunning, D. (1999). Unskilled and unaware of it: How difficulties in recognizing one's own incompetence lead to inflated self-assessments. *Journal of Personality and Social Psychology*, *77*, 1121–1134. doi:10.1037/0022-3514.77.6.1121

Kunda, Z. (1990). The case for motivated reasoning. *Psychological Bulletin*, *108*, 480–498. doi:10.1037/0033-2909.108.3.480

Lichtenstein, S., Fischhoff, B., & Phillips, L. D. (1982). Calibration of probabilities: The state of the art to 1980. In D. Kahneman, P. Slovic, & A. Tversky (Eds.), *Judgment under uncertainty*: *Heuristics and biases* (pp. 306–334). Cambridge, England: Cambridge University Press.

Lusardi, A., & Mitchell, O. S. (2011). Financial literacy and planning: Implications for retirement wellbeing. In O. S. Mitchell & A. Lusardi (Eds.), *Financial literacy*: *Implications for retirement security and the financial marketplace* (pp. 17–39). Oxford, England: Oxford University Press.

Lusardi, A., & Mitchell, O. S. (2014). The economic importance of financial literacy: Theory and evidence. *Journal of Economic Literature*, *52*, 5–44. doi:10.1257/jel.52.1.5

Metcalfe, J. (2009). Metacognitive judgments and control of study. *Current Directions in Psychological Science*, *18*, 159–163. doi:10.1111/j.1467-8721.2009.01628.x

Moore, D. A., & Healy, P. J. (2008). The trouble with overconfidence. *Psychological Review*, *115*, 502–517. doi:10.1037/0033-295X.115.2.502

Nelson, T. O. (1984). A comparison of current measures of the accuracy of feeling-of-knowing predictions. *Psychological Bulletin*, *95*, 109–133. doi:10.1037/0033-2909.95.1.109

Nelson, T. O., & Narens, L. (1990). Metamemory: A theoretical framework and new findings. In G. H. Bower (Ed.), *The psychology of learning and motivation* (Vol. 26, pp. 125–141). New York, NY: Academic Press.

Paulhus, D. L., Harms, P. D., Bruce, M. N., & Lysy, D. C. (2003). The over-claiming technique: Measuring self-enhancement independent of ability. *Journal of Personality and Social Psychology*, *84*, 890–904. doi:10.1037/0022-3514.84.4.890

Phillips, D. L., & Clancy, K. J. (1972). Some effects of "social desirability" in survey studies. *American Journal of Sociology*, *77*, 921–940. doi:10.1086/225231

Reder, L. M., & Ritter, F. E. (1992). What determines initial feeling of knowing? Familiarity with question terms, not with the answer. *Journal of Experimental Psychology*: *Learning*, *Memory*, *and Cognition*, *18*, 435–451. doi:10.1037/0278-7393.18.3.435

Rozenblit, L., & Keil, F. C. (2002). The misunderstood limits of folk science: An illusion of explanatory depth. *Cognitive Science*, *26*, 521–562. doi:10.1016/S0364-0213(02)00073-3

Yaniv, I., & Meyer, D. E. (1987). Activation and metacognition of inaccessible stored information: Potential bases for incubation effects in problem solving. *Journal of Experimental Psychology*: *Learning*, *Memory*, *and Cognition*, *13*, 187–205. doi:10.1037/0278-7393.13.2.187

Zarnoth, P., & Sniezek, J. A. (1997). The social influence of confidence in group decision making. *Journal of Experimental Social Psychology*, *33*, 345–366. doi:10.1006/jesp.1997.1326

Appendix A: Study 1 Stimuli

Foils indicated by *

tax bracket

Fixed-rate mortgage

home equity

revolving credit

annualized credit*

vesting

Retirement

stock options

Inflation

private equity fund

interest rate

pre-rated stocks*

Roth IRA

whole life insurance

fixed-rate deduction*

Appendix B: Study 2 and 3 Stimuli

Foils indicated by *

Biology	**Philosophy**
mammal	logistic heresy*
adrenal gland	creationism
sciatica	Goedel's theorem
insulin	social constructionism
meta-toxins*	Platonic sense*
intestine	hermeneutics
bio-sexual*	esoteric deduction*
meiosis	ghost in the machine
ribonucleic acid	Hegel
electrocardiograph	Socrates
amniotic sac	categorical imperative
hemoglobin	free will
retroplex*	Ayn Rand
antigen	situational ethics
recessive trait	Principia Mathematica

Literature	**History**
Antigone	Napoleon
Murphy's Last Ride*	Robespierre
Catcher in the Rye	El Puente*
The Bible	My Lai
Hiawatha	The Lusitania
Trapnell Meets Katz*	Ronald Reagan
Mein Kampf	Prince Lorenzo*
The Aeneid	The Luddites
Faustus	Neville Chamberlain
The Boy Who Cried Wolf	Vichy Government
Pygmalion	Queen Shattuck*
Hickory Dickory Dock	Bay of Pigs
The Divine Comedy	Torquemada
Windermere Wild*	Wounded Knee
The Raven	Clara Barton

Appendix C: Study 4 North American Geography Quiz Questions

Easy	**Difficult**
Have you ever been to New York?	Have you ever been to Mississippi?
Have you ever been to Washington, D.C.?	Have you ever been to Missouri?
Have you ever been to California?	Have you ever been to Wyoming?
Have you ever been to Pennsylvania?	Have you ever been to North Dakota?
Have you ever been to Florida?	Have you ever been to Nebraska?
	Have you ever been to Oregon?
Can you name the two most populated states?	Can you name the 15 most populated states?

Question	**Options**	
	Easy condition	**Difficult condition**
How many states can you name?	5; 6-14; 15 or more	0-40; 41-49; 50
How many state capitals can you name?	1-2; 3-4; 5 or more	1-10; 11-30; 31 or more
How many states have you traveled in?	1-2; 3-4; 5 or more	1-10; 11-30; 31 or more
For how many states can you name the biggest city?	1-2; 3-4; 5 or more	1-10; 11-30; 31 or more
How many Canadian provinces can you name?	0; 1-2; 3 or more	1-4; 5-8; 9 or more
How many cities in Canada can you name?	0; 1-2; 3 or more	1-5; 6-10; 11 or more
How many cities in Mexico can you name?	0; 1-2; 3 or more	1-5; 6-10; 11 or more

Appendix D: Study 4 Stimuli

Foils indicated by *

Philadelphia, Pennsylvania

Monroe, Montana*

The national mall

Hoover Dam

Acadia National Park

Buffalo, New York

Lake Othello, Wisconsin*

Davis, California

Aleutian Range, Alaska

Springfield, Massachusetts

Jefferson City, Missouri

Lake Erie

Arlington, Texas

Cheyenne, Wyoming

Cashmere, Oregon*

CHAPTER III

EXPERTS BETTER KNOW WHAT THEY DON'T KNOW: GENUINE EXPERTISE PREDICTS LESS OVERCLAIMING OF IMPOSSIBLE KNOWLEDGE[5]

Abstract

Does true expertise make it easier or harder to assess one's knowledge? People sometimes mistakenly believe they have knowledge of nonexistent terms, a phenomenon called overclaiming. One hypothesis is that genuinely knowledgeable people are more prone to overclaiming within their domain of expertise because they have more information with which to construct apparent meaning for nonexistent terms. Alternatively, genuine knowledge may aid in the detection of ignorance. In 4 studies, we find support for the latter. Doctors and medical students were less likely than premedical students to overclaim knowledge of invented medical terms (Study1). Participants who knew more about science and U.S. civics were less likely to overclaim knowledge within those domains, holding constant self-perceived knowledge (Studies 2-4). The negative relationship between genuine knowledge and overclaiming was mediated by a larger gap between the number of associations knowledgeable people generated about real and bogus terms (Study 3). Also, people with more knowledge followed a more automatic (vs. deliberative) thinking process when judging their knowledge, which in turn predicted less overclaiming (Study 4).

Keywords: knowledge, expertise, self-knowledge, judgment, inference

[5] Based on a manuscript of the same title authored by Atir, S., Rosenzweig, E., & Dunning, D. and prepared for submission to *Psychological Science*.

"Ignorance more frequently begets confidence than does knowledge"

Charles Darwin, English naturalist, geologist and biologist, 1809-1882

It is important to recognize what one knows, but no less important to identify the limits of one's knowledge; an overly optimistic assessment of where those limits lie can lead to confident but ill-advised decisions (e.g., Barber & Odean, 2000; Gervais, & Odean, 2001; Manners & Moore, 2013), and can hinder the pursuit of necessary information (Christensen-Szalanski & Bushyhead, 1981). Previous work, however, suggests that teasing apart what one knows from what one does not can be tricky (e.g., Benjamin, Bjork, & Schwartz, 1998; Koriat, 1995; Rozenblit & Keil, 2002). People often misestimate their knowledge, sometimes even believing themselves to be familiar with concepts, terms, people, etc. that were invented by the researcher for the purpose of the study and therefore cannot be known (Nathanson, Westlake, & Paulhus, 2007; Paulhus, Harms, Bruce, & Lysy, 2003). This phenomenon is called *overclaiming*. For example, when participants are asked how knowledgeable they are about specific terms related to personal finance, some will know they are familiar with what *inflation* means, but some will also think they know something about *fixed-rate deduction*, a bogus term that does not exist (Atir, Rosenzweig, & Dunning, 2015).

Previous work has demonstrated that overclaiming is partly determined by self-perceptions of knowledge; people are more likely to mistakenly believe they know specific terms or information when they consider themselves to be generally knowledgeable in the domain from which the terms are drawn (Atir et al., 2015; Bradley, 1981). For instance, to revisit the personal finance example, those who consider themselves knowledgeable about personal finance are more likely to incorrectly indicate they know what a *fixed-rate deduction* is. But what about people

who truly *are* knowledgeable? People's perceptions of their knowledge can diverge sharply from their genuine knowledge (e.g., Critcher & Dunning, 2009; Koriat, 1995), such that there is often a mismatch between how knowledgeable people believe they are and how knowledgeable they demonstrably are. Do people who genuinely know more overclaim more, or does possessing knowledge provide some protection against overclaiming?

In this paper we explore the link between genuine knowledge and overclaiming. Does being truly knowledgeable within a domain make people better able to determine whether they know specific items within it? One hypothesis is that genuine knowledge increases the likelihood that people will overclaim, because their richer base of knowledge provides more material with which to confabulate and thereby arrive at a false sense of familiarity with the term (Graeff, 2003; Schuman & Presser, 2980). Alternatively, knowledgeable people may be better able to recognize their own ignorance. Having more knowledge may sharpen one's ability to distinguish between what is known and what is unknown (Cantor & Marsh, 2017). In the context of overclaiming, genuine knowledge may protect against mistakes by creating a clearer, more detailed picture of the information that one knows well, which stands in contrast to the relatively sparse information that comes to mind when one encounters an unknown term. This prediction is consistent with previous work on contrasts effects (e.g., Dijksterhuis et al., 1998).

Judgments of knowledge, even about specific terms or questions, can be informed by heuristics (Koriat & Levy-Sadot, 2001). People often use meta-cognitive cues to infer their knowledge, and may or may not engage in a full-blown retrieval attempt (Whittlesea, Jacoby & Girard, 1990). One potential cue for assessing knowledge may be the relative sense of familiarity or knowledge one feels about a given term or piece of information within their domain of expertise. If terms feels much less familiar, either relative to the other terms at hand or to

previous experience with information within the domain, the person can infer that they have little knowledge of it. Critically, relying on a relative sense of familiarity with the terms will be more effective for people who are more knowledgeable. For them, the experience with a bogus item will be very different from the experience with real terms, which they know well. Knowledgeable people have a large store of relevant knowledge that allows them to more fluently recognize the terms that do not immediately evoke a sense of familiarity. The relative gap in familiarity will act as a meta-cognitive cue, which will make it immediately clear that they do not know the bogus term. Thus, their experience of familiarity and lack thereof that guides knowledge judgments is likely to happen automatically and without deliberation. For people who are less knowledgeable, the experience with the bogus and real terms will be more similar because they do not have a large base of knowledge to firmly support a sense of familiarity with the real terms. They do have world knowledge that they can use to cobble together some notion of both real and bogus terms to decide whether or not they know them (Graeff et al., 2016; Koriat, 2018; van Loom de Bruin, van Gog, & van Merriënboer, 2013). Thus, their experience of real and bogus terms will be less distinguishable, making it more difficult for them to automatically recognize their ignorance of the bogus terms.

In this paper we aimed to achieve three goals. First, we sought to explore the relationship between genuine knowledge and overclaiming; in Study 1 we made use of ecologically valid group differences in expertise by recruiting medical doctors, medical students, and pre-medical students to test whether expertise was associated with less overclaiming of knowledge about bogus medical terms. Study 2 provided another test of the relationship between knowledge and overclaiming, this time in a different domain (science) with a population of lay people and using a continuous measure of knowledge (a knowledge quiz). Our second goal was to test a potential

mechanism for the impact of genuine knowledge on overclaiming. Specifically, we wanted to test what meta-cognitive cue more knowledgeable people use to arrive at a more accurate assessment of knowledge. In Study 3 we elicited participants' associations with real and bogus terms in a new domain (U.S. civics) and tested whether more knowledgeable participants overclaimed less knowledge of the bogus terms, and whether this was mediated by the *relative* number of associations they generated to the real items vs. the bogus ones. In Study 4, we tested whether more knowledgeable people used a more automatic thinking process to assess their knowledge, which would be consistent with reliance on relative familiarity as a meta-cognitive cue, and whether this mediated the effect of genuine knowledge on overclaiming.

Our third goal was to gain a better understanding of the relationship between self-perceived knowledge and overclaiming, and its interplay with genuine knowledge. To that end, in Studies 2-4 we measured people's self-perceived knowledge in addition to their genuine knowledge. This allowed us to replicate the documented positive relationship between self-perceived knowledge and overclaiming, while also testing whether self-perceived knowledge and genuine knowledge have independent effects on overclaiming.

Study 1

We focused our investigation of genuine knowledge on overclaiming, i.e., mistakenly claiming to have knowledge that one cannot have because it does not exist. Overclaiming is measured by asking participants to rate their knowledge of real and bogus items (terms, concepts, etc.) related to a particular domain (Paulhus et al., 2003). Indicating knowledge of a bogus item - called a foil – clearly marks a mistake in knowledge assessment.

In Study 1 we tested whether expertise was associated with any protection against knowledge mistakes – specifically, whether people with different levels of expertise differed in

their overclaiming. In this study we looked at the medical domain, capitalizing on naturally occurring differences in levels of expertise within this domain. We recruited three groups of participants who share an interest in medicine but vary in how much medical knowledge they have: medical doctors, medical students, and premedical students. The knowledge that medical doctors and medical students gained in the course of their training may help them differentiate between what they know and what they do not know better than premedical students who have yet to undergo that training. If so, we would expect the two groups with greater knowledge to overclaim less.

Method

Participants. Medical doctors at a university medical hospital were recruited via email. Medical students at that same university were recruited via a post on a private Facebook group only accessible to those students. Two hundred and seventy emails were sent out to doctors, and forty-nine doctors (20 women and 29 men, M_{age} = 48.81, SD_{age} = 9.98, 1 did not report age) responded. One hundred and fourteen medical students (63 women and 50 men, M_{age} = 24.78, SD_{age} = 2.34, 1 did not report gender and age) responded. An additional 14 participants did not complete the entire survey and were excluded from all analyses. For the low-expertise group, we recruited undergraduates who were interested in studying medicine through the listserv of a pre-medical student organization. Sixty-seven undergraduates (44 women and 20 men, M_{age} = 19.12, SD_{age} = 0.86, 3 did not report gender and 3 did not report age) completed the survey; of those, 67 were either pre-med (N = 47) or considering being pre-med (N = 20); these were grouped together in our analyses. Two additional respondents described themselves as neither and were excluded from all analyses. In exchange for their participation, medical students and undergraduates were entered into a lottery to win one of nine Amazon gift cards. An additional

11 participants did not complete the entire survey and were excluded from all analyses.

Procedure. The overclaiming measure consisted of 30 items related to medicine: 23 real medical terms[6] (e.g., *digoxin*, *Kawasaki disease*) and 6 foils (e.g., *Apoliosis*, *Partin's arteritis*; full list of items available in Appendix A). Participants rated their knowledge of each term by choosing the appropriate number from 1 (never heard of it) to 7 (very knowledgeable about it). The order of the items was randomized for each participant, and each participant saw the items presented one at a time. Finally, participants provided demographic information.

Results and Discussion

Overall, 50% of participants indicated they had at least some knowledge of at least one of the bogus medical terms (i.e., they rated their knowledge of one or more of the foils as higher than *1 = never heard of it*); by group, 52.2% of premedical students, 53.4% of medical students, and 38.8% of medical doctors mistakenly indicated they had at least some knowledge of at least one bogus medical term.

In this and all future studies in the paper, we computed overclaiming by averaging each participant's knowledge ratings of the foils ($M = 1.41$, $SD = .71$)[7]. We also averaged each participant's knowledge ratings of the real items ($M = 3.45$, $SD = 1.48$). An analysis of variance indicated that overclaiming differed marginally significantly between participants with different

[6] There were 24 real medical terms in the survey. However, one term was misspelled and was therefore excluded from all analyses.

[7] We have considered other ways of computing overclaiming that have been used in previous work (Paulhus et al., 2003). One option is to look at bias, which essentially averages the ratings of both the real items and the bogus items. However, genuinely knowledgeable people naturally score higher on this measure because they rate their knowledge of the real items higher. An alternative is to look at bias controlling for accuracy, but this doesn't make sense in our case because controlling for accuracy also controls for genuine knowledge. We have therefore chosen to look at the ratings of the foils, which we think operationalizes overclaiming in an intuitive way that is easy to understand. An alternative is to look at the proportion of false alarms averaged across the different cutoffs; doing so results in a different beta than the beta for the ratings of the foils but the statistical significance in identical.

levels of expertise (premedical student, medical student, medical doctor), $F(2,227) = 2.71$, $p = .069$. Follow up analyses revealed that medical doctors ($M = 1.28$, $SE = .07$) and medical students ($M = 1.36$, $SE = .05$) overclaimed to a statistically similar degree, $p = .303$, whereas premedical students overclaimed marginally more than both doctors ($M = 1.57$, $SE = .12$), $t(114) = 1.81$, $p = .073$, and medical students, $t(179) = 1.76$, $p = .081$. When medical doctors and medical students were combined into a single group ($M = 1.34$, $SE = .04$), the analysis indicated that they overclaimed significantly less than premedical students, $b = .23$, $t(228) = 2.23$, $p = .027$. These results suggest that genuine knowledge is associated with decreased overclaiming.

The relationship between level of expertise and overclaiming may plateau at high levels of expertise; even though medical students and medical doctors have different levels knowledge and expertise, they did not differ in their overclaiming to a statistically significantly degree. Alternatively, the relatively small sample size may have rendered our statistical power insufficient to detect a difference in overclaiming between medical students and doctors, who are likely closer in knowledge than they are to the premedical students.

Study 2

Study 1 provided initial support for the hypothesis that expertise in a domain is associated with less overclaiming within that domain. In Study 2 we assessed in a more controlled way the relationship between being knowledgeable and the likelihood of mistakenly claiming knowledge within that domain that one cannot have. This study allowed us to expand on the findings from Study 1 in several ways. First, we tested the hypothesis in a new domain, science. Second, we recruited a sample from the lay population, and obtained a continuous measure of their genuine knowledge in science using the National Science Foundation's scientific literacy quiz (National Science Board, 2010). Finally, we also collected people's perceptions of their own overall

knowledge in science. Given that self-perceived knowledge is related to genuine knowledge and also associated with greater overclaiming (Atir et al., 2015), the effect of self-perceived knowledge may confound and suppress that of genuine knowledge. Adding this measure allowed us to replicate previous work and, critically, to partial out the true effect of genuine knowledge on overclaiming.

Method

Participants. Participants were recruited through Amazon's Mechanical Turk and were restricted to respondents within the United States. One hundred and one people participated (43 women and 57 men, M_{age} = 35.21, SD = 12.49, 1 did not report gender or age). Four additional participants failed to complete the entire study and were excluded from all analyses.

Procedure. Participants completed measures of overclaiming, self-perceived knowledge, and genuine knowledge. The measures of overclaiming and self-perceived knowledge appeared first, in counterbalanced order across participants; genuine knowledge was always measured last.

Overclaiming was assessed using a subset of 30 science-related items from The Overclaiming Questionnaire (Paulhus et al., 2003). Participants were first presented with 15 terms related to the physical sciences, of which 12 were real terms (e.g., *Manhattan Project, nuclear fusion*) and 3 were invented (e.g., *cholarine*). The items were presented individually on the screen in a random order digitally determined for each participant. For each item, participants were asked to rate their knowledge of the term by choosing the appropriate number from 1 (never heard of it) to 7 (very knowledgeable). The procedure was then repeated with 15 items related to the life sciences (full lists available in Appendix B).

To assess self-perceived knowledge, we had participants rate their overall knowledge in science by answering the following two questions: "In general, how knowledgeable would you

say you are about science?" (from 1 = *not knowledgeable at all* to 7 = *extremely knowledgeable*) and "How would you rate your general knowledge of science compared to the average American?" (from 1 = *much less knowledgeable* to 7 = *much more knowledgeable*). The two questions assessing self-perceived knowledge were averaged (α = .91) to create a single variable.

Finally, to assess genuine knowledge in science, participants completed the National Science Foundation's 11-question scientific literacy quiz (National Science Board, 2010). Questions were multiple-choice and included, for example, the following: "All radioactivity is man made" (*Yes* or *No*). The full list of questions in available in Appendix C. Participants also provided demographic information, reported whether they had completed a similar study in the past, and, if they wished, provided a guess as to what the study was trying to test (none correctly guessed our hypothesis or described anything resembling overclaiming).

Results and Discussion

Eighty-six percent of participants indicated they had at least some knowledge of at least one of the six bogus science terms. On average, participants rated their knowledge of the bogus terms as 2.17 (*SD* = 1.01) and their knowledge of the real items as 4.60 (*SD* = .98). A single measure of genuine knowledge was obtained by summing the number correct answers on the scientific literacy quiz (*M* = 9.42 out of 11, *SD* = 1.56).

As predicted, a multiple regression analysis with genuine knowledge indicated that genuine knowledge negatively predicted overclaiming, $b = -.17$, $t(99) = -2.68$, $p = .009$. The more participants knew about science, the less likely they were to mistakenly claim knowledge of the bogus science terms. Replicating past research (Atir et al., 2015), participants' self-perception of their overall knowledge in science (*M* = 4.51, *SD* = 1.17) also predicted overclaiming, but in the opposite direction, $b = .20$, $t(99) = 2.41$, $p = .018$. That is, the more

participants viewed themselves as knowledgeable about science, the more knowledge they mistakenly claimed of the bogus science terms.

Genuine knowledge and self-perceived knowledge are related constructs that are often positively correlated (e.g., Nelson, Gerler, & Narens, 1984; Sporer, Penrod, Read, & Cutler, 1995). As a result, genuine knowledge may be capturing some variance that is explained by self-perceived knowledge. Because self-perceived knowledge exerts a positive effect on overclaiming and genuine knowledge is predicted to exert a negative effect, self-perceived knowledge may be a suppressor variable for genuine knowledge (Tzelgov & Henik, 1991). To isolate the true effect of genuine knowledge, we entered both genuine knowledge and self-perceived knowledge into a regression model simultaneously. We found that, as predicted, genuine knowledge negatively predicted overclaiming, $b = -.19$, $t(98) = -3.19$, $p = .002$ (see Figure 1), and self-perceived knowledge positively predicted it, $b = .24$, $t(98) = 2.96$, $p = .004$. Thus, genuine knowledge and self-perceived knowledge appear to have independent and opposite effects on overclaiming: genuine knowledge predicts less overclaiming, and self-perceived knowledge predicts more overclaiming.

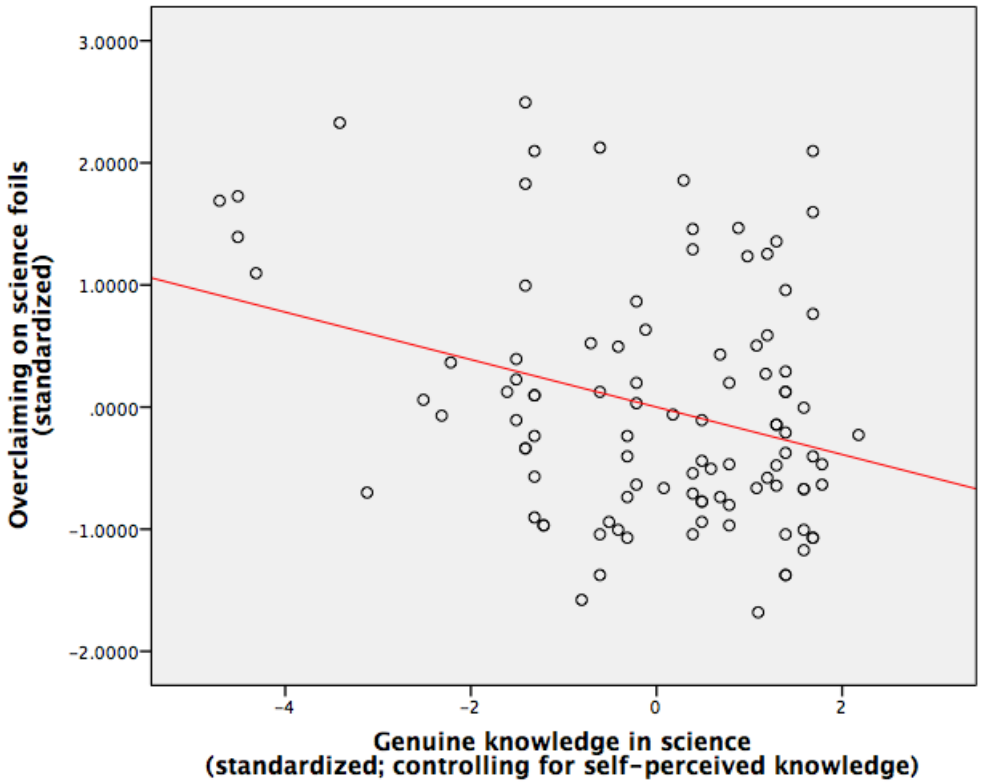

Figure 1. Overclaiming as a function of genuine knowledge after controlling for self-perceived knowledge.

Study 3

The first two studies provided evidence that possessing knowledge is associated with protection against overclaiming knowledge within that domain. We next examined what might explain this relationship. What meta-cognitive cue facilitates knowledgeable people's accurate recognition of their ignorance? We suggest that knowledgeable people have a greater gap between their experience of bogus vs. real information, which helps them identify what they don't know. Knowledgeable people are likely to experience the real items as highly fluent because the items are clearly embedded in a large store of mental content. As a result, the foils will feel relatively less fluent, and the relative lack of content associated with them will be more noticeable.

In Study 3, we elicited participants' associations to each of the items they rated, and tested whether genuine knowledge predicted a greater difference in number of associations to the foils vs. the real items. Though greater knowledge may allow participants to generate more associations to the foils based on the real information in their memory (as postulated in the introduction), we expect that the critical factor will be the *difference* between how many associations people generate about the foils and the real items. This study also provided an opportunity to replicate the effects from the previous studies in a different domain, U.S. civics. Finally, we were also able to test whether the associations could explain the positive relationship between overclaiming and self-perceived knowledge.

Method

Participants. Participants were recruited through Amazon's Mechanical Turk and were restricted to respondents within the United States. Three hundred and five people participated (154 women, 150 men, 1 did not report gender, M_{age} = 34.12, SD_{age} = 10.68). Ninety-nine additional participants failed to complete the entire study and were excluded from all analyses.

Procedure. Participants first reported their self-perceived knowledge in U.S. civics by answering two questions: "In general, how knowledgeable would you say you are about U.S. civics?" (from 1 = *not knowledgeable at all* to 7 = *extremely knowledgeable*) and "How would you rate your general knowledge of U.S. civics compared to the average American?" (from 1 = *much less knowledgeable* to 7 = *much more knowledgeable*). The two ratings were averaged to create a single measure of self-perceived knowledge (α = .89). We then elicited their associations for each of 15 items related to U.S. civics. Of these, 12 were real (e.g., *president pro tempore*, *caucus*) and 3 were foils (e.g., *implied impeachment*, *reserved pardon*; full list available in Appendix D). Participants read the following instructions for generating associations:

> In this part of the study, we are interested in what people know and think about U.S. history and civics.
> You will now see 15 items (terms, concepts, names, etc.) that are related to U.S. history and civics. For each one, please type in any associations or information you have about what that item means or refers to.
> For example, if the item were **President of the United States,** you might type in **commander in chief, Barack Obama[8], head of state,** and/or anything else that comes to your mind related to the term President of the United States.
> Separate your thoughts with a comma.
> If nothing comes to mind—no information or associations about the item—simply enter "N/A" in the textbox.

Each term was presented individually on the screen in an order randomly determined for each participant. After completing the free association task, participants completed an overclaiming measure – reporting how knowledgeable they were about each of those same 12 terms, again presented individually in an order randomly determined for each participant.

Finally, to measure their genuine knowledge, we presented participants with 25 multiple-choice questions about U.S. civics. Questions were borrowed from the American Naturalization Test and similar civics tests available online and included, for example, "Who was President during World War I?" (*Woodrow Wilson, Theodore Roosevelt, Franklin Roosevelt, Warren Harding,* or *John Adams*; correct answer: Woodrow Wilson). Full list available in Appendix E. The order of the questions was digitally randomized for each participant, as was the order of the answer choices. Participants also provided demographic information, reported whether they had taken the U.S. naturalization test, reported whether they had looked up any of the answers and, if they wished, provided a guess as to what hypothesis our research was trying to test (none correctly guessed our hypothesis or described anything resembling overclaiming).

Results and Discussion

Ninety-seven percent of participants indicated they had at least some knowledge of at

[8] When the study was conducted, Barack Obama was the U.S. President.

least one of the three bogus civics terms. On average, participants rated their knowledge of the bogus terms as 2.91 (*SD* = 1.03) and their knowledge of the real items as 4.04 (*SD* = 1.03). A single measure of genuine knowledge was obtained by summing the number correct answers on the U.S. civics quiz (*M* = 15.22 out of 25, *SD* = 1.56).

Unlike in Studies 1 and 2, genuine knowledge in U.S. civics was not a significant predictor of overclaiming on its own, $b = -.008$, $p = .52$. As in Study 2, because self-perceived knowledge is related to and positively correlated with genuine knowledge ($r = .34$, $p < .001$ in this study) and may be suppressing its effect, we entered both variables into a regression model simultaneously to partial out the effect of genuine knowledge. As predicted, genuine knowledge negatively predicted overclaiming when controlling for self-perceived knowledge, $b = -.03$, $t(302) = -2.07$, $p = .039$. In the same analysis, self-perceived knowledge (*M* = 4.21, *SD* = 1.23) positively predicted overclaiming, $b = .21$, $t(302) = 4.23$, $p < .001$. Self-perceived knowledge also positively predicted overclaiming on its own, $b = .18$, $t(303) = 3.73$, $p < .001$.

Two research assistants who were unaware of our hypothesis were provided with the participants' associations and coded the number of associations per response. Intercoder reliability was high (97.5% agreement, the remaining 2.5% of cases were resolved by the experimenter). For each participant, we created a measure of associations to each item type by averaging the number of associations generated for the foils (*M* = 1.01, *SD* = .78; $\alpha = .67$) and the real items (*M* = 1.29, *SD* = .81; $\alpha = .89$), respectively. A multiple regression analysis revealed that genuine knowledge (controlling for self-perceived knowledge) significantly predicted a greater number of associations to the real items, $b = .07$, $t(302) = 6.68$, $p < .001$, and marginally significantly predicted a greater number of associations to the foils, $b = .02$, $t(302) = 1.84$, $p = .067$. These results suggest that people who are more knowledgeable generated more

associations to both the real items and the foils. However, the number of associations in itself did not appear to be the key mediator to explain the protective effect of genuine knowledge on overclaiming; the number of associations to the foils *positively* predicted overclaiming, $b = .43$, $t(303) = 5.93$, $p < .001$, and the number of associations to the real items was in the direction of predicting less overclaiming, though the result did not reach statistical significance, $p = .15$.

Consistent with our hypothesis regarding the mechanism by which genuine knowledge protects against overclaiming, genuine knowledge positively predicted the difference between the number of associations to real items and to the foils (number of associations to the real items minus number of associations to the foils; $M = .29$, $SD = .72$), $b = .05$, $t(303) = 5.78$, $p < .001$ (the result was essentially the same when holding self-perceived knowledge constant, $b = .05$, $t(302) = 5.13$, $p < .001$). That is, more knowledgeable people had a larger difference between the number of associations they generated to the real items and the number of associations they generated to the foils. That difference was driven by the fact that having genuine knowledge allowed people to generate more associations with real items (see Figure 2). The gap between how many associations participants were able to generate for real items versus foils negatively predicted overclaiming, $b = -.37$, $t(303) = -4.60$, $p < .001$. Supporting a mediation hypothesis, genuine knowledge (controlling for self-perceived knowledge) was no longer a significant predictor of overclaiming when controlling for the difference in associations, $p = .52$. Indeed, a mediation analysis found that the negative relationship between genuine knowledge and overclaiming was mediated by the increased gap between the number of associations generated to the real items and the foils, $ab = -.02$, BCa CI [-.03, -.01]. This result suggests that genuinely knowledgeable people may be better able to differentiate between real and bogus terms based on the difference in the amount of related content that comes to mind for each one. They are

consequently less likely to mistakenly indicate they have knowledge they cannot, in fact, have.

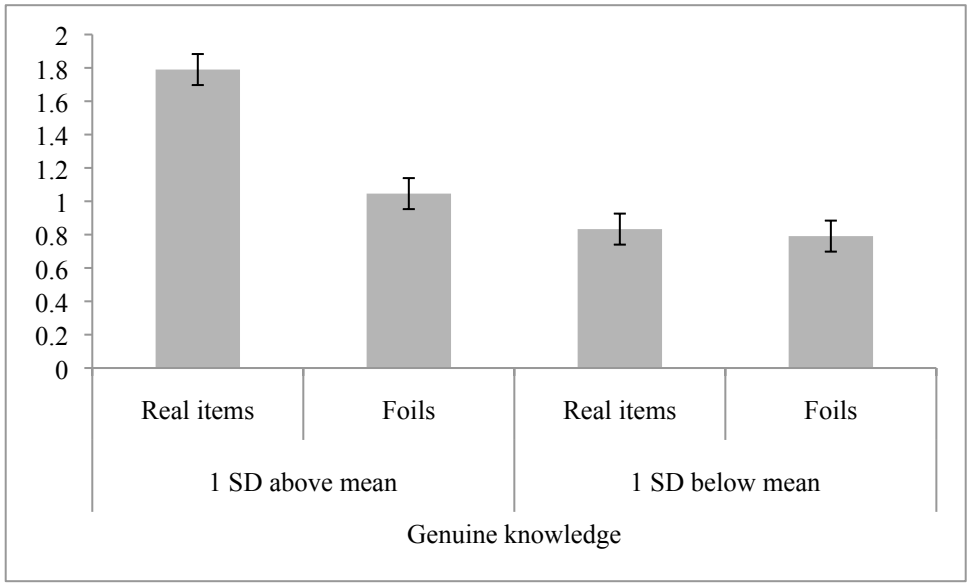

Figure 2. Average number of associations about the items as a function of level of genuine knowledge and item type[9].

We also examined whether the number of associations people generated could help shed light on the positive relationship between self-perceived knowledge and overclaiming. We found that self-perceived knowledge (controlling for genuine knowledge) was not significantly related to the difference between the number of associations to real items and to the foils, $p = .38$, which suggests that people who only view themselves as knowledgeable are no better or worse at differentiating between the foils and the real items, in terms of the associations they generated, than people who view themselves as less knowledgeable. We also found no evidence that self-perceived knowledge (controlling for genuine knowledge) was related to the number of associations people generated to the real items, $p = .14$, or to the foils, .52. Thus, while we consistently find that people who only view themselves as knowledgeable are more likely to

[9] The figure looks similar when controlling for self-perceived knowledge.

overclaim, we find no evidence that people's associations to the items can account for this link.

Study 4

The results of Study 3 suggest that more knowledgeable people may use the sharper contrast between their experiences with the foils vs. the real items as a meta-cognitive cue for determining their knowledge. If highly knowledgeable people are using the contrast in their experience between real items and foils as a meta-cognitive cue, they are likely approaching the items more automatically and heuristically – a processing style associated with the use of meta-cognitive cues. To test for this, In Study 4 we asked participants to report the thinking process that guided their response after rating each item – whether it was deliberative or automatic. We predicted that more knowledgeable people would use a more automatic thinking process to assess their knowledge of the foils, and that his would be associated with less overclaiming. This result would also be consistent with previous work showing that accuracy in certain judgments is associated with automatic thinking and speed (Dunning & Perretta, 2002; Dunning & Stern 1994; Pettibone, 2012; Zhu, Ritter, Muller, & Dijksterhuis, 2017). This study also allowed us to test whether thinking process is related to the relationship between self-perceived knowledge and overclaiming; people who view themselves as knowledgeable may employ a more deliberative thinking process to determine their knowledge of the foils, thereby creating a false sense of familiarity.

Method

Participants. Participants were recruited through Amazon's Mechanical Turk and were restricted to respondents within the United States. Three hundred and fifty two people participated (187 women, 164 men, 1 did not report gender, M_{age} = 37.37, SD_{age} = 12.02). Fifteen additional participants failed to complete the entire study and were excluded from all analyses.

Procedure. Participants first reported their self-perceived knowledge in science by answering the same questions as in Study 2. These two ratings were averaged to create a single measure of self-perceived knowledge *(*α = .88). We then measured their overclaiming in the same way described in Study 2 with one important difference: after rating their knowledge of each science item, participants were asked, on a new page, to describe their overall thinking process as they were making their knowledge judgment, on a scale ranging from 1 = *very automatic; the answer just came to me without any thinking* to 5 = *very deliberative; I devoted some effort and time thinking before answering*. As a measure of their genuine knowledge in science, participants completed the same science literacy quiz described in Study 2. Finally, participants provided demographic information

Results and Discussion

Overall, 88.1% of participants indicated they had at least some knowledge of at least one of the six science foils. On average, participants rated their knowledge of the foils as 2.17 (*SD* = 1.14) and their knowledge of the real items as 4.75 (*SD* = 1.05). A single measure of genuine knowledge was obtained by summing the number correct answers on the scientific literacy quiz (*M* = 9.55 out of 11, *SD* = 1.74).

Replicating the results of the previous studies, genuine knowledge in science negatively predicted overclaiming, $b = -.22$, $t(350) = -6.75$, $p < .001$, and self-perceived knowledge in science (*M* = 4.60, *SD* = 1.13) positively predicted overclaiming, $b = .27$, $t(350) = 5.16$, $p < .001$. As in the preceding studies, we next entered both variables into the regression model simultaneously to isolate the true effect of genuine knowledge. As predicted, genuine knowledge negatively predicted overclaiming, $b = -.26$, $t(349) = -8.30$, $p < .001$, and self-perceived knowledge positively predicted overclaiming (*M* = 2.17, *SD* = 1.14), $b = .34$, $t(349) = 7.00$, $p <$

.001.

We averaged participants' ratings of their thinking process across the six foils to obtain a single measure of their thinking process about the foils (α = .87; M = 2.79, SD = 1.17). Consistent with our hypothesis, genuine knowledge marginally significantly predicted more automatic thinking about the foils, b = .07, $t(349)$ = 1.89, p = .06 (this result was significant when controlling for self-perceived knowledge, b = .10, $t(348)$ = 2.70, p = .007). We next explored the relationship between genuine knowledge, thinking process about the foils, and overclaiming. We found that thinking automatically about the foils predicted lower overclaiming, b = -.31, $t(349)$ = 6.30, $p < .001$; people who thought about the foils more automatically were less likely to indicate that they had knowledge of them. Put differently, thinking about the foils in a more deliberative fashion was associated with finding the bogus science terms more familiar. Deliberation may have prompted participants to work harder to generate associations that seemed related to the foils, which (as we have seen in Study 3) may have induced a false sense of knowledge. We next tested whether participants' thought process mediated the relationship between genuine knowledge and overclaiming; as predicted, there was a significant negative indirect effect of genuine knowledge on overclaiming through automatic thinking about the foils, ab = -.02, BCa CI [-.04, -.001] (the result was similar when controlling for self-perceived knowledge, ab = -.02, BCa CI [-.04, -.007]). These results suggest that more knowledgeable people think about the foils in a more automatic fashion, which is in turn associated with reduced overclaiming.

The pattern was different for self-perceived knowledge. Self-perceived knowledge predicted more *deliberative* thinking about the foils, b = .22, $t(349)$ = 4.12, $p < .001$. That is, people who only viewed themselves as knowledgeable thought about the foils in a more

deliberative fashion. We next tested whether deliberativeness mediated the positive relationship between self-perceived knowledge and overclaiming; we found that indeed, more deliberative thinking about the foils mediated this relationship, *ab* = .06, BCa CI [.03, .10]. The result suggests that people who only see themselves as knowledgeable deliberate more about the foils, perhaps spending time bringing to mind more content that feels related to the bogus term, and this leads them to feel more knowledgeable about the bogus terms.

In sum, the thinking process people use to assess their knowledge of the foils mediates the opposite effects of genuine knowledge and self-perceived knowledge on overclaiming. Genuine knowledge was associated with more automatic thinking, which in turn was associated with better recognition of the foils as unfamiliar. This suggests that knowledgeable people's superior ability to differentiate between the real and bogus terms occurs automatically. Self-perceived knowledge, on the other hand, is associated with more deliberative thinking, which may lead to cognitive expounding on the bogus items, and therefore to an increased sense of knowledge about the foils.

General Discussion

Recognizing one's own ignorance can be both important and difficult. Across four studies in three different domains, we find evidence that knowledgeable people may have an advantage at telling apart what they know from what they do not. In Study 1, we found that true medical experts – specifically, medical doctors and medical students – were less likely to mistakenly believe they had knowledge of bogus medical terms than did premedical students who had yet to be trained. In Study 2, we found that genuine knowledge and self-perceived knowledge diverge in their relationship with overclaiming; lay people's genuine knowledge in science predicted less overclaiming within that domain, whereas people's perceptions of their

knowledge in science predicted greater overclaiming.

In Study 3, we explored the associations that came to people's minds when they thought about real items and foils in the domain of U.S. civics; the results suggested that more knowledgeable people in that domain had a larger gap between the number of associations they came up with for the foils vs. the real items. The protective effect of genuine knowledge on mistaken claims of knowledge was mediated by this larger difference. Knowledgeable people may rely on relative familiarity as a cognitive cue to assess their knowledge. Finally, in Study 4, we found that more genuinely knowledgeable people thought about the foils and the real items in a more automatic fashion, which is consistent with reliance on a cognitive cue. Automatic thinking also mediated knowledgeable people's relatively lower rate of mistaken recognition of the foils. People who only viewed themselves as knowledgeable, on the other hand, thought about the items in a more deliberative way, and this mediated their greater tendency to find bogus terms familiar.

People rely on various meta-cognitive cues as mental shortcuts to assess their own knowledge (Koriat & Levy-Sadot, 2001); in this work, we find that one such cue may not be just how many associations people have with an object of judgment, but how that number compares with the associations they have with other items in the same domain. If that relative number is low, this can be a valuable cue to correctly determining one's own lack of knowledge, and to avoiding overclaiming.

It may seem intuitively true that careful deliberation always reduces mistakes and improves the accuracy of knowledge judgments. Doubtlessly, further contemplation can curtail some mistakes (e.g., Finucane, Alhakami, Slovic, & Johnson, 2000). However, as evidenced by the results of Study 4, more deliberation may also, under some circumstances, be associated with

more error in meta-knowledge judgment. This may occur because deliberation can stoke an initial false sense of familiarity. In the case of overclaiming, perhaps further deliberation leads down the wrong path and creates an even stronger sense of knowledge where, in fact, none exists.

Across the four studies, we consistently find that genuine knowledge is associated with less overclaiming. In future work, it will be important to experimentally manipulate genuine knowledge to test whether gaining knowledge can decrease overclaiming. It will also be interesting to explore more closely the contributions of genuine knowledge and self-perceived knowledge to overclaiming. As genuine knowledge increases, so does self-perceived knowledge. Given the evidence that both of these factors are associated with opposing effects on overclaiming, it remains to be determined what their relative power is in predicting overclaiming. For instance, when happens in the course of learning, as both genuine knowledge and self-perceived knowledge increase? Is the net effect to make it easier to tell apart what one knows from what one doesn't, or to exacerbate the difficulty of this assessment? Future research should investigate these questions.

References

Atir, S., Rosenzweig, E., & Dunning, D. (2015). When knowledge knows no bounds: Self-perceived expertise predicts claims of impossible knowledge. *Psychological Science*, *26*, 1295–1303.

Barber, B. M., & Odean, T. (2000). Trading is hazardous to your wealth: The common stock investment performance of individual investors. *The journal of Finance*, *55*, 773-806.

Barber, L. K., Barnes, C. M., & Carlson, K. (2013). Random and systematic error effects of insomnia on survey behavior. *Organizational Research Methods*, *16*, 616–649.

Benjamin, A. S., Bjork, R. A., & Schwartz, B. L. (1998). The mismeasure of memory: when retrieval fluency is misleading as a metamnemonic index. *Journal of Experimental Psychology: General*, *127*(1), 55.

Bradley, J. V. (1981). Overconfidence in ignorant experts. *Bulletin of the Psychonomic Society*, *17*(2), 82-84.

Cantor, A. D., & Marsh, E. J. (2017). Expertise effects in the Moses illusion: detecting contradictions with stored knowledge. *Memory, 25*, 220-230.

Christensen-Szalanski, J. J., & Bushyhead, J. B. (1981). Physicians' use of probabilistic information in a real clinical setting. *Journal of Experimental Psychology: Human perception and performance*, *7*, 928.

Critcher, C. R., & Dunning, D. (2009). How chronic self-views influence (and mislead) self-assessments of task performance: Self-views shape bottom-up experiences with the task. *Journal of Personality and Social Psychology*, *97*, 931–45.

Dijksterhuis, A., Spears, R., Postmes, T., Stapel, D., Koomen, W., Knippenberg, A. V., & Scheepers, D. (1998). Seeing one thing and doing another: Contrast effects in automatic behavior. *Journal of Personality and Social Psychology*, *75*, 862-871.

Dunning, D., & Perretta, S. F. (2002). Automaticity and eyewitness accuracy: A 10- to 12-second rule for distinguishing accurate from erroneous positive identifications. *Journal of Applied Psychology, 87*, 951-962.

Dunning, D., & Stern, L. B. (1994). Distinguishing accurate from inaccurate eyewitness identifications via inquiries about decision processes. *Journal of Personality and Social Psychology, 67*, 818-835.

Finucane, M. L., Alhakami, A., Slovic, P., & Johnson, S. M. (2000). The affect heuristic in judgments of risks and benefits. *Journal of behavioral decision making, 13*, 1-17.

Gervais, S., & Odean, T. (2001). Learning to be overconfident. *Review of Financial Studies, 14*, 1–27.

Graeff, T. R. (2003). Exploring consumers' answers to survey questions: Are uninformed responses truly uninformed? *Psychology & Marketing, 20*, 643–667.

Graeff, Schuman & Presser, Marsh, E. J., Cantor, A. D., & Brashier, N. M. (2016). Believing that humans swallow spiders in their sleep: False beliefs as side effects of the processes that support accurate knowledge. *Psychology of Learning and Motivation, 64*, 93–132.

Koriat, A. (1995). Dissociating knowing and the feeling of knowing: Further evidence for the accessibility model. *Journal of Experimental Psychology: General*, *124*, 311-333.

Koriat, A. (2018). When reality is out of focus: Can people tell whether their beliefs and judgments are correct or wrong? *Journal of Experimental Psychology: General, 147,* 613-631.

Koriat, A., & Levy-Sadot, R. (2001). The combined contributions of the cue-familiarity and accessibility heuristics to feelings of knowing. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 27*, 34–53.

Mannes, A. E., & Moore, D. A. (2013). A behavioral demonstration of overconfidence in judgment. *Psychological science*, *24*, 1190-1197.

Meade, A. W., & Craig, S. B. (2012). Identifying careless responses in survey data. *Psychological Methods*, *17*, 437–455.

National Science Board. 2010. *Science and Engineering Indicators 2010*. Arlington, VA: National Science Foundation (NSB 10-01).

Nathanson, C., Westlake, B., & Paulhus, D. L. (2007). Controlling response bias in the measurement of consumer knowledge. Presented at *The Meeting of the Association for Psychological Science, Washington, DC.*

Nelson, T. O., Gerler, D., & Narens, L. (1984). Accuracy of feeling-of-knowing judgment for predicting perceptual identification and relearning. *Journal of Experimental Psychology: General, 113,* 282-300.

Paulhus, D. L., Harms, P. D., Bruce, M. N., & Lysy, D. C. (2003). The over-claiming technique: Measuring self-enhancement independent of ability. *Journal of Personality and Social Psychology*, *84*(4), 890-904.

Pettibone, J. C. (2012). Testing the effect of time pressure on asymmetric dominance and compromise decoys in choice. *Judgment and Decision Making, 7,* 513.

Rozenblit, L., & Keil, F. C. (2002). The misunderstood limits of folk science: an illusion of explanatory depth. *Cognitive Science*, *26*, 521-562.

Schuman, H., & Presser, S. (1980). Public opinion and public ignorance: The fine line between

attitudes and nonattitudes. *The American Journal of Sociology, 85*, 1214–1225.

Sporer, S. L., Penrod, S., Read, D., & Cutler, B. (1995). Choosing, confidence, and accuracy: A meta-analysis of the confidence-accuracy relation in eyewitness identification studies. *Psychological Bulletin*, *118*, 315.

Tzelgov, J., & Henik, A. (1991). Suppression situations in psychological research: Definitions, implications, and applications. *Psychological Bulletin*, *109*, 524-536.

van Loon, M. H., de Bruin, A. B. H., van Gog, T., & van Merriënboer, J. J. G. (2013). Activation of inaccurate prior knowledge affects primary-school students' metacognitive judgments and calibration. *Learning and Instruction, 24*, 15–25.

Whittlesea, B. W., Jacoby, L. L., & Girard, K. (1990). Illusions of immediate memory: Evidence of an attributional basis for feelings of familiarity and perceptual quality. *Journal of Memory and Language, 29,* 716-732.

Zhu, Y., Ritter, S. M., Müller, B. C., & Dijksterhuis, A. (2017). Creativity: Intuitive processing outperforms deliberative processing in creative idea selection. *Journal of Experimental Social Psychology*, *73*, 180-188.

Appendix A: Study 1 Stimuli

Foils indicated by *

xerostomia

cerumen

shock wave lithotripsy

hyperventilation syndrome

familial Mediterranean fever

graft-versus-host disease

histiocyte

grand multipara

nephritis, lipomatous

echovirus

dysphagia

Klumpke's palsy

Cushing's syndrome

prolactinoma

type 1 diabetes mellitus

midsystolic click

somatic symptom disorder

digoxin

erythrocyte enzyme deficiency

wegener granulomatosis

schistosomiasis

Kawasaki disease

Meckel's diverticulum

Apoliosis*

Partin’s arteritis*

mebamectin*

acute diatresia*

acromicrola syndrome*

pesothelioma*

Appendix B: Studies 2, 4 Stimuli

Foils indicated by *

Physical sciences

Manhattan Project

planets

nuclear fusion

cholarine*

atomic number

hydroponics

alloy

plate tectonics

photon

ultra-lipid*

centripetal force

plates of parallax*

nebula

particle accelerator

satellite

Life sciences

See Appendix B in Chapter II.

Appendix C: Science Literacy Quiz

Questions from the National Science Board quiz (2010).

1. The center of the earth is very hot. (True, False)
2. All radioactivity is man-made. (True, False)
3. Lasers work by focusing sound waves. (True, False)
4. Electrons are smaller than atoms. (True, False)
5. The continents have been moving their location for millions of years and will continue to move. (True, False)
6. Does the Earth go around the Sun, or does the Sun go around the Earth? (Earth around Sun, Sun around Earth)
7. It is the father's gene that decides whether the baby is a boy or a girl. (True, False)
8. Antibiotics kill viruses as well as bacteria. (True, False)
9. A doctor tells a couple that their genetic makeup means that they've got one in four chances of having a child with an inherited illness.

   Does this mean that if their first child has the illness, the next three will not? (Yes, No)
10. Does this mean that each of the couple's children will have the same risk of suffering from the illness? (Yes, No)
11. Two scientists want to know if a certain drug is effective against high blood pressure.

    The first scientist wants to give the drug to 1,000 people with high blood pressure and see how many of them experience lower blood pressure levels.

    The second scientist wants to give the drug to 500 people with high blood pressure and not give the drug to another 500 people with high blood pressure, and see how many in

both groups experience lower blood pressure levels.

Which is the better way to test this drug? (First way, Second way)

Appendix D: Study 4 Stimuli

Foils indicated by *

filibuster

caucus

pocket veto

constituents

majority rule

party whip

congressional veto*

reserved pardon*

ex post facto law

implied impeachment*

jury duty

amendment

floor leader

president pro tempore

civil law

Appendix E: U.S. Civics Quiz

Correct answers in bold font.

1. **What do we call the first ten amendments to the Constitution?**

the Articles of Confederation
the inalienable rights
the Declaration of Independence
the Bill of Rights
the navigation acts

2. **What did Susan B. Anthony do?**

made the first flag of the United States
the first woman elected to the House of Representatives
fought for women's rights
founded the Red Cross
first female Speaker of the House
wrote about her life as a pioneer out west

3. **The Federalist Papers supported the passage of the U.S. Constitution. Choose one of the writers.**

James Madison
John Adams
Thomas Jefferson
George Washington
Benjamin Franklin
James Monroe

4. **How many justices are on the Supreme Court?**

eleven (11)
twelve (12)
ten (10)
nine (9)
three (3)
five (5)

5. **The House of Representatives has how many voting members?**

four hundred thirty-five (435)

one hundred (100)
four hundred forty-one (441)
two hundred (200)
two hundred fifty (250)
four hundred fifty (450)

6. **If both the President and the Vice President can no longer serve, who becomes President?**

the Secretary of State
the Speaker of the House
the President Pro Tempore
the Secretary of the Treasury
a new President is elected by the House
the Chief Justice

7. **Which of the following CAN'T green card holders do? Choose all correct answers.**

serve on a jury
vote in federal elections
run for government office
get an education in a state school

donate to a political party
join the military
pay taxes

8. **Why does the flag have 13 stripes?**

because the stripes represent the members of the Second Continental Congress
because the stripes represent the original colonies
because it was considered lucky to have 13 stripes on the flag
because the stripes represent the number of signatures on the U.S. Constitution
because the stripes represent the "patriots of '76" who died fighting the British
because the stripes represent the Founding Fathers
because the stripes represent major rivers of the United States

9. **Which of the following things did Abraham Lincoln do?**

declared war on Great Britain
purchased Alaska
preserved the Union
established the United Nations
fought in World War I
won U.S. independence from the British

wrote the pamphlet, "Common Sense"

10. Choose <u>one</u> U.S. territory.

Bermuda
Guam
Haiti
Cayman Islands
District of Columbia
Dominican Republic
Panama canal

11. How many U.S. Senators are there?

fifty (50)
four hundred thirty-five (435)
one hundred (100)
fifty-two (52)
one hundred two (102)
two hundred (200)

12. Which of the following is/are rights or freedoms from the First Amendment?

Choose all correct answers.

right to a trial by jury
right to bear arms
right to vote
freedom of speech
freedom of religion
freedom of the press

13. What is the supreme law of the land?

the Declaration of Independence
the Constitution
the Articles of Confederation
the Emancipation Proclamation
the bill of rights
state law

14. When was the Constitution written?

1787
1790
1789
1776

1781
1788

15. When was the Declaration of Independence adopted?

March 4, 1789
July 4, 1789
July 4, 1776
December 7, 1787
January 14, 1784
September 17, 1787
June 5, 1789

16. We elect a U.S. Representative for how many years?

six (6)
two (2)
four (4)
eight (8)
It’s a lifelong appointment

17. How many amendments does the Constitution have?

ten (10)
twenty-seven (27)
twenty-three (23)
twenty-one (21)
twenty-eight (28)

18. We elect a U.S. Senator for how many years?

ten (10)
four (4)
six (6)
two (2)
they are appointed, not elected

19. Who was President during World War I?

Theodore Roosevelt
Franklin Roosevelt
Woodrow Wilson
Warren Harding
John Adams

20. What stops one branch of government from becoming too powerful?

the filibuster

the people

the President

the separation of powers

21. Who did the United States fight in World War II?

the British Empire

Germany, Japan, and Austria-Hungary

Germany, Japan and Italy

Germany, Japan and the Soviet Union

22. Who is the President pro tempore?

the highest-ranking senator

the House minority leader

a member of the President's Cabinet

the Vice President

23. Which future president was among the delegates who signed the Constitution?

Benjamin Franklin

George Washington

John Adams

Thomas Jefferson

24. What are the two parts of the U.S. Congress?

the committee and the subcommittee

the legislative and the judicial

the Lower and Upper House

the Senate and House of Representatives

25. There were 13 original states. Choose three.

Delaware, Georgia, and Iowa

New Jersey, Connecticut, and Washington DC

New York, Pennsylvania, and Rhode Island

South Carolina, Florida, and Alabama

CHAPTER IV

HOW GENDER DETERMINES THE WAY WE SPEAK ABOUT PROFESSIONALS[10]

Abstract

Gender inequality persists in many professions, particularly in high-status fields such as science, technology, engineering, and math (STEM). We report evidence of a new form of gender bias that may contribute to this state: gender influences the way people speak about professionals. When discussing professionals or their work, it is common to refer to them by surname alone (e.g., "Darwin developed the theory of evolution"). We present evidence that people are more likely to refer to male than female professionals in this way. This gender bias emerges in archival data across domains; students reviewing professors online and pundits discussing politicians on the radio are more likely to use surname when speaking about a man (vs. a woman). Participants' self-reported references also indicate a preference for using surname when speaking about male (vs. female) scientists, authors, and others. Finally, experimental evidence provides convergent evidence: participants writing about a fictional male scientist are more likely to refer to him by surname than participants writing about an otherwise-identical female scientist. We find that, on average, people are over twice as likely to refer to male than female professionals by surname. Critically, we identified consequences of this gender bias in speaking about professionals. Researchers referred to by surname are judged as more famous and eminent. They are consequently seen as higher-status and more deserving of eminence-related

[10] Published as: Atir, S., & Ferguson, M. J. (2018). How gender determines the way we speak about professionals. *Proceedings of the National Academy of Sciences, 115*, 7278-7283.

benefits and awards. For instance, scientists referred to by surname were seen as 14% more deserving of a National Science Foundation career award.

Keywords: gender, bias, reference

Despite significant strides over the past decades, gender inequality in professional contexts persists. Men still outearn women in the U.S. (Proctor, Semega, & Kollar, 2016), and women remain underrepresented in many high-status professional fields, including science, technology, engineering, and math (STEM; National Science Board, 2016). Potentially contributing to this unequal state is gender bias in implicit and explicit forms (Mo, 2015; Nosek et al., 2009; Leslie, Cimpian, Meyer, & Freeland, 2015; Riffkin, 2014; Moss-Racusin, Dovidio, Brescoll, Graham, & Handelsman, 2012). The present work provides evidence of a novel form of gender bias that manifests in the way people refer to professionals when speaking about them and identifies the consequences of this bias.

In many countries, it is common to refer to professionals in certain fields by surname alone when speaking about them or their work; for example, scientists, politicians, authors, and others are frequently referred to by surname (e.g., Einstein, Obama, Hemingway). Might the gender of the professional influence the speaker's choice to refer to her or him by surname? And does this choice, in turn, have consequences for how others perceive the professional? Previous research hints at a gender bias in the use of surname references: first, during the 2008 Democratic primary in the U.S., television newspeople were more likely to refer to Barack Obama than to Hillary Clinton by surname (Uscinski & Goren, 2011). However, this difference may be explained by Hillary Clinton's more frequent use of her first name in her campaign, possibly as a way of distinguishing herself from her husband. Second, qualitative, descriptive work in sociology suggests that sports commentators are more likely to refer to male (vs. female) players by surname (Messner, Duncan, & Jensen,1993; Halbert & Latimer, 1994*)*. In the current work, over a series of eight studies, we test whether this gender bias exists and examine its consequences.

Across the first four studies, using archival and experimental approaches, we find that people more commonly refer to men than to women by surname in a variety of fields, including the academic, political, and scientific. The results of four additional experiments reveal the consequences of this gender bias. Specifically, professionals who are referred to by surname are perceived as more famous and eminent. When fame is brought to mind, such professionals enjoy an advantage with regards to eminence-related benefits; they are judged to hold higher status, to be more likely to win an award for their work, and to be more deserving of awards such as the National Science Foundation (NSF) career award and its associated funding.

Note that we use "surname" to mean a reference by surname *alone*, without a first name or a professional or common title (e.g., Dr., Ms.). In this work we focus on third-person references (speaking *about* a target) rather than forms of address (speaking *to* a target), and, for the sake of simplification, examine only cases in which the speaker is not personally acquainted with the target.

Study 1

In Study 1 we tested whether people are more likely to refer to men than women by surname in the academic domain. Data were obtained from the website *Rate My Professors*, which allows students to rate and review their professors (e.g., "I love [surname redacted]'s lectures. He's a funny guy"). Data were collected for all professors in five departments (biology, psychology, computer science, history, and economics) from fourteen universities chosen for their academic and geographic diversity. For each of the 4,494 comments that included a reference to the professor, we recorded the professor's gender, ratings (helpfulness, clarity, course interest, and course ease) and the form of reference used to refer to the professor in the

review. Specifically, surname was contrasted with any other form of reference: full name, first name only, Prof/Dr. full name/surname, Mr./Ms/Mrs./Miss surname, or other.

Supporting our prediction, students were 55.9% more likely to refer to a male versus female professor by surname, $\chi^2 = 35.76$, $p < .001$ (Figure 3). The gender bias in use of surname remained significant when controlling for university, department, and year in which the comment was posted, $\chi^2 = 34.13$, $p < .001$ (for this and all studies, see the methods and materials section and the appendices for more information). The gender bias was not explained by differences in students' favorability toward male and female professors or their courses; when the ratings of the professor's helpfulness, clarity, course ease, and course interest were added to the model, the professor's gender remained a significant predictor of use of surname, $\chi^2 = 43.46$, $p < .001$. We also assessed professor seniority using the difference in years between the oldest and most recent reviews. Though this is a very rough index of seniority, it positively predicted use of surname alone, $b = .033$, $\chi^2 = 8.13$, $p = .004$. However, gender remained a significant predictor of surname use controlling for seniority, $\chi^2 = 33.55$, $p < .001$. We also examined each of the five departments separately. Within each department, students were numerically more likely to refer to male than female professors by surname; this difference was statistically significant in psychology (*M*s = 39.7% vs. 21.7%, $\chi^2 = 15.7$, $p < .001$), history (*M*s = 31.1% vs. 23.0%, $\chi^2 = 4.03$, $p = .045$), and computer science (*M*s = 48.4% vs. 18.1%, $\chi^2 = 16.7$, $p < .001$), statistically directional in biology (*M*s = 32.6% vs. 18.1%, $\chi^2 = 2.6$, $p = .106$), and did not reach statistical significance in economics (*M*s = 52.1% vs. 48.6%, $p = .52$).

We next tested whether even male professors who were perceived as having more feminine traits were less likely to be called by their surname. Indeed, controlling for gender, reviews that contained a stereotypically female trait (pretty, cute, helpful, understanding, kind,

supportive, emotional, and meek; adjectives chosen based on previous research, Cejka & Eagly, 1999; Cuddy, Fiske & Glick, 2008; Langford & MacKinnon, 2000*)* were less likely to contain a surname reference, *M*s = .32 (*SE* = .012) vs. .25 (*SE* = .019), $\chi^2 = 13.21$, $p < .001$. The converse was true as well; controlling for gender, reviews that contained a stereotypically male trait (analytical, easygoing, brilliant, tough, arrogant, rude; Cejka & Eagly, 1999; Cuddy et al., 2008; Langford & MacKinnon, 2000) were more likely to contain a surname reference, *M*s = .38 (*SE* = .034) vs. .31 (*SE* = .012), $\chi^2 = 4.25$, $p = .031$. These results provide convergent evidence for the importance of gender in differentiating surname usage.

Study 2

Does this gender bias extend beyond the academic domain? We hypothesized that when discussing politics, pundits and other commentators would more commonly refer to male than female politicians by surname. Data were obtained from transcripts of the following popular, politically diverse American radio programs that regularly discuss current events: *All Things Considered*, *Fresh Air*, *Morning Edition*, *The Rush Limbaugh Show*, and *The Sean Hannity Show*. Overall, 9,572 references from 336 segments from 2014 and 2015 were coded. Speakers included the shows' hosts and various guests and correspondents, and the targets included mainly politicians, as well as other individuals connected with the relevant news story.

Consistent with our hypothesis, speakers were more than twice as likely (126.42%) to use a surname when speaking about a man than when speaking about a woman (Figure 1), $z = -4.60$, $p < .001$, *OR* = .21 [0.11, 0.41]. The same pattern emerged excluding references to Hillary Clinton, whose campaign slogans often referred to her by first name (e.g., *Hillary for President*), suggesting that the result was not driven solely by references to her, $z = -4.66$, $p < .001$, *OR* = .19

[0.10, 0.39]. The result also remained significant when controlling for speaker gender and for the political affiliation of both target and speaker, $z = -4.51$, $p < .001$, $OR = .21$ [0.11, 0.41].

Study 3

In Study 3 we investigated the gender bias in surname use in a broader range of domains and using a different design and sample. One hundred and ninety participants[11] were shown two lists of well-known individuals in counterbalanced order. One list consisted of figures in American politics (e.g., *Susan Rice, Carly Fiorina, Joe Biden, Antonin Scalia*), and the second of well-known figures in various non-political domains, including literature, science, and sports (e.g., *Jane Austen, Charles Dickens, Carl Sagan, Marie Curie*). Each list included an equal number of women and men roughly matched on average in terms of age, years active, position, and profession. Participants were asked to consider how they refer to each figure when talking about her or him in casual conversation, and then estimate the percentage of time they refer to each individual by surname, full name, first name, or by some other form of reference, adding up to 100% (if participants did not who the figure was, they were instructed to choose N/A instead). We found that participants were 74.18% more likely to report using a surname when referring to male than female figures, $\chi^2 = 223.62$, $p < .001$ (Figure 3). We found no evidence that the effect differed depending on participant gender or political affiliation, ps > .34.

Participants' estimates suggest that the target's gender may influence the form of reference they use. However, we tested several alternative explanations: people may perceive the men on our list to hold more influential positions than the women, be better known, have more distinctive surnames or less distinctive first names, and/or be less likely to share their surname with a well-known family member. Any of these may in turn increase use of a surname reference

[11] Unless noted otherwise, participants in all studies were U.S.-based Amazon Mechanical Turk workers.

and thus may account for the result without directly implicating gender. To test these accounts, 217 participants in a new sample were randomly assigned to provide one of the following ratings about the figures used in Study 3: how well known each figure was, how distinctive each first name was, or how distinctive each surname was. An additional group of 44 students at Cornell University rated how influential each political figure's position was (e.g., *attorney general, governor*).

We found that targets who were better known and whose positions were judged to be more influential were more likely to be referred to by surname, $\chi^2_{\text{well known}} = 148.35$, $p < .001$, $\chi^2_{\text{influential}} = 10.29$, $p = .001$. Targets whose first names and surnames were perceived to be more common were more likely to be referred to by surname, $\chi^2_{\text{first name}} = 9.37$, $p = .002$, $\chi^2_{\text{surname}} = 29.81$, $p < .001$. Critically, when these variables were added to the model (either individually or simultaneously), gender remained a significant predictor of surname use, $\chi^2_{\text{full model}} = 49.19$, $p < .001$. Finally, gender remained a significant predictor of surname use when we excluded women who shared a surname with a well-known family member (e.g., Hillary Clinton), $\chi^2 = 216.30$, $p < .001$, as well as when we both excluded these figures and added the previously discussed variables into the model, $\chi^2_{\text{full model}} = 47.20$, $p < .001$

Study 4

In Study 4 we sought to experimentally test whether people are more likely to use a surname to refer to a man than a woman even when gender is the only dimension on which they differ. Participants ($N = 184$) read information about a scientist, presented in bullet-point form (e.g., *Chemist and X-ray crystallographer, Years: July 25th, 1920 – April 16th, 1958*). The scientist's name, which appeared in bold font at the top, was either *Dolores Berson* (female condition) or *Douglas Berson* (male condition), determined randomly. All other information was

identical between gender conditions. Participants were asked to rewrite the bullet points in full sentences, incorporating all the information. We also tested whether the gender bias in use of surname might be limited to specific types of interpersonal exchanges; for instance, whether it would be eliminated in formal contexts, when people potentially employ more rigid rules with regards to forms of reference. To test this possibility, participants were randomly assigned to imagine that they were either lecturing about the scientist (formal expression condition) or telling a friend about the scientist in casual conversation (casual expression condition).

Supporting our predictions, participants writing about a male scientist were more than four times as likely (309.84%) to refer to him by surname than were participants writing the same information about a female scientist, $\chi^2 = 11.19$, $p < .005$ (Figure 3). This pattern did not significantly differ between participants who were writing formally and those writing casually, $p = .43$. We found no evidence that male and female participants differed in their likelihood of exhibiting the gender bias, $p = .93$.

In Studies 1-4 we found that people are less likely to refer to women by surname, but how *do* people refer to women? There was no single form of reference that was consistently applied to women more often than to men. Students (Study 1) more commonly referred to women (vs. men) by a common title (Ms., Mrs., Miss, and Mr.; one potential explanation is that female instructors were less likely to hold a PhD and were therefore more often referred to by a common title), *M*s = .06 (*SE* = .01) vs. .01 (*SE* = .003), $\chi^2 = 16.33$, $p < .001$ and by "Prof. [surname]" (see Cowan & Kasen, 1984; Files et al., 2017; Takiff, Sanchez, & Stewart, 2001 for work on gender differences in use of professional titles), *M*s = .28 (*SE* = .02) vs. .22 (*SE* = .01), $\chi^2 = 8.86$, $p = .002$, pundits (Study 2) – by first name, raw percentages = 29.51% vs. 11.3%, $z =$

4.35, $p < .001$, $OR = .21$, and participants (Study 3) – by full name, Ms = .72 (SE = .02) vs. .55 (SE = .02) $\chi^2 = 208.84$, $p < .001$.

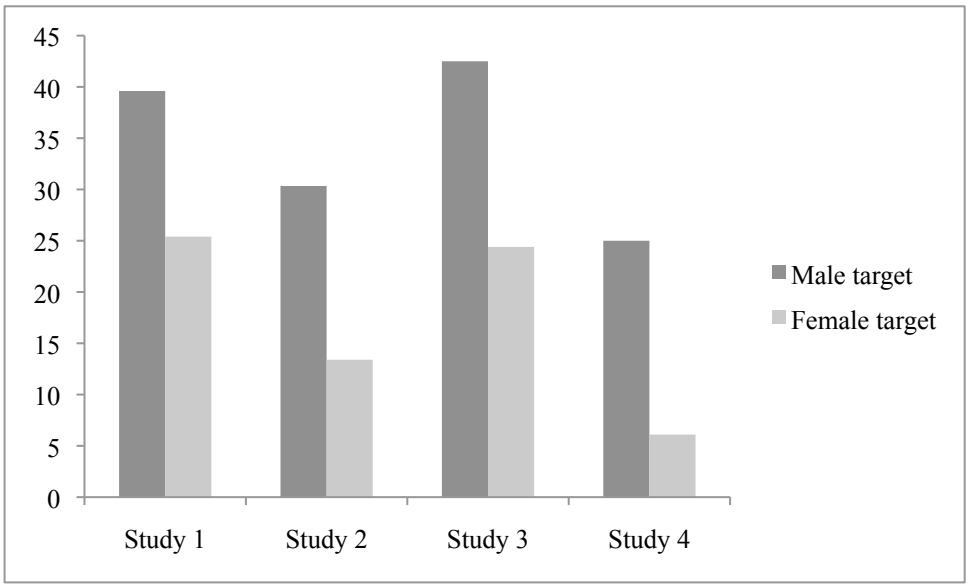

Figure 3. Percentage of responses containing or reporting a reference by surname to male and female targets (Studies 1-4). Numbers represent raw percentages (therefore no error bars are included). Across the 4 studies, people were, on average, 141.58% more likely to refer (or report referring) to male professionals than female professionals by surname (averaged at the level of study).

Studies 5a-b

Does the choice of reference have consequences for how the target is perceived and judged? Referring to a target by surname may imply a certain level of fame and eminence; the more famous a target, the fewer identifying details are needed. Indeed, in Study 3 we found that well-known and influential individuals are more likely to be referred to by surname (though this did not fully explain the gender difference). Thus, people might make the converse inference: that a surname reference signifies fame and eminence. In Studies 5a and 5b we tested this prediction.

In Studies 5a and 5b participants were presented with two pairs of one-paragraph research proposals. Within each pair, one proposal referred to the researcher by surname and the other by full name (first name was gender-neutral), and the text associated with each condition was counterbalanced across participants. Participants were asked which of the two researchers within each pair was better known (Study 5a, $N = 402$), more distinguished (Study 5a, $N = 399$), or more eminent (Study 5b, $N = 530$).

As predicted, researchers who were referred to by surname were selected as better known, $\chi^2 = 18.50$, $p < .001$, $OR = 1.91$ [1.44, 2.52], and more eminent, $\chi^2 = 9.79$, $p = .002$, $OR = 1.50$ [1.18, 1.91], but not significantly more distinguished, $p = .26$, than researchers referred to by full name (Figure 4). These results suggest that surname references are associated with fame and eminence[12][13].

[12] To test the possibility that participants evaluated the researchers differently because they thought the quality of the writing was worse and sounded less smooth when the researchers were referred to by full name, we ran a new study in which we presented participants with the same research proposals and asked them to rate how much the description of each proposal "read smoothly and made sense" on a 1 (Not at all smooth) to 7 (Very smooth) scale. They also evaluated which researcher was better known. Replicating the result of Study 5, we found that researchers who were referred to by surname (vs. full name) were judged as better-known, $\chi^2 = 8.54$, $p = .003$, $OR = 1.51$ [1.17, 1.93]. We found no evidence, however, that participants' judgments of smoothness differed depending on the type of reference, $p = .95$.

[13] In this and the rest of the studies in this paper we used gender-neutral names. To test whether the reference effect on fame would emerge when the target is clearly female, we ran a study where we identified the researchers with the gender-neutral names as female (preregistration details at http://aspredicted.org/blind.php?x=8td8sw). Participants read the same proposals as the ones used in Study 5, except that the researchers with gender-neutral names were identified using a female pronoun. Participants were asked which researcher was better known within each pair. We found that, replicating the result of Study 5a, researchers who were referred to by surname were selected as better known, $\chi^2 = 18.53$, $p < .001$, $OR = 1.71$ [1.37, 2.14].

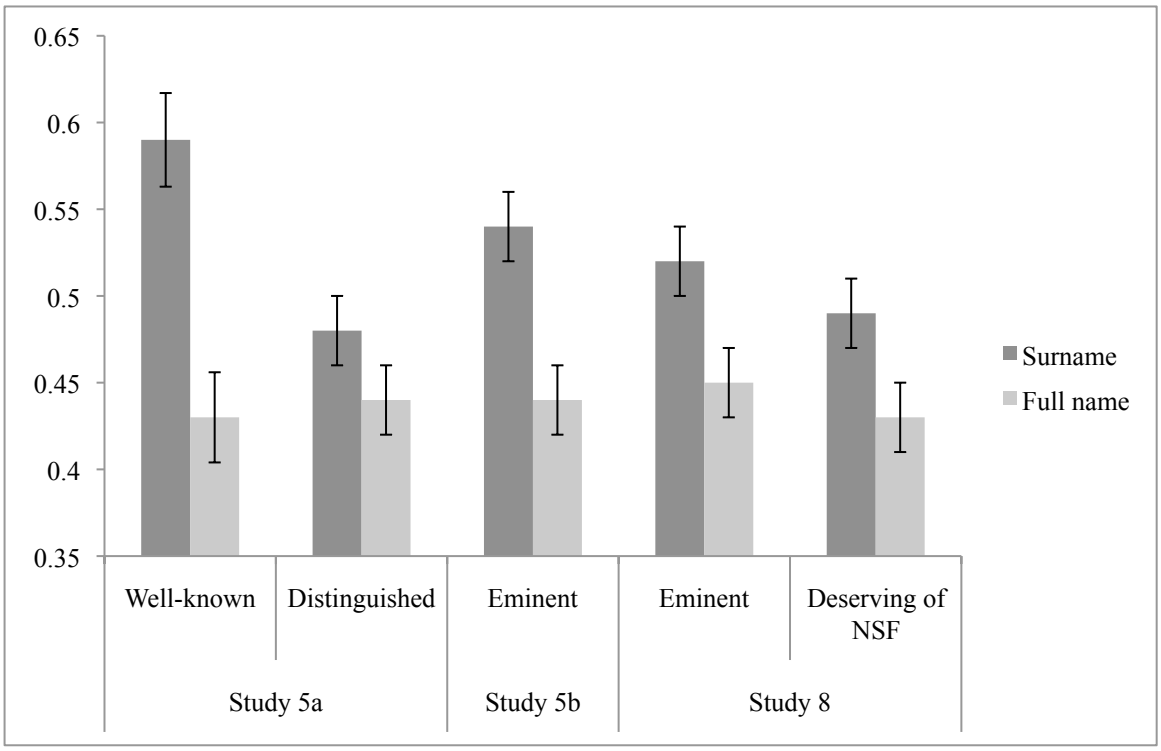

Figure 4. Likelihood of selecting a given scientist referred to by surname vs. full name in Studies 5a, 5b, and 8 (estimated marginal means). Error bars represent standard errors.

Study 6

We next tested whether the effect of reference type of fame would emerge when participants are not asked to directly compare the two types of reference, and are only exposed to one type of reference. To that end, we ran a study that was similar to Study 5a but which employed a between-subjects design. Participants were assigned to read two research proposals in which both researchers were referred to by either surname ($N = 463$) or full name ($N = 470$), varied randomly between participants. They were asked to estimate how well known each of the researchers was on a scale of 1 (not at all well known) to 9 (extremely well known). Consistent with our predictions, we replicated the result of Study 5a; specifically, even when evaluating the researchers individually and without being exposed to both types of reference, participants who read about researchers referred to by surname rated them as significantly more well-known than

participants who read about researchers referred to by full name, *M*s = 5.08 (*SD* = 1.59) vs. 4.57 (*SD* = 1.67), $t(931) = 4.81$, $p < .001$.

Study 7

Impressions of fame and eminence can have critical consequences. The Matthew Effect (Merton, 1968) refers to a rich-get-richer phenomenon in science and academia, such that "recognition is awarded partly on the basis of past recognition" (Vazire, 2017). This notion was first suggested in 1968 and has since been demonstrated repeatedly in various disciplines (Azoulay, Stuart, & Wang, 2013; Gush, Jaffe, Larsen, & Laws, 2018; Tol, 2009; Tomkins, Zhang, & Heavlin, 2017; Costas, Bordons, Van Leeuwen, & Van Raan, 2009; Bol, de Vaan, & van de Rijt, 2018). For instance, reviewers are more likely to accept papers by famous authors when they know the authors' identity than when they are blind to it (Tomkins et al., 2017). Given that a surname reference leads to perceptions of fame and eminence, targets may also enjoy the benefits that follow from such judgments. In the final two studies we tested whether surname (vs. full name) references influence judgments relating to the researcher's work, as well as whether these judgments are tied specifically to initial inferences about fame and eminence.

In Study 7 ($N = 517$) we tested whether researchers referred to by surname (vs. full name) are judged as being of higher status and as being more likely to win an award for their work. We hypothesized that these effects would emerge more strongly when the researcher's fame is explicitly brought to mind. The design was identical to that used in Studies 5a-b, except participants were asked which of the two researchers was of higher status in the relevant field, and who was more likely to win a prize for their work. Either before or after these measures (randomly determined), participants were asked which of the two researchers was better known.

Replicating the results of Study 5a, researchers referred to by surname were perceived as better known than those referred to by full name, *M*s = .58 (*SE* = .02) vs. .48 (*SE* = .02), χ^2 = 8.95, *p* = .003, *OR* = 1.51 [1.18, 1.93], and the effect did not significantly interact with question order, *p* = .13. As predicted, the effect of reference type on judgments of status and likelihood of winning a prize was qualified by question order (χ^2_{status} = 9.13, *p* = .003; χ^2_{prize} = 6.00, *p* = .014); when fame was brought to mind first, researchers who were referred to by surname (vs. full name) were perceived as holding higher status, *M*s = .71 (*SE* = .03) vs. .43 (*SE* = .03), χ^2 = 30.24, *p* < .001, *OR* = 3.12 [2.15, 4.51], and as more likely to win a prize for their work, *M*s = .64 (*SE* = .03) vs. .49 (*SE* = .03), χ^2 = 10.85, *p* = .001, *OR* = 1.86 [1.30, 2.67]. When fame was not first brought to mind, the effects on status and prize-winning-likelihood were statistically marginal, *M*s = .54 (*SE* = .03) vs. .46 (*SE* = .03), χ^2 = 3.20, *p* = .074, *OR* = 1.37 [.98, 1.93], and nonsignificant, *p* = .91, respectively.

Study 8

Some professional outcomes, such as career awards, are particularly tied to eminence. In Study 8 (*N* = 554) we tested whether researchers referred to by surname (vs. full name) are judged as more deserving of a fictitious National Science Foundation career award and of award-related funds. The design was similar to that used in Studies 5a-b, with two differences. First, because deservingness of award is based on research findings rather than proposed research, the research proposals from Studies 5a-b were altered to describe research findings instead of research proposals (e.g., "X hypothesizes" was changed to "X found"). Second, participants were asked which of the two researchers was more eminent, which should receive the prestigious and lucrative NSF career award given to the most eminent scientists, and, if the $500K prize money

were to be distributed to more than one researcher, how the money should be allocated between the two researchers.

We found that researchers referred to by surname were perceived as being more eminent (replicating the result of Study 5b), $\chi^2 = 4.76$, $p = .029$, $OR = 1.31$ [1.04, 1.66], and as 14% more deserving of an NSF career award, $\chi^2 = 4.34$, $p = .037$, $OR = 1.29$ [1.02, 1.64], than those referred to by full name (Figure 4). Researchers referred to by surname were also allocated a larger share of the award money by 6%, $F(1,551) = 4.61$, $p = .032$ (Figure 5).

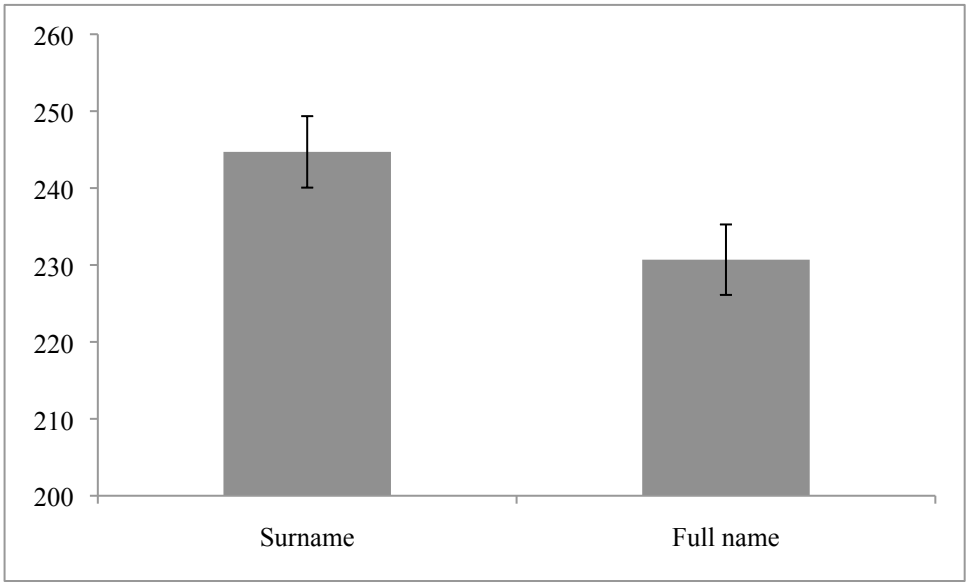

Figure 5. Amount of NSF career award allocated to a given scientist referred to by surname vs. full name in Study 8 (estimated marginal means). Error bars represent standard errors.

General Discussion

The way we speak about others influences and is influenced by the way we think about them (McConnell & Fazio, 1996). Across four studies using diverse research methods, we find evidence of a gender bias in the way we speak about professionals in a variety of domains. Specifically, analyses of archival data revealed that students reviewing their professors online were more likely to refer to their male professors than their female professors by surname alone

(Study 1) and that pundits and other commentators speaking about politics on the radio were more likely to refer to male than female targets by surname alone (Study 2). Participants' reports regarding the way they speak about well-known figures, including authors, athletes, politicians and others, showed the same pattern (Study 3). Finally, participants paraphrasing biographical information about a fictional male scientist were more likely to refer to him by surname alone than participants writing about an otherwise identical female scientist (Study 4). Taken together, the results suggest that gender predicts the way we speak about professionals, such that men are more likely than women to be referred to by their surname. The results of four additional studies suggest that this gender bias may be consequential: participants judged fictional researchers referred to by surname as better known and more eminent in their field than researchers referred to by full name, both when making a direct comparison (Studies 5a-b) and when evaluating the researchers individually (Study 6). Evidence suggests that this inference of fame and eminence, in turn, led to increased judgments of status, likelihood of winning an award, and deservingness of the NSF career award and associated funding (Studies 7-8).

What might explain the gender bias in use of surname references? First, surname may be more associated with men because in many cultures women's surnames are traditionally less permanent, commonly changing to a male partner's upon marriage (Goldin & Shim, 2004). Second, including a first name can often be used to mark the target's gender. Male is the assumed default (Hamilton, 1991; Miller, Taylor, & Buck, 1991), perhaps particularly in high status professions such as science, which are often male-dominated (National Science Board, 2016); thus, this gender marking may (intentionally or unintentionally) be deemed more necessary for a female target. Third, people may be more likely to attend to a woman's first

name because it marks her atypical gender in male-dominated professions; women's first names may therefore come to mind more easily and be used more often.

The evidence reported here, across experimental and archival data, documents a novel gender bias in professional reference. The implications may include biased, unwarranted judgments of female professionals as less well-known and eminent than their equivalent male counterparts, and consequently less deserving of the associated benefits. If people use the full name of female professionals to highlight women's participation and contribution, the consequences may be ironic, leading to lower judgments of eminence, status, and deservingness.

Materials and Methods

All data are available on Open Science Framework at https://osf.io/4a44t/. Additional details on the materials and methods are available in the SI.

Study 1.

Data. Data were obtained from Rate My Professors (www.ratemyprofessors.com), a website that allows students to evaluate their professors and the classes they teach on several dimensions (detailed below), and to post an accompanying open-ended comment. Data were collected for all professors (for whom reviews existed) in 5 departments (biology, psychology, computer science, history, and economics) from 14 universities. We chose universities that are considered academically rigorous (Cornell University, Columbia University, Brown University, and The Massachusetts Institute of Technology), universities with an active social scene (Bucknell University, Colgate University, Tulane University, Lehigh University, and University of Mississippi), and universities that are relatively conservative (Hillsdale University, Houston Baptist University, Harding University, Texas A&M University, and Liberty University). We did not collect data for any other universities.

For each of 18,046 reviews (of 1,674 professors), we recorded the professor's gender, her/his university and department affiliations, the state and city in which the university is located, the reviewer's ratings of helpfulness, clarity, easiness, and interest (included in most but not all comments), and the year the review had been posted. If the review was accompanied by an open-ended comment (94.5% of reviews, $N = 17{,}055$), we checked whether the comment contained a reference to the professor. If it did (24.9% of reviews, $N = 4{,}494$), the comment was coded according to the type of reference used: by surname, full name, first name, Prof. full name, Prof. surname, Dr. full name, Dr. surname, common titles (Mr./Ms/Mrs./Miss surname), or other reference. Seniority was computed by taking the difference between the years of the most recent and oldest reviews of a professors and adding one. We also recorded whether the reviewer described the professor using any stereotypically female or male traits. This coding, as well as the coding of common titles, was conducted at a later date than the initial coding; these variables were coded for all reviews except those of professors in five groups (biology at Cornell University, economics at Brown University, and computer science, economics, and biology at Columbia University), for which technical difficulties prevented us from obtaining the comments' texts again. Of the reviews that included a reference to the professor, 71.5% ($N = 3{,}212$) were to male and 28.5% ($N = 1{,}282$) were to female professors.

Analysis. We focused on the comments that contained a reference to the professor ($N = 4{,}494$). In order to compare use of surname references to male and female professors, we created a dummy variable for each reference type; for example, in the main analysis surname references were coded as 1 and all other references (full name, first name, etc.) were coded as 0. We used a Generalized Linear Model to determine the effect of professor gender on the use of a particular reference type with repeated reviews of the same professor.

Study 2.

Data. Data were obtained from transcripts of the following politically diverse (i.e., conservative and liberal) American radio programs that regularly discuss politics and current events: *All Things Considered,* (1,982 references from 124 segments, 9.20.2013 – 11.8.2015), *Fresh Air* (1389 references from 17 segments, 8.13.2014 – 10.15.2015), *Morning Edition* (1256 references from 101 segments, 1.22.2015 – 11.30.2015), *The Rush Limbaugh Show* (2022 references from 66 segments, 10.5.2015 – 12.4.2015), and *The Sean Hannity Show* (2923 reference from 28 segments, 8.25.2015 – 12.1.2015). Overall, 9,572 references were coded.

For programs that broadcast both political and non-political content, only political segments were coded. Clearly scripted speech was not coded; we instead focused on interview segments, which are less scripted and thus less likely to be constrained by journalistic reference conventions. For each reference to a third-party, research assistants recorded the type of reference used (surname, full name, first name, title + full, title + surname, or other), whether it was the first reference to the target in that segment, as well as the following information about both the target and speaker: full name, gender, political affiliation (if known), and position or title.

Analysis. We used a multilevel, crossed-effects logistic regression to determine the effect of target gender on reference use. Our statistical model included a fixed effect of target gender and random effects of target and speaker identity to account for the non-independence of each speaker's and target's observations in our data set. Reference use was treated as a binary dependent variable (e.g., surname = 1 and any other reference = 0).

Participants (Studies 3-8).

Participants in Studies 3-8 completed the study through Amazon's Mechanical Turk in exchange for monetary compensation. Participation was restricted to respondents within the United States. The only exception was one of the two separate samples of participants in Study 3 who were recruited later to rate the stimuli; they were Cornell undergraduate students. Data were obtained from 190 participants in Study 3 (96 women, 94 men, mean age = 35.17, *SD* = 11.48; two additional samples were recruited later to rate the stimuli in this study; one sample of 217 participants (120 men, 97 women, mean age = 38.34, *SD* = 12.70) and another of 44 participants students (9 men, 35 women, mean age = 19.73, *SD* = 1.33)), 183 participants in Study 4 (93 women, 87 men, 2 other, mean age = 33.27, *SD* = 10.09, 1 did not report age and gender), 801 participants in Study 5a (376 women, 419 men, 3 other, mean age = 36.21, *SD* = 11.45, 2 did not report both gender and age, 1 did not report only gender, 1 did not report only age), 530 participants in Study 5b (287 women, 231 men, 2 other, mean age = 35.06, *SD* = 11.29, 10 did not report gender, of which 5 also did not report age), 933 participants in Study 6 (544 women, 378 men, 7 other, mean age = 36.49, *SD* = 11.91, 4 did not report gender, of which 1 also did not report age), 517 participants in Study 7 (302 women, 209 men, 4 other, mean age = 34.92, *SD* = 11.26, 5 did not report gender, of which 3 also did not report age), and 554 participants in Study 8 (333 women, 215 men, 3 other, mean age = 34.64, *SD* = 10.57, 3 did not report gender). Additional participants failed to complete the entire study (42, 114, 41, 50, 30, 45, and 64 participants in Studies 3, 4, 5a, 5b, 6, 7, and 8, respectively), failed the attention check (82, 3, 89, and 54 participants in Studies 5b, 6, 7, and 8, respectively), or did not follow instructions (in Study 4, three participants copied the bullet points instead of rephrasing the information in their own words, and one participant did not include any of the information from the bullet points in the response) and were excluded from all analyses.

Procedures Common to Studies 3-8.

For all studies with participants (Studies 3-8), informed consent was obtained at the beginning of the study. After the study, participants filled out a demographic questionnaire and provided information for payment. All procedures for these studies were approved by the Institutional Review Board of Cornell University.

Study 3.

Procedure and Materials. Participants were presented with two lists of well-known individuals. One list consisted of figures in American politics, and the other consisted of well-known individuals in various other domains, e.g., literature, science, and sports. Half of the individuals in each list were women, and half were men. The lists were presented individually on consecutive pages in counterbalanced order, and the order of the individuals within each list was digitally randomized for each participant. See Appendix A for a list of the figures.

Participants were asked to think about how they refer to each figure when talking about her or him in casual conversation and to estimate what percentage of the time they refer to her or him by surname, full name, first name, or in some other way. The total had to add up to 100. If they did not know who the individual was, they chose the N/A option instead.

Two separate groups of participants provided additional ratings of the stimuli. Participants in one group were randomly assigned to provide one of the following ratings: how well known each individual was; how distinctive each first name was, or how distinctive each surname was. An additional group rated how influential each position of the political figures was (e.g., *attorney general, governor*). Finally, the researcher noted any women who shared a surname with a well-known family member (Hillary Clinton, Sarah Palin, Nancy Pelosi, Louisa May Alcott, Marie Curie, Serena Williams, Virginia Woolf). These variables were used as

controls in subsequent analyses. Participants then reported their political identity on a scale from 1 (*very liberal*) to 5 (*very conservative*).

Analysis. The modal response for reported use of references was a 100% for one of the options (with full name being the most common choice) and zeroes for the other options. Thus, the percentage of surname use was severely right skewed, with mostly 0 choices. We therefore created a binary distinction between 0% and any other percentage. We used Generalized Estimating Equations to determine the effect of figure gender on the use of a particular reference type while controlling for differences in surname use between figures and between participants. We also ran all analyses with the raw data in which reported use of each reference was continuous, using a mixed linear model. The results were very similar across approaches. All statistics reported in the main text and Appendix B reflect the results of the binary approach; however, the results were statistically significant for both unless noted otherwise.

Study 4.

Procedure and Materials. Participants read information about a scientist, presented in the form of a list of bullet points (e.g., *Chemist and X-ray crystallographer, Worked in physical chemistry lab, 1945: Earned PhD*; the scientist was inspired by Rosalind Franklin). The scientist's name, which appeared in bolded font at the top of the list, was either *Dolores Berson* (female condition) or *Douglas Berson* (male condition). Both names are uncommon today but were roughly equally popular early in the 20th century when the fictional scientist was active (based on the online Name Voyager tool (http://www.babynamewizard.com/voyager), which uses information reported by the Social Security Administration). All other information was identical between gender conditions. Participants were then randomly assigned to either imagine that they were asked to give a lecture about the scientist (formal expression condition) or to

imagine that they were telling a friend about the scientist in the course of casual conversation (casual expression condition). They were asked to type in their lecture or conversation, making sure to incorporate all the information from the bullet points in full sentences.

***Analysis*.** The distributions of the average number of each type of reference in an essay were all right skewed, with most essays containing either one or no references of a given type. Thus, for each type of reference, we created a binary distinction between any instances of that reference in an essay (coded as 1) and no instances of that reference in an essay (coded as 0). A logistic regression was performed to test the effects of researcher gender on the likelihood that participants used a surname reference in their essay.

Studies 5a-b.

Procedure and Materials. Study 5a was preregistered on Open Science Framework (DOI 10.17605/osf.io/bh7aa). Participants in this study were asked to imagine themselves as an employee of a funding agency evaluating research proposals by academics. They then read four 1-pararaph summaries of research proposals (see Appendix A for an example proposal). The proposals comprised two pairs (4 proposals) pretested to differ minimally on any of the dependent variables when no names were mentioned. The critical difference was that in one of the proposals in each pair the researcher was referred to by surname (Boland, Hastings, Wiggins, and Hirst), whereas in the other the researcher was referred to by full name with a unisex first name (Jamie and Casey; chosen based on a pretest testing the extent to which different names are seen as unisex). The assignment of reference type to any given proposal within a pair was counterbalanced between participants. After reading all proposals once, participants chose between proposals within each pair; they *either* answered "In your estimation, which of the two

researchers is better-known?" *or* "In your estimation, which of the two researchers is more distinguished?"

Study 5b was very similar to Study 5a, except participants were asked about eminence instead of fame: "What is your best guess as to which of the two researchers is more eminent?" Eminence was defined for participants as "fame or recognized superiority, especially within a particular sphere or profession."

Analysis. We used Generalized Estimating Equations to determine the effect of type of reference on target evaluations with repeated measures (each participants chose between two pairs of proposals). In all analyses we controlled for the specific proposal and researcher name.

Study 6.

Procedure and Materials. The study was similar to Studies 5a and 5b but employed a between-subjects manipulation of reference type: participants read two short research proposals, and in both research proposals the researcher was referred to using the same reference type; for half of the participants – by surname (Berson and Boland), and for half – by a gender-neutral full name (first names: Alex and Riley). After reading each proposal, participants were asked "in your estimation, how well known is this researcher?" on a scale of 1 (not at all well known) to 9 (extremely well known).

Analysis. The ratings of the two researchers were averaged for each participant. The ratings of researchers referred to by surname vs. full name were compared using a 2-tailed *t*-test.

Study 7.

Procedure and Materials. The study was very similar to Study 5b, except that participants were asked about each pair "What is your best guess as to which of the two researchers is of higher-status in their field?" and "What is your best guess as to which of the two

researchers is more likely to win a prize for their work?" They were also asked "What is your best guess as to which of the two researchers is better known?" and this question was presented either first or last, randomly determined.

Analysis. Identical to Studies 5a-b.

Study 8.

Procedure and Materials. The study was similar to Study 5b, except that participants imagined working specifically for the National Science Foundation (NSF). The proposals were altered to reflect research findings (e.g.," Hastings is interested in exploring" was changed to "Hastings explored"). Participants were asked three questions about the researchers within each pair: who was more eminent, who "should receive the prestigious and lucrative career National Science Foundation award, given to the most eminent scientists in the country?" and how much of the $500K award money they would allocate to each researcher if they didn't have to give the award to just one person.

Analysis. Identical to Studies 5a-b.

References

1. Proctor, B. D., Semega, J. L., Kollar, M. A. (2016). *Income and poverty in the United States: 2015.* Washington, DC: US Government Printing Office.
2. National Science Board. (2016). *Science and engineering indicators 2016.* Arlington, VA: National Science Foundation.
3. Mo, C. H. (2015). The consequences of explicit and implicit gender attitudes and candidate quality in the calculations of voters. *Political Behavior*, *37*, 357-395.
4. Nosek, B. A., Smyth, F. L., Sriram, N., Lindner, N. M., Devos, T., Ayala, A., ... & Kesebir, S. (2009). National differences in gender–science stereotypes predict national sex differences in science and math achievement. *Proceedings of the National Academy of Sciences*, *106*, 10593-10597.
5. Leslie, S. J., Cimpian, A., Meyer, M., & Freeland, E. (2015). Expectations of brilliance underlie gender distributions across academic disciplines. *Science*, *347*, 262-265.
6. Riffkin, R. (2014). Americans still prefer a male boss to a female boss. *Gallup News Service.*
7. 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*, 16474-16479.
8. Uscinski, J. E., & Goren, L. J. (2011). What's in a name? Coverage of Senator Hillary Clinton during the 2008 Democratic Primary. *Political Research Quarterly*, *64*, 884-896.
9. Messner, M. A., Duncan, M. C., & Jensen, K. (1993). Separating the men from the girls: The gendered language of televised sports. *Gender & Society*, *7*, 121-137.

10. Halbert, C., & Latimer, M. (1994). “Battling” gendered language: An analysis of the language used by sports commentators in a televised coed tennis competition. *Sociology of Sport Journal*, *11*, 298-308.
11. Cejka, M. A., & Eagly, A. H. (1999). Gender-stereotypic images of occupations correspond to the sex segregation of employment. *Personality and Social Psychology Bulletin*, *25*, 413-423.
12. 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.
13. Langford, T., & MacKinnon, N. J. (2000). The affective bases for the gendering of traits: Comparing the United States and Canada. *Social Psychology Quarterly*, *63*, 34-48.
14. Cowan, G., & Kasen, J. H. (1984). Form of reference: Sex differences in letters of recommendation. *Journal of Personality and Social Psychology*, *46*, 636-645.
15. Files, J. A., Mayer, A. P., Ko, M. G., Friedrich, P., Jenkins, M., Bryan, M. J., ... & Duston, T. (2017). Speaker introductions at internal medicine grand rounds: forms of address reveal gender bias. *Journal of women's health*, *26*, 413-419.
16. Takiff, H. A., Sanchez, D. T., & Stewart, T. L. (2001). What's in a name? The status implications of students' terms of address for male and female professors. *Psychology of Women Quarterly*, *25*, 134-144.
17. Merton, R. K. (1968). The Matthew effect in science: The reward and communication systems of science are considered. *Science*, *159*, 56-63.
18. Vazire, S. (2017). Our obsession with eminence warps research. *Nature News*, *547*, 7.

19. Azoulay, P., Stuart, T., & Wang, Y. (2013). Matthew: Effect or fable?. *Management Science*, *60*, 92-109.
20. Gush, J., Jaffe, A., Larsen, V., & Laws, A. (2018). The effect of public funding on research output: the New Zealand Marsden Fund. *New Zealand Economic Papers*, *52*, 227-248.
21. Tol, R. S. (2009). The Matthew effect defined and tested for the 100 most prolific economists. *Journal of the American Society for information Science and Technology*, *60*, 420-426.
22. Tomkins, A., Zhang, M., & Heavlin, W. D. (2017). Reviewer bias in single-versus double-blind peer review. *Proceedings of the National Academy of Sciences*, *114*, 12708-12713.
23. Costas, R., Bordons, M., Van Leeuwen, T. N., & Van Raan, A. F. (2009). Scaling rules in the science system: Influence of field-specific citation characteristics on the impact of individual researchers. *Journal of the American Society for Information Science and Technology*, *60*, 740-753.
24. Bol, T., de Vaan, M., & van de Rijt, A. (2018). The Matthew effect in science funding. *Proceedings of the National Academy of Sciences*, *115*, 4887-4890.
25. McConnell, A. R., & Fazio, R. H. (1996). Women as men and people: Effects of gender-marked language. *Personality and Social Psychology Bulletin*, *22*, 1004-1013.
26. Goldin, C., & Shim, M. (2004). Making a name: Women's surnames at marriage and beyond. *Journal of Economic Perspectives*, *18*, 143-160.
27. Hamilton, M. C. (1991). Masculine bias in the attribution of personhood: People= male, male= people. *Psychology of Women Quarterly*, *15*, 393-402.
28. Miller, D. T., Taylor, B., & Buck, M. L. (1991). Gender gaps: Who needs to be explained?. *Journal of Personality and Social Psychology*, *61*, 5-12.

Appendix A: Additional Details on Materials and Methods

Study 1. In addition to the review-level variables described in the main paper, for each of 1,674 professors, we recorded the following variables from Rate My Professors: average ratings of helpfulness, clarity, easiness, and interest (all on a 1-5 scale), and hotness (illustrated on Rate My Professors using a picture of a hot pepper; 0 = no pepper to 3 = a red, flaming pepper). Gender was inferred from the professor's name and from the pronouns used in the reviews or on the university's website.

Since these data were collected, Rate My Professors has changed the information included in a review and presented on a professor's page. Some of the variables we have collected are no longer available on the website.

Of the comments that contained a reference to the professor, a small percentage of comments contained more than one reference (1.58%), of which most (69.0%) were the same reference type as the first reference. Thus, in the case of multiple references appearing in the same comment, only the first reference was included in the analyses.

Additional data: Our initial plan was to collect data for all reviews in all departments at Cornell, Columbia, and Brown Universities. We began this process by going through the departments in alphabetical order, starting with accounting. We coded some data in this way before we decided to include a broader range of universities and focus instead on a smaller number of departments (biology, psychology, computer science, history, and economics). Including these initial data leaves all the results reported in this paper essentially unchanged.

Study 2. For programs that broadcast both political and non-political content, only political segments were coded. Clearly scripted speech was not coded. For example, political reports were not coded because their wording was presumably written in advance in accordance

with journalism reference conventions (full name is used for the first reference, and surname for all subsequent references; Associated Press, 2016). Instead, we focused on interview segments, which are less scripted. *The Rush Limbaugh Show* is unscripted throughout and was thus coded in its entirety.

Study 3. The list of figures in American politics: *Susan Rice, Loretta Lynch, Hillary Clinton, Madeleine Albright, Nancy Pelosi, Sonia Sotomayor, Ruth Bader Ginsburg, Elizabeth Warren, Wendy Davis, Ann Kirkpatrick, Barbara Boxer, Dianne Feinstein, Janet Yellen, Sarah Palin, Joni Ernst, Carly Fiorina, Nikki Haley, Michele Bachmann, Sandra Day O'Connor, Susan Collins, Susana Martinez, Mia Love, Condoleezza Rice, Ashton Carter, Thomas Perez, Mitt Romney, Colin Powell, John Boehner, John Roberts, Antonin Scalia, Ted Cruz, Chris Murphy, Adam Schiff, Harry Reid, Dick (Richard) Durbin, Ben Bernanke, Joe Biden, Ed Markey, Bernie Sanders, Josh Earnest, Paul Ryan, Clarence Thomas, Bill Nelson, Mike Pence, John Edwards,* and *John Kerry*.

The list of well-known individuals in various other domains: *Maya Angelou, Toni Morrison, Jane Austen, Susan B. Anthony, Danielle Steel, Stephenie Meyer, Florence Nightingale, Emily Dickinson, Sally Ride, Louisa May Alcott, Marie Curie, Helen Keller, Serena Williams, Lindsey Vonn, Virginia Woolf, Amelia Earhart, Margaret Mead, Diane Sawyer, Dorothy Day, Margaret Thatcher, Jane Goodall, Angela Merkel, Barbara Walters, Robert Frost, Stephen King, Charles Dickens, Martin Luther King Jr., John Grisham, Neil Gaiman, Francis Galton, Mark Twain, Neil Armstrong, William Faulkner, Robert Oppenheimer, John Steinbeck, David Beckham, Ryan Lochte, Ernest Hemingway, Howard Hughes, Charles Darwin, David Muir, Thomas Merton, Winston Churchill, Carl Sagan, Vladimir Putin,* and *Peter Jennings*).

The list of women was generated first, by brainstorming well-known female figures. The list of men is made up of well-known male figures that are roughly matched, on average across the lists, on age, years active, and position or profession.

Studies 5a-b. Participants in this study were asked to imagine themselves as an employer at a funding agency:

> Imagine that you work for a funding agency. Your job is to evaluate research proposals by academics who wish to have their projects funded.
>
> Submission is open to researchers at all levels; the funding agency receives proposals from academics who are well-established in their field, as well
>
> as from researchers with little experience, who are still learning. Thus, the research proposals vary in their quality and in how much they deserve funding.

They then read four 1-pararaph summaries of research proposals. The following is an example proposal:

> Hastings is interested in exploring the causes and consequences of biodiversity.
>
> Neighboring populations of animals may be similar, and yet differ in important ways.
>
> Several hypotheses have been suggested regarding the potential benefits of
>
> biodiversity. Hastings plans to use neighboring salmon populations to shed light on how such diversity emerges, and why it is maintained.

In Study 5b, Prior to answering questions, participants read: "You may feel like you do not know enough to make these judgments, but we would like you to base your answers on what you have read and give us your best guess about the questions."

Appendix B: Supplementary Results

Study 1. Reviews of female professors (vs. those of male professors) were marginally less likely to include an open-ended comment, *M*s = .94 (*SE* = .003) vs. .95 (*SE* = .005), χ^2 = 3.09, *p* = .079, and were less likely to contain a reference, *M*s = .24 (*SE* = .009) vs. .27 (*SE* = .006), χ^2 = 6.79, *p* < .01.

In addition to the effects of target gender on references described in the main text, female professors were marginally more likely to be referred to as Dr. surname (*M* = .27, *SE* = .021) than were men (*M* = .23, *SE* = .013), χ^2 = 3.75, *p* = .053. None of the other types of reference differed significantly by professor gender except as described in the main text, *p* > .19.

Female professors were rated as less helpful, *M*s = 3.61 (*SE* = .04) vs. 3.76 (*SE* = .03), *t*(1548) = 2.85, *p* < .005, and clear, *M*s = 3.56 (*SE* = .04) vs. 3.72 (*SE* = .03), *t*(1548) = 2.85, *p* < .005, than male professors. However, helpfulness and clarity, as well as interest in the course, *negatively* related to use of surname, $\chi^2_{\text{helpfulness}}$ = 70.39, *p* < .001, χ^2_{clarity} = 48.26, *p* < .001, χ^2_{interest} = 31.39, *p* < .001. Thus, target gender appears to have independent effects on use of surname and on ratings of helpfulness and clarity. Ratings of how easy the course was were not significantly related to gender and to use of surname, *p* = .79.

Study 2. Because the first reference to a given target within a segment may be more likely to include a full name as a way to introduce the target, we also analyzed the data without these first references; as predicted, speakers referred to men more than to women by surname *M*s = 38.25% vs. 17.66% , *z* = -4.55, *p* < .001, *OR* = .15 [0.06, 0.33].

The effect of target gender on surname reference was qualified by an interaction with speaker gender, *z* = 3.82, *p* < .001; only male speakers were significantly more likely to refer to men than women by surname, *M*s = 32.76% vs. 12.65%, *z* = -4.53, *p* < .001, *OR* = .17 [0.08, 0.36], whereas female speakers were not, *M*s = 18.26% vs. 15.62%, *p* = .11. It is worth noting,

however, that there were many more observations for male speakers than female speakers (*N*s = 7,849 vs. 1,723), making the analyses with only male speakers better powered; it would be important to analyze a dataset with a greater number of female speakers to determine whether they also show the gender effect. The effect of gender on surname reference did not significantly interact with target or speaker political affiliation, *p*s > .07. Pundits were also more likely to refer to women (29.51%) than to men (11.30%) by first name, $z = 4.35$, $p < .001$, $OR = 10.65$ [3.67, 30.95], and were more likely to refer to men (37.99%) than to women (17.78%) by their title along with their surname (e.g., Senator Schumer), $z = -4.72$, $p < .001$, $OR = .26$ [0.15, 0.46]. Pundits did not differ significantly in their likelihood of referring to men and women by full name or by title along with full name, *p*s > .08.

Study 3. The gender bias in surname use emerged for both political, $\chi^2 = 32.13$, $p < .001$, $M_{men} = .38$ (.02) vs. $M_{women} = .31$ (.02), and non-political figures, $\chi^2 = 213.55$, $p < .001$, $M_{men} = .45$ (.02) vs. $M_{women} = .18$ (.02), but was larger for non-political figures, $\chi_{interaction}^2 = 120.93$, $p < .001$, perhaps because the latter were better known overall. The results were slightly different when we used a mixed linear model (see Methods section for details on why we used two different methods of analysis); the bias was only statistically significant for non-political figures, $F(1, 43.96) = 18.73$, $p < .001$, and not statistically significant for political figures, $p = .37$.

The effect of gender on surname was qualified by an interaction with party affiliation of target, $\chi^2 = 22.68$, $p < .001$: participants were significantly more likely to use surname when referring to Republican men vs. women, $\chi^2 = 49.88$, $p < .001$, but not when referring to Democratic politicians, $p = .53$. As reported in the main text, participants were more likely to report using full name when referring to women than to men, $\chi^2 = 208.84$, $p < .001$. They were also more likely to use first name only when referring to women than to men, $\chi^2 = 10.69$, $p <$

.005. However, this latter result was not significant using a mixed linear model, $p = .24$, nor when Hillary Clinton was excluded from the analysis, $p = .80$, suggesting this result may have been driven by a single politician rather than a broader trend. As reported in the main text, targets whose first names were perceived to be more common were more likely to be referred to by surname. However, this result was not significant when using a mixed linear model, $p = .55$.

Study 4. Participants most commonly referred to the scientist in their essays by full name (60.27% of references), followed by surname only (17.47%) and first name only (16.78%). A small percentage of references referred to the scientist by professional title (Dr. Surname/Full Name; 3.43%) or common title (Mr./Mrs. Surname/Full Name; 2.05%).

A logistic regression was performed to test the effect of researcher gender on the likelihood that participants used a surname reference in their essay. The model explained 12.5% (Nagelkerke R^2) of the variance in use of surname in the essay and correctly classified 85.2% of cases. Similar logistic regressions were run for full name and first name references. Researcher gender did not significantly predict use of these references, ps > .88.

Four participants in the female condition referred to the scientist by male pronouns. Excluding these participants leaves the results essentially unchanged.

Studies 5a and 5b. The effect of surname on judgments of fame was qualified by a marginal interaction with participant gender, $\chi^2 = 3.68$, $p = .055$; follow up analyses revealed that female participants were more likely to select the researcher referred to by surname (vs. full name) as better known, ($M = .65$ ($SE = .04$) vs. $M = .41$ ($SE = .04$)), $\chi^2 = 17.77$, $p < .001$; male participants were only marginally more likely to do so, ($M = .55$ ($SE = .04$) vs. $M = .46$ ($SE = .04$)), $\chi^2 = 3.36$, $p = .067$. The was no evidence that the effect of surname on judgments of being distinguished or eminent differed between male and female participants; the effect on judgments

of being distinguished or eminent did not significantly interact with participant gender, *p*s = .52 and .25, respectively.

Study 6. We found no evidence that the effect of reference type on judgments of fame differed between male and female participants, $p = .19$. To determine whether the difference between the conditions was the result of demand effects, we included a hypothesis probe at the end of the study, asking participants what they thought "this study was about or what it was trying to test". Five hundred and thirteen participants responded to the hypothesis probe; of those, only 3.9% guessed that the study had anything to do with names, and none mentioned a difference in reference type. Given that none of the participants who responded guessed our hypothesis or identified the independent variable, it does not appear that the difference between the conditions is the result of demand effects.

We also ran two pre-registered replications of this study (preregistrations details available at http://aspredicted.org/blind.php?x=ue2d2m and http://aspredicted.org/blind.php?x=ts45du). These studies were identical to the original study, except that we modified our attention check in the second replication study to make it slightly more difficult (almost all participants passed the attention check in the previous two studies, so we reasoned that it may have been too easy) and increased our sample size by 15% to make up for the potential increase in excluded participants. We found that in both studies, participants rated researchers referred to by surname as better known than researchers referred to by full name, and this difference was statistically significant in one of the studies, *M*s = 4.85 (*SD* = 1.71) vs. 4.53 (*SD* = 1.71), $t(1078) = .001$ (one-tailed), and was directional but did not reach statistical significance in the other study, *M*s = 4.89 (*SD* = 1.67) vs. 4.72 (*SD* = 1.75), $t(932) = .075$ (one-tailed).

Study 7. We found no evidence that the effect of reference type on judgments of fame, status, or likelihood of wining a prize differed between male and female participants, *p*s > .69.

Study 8. We found no evidence that the effect of reference type on eminence, choice of the NSF winner, or the allocation of the NSF cash award differed between male and female participants, *p*s > .53

Appendix C: Additional Studies

We conducted several additional studies in which our main goal was to test different questions than the ones addressed in the paper; however, the design sometimes also allowed us to re-test the questions that were the focus of this paper. We report these results here. We ran five studies using a similar design to Studies 5-8 (participants read research proposals about scientists who were referred to by full name or surname). The results were all consistent with our reported results and conclusions: in all five studies, participants judged the researcher referred to by surname to be better known than the researcher referred to by full name, and this difference was statistically significant (in four studies) or directional (in one study). One additional study was similar but employed a between-subjects design; its sample size was small compared to our other studies ($N = 209$), and though the result did not reach statistical significance ($p = .21$, two-tailed), it was in the predicted direction, i.e., participants rated researchers referred to by surname as better known than researchers referred to by full name, Ms = 4.73 ($SD = 1.61$) vs. 4.44 ($SD = 1.72$). We used the data from this latter study to compute the sample size for Study 6 and its replications. Consistent with the results of Study 7, in another study we found that reference type did not significantly affect judgments of deservingness of a career award when fame was not first brought to mind, which motivated us to run Study 7. We are also in the process of designing a

paradigm for a follow-up project on the potential behavioral consequences of the gender reference bias in a new domain. To that end, we have conducted three exploratory studies in a different domain; in these studies (two using a within-subjects and one a between-subjects design), participants read reviews of popular science books by authors referred to by surname or full name. We found no statistically significant differences between conditions in how well known participants judged the authors to be. Note, however, that each study had some methodological issue (in one study the full name appeared once in the surname condition, in another there was only one reference, and it did not appear until halfway through the text). There were also considerable differences between these studies and the studies reported in this paper; for example, they were in a different domain where fame may play a different role, and the descriptions were longer and may not have engaged the participants. We are still improving the design of this new paradigm. These studies were not run with the purpose of being included in this paper, but we are reporting them here for the sake of transparency.

CHAPTER V

GENERAL DISCUSSION

Overall, the research described here examines the way people think about themselves and others with regards to knowledge and expertise from two different directions. The first direction elucidates how people go about assessing their own knowledge, and why this process does not always result in perfectly accurate judgments. The second direction spotlights one gender bias in the way people speak about others who have expertise, i.e., professionals, and explores the bias's potential consequences.

Self-judgments of knowledge

The results of 9 studies across two papers add to our understanding of how people make self-judgments of knowledge in three primary ways. First, they suggest that assessing one's knowledge is not as straightforward as intuition may lead us to suspect. It seems to be a more complicated matter than rifling through a mental knowledge catalogue for a match and then carefully and accurately weighing the contents against a standard of knowledge to arrive at the appropriate assessment. Instead, the conclusion that emerges is that self-judgments of knowledge are fairly intuitive, and this intuition is informed by myriad heuristics and cognitive cues. The cues are of varying levels of reliability and, what's more, their reliability changes depending on circumstances.

Second, the work reported here provides evidence of two specific cues that play a role in self-judgments of knowledge. Chapter II reported a set of studies showing that the more people view themselves as knowledgeable within a particular domain, the more likely they are to make mistakes of a particular kind in self-judgments of knowledge: they are more likely to mistakenly feel they have some knowledge of terms, places, people, etc. that are ostensibly related to that

domain but, in fact, do not exist. Study 4 showed that when people's self-perceived knowledge was experimentally inflated, their overclaiming shifted upwards accordingly. Study 3 hinted that overclaiming at least partly reflects honest mistakes, given that overclaiming did not decrease substantially when participants were warned that some of the terms were bogus; thus, people still reported being familiar with the terms even when they were given a self-presentational motivation not to pretend familiarity. Importantly, self-perceived knowledge remained a significant predictor of overclaiming. Taken together, the results reported in Chapter II uncover one heuristic that people use to assess their knowledge of specific items: their self-perceived global level of knowledge on the relevant topic. If people believe they are generally knowledgeable about a topic, they are more likely to overestimate their knowledge of specific terms drawn from that topic – even ones that do not exist. As is the case with most heuristics, this one may be useful in some contexts. It is not unreasonable to rely on a general assessment of knowledge within a domain to infer that one is likely or unlikely to be familiar with a given domain-related term.

Chapter III reported evidence of a different cognitive cue used to assess knowledge: the *relative* volume of content that comes to mind when considering a given term. If one is highly knowledgeable within a given domain, the information that would come to mind about a bogus term will appear relatively sparse, and the person will use this cue to infer that they have little or no knowledge about the term. For a less knowledgeable person, using this cue may lead to errors, such that unfamiliar terms – bogus or not – will still appear *relatively* more similar to the terms the person does know in terms of the amount of mental content they prompt, and the person may incorrectly infer that she has some knowledge of them.

Third, this research sheds some light on the nuanced relationship between overclaiming, self-perceived knowledge and genuine knowledge. The studies reported in Chapter III make it clear that, although self-perceived knowledge is often correlated with genuine knowledge, self-perceived knowledge on a topic is far from a perfect measure of how much someone actually knows about that topic. While flattering self-perceptions of knowledge lead to more difficulty in differentiating what one knows from what one doesn't, high levels of real expertise are associated with an improved ability to make this differentiation, at least in the overclaiming task. Thus, the two related constructs – self-perceived knowledge and genuine knowledge – exert independent, opposing effects on overclaiming.

Future Directions

The final study in Chapter II demonstrated that experimentally manipulating self-perceived knowledge shifted overclaiming in the predicted direction, which strongly suggests that the relationship between self-perceived knowledge and overclaiming is not merely correlational but likely causal. However, the mechanism by which self-perceived knowledge influences overclaiming remains to be discovered. One hypothesis is that one's overall sense of knowledgeability in a topic is directly fed into the quick, intuitive feeling of knowing on which knowledge judgments often appear to be based. It may also be the case that self-perceived knowledge increases one's motivation to find the terms familiar. If a person views himself as knowledgeable in biology, for instance, he may want the assessments of his knowledge of specific terms within biology to align with this self-perception. This increased motivation may lead him to alter his assessment process in various ways to achieve the desired outcome. For example, he may lower his threshold for what counts as knowledge; perhaps a mild feeling of

familiarity with the term will be deemed as sufficient justification for claiming some knowledge when one is motivated to know, but not when one is indifferent about it.

This motivated person may instead or in addition spend greater effort and more time probing his memory for any shred of evidence that the term is indeed familiar. This idea is consistent with the results presented in Chapter III, that people with greater self-perceived knowledge reported using a more deliberative thinking process when assessing their knowledge of the terms. Thus, those who view themselves as knowledgeable may exert extra effort deliberatively considering the foils, which may shake loose more content that could conceivably be related to the foil. This process may plausibly result in a stronger feeling of knowledge about the foils. It is interesting to note, however, that another study in the same chapter found that self-perceived knowledge was not significantly related to the number of associations people generated about the foils. One way to reconcile the two findings is by taking into account that in this study, participants were explicitly tasked with generating associations. When asked to try to make mental connections with the terms, perhaps people can come up with a similar number of associations, but they may naturally undertake this process to varying degrees. In line with previous research on motivated reasoning and confirmation bias (e.g., Ditto & Lopez, 1992), those who are motivated to feel familiarity may work harder to find it.

The relationship between genuine knowledge and overclaiming also leaves questions for future exploration. Although the studies in Chapter III provide hints for what mediates the link between the two, the question of causality remains to be addressed. Does genuine knowledge *decrease* overclaiming, or are the two associated in a different way? For instance, it is possible that people who have a good sense of what they know and what they do not know are the ones who seek out knowledge and end up experts. To make a convincing argument for causality,

future work would have to randomly assign participates to gain knowledge in one topic or the other, and measure their overclaiming of terms related to both topics, both before and after the learning phase. A causal account would predict that people would overclaim less *after* they had gained genuine knowledge than before, but only about terms drawn from the topic in which they became more knowledgeable. This would suggest that learning – gaining genuine knowledge – truly protects against certain mistakes of self-judgments of knowledge. One potential issue with the proposed design is that learning may enhance not only genuine knowledge but also self-perceived knowledge, which, as demonstrated in Chapter II, has a positive effect on overclaiming. It would therefore be critical to minimize the increase in self-perceived knowledge and to control for this variable in the analyses.

Though the investigation presented in this packet has focused on overclaiming in the research laboratory, so to speak, is it interesting to consider overclaiming in everyday life. People are unlikely to frequently encounter social psychologists asking them what they know about meta-toxins, but people *are* likely to find themselves in situations where they are asked to assess their knowledge of things that they have never actually heard of. Claiming to have knowledge of such things would functionally be overclaiming. Under what circumstances might people overclaim in the wild, and to what end or ends?

The social cost of overclaiming – of proclaiming to know something you do not know – can be high. It is probably quite embarrassing to realize that you have overestimated your knowledge and be forced to admit both your mistake and your ignorance. The segment *Lie Witness News* on the late-night talk show *Jimmy Kimmel Live!* provides an extreme example of the embarrassing results of overclaiming. In this segment, a camera crew interviews people about their opinions on ostensibly real issues that are, in fact, entirely made up. In one episode, for

example, people on the streets of Los Angeles were asked to give their opinion on the crisis in Wakanda, a fictional country created by *Marvel Comics*. In another, attendees of the music festival Coachella were asked for their thoughts on made up bands such as *The Hillary Clintons* and *The Obesity Epidemic*. The interviewees who provided answers – who overclaimed – were then mocked on television for the amusement of millions of viewers.

If overclaiming is partly the result of a motivational process, why do people not err on the side of underclaiming and save themselves the embarrassment? There are at least two potential explanations worth exploring in future work. First, although the cost of overclaiming is potentially high, it is also uncertain. The probability of being called on to actually provide proof of the claimed knowledge may often be small or perceived as such. In addition, there is a more imminent risk of embarrassment by proclaiming ignorance. If people in your social circle are discussing a book that sounds vaguely familiar, your desire to avoid embarrassment or exclusion may lead you to rely on that faint sense of familiarity to justify nodding along. Second, in some instances overclaiming may be useful for social interactions. When your friend asks you if you're familiar with a particular issue, she likely want to discuss it with you, perhaps to create a connection with you. Trying to find a way in which you *are* familiar with the issue greases the wheel of social interaction.

Speaking About Others in Professional Contexts

Chapter IV examined one way in which a person's gender influences the way we speak about her or him. Four studies looking at diverse professional fields, including academia, politics, and science, among others, consistently found that people are more likely to refer to male professionals than to female professionals by just their surname. Thus, there may be a gender bias in the use of surname to refer to professionals. An alternative explanation I explore

in the chapter is that use of surname may be determined not by gender per se but by factors that tend to co-vary with gender. For example, male professionals are potentially, on average, better known than female professionals, and it may be their fame that is driving the differential use of surname. However, two pieces of information suggest that this is unlikely to entirely account for the gender differences. First, even when such variables – including fame, position, and others – were measured and controlled for, either individually or simultaneously, men were still more likely to be referred to by surname than were women. Second, the gender bias emerged even in an experimental study where the only factor that differed between two fictional professionals was their gender. Though we should be careful in giving too much weight to the effect size of any individual study, it is worth noting that in this study, use of surname was four times larger when the professional was a man vs. a woman.

As noted in the introduction, the way we think about people shapes the way we talk about them, but the reverse can also be true. Five additional studies in Chapter IV suggest that the subtle difference in reference type may have career-relevant consequences. Fictional researchers who were referred to by surname were judged as more famous and eminent than those referred to by full name, and were consequently seen as higher status, more likely to win a prize for their work, and more deserving of a career award. Taken together, the two sets of studies found that people more often refer to male vs. female professionals by surname, and that surname use is associated with eminence, fame, and possibly other favorable judgments. Clearly, additional research is needed to better understand how reference type influences perceptions of real male and female professionals. However, the extant data naturally bring to mind the possibility that this subtle gender bias in use of surname may give male professionals an unfair advantage over

female professionals and may thus contribute to the unequal state of women in certain professional fields.

There may be small but accumulating effects of hearing people referred to by surname or full name, as happens in casual conversations with colleagues, in meetings, at conferences, and in talks and presentations. The data suggest that the professional who is referred to by surname – who is more likely to be a man – may be perceived as more famous and eminent. Research shows that perceived eminence often goes hand in hand with various career-related benefits. For example, reviewers who are aware of the authors' identities accept more papers from famous and high status authors than reviewers in a double-blind review process (Tomkins, Zhang, Heavlin, 2017; Simcoe & Waguespack, 2011). Interestingly, reviewers who are aware of authors' gender are also less likely to accept papers by female authors than reviewers who are double blind (Roberts & Verhoef, 2016).

In addition to paper acceptances, it is easy to think of other situations in academia in which fame could grant a boost. For example, when deciding on colloquia talks, organizers may tend to invite the people who come to mind, and people who are seen as famous may be more likely to come to mind. Their names may come up more in conversations about who to invite. Colloquia talks add to a researcher's visibility, can help foster new collaborations, and can even lead to job offers. One decision can make a large impact on one's career. An analysis of a large academic funding program provides a striking demonstration of this: it found that researchers who were *just* above the funding threshold went on to receive more than twice as much funding in the subsequent eight years than those who were *just* below the funding threshold (Bol, de Vaan, van de Rijt, 2018). A small disparity in the review scores can snowball into a major difference in funding, which can have far-reaching implications for a professional's career.

In addition to exploring the potential consequences of the gender bias in use of surname, it will also be important to discover what underlies the bias; in other words, why do people use surname references more often for men than for women? One hypothesis is that surnames are associated with maleness – that surnames just sound more male to people in general; calling a woman by her surname therefore does not come naturally. If this is the case, it pushes the question back: why are surnames associated with maleness? Perhaps it is because in the U.S. and in many other countries, men typically keep their surname throughout their lives, whereas women often assume their male partner's surname upon marriage (Goldin & Shim, 2004). Women's surnames may therefore be seen as more changeable and less closely related to a woman's identity than is the case for a man. Potentially another piece of the puzzle is that children of both genders traditionally receive their surname from the father (Johnson & Scheuble, 2002). Thus, even women's surnames originate with their fathers – with men.

Another hypothesis rests on the observation that man is the assumed default gender (Hamilton, 1991; Miller, Taylor, & Buck, 1991). Thus, discussing a professional using surname alone can create the incorrect impression that the professional is a man. Given that first names are often gendered, including a professional woman's first name communicates her gender and sidesteps this potential misperception. It remains an empirical question whether man is *always* implicitly assumed to be the default or whether the default can change – for example, in a female-dominated field. If women make up the vast majority of professionals within a field, people may not mark a woman's gender by using her first name because it would be unnecessary, and the gender bias may be reduced or eliminated.

It is also an open question whether – or when – the decision to communicate gender is deliberate vs. automatic and non-conscious. In some cases, egalitarian-minded people may wish

to deliberately clarify a female professional's gender to call attention to women's representation, especially if she is from a male-dominated field. Using the full name may seem like one effective way to do so. However, the attempt to highlight women's contributions by using their full name may ironically lead them to appear less eminent and high status.

A final point that is important to keep in mind is that the gender bias in use of surname may vary from country to country. The norms of calling professionals by their surnames vary by country and culture, as do the norms that dictate how surnames are assumed and under what circumstances they change. The preceding discussion about the consequences and mechanism of this gender bias will not be equally relevant for all counties and cultures. It will be interesting in future work to test which countries show the gender bias. The results may also inform our understanding of the effect's mechanism; for instance, if people use full name in order to identify a woman's gender, then the gender bias may be smaller in countries with a more gendered language than English because the gender would be clear even when surname alone is used.

Conclusion

The package of papers presented here is broadly about the way people think and speak about themselves and others in the context of knowledge and expertise. First, the results shed light on the way people assess their own knowledge, and imply that the intuitive view of knowledge assessment as straightforward and as based on direct access to stored knowledge may not adequately describe the process. Second, the results reveal a novel gender bias in the way people refer to others who have gained a high degree of expertise: professionals. Specifically, the results indicate that people refer to male and female professionals differently, which may affect important career-related judgments and evaluations.

REFERENCES

Apple (2017). 2017 employer information report. Available at https://www.apple.com/diversity/pdf/2017-EEO-1-Consolidated-Report.pdf

Associated Press (2016). *The 2016 associated press stylebook and briefing on media law*. New York, NY: Basic Books.

Benjamin, A.S., Bjork, R.A., & Schwartz, B.L. (1998). The mismeasure of memory: when retrieval fluency is misleading as a metamnemonic index. Journal of Experimental Psychology: General, 127, 55–68.

Blau, F. D., & Kahn, L. M. (2017). The gender wage gap: Extent, trends, and explanations. *Journal of Economic Literature, 55*, 789-865.

Bol, T., de Vaan, M., & van de Rijt, A. (2018). The Matthew effect in science funding. *Proceedings of the National Academy of Sciences*, *115*, 4887-4890.

Boroditsky, L., Schmidt, L., & Phillips, W. (2003). Sex, syntax, and semantics. In D. Gentner & S. Goldin-Meadow (Eds.), *Language in mind: Advances in the study of language and cognition*. Boston: MIT Press.

Brown, A. S. (1991). A review of the tip-of-the-tongue experience. *Psychological Bulletin, 109,* 204-223.

Catalyst (2018). *Women CEOs of the S&P 500*. New York: Catalyst. Available at http://www.catalyst.org/knowledge/women-ceos-sp-500.

Colom, R., García, L. F., Juan-Espinosa, M., & Abad, F. J. (2002). Null sex differences in general intelligence: Evidence from the WAIS-III. *The Spanish Journal of Psychology*, *5*, 29-35.

Ditto, P. H., & Lopez, D. F. (1992). Motivated skepticism: Use of differential decision criteria for preferred and nonpreferred conclusions. *Journal of Personality and Social Psychology*, *63*(4), 568-584.

Duncan, B. L. (1976). Differential social perception and attribution of intergroup violence: Testing the lower limits of stereotyping of blacks. *Journal of personality and social psychology*, *34*, 590-598.

Facebook (2018). Facebook 2018 diversity report. Available at https://newsroom.fb.com/news/2018/07/diversity-report/

Fortune (2018). Fortune 500. Available at http://fortune.com/fortune500/.

Goldin, C., & Shim, M. (2004). Making a name: Women's surnames at marriage and beyond. *Journal of Economic Perspectives*, *18*, 143-160.

Google (2018). Google diversity annual report 2018. Available at https://static.googleusercontent.com/media/diversity.google/en//static/pdf/Google_Diversity_annual_report_2018.pdf.

Gorham, B. W. (2006). News media's relationship with stereotyping: The linguistic intergroup bias in response to crime news. *Journal of communication*, *56*, 289-308.

Gu, J., Zhong, C. B., & Page-Gould, E. (2013). Listen to your heart: When false somatic feedback shapes moral behavior. *Journal of Experimental Psychology: General*, *142*, 307-312.

Hamilton, M. C. (1991). Masculine bias in the attribution of personhood: People= male, male= people. *Psychology of Women Quarterly*, *15*, 393-402.

Hart, J. T. (1967). Memory and the memory-monitoring process. *Journal of Verbal Learning and Verbal Behavior, 6,* 685-691.

Hegarty, M., & Sims, V. K. (1994). Individual differences in mental animation during mechanical reasoning. *Memory & Cognition*, *22*(4), 411-430.

Hoffman, M., Gneezy, U., & List, J. A. (2011). Nurture affects gender differences in spatial abilities. *Proceedings of the National Academy of Sciences*, 201015182.

Hogg, M. A., & Reid, S. A. (2006). Social identity, selfcategorization, and the communication of group norms. *Communication Theory, 16*, 7–30.

Hodson, G., Dovidio, J. F., & Gaertner, S. L. (2002). Processes in racial discrimination: Differential weighting of conflicting information. *Personality and Social Psychology Bulletin*, *28*, 460-471.

Johansson, P., Hall, L., Sikström, S., & Olsson, A. (2005). Failure to detect mismatches between intention and outcome in a simple decision task. *Science*, *310*, 116-119.

Johnson, D. R., & Scheuble, L. K. (2002). What should we call our kids? Choosing children's surnames when parents' last names differ. *The Social Science Journal*, *39*, 419-429.

Koriat, A. (1993). How do we know that we know? The accessibility model of the feeling of knowing. *Psychological Review*, *100*, 609-639.

Koriat, A. (1995). Dissociating knowing and the feeling of knowing: Further evidence for the accessibility model. *Journal of Experimental Psychology: General*, *124*, 311-333.

Koriat, A. (2008). Easy comes, easy goes? The link between learning and remembering and its exploitation in metacognition. *Memory & Cognition*, *36*, 416-428.

Koriat, A., Ackerman, R., Lockl, K., & Schneider, W. (2009a). The memorizing-effort heuristic in judgments of learning: A developmental perspective. *Journal of Experimental Child Psychology*, *102*, 265–279.

Koriat, A., Ackerman, R., Lockl, K., & Schneider, W. (2009b). The easily learned, easily-remembered heuristic in children. *Cognitive Development*, *24*, 169-182.

Koriat, A., & Lieblich, I. (1977). A study of memory pointers. *Acta Psychologica*, *41*, 151-164.

Loftus, E. F., & Palmer, J. C. (1974). Reconstruction of automobile destruction: An example of the interaction between language and memory. *Journal of verbal learning and verbal behavior*, *13*, 585-589.

Loftus, E. F., & Pickrell, J. E. (1995). The formation of false memories. *Psychiatric annals*, *25*, 720-725.

Maass, A., Salvi, D., Arcuri, L., & Semin, G. R. (1989). Language use in intergroup contexts: The linguistic intergroup bias. *Journal of Personality and Social Psychology, 57,* 981–993.

Matvey, G., Dunlosky, J., & Guttentag, R. (2001). Fluency of retrieval at study affects judgments of learning (JOLs): an analytic or nonanalytical basis for JOLs? Memory & Cognition, 29, 222–233.

Miller, D. T., Taylor, B., & Buck, M. L. (1991). Gender gaps: Who needs to be explained?. *Journal of Personality and Social Psychology*, *61*, 5-12.

National Conference of State Legislatures (2018). Women in state legislatures for 2018. Available at http://www.ncsl.org/legislators-staff/legislators/womens-legislative-network/women-in-state-legislatures-for-2018.aspx

National Science Board (2016). Science and engineering indicators 2016, *National Science Foundation, NSB-2016-1*. Arlington, VA. Available at https://www.nsf.gov/statistics/2016/nsb20161/#/data

Nelson, T. E., Biernat, M. R., & Manis, M. (1990). Everyday base rates (sex stereotypes): Potent and resilient. *Journal of Personality and Social Psychology*, *59*, 664-675.

Pronin, E. (2009). The introspection illusion. *Advances in experimental social psychology*, *41*, 1-67.

Roberts, S. G., & Verhoef, T. (2016). Double-blind reviewing at EvoLang 11 reveals gender bias. *Journal of Language Evolution*, *1*, 163-167.

Roediger, H.L., & McDermott, K. B. (1995). Creating false memories: Remembering words not presented in lists. Journal of Experimental Psychology: Learning, Memory and Cognition, 24, 803–814.

Simcoe, T. S., & Waguespack, D. M. (2011). Status, quality, and attention: What's in a (missing) name?. *Management Science*, *57*, 274-290.

Swim, J. K., Aikin, K. J., Hall, W. S., & Hunter, B. A. (1995). Sexism and racism: Old-fashioned and modern prejudices. *Journal of personality and social psychology*, *68*, 199.

Swim, J. K., & Sanna, L. J. (1996). He's skilled, she's lucky: A meta-analysis of observers' attributions for women's and men's successes and failures. *Personality and Social Psychology Bulletin*, *22*, 507-519.

Taylor, D. M., & Jaggi, V. (1974). Ethnocentrism and causal attribution in a South Indian context. *Journal of cross-cultural Psychology*, *5*, 162-171.

The American Association of University Women (2016). Barriers and bias: The status of women in leadership. Available at https://www.aauw.org/research/barriers-and-bias/

Tomkins, A., Zhang, M., & Heavlin, W. D. (2017). Reviewer bias in single-versus double-blind peer review. *Proceedings of the National Academy of Sciences*, *114*, 12708-12713.

United Nations Educational, Scientific, and Cultural Organization (2018). Women in Science. Available at http://uis.unesco.org/sites/default/files/documents/fs51-women-in-science-2018-en.pdf

Valins, S. (1966). Cognitive effects of false heart-rate feedback. *Journal of personality and social psychology*, *4*, 400-408.

Wade, K. A., Garry, M., Read, J. D., & Lindsay, D. S. (2002). A picture is worth a thousand lies: Using false photographs to create false childhood memories. *Psychonomic Bulletin & Review*, *9*, 597-603.

Yaniv, I., & Meyer, D. E. (1987). Activation and metacognition of inaccessible stored information: Potential bases for incubation effects in problem solving. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 13,* 187-205.