

THE DUNNING–KRUGER EFFECT: ON BEING IGNORANT OF ONE’S OWN IGNORANCE

David Dunning

Contents

1. Two Assertions About Ignorance	249
1.1. Ignorance is prevalent in everyday life	249
1.2. Ignorance is often invisible to those to suffer from it	250
1.3. Overview of chapter	251
2. Why Ignorance is Invisible	251
2.1. Ignorance lies in the realm of unknown unknowns	252
2.2. Ignorance is disguised by domain-specific misbeliefs	256
2.3. Ignorance is disguised by “reach-around” knowledge	257
2.4. A threshold condition for lack of recognition	259
3. The Dunning–Kruger Effect	259
3.1. Definition	260
3.2. The double burden of incompetence	260
3.3. Expertise and metacognitive judgment	261
3.4. Empirical demonstrations	262
4. Alternative Accounts	265
4.1. Regression to the mean	265
4.2. Noise plus bias	266
4.3. Lack of incentives	268
5. The Errors of Top and Bottom Performers Compared	270
5.1. Counterfactual comparisons	271
5.2. Impact of social comparison information	272
5.3. The paradox of gaining expertise	274
6. Sources of Self-evaluation	275
6.1. The issue of indirect indicators	276
6.2. The problem of “rational errors”	277
6.3. The impact of preconceived notions of skill	280
6.4. Preconceived notions “versus” bottom-up experience	282
7. Outstanding Issues	285
7.1. Individual differences in meta-ignorance	285
7.2. Perseverance in ignorance	286

Department of Psychology, Cornell University, Ithaca, New York, USA

Advances in Experimental Social Psychology, Volume 44
ISSN 0065-2601, DOI: 10.1016/B978-0-12-385522-0.00005-6

© 2011 Elsevier Inc.
All rights reserved.

7.3. Boundary conditions to the invisibility of incompetence	288
7.4. Can ignorance be bliss?	288
8. Concluding Remarks	289
Acknowledgments	290
References	290

Abstract

In this chapter, I provide argument and evidence that the scope of people's ignorance is often invisible to them. This *meta-ignorance* (or ignorance of ignorance) arises because lack of expertise and knowledge often hides in the realm of the “unknown unknowns” or is disguised by erroneous beliefs and background knowledge that only appear to be sufficient to conclude a right answer. As empirical evidence of meta-ignorance, I describe the Dunning–Kruger effect, in which poor performers in many social and intellectual domains seem largely unaware of just how deficient their expertise is. Their deficits leave them with a double burden—not only does their incomplete and misguided knowledge lead them to make mistakes but those exact same deficits also prevent them from recognizing when they are making mistakes and other people choosing more wisely. I discuss theoretical controversies over the interpretation of this effect and describe how the self-evaluation errors of poor and top performers differ. I also address a vexing question: If self-perceptions of competence so often vary from the truth, what cues are people using to determine whether their conclusions are sound or faulty?

Allow me to begin this chapter with a stipulation that I hope will not be too controversial. That stipulation is that people conduct their daily affairs under the shadow of their own inevitable ignorance. People simply do not know everything about everything. There are holes in their knowledge, gaps in their expertise. I, for example, can name many areas in which my knowledge is incomplete, if it even begins at all. I am not up on the latest developments in hydrostatics and hydraulic circuitry design. I do not know much about the highlights of twentieth century Zimbabwean sculpture. I am not your “go to” guy when it comes to good restaurants in Düsseldorf, Germany.

Of course, one might concede the inevitability of ignorance, but argue that most—if not all—of people's ignorance covers obscure topics that carry no implications for their everyday lives. Much like ants fail to suffer because they do not know, or even conceive of, such topics as bebop jazz or quantum mechanics, people may not suffer because the topics they fail to know fall well beyond the issues that actually influence their outcomes in life. Economists, for example, have argued that most ignorance is rational, in that there are several topics for which gaining expertise would just not provide the tangible benefit to make it worthwhile (Downs, 1957).

1. TWO ASSERTIONS ABOUT IGNORANCE

But I believe this stance toward ignorance is mistaken. Instead, I wish to make two assertions about people's inevitable ignorance that makes it a quite relevant issue for their daily lives. Of course, making those assertions convincingly takes some argumentation and, more importantly, data.

1.1. Ignorance is prevalent in everyday life

First, I wish to argue that the boundary where people's knowledge ends and their ignorance begins frequently arrives far sooner than one would expect. That boundary often insinuates itself well within the geography of everyday tasks that determine whether people live happy and effective lives—certainly within the circle of challenges that people typically face over the course of a lifetime.

For example, in contemporary society, people must filter a good deal of news about scientific facts on such important issues as the environment, medical treatment, and biotechnology. In that regard, the National Science Foundation, in its biannual survey of scientific knowledge, finds large gaps in the basic facts of what people know. In its 2008 survey of roughly 1500 United States adults, only about 53% of respondents knew that electrons are smaller than atoms and only 51% could successfully identify that it was the earth that revolved around the sun (rather than the other way around), taking a year for the earth to do it. When asked whether it was better to test a new high blood pressure drug by giving it (a) to 1000 participants or (b) to 500 participants, with an additional 500 receiving a placebo, only 38% gave the correct answer with an appropriate rationale¹ (National Science Board, 2010).

But perhaps science is not a day-to-day activity for typical citizens, so they can be excused for not having basic knowledge about topics they make no direct decisions about. They do, however, make decisions in every election; thus, it is important for citizens to have a basic working knowledge of their government. In a 2009 survey of roughly 2500 American citizens, only half of respondents could name all three branches of the Federal government, only 54% knew that the power to declare war rests with Congress rather than the President, and only 57% could properly identify the role played by the electoral college, with many thinking it “trains those aspiring for higher office” or “supervised the first television debates” (Intercollegiate Studies Institute, 2008).

¹ For a response to be coded as accurate, the respondent must provide an appropriate rationale. Many respondents opt for the placebo group, but do so, for example, to keep the fatality rate down if the drug should prove deadly. This is not coded as accurate (Miller, 1998).

Or, perhaps, people have more expertise about decisions that carry specific and concrete consequences for them, like saving for retirement. Over recent decades, many companies and institutions have moved from “defined benefit” plans, in which the benefits people receive once they retire are fixed, to “defined contribution” plans, in which employers hand over lump sums of money for their employees to invest as they see fit. For defined contribution plans to be successful, employees must be savvy about how to invest. With this as background, studies of financial literacy provide cause for concern. [Lusardi and Mitchell \(2009\)](#), for example, presented respondents with the following two questions, the first probing people’s understanding of interest rates and the second their understanding of inflation.

- Suppose you had \$100 in a savings account and the interest rate was 2% per year. After 5 years, how much do you think you would have in the account if you left the money to grow: more than \$102, exactly \$102, less than \$102.
- Imagine that the interest rate on your savings account was 1% per year and inflation was 2% per year. After 1 year, would you be able to buy more, exactly the same as, or less than today with the money in this account?

Only 56% answered both questions correctly.

In matters of health literacy, an area certainly within everyday concern, the picture remains the same. The Institute of Medicine reports that 90 million people in the United States (with a population of just over 300 million) have substantial difficulty understanding and following health information, thus taking drugs erratically or in ways that undercut their effectiveness ([Nielsen-Bohlman, Panzer, & Kindig, 2004](#)). In one specific study, asthma sufferers were asked to demonstrate how to use an inhaler, with researchers noting whether respondents followed six steps essential for inhalers to be effective (e.g., did the respondent shake the inhaler before using, exhale before taking a puff, wait at least 30 s between puffs). Respondents did not show a high degree of competence, with 48% of those reading at a high school level and 89% of those reading below a third grade level failing to follow three or more of the crucial steps identified ([Williams, Baker, Honig, Lee, & Nowlan, 1998](#)).

1.2. Ignorance is often invisible to those to suffer from it

But it is the second assertion that may be more important, and to which the bulk of this chapter is devoted. That assertion is that people are destined not to know where the solid land of their knowledge ends and the slippery shores of their ignorance begin. In perhaps the cruelest irony, the one thing people are most likely to be ignorant of is the extent of their own ignorance—where it starts, where it ends, and all the space it fills in-between. This is not a matter of trying. It is reasonable to assume that people are a lot like Marcus Tullius Cicero, the eminent Roman orator, who once admitted

that he was not ashamed to confess he was ignorant of what he did not know. The trick is if only he, and we, could figure out what that “what” is.

In the discussion that follows, I will argue that it is nearly impossible, left to one’s own devices, for one to surmise what one does not know. It is an intrinsically difficult task and one that people fail repeatedly (Carter & Dunning, 2008). As such, we should not demand of people that they have some magical awareness of all that they do not know. To be sure, people occasionally can identify pockets of their own incompetence, but they are far from perfect in identifying all of them. Instead, they often believe they act with adequate if not excellent expertise, when instead they misstep out of misunderstanding and miscalculation that they fail to recognize as such. They may think that they are doing just fine when they are, instead, doing anything but.

1.3. Overview of chapter

In this chapter, I begin by describing why ignorance so often slinks around invisibly to those who suffer from it, covering a number of issues that arise because people act out of an inevitable ignorance that they are not in a position to recognize. I then turn to an instance in which this predicament is its most visible and flamboyant—namely, the Dunning–Kruger effect, in which people suffering the most among their peers from ignorance or incompetence fail to recognize just how much they suffer from it. I describe the phenomenon, report the empirical evidence for it, discuss alternative theoretical accounts for it, and lay out some of its many implications. I also discuss the types of “errors” made by top performers—that is, those imbued with ample competence and expertise—and show how they differ from those of poor performers.

Next, I note that although people may have little insight into their own ignorance, they do go ahead with a firm sense that they are knowledgeable about certain topics and tasks. Where do these self-impressions of skill come from? I discuss empirical work in my lab that has documented two sources of people’s self-impressions—sources that are, regrettably, not necessarily tied closely to actual skill. Finally, I end by discussing the open issues that deserve further empirical study. Among those questions is whether people ultimately learn about their deficits? And if not, why not?



2. WHY IGNORANCE IS INVISIBLE

The central assertion of this chapter is that people’s ignorance is often invisible to them—that they suffer, for lack of a better term, a *meta-ignorance*, remaining ignorant of the multitude of ways they demonstrate gaps in knowledge. To be sure, people are often successful in identifying a few

areas where their expertise is lacking, or topics they wish they knew more about—but I would assert that any individual’s mental catalogue of their areas of ignorance is likely to be very incomplete. People’s catalogs are likely to be imperfect because many of their deficits are camouflaged in one of two ways. First, many instances of an ignorance fall into the category of *unknown unknowns*. Second, many instances of ignorance may be obscured because they are hidden behind misbeliefs that people mistake for valid knowledge in the domain in question. Third, people may be able to construct responses on general world knowledge, or “reach-around” knowledge appears to be relevant and reasonable when it really is not.

2.1. Ignorance lies in the realm of unknown unknowns

Consider any complex project, whether it be building a functional building, crafting a winning legal argument, or protecting one’s country from terrorist attack. Information relevant to that project can be broken down into three different categories. First, there are *known knowns*, information that people have and know that they have. Second, there are *known unknowns*, information that people do not have and know that they lack.

But most important to our analysis of ignorance is a third category of information, *unknown unknowns*, information that is relevant to the project but that people do not know they lack. These are considerations that the person does not even conceive of. Questions that people do not know enough to ask. The notion can refer to any piece of information that lies outside a person’s ken. It can refer to potential problems or risks that the person does not anticipate, actions that are essential to attain success that the person does not know about, possible moves or strategies that a decision-maker might make if only that decision-maker knew of their existence, contingencies that one should prepare for if one were forewarned, or even solutions that a decision-maker might arrive at if only they could be intuited.²

It is likely that people are not aware of their ignorance because much of it is stashed in the realm of unknown unknowns. A hypothetical example may make this clearer. Suppose a couple were bringing their newborn baby home and knew that now was the time to childproof their house against risks to their infant. As part of their known knowns, they may know that they have to place gates in front of stairways and barriers around fireplaces. As part of their known unknowns, they may have questions about other potential precautions that they are suspect they may have to take. Should they, for example, do something about their electrical outlets? (Yes, they

² The careful reader will notice one last category of knowledge that has been omitted—unknown knowns. Such a category likely exists, and one would probably have to start any treatment of it with the philosophy of Zizek (2004), who aligns it with ideology. But the notion of unknown knowns is a topic that deserves its own focused discussion. Thus, it lies outside the scope of the present chapter.

should cover them with child-resistant covers but not plastic plugs, which can be pried off.) And once those known unknowns are addressed, the couple breathes a sigh of relief and brings their infant home, confident that they have done an adequate job and that the house is safe. But, worryingly, beyond the couple's realm of awareness may lie that extensive class of unknown unknowns—precautions that the parents should have taken but have no conception of—such as raising the cords of drapes and mini-blind so that the baby does not accidentally become strangled, or moving all household plants to where the baby cannot reach them, lest they turn into a poisonous snack.

The notion of unknown unknowns was made notorious in 2002 in a press conference by the United States Secretary of Defense Donald Rumsfeld when he noted that his department carried the burden of not necessarily knowing all they did not know about terrorist risks facing the United States (Kamen, 2002), but the concept has a long history in design and engineering (e.g., Kossiakoff & Sweet, 2003). Engineers are taught to be vigilant against unknown unknowns, and to test any system they create against any contingency they can think of to best flush out as many unknown unknowns as possible. Architects are asked to calculate the amount of concrete a building needs to remain stable, and then use eight times that amount to guard against unknown unknown dangers that would otherwise identify themselves only after it is too late (Heath, Larrick, & Klayman, 1998).

The notion of unknown unknowns lies also at the center of an emerging yet still unconventional strain in economics and decision theory (e.g., Schipper, 2010). In this area, scholars recognize that decision-makers may not live in the world portrayed in traditional economic analysis, where rational actors have complete information of all possible contingencies and outcomes that may befall them. Instead, actors are left unaware of possible states of the world that might obtain. For example, decision-makers may be asked to play a game in which they have to discover for themselves what the parameters of the game really are (e.g., Halpern & Rego, 2006) rather than having the game explained completely to them.

Given the existence of unknown unknowns, it is not surprising that an accumulating body of evidence from far-flung corners of psychology shows that people seem to know nothing about the gaps in their knowledge. For example, readers often claim to have reached a deep comprehension of a narrative passage yet fail to recognize the direct contradictions contained within (Epstein, Glenberg, & Bradley, 1984; Glenberg, Wilkinson, & Epstein, 1982). They can claim they know how helicopters, flush toilets, and cylinder locks work, but have to back off those claims once they take a stab explaining how those gadgets work (Rozenblit & Keil, 2002). They similarly claim they can explain their favorite political candidate's position on an important social issue but often cannot do so when asked (Alter, Oppenheimer, & Zemla, 2010).

In our own work, we have found that graduate students pursuing degrees in psychology fail to notice shortcomings in their knowledge of research methods. Via e-mail, we presented a national sample of graduate students a task in which they had to critique the methods of four separate studies and then self-evaluate how well they had done. We varied the number of methodological flaws we wove into those studies to see if respondents gave weight to the number of flaws they missed in their self-evaluations of performance. They did not. Respondents appeared to have no magical awareness of “unknown unknown” methodological flaws that were in the materials to spot but that they had missed. Indeed, informing them of the flaws, they had missed caused respondents to significantly lower their self-ratings on their methodological skills—except, interestingly, for skills related specifically to their own research (Caputo & Dunning, 2005, Study 4).

Unknown unknown gaps in knowledge may go unrecognized in everyday life because people fail to have outside agents hovering over them, peppering them with exams that could impolitely expose holes in their knowledge. Students in medical schools, however, often do have such agents hovering around them, eager to assess skills with well-honed, objectively structured exercises. In these circumstances, how often do medical and nursing students show gaps in knowledge that they appear to know nothing about?

The answer appears to be often. [Barnsley et al. \(2004\)](#) asked student interns to perform seven common clinical procedures while being watched by their tutors. The tutors graded the interns along an assessment instrument that had been carefully crafted by consensus among the hospital’s experienced doctors and nurses to contain standards indicating that the intern still needed supervision on the relevant procedure or was so competent that he or she could now teach it to others. The evaluations of the interns and the tutors dramatically disagreed. All the interns, for example, felt they knew venipuncture well enough to teach others, but only 10% of their tutors agreed—with nearly 50% of interns judged as still needing supervision. On bladder catheterization of a male patient, 80% of interns thought they knew the procedure well enough to teach—but none of their tutors concurred, judging that half of the interns were still in need of supervision.

Other studies have discovered similar unknown unknown gaps in clinical knowledge. [Watts, Rush, and Wright \(2009\)](#) asked first-year nursing students to complete an exercise in which they dressed a wound. The nursing student then watched a videotape of their performance along with an instructor. On average, students saw roughly three mistakes in which they could have contaminated the wound, but their instructors on average saw more than six. Students were knowledgeable about how misuse of gloves could have contaminated the wound,

catching 92% of the instances in which their instructors saw an error, but recognized only 15% of the errors coming from mishandling of swabs and 24% from the handling of cleaning solutions. [Vnuk, Owen, and Plummer \(2006\)](#) asked 95 first-year medical students to complete a CPR (cardio-pulmonary resuscitation) exercise and then asked them how well they had done. Only three felt they had “failed” the exercise (which they knew meant that they had missed steps, put them in the wrong order, executed them incorrectly, or moved too slowly), but an expert examiner judged that a full 36 had failed.

People also demonstrate unknown unknown gaps in the possible solutions they can generate to problems. For example, [Deanna Caputo and I \(2005\)](#) presented participants with a popular word puzzle called Boggle, in which participants look over a 4×4 array of letters and try to find strings of letters that form English words. An example of a Boggle array is given in [Fig. 5.1](#), with the word *knife* indicated as it is found in the puzzle.³ We asked participants to find as many words as they could in three Boggle puzzles, spending 3 min on each, and then to rate how well they had thought they had done. We varied the specific puzzles participants confronted, so some participants faced puzzles with many more solutions than did others.

We were interested in whether participants displayed any insight into gaps in their performance. Would they have an adequate understanding of the solutions they had missed? The answer was a clear no, as indicated in [Table 5.1](#), which shows how much weight participants gave to solutions found and missed in their self-evaluations. However, once explicitly informed of the number of solutions they had missed, participants were quite willing to give weight to that number (see [Table 5.1](#)).

A	B	C	D
E	F	G	H
I	J	K	L
M	N	O	P

Figure 5.1 Example of a 4×4 Boggle puzzle array like those used in [Caputo & Dunning \(2005, Study 1\)](#). The letters comprising the word *knife* are highlighted.

³ If you continue reading, I will reveal the number of three-letter plus English words contained in the Boggle puzzle in [Fig. 5.1](#).

The lack of weight given to the number of solutions missed is more comprehensible when other statistics are considered. Participants' estimates of how many solutions they missed bore no correlation ($r = -0.15$, *ns*) with the truth. More importantly, participants tended, on average, to think they had missed 18 possible solutions when, in fact, they missed 154. Not surprisingly, participants lowered their evaluation of their Boggle acumen after seeing the long list of all the solutions they had shown no awareness of (Caputo & Dunning, 2005, Study 1). In a follow-up study, participants bet less money that they could beat another student in a Boggle competition once their errors of omission were pointed out to them (Study 5).⁴

2.2. Ignorance is disguised by domain-specific misbeliefs

Ignorance is also hidden because it is often in disguise. People may believe they possess accurate knowledge in a domain that is, in fact, misguided and misinformed. For example, despite a lifetime of interactions with objects, people possess an intuitive physics that contains many mistakes and misperceptions about how everyday objects move (Proffitt, 1999). Roll a ball into the coiled tube and many people believe that the ball will roll out of the tube in a curving trajectory, when it, in fact, will just go straight (Kaiser, Jonides, & Alexander, 1986). People display common misconceptions about how such everyday items as mirrors (e.g., Hecht, Bertamini, & Gamer, 2005) and bicycles (Lawson, 2006) work. They also display erroneous beliefs about how emotions work in humans (Wilson & Gilbert, 2003; Woodzicka & LaFrance, 2001).

Table 5.1 Weight given (*B*) to solutions found versus missed in self-evaluations before and after participants were informed of all possible solutions to the Boggle puzzles (Caputo & Dunning, 2005, Study 1)

	Self-evaluation	
	Ability	Performance
<i>Before</i>		
Solutions found	0.63***	0.28**
Solutions missed	-0.00	0.06
<i>After</i>		
Solutions found	0.59**	0.36**
Solutions missed	-0.29**	-0.23*

* $p < 0.06$, ** $p < 0.02$, *** $p < 0.01$.

⁴ There are 15 English words in the Boggle puzzle. So that they all may be removed from the realm of the unknown unknowns, they are plonk, knife, mink, knop, jink, glop, fink, pol, nim, lop, kop, jin, ink, fin, and fie.

People also display remarkable misbeliefs about social conditions, with people who are most wrong sometimes expressing the greatest confidence in their beliefs. For example, in a survey of opinions about welfare, Kuklinski and colleagues found that the most confident respondents thought that 25% of families received welfare in the United States (the figure is closer to 7%) and that 80% of those receiving welfare were African-American (the reality is less than half). Respondents who thought that 15% of the federal budget went to welfare were just as confident as those who expressed the truth (1%) (Kuklinski, Quirk, Jerit, Schwieder, & Rich, 2000).

2.3. Ignorance is disguised by “reach-around” knowledge

People also display their ignorance in other remarkable ways that provide a third explanation as to why people often fail to know what they fail to know. In short, researchers have caught people expressing knowledge about topics that researchers know with certainty people cannot know anything about. Why are they certain? They are certain because these topics do not exist.

2.3.1. Over-claiming

In 2003, Paulhus and colleagues asked respondents to rate their knowledge in 150 different topics, ranging from *Napoleon* to *double entendres* to *The Divine Comedy* to *behaviorism* and so on. Sprinkled within those topics were 30 that were merely the invention of the experimenters, such as *El Puente*, *La Neige Jaune*, *choramine*, and *esoteric deduction*. Of the real topics, respondents claimed at least some knowledge of 44% of them. Of the nonexistent topics, respondents claimed the same for roughly 25% of them. Paulhus and colleagues referred to this tendency as *over-claiming*, and described it as a form of self-enhancement that was independent of actual intellectual ability.

2.3.2. Nonattitudes

But Paulhus’s work followed a long tradition in sociological research showing that people frequently express opinions about nonexistent social groups (e.g., the Wallonians), political figures, and government agencies and policies (e.g., the Metallic Metals Act) (Bishop, Tuchfarber, & Oldendick, 1986). These are all topics about which participants, by definition, cannot have any actual knowledge—but substantial numbers of people claim enough background to have formed an opinion. For example, 11% of respondents will provide an opinion about a fictitious “agricultural trade act” and 14% of a “monetary control act” even if given the explicit option of saying they do not know what the act is. If an explicit “don’t know” response option is withheld, the percentages offering an opinion rise to around 36% for each “act” (Bishop et al., 1986). This tendency does not seem to arise

entirely from mere deceit on the part of respondents (Bishop, Oldendick, Tuchbarber, & Bennett, 1980), but rather from some other process.

2.3.3. Reach-around knowledge defined

This other process appears to be an important one, for it may provide yet another explanation for why people claim knowledge for topics about which they are really uninformed or misinformed. The process is that people take cues from the social situation they are in and their general world knowledge to cobble together enough apparent information to form an impression. That is, people reach back or around to any knowledge they have that might appear to be relevant, and then use it to impose some meaning on the questions they are asked and then to form a judgment. That is, they do not use domain-specific information to inform their judgments (how could they, for no domain exists), but instead use more general knowledge—*reach-around knowledge*—that seems like it might be relevant to the task at hand. For example, when asked about a fictitious Agricultural Trade Act, survey respondents frequently make comments about issues that were plausibly relevant (e.g., “Shipments from Japan are killing our products here”), or made responses consistent with their more general attitudes toward the government (Schuman & Presser, 1980).

This reaching back to more general knowledge might also be behind over-claiming. Graeff (2003) asked respondents their impressions of consumer products that did not exist, such as Thompson drill bits, Yamijitsu stereo equipment, and Barjolet cheeses. He found that respondents were more willing to claim knowledge for brands for which there was broad knowledge that they could refer to—such as Yamijitsu stereos, for which respondents could fall back to their general impression of Japanese stereo equipment, and Barjolet cheeses, in which they could rely on any general knowledge they had of French cheeses. Such general knowledge was not available for other brands (such as Thompson drill bits), and he found that people were much less likely to claim any knowledge in those situations.

This reaching around back toward general knowledge may also be behind other instances in which respondents adhere to beliefs even if those beliefs come under presumably definitive challenge. Prasad and colleagues identified respondents who believed that Saddam Hussein had played a role in the attacks of 9/11 and then confronted them with the fact that the federal commission had concluded that he had played no role (Prasad et al., 2009). Of those confronted, 10% directly refuted the commission’s conclusion, arguing from more general knowledge rather than any specific knowledge about the events of 9/11, making assertions such as: “I believe he was definitely involved with it because he was definitely pumping money into the terrorist organizations every way he could. And he would even send \$25,000 to somebody who committed suicide to kill another person, to their family” (Prasad et al., p. 153).

2.4. A threshold condition for lack of recognition

This ability to reach back to general knowledge, importantly, may also provide a boundary condition for when people will claim knowledge they do not have versus profess (correctly) their ignorance. If there is no domain-specific belief or general reach-around knowledge to fall back on, people may rightly claim no knowledge or opinion. Very few people, for example, would volunteer to stand in for a cardiothoracic surgeon to perform a triple-bypass on a friend, presumably because people have no background information, no intellectual scaffolding, with which they can construct a mental model about how to proceed. The same presumably holds true for such esoteric topics as building a rocket engine or reciting the Icelandic Sagas.

However, pubs the world over are filled with football fans (whether it be the American, Canadian, European, or Australian game) who think they can do a better job managing their favorite team than its manager, presumably because a number of informal conversations and arguments with bar mates over the years has led them to conceive of some intellectual scaffolding of thoughts and intuitions that may or may not constitute actual expertise. Likewise, there are likely many DIY (do-it-yourself) home repair enthusiasts who happily go about rewiring the electrical circuits in their house based on watching their neighbor do it once or seeing some program last year on the Home and Garden channel. They may succeed, but safety experts suggest hiring a professional to take care of electrical, plumbing, or roof repairs. The issue is not so much the added cost of having that expert correct any mistakes the DIY'er might make. Rather, the issue is the number of trips to the hospital emergency room that these episodes of home repairs might end in (Leamy & Weber, 2009).

That is, there is a threshold that has to be met for people to make inappropriate claims of expertise. They have to have some fragments of information, enough scaffolding based on domain-specific or general world knowledge, to allow them to cobble together a plausible response. If they cannot do that, they will not make an inappropriate claim. The question, thus, is how commonly can people gather enough information or argument to feel like they have passed that threshold?

3. THE DUNNING-KRUGER EFFECT

Arguing that ignorance tends to be invisible is somewhat difficult, in that people listening to the contention have a hard time resonating with it. If they try to introspect about any unknown unknowns or invisible pockets of ignorance in their own life, they will, by definition, come up empty—leaving the contention to feel a little alien or disputable. But there is a

manifestation of the argument that is quite visible in everyday life and that people do resonate with. It is not the meta-ignorance they witness in themselves; rather, it is the meta-ignorance they witness in others.

3.1. Definition

Specifically, for any given skill, some people have more expertise and some have less, some a good deal less. What about those people with low levels of expertise? Do they recognize it? According to the argument presented here, people with substantial deficits in their knowledge or expertise should not be able to recognize those deficits. Despite potentially making error after error, they should tend to think they are doing just fine. In short, those who are incompetent, for lack of a better term, should have little insight into their incompetence—an assertion that has come to be known as the Dunning–Kruger effect (Kruger & Dunning, 1999). This is the form of meta-ignorance that is visible to people in everyday life.

Thus, the central question is whether the people they spot really do remain innocent of their own deficits even when those deficits are relatively severe. In 1999, Justin Kruger and I decided to examine the extent to which poor performers in knowledge domains reveal any insight about the depth of their shortcomings and lackluster performance. Our strategy was to ask participants to take tests assessing intellectual expertise in such domains as logical reasoning and grammar, as well as tasks assessing social skill. We then asked participants to rate how well they thought they were doing. Over the years, we have done so in two different ways. First, we have asked participants to provide comparative self-evaluations, rating how well they think they are doing relative to their peers. Second, we have asked participants to provide self-evaluations along more “absolute” scales involving no social comparison, such as estimating how many specific questions they think they are getting right on the test presented to them.

Would poor performers understand how badly they did? We predicted that they would not. Much like their more skilled peers, these individuals would select the answers that looked the most sensible to them—and so at the end of the day would think that their overall performance was rather reasonable. Of course, operating from incomplete and corrupted knowledge, they would make many mistakes and not recognize those mistakes as they made them.

3.2. The double burden of incompetence

In essence, we proposed that when it came to judgments of performance based on knowledge, poor performers would face a double burden. First, deficits in their expertise would lead them to make many mistakes. Second, those exact same deficits would lead them to be unable to recognize when

they were making mistakes and when other people were choosing more wisely. As a consequence, because poor performers were choosing the responses that they thought were the most reasonable, this would lead them to think they were doing quite well when they were doing anything but.

This double-curse arises because, in many life domains, the act of *evaluating* the correctness of one's (or anyone else's) response draws upon the exact same expertise that is necessary in *choosing* the correct response in the first place. That is, in the parlance of psychological research, the skills needed to execute the *meta-cognitive* task of judging the accuracy of a response (cf. [Everson & Tobias, 1998](#); [Maki, Jonas, & Kallod, 1994](#)) are precisely the same as those necessarily for the *cognitive* task of producing an accurate response. Need to judge whether one (or someone else) has written a grammatically correct sentence? That act of judgment relies on the same set of skills as the act of forming a grammatically correct sentence in the first place. Want to know if one has constructed a logically sound argument? That act of evaluation depends on the exact same know-how needed to construct a sound argument. Thus, if poor performers suffer deficits in knowledge that failed them when it came time to form correct responses, those exact same deficits would similarly fail them when it came time to judge the worth of those responses. They would not know when their responses were incorrect; they would not know when others formed better ones.

3.3. Expertise and metacognitive judgment

We knew going into this work that previous research supported our analysis. Previous work has shown that strong and poor performers differ in their success at the metacognitive task of evaluating their performance. When people are asked to evaluate responses to individual test items, strong performers anticipate better which individual items they are likely to get right versus wrong than do poor performers. This difference in metacognitive achievement has been discovered in a wide range of tasks, such as students taking an exam ([Shaughnessy, 1979](#); [Sinkavich, 1995](#)), readers indicating how well they comprehend a narrative passage ([Maki & Berry, 1984](#); [Maki et al., 1994](#)), clinicians making mental illness diagnoses ([Garb, 1989](#); [Levenberg, 1975](#)), bridge players indicating their best versus worst moves ([Keren, 1987](#)), pharmacy school graduates seeking licensure ([Austin, Gregory, & Galli, 2008](#)), physics experts knowing which problems will be more difficult ([Chi, Glaser, & Rees, 1982](#)), and tennis players knowing which shots are more likely to be winners ([McPherson & Thomas, 1989](#)). In each case, the judgments of strong performers about which individual responses would meet with success versus failure were more accurate than

the judgments of their less competent peers (although see [Glenberg & Epstein, 1987](#); [Wagenaar & Keren, 1985](#), for null results).

3.4. Empirical demonstrations

But how would this difference between strong and poor performers translate from judgments of individual items to evaluations of overall performance? [Figure 5.2](#) shows the results of one such study examining whether poor performers show any insight into the weakness of their performance. In this particular study, 141 students who had just completed an exam in one of their college courses were asked to evaluate their “mastery of course material” as well as their performance on the specific exam they had just completed. Participants estimated their performances along percentile scales; that is, they estimated the percentage of other students in the course they thought they had outperformed. They also gave us permission to retrieve their actual exam score, so that we could compare their perception of their performance against the reality of it ([Dunning, Johnson, Ehrlinger, & Kruger, 2003](#)).

As [Fig. 5.2](#) shows, there are many observations one can make about how well perceived performance tracks actual performance. In the figure, we used participants’ objective performance to divide them into four groups—from bottom quartile performers up to top quartile performers. As can be seen in figure, three main findings emerge (actually four, but we will withhold discussion of the fourth until later). First, whether one is talking about mastery of course material or performance on the test, respondents tended to think of their performance, on average, as anything but average. Respondents in all four performance groups tended to think they scored above the 50th percentile, or rather the average of the class. Overall, participants thought their mastery of course material lay in the 70th percentile and their test performance in the 68th—well above that which is statistically possible. When asked to estimate their raw score, they overestimated on average by 3 points—perceiving a score of 37 versus a reality of achieving 34 ($p < 0.001$). These findings are not news. People typically tend to hold overly inflated views of their competence and performance—thinking on average that they are outperforming their peers when it is statistically impossible for a group to post, on average, “above-average” performances (see [Alicke & Govorun, 2005](#); [Dunning, 2005](#); [Dunning, Heath, & Suls, 2004](#); [Dunning, Meyerowitz, & Holzberg, 1989](#); [Weinstein, 1980](#)). [Zenger \(1992\)](#), for example, found in a survey of several hundred engineers in two companies that 32% in one company and 42% in the other thought their skill put them in the top 5% of performers in that company—a statistically absurd result.

That bias aside, there was a statistically observable relation between perceived and actual performance ($r = 0.47$ and 0.60 , $p < 0.01$, for percentile and raw score estimates). People who did poorly on the exam

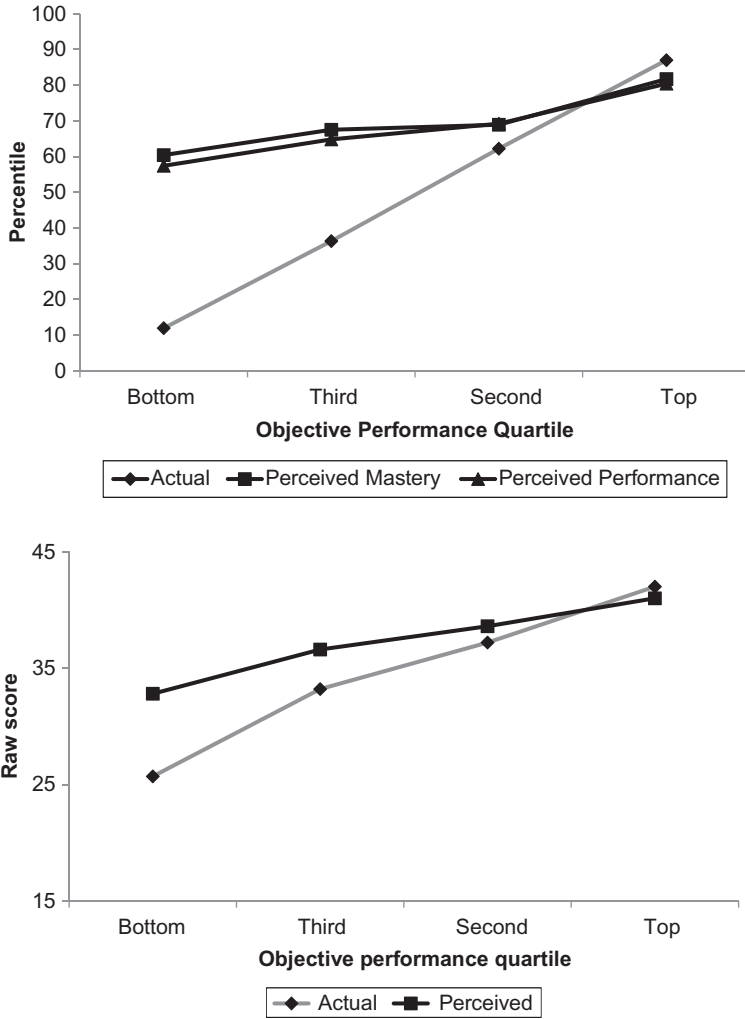


Figure 5.2 Perceived performance as a function of objective performance on a class exam. Top panel presents percentile ratings for perceived mastery of course material and performance on the exam. The bottom panel presents perceived raw score on the exam (out of 45 points). From [Dunning et al. \(2003\)](#), pp. 83–86, by SAGE Publications. Adapted with permission.

intuited that they were doing worse on the exam than those who did well. That said, although it was statistically significant, the relation was quite shallow—with bottom performers, on average, thinking they performed roughly 15–20 percentile points (8 raw score points) worse than top performers’ self-judgments. This finding also replicates a raft of past work,

showing that there tends to be a statistically significant, albeit only weak to meager, relation between what people believe about their skill and the reality as revealed by actual performance (for reviews, see [Dunning, 2005](#); [Dunning et al., 2004](#); [Mabe & West, 1982](#))—whether in the classroom ([Camerer & Hogarth, 1999, 1982](#)), workplace ([Harris & Schaubroeck, 1988](#); [Stajkovic & Luchins, 1998](#)), or the doctor’s office ([Davis et al., 2006](#)).

All this leads to the third—and most central—finding, that people in the bottom 25% of performers, whose actual performance lies in the 12th percentile, thought that their mastery of course material and test performance lay closer to the 60th percentile—a misjudgment of over 45 percentile points. If we look at their raw score estimates, we find that people at the bottom overestimated their raw performance on the test by nearly 30% ([Dunning et al., 2003](#)).

We have observed this pattern of dramatic overestimation by bottom performers across a wide range of tasks in the lab—from tests of logical reasoning and grammar skills ([Kruger & Dunning, 1999](#)) to more social abilities like emotional intelligence ([Sheldon, Ames, & Dunning, 2010](#)) and discerning which jokes are funny ([Kruger & Dunning, 1999](#)). We and others have also observed similar overestimation in real world settings as people tackle everyday tasks, such as hunters taking a quiz on firearm use and safety, based on one created by the National Rifle Association, at a Trap and Skeet competition ([Ehrlinger, Johnson, Banner, Dunning, & Kruger, 2008](#)), and laboratory technicians taking an exam about medical lab procedures and knowledge ([Haun, Zerinque, Leach, & Foley, 2000](#)). In all cases, top to bottom performers provide self-evaluations along percentile scales that largely replicate ([Fig. 5.2](#)).

Similar data have been observed in real world settings on measures other than percentile rankings. Among students taking part in a regional collegiate debate tournament, those among the bottom quartile during preliminary rounds dramatically overestimated the likelihood that they were winning their matches. They thought that they were winning nearly 60% of their matches, when they actually won only 22% of them ([Ehrlinger et al., 2008, Study 2](#)). Of individuals entering chess tournaments, people who possess less skill, as indicated by their Elo rating, mispredicted their tournament performance more than those with greater skill, irrespective of previous experience with tournament chess ([Park & Santos-Pinto, 2010](#)). Novice drivers in the Netherlands and Finland who failed their first driver’s test overestimated how their examiners would rate them to a greater degree than did those who passed the test ([Mynttinen et al., 2009](#)). Of international pharmacy school graduates seeking licensure in Canada, those performing in the bottom 25% provided overly inflated views of how well they performed relative to their peers ([Austin et al., 2008](#)). Medical students receiving the lowest grades (D+) among their peers in clinical clerkships in obstetrics and gynecology overestimated their grades by two full grades, thinking on

average they would receive a B+. The self-underestimates of students receiving an A were trivial by comparison (Edwards, Kellner, Siström, & Magyari, 2003). The worst performers among medical students conducting an exercise in which they interviewed a parent who may be abusing her child rated themselves much more positively than their instructors did (Hodges, Regehr, & Martin, 2001). A survey across 34 countries of the math skills of 15-year-olds discovered that higher math performance was associated with more accurate self-perceptions of math skill (Chiu & Klassen, 2010).

4. ALTERNATIVE ACCOUNTS

My colleagues and I have laid blame for this lack of self-insight among poor performers on a double-course—their deficits in expertise cause them not only to make errors but also leave them unable to recognize the flaws in their reasoning. However, over the past decade, a few psychologists have suggested alternative accounts for the data we have observed in Fig. 5.2 and elsewhere. Seeing such critiques of our work has sharpened our thinking and led us to collect data testing our account more discerningly.

But, perhaps unknown to our critics, these responses to our work have also furnished us moments of delicious irony, in that each critique makes the basic claim that our account of the data displays an incompetence that we somehow were ignorant of. Thus, at the very least, we can take the presence of so many critiques to be *prima facie* evidence for both the phenomenon and theoretical account we made of it, whoever turns out to be right.

That said, the major critiques all deserve discussion.

4.1. Regression to the mean

The most common critique of our metacognitive account of lack of self-insight into ignorance centers on the statistical notion of regression to the mean. Recall from elementary statistics classes that no two variables are ever perfectly correlated with one another. This means that if one selects the poorest performers along one variable, one will see that their scores on the second variable will not be so extreme. Similarly, if one selects the best performers along a variable, one is guaranteed to see that their scores on the second variable will be lower. Stir that observation with the well-known fact that people tend to rate themselves as above average, and one gets the graph displayed in Fig. 5.2 (Krueger & Mueller, 2002).

There are actually two different versions of this “regression effect” account of our data. Some scholars observe that Fig. 5.2 looks like a regression effect, and then claim that this constitutes a complete explanation

for the Dunning–Kruger phenomenon. What these critics miss, however, is that just dismissing the Dunning–Kruger effect as a regression effect is not so much explaining the phenomenon as it is merely relabeling it. What one has to do is to go further to elucidate why perception and reality of performance are associated so imperfectly. Why is the relation so regressive? What drives such a disconnect for top and bottom performers between what they think they have achieved and what they actually have?

The second version of this regression account goes further to explain why the relation between perceived and actual performance is so regressive. One factor that may prompt such an imperfect perception/reality correlation is measurement error. Whenever estimating someone’s level of expertise, there are always errors in the estimate; sometimes people get lucky and post a performance that overstates their true level of know-how; sometimes they get unlucky and post too low a performance. This error, or rather lack of measurement reliability, may cause performance to become untethered to perceptions of expertise, and thus cause [Fig. 5.2](#) ([Krueger & Mueller, 2002](#)).

Fortunately, there are ways to estimate the degree of measurement unreliability and then correct for it. One can then assess what the relation is between perception and reality once unreliability in measuring actual performance has been eliminated. See [Fig. 5.3](#), which displays students’ estimates of exam performance, in both percentile and raw terms, for a different college class ([Ehrlinger et al., 2008](#), Study 1). As can be seen in the figure, correcting for measurement unreliability has only a negligible impact on the degree to which bottom performers overestimate their performance (see also [Krueger & Dunning, 2002](#)). The phenomenon remains largely intact.

4.2. Noise plus bias

[Burson, Larrick, and Klayman \(2006\)](#) extended the regression account to construct a “noise plus bias” explanation for the Dunning–Kruger effect. They accepted the presence of regression effects and suggested that people’s percentile ratings of their performance could be pushed up or down depending on how easy or difficult they perceived the overall task to be. For tasks perceived to be easy, most participants would rate their performance high—thus producing the typical Dunning–Kruger effect of low performers grossly overestimating their performance. However, for tasks perceived as difficult, people would rate their performances much more negatively, causing low performers to rate their performance low—and accurately—whereas high performers would also rate their performance more negatively—and thus provide unduly unfavorable ratings of their performance. In a sense, this would “flip” the typical Dunning–Kruger effect, with high performers now being the ones grossly misestimating their achievements.

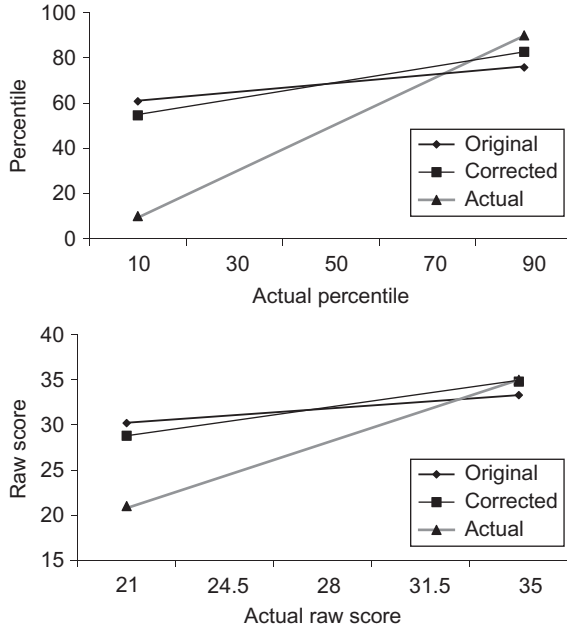


Figure 5.3 Relationship between perceived and actual performance on a course exam before and after correcting for measurement unreliability. The top panel displays percentile estimates of performance as a function of actual percentile. The bottom panel displays perceived raw score as a function of actual raw score (out of 40 points). From Ehrlinger et al. (2008), by Elsevier Publications. Adapted with permission.

In three studies, Burson et al. (2006) gave participants easy versus difficult tasks and obtained data that were largely supportive of their analysis. However, two main issues prevent their analysis from being a more plausible—and accurate—account of the Dunning–Kruger effect. First, instead of using a broad range of tasks, Burson et al. focused on performance on trivia questions (and in a last study a *word prospector* puzzle). As Burson et al. themselves noted, participants may not commonly have had enough intellectual scaffolding to believe that their answers were reasonable ones—an important precondition for the Dunning–Kruger effect to emerge (as noted above and in Kruger & Dunning, 1999).

To be sure, Burson et al. (2006) worried about this issue and showed that their participants performed above chance levels, but it still could have left participants with many experiences in which there were questions they knew they could not answer. Consistent with this interpretation, participants tended to rate themselves as below average across all tasks Burson et al. presented to them—a finding that is quite atypical relative to other research in this area (for reviews, see Dunning, 2005; Dunning et al., 2004). A better set of tasks would have been the types of problem-solving tasks (e.g., logical

reasoning) that Kruger and Dunning (1999) presented to their participants—and which the world tends to present to its inhabitants on a day-to-day basis.

Second, the quarrel that Burson et al. (2006) had with the Dunning–Kruger effect centered on the use of percentile scales, suggesting that people typically did not know how to translate a difficult or easy experience to a percentile evaluation. That is, it was a bit problematic for participants to assess how well they were doing relative to their peers without really knowing how well those peers were doing. Consequently, participants' percentile estimates were “biased” by perceptions of overall task difficulty or ease.

This analysis suggests that people would get their raw scores on any test right. Where they have problems is with translating a raw score into a “social” score in which they compare their achievement to those of their peers. Thus, this reasoning would predict that people would show little Dunning–Kruger effect if they rated themselves along objective scales rather than social or comparative ones. However, we have found that the Dunning–Kruger effect arises even on estimates along more objective scales, as enumerated above. Bottom performers' reliability, and dramatically, overestimate their performances even on assessments that require no social comparison. Top performers, in contrast, tend to neither over- or underestimate how well they are doing on these objective measures—as predicted by the original metacognitive analysis of the Dunning–Kruger effect. To be sure, their estimates are not perfect, but they do not show the type of overwhelming bias found in the estimates of poor performers (Ehrlinger et al., 2008).

4.3. Lack of incentives

A lament that one often hears from economic theorists is that psychologists typically provide no incentives to participants to reach careful, serious, or accurate judgments (e.g., Ariely & Norton, 2007; Camerer & Hogarth, 1999; Hertwig & Ortmann, 2001). As a consequence, participants may provide sloppy estimates, or have those estimates contaminated by motives such as looking good in the eyes of the experimenter. Thus, the inflated self-ratings that poor performing participants provide may fail to reflect what participants really think about themselves. Poor performers may actually have ample insight into the inferior quality of their performance; they just do not want to admit it, either to themselves or to the experimenter.

4.3.1. Money

We have found, however, that providing ample incentives for accurate self-judgments does nothing to enhance the truthfulness of people's assessments of their competence. In one example study, we brought participants in to take a pop quiz on logical reasoning. Roughly half of the participants assessed their performance after being told that they would be paid \$30 if their

estimate of their raw score on the test came within 5% of their true score; \$100 if their estimate was exactly right. No such incentives were mentioned for the remaining participants. As seen in Fig. 5.4, comparing the accuracy of both groups revealed no enhancement in accuracy for the incentive group (Ehrlinger et al., 2008, Study 4). (And no one won the \$100.)

4.3.2. Accountability

Increasing participants' accountability for their self-ratings is a way in which a social incentive can be added. Specifically, asking participants to justify their responses to an authority has been shown to cause people to make more careful and considered judgments (Lerner & Tetlock, 1999) that they




Figure 5.4 Relationship between perceived and actual performance on a course exam with and without financial incentives for accuracy. The top panel displays percentile estimates of performance as a function of actual percentile and financial incentive. The bottom panel displays perceived raw score as a function of actual raw score (out of 20 points) and financial incentive. From Ehrlinger et al. (2008), by Elsevier Publications. Adapted with permission.

imbue with less overconfidence (Tetlock & Kim, 1987). Thus, in a study involving a logical reasoning test, we told roughly half the participants that a supervising professor would interview them for up to 10 min about their answers to the test. This manipulation did nothing to improve the accuracy of participants' views about how well they had done on the test, with poor performers still grossly overestimating how well they had done (Ehrlinger et al., 2008, Study 5).

4.3.3. Behavioral choices

Finally, the choices of poor performers reveal that they do believe their overly optimistic assessments of achievement. Ferraro (2010) offered students some “insurance” about their final exam performance before the exam began. For the price of 10 exam points, participants could purchase insurance that would add 20 points to their exam score if their final exam score fell within the bottom 50%. For 4 exam points, participants could instead buy a 8-point bump in their exam score should they fall between the 50th and 75th performance percentile.

A quick analysis suggests that the first insurance contract should be more popular than the second, in that twice as many people would be eligible to profit from it—with the profit being much larger (10 points) than the alternative (4 points). However, Ferraro found that twice as many students bought the second contract than the first. And of those students buying the second contract, over 80% fell below the 50th percentile in their performance. In essence, they bought insurance thinking that they would place somewhat above the course average in their performance, but their actual performance failed to reach that mark.

5. THE ERRORS OF TOP AND BOTTOM PERFORMERS COMPARED

A discerning reader is most likely already to have discovered the fourth finding inherent in Fig. 5.2. It is not only poor performers who misestimate how well they do. Top performers, as well, tend to underestimate their performances—a finding we have replicated across many settings. However, our data suggest that these misjudgments come from a different source from the misjudgments of poor performers.

Essentially, bottom performers overestimate their proficiency because their intellectual deficits deprive them of the resources necessary to recognize that they are choosing incorrectly. They make the mistake of thinking that all their choices are at least reasonable, or at least the most reasonable they can detect. The problem for top performers is different. They have ample resources to know when they are most likely to be right or wrong in

their choices. They get themselves right. What they get wrong is other people. Because correct answers come relatively easy to them, they mistakenly believe that other people must be coming to the same correct choices. As a consequence, their own performances, albeit good, are not that special relative to how well they think other people are doing.

In a phrase, top performers suffer from a false consensus effect (Ross, Greene, & House, 1977), thinking that other people are responding similarly to themselves much more than other people really are. This assertion is consistent with past work on the attribution of knowledge, which has shown that people, once privy to knowledge, tend to overestimate how much other people possess the same knowledge (Fischhoff, 1975, 1977; Fussell & Krauss, 1991, 1992; Nickerson, Baddeley, & Freeman, 1987). It is also consistent with one pattern of data we have observed across the numerous studies we have conducted on the Dunning–Kruger effect. Top performers consistently underestimate how well they perform on percentile scales—in essence, underestimating how well they are doing relative to their peers. However, on objective or absolute scales (e.g., how many test items answered correctly), we see no consistent evidence of underestimation or overestimation (Ehrlinger et al., 2008).

5.1. Counterfactual comparisons

We have also conducted statistical analyses showing that the errors of top and bottom performers come from different sources, focusing on those percentile estimates that participants provide when they judge how well they are doing relative to their peers. When we give participants a test, what leads them to their best guess about how many of their peers they have outperformed? To begin the analysis, we examined how participants, across several of our studies, combined two other estimates we had asked for to get to their overall percentile estimate. Those two underlying estimates were their perceptions of the raw test score they thought they had achieved and the raw score they believed the average participant had obtained. Not surprisingly, we found that participants tended to provide higher percentile ratings to the extent they thought their underlying raw score was high, and also to the extent that they thought the average participant had done poorly (Ehrlinger et al., 2008). Conducting regression analyses allowed us to gauge the exact weight participants gave to both types of underlying estimates when evaluating their performance in percentile terms.

We then asked a “what if” question: What if participants actually knew the truth of how well (or how poorly) they or the average person had objectively done on the test in that study: how much more accurate would their self-judgment on the percentile measure have been? That is, knowing how much weight participants gave to their own versus the average person’s scores in their percentile estimates, we could estimate how much

participants' percentile self-estimates would change if we replaced their subjective guess about their own (or the average person's score) with the truth. This statistical approach, borrowed from sociology, essentially asks what participants' self-judgments would have been if they lived in a counterfactual world in which they accurately knew either their own objective performance or that of the average peer. Such a technique, called *counterfactual regression analysis*, is commonly used to address such questions as how much a person's IQ would have increased had he or she had hypothetically stayed in high school for one more year (Winship & Korenman, 1997; Winship & Morgan, 2000).

Our statistical exploration showed that the self-rating errors of bottom performers differed in their source from those of top performers. We knew at the start that bottom performers grossly overestimated their own test score performance. Thus, it was no surprise that counterfactually correcting for this overestimation led to significantly more accurate percentile self-ratings, as seen in Fig. 5.5. Bottom-performing participants, who overestimated their performance by 45 percentile points in the original data, would have overestimated their performance by only 15 points had they known their true objective score. Interestingly, bottom performers also tended to overestimate how well their peers, on average, had done. Thus, correcting for this social error alone led to increased error in how bottom performers would have rated themselves—from an overestimate of 45 percentile points to one of 50 points (Ehrlinger et al., 2008).

A similar analysis for top performers produced a different set of conclusions, also as seen in Fig. 5.5. Correcting for top performers' misestimates of their own objective performance would have improved the accuracy their percentile ratings from an underestimate of 14 percentile points to 9 points. However, unlike bottom performers, correcting top performers' beliefs about their peers (they tended to overestimate how well their peers did by an average of 26%) also improved their ratings, from an underestimate of 14 to 8 percentile points (Ehrlinger et al., 2008). That is, the self-evaluation errors of top performers were associated with a mix of mistaken impressions of both self- and peer-performance, whereas the errors of bottom performers were entirely associated with faulty impressions of self-performance.

5.2. Impact of social comparison information

In a sense, calling the phenomenon the Dunning–Kruger effect is a misnomer, in that there is no single phenomenon but rather a family of effects flowing from the fact that people with surfeits of ignorance suffer a double-curse. One additional effect in this family is that poor performers tend to be worse judges of other peoples' competence than top performers. Indeed, when top and bottom performers in grammatical skill are asked to judge the




Figure 5.5 Impact of counterfactual regression analysis in which errors in self- and average peer estimates are corrected. Top panel displays impact of corrections for bottom quartile performers. Bottom panel displays impact of corrections for top quartile performers (Ehrlinger et al., 2008).

grammar of others, top performers provide much more accurate judgments than do bottom performers (Kruger & Dunning, 1999, Study 3).

But from this fact, we can surmise yet another effect of the double-curse. One way to learn about one's own incompetence is by observing the behavior of other people—that is, using social comparison information. One merely has to see when other people approach a task differently, judge when those other approaches are superior or inferior to one's own, and adjust one's self-view of competence accordingly. But there is a hitch for the bottom performer. What if you cannot reliably intuit which approaches are inferior or superior? If that is the case, then such social comparison information, although it may be abundant, is less useful for the task of gaining self-knowledge.

Table 5.2 Impact of social comparison information on perceived percentile performance of top and bottom quartile performers

Quartile/measure	Before	After	Change
<i>Top quartile</i>			
Grammar ability	71.6	77.2	5.6*
Test performance	69.5	79.7	10.2**
<i>Bottom quartile</i>			
Grammar ability	66.8	63.2	−3.5
Test performance	60.5	65.4	4.9

* $p < 0.05$, ** $p < 0.01$.

From Kruger and Dunning (1999), by the American Psychological Association. Adapted with permission.

We have found this to be the case. As seen in Table 5.2, when bottom performers are shown how other people have responded to a quiz on grammar skill, they fail to revise their opinions of their own aptitude on grammar. The experience of top performers is quite different. They accurately see that their peers are performing less well than they themselves are—that is, their false consensus error is corrected—and thus increase how special or distinctive they believe their own performance and skills to be. Bottom performers, unable to recognize superior performance, do not receive such a corrective benefit (Kruger & Dunning, 1999, Study 3). This difference between top and bottom performers has been replicated among medical students judging their own interviewing skills. After seeing videotapes of other medical students conducting interviews, top performers raise their self-evaluations to better match what their supervisors are saying about them. Bottom performers adjust not a whit (Hodges et al., 2001).

5.3. The paradox of gaining expertise

One final prediction follows from our analysis of the Dunning–Kruger effect. There is an avenue by which bottom performers can be guided toward more accurate self-judgments. If they misjudge themselves because they do not have the intellectual resources to judge superior versus inferior performance, one has merely to provide them with those resources. Of course, this procedure leads to a paradox, in that it renders bottom performers no longer ignorant or incompetent. That is, one way to train incompetent people to recognize their incompetence is to rid them of that incompetence.

We have shown that once poor performers are educated out of their incompetence, they show ample ability and willingness to recognize the errors of their past ways. In one such study, we asked participants to complete a number of Wason selection tasks—a logical reasoning task familiar to students of psychology (Wason, 1966). Not surprisingly,

we found that bottom performers grossly overestimated their performance, thinking that their score on the task lay in the 55th percentile when it, in fact, lay in the 12th.

However, next we took roughly half of our participants and gave them a 20-min training session on how to solve Wason tasks—and then asked them to re-rate how well they had done on the original test. As seen in Fig. 5.6, participants at the bottom dramatically revised their self-judgments. They rated their test performance 19 percentile points more harshly and their overall skill at logical reasoning 10 points more negatively—an irony, in that, if anything, the 20-min lesson we provide participants had led them to be more, not less, skilled in logical reasoning. But with adequate intellectual resources in place, participants proved quite willing to rate themselves negatively when faced with a deficient performance (Kruger & Dunning, 1999, Study 4).

6. SOURCES OF SELF-EVALUATION

So far, in making the case that people do not necessarily know the scope of their ignorance, I have been making a “negative” account, showing why people cannot be expected to know when their responses




Figure 5.6 Self-perceived logical reasoning ability (percentile rating) as a function of objective performance before and after being given a training session on how to address Wason selection tasks. From Kruger and Dunning (1999), by the American Psychological Association. Adapted with permission.

are misguided or misinformed. But this negative account leaves open an important question. People often think their responses are reasonable; they often have some level of confidence in the answers they provide. If people cannot recognize when their responses are mistaken (the negative account), what is the “positive process” that leads people to think generally that their responses are correct (and in a few cases that their responses are suspect)?

To begin the positive account of how people reach their self-evaluative judgments, I must first make clear what people fail to have at their disposal when judging the wisdom of their judgments and choices. What people do not have is a direct-access cue that tells them when they are right or wrong in their conclusions. They possess no grand answer sheet that informs them of the accuracy of their judgments. There exists no Pinocchio’s nose to indicate unequivocally when an answer is a truth or a lie; when it comes to gauging the accuracy of many of life’s decisions, there is no iPhone app for that.

6.1. The issue of indirect indicators

Instead, what people have are indirect cues that are correlated with accuracy, albeit in only an imperfect way (Koriat, 2008a). Across many domains, for example, people are more confident when they reach answers quickly rather than more slowly (Dunning & Stern, 1994; Kelley & Lindsay, 1993; Schwarz, 2004), and judgment speed appears to be a valid indicator of accuracy under usual circumstances (Dunning & Perretta, 2002; Koriat, 2008a; Koriat, Ma’ayan, & Nussinson, 2006). The same can be said for familiarity with the task at hand—either the general topic or the specific elements included in the task. People who consider the overall domain or task elements to be familiar also are more confident, and accurate, in their responses (Griffin, Jee, & Wiley, 2009; Koriat, 2008b).

However, under other circumstances, these usually valuable tealeaf indicators of accuracy can mislead. Decision speed, for example, can be increased by exposing people to answers—both correct and incorrect—making them more confident in whatever conclusion they reach without any concomitant increase in accuracy (Kelley & Lindsay, 1993). Describing a task in a tiny and unfamiliar font makes people less confident that they can successfully complete a task (Song & Schwarz, 2008), irrespective of actual ability. In a similar vein, making a topic or its elements more familiar by exposing participants to them also leads people to be more confident that they can provide correct answers, irrespective of actual accuracy. For example, exposing participants to the equation $45 + 56$ makes people more confident they can calculate the equation 45×56 (Schwartz & Metcalfe, 1992). Asking people questions about China makes the topic more familiar to people, and they become more confident that they can answer other questions about that country (Arkes, Boehm, & Xu, 1991).

6.2. The problem of “rational errors”

In work in our laboratory, we have found that another tealeaf indicator that people rely on is how rational their decisions are. By “rational,” I mean that people follow some overall rule or algorithm to compute their response across similar problems. The more they systematically apply that overall rule, they end up more confident in the quality of those responses. To the extent that they approach each problem with a different rule or strategy, they are less confident.

For many tasks, this makes sense. Mathematics, for example, is a skill that exactly asks people to apply systematic operations to numbers across similar problems to achieve some sort of calculative result. And if one is applying the same overall rule to solve similar sorts of math problems, then one does have evidence that one is solving those math problems correctly. There is, however, a problem. People may be applying the right algorithm or rule to solve a math or logic problem, or any sort of puzzle, but how they be sure they have applied the right algorithm or one fraught with error?

An observation in educational psychology is that schoolroom errors of children are often not haphazard, but are frequently rational in nature. Students are conscientiously following systematic rules, just the wrong ones. For example, if asked to solve the equation $33 - 17$, many students state the answer is 24. They are wrong not because they are sloppy, but because they have an algorithm in their head about what subtraction is; it is just a mistaken algorithm. They assume you take the smaller number in each column and subtract it from the larger one, and so the 1 is correctly subtracted from the first 3, but the second 3 is subtracted from the larger 7. In short, their mistakes are rational in that they follow a rule or algorithm that contains some misunderstanding or glitch that is systematically applied (Ben-Zeev, 1995, 1998).

Other work has connected rationality with positive evaluations of performance, even if those favorable evaluations are unwarranted. This notion of rationality is reminiscent of the distinction made by Tetlock (2005) between *foxes* and *hedgehogs* in his study of expert decision-making. Foxes are flexible and nuanced in their thinking when they strive to predict future events. Hedgehogs approach all predictions with a grand (i.e., rational) theory that they are unwilling to deviate from. Tetlock found that foxes tended to be more accurate in their predictions of future world events—and expressed less exuberant but more appropriate levels of confidence—than did hedgehogs.

In our investigations, we have looked to see whether following a rigid algorithm leads to more favorable perceptions of performance, irrespective of whether that algorithm was right or wrong. In one such investigation, we reanalyzed the data from Kruger and Dunning (1999, Study 4, $n = 140$), in which participants struggled with Wason selection tasks. An example item is

presented in Fig. 5.7. On this logical reasoning task, we assessed how consistently participants approached the task, and observed two interesting patterns (Williams & Dunning, 2010). First, participants who got nearly every item right approached each item in a systematic, rule-based way. This stands to reason: Given that the Wason task is a logical reasoning task, each individual instance of it *should* be approached in the same way. However, we also found that participants who got nearly every item wrong *also* approached the Wason task in an exacting algorithmic way. They had just applied the wrong algorithm (see, e.g., Fig. 5.7), leading them to be mistaken in every single answer they gave.

Figure 5.8 illustrates one consequence of these two patterns, that consistency in approaching the Wason selection task was associated with extreme performance—both good and bad. As seen in the curvilinear trend line from a polynomial regression analysis, participants who were mostly right *and* mostly wrong tended to be the most systematic in their approach to the task.




Figure 5.7 An example item used in the Wason selection task. *Note:* The correct answer is turning over the “A” and “7” cards. When participants make consistent (i.e., rational) errors, they typically turn over only the “A” card, or the “A” and “4” card.




Figure 5.8 Relationship between actual performance (measured in percentiles) and consistency in responding (Williams & Dunning, 2010).

Those posting less extreme performances tended to be more haphazard in their approach the Wason task.

As expected, we observed is that confidence in one's performance more closely followed how systematic participants were in their responses than how accurate they were. This led to an irony that is displayed in [Table 5.3](#), in which we looked at all participants who were completely consistent in their answers, that is, they responded to all items in exactly the same way. Of these participants, 28 solved all items correctly; 8 all items incorrectly. As seen in the table, both groups of participants were indistinguishable in how favorably they viewed their performance. Both, for example, thought they had solved 8–9 items (out of 10) correctly when in fact one group had a perfect score and the other a perfectly opposite result ([Williams & Dunning, 2010](#)).

Further evidence implicates consistency with favorable views of performance with no commensurate rise in accuracy. In one study, we presented participants with a series of caricatures drawn by Hirschfeld, who was famous for embedding the name of his daughter, Nina, into his drawings. Participants were asked to find all the Nina's they could in each drawing. One group was forced to approach the task consistently. The computer covered each drawing with a grid of 20 squares, and then exposed each square of the grid in the same regular sequence. The other group was impelled to approach the task more haphazardly. For each drawing, the sequence in which the squares were exposed was different. Later, participants in the first group were rated their performance more positively than the second group did, even though each performed equivalently—thinking,

Table 5.3 Perceived and actual performance of completely consistent participants comparing those answering all items right versus all items wrong (adapted from [Williams & Dunning, 2010](#))

	Performance		<i>t</i>
	All wrong	All right	
Perceived			
Ability percentile	68.1	76.0	1.03
Test score percentile	66.1	79.2	1.73*
Raw score	8.3	8.9	1.01
Aggregate (standardized)	0.8	0.4	1.50
Actual			
Percentile	8.9	90.0	
Raw score	0	10.0	
<i>n</i>	8	28	

* $p < 0.10$.

for example, that they were more likely to have found all the Ninas imbedded in the caricatures (Williams & Dunning, 2010).

6.3. The impact of preconceived notions of skill

But another cue people rely on suggests that a good chunk of everyone's performance evaluation on any specific task is formed well before they ever hear about that task. People carry with them preconceived notions about whether they are good or bad at math, logic, counseling others, public speaking—the list is endless. And those preconceived notions color people's evaluations of their performances—even their guesses about how well they have objectively done. For example, we gave participants a logical reasoning test and, at the end, asked participants to estimate their raw score on the test as well as their percentile ranking among their peers. Their estimates of their raw score were just as strongly correlated with their preexisting notions of their logical reasoning skill as with their actual raw score. Their percentile self-estimates were even more closely associated with their preconceived notions of ability than with their actual performance (Ehrlinger & Dunning, 2003, Study 1).

The strategy of consulting “top-down” self-views of competence would seem—at first—to be a rational and appropriate strategy to use. And it would be, if people's preconceived notions were strongly correlated with actual performance. As mentioned above, however, decades of research suggest that the impressions people have of their skill are only weakly to modestly correlated with objective performance (for reviews, see Davis et al., 2006; Dunning, 2005; Dunning et al., 2004; Mabe & West, 1982), making the strategy a suspect one on which to rely heavily.

6.3.1. The impact of altering preexisting self-views

Our work has shown other ways in which relying on top-down self-views may influence performance evaluations that have nothing to do with objective performance. In one study, we took a reasoning task, based on GRE analytical items, and found that switching which self-view was relevant to the test significantly altered how well participants thought they performed on it. As seen in Fig. 5.9, when the 10-item test was described as focusing on “abstract reasoning,” a trait our participants stated they had in abundance, participants estimated that they answered 10% more items correctly and ranked themselves 12 percentile points more favorably than when the test was described as an examination of “computer programming skills,” a trait our participants denied having to any positive degree. These differences arose despite the fact that the test was identical regardless of its label, and despite the fact that participants achieved the same scores regardless of label.




Figure 5.9 Perceived and actual performance as a function of label given to test (abstract reasoning vs. computer programming) (Ehrlinger & Dunning, 2003).

In another study, causing participants to question their knowledge of North American geography by asking them, for example, whether they had ever visited Wyoming or Nebraska, made them think they did worse on a subsequent geography quiz compared to a group asked more benign questions, such as whether they had ever visited New York City or California. Such differences in performance estimates arose irrespective of actual performance on the test (Ehrlinger & Dunning, 2003).

6.3.2. Implications for gender and science

This reliance on preconceived self-notions may prevent people from realizing competencies that they have—or at least inhibit them from recognizing that they are doing just as well as their peers. Consider the fact, of tremendous current interest and importance, that men and women enter and stick to careers in computer science, chemical engineering, and earth sciences at stunningly different rates, with men overrepresented relative to women (National Science Foundation, 2000). Women comprise only 22% of the labor force in science and engineering, despite being 56% of the labor force overall (National Science Foundation, 2000), and despite no apparent differences in ability to handle such careers (Seymour, 1992).

Could men and women diverge in their enthusiasm for science because they hold different preconceived notions of their scientific talent that bear no relation to the truth? There is evidence that women tend to think less of their scientific aptitude than men think of theirs (Eccles, 1987)—a finding we replicated within a sample from our own university. Could that different self-impression lead to a cascade of psychological events that cause men and women to diverge on different career paths? To test this idea, we brought




Figure 5.10 Perceived and actual performance of male and female students on a science quiz (Ehrlinger & Dunning, 2003).

men and women college students into the laboratory, gave them a test on scientific concepts, afterward asking them to judge how well they had objectively done. Men and women diverged in their assessments, with women thinking they answered 13% fewer questions right than the men thought, and also believing that their performance lay 17 percentile points lower than what the men thought of theirs (see Fig. 5.10). Both these differences were traceable back to differences in preconceived beliefs about scientific talent, and arose despite the fact that male and female participants performed equally well on the test (Ehrlinger & Dunning, 2003, Study 4).

And these differences in perception mattered. At the end of the session, all participants were asked if they wanted to take part in a “science jeopardy” game show competition being held later in the session by the chemistry and psychology departments. A full 70% of male participants expressed some interest; only 49% of female participants did likewise. This difference was traceable back to the perception but not the reality of how well participants thought they had done on the test just completed (Ehrlinger & Dunning, 2003, Study 4). One can speculate about how many life and career decisions are guided by a similar psychological process that bears no relation to actual ability or achievement.

6.4. Preconceived notions “versus” bottom-up experience

In a sense, the impact of preconceived self-notions presents two mysteries. The first is the exact psychological mechanism that allows such views to influence impressions of objective performance. The second is why the

impact of such self-views is not swamped by concrete “bottom-up” experiences people have as they complete a task—such as whether the concepts in the test seem familiar, whether people answer questions quickly or only after considerable effort, or whether they struggle between different response options they are considering. As noted above, these concrete experiences all influence people’s confidence in their performances, so why do these signals not “crowd out” the impact of more abstract signals coming from top-down notions of self?

We have discovered that these two questions can be addressed by the same answer. The impact of top-down views is hardly drowned out by bottom-up experiences. Instead, top-down views set up expectations that actually *change* people’s bottom-up experiences with a task. People who think they are skilled at a task, for example, think they come to answers more quickly and with less struggle than people who believe they are less skilled. People who think they are skilled feel the concepts and questions they confront are more familiar than do those who are less confident in their ability.

In this way, top-down views of competence act much like other abstract labels that alter the concrete phenomenological experiences people have as they complete a task. Yogurt labeled as “full fat” rather than “low fat” is rated as tastier (Wardle & Solomons, 1994). A bottle of wine is rated as more pleasant, and activates more of the orbitofrontal cortex, when its price is described as \$90 rather than \$10 (Plassman, O’Doherty, Shiv, & Rangel, 2008). People literally see the skin color in a face as darker when it is labeled as an African-American face rather than a European American one (Levin & Banaji, 2006).

Across several studies, we have shown that people’s top-down self-views influence their experiences with a task, which in turn influence their impressions of objective performance. In one such study, students completed an interpersonal perception task after rating their “social perception ability.” For each item on the test, they also described their experience in coming to an answer—such as whether they knew the answer immediately or had to go back and forth between possible answers. At the end of the test, they also indicated how many items they thought they got right. Statistical analysis subsequently revealed that participants’ confidence in their social perception ability significantly predicted how they rated their bottom-up experience with the task, which in turn predicted how well they thought they had objectively performed (Critcher & Dunning, 2009, Study 2).

Other data confirm that top-down self-views color bottom-up experience, and thus impressions of objective performance. In one study, participants were asked to take two different history tests—one designed for the high school level and one for the graduate school level. In fact, the two tests were equivalent and participants did not differ in their performance between the two tests. (Indeed, we counterbalanced across participants which exact test was given which label.) However, participants held a

top-down expectation that they could better handle the high school test, and described the experience of taking the high school test as more benign and familiar (e.g., “This question deals with material I’ve learned before”) than they did the graduate school test. As a consequence of these different “experiences,” participants estimated that they performed significantly better on the high school test than they did the graduate school version (Cricher & Dunning, 2009, Study 4).

One final study firmly established that the capacity of top-down views to influence bottom-up experiences was essential to ultimately shape performance estimates. We replicated the study in which participants completed a test we described as focused either on abstract reasoning or on computer programming skills. However, we varied the timing of this label. Roughly half of participants were given the label before they started the test. The remainder were given the label only after they had completed the test but before they judged their performance. If top-down views influence performance estimates only because they first mold bottom-up experience with the task, the impact of the label should arise only if participants were informed of that label *before* they started the test. Only then did the label have the capacity to influence their concrete experiences with the test. And in this replication, as evidenced in Fig. 5.11, this was exactly what we found.




Figure 5.11 Perceived performance (standardized composite measure) as a function of test label (abstract reasoning versus computer programming) and timing that the label is applied (before or after the test) (Cricher & Dunning, 2009).

Informing participants of the label after the test had no impact on their performance estimates (Cricher & Dunning, 2009, Study 1).

7. OUTSTANDING ISSUES

There are numerous outstanding issues that deserve future research attention regarding people's inability to spot their ignorance in general and the Dunning-Kruger effect more specifically.

7.1. Individual differences in meta-ignorance

First, are there general individual differences in meta-ignorance, or is meta-ignorance a phenomenon that arises in a more domain-specific way? In our original treatment of the Dunning-Kruger effect, we proposed that the phenomenon was best understood as domain-specific. Each individual has his or her own personal pockets of ignorance of which he or she will be unaware (Kruger & Dunning, 1999). I still feel this is a useful way to think of the phenomenon. To be sure, there is such a thing as *g*, that is, general intelligence, but much research shows that people's performance on intellectual tasks can vary greatly from setting to setting (e.g., Ceci & Liker, 1986). Thus, it is likely that pockets of incompetence arise quite independently from general intellectual skill, and people should be prepared accordingly.

7.1.1. Intellectual characteristics

That said, it might be useful to pursue work exploring whether there are any general characteristics that tend to provide or deprive people of insight into their shortcomings. Some of these characteristics may be intellectual in nature, and may involve practical competencies necessary to make it in the contemporary world. For example, literacy has been shown to influence how people perform in a wide variety of settings, from health behavior to job settings to financial decision-making (UNESCO, 2002). Its close cousin, numeracy, or the ability to reason with numbers and mathematical concepts, has been similarly linked to health and economic outcomes (e.g., Reyna, Nelson, Han, & Dieckmann, 2009). It might be the case that those who are less literate or numerate may suffer not only from lack of skill but also from not knowing that there is information they need to seek out.

Some empirical evidence already suggests that people who are more educated (which we can take as a proxy for literacy) are better able to separate what they know from what they do not. In research on nonattitudes, highly educated people are more likely to offer opinions on real topics but to claim ignorance on nonexistent ones, relative to their less

educated peers. Less educated peers, paradoxically, tend to claim greater ignorance on real topics but to offer more opinions on nonexistent ones, suggesting they have a more difficult time separating knowledge from ignorance (Schuman & Presser, 1980; Bishop et al., 1980, 1986).

7.1.2. Motivational characteristics

Other potential characteristics preventing people from recognizing their incompetence may be more motivational in nature, centering on people's tendency to defend their sense of self-worth (see Kunda, 1990; Mele, 1997). To date, there have been some explorations of individual differences associated with self-esteem defense—and these explorations show that people prone to defensiveness do bolster themselves more when given a chance. Narcissism and self-deceptive enhancement predict over-claiming of knowledge about nonexistent concepts (Paulhus, Harms, Bruce, & Lysy, 2003). Narcissism, as well, predicts how well people think they can mind-read the intentions and emotions of others, irrespective of actual performance (Ames & Kammrath, 2004).

However, narcissism appears to have an impact that is independent of competence. High narcissists rate themselves more positively, but their judgments are not less sensitive to actual level of performance. That is, it is not uniquely the high narcissists who miss how poorly they are doing when they do badly (Ames & Kammrath, 2004). They are not the ones responsible for the Dunning–Kruger effect; all poor performers are.

7.2. Perseverance in ignorance

But there may be a way in which motivational or self-defensive characteristics matter. When talking about the Dunning–Kruger effect with laypeople, it often becomes apparent that when people express frustration about the effect, it is not so much the incompetence that bothers them as it is the blowback they receive when they try to intervene. Many poor performers push back. They rebel against the advice; they argue points of view that contradict their own.

We have found that pointing out people's deficits does necessarily induce them to strive to overcome those limitations. In a recent study on emotional intelligence, we revealed to business school students their score relative to national norms and asked if they wanted a book on the “emotionally intelligent manager” that we could sell them at a 50% discount. Of those scoring in the top quartile, 64% wanted the book. Of those in the bottom quartile, only 19% did (Sheldon et al., 2010). In a similar vein, Prasad et al. (2009) found that confronting people with evidence did not necessarily lead them to reconsider their misbeliefs about Saddam Hussein's involvement in the 9/11 tragedy. Among the 49 respondents confronted, only 1 changed his mind, and 7 denied they had ever claimed the link in the

first place. The remaining 41 all refused to change their mind, instead either counterarguing the confronter's evidence, refusing to believe in the evidence's validity, bolstering the attitudes they already had, or simply refusing to engage in any discussion on the matter.

Other work has shown that people do not necessarily learn to anticipate their incompetence even after repeated feedback. Although high performing students in a psychology course became more accurate in predicting their test performance in a class from test to test, low performing students did not—remaining stubbornly optimistic about how well they would do on the next test (Hacker, Bol, Horgan, & Rakow, 2000), a result replicated among students attending undergraduate economics courses (Ferraro, 2010).

What processes might be the sources of people's resistance to recognizing their own ignorance even in the face of direct feedback? Motivational defenses aimed at keeping self-esteem high may very well be behind a high level of pushback. Another source may be people's central worldviews. Nyhan and Reifler (2010) presented voters with newspaper articles that contained false claims, such as that the Bush administration had found weapons of mass destruction in Iraq during the Iraq war. Introducing a correction of that false fact into a newspaper article altered the beliefs of liberal voters but not conservative ones, who were known to support the war more. Similarly, when reading about a false claim that the Bush administration had banned stem cell research, the introduction of a correction changed the belief of conservative readers but not of liberal ones, who maintained their belief in the existence of this nonexistent ban. Tying this resistance more directly to self-esteem concerns, Nyhan and Reifler (2009) found that conservatives were more likely to accept facts and arguments about withdrawing the military from Iraq after completing a self-affirmation exercise designed to quell self-esteem concerns.

But sources of resistance need not all be motivational in nature. Preexisting knowledge itself might be a source of people's pushback. People, for example, counterargue political stances that oppose their own more to the extent that they are politically sophisticated and have more political knowledge (Tabor, Cann, & Kucsova, 2009). Knowledge may make it more difficult for people to assimilate new arguments and tasks. In a recent study, London cab drivers were asked to learn about a hypothetical new area that existed in the middle of London. Their prior knowledge of London greatly interfered with their ability to learn routes through this new district, and they underperformed matched controls (Woollett & Maguire, 2010).

To date, when researchers have looked at how preexisting knowledge might lead to resistance in learning, they have looked at accurate knowledge. One might presume that resistance may be promoted by "knowledge" that is inaccurate in nature as well. That is, if a person has a mistaken

idea of how streets are laid out in London, would that prevent him or her from learning the correct layout? To date, no work has been completed on this issue, but one can predict that any sort of knowledge—accurate or erroneous—may interfere with people’s ability to update that knowledge. Once people believe something, for better or worse, it may be more difficult to alter that belief than if the person knew nothing at all.

7.3. Boundary conditions to the invisibility of incompetence

Key to the analysis guiding this chapter is that often the expertise needed to evaluate knowledge is exactly the same expertise needed to act expertly. But sometimes, one does not have to rely on the same expertise to judge performance as one does to attain it. Could those instances be exceptions to the rule, when people become quite competent at spotting their incompetence? For example, the skills needed to evaluate one’s free throw shooting ability in basketball (e.g., an adequate pair of eyes) are quite distinct from those needed to produce good free throws (e.g., good hand-eye coordination and proper technique).

Past work tends to show that evaluations of performance correlate more highly with reality in those areas in which the skills needed to evaluate performance are clearly different from those needed to produce performance. When it comes to athletic tasks, for example, the correlation between perception and reality of performance tends to hover around 0.47. However, as one moves to domains that are more knowledge-based, the correlation tends to dissolve—to 0.33 for skilled technical knowledge, 0.17 for medical related tasks, 0.28 for job interview skills, 0.20 for general mechanical expertise, 0.17 for interpersonal ability, and 0.04 for managerial skills (Mabe & West, 1982). In one illustrative study, varsity college football players did not differ from their coaches in how they evaluated their strength, speed, and size—arguably because the manner in which players evaluated those qualities differed from the way that strength, speed, and size were produced. However, when it came to traits in which one could argue that the same skills were needed to produce and evaluate performance—such as mental toughness, coordination, and “football sense”—varsity players tended to rate themselves more favorably than did their coaches (Felson, 1981).

7.4. Can ignorance be bliss?

We can also rely on the discussion above to address one enduring and unsettled dispute—whether the optimism and overconfidence that people so often exhibit is beneficial or costly to them (Colvin & Block, 1994; Dunning, 2005; Kurt & Paulhus, 2008; Taylor & Brown, 1988). Often, overconfidence is taken to be an energizer that spurs people on to their

goals, helping them to achieve even unrealistic ones (e.g., Taylor & Brown, 1988). Perhaps meta-ignorance—not knowing all the obstacles and complications along the path to one’s goal—is an advantage. That is, at the moment when people need to motivate themselves to action, it may just be folly to be wise.

Perhaps, it may be important to make a distinction about when ignorance—and the overconfidence it engenders—may be beneficial or costly. The road to a goal often contains two phases. The first is a planning and preparation phase, in which people must map out how they can reach their goal. The second is the actual execution of a plan. Overconfidence may be beneficial in the second phase, when people potentially must energize and persevere to press on to their goals, but it may be deadly in the first phase.

For example, it may be appropriate for a general to incite his or her troops to supreme confidence when the day of battle arrives. However, one would not want that general to be incompetent or overconfident in the weeks of planning leading up to that battle. One would not want that general to deny that more reserve troops are needed, or that protective gear is not necessary, or that the troops have enough ordnance. One would want to make sure that the general has thought out all the contingencies of battle, so that he or she can change plans if the circumstances of engagement change. Thus, it is possible for ignorance and overconfidence to be both an advantage and a disadvantage, depending on whether one is talking about planning and preparation versus execution.

8. CONCLUDING REMARKS

Plato, in his enduring classic, *Apology*, describes a puzzle that his mentor Socrates once had to solve. Socrates’s friend, Chaerephon, had gone to the Oracle of Delphi and asked whether there was anyone wiser than Socrates, to which the Oracle had replied that there was no one else. This both surprised and vexed Socrates, who felt that there were many more citizens who knew more than he did, and so he went out on a search to find the one wiser person that he could bring to the Oracle as a counterexample. He interviewed the politicians, poets, and artisans of Athens, and although he found them all knowledgeable and quite skilled, he also found them to be supremely confident in their expertise and unwilling to acknowledge when their intelligence was either faulty or valueless. In this, Socrates discovered what the Oracle had been talking about. He, alone among all other citizens, recognized that his knowledge and wisdom was trivial next to that of the gods. He knew of his limits, and this insight gained him the slightest of advantages in wisdom over others.

In this chapter, I have asserted the inevitability of each individual's ignorance—and have argued that when this ignorance visits people's decisions and actions, they are likely not to know it. Nowhere is this blindness more perceptible than in the impressions that incompetent performers have of their own intellectual and social achievements, and it is a cautionary tale for the rest of us, because, at times, we are the ones who exchange roles with them. Ignorance makes a habit of sly and artful invisibility. But, perhaps, once we know of the trick, we become a little bit wiser in how to look out for and deal with this mischievous, significant, and hopefully not-too-frequent companion.

ACKNOWLEDGMENTS

I acknowledge the many and essential contributions of coauthors on previous chapters, cited throughout this chapter, who helped breathe life into much of the research described herein. Special thanks go to Justin Kruger, who collaborated on the initial series of studies upon which this review is based.

The writing of this chapter was supported by National Science Foundation Grant 0745806. The views expressed in this chapter do not necessarily reflect those of the Foundation.

REFERENCES

- Alicke, M. D., & Govorun, O. (2005). The better-than-average effect. In M. D. Alicke, D. A. Dunning, & J. I. Krueger (Eds.), *The self in social judgment* (pp. 85–106). New York: Psychology Press.
- Alter, A. L., Oppenheimer, D. M., & Zemla, J. C. (2010). Missing the trees for the forest: A construal level account of the illusion of explanatory depth. *Journal of Personality and Social Psychology, 99*, 436–451.
- Ames, D. R., & Kammrath, L. K. (2004). Mind-reading and metacognition: Narcissism, not actual competence, predicts self-estimated ability. *Journal of Nonverbal Behavior, 28*, 187–209.
- Ariely, D., & Norton, M. I. (2007). Psychology and experimental economics: A gap in abstraction. *Current Directions in Psychological Science, 16*, 336–339.
- Arkes, H. R., Boehm, L. E., & Xu, G. (1991). Determinants of judged validity. *Journal of Experimental Social Psychology, 27*, 576–605.
- Austin, Z., Gregory, P. A. M., & Galli, M. (2008). “I just don’t know what I’m supposed to know”: Evaluating self-assessment skills of international pharmacy graduates in Canada. *Administrative Pharmacy, 4*, 115–124.
- Barnsley, L., Lyon, P., Ralson, S., Hibbert, E., Cunningham, I., Gordon, F., et al. (2004). Clinical skills in junior medical officers: A comparison of self-reported confidence and observed competence. *Medical Education, 38*, 358–367.
- Ben-Zeev, T. (1995). The nature and origin of rational errors in arithmetic thinking: Induction from examples and prior knowledge. *Cognitive Science, 19*, 341–376.
- Ben-Zeev, T. (1998). Rational errors and the mathematical mind. *Review of General Psychology, 2*, 366–383.
- Bishop, G. F., Oldendick, R. W., Tuchbarber, A. J., & Bennett, S. E. (1980). Pseudo-opinions on public affairs. *Public Opinion Quarterly, 44*, 198–209.

- Bishop, G. F., Tuchfarber, A. J., & Oldendick, R. W. (1986). Opinions on fictitious issues: The pressure to answer survey questions. *Public Opinion Quarterly*, *50*, 240–250.
- Burson, K. A., Larrick, R. P., & Klayman, J. (2006). Skilled or unskilled, but still unaware of it: How perceptions of difficulty drive miscalibration in relative comparisons. *Journal of Personality and Social Psychology*, *90*, 60–77.
- Camerer, C., & Hogarth, R. (1999). The effects of financial incentives in experiments: A review and capital-labor production framework. *Journal of Risk and Uncertainty*, *19*, 7–42.
- Caputo, D. D., & Dunning, D. (2005). What you don't know: The role played by errors of omission in imperfect self-assessments. *Journal of Experimental Social Psychology*, *41*, 488–505.
- Carter, T. J., & Dunning, D. (2008). Faulty self-assessment: Why evaluating one's own competence is an intrinsically difficult task. *Personality and Social Psychology Compass*, *2*, 346–360.
- Ceci, S. J., & Liker, J. (1986). A day at the races: IQ, expertise, and cognitive complexity. *Journal of Experimental Psychology: General*, *115*, 255–266.
- Chi, M. T. H., Glaser, R., & Rees, E. (1982). Expertise in problem solving. In R. Sternberg (Ed.), *Advances in the psychological of human intelligence* (Vol. 1, pp. 17–76). Hillsdale, NJ: Erlbaum.
- Chiu, M. M., & Klassen, R. M. (2010). Relations of mathematics self-concept and its calibration with mathematics achievement: Cultural differences among fifteen-year-olds in 34 countries. *Learning and Instruction*, *20*, 2–17.
- Colvin, C. R., & Block, J. (1994). Do positive illusions foster mental health? An examination of the Taylor and Brown formulation. *Psychological Bulletin*, *116*, 3–20.
- Critcher, C. R., & Dunning, D. (2009). How chronic self-views influence (and mislead) self-evaluations of performance: Self-views shape bottom-up experiences with the task. *Journal of Personality and Social Psychology*, *97*, 931–945.
- Davis, D. A., Mazmanian, P. E., Fordis, M., Harrison, R. V., Thorpe, K. E., & Perrier, L. (2006). Accuracy of physician self-assessment compared with observed measures of competence. *JAMA*, *296*, 1094–1102.
- Downs, A. (1957). *An economic theory of democracy*. New York: Harper & Brothers.
- Dunning, D. (2005). *Self-insight: Roadblocks and detours on the path to knowing thyself*. New York: Psychology Press.
- Dunning, D., Heath, C., & Suls, J. (2004). Flawed self-assessment: Implications for health, education, and the workplace. *Psychological Science in the Public Interest*, *5*, 69–106.
- Dunning, D., Johnson, K., Ehrlinger, J., & Kruger, J. (2003). Why people fail to recognize their own competence. *Current Directions in Psychological Science*, *12*, 83–87.
- Dunning, D., Meyerowitz, J. A., & Holzberg, A. D. (1989). Ambiguity and self-evaluation: The role of idiosyncratic trait definitions in self-serving assessments of ability. *Journal of Personality and Social Psychology*, *57*, 1082–1090.
- Dunning, D., & Perretta, S. F. (2002). Automaticity and eyewitness accuracy: A 10- to 12-second rule for distinguishing accurate from erroneous positive identifications. *The 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.
- Eccles, J. S. (1987). Gender roles and women's achievement-related decisions. *Psychology of Women Quarterly*, *18*, 585–609.
- Edwards, R. K., Kellner, K. R., Siström, C. L., & Magyari, E. J. (2003). Medical student self-assessment of performance on an obstetrics and gynecology clerkship. *American Journal of Obstetrics and Gynecology*, *188*, 1078–1082.
- Ehrlinger, J., & Dunning, D. (2003). How chronic self-views influence (and potentially mislead) assessments of performance. *Journal of Personality and Social Psychology*, *84*, 5–17.

- Ehrlinger, J., Johnson, K., Banner, M., Dunning, D., & Kruger, J. (2008). Why the unskilled are unaware? Further explorations of (lack of) self-insight among the incompetent. *Organizational Behavior and Human Decision Processes*, *105*, 98–121.
- Epstein, W., Glenberg, A. M., & Bradley, M. (1984). Coactivation and comprehension: Contribution of text variables to the illusion of knowing. *Memory & Cognition*, *12*, 355–360.
- Everson, H. T., & Tobias, S. (1998). The ability to estimate knowledge and performance in college: A metacognitive analysis. *Instructional Science*, *26*, 65–79.
- Felson, R. B. (1981). Ambiguity and bias in the self-concept. *Social Psychology Quarterly*, *44*, 64–69.
- Ferraro, P. J. (2010). Know thyself: Competence and self-awareness. *Atlantic Economic Journal*, *38*, 183–196.
- Fischhoff, B. (1975). Hindsight \neq foresight: The effect of outcome knowledge on judgment under uncertainty. *Journal of Experimental Psychology: Human Perception and Performance*, *1*, 288–299.
- Fischhoff, B. (1977). Perceived informativeness of facts. *Journal of Experimental Psychology: Human Perception and Performance*, *3*, 349–358.
- Fussell, S. R., & Krauss, R. M. (1991). Accuracy and bias in estimates of others' knowledge. *European Journal of Social Psychology*, *21*, 445–454.
- Fussell, S. R., & Krauss, R. M. (1992). Coordination of knowledge in communication: Effects of speakers' assumptions about what others know. *Journal of Personality and Social Psychology*, *62*, 378–391.
- Garb, H. N. (1989). Clinical judgment, clinical training, and professional experience. *Psychological Bulletin*, *105*, 387–396.
- Glenberg, A. M., & Epstein, W. (1987). Inexpert calibration of comprehension. *Memory & Cognition*, *15*, 84–93.
- Glenberg, A. M., Wilkinson, A., & Epstein, W. (1982). The illusion of knowing: Failure in the self-assessment of comprehension. *Memory & Cognition*, *10*, 597–602.
- Graeff, T. R. (2003). Exploring consumers' answers to survey questions: Are uninformed responses truly uninformed? *Psychology & Marketing*, *20*, 643–667.
- Griffin, T. D., Jee, B. D., & Wiley, J. (2009). The effects of domain knowledge on metacomprehension accuracy. *Memory & Cognition*, *37*, 1001–1013.
- Hacker, D. J., Bol, L., Horgan, D. D., & Rakow, E. A. (2000). Test prediction and performance in a classroom context. *Journal of Educational Psychology*, *92*, 160–170.
- Halpern, J. Y., & Rego, L. C. (2006). Extensive games with possibly unaware players. *Proceedings of the fifth international joint conference on autonomous agents and multiagent systems* (pp. 744–751). New York: Association for Computing Machinery.
- Hansford, B. C., & Hattie, J. A. (1982). The relationship between self and achievement/performance measures. *Review of Educational Research*, *52*, 123–142.
- Harris, M. M., & Schaubroeck, J. (1988). A meta-analysis of self-supervisor, self-peer, and peer-supervisor ratings. *Personnel Psychology*, *41*, 43–62.
- Haun, D. E., Zeringue, A., Leach, A., & Foley, A. (2000). Assessing the competence of specimen-processing personnel. *Laboratory Medicine*, *31*, 633–637.
- Heath, C., Larrick, R. P., & Klayman, J. (1998). Cognitive repairs: How organizational practices can compensate for individual shortcomings. In B. M. Staw & L. L. Cummings (Eds.), *Research in Organizational Behavior* (Vol. 20, pp. 1–37).
- Hecht, H., Bertamini, M., & Gamer, M. (2005). Naive optics: Acting upon mirror reflections. *Journal of Experimental Psychology: Human Perception and Performance*, *31*, 1023–1038.
- Hertwig, R., & Ortmann, A. (2001). Experimental practices in economics: A methodological challenge for psychologists? *The Behavioral and Brain Sciences*, *25*, 383–403.
- Hodges, B., Regehr, G., & Martin, D. (2001). Difficulties in recognizing one's own incompetence: Novice physicians who are unskilled and unaware of it. *Academic Medicine*, *76*, S87–S89.

- Intercollegiate Studies Institute (2008). *The diverging influences of the college degree & civic learning on American beliefs*. Wilmington, DE: Intercollegiate Studies Institute.
- Kaiser, M. K., Jonides, J., & Alexander, J. (1986). Intuitive reasoning about abstract and familiar physics problems. *Memory and Cognition*, *14*, 308–312.
- Kamen, A. (2002). Tent rent: \$10 a day. *Washington Post*, A25.
- Kelley, C. M., & Lindsay, D. S. (1993). Remembering mistaken for knowing—Ease of retrieval as a basis for confidence in answers to general knowledge questions. *Journal of Memory and Language*, *32*, 1–24.
- Keren, G. B. (1987). Facing uncertainty in the game of bridge: A calibration study. *Organizational Behavior and Human Decision Processes*, *139*, 98–114.
- Koriat, A. (2008a). Subjective confidence in one's answers: The consensuality principle. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *34*, 945–959.
- Koriat, A. (2008b). When confidence in a choice is independent of which choice is made. *Psychonomic Bulletin & Review*, *15*, 997–1001.
- Koriat, A., Ma'ayan, H., & Nussinson, R. (2006). The intricate relationships between monitoring and control in metacognition: Lessons for the cause-and-effect relation between subjective experience and behavior. *Journal of Experimental Psychology: General*, *135*, 36–69.
- Kossiakoff, A., & Sweet, W. N. (2003). *Systems engineering: Principles and practice*. New York: Wiley.
- Krueger, J., & Mueller, R. A. (2002). Unskilled, unaware, or both? The contribution of social-perceptual skills and statistical regression to self-enhancement biases. *Journal of Personality and Social Psychology*, *82*, 180–188.
- Kruger, J. M., & 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.
- Kruger, J., & Dunning, D. (2002). Unskilled and unaware—But why? A reply to Krueger and Mueller. *Journal of Personality and Social Psychology*, *82*, 189–192.
- Kuklinski, J. H., Quirk, P. J., Jerit, J., Schwieder, D., & Rich, R. F. (2000). Misinformation and the currency of democratic citizenship. *Journal of Politics*, *62*, 79–816.
- Kunda, Z. (1990). The case for motivated reasoning. *Psychological Bulletin*, *108*, 480–498.
- Kurt, A., & Paulhus, D. L. (2008). Moderators of the adaptiveness of self-enhancement: Operationalization, motivational domain, adjustment facet, and evaluator. *Journal of Research in Personality*, *42*, 839–853.
- Lawson, R. (2006). The science of cycology: Failures to understand how everyday objects work. *Memory and Cognition*, *34*, 1667–1675.
- Leamy, E., & Weber, V. (2009). *Homeowners beware of DIY repair dangers: Do-it-yourself home projects often lead to injuries and increased costs*. <http://abcnews.go.com/GMA/story?id=8373975> (accessed February 2, 2011).
- Lerner, J. S., & Tetlock, P. E. (1999). Accounting for the effects of accountability. *Psychological Bulletin*, *125*, 255–275.
- Levenberg, S. B. (1975). Professional training, psychodiagnostic skill, and kinetic family drawings. *Journal of Personality Assessment*, *39*, 389–393.
- Levin, D. T., & Banaji, M. R. (2006). Distortions in the perceived lightness of faces: The role of race categories. *Journal of Experimental Psychology: General*, *135*, 501–512.
- Lusardi, A., & Mitchell, O. S. (2009). *How ordinary consumers make complex economic decisions: Financial literacy and retirement reactions*. Dartmouth College, Unpublished manuscript.
- Mabe, P. A. III, & West, S. G. (1982). Validity of self-evaluation of ability: A review and meta-analysis. *The Journal of Applied Psychology*, *67*, 280–286.
- Maki, R. H., & Berry, S. L. (1984). Metacomprehension of text material. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *10*, 663–679.

- Maki, R. H., Jonas, D., & Kallod, M. (1994). The relationship between comprehension and metacomprehension ability. *Psychonomic Bulletin & Review*, *1*, 126–129.
- McPherson, S. L., & Thomas, J. R. (1989). Relation of knowledge and performance in boys' tennis: Age and expertise. *Journal of Experimental Child Psychology*, *48*, 190–211.
- Mele, A. R. (1997). Real self-deception. *The Behavioral and Brain Sciences*, *20*, 91–136.
- Miller, J. D. (1998). The measurement of civic scientific literacy. *Public Understanding of Science*, *7*, 203–223.
- Mynttinen, S., Sundstrom, A., Vissers, J., Koivukoski, M., Hakuli, K., & Keskinen, E. (2009). Self-assessed driver competence among novice drivers: A comparison of driving test candidate assessments and examiner assessments in a Dutch and Finnish sample. *Journal of Safety Research*, *40*, 301–309.
- National Science Board (2010). *Science and engineering indicators 2010*. Washington, DC: National Science Foundation.
- National Science Foundation (2000). *Women, minorities, and persons with disabilities in science and engineering*. Arlington, VA: National Science Foundation.
- Nickerson, R. S., Baddeley, A., & Freeman, B. (1987). Are people's estimates of what other people know influenced by what they themselves know? *Acta Psychologica*, *64*, 245–259.
- Nielsen-Bohman, L., Panzer, A. M., & Kindig, D. A. (2004). *Health literacy: A prescription to end confusion*. Washington, DC: Institute of Medicine.
- Nyhan, B., & Reifler, J. (2009). *Opening the partisan mind: Self-affirmation and information processing about the surge in Iraq*. Unpublished manuscript, University of Michigan.
- Nyhan, B., & Reifler, J. (2010). When corrections fail: The persistence of political misconceptions. *Political Behavior*, *32*, 303–330.
- Park, Y.-J., & Santos-Pinto, L. (2010). Overconfidence in tournaments: Evidence from the field. *Theory and Decision*, *69*, 143–166.
- 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.
- Plassman, H., O'Doherty, J., Shiv, B., & Rangel, A. (2008). Marketing actions can modulate neural representations of experienced pleasantness. *Proceedings of the National Academy of Science United States of America*, *105*, 1050–1054.
- Prasad, M., Perrin, A., Bezila, K., Hoffman, S., Kindleberger, K., Manturuk, K., et al. (2009). There must be a reason: Osama, Saddam, and inferred justification. *Sociological Inquiry*, *79*, 142–162.
- Proffitt, D. R. (1999). Naive physics. In R. Wilson & F. Keil (Eds.), *The MIT encyclopedia of the cognitive sciences* (pp. 577–579). Cambridge, MA: MIT Press.
- Reyna, V. F., Nelson, W. L., Han, P. K., & Dieckmann, N. F. (2009). How numeracy influences risk comprehension and medical decision making. *Psychological Bulletin*, *135*, 943–973.
- Ross, L. D., Greene, D., & House, P. (1977). The “false consensus effect”: An egocentric bias in social perception and attribution processes. *Journal of Experimental Social Psychology*, *13*, 279–301.
- Rozenblit, L., & Keil, F. C. (2002). The misunderstood limits of folk science: An illusion of explanatory depth. *Cognitive Science*, *26*, 521–562.
- Schipper, B. C. (2010). *The unawareness bibliography*. <http://www.econ.ucdavis.edu/faculty/schipper/unaw.htm> (accessed November 30, 2010).
- 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.
- Schwartz, B. L., & Metcalfe, J. (1992). Cue familiarity but not target retrievability enhances feeling-of-knowing judgments. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *18*, 1074–1083.

- Schwarz, N. (2004). Metacognitive experiences in consumer judgment and decision making. *Journal of Consumer Psychology, 14*, 332–348.
- Seymour, E. (1992). Undergraduate problems with teaching and advising in SME majors: Explaining gender differences in attribution rates. *Journal of College Science Teaching, 21*, 284–292.
- Shaughnessy, J. J. (1979). Confidence judgment accuracy as a predictor of test performance. *Journal of Research in Personality, 13*, 505–514.
- Sheldon, O., Ames, D., & Dunning, D. (2010). *Self-assessments of emotional intelligence*. Unpublished data, Rutgers University.
- Sinkavich, F. J. (1995). Performance and metamemory: Do students know what they don't know? *Instructional Psychology, 22*, 77–87.
- Song, H., & Schwarz, N. (2008). If it's hard to read, it's hard to do: Processing fluency affects effort prediction and motivation. *Psychological Science, 19*, 986–988.
- Stajkovic, A. D., & Luchins, F. (1998). Self-efficacy and work-related performance: A meta-analysis. *Psychological Bulletin, 124*, 240–261.
- Tabor, C. S., Cann, D., & Kucsova, S. (2009). The motivated processing of political arguments. *Political Behavior, 31*, 137–155.
- Taylor, S. E., & Brown, J. D. (1988). Illusion and well-being: A social psychological perspective on mental health. *Psychological Bulletin, 103*, 193–210.
- Tetlock, P. E. (2005). *Expert political judgment*. Princeton, NJ: Princeton University Press.
- Tetlock, P. E., & Kim, J. I. (1987). Accountability and judgment processes in a personality prediction task. *Journal of Personality and Social Psychology, 52*, 700–709.
- UNESCO, J. I. (2002). *Education for all: Is the world on track? EFA global monitoring report*. Paris: UNESCO.
- Vnuk, A., Owen, H., & Plummer, J. (2006). Assessing proficiency in adult basic life support: Student and expert assessment and the impact of video recording. *Medical Teacher, 28*, 429–434.
- Wagenaar, W. A., & Keren, G. B. (1985). Calibration of probability assessments by professional blackjack dealers, statistical experts, and lay people. *Organizational Behavior and Human Decision Processes, 36*, 406–416.
- Wardle, J., & Solomons, W. (1994). Naughty but nice: A laboratory study of health information and food preferences in a community sample. *Health Psychology, 13*, 180–183.
- Wason, P. C. (1966). Reasoning. In B. M. Foss (Ed.), *New horizons in psychology* (pp. 135–151). Baltimore: Penguin Press.
- Watts, W. E., Rush, K., & Wright, M. (2009). Evaluating first-year nursing students' ability to self-assess psychomotor skills using video. *Nursing Education Perspectives, 30*, 214–219.
- Weinstein, N. D. (1980). Unrealistic optimism about future life events. *Journal of Personality and Social Psychology, 58*, 806–820.
- Williams, N. V., Baker, D. W., Honig, E. G., Lee, T. M., & Nowlan, A. (1998). Inadequate literacy is a barrier to asthma knowledge and self-care. *Chest, 114*, 1008–1015.
- Williams, E., & Dunning, D. (2010). *From formulae to faith: A consistency heuristic underlies mistaken confidence in self-performance*. University of Florida under review.
- Wilson, T. D., & Gilbert, D. T. (2003). Affective forecasting. In M. P. Zanna (Ed.), *Advances in experimental social psychology* (Vol. 35, pp. 345–411). San Diego: Academic Press.
- Winship, C., & Korenman, S. (1997). Does staying in school make you smarter? The effect of education on IQ. In B. Devlin, S. E. Fienberg, D. Resnick, & K. Roeder (Eds.), *Intelligence and success: Is it all in the genes? Scientists respond to the bell curve* (pp. 215–234). New York: Springer-Verlag.
- Winship, C., & Morgan, S. L. (2000). The estimation of causal effects from observational data. *Annual Review of Sociology, 25*, 659–706.

- Woodzicka, J. A., & LaFrance, M. (2001). Real versus imagined gender harassment. *Journal of Social Issues, 57*, 15–30.
- Woollett, K., & Maguire, E. (2010). The effect of navigational expertise on wayfinding in new environments. *Journal of Environmental Psychology, 30*, 565–573.
- Zenger, T. R. (1992). Why do employers only reward extreme performance? Examining the relationships among performance, pay, and turnover. *Administrative Science Quarterly, 37*, 198–219.
- Zizek, S. (2004). Iraq's false promise. *Foreign Affairs, 42–49*, January/February.