Overprecision is a Property of Thinking Systems

Don A. Moore

University of California at Berkeley

Draft of March 11, 2022

In press at *Psychological Review*

© 2022, American Psychological Association. This paper is not the copy of record and may not exactly replicate the final, authoritative version of the article. Please do not copy or cite without authors' permission. The final article will be available, upon publication, via its DOI: 10.1037/rev0000370

Abstract:

Overprecision is the excessive certainty in the accuracy of one's judgment. This paper proposes a new theory to explain it. The theory holds that overprecision in judgment results from neglect of all the ways in which one could be wrong. When there are many ways to be wrong, it can be difficult to consider them all.

Overprecision is the result of being wrong and not knowing it. This explanation can account for why question formats have such a dramatic influence on the degree of overprecision people report. It also explains the ubiquity of overprecision not only among people but also among artificially intelligent agents.

Keywords: Overconfidence, overprecision, unknown unknowns, hypothesis testing, artificial intelligence, bias, decision making, judgment, error neglect

Author note:

The data, analysis code, supplementary results, and preregistrations for the studies presented here are available: <https://osf.io/fcrj7/> Proprietary Watson data supplied by J. William Murdock at IBM Watson Research Center. This paper benefited from helpful comments from Ned Augenblick, Max Bazerman, Clayton Critcher, Amelia Dev, Craig Fox, Uriel Haran, Leif Nelson, Terry Odean, Dmitry Taubinsky, Tory Taylor, Liz Tenney, and Dan Walters as well as seminar participants at Carnegie Mellon University, Washington University, UC Berkeley, the Data Colada Seminar, Utah's JDM Winter Symposium, the University of Wisconsin at Madison, and the JDM preconference of the Society for Personality and Social Psychology. Thanks to Amelia Dev, Christina Carr, Karin Garrett, Randall Hale, Josh Moore, Tory Taylor, and Josh Wu for valuable research assistance. Address correspondence to dm@berkeley.edu

Overprecision is a Property of Thinking Systems

The certainty with which we hold our beliefs—beliefs about political issues, about scientific hypotheses, or about our own abilities—exceeds their accuracy (Moore, Tenney, et al., 2015). This bias inflates the confidence judgments on which decisions depend. Excess certainty also enables other biases by making people too confident in their flawed intuitions. Researchers have written about this overprecision in judgment as if it resulted from a heuristic such like availability or representativeness (Brenner, Koehler, Liberman, et al., 1996; Hoffrage, 2004; Malmendier & Tate, 2005). I argue here that there is more to it than that. It results from a fundamental epistemological conundrum facing every finite thinking agent: It is impossible know all the ways to be wrong. By neglecting the ways in which they could be wrong, people wind up too sure that their beliefs are correct.

Forms of Overconfidence

Overconfidence may be the most consequential of the many biases to which human judgment is vulnerable (Bazerman & Moore, 2013; Kahneman, 2011). Scholars have implicated overconfidence in some of the greatest follies in history, including financial disasters, religious extremism, legal disputes, and wars (Babcock et al., 1995; Daniel et al., 2001; Johnson & Fowler, 2011; Ortoleva & Snowberg, 2015). The confidence people have in their political predictions, ideological positions, economic forecasts, or just about anything else consistently exceeds their actual accuracy (Moore, Tenney, et al., 2015; Soll et al., 2021). Excessive certainty can make people reluctant to consider the views of others, heed advice, or seek compromise to conflict (Minson & Mueller, 2012; Yaniv & Kleinberger, 2000). This

makes disagreements more intense, protracted, and costly (Johnson et al., 2006; Thompson & Loewenstein, 1992). It can also make people too willing to act on their poorly-informed views, even when this makes them vulnerable to exploitation by better-informed counterparties (Barber & Odean, 2000; Lyons et al., 2021).

Excessive certainty in the accuracy of one's knowledge, also known as overprecision, is the most ubiquitous form of overconfidence (Moore & Healy, 2008). The robustness and pervasiveness of overprecision stands in marked contrast to the inconsistencies in other forms of overconfidence: overestimation and overplacement. In other words, *underestimation* and *underplacement* are common but *underprecision* is almost nonexistent.

Overestimation describes exaggerations of performance, ability, or the likelihood of success. However, the well-documented "hard-easy" effect is a consequence of underestimation on easy tasks and overestimation on hard ones (Gigerenzer et al., 1991; Juslin, 1993). Although there are some circumstances in which people appear stubbornly optimistic (Simmons & Massey, 2012), pessimism is easy to find (Beck et al., 1974; Dolinski et al., 1987; Scheier et al., 2001). Sandbagging, in which people claim to be less than they are, is common (Gibson & Sachau, 2000). Defensive pessimism, in which people ruminate on the risk of failure, is widespread (Norem & Cantor, 1986).

Overplacement is a form of overconfidence in which people exaggerate their placement relative to others, also known as the "better-than-average" effect (Alicke & Govorun, 2005). On easy tasks or where people feel capable, they consistently claim to be better than others (Hoelzl & Rustichini, 2005; Williams & Gilovich,

2008). By contrast, underplacement emerges for hard tasks and rare events. People report they are worse than others on difficult tasks like juggling (Kruger, 1999). Just so, they believe they are less likely than others to experience rare events like winning the lottery or living past 100 (Chambers et al., 2003; Kruger & Burrus, 2004). Moore and Healy (2008) offer an explanation for the apparent inconsistency of underestimation on easy tasks and underplacement on hard tasks.

Overprecision is different from overestimation. It is possible to underestimate but still be overprecise, such as the person who is convinced that she failed an exam that she actually passed. Evidence suggests that, although people tend to overestimate on hard but underestimate on easy tasks, overprecision persists across all levels of difficulty (McKenzie et al., 2008). In fact, overprecision has proven itself remarkably robust across people, cultures, tasks, and judgments. However, the degree of overprecision varies dramatically with the way in which it is measured (Haran et al., 2010; Hoffrage, 2004).

Measuring Overprecision

The most common method researchers use to assess reported confidence is by comparing their confidence intervals with actual accuracy. For example, when business executives offered 80% confidence intervals around their market forecasts, those forecasts actually included the truth only 36% of the time (Ben-David et al., 2013). When epidemiologists attempted to forecast the course of the Covid-19 pandemic, the truth fell inside their 75% confidence intervals only 45% of the time (Recchia et al., 2021). And a survey of expert judgments, in domains as diverse as dispersion of radioactive materials to risks for construction cranes, found that 90%

confidence intervals had hit rates between 30% and 80%, with a median around 60% (Lin & Bier, 2008).

But there are reasons to question the degree to which ordinary people understand confidence intervals (Cesarini et al., 2006). Confidence intervals are an unfamiliar way to express uncertainty, and daily life rarely requires us to specify confidence intervals. Confidence intervals provided by research participants are shockingly unresponsive to important details like confidence level. Teigen and Jørgensen (2005) found that 75% and 99% confidence intervals have 50% and 57% hit rates, respectively. That is, instead of providing nuanced intervals responsive to the specified confidence level, people instead provide roughly plausible upper and lower bounds (Soll & Klayman, 2004).

Other confidence elicitation ask, in one form or another, “How sure are you that you are correct?” For instance, many studies ask participants to select the best answer from a set of options and then estimate the probability they selected the correct one (Dawes & Mulford, 1996; Fischhoff et al., 1977; Gigerenzer et al., 1991; Oskamp, 1965). These item-confidence judgments have the advantage that people seem to understand them better than confidence-intervals. Probability judgments also seem more intuitive than confidence intervals, percentiles, or odds (Seaver et al., 1978; Wallsten et al., 2016; Wallsten & Budescu, 1983). However, they are vulnerable to the critique that they stack the deck in favor of finding overprecision. Noise in people’s confidence should make accuracy regressive relative to confidence. By conditioning on high confidence, accuracy is likely to be lower (Erev

et al., 1994). Put another way, overconfidence is most likely to prevail where people are most confident.

More ambitious researchers have attempted to elicit full subjective probability distributions. Usually, this requires the researcher to break the range of possibilities into a set of mutually exclusive and exhaustive categories, or bins, and then ask for the probability in each of them. For example, Haran and colleagues (2010, Study 2) asked participants to estimate the high temperature in Washington, DC one month in the future. Participants who provided 90% confidence intervals only got the right answer inside 29% of the time. By contrast, responses appeared better calibrated when participants estimated the probabilities of each bin in a histogram. In that case, 90% confidence captured the truth 60% to 80% of the time.

Similarly, Goldstein and Rothschild (2014) report that people show impressively good calibration using a graphical “distribution builder” that asked for estimates of the likelihood of each bin in a histogram. The researchers tested the accuracy of participants’ reports with a variety of symmetrical, skewed, and uniform distributions. Histograms proved nearly 80% accurate inside the range that included 80% of the respondent’s confidence. The researchers also asked their participants for 80% confidence intervals using the conventional method of reporting the 11th and 90th percentiles; these estimates were only about 50% accurate. It is striking that histogram reports were so much better calibrated, given that they too represent an unfamiliar way to report uncertainty.

Other studies employ elicitation methods that are closer to everyday life. Life, after all, tests our confidence every day. For instance, because it is worse to hit the car in

front of you than to leave a little extra space, drivers will err on the side of stopping sooner. Mannes and Moore (2013) inferred confidence by observing how much people respond to asymmetric penalties for error. In each of several rounds, research participants estimated the high temperature in Pittsburgh, Pennsylvania on randomly selected days in the past. Some rounds rewarded responses that were close to or *higher* than the truth. Other rounds, by contrast, only rewarded responses close to or *lower* than the truth. These shifting incentives should have bigger effects on those less confident in their accuracy. The results revealed, however, that responses moved less than they should. This lack of movement implies overprecision: People excessively sure of their estimates move too little in response to asymmetric payoffs (Mamassian, 2008). The evidence does not suggest, then, that the familiarity of some elicitation method or similarity to everyday judgment is exclusively responsible for overprecision.

Most everyone reports higher confidence when they focus only on the one answer than when they rate the likelihood of all the possible answers (Moore, Carter, et al., 2015). Other elicitations that focus respondents on the likelihood of their favored response, such as confidence intervals, also show dramatic overprecision (Alpert & Raiffa, 1982). Again, small modifications of this procedure that get people to think about the possibility that they might be wrong tend to reduce overprecision (Soll & Klayman, 2004). And it must be noted that while variations in question format have large effects on the degree of overprecision, there is no question format that eliminates it, let alone reverses it. Here, I will argue

that what accounts for differences between elicitations is the degree to which each one forces respondents to consider why they might be wrong.

A New Explanation

I propose that overprecision arises from a ubiquitous limitation in our understanding of the world. This epistemological constraint afflicts any fallible agent of limited processing capacity: it is the impossibility of identifying all the ways in which one may be wrong. Some information will always be unknown to agents of finite cognitive capacity; many situations will also include unknown unknowns—that is, useful knowledge that the agent fails to realize they lack. The “error neglect” theory of overprecision presented here posits that when there are many ways to be wrong, finite agents will fail to consider them all and will therefore come away too sure of their beliefs.

The probability that an hypothesis is correct is proportional to the strength of the evidence for it, as a percentage of all the available evidence. In the language of Bayes’ law, the probability of a focal hypothesis (H_i) given the evidence (E) is equal to the base-rate of the hypothesis, $p(H_i)$, multiplied by the strength of the evidence given the hypothesis, $p(E|H_i)$, divided by that product, summed over all possible k hypotheses:

$$p(H_i|E) = \frac{p(H_i)p(E|H_i)}{\sum_1^k p(H_k)p(E|H_k)}$$

Savage (1954) noted the vexing fact is that the world’s messy complexity rarely fits neatly into the structure required to apply Bayes’ formula. Textbook examples and stylized decision problems imagine a small world with a well-defined set of

alternatives. But the large world in which we actually live challenges boundedly rational agents of limited cognitive capacity (Felin et al., 2017; Simon, 1955).

The error neglect account posits that overprecision arises from a failure to consider all k alternatives. This can happen for many reasons, including limited computational power, inattention, or mere laziness (Thomas et al., 2014; Walters et al., 2016). It also results from unknown unknowns—important considerations, possibilities, or hypotheses that have not even entered the agent’s awareness. If, instead of considering all k possible answers, the individual instead considers only some subset j ($j \subseteq k$), the consequence will be an underestimate of the Bayesian denominator and an overestimate of the probability that the chosen answer is the right one (for an axiomatic justification of this claim, see Backus et al., 2022).

Evidence consistent with this explanation traces itself back to the early work of Fischhoff, Slovic, and Lichtenstein (1978), who noted people’s insensitivity to the omission of likely hypotheses. They told participants in their study, including experienced auto mechanics, a story about a car that failed to start. The researchers invited them to assign likelihoods to the branches of a “fault tree,” enumerating possible causes. Those who got the full fault tree assigned, on average, a 26% chance to the battery as a likely culprit, given the facts in the story. However, participants who got a pruned fault tree lacking the “battery” category did not appear to recognize its omission. Instead of adding 26% probability to the residual “other causes” category, their responses inflated the probability assigned to those categories they *did* see. Fischhoff *et al.* offered an “availability hypothesis” to explain

their results, arguing that “people may not realize what is missing and, therefore, may overestimate” the likelihoods of the given options.

Juslin, Winman, and Hansson (2007) offer a structural justification for availability based on the limits of human working memory. They describe people as naïve intuitive statisticians who sample facts from memory to estimate the width of subjective confidence intervals (Peterson & Beach, 1967; Tong & Feiler, 2017). However, consideration of evidence is constrained by the limitations of working memory. The variance in small samples, Juslin *et al.* note, is smaller than the variance in the population of relevant facts. Because they are not aware of the degree to which their small sample fails to represent the population, their subjective confidence intervals are too narrow.

The error neglect account presented here builds on the foundations laid by this prior work. However, I argue that the causes of overprecision in a favored hypothesis go beyond biases of human attention. If all hypotheses are equally salient or available, then the availability explanation would not predict overprecision. But if cognitive constraints limit human abilities to consider available hypotheses, then overprecision should grow with the number of candidate hypotheses. The first study in this paper illustrates precisely this effect in human judgment.

Moreover, if overprecision were a uniquely human bias, availability bias should not affect rational Bayesian judgment or computer systems not subject to human motivations. On the other hand, my formalization of the error neglect explanation implies that even good Bayesians, if they cannot think of all the relevant hypotheses, will wind up being too sure of the hypotheses they do consider. While

narrow human attention and availability biases may exacerbate overprecision, I predict it should persist even in the absence of availability biases and even for artificially intelligent agents. Study 2 tests whether an artificially intelligent agent is also overprecise in its judgment.

Evidentiary Support

Since Bayes' rule provides a poor description of intuitive human judgment, support theory (Tversky & Koehler, 1994) seeks greater descriptively accuracy. Support theory is not fundamentally incompatible with Bayesian logic. It stipulates that, given a pair of mutually exclusive outcomes, A and B , the judged probability of A will depend on support for A divided by the summed total support for both A and B :

$$p(A, B) = \frac{s(A)}{s(A) + s(B)}$$

When A and B represent the union of a larger set of subsidiary outcomes, the theory predicts disaggregation will increase perceived support, and here is where it parts ways with Bayes. When these subsidiary outcomes are left implicit, their estimated probability will be below ("subadditive" with regard to) their explicit disjunction. If B is the implicit union of the mutually exclusive outcomes C and D , then $s(B) < s(C) + s(D)$. For example, people judge the probability of death due to all natural causes as lower than the probability of death due to heart disease, cancer, or other natural causes (Tversky & Koehler, 1994). Put another way, people easily neglect or underestimate alternative hypotheses that are not specified, just as they neglected consideration of the battery when Fishchhoff *et al.* (1978) left it off their fault tree.

In terms of support theory, error neglect merely notes that, when considering alternative hypotheses, it is easy, especially when the number of alternatives taxes cognitive capacity, to lump together non-focal hypotheses. Although faithful Bayesians would contend that they can always specify an exhaustive set of alternative hypotheses by just including an “everything else” residual category, support theory questions the fidelity of this computation. It is, after all, the key contention of support theory that lumping together all other possibilities outside the focal hypothesis H_i into a catch-all $\neg H_i$ residual will underestimate its likelihood thanks to the subadditive nature of such unions (Rottenstreich & Tversky, 1997).

Many experimental confidence elicitation fail to specify alternative hypotheses and so make it easy to neglect them. For instance, item-confidence elicitation ask people to estimate the probability a favored answer is correct. The error neglect explanation holds that by focusing attention on a favored hypothesis, these measures exacerbate overprecision: People will be excessively sure that their hypothesis is correct if they neglect all the ways in which it could be wrong. Error neglect also predicts that confidence elicitation that force people to consider more ways they could be wrong will produce less overprecision.

I present illustrative evidence from two studies: one with human participants and one study of artificial intelligence. Study 1 shows how proliferating hypotheses contributes to their neglect, and ultimately to overprecision. Study 2 documents artificially intelligent agents’ vulnerability to overprecision in judgment.

Preregistered research plans, along with materials, data, and analysis code are available online: <https://osf.io/fcrj7/> (Moore, 2022).

STUDY 1: NEGLECT OF ALTERNATIVES GROWS WITH THEIR NUMBER

Study 1 employs an experimental manipulation that varies k , the number of alternative hypotheses. The smaller the set of possible alternatives the easier it is for individuals to think through them. Experimental conditions featuring twenty possible answers are likely to prove challenging for participants whose working memory is limited to around seven hypotheses (Miller, 1956). On the other hand, when there are just two possibilities, the set of alternative hypotheses is as small as it can be; it should be feasible for most people to keep track of and remember both the focal and the alternative hypothesis.

Method

Participants. I preregistered a sample size of 200, recruited online through MTurk (via Cloud Research; see Litman et al., 2017).

Procedure. Participants underwent a brief training procedure designed to verify their understanding of the elicitation procedures.

Participants then saw five dot images, selected at random (with replacement) from a set of 40. Each image appeared for ten seconds. The dot counts on the images came from a uniform distribution between 1 and 1000. The subsequent screen asked, “Please indicate how confident you are that the true number of dots in the image you just saw lies within each of the following ranges:” followed by 2, 4, 6, 8, 10, 12, 14, 16, 18 or 20 ranges of equal size, i.e. ‘bins,’ covering the 1 to 1000 range of possible answers. In the 2-bin condition participants estimated the probability

that there were 0-500 dots or 501-1000 dots. In the 20-bin condition, participants estimated the probability there were 0-50 dots, 51-100 dots...951-1000 dots. Participants indicated their subjective probability, on a scale of 0-100, that the correct number of dots fell within each of the bins listed. After submitting their responses, participants proceeded to the next image and this process repeated until the participant completed five images. Binning varied within-participant, such that each individual saw five different binning arrangements, selected at random (without replacement) from the ten possibilities.

In addition to the base payment, participants were eligible to win one of four \$20 prizes. The better calibrated a participant's estimates, the better their chances of winning a prize. I scored calibration using the quadratic scoring rule (Selten, 1998). The survey did not present the rule itself, but informed participants that the likelihood of winning a prize increased with accuracy and that they should report their true beliefs.

Measures. The key dependent measure is the percentage likelihood assigned to the "most likely" bin. The response scale was a slider on a 0-100% scale. Note that responses were forced to sum to 100%.

My preregistered measure of overprecision is the difference between peak confidence (the bin assigned highest probability) and accuracy (the frequency with which the truth falls inside that bin). Tied peaks received partial credit. In other words, if two bins each got 50% probability, an answer in either one earned half a hit. Overprecision is peak confidence minus the mean hit rate (across all participants) for that condition (that binning arrangement).

Results and Discussion

Planned analyses. I conducted a linear regression at the level of the round (5 observations per participant) with fixed effects for participant, in which the number of bins (B) served as the independent variable. In order to accommodate the possibility that the number of bins bears a curvilinear relationship with overprecision, I also preregistered $1 - (\frac{1}{B})$ as an independent variable. The error neglect explanation predicts that overprecision would grow with the number of bins, but at a declining rate, producing a statistically significant $1 - (\frac{1}{B})$ term. The results do indeed confirm that the asymptotic term $1 - (\frac{1}{B})$ is a predictor of overprecision, $t = 10.02$, $p < 10^{-16}$, as shown in Figure 1.

Figure 1. Violin plot showing confidence assigned to the focal bin, by number of bins, Study 1. Error bars show standard errors. Red Xs indicate actual hit rates inside the focal bin.

Additional analyses. The results are consistent with error neglect: Judgments are overprecise. That is, confidence ($M = 67.7\%$) exceeds accuracy ($M = 32.9\%$), $t(199) = 32.1, p < 10^{-16}, d = 2.3$. More importantly, the significance of the term $1 - (\frac{1}{B})$ reflects the fact that confidence in a focal hypothesis declines with the addition of alternative hypotheses, but at a declining rate. Recall the support theory formulation of subjective probability:

$$p(A, B) = \frac{s(A)}{s(A) + s(B)}$$

If, by adding bins, Study 1 was disaggregating alternative hypotheses, doing so would increase the denominator by breaking $s(B)$ apart.

However, it is notable that accuracy was so responsive to the binning manipulation. Since the range of possible answers remained constant while the binning changed (from 2 to 20 bins), the 2-bin condition made the task easier. Since the binning variation affected accuracy more than it affected confidence, results reveal the most overprecision with 20 bins but the least with two bins. Confidence exceeded accuracy by 39% in the 20-bin condition, $t(99) = 17.7, p < 10^{-16}, d = 2.56$. By contrast, confidence exceeded accuracy by only 8.2% in the 2-bin condition, $t(100) = 5.8, p < 10^{-6}, d = .53$.

The results do not suggest that information leakage, demand effects, or the ignorance prior drove the results. If they had, the simple provision of more bins would have led participants to assign probabilities to them, driving down confidence in the focal bin. In fact, we observe that overprecision grows as the number of alternative hypotheses proliferate. Because participants had useful

knowledge and focused on a most likely bin, as the number of bins went up and the bins got smaller, the task got harder and accuracy fell. However, participants' confidence did not fall nearly as much, exacerbating overprecision.

STUDY 2: WATSON

When there are an infinite number of ways to be wrong, even extremely intelligent agents cannot consider them all. Artificially intelligent (but finite) agents are also susceptible to overconfidence (Guo et al., 2017; Oelrich et al., 2020; Tesauro et al., 2013), especially when assumptions underlying their models are violated (Attenberg et al., 2011). AI systems' vulnerability to overprecision is a well-known problem in the study of deep learning (Abdar et al., 2021; Kivlichan et al., 2021). The failure to account for unknown unknowns represents an important reason why they are overconfident (Maurer & Bennette, 2018). Any set of training data, however large, will be smaller than the world. This creates the possibility that some new instance is different enough from the training data that the lessons from training apply imperfectly.

In principle, with a sufficient training set, an agent could implement a brute force correction to the problem of overconfidence (Angelopoulos & Bates, 2021). That is, it could take the "outside view" (Kahneman & Lovallo, 1993) and compute the difference between confidence and accuracy across the entire sample of prior instances. In practice, this solution depends on the set of training questions being representative of future questions. Insofar as a specific novel question is distinctive, the universal correction factor may not apply.

When programmers at IBM built the Watson system to play the game show Jeopardy, Watson needed confidence estimates to determine whether to buzz in on a round (Ferrucci, 2012). In the 2011 exhibition match in which Watson bested Jeopardy champions Ken Jennings and Brad Rutter, Watson displayed its top three guesses for each round. Across all 354 rounds, Watson's confidence averaged 32%. Watson's confidence exceeded accuracy; these guesses were only correct 29% of the time, a marginally significant difference, one-sample $t(353) = 1.84, p = .07, d = .097$.

Seeking more data, I reached out to IBM and received a (proprietary) file listing 332,111 guesses to 3477 Jeopardy clues. These clues represent part of Watson's training set. They came from 66 randomly-chosen episodes of Jeopardy between 2003 and 2010. The set provides an average of 95.5 candidate guesses per clue, at most one of which is correct. So most of these guesses are wrong, and are associated with low confidence. Watson's hit rate on these guesses is 1.03%. Watson's confidence, averaged across all its guesses, 1.20%, is significantly above this hit rate according to one-sample t -test, $t(332110) = 12.2, p < 10^{-16}, d = .021$.

Watson's overprecision emerges despite the fact that Watson was programmed to adjust its confidence after considering the rate at which it had gotten prior answers correct when the evidence had been similarly strong (Murdock & Tesauro, 2012). Yet because this training set was of finite size, it could not capture all the idiosyncrasies of every possible question, and Watson remained unaware of all the ways in which it could be wrong.

GENERAL DISCUSSION

The philosophical tradition of fallibilism reminds us that every belief is provisional and that we cannot know anything with certainty (Peirce, 1955). Indeed, Popper (1959) cautioned scientists against excessive confidence in their own theories. While evidence can sometimes allow us to rule out some hypotheses, it cannot prove any theory permanently and irrefutably true. It is always possible that another theory, as yet unimagined, could better explain the data.

Contravening this wisdom, overprecision is ubiquitous in judgments by both people and artificially intelligent agents. The error neglect explanation presented here can explain three important features of the empirical record. First, the theory provides an explanation for the ubiquity of overprecision: Overprecision is a necessary consequence of imperfect knowledge and finite cognitive capacity. Second, it comports better with the evidence than do other popular explanations (see the “Prior Explanations” section below). Third, it can help explain why small differences in elicitation can produce large differences in confidence.

The highest levels of overprecision are evident in unstructured open-ended elicitations for which the question itself does not force respondents to reflect on other possible answers or on why they might be wrong (Wallsten et al., 2016). Confidence intervals and best-guess elicitations produce the most overprecise judgments (Soll & Klayman, 2004). Other formats, such as full histograms, help people improve calibration between judgment and reality by forcing them to think about ways in which their favored answer might be incorrect and explicitly estimate the probability of alternatives (Goldstein & Rothschild, 2014; Haran et al., 2010).

Constraints on human cognitive capacity will exacerbate overprecision by limiting the set of alternative hypotheses people spontaneously generate (Thomas et al., 2014). The fewer hypotheses generated, the higher the probability assigned to each one. Other features of human psychology can also contribute to an overestimate of the strength of evidence for a focal hypothesis. People naturally tend to search for hypothesis-consistent over *-inconsistent* information (Klayman & Ha, 1987). When considering a favored answer or hypothesis, the evidence for it is more memorable than the evidence against it (Nickerson, 1998). In addition, people routinely neglect alternative hypotheses when they are unspecified or lumped together. Fox and Tversky (1998, Study 1) approached basketball fans during the NBA championships when there were only eight teams left in contention. The researchers asked fans to consider each of the teams, one at a time, and estimate the probability of that team winning the championships. The total probability summed to 240%, on average. Respondents were overconfident about each of eight hypotheses they considered because each judgment focused on one team at a time and neglected the other teams.

Varieties of Neglect

What exactly does it mean to neglect alternative hypotheses? I see at least three possibilities. First, the set of alternative hypotheses may be easily defined but taxing to think about (Sims, 2003). Participants in Fox and Tversky's (1998) study knew that there were eight teams but simplified the task by focusing on one at a time. That describes the challenge facing participants in Study 1. They knew that there were between 1 and 1000 dots in each image they saw, but did they bother to

think through each of the 1000 possibilities? Probably not. It is worth noting the parallel with Fischhoff *et al.*'s famous (1978) "fault trees" study. Most of the study's participants (especially the automobile mechanics) knew that the battery is often to blame when a car fails to start, but few spontaneously noted its absence when they received the incomplete fault tree. It is rare for people to ask themselves what is missing. It is easier to take evidence at face value than to ask how it might be incomplete or biased (Brenner, Koehler, & Tversky, 1996).

Second, the space of possible hypotheses may be difficult to define. "Who will win the next Presidential election, and how likely is that person to win?" Most people can formulate an opinion on the subject, but few spontaneously consider the thousands or millions of unlikely candidates for office. Instead, they treat those probabilities as if they were zero and come away too confident in their prediction.

Third, some viable alternatives may never even enter the consideration set. Strong *a priori* assumptions may make it difficult for people to even imagine viewing the same set of facts differently, representing an impediment to accurately calibrate their confidence. There is an analogy worth noting here to the challenges of scientific discovery. Scientists bring extant theories to their attempts to understand empirical data. But since they cannot consider all possible theories, perspectives, or causal models, they cannot define the entire set of possible hypotheses. I return to this issue in the paper's conclusion.

What You See Is All There Is

Numerous psychological biases arise from the failure to consider alternative hypotheses. Kahneman (2011) called it WYSIATI—what you see is all there is—to

describe the simplistic acceptance of the facts before you at face value. He describes the persuasive power of biased evidence, even when it is obviously and transparently cherry-picked: “The confidence that individuals have in their beliefs depends mostly on the quality of the story they can tell about what they see, even if they see little. We often fail to allow for the possibility that evidence that should be critical to our judgment is missing—what we see is all there is.” We are too accepting of one-sided evidence, often failing to spontaneously consider other interpretations or perspectives (Brenner, Koehler, & Tversky, 1996).

Ross and Ward (1996) called it naïve realism: the simplistic assumption that one’s own view represents the only sensible interpretation of the facts. Its wide-ranging consequences include people’s tendency to egocentrically project their own beliefs, thoughts, and sentiments onto others (Krueger & Clement, 1994). This “false consensus” effect makes it easy for us to believe that others see the world much as we do (Ross et al., 1977). Failures of perspective-taking mean that people can neglect others’ points of view (Keysar et al., 1995) and have trouble coming to agreement (Thompson & Loewenstein, 1992). This egocentrism produces the curse of knowledge, in which we fail to understand the perspectives of others, even when we know they hold different information than we do (Camerer et al., 1989).

Prior Explanations

Four prior attempts at explaining overprecision have received enough attention to merit mention here: anchoring, conversational norms, motivation, and individual differences. None of these four explanations can account for the variety of evidence in the literature on overprecision. None of them provide compelling

explanations for the massive effects of question formats. And none of them can account for the evidence presented here.

The anchoring explanation hypothesizes that people anchor on a single best guess and then adjust too little from it when considering alternative answers. Tversky and Kahneman (1974) offered anchoring and insufficient adjustment as an explanation for overly narrow confidence intervals. They drew the analogy to a judgment process in which people first consider some anchor value and then adjust too little from it when forming a final judgment. In their classic 1974 paper, they present the results of their now-famous anchoring study in which participants attempted to estimate the percentage of African countries in the United Nations. Before they answered, they had to spin a wheel of fortune that came up either 10 or 65. Then participants answered whether they thought the right answer as higher or lower than that number. Those who got the low anchor estimated 25%, on average. The mean among those who got the high anchor was 45%.

If anchoring on a best guess is to blame for overly narrow confidence intervals, then, then providing that best guess first should produce narrower confidence intervals than specifying the intervals' endpoints first. In fact, studies that test this hypothesis find that asking people for a best guess actually widens subsequent confidence intervals (Block & Harper, 1991; Selvidge, 1980; Soll & Klayman, 2004). This evidence directly contradicts the anchoring explanation.

The second explanation for overprecision hypothesizes that people are motivated to claim excessive certainty because it feels good to be right. In other words, overprecision might be its own reward (Burton, 2008). This is plausible,

given sanctimonious satisfaction that accompanies saying “I told you so” (Sezer et al., 2020). Some theories have posited that feeling confident may be an end unto itself, delivering what might be called belief utility—pleasure that comes from holding a belief, independent of its accuracy (Benabou & Tirole, 2002). While the idea is plausible, there is a distinct shortage of evidence that expressing confidence by, say, reporting narrow confidence intervals actually delivers any meaningful feelings satisfaction.

Moreover, the claim that confidence feels good must contend with the disutility of being confident and mistaken; the more sure you claim to be, the more embarrassing it is when the facts prove you wrong (Tenney et al., 2019). It is hard to argue that false certainty is self-serving if it leads to costly and embarrassing errors (McGraw et al., 2004). Motivations toward self-aggrandizement turn out to be rather self-defeating if they lead to erroneous public stances, or to costly losing bets. The motivation to avoid such errors ought to operate in direct opposition to any motivation to claim certainty.

Some have suggested that the value of being right increases with the importance of the topic, and there is indeed some correlational evidence suggesting that people are more overconfident in more important domains (Brown, 2011). However, correlational evidence cannot rule out other causal paths. In this case, it is plausible that people work to get better at tasks they regard as important, contributing to their confidence. Notably, however, studies documenting this relationship usually measure placement rather than precision. Manipulations of

motivation or importance have strikingly little effect on precision in judgment (Logg et al., 2018).

The third explanation posits that people express excessive precision in judgment because they think others expect it of them. This conversational norms explanation notes that people prefer to receive precise communications at the expense of accuracy. For example, when estimating the distance from Chicago to New York, people prefer the advisor who says it is 730-780 miles than the one who says it is 700-1500 miles, even knowing the correct answer is 713 miles (Yaniv & Foster, 1995). If audience preferences drive overprecise claims, then the audience should matter; that is, people should express greater precision when they are communicating with an audience who expects it than when they are communicating privately with an experimenter who promises to pay them for accuracy. Yet there is a conspicuous lack of evidence suggesting that overprecision depends on the audience or their interests (Moore et al., 2010).

The best evidence for the conversational norms explanation comes from studies providing social rewards to confident credibility. Radzevick and Moore (2011) put participants in the role of advisors competing with one another to attract customers to buy their advice. Over time, advisors saw that customers were more likely to buy advice from confident advisors and, in turn, expressed more confidence. So it may be that social dynamics can encourage confident claims (Anderson et al., 2012; Kennedy et al., 2013; Schwardmann & van der Weele, 2019; Van Zant, 2021). However, this represents a poor explanation for overprecision

expressed privately, even in the presence of incentives to punish poor calibration (Harvey, 1997; Hoffrage, 2004).

The fourth explanation is individual differences. This explanation posits that there are just some kinds of people who are more overprecise than others (Cesarini et al., 2009). For instance, it seems plausible that narcissistic individuals might be more overprecise (Campbell et al., 2004), or that some cultures might encourage the greater expression of certainty (Moore et al., 2018; Muthukrishna et al., 2018; Yates et al., 1997). There is some evidence of male overplacement (Niederle & Vesterlund, 2007) and gender stereotypes might lead us to expect that men behave as if they are more certain of themselves (Kay & Shipman, 2014). However, empirical evidence does not support the stereotype of male overprecision. Indeed, gender differences in all forms of overconfidence have proven difficult to replicate (Moore & Dev, 2018). If men are, in fact, sometimes more overconfident, this effect is contextually dependent and does not hold generally (Bordalo et al., 2019). Similar limitations seem to apply to other individual difference correlates of overconfidence (Moore & Swift, 2010).

Having focused on the limitations of these four explanations, I must concede that the evidence I present in this paper does little to undermine or refute them. Indeed, it is entirely possible that overprecision is multiply determined and that these other causes do operate in some conditions, even though none of them offer a general explanation that accounts for all the evidence. Next, I consider evidence that is most problematic for the explanation presented in this paper.

Contrary Evidence

I note two findings that appear problematic for the error neglect explanation for overprecision: underprecision and dud-alternatives. First, some studies have found evidence of underprecision (Moore et al., 2018; Prims & Moore, 2017), which error neglect does not predict. Two features of the underprecision evidence merit mention. First, it comes from histogram elicitations for which reported probability distributions are less concentrated than the true distribution. Less overprecision on histogram elicitations is, of course, quite compatible with error neglect. Second, it usually features aleatory uncertainty about random events in the world, which, as a rule, is associated with less confidence and less precision than is epistemic uncertainty (Tannenbaum et al., 2017). One particular error people make thinking about aleatory uncertainty is underconfidence regarding large samples. The probability of obtaining between 451 and 549 heads in 1000 coin flips is over 99%. In fact, people think the chance is below 20%, whereas they estimate the probability of getting less than 50 heads in 1000 flips is greater than 5% when in fact is nearly impossible (Benjamin et al., 2016).

A second finding that appears problematic for error neglect is the dud-alternatives effect: the inclusion of an alternative hypothesis can actually enhance the plausibility of a focal hypothesis (Windschitl & Chambers, 2004). This happens when the new alternative is preposterous—that is, a dud. This is problematic because it appears to be an instance in which adding more hypotheses to the Bayesian denominator *increases* the judged probability. Obviously, this result challenges not only the error neglect account but any Bayesian or quasi-Bayesian

theory of subjective probability assessment. The way to accommodate this result is to concede that additional information can inform not only neglected hypotheses but also focal hypotheses. If consideration of a dud provides false assurance at having ruled out an alternative hypothesis, it can enhance confidence in a focal hypothesis.

Scientific Uncertainty

The problem of unknown unknowns is a profound epistemological problem. It leads to what has been called theory uncertainty in physics (Levine et al., 2002) or model uncertainty in statistics (Draper, 1995) and in economics (Hansen & Sargent, 2001). In short, researchers can only test a subset of all possible models or theories. Consequently, we will always be vulnerable to the possibility that the true model will be among those left untested. In other words, we will be wrong for reasons we fail to anticipate. Some analyses provide helpful goodness-of-fit statistics that can help distinguish between rival models. However, our statistics rarely tell us which model we ought to be testing or which variables we have failed to consider. Some imagine that machine learning can help solve this problem, but machines also omit unmeasured variables and unimagined models.

Scientific estimates come with confidence intervals that acknowledge their imprecision. Published estimates of physical constants, from the speed of light to Avogadro's number, include confidence intervals. These confidence intervals include the true value less often than they ought to (Croft & Dailey, 2015; Henrion & Fischhoff, 1986). Confidence intervals around published effects look overconfident in part because the researchers cannot foresee all the boundary conditions and

reasons why their studies will fail to replicate (Camerer et al., 2016; Cameron et al., 2020; Open Science Collaboration, 2015). Statistics compute confidence intervals assuming that measurement error is the only source of error, thereby neglecting other sources of error that are more difficult to anticipate or quantify. As scientists, we are always vulnerable to being too sure our theories are accurate since we cannot know how the next scientific paradigm shift will upend our assumptions.

Scientific theories are based on existing data but designed to generalize beyond them. Good theories make novel predictions. However, scientists are always at risk of being overconfident in a favored hypothesis. This is the same as overfitting the data, relying on the observed evidence and neglecting counterexamples or contrary evidence, perhaps because they lie outside one's experience. Every theory is at risk of falsification by facts and evidence unknown to the scientist formulating the theory, including the theory presented here. Lacking omniscience, we are all vulnerable to neglecting ways in which we are wrong and therefore being overly certain in the accuracy of our beliefs.

References

- Abdar, M., Pourpanah, F., Hussain, S., Rezazadegan, D., Liu, L., Ghavamzadeh, M., Fieguth, P., Cao, X., Khosravi, A., Acharya, U. R., Makarenkov, V., & Nahavandi, S. (2021). A Review of Uncertainty Quantification in Deep Learning: Techniques, Applications and Challenges. *Information Fusion*, 76, 243–297. <https://doi.org/10.1016/j.inffus.2021.05.008>
- Alicke, M. D., & Govorun, O. (2005). The better-than-average effect. In M. D. Alicke, D. Dunning, & J. Krueger (Eds.), *The self in social judgment* (pp. 85–106). Psychology Press.
- Alpert, M., & Raiffa, H. (1982). A progress report on the training of probability assessors. In D. Kahneman, P. Slovic, & A. Tversky (Eds.), *Judgment under uncertainty: Heuristics and biases*. Cambridge University Press.
- Anderson, C., Brion, S., Moore, D. A., & Kennedy, J. A. (2012). A status-enhancement account of overconfidence. *Journal of Personality and Social Psychology*, 103(4), 718–735. <https://doi.org/10.1037/a0029395>
- Angelopoulos, A. N., & Bates, S. (2021). A Gentle Introduction to Conformal Prediction and Distribution-Free Uncertainty Quantification. *ArXiv:2107.07511 [Cs, Math, Stat]*. <http://arxiv.org/abs/2107.07511>
- Attenberg, J., Ipeirotis, P. G., & Provost, F. J. (2011). Beat the Machine: Challenging Workers to Find the Unknown Unknowns. *Human Computation: Papers from the 2011 AAAI Workshop*.

- Babcock, L., Loewenstein, G., Issacharoff, S., & Camerer, C. F. (1995). Biased judgments of fairness in bargaining. *American Economic Review*, 85(5), 1337–1343.
- Backus, M., Little, A. T., Moore, D. A., & Augenblick, E. (2022). Constraints on thinking cause overprecision. *Unpublished Manuscript*. <https://psyarxiv.com/evcx2/>
- Barber, B. M., & Odean, T. (2000). Trading is hazardous to your wealth: The common stock investment performance of individual investors. *Journal of Finance*, 55(2), 773–806.
- Bazerman, M. H., & Moore, D. A. (2013). *Judgment in managerial decision making* (8th ed.). Wiley.
- Beck, A. T., Weissman, A., Lester, D., & Trexler, L. (1974). The measurement of pessimism: The hopelessness scale. *Journal of Consulting and Clinical Psychology*, 42(6), 861.
- Benabou, R., & Tirole, J. (2002). Self-confidence and personal motivation. *Quarterly Journal of Economics*, 117(3), 871–915.
- Ben-David, I., Graham, J. R., & Harvey, C. R. (2013). Managerial miscalibration. *Quarterly Journal of Economics*, 128(4), 1547–1584.
- Benjamin, D., Moore, D. A., & Rabin, M. (2016). Misconceptions of chance: Evidence from an integrated experiment. *Unpublished Manuscript*.
- Block, R. A., & Harper, D. R. (1991). Overconfidence in estimation: Testing the anchoring-and-adjustment hypothesis. *Organizational Behavior and Human Decision Processes*, 49(2), 188–207.

- Bordalo, P., Coffman, K., Gennaioli, N., & Shleifer, A. (2019). Beliefs about Gender. *American Economic Review*, 109(3), 739–773.
<https://doi.org/10.1257/aer.20170007>
- Brenner, L., Koehler, D. J., Liberman, V., & Tversky, A. (1996). Overconfidence in probability and frequency judgments: A critical examination. *Organizational Behavior and Human Decision Processes*, 65(3), 212–219.
- Brenner, L., Koehler, D. J., & Tversky, A. (1996). On the evaluation of one-sided evidence. *Journal of Behavioral Decision Making*, 9(1), 59–70.
- Brown, J. D. (2011). Understanding the Better Than Average Effect: Motives (Still) Matter. *Personality & Social Psychology Bulletin*, 38(2), 209–219.
<https://doi.org/10.1177/0146167211432763>
- Burton, R. A. (2008). *On being certain: Believing you are right even when you're not*. St. Martin's press.
- Camerer, C. F., Dreber, A., Forsell, E., Ho, T.-H., Huber, J., Johannesson, M., Kirchler, M., Almenberg, J., Altmejd, A., & Chan, T. (2016). Evaluating replicability of laboratory experiments in economics. *Science*, 351(6280), 1433–1436.
- Camerer, C. F., Loewenstein, G., & Weber, M. (1989). The curse of knowledge in economic settings: An experimental analysis. *Journal of Political Economy*, 97(5), 1232–1254.
- Cameron, E., Angus, G. W., & Burgess, J. M. (2020). Overconfidence in Bayesian analyses of galaxy rotation curves. *Nature Astronomy*, 4(2), 132–133.
<https://doi.org/10.1038/s41550-019-0998-2>

- Campbell, W. K., Goodie, A. S., & Foster, J. D. (2004). Narcissism, confidence, and risk attitude. *Journal of Behavioral Decision Making*, 17(4), 481–502.
- Cesarini, D., Lichtenstein, P., Johannesson, M., & Wallace, B. (2009). Heritability of overconfidence. *Journal of the European Economic Association*, 7(2-3), 617–627.
- Cesarini, D., Sandewall, Ö., & Johannesson, M. (2006). Confidence interval estimation tasks and the economics of overconfidence. *Journal of Economic Behavior and Organization*, 61(3), 453–470.
- Chambers, J. R., Windschitl, P. D., & Suls, J. (2003). Egocentrism, event frequency, and comparative optimism: When what happens frequently is “more likely to happen to me.” *Personality and Social Psychology Bulletin*, 29(11), 1343–1356.
- Croft, R. A. C., & Dailey, M. (2015). On the measurement of cosmological parameters. *ArXiv:1112.3108 [Astro-Ph, Physics:Gr-Qc]*. <http://arxiv.org/abs/1112.3108>
- Daniel, K. D., Hirshleifer, D. A., & Subrahmanyam, A. (2001). Overconfidence, arbitrage, and equilibrium asset pricing. *Journal of Finance*, 56(3), 921–965.
- Dawes, R. M., & Mulford, M. (1996). The false consensus effect and overconfidence: Flaws in judgment or flaws in how we study judgment? *Organizational Behavior and Human Decision Processes*, 65(3), 201–211.
- Dolinski, D., Gromski, W., & Zawisza, E. (1987). Unrealistic pessimism. *Journal of Social Psychology*, 127(5), 511–516.
- Draper, D. (1995). Assessment and propagation of model uncertainty. *Journal of the Royal Statistical Society. Series B (Methodological)*, 57(1), 45–97.

- Erev, I., Wallsten, T. S., & Budescu, D. V. (1994). Simultaneous over- and underconfidence: The role of error in judgment processes. *Psychological Review*, 101(3), 519–527.
- Felin, T., Koenderink, J., & Krueger, J. I. (2017). Rationality, perception, and the all-seeing eye. *Psychonomic Bulletin & Review*, 24(4), 1040–1059.
- Ferrucci, D. A. (2012). Introduction to “this is watson.” *IBM Journal of Research and Development*, 56(3.4), 1.
- Fischhoff, B., Slovic, P., & Lichtenstein, S. (1977). Knowing with certainty: The appropriateness of extreme confidence. *Journal of Experimental Psychology: Human Perception and Performance*, 3(4), 552–564.
- Fischhoff, B., Slovic, P., & Lichtenstein, S. (1978). Fault trees: Sensitivity of estimated failure probabilities to problem representation. *Journal of Experimental Psychology: Human Perception and Performance*, 4(2), 330–344.
- Fox, C. R., & Tversky, A. (1998). A belief-based account of decision under uncertainty. *Management Science*, 44(7), 879–895.
- Gibson, B., & Sachau, D. (2000). Sandbagging as a self-presentational strategy: Claiming to be less than you are. *Personality and Social Psychology Bulletin*, 26(1), 56–70.
- Gigerenzer, G., Hoffrage, U., & Kleinbölting, H. (1991). Probabilistic mental models: A Brunswikian theory of confidence. *Psychological Review*, 98(4), 506–528.
- Goldstein, D. G., & Rothschild, D. (2014). Lay understanding of probability distributions. *Judgment and Decision Making*, 9(1), 1–14.

- Guo, C., Pleiss, G., Sun, Y., & Weinberger, K. Q. (2017). On Calibration of Modern Neural Networks. *ArXiv:1706.04599 [Cs]*. <http://arxiv.org/abs/1706.04599>
- Hansen, L., & Sargent, T. J. (2001). Robust control and model uncertainty. *American Economic Review*, 91(2), 60–66.
- Haran, U., Moore, D. A., & Morewedge, C. K. (2010). A simple remedy for overprecision in judgment. *Judgment and Decision Making*, 5(7), 467–476.
- Harvey, N. (1997). Confidence in judgment. *Trends in Cognitive Sciences*, 1(2), 78–82.
- Henrion, M., & Fischhoff, B. (1986). Assessing uncertainty in physical constants. *American Journal of Physics*, 54(9), 791–798.
- Hoelzl, E., & Rustichini, A. (2005). Overconfident: Do you put your money on it? *Economic Journal*, 115(503), 305–318.
- Hoffrage, U. (2004). Overconfidence. In R. F. Pohl (Ed.), *Cognitive illusions: Fallacies and biases in thinking, judgment, and memory* (pp. 235–254). Psychology Press.
- Johnson, D. D. P., & Fowler, J. H. (2011). The evolution of overconfidence. *Nature*, 477(7364), 317–320.
- Johnson, D. D. P., McDermott, R., Barrett, E. S., Cowden, J., Wrangham, R., McIntyre, M. H., & Rosen, S. P. (2006). Overconfidence in wargames: Experimental evidence on expectations, aggression, gender and testosterone. *Proceedings of the Royal Society B*, 273, 2513–2520.
- Juslin, P. (1993). An explanation of the hard/easy effect in studies of realism of confidence of one's general knowledge. *European Journal of Cognitive Psychology*, 5(1), 55–71.

- Juslin, P., Winman, A., & Hansson, P. (2007). The naïve intuitive statistician: A naïve sampling model of intuitive confidence intervals. *Psychological Review*, 114(3), 678–703.
- Kahneman, D. (2011). *Thinking fast and slow*. Farrar, Straus and Giroux.
- Kahneman, D., & Lovallo, D. (1993). Timid choices and bold forecasts: A cognitive perspective on risk and risk taking. *Management Science*, 39, 17–31.
- Kay, K., & Shipman, C. (2014). *The Confidence Code: The Science and Art of Self-Assurance—What Women Should Know*. HarperCollins.
<https://books.google.com/books?id=U7hXAgAAQBAJ>
- Kennedy, J. A., Anderson, C., & Moore, D. A. (2013). When overconfidence is revealed to others: Testing the status-enhancement theory of overconfidence. *Organizational Behavior and Human Decision Processes*, 122(2), 266–279.
- Keysar, B., Ginzler, L. E., & Bazerman, M. H. (1995). States of affairs and states of mind: The effect of knowledge of beliefs. *Organizational Behavior and Human Decision Processes*, 64(3), 283–293.
- Kivlichan, I., Lin, Z., Liu, J., & Vasserman, L. (2021). Measuring and Improving Model-Moderator Collaboration using Uncertainty Estimation. *Proceedings of the 5th Workshop on Online Abuse and Harms (WOAH 2021)*, 36–53.
<https://doi.org/10.18653/v1/2021.woah-1.5>
- Klayman, J., & Ha, Y. (1987). Confirmation, disconfirmation, and information in hypothesis testing. *Psychological Review*, 94(2), 211–228.
- Krueger, J. I., & Clement, R. W. (1994). The truly false consensus effect: An ineradicable and egocentric bias in social perception. *Journal of Personality*

- and Social Psychology*, 67(4), 596–610. <https://doi.org/10.1037/0022-3514.68.4.579>
- Kruger, J. (1999). Lake Wobegon be gone! The “below-average effect” and the egocentric nature of comparative ability judgments. *Journal of Personality and Social Psychology*, 77(2), 221–232.
- Kruger, J., & Burrus, J. (2004). Egocentrism and focalism in unrealistic optimism (and pessimism). *Journal of Experimental Social Psychology*, 40(3), 332–340. <https://doi.org/10.1016/j.jesp.2003.06.002>
- Levine, E. S., Schulz, A. E., & White, M. (2002). Future galaxy cluster surveys: The Effect of theory uncertainty on constraining cosmological parameters. *The Astrophysical Journal*, 577(2), 569.
- Lin, S.-W., & Bier, V. M. (2008). A study of expert overconfidence. *Reliability Engineering & System Safety*, 93(5), 711–721.
- Litman, L., Robinson, J., & Abberbock, T. (2017). TurkPrime.com: A versatile crowdsourcing data acquisition platform for the behavioral sciences. *Behavior Research Methods*, 49(2), 433–442.
- Logg, J. M., Haran, U., & Moore, D. A. (2018). Is overconfidence a motivated bias? Experimental evidence. *Journal of Experimental Psychology: General*.
- Lyons, B. A., Montgomery, J. M., Guess, A. M., Nyhan, B., & Reifler, J. (2021). Overconfidence in news judgments is associated with false news susceptibility. *Proceedings of the National Academy of Sciences*, 118(23). <https://doi.org/10.1073/pnas.2019527118>

- Malmendier, U., & Tate, G. (2005). CEO overconfidence and corporate investment. *Journal of Finance*, 60(6), 2661–2700.
- Mamassian, P. (2008). Overconfidence in an objective anticipatory motor task. *Psychological Science*, 19(6), 601–606.
- Mannes, A. E., & Moore, D. A. (2013). A behavioral demonstration of overconfidence in judgment. *Psychological Science*, 24(7), 1190–1197.
- Maurer, K., & Bennette, W. (2018). Facility Locations Utility for Uncovering Classifier Overconfidence. *ArXiv*.
- McGraw, A. P., Mellers, B. A., & Ritov, I. (2004). The affective costs of overconfidence. *Journal of Behavioral Decision Making*, 17(4), 281–295.
- McKenzie, C. R. M., Liersch, M. J., & Yaniv, I. (2008). Overconfidence in interval estimates: What does expertise buy you? *Organizational Behavior & Human Decision Processes*, 107, 179–191.
- Miller, G. A. (1956). The Magical Number Seven, Plus or Minus Two: Some Limits on Our Capacity for Processing Information. *Psychological Review*, 63, 81–97.
- Minson, J. A., & Mueller, J. (2012). Two is company and four is a crowd? Advice under-utilization by individuals and dyads. *Psychological Science*.
- Moore, D. A. (2022). *Overprecision is a Property of Thinking Systems*. OSF. <https://osf.io/fcrj7/>
- Moore, D. A., Carter, A., & Yang, H. H. J. (2015). Wide of the mark: Evidence on the underlying causes of overprecision in judgment. *Organizational Behavior and Human Decision Processes*, 131, 110–120.

- Moore, D. A., & Dev, A. S. (2018). Individual differences in overconfidence. In V. Zeigler-Hill & T. Shackelford (Eds.), *Encyclopedia of Personality and Individual Differences*. Springer. <http://osf.io/hzk6q>
- Moore, D. A., Dev, A. S., & Goncherova, E. Y. (2018). Overconfidence across cultures. *Collabra: Psychology*, 4(1), 36.
- Moore, D. A., Haran, U., & Radzevick, J. R. (2010). *Overprecision and Manipulating the Audience*. OSF. <https://osf.io/fts25/>
- Moore, D. A., & Healy, P. J. (2008). The trouble with overconfidence. *Psychological Review*, 115(2), 502–517.
- Moore, D. A., & Swift, S. A. (2010). The Three Faces of Overconfidence in Organizations. In R. Van Dick & J. K. Murnighan (Eds.), *Social Psychology of Organizations* (pp. 147–184). Taylor & Francis.
- Moore, D. A., Tenney, E. R., & Haran, U. (2015). Overprecision in judgment. In G. Wu & G. Keren (Eds.), *Handbook of Judgment and Decision Making* (pp. 182–212). Wiley.
- Murdock, J. W., & Tesauro, G. (2012). *Statistical Approaches to Question Answering in Watson*. Joint Poicy Board for Mathematics. http://www.mathaware.org/mam/2012/pdfs/StatQA_v3.pdf
- Muthukrishna, M., Heine, S. J., Toyakawa, W., Hamamura, T., Kameda, T., & Henrich, J. (2018). Overconfidence is universal? Depends what you mean. *PLoS One*, 13(8). <https://doi.org/10.1371/journal.pone.0202288>
- Nickerson, R. S. (1998). Confirmation bias: A ubiquitous phenomenon in many guises. *Review of General Psychology*, 2, 175–220.

- Niederle, M., & Vesterlund, L. (2007). Do women shy away from competition? Do men compete too much? *Quarterly Journal of Economics*, 122(3), 1067–1101.
- Norem, J. K., & Cantor, N. (1986). Defensive pessimism: Harnessing anxiety as motivation. *Journal of Personality and Social Psychology*, 51(6), 1208–1217.
- Oelrich, O., Ding, S., Magnusson, M., Vehtari, A., & Villani, M. (2020). When are Bayesian model probabilities overconfident? *ArXiv:2003.04026*.
<http://arxiv.org/abs/2003.04026>
- Open Science Collaboration. (2015). Estimating the reproducibility of psychological science. *Science*, 349(6251), aac4716.
- Ortoleva, P., & Snowberg, E. (2015). Overconfidence in political behavior. *American Economic Review*, 105(2), 504–535.
- Oskamp, S. (1965). Overconfidence in case-study judgments. *Journal of Consulting Psychology*, 29(3), 261–265.
- Peirce, C. S. (1955). The scientific attitude and fallibilism. In J. Buchler (Ed.), *Philosophical Writings of Peirce*. Dover.
- Peterson, C. R., & Beach, L. R. (1967). Man as an intuitive statistician. *Psychological Bulletin*, 68(1), 29–46.
- Popper, K. (1959). *The Logic of Scientific Discovery*. Routledge.
- Prims, J. P., & Moore, D. A. (2017). Overconfidence over the lifespan. *Judgment and Decision Making*, 12(1).
- Radzevick, J. R., & Moore, D. A. (2011). Competing to be certain (but wrong): Market dynamics and excessive confidence in judgment. *Management Science*, 57(1), 93–106. <https://doi.org/10.1287/mnsc.1100.1255>

- Recchia, G., Freeman, A. L. J., & Spiegelhalter, D. (2021). How well did experts and laypeople forecast the size of the COVID-19 pandemic? *PLOS ONE*, 16(5), e0250935. <https://doi.org/10.1371/journal.pone.0250935>
- Ross, L., Greene, D., & House, P. (1977). The false consensus effect: An egocentric bias in social perception and attribution processes. *Journal of Experimental Social Psychology*, 13(3), 279–301.
- Ross, L., & Ward, A. (1996). Naive realism in everyday life: Implications for social conflict and misunderstanding. In E. Reed, E. Turiel, & T. Brown (Eds.), *Values and knowledge* (pp. 103–135). Lawrence Erlbaum Associates.
- Rottenstreich, Y., & Tversky, A. (1997). Unpacking, repacking, and anchoring: Advances in support theory. *Psychological Review*, 104(2), 406–415.
- Savage, L. J. (1954). *The foundations of statistics*. Dover.
- Scheier, M. F., Carver, C. S., & Bridges, M. W. (2001). Optimism, pessimism, and psychological well-being. In E. C. Chang (Ed.), *Optimism and pessimism: Implications for theory, research, and practice* (pp. 189–216). American Psychological Association.
- Schwardmann, P., & van der Weele, J. (2019). Deception and self-deception. *Nature Human Behaviour*, 3, 1055–1061. <https://doi.org/10.1038/s41562-019-0666-7>
- Seaver, D. A., von Winterfeldt, D., & Edwards, W. (1978). Eliciting subjective probability distributions on continuous variables. *Organizational Behavior & Human Decision Processes*, 21, 84–104.

- Selten, R. (1998). Axiomatic characterization of the quadratic scoring rule. *Experimental Economics*, 1(1), 43–61.
- Selvidge, J. E. (1980). Assessing the extremes of probability distributions by the fractile method. *Decision Sciences*, 11(3), 493–502.
- Sezer, O., Affinito, S., & Staats, B. (2020). The I Told You So Effect. *ACR North American Advances*.
- Simmons, J. P., & Massey, C. (2012). Is Optimism Real? *Journal of Experimental Psychology: General*. <https://doi.org/10.1037/a0027405>
- Simon, H. A. (1955). A behavioral model of rational choice. *Quarterly Journal of Economics*, 69, 99–118.
- Sims, C. A. (2003). Implications of rational inattention. *Journal of Monetary Economics*, 50(3), 665–690. [https://doi.org/10.1016/S0304-3932\(03\)00029-1](https://doi.org/10.1016/S0304-3932(03)00029-1)
- Soll, J. B., & Klayman, J. (2004). Overconfidence in interval estimates. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 30(2), 299–314.
- Soll, J. B., Palley, A., Klayman, J., & Moore, D. A. (2021). Overconfidence in probability distributions: People know they don't know but they don't know what to do about it. *Unpublished Manuscript*.
- Tannenbaum, D., Fox, C. R., & Ülkümen, G. (2017). Judgment extremity and accuracy under epistemic versus aleatory uncertainty. *Management Science*, 63(2), 497–518.

- Teigen, K. H., & Jørgensen, M. (2005). When 90% confidence intervals are 50% certain: On the credibility of credible intervals. *Applied Cognitive Psychology*, 19(4), 455–475.
- Tenney, E. R., Meikle, N. L., Hunsaker, D., Anderson, C., & Moore, D. A. (2019). Is overconfidence a social liability? The effect of verbal versus nonverbal expressions of confidence. *Journal of Personality and Social Psychology*, 116(3), 396–415.
- Tesauro, G., Gondek, D. C., Lenchner, J., Fan, J., & Prager, J. M. (2013). Analysis of Watson's Strategies for Playing Jeopardy! *Journal of Artificial Intelligence Research*, 21, 205–251.
- Thomas, R., Dougherty, M. R., & Buttaccio, D. R. (2014). Memory constraints on hypothesis generation and decision making. *Current Directions in Psychological Science*, 23(4), 264–270.
- Thompson, L., & Loewenstein, G. (1992). Egocentric interpretations of fairness and interpersonal conflict. *Organizational Behavior and Human Decision Processes*, 51(2), 176–197.
- Tong, J., & Feiler, D. (2017). A behavioral model of forecasting: Naive statistics on mental samples. *Management Science*, 63(11), 3609–3627.
- Tversky, A., & Kahneman, D. (1974). Judgment under Uncertainty: Heuristics and Biases. *Science*, 185(4157), 1124–1131.
<https://doi.org/10.1126/science.185.4157.1124>

- Tversky, A., & Koehler, D. J. (1994). Support theory: A nonextensional representation of subjective probability. *Psychological Review*, 101(4), 547–567.
- Van Zant, A. B. (2021). Strategically overconfident (to a fault): How self-promotion motivates advisor confidence. *Journal of Applied Psychology*.
<https://doi.org/10.1037/apl0000879>
- Wallsten, T. S., & Budescu, D. V. (1983). State of the art—Encoding subjective probabilities: A psychological and psychometric review. *Management Science*, 29(2), 151–173.
- Wallsten, T. S., Shlomi, Y., Nataf, C., & Tomlinson, T. (2016). Efficiently Encoding and Modeling Subjective Probability Distributions for Quantitative Variables. *Decision*, 3(3), 169–189.
- Walters, D. J., Fernbach, P. M., Fox, C. R., & Sloman, S. A. (2016). Known unknowns: A critical determinant of confidence and calibration. *Management Science*, 63(12), 4298–4307.
- Williams, E. F., & Gilovich, T. (2008). Do people really believe they are above average? *Journal of Experimental Social Psychology*, 44(4), 1121–1128.
- Windschitl, P. D., & Chambers, J. R. (2004). The dud-alternative effect in likelihood judgment. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 30(1), 198.
- Yaniv, I., & Foster, D. P. (1995). Graininess of judgment under uncertainty: An accuracy-informativeness trade-off. *Journal of Experimental Psychology: General*, 124(4), 424–432.

- Yaniv, I., & Kleinberger, E. (2000). Advice taking in decision making: Egocentric discounting and reputation formation. *Organizational Behavior and Human Decision Processes*, 84(2), 260–281.
- Yates, J. F., Lee, J.-W., & Bush, J. G. (1997). General knowledge overconfidence: Cross-national variations, response style, and "reality ". *Organizational Behavior and Human Decision Processes*, 70(2), 87–94.