Constraints on Thinking Cause Overprecision*

Matthew Backus[†] Andrew T. Little[‡] Don A. Moore^{*§}

April 26, 2022

Abstract

Overprecision—the excessive faith in the accuracy of one’s beliefs—may be the most consequential of the many biases in judgment. We offer a theory to explain the ubiquity of overprecision as a consequence of finite cognitive capacity. The individual who considers only a subset of all possibilities or hypotheses will come away excessively confident in focal hypotheses. In our incomplete and simplified model of cognition, individuals employ incomplete and simplified models of the world in they consider only a subset of the “true” (modeled) probability space, but correctly apprehend relative probabilities. The resulting beliefs follow a simple formula in which the probability assigned to all considered hypotheses are scaled up by a constant factor. We show that this framework rationalizes five well-known behavioral phenomena concerning overprecision of beliefs and identifies the kinds of interventions which can improve calibration in judgment.

*Thanks to Ned Augenblick, Stefano DellaVigna, David Broockman, Jacopo Perego, Jett Pettus, Daniel Stone, and Steven Tadelis for thoughtful comments. All remaining errors are our own.

[†]Columbia, NBER, and CEPR

[‡]UC Berkeley

[§]UC Berkeley, corresponding author

Although Bayes' Theorem is a cornerstone of rational judgment, it ranks among the more challenging lessons in any course on decision making. It isn't just our students who struggle, as the many violations of Bayesian reasoning attest (e.g., Bazerman and Moore, 2012; Kahneman and Tversky, 1973). The challenge of applying Bayes' Theorem lies not in the computation itself, which simply calls for the division of two numbers. The hard part is knowing what numbers to use. To begin with, these numbers are probabilities, which requires some infrastructure. For instance, a first essential step to constructing probabilities is enumerating the probability space of everything that could happen. This may require thinking through a nearly or actually infinite number of possibilities.

While there is room for debate about many features of human and animal cognition, finitude is not among them. Attention illuminates a subset of all possible considerations and thus focal hypotheses are more prominent or salient than others. For instance, when considering the question, "How tall is Mount Everest?" it is impractical to consider all possible answers. And this is just a single, univariate question. Correctly applying Bayes' Theorem in realistic multivariate environments ("How much food will each of various possible foraging sites or hunting grounds provide?") with all of the relevant evidence ("...given everything I know or could learn about them") is well beyond human (or other animal's) capacity.

We offer a new formalization of these limitations by focusing on the foundations of probabilistic reasoning. To endow agents with beliefs as probabilities, as is the custom in social sciences, it is necessary to endow them with a model—a *probability space*—that gives them meaning.¹ We impose finitude as a constraint on the enumeration of the elements of the probability space that an individual agent "considers." Combined with an assumption that agents correctly assess relative probabilities, we deliver an intuitive explanation for overprecision in judgment and decision making. This overprecision does not result from misapplying Bayes' Theorem. Rather, we show it can arise from correct application of Bayes' Theorem to a model of the world that is incomplete, reflecting the finite cognitive capacity of its author.² In this sense our model is an example of attribute substitution: a heuristic substituting a computationally feasible attribute judgment for a more accurate but complex one (Kahneman and Frederick,

¹Early probability theory was developed for games of chance, in which the event space and unconditional probabilities are reasonably well-defined. The modern treatment of uncertainty proposes to characterize what we do not know, which may or may not be random in any meaningful sense, as the realization of a fictional stochastic process—that is, a model. Hacking (1975) contends that this leap, which was neither obvious nor uncontroversial, facilitated the development of much of modern science.

²Our theory is based on a model of Bayesian belief updating. However, we confess that we do not believe that everyone is a perfect Bayesian. The important thing for our purposes is that their judgments adhere to the basic logic, if not the precise formulas, of Bayes's Law. We have chosen to base the formalization of our theory on Bayesian reasoning not because we believe people always reason in perfect Bayesian fashion, but because it provides a clearer, less ambiguous, and more parsimonious articulation of our theory than alternative approaches could have provided, while providing a reasonable approximation of the psychological processes involved.

2002). Behavioral bias results when agents employ simplified models of the world.³

Of the different manifestations of overconfidence, overprecision—the excessive certainty in the accuracy of one’s beliefs—is the most pervasive and robust (Moore et al., 2015).⁴ Subjective confidence almost invariably exceeds objective measures of accuracy. Its consequences are widespread, including susceptibility to false news, excessive stock market trading, excessive entrepreneurial entry, insufficient risk hedging, neglect of dissenting perspectives, and political partisanship (Hogarth and Karelaia, 2012; Lyons et al., 2021; Odean, 1998; Ortaleva and Snowberg, 2015; Yaniv, 2004).

In addition to highlighting human limitations in probabilistic reasoning, our approach is also consistent with what we know people are good at: *relative* probability assessments. In general, people are better at comparative than absolute judgment (Stewart et al., 2005). It is easier to determine which of two lights is brighter than it is to estimate their brightness in lumens (Eriksen and Hake, 1955). It is easier to identify which sound has a higher pitch than to report their frequencies in Hertz (Hartman, 1954). Just so, people are better at estimating which event is more likely than they are at estimating the probability of events (Por and Budescu, 2017). Consequently, our theory supposes that agents’ probability assessments for the subset of considered possibilities match the true relative probabilities of those events.

We build on a literature in behavioral economics studying how cognitive constraints can lead to systematic biases in beliefs (Benabou and Tirole, 2002; Rayo and Becker, 2007; van den Steen, 2004). One close predecessor is Gennaioli and Shleifer (2010), who also model agents who only consider a subset of possible outcomes. They focus on how the representativeness heuristic guides consideration of possibilities in a multivariate environment, and do not address overprecision. López-Pérez et al. (2021) discusses how not considering all possibilities can produce narrow confidence intervals. We provide a general approach, which enables us to account for a wider variety of empirical results. Our general approach can also be construed as a special case of behavioral inattention (Gabaix, 2019), and in this way enlarges the set of findings that can be rationalized by that framework.⁵

We begin with a fable of a hominid, Ug, who catches a glimpse of an animal. If it is type *A*, then it will eat him. If it is Type *B*, he can eat it. If it is Type *C*, neither will eat the other. Type *C* stands in for a potentially large set of creatures, from crows to caterpillars.

³In contrast, when social scientists apply simplified models to describe a rich environment, as we do here, it is called “research.” Perhaps more to the point, the rational model, in which it is supposed that agents “correctly” adopt the necessarily simplified and provisional model of reality that the researcher has chosen, is an excellent example of the fallacy of misplaced concreteness.

⁴Overestimation and overplacement (the exaggerated belief in one’s superiority) represent other forms of overconfidence (Johnson and Fowler, 2011).

⁵For the interested reader, the relationship between our model and the “*m*” formulation of behavioral inattention in Gabaix (2019) is explored in the Appendix.

Consistent with standard approaches to modeling beliefs, we treat the evidence that Ug observes in his glimpse as the realization of a random variable, call this E . The relevant probabilities for Ug are then the conditional probabilities $Pr(J|E)$, for $J \in \{A, B\}$. In words, “among the animals that look like E , what proportion are J ’s?” Answering this question will determine the direction in which Ug should run.

If Ug regularly encounters animals with properties E , he may have a good sense of $Pr(J|E)$. That is, this may already be well-formed subjective belief; the fact that he spotted an E in the distance does not mean he needs to update his belief about the relative prevalence of animals on the savanna, just his belief about this particular animal.

If this belief is not already filed away, Ug can construct it using Bayes’ rule:

$$Pr(J|E) = \frac{Pr(J, E)}{Pr(E)}. \quad (1)$$

In addition to assuming that Ug knows how to use Bayes’ rule, standard analysis also assumes that he knows the constituent probabilities required to do this calculation. Is this realistic? $Pr(J, E)$ means “of all the animals I might see on the savanna, what proportion are animal J and look like E ”? Ug could plausibly just have a general sense or approximation of how many J ’s he has seen with properties E , or a sense that, say, 1/12 of the animals he sees are A ’s that look like E .

How about $Pr(E)$, i.e., “what proportion of animals on the savanna look like E ”? Again, if E is something Ug sees frequently, he may have this filed away and ready to use for equation (1). If not, assuming he can do the calculations in the above paragraph to get each $Pr(J, E)$ (and again assuming that A , B , and C are the only possibilities), one path is to use the law of total probability to compute:

$$Pr(E) = Pr(A, E) + Pr(B, E) + Pr(C, E). \quad (2)$$

At this point, the calculations are becoming taxing. Ug may simplify the computations by ignoring C ’s. He may do this because C stands in for a set of possibilities too large to contemplate; alternatively, C is “less salient” since he can’t eat them and they won’t eat him, or perhaps C ’s are the least likely to display property E . Let’s say that, in the true model, $Pr(A, E) = 1/12$, $Pr(B, E) = 1/24$, and $Pr(C, E) = 1/36$. Our first unconventional assumption is that Ug “ignores” C when thinking through what animal he could be facing. In other words, he reasons “as if” $Pr(C, E) = 0$.

This brings us to the question of how ignoring C affects his assessments of the other possibilities $Pr(A, E)$ and $Pr(B, E)$. It may seem most straightforward to assume ignoring C does not affect how he thinks about these possibilities. This could raise two problems. First, as long

as C is possible, dropping $P(C, E)$ from (2) in the computation of assessments by (1) will not yield valid probabilities, since the sum of the assessments he assigns to all events will end up being less than 1. Second, particularly if E is rare or specific, it may be hard to grapple with these assessments in absolute terms: is the proportion of animals that are A s with four legs, a particular hue, and weighing around 400 pounds one in a hundred or one in a thousand? A simple solution to both problems is to suppose that Ug correctly assesses relative probabilities, which is equivalent to assuming his assessment of all probabilities is equal to the truth times a common factor $\lambda > 0$.⁶

Combining these assumptions, he computes his estimate of $Pr(E)$ as:

$$\hat{Pr}(E) = \lambda Pr(A, E) + \lambda Pr(B, E) = 3\lambda/24. \quad (3)$$

and uses this to estimate that the probability of A is:

$$\hat{Pr}(A|E) = \frac{\lambda Pr(A, E)}{\lambda \hat{Pr}(E)} = \frac{1/12}{3/24} = 2/3, \quad (4)$$

and, similarly, $\hat{Pr}(B|E) = 1/3$. We refer to these (potentially incorrect) beliefs as *estimates*, and adorn them with hats. Importantly, the exact choice of the factor λ does not affect this estimate (though, as elaborated in the Appendix, it must be that $\lambda = 1/(Pr(A, E) + Pr(B, E))$ if Ug's estimates are to be interpreted as probabilities; see assumption 3).

Alternatively, if Ug performed these calculations under the true model, he would conclude that the probability of A , conditional on E , was lower:

$$Pr(A|E) = \frac{1/12}{1/12 + 1/24 + 1/36} = 6/11 < 2/3 = \hat{Pr}(A|E). \quad (5)$$

And similarly $Pr(B|E) = 3/11 < 1/3 = \hat{Pr}(B|E)$. By not considering all of the possibilities, he ends up overestimating the probabilities of everything he does consider. The more possibilities he neglects, the larger the distortion.

To generalize our theory, let an agent who wants to form beliefs about a set of hypotheses $\mathcal{X}$ observe some “evidence” $E \in \mathcal{E}$, where $\mathcal{E}$ is the set of all kinds of evidence she could observe. To simplify, we suppose that $\mathcal{X}$ and $\mathcal{E}$ are discrete, though both could potentially be very large or even infinite.

The standard approach to modeling beliefs as probabilities starts by defining the full space of *outcomes* as $\Omega = \mathcal{X} \times \mathcal{E}$; i.e., any combination of a potentially infinite set of hypothesis and a potentially infinite set of realizations of evidence. A fully specified probability model requires probabilities for each of these outcomes sum up to one, i.e. $P(X, E)$, $X \in \mathcal{X}$ and

⁶See assumption 2 in the Appendix.

Figure 1: Illustration of the key mechanism leading to overprecision. In each panel, the thick grey bars represent the true probability distribution if all hypotheses are considered (10 total in the top panels, 20 total in the bottom panels). The thin black bars correspond to the estimate with only a subset considered (three in the left panels, six in the right panels). The relative heights of the considered hypotheses match the truth, but they are inflated upwards by a factor of λ' .

$E \in \mathcal{E}$ such that $\sum_{X \in \mathcal{X}} \sum_{E \in \mathcal{E}} P(X, E) = 1$.

If, instead, the agent only considers a subset of the hypotheses $\hat{\mathcal{X}} \subset \mathcal{X}$ (such as when Ug neglects C), she will wind up with an incomplete representation of the probability space. This could be because the agent does not have time to consider all possible hypotheses that she knows of, or because she is unaware of some hypotheses. The agent may also neglect some kinds of evidence (and, again, may not even be aware of all possibilities): let $\hat{\mathcal{E}} \subseteq \mathcal{E}$ denote the ones she considers.

As in our opening fable, the only constraint we place on the probabilities that the agent assigns to possibilities she does consider is that the relative probabilities match the true model.

It follows, as detailed in the appendix, that such an imperfectly rational Bayesian will overestimate the relative likelihood of the hypotheses she considers:

$$\hat{Pr}(X|E) = \lambda' Pr(X|E), \quad (6)$$

Figure 2: Increasing the number of hypotheses considered decreases confidence in the focal hypothesis, moving it closer to the truth. For each curve the real probability of the focal hypothesis is 50%, with uniform probability among the remaining $n - 1$ hypotheses. For the black curve, there are only four possible hypotheses, so once the agent considers all four she has a correct belief. In the grey curves there are 10 and 20 hypotheses to consider, which increases overprecision and requires the agent consider more hypotheses before her beliefs converge to the truth.

where $\lambda' > 1$. Figure 1 illustrates, comparing true probability distributions (grey bars) to estimates where only a subset of possibilities are considered. The shape of the distribution among considered hypotheses parallels the truth, but all estimates are inflated upwards.

This, then, is the essence of overprecision: Overestimating the likelihood of focal hypotheses. In particular, any individual will end up with overprecise estimates about whatever she views to be the most likely possibility, as by definition this is something she considers. To paraphrase Kahneman (2011), nothing you think about is as likely as you think it is, when you are thinking about it. Our theory may help account for at least five important empirical results.

First, overprecision in judgment is ubiquitous. In virtually every situation, confidence in a focal hypothesis exceeds accuracy, as measured by the actual rate at which people are correct (Harvey, 1997). This fact holds across cultures, races, and genders (Moore and Dev, 2018). It holds across levels of expertise (McKenzie et al., 2008). And it is even evident across scientific fields in published estimates of physical constants (Croft and Dailey, 2015; Henrion and Fischhoff, 1986).

Second, directing attention toward a wider range of possibilities decreases the subjective probability assigned to hypotheses which would have been considered regardless (see figure 2). Indeed, one of the most time-honored and effective debiasing strategies invites people to

Figure 3: Directing attention to hypotheses within a group increases the estimate of the likelihood of the group. In the left panel, only one element of the red group is considered, and in the middle panel all ten elements are considered (the same three black hypotheses are considered in both). While considering more of the red hypotheses decreases the estimate of each individually, their collective estimate increases, which can be seen by the decrease in estimates of the black hypotheses. The right panel shows that when few red hypotheses are considered, the agent is underprecise about the group as a whole, but as more are considered she is eventually overprecise.

explicitly consider alternative hypotheses (Lord et al., 1984).⁷

Similar logic accounts for a third empirical result: that people’s estimates of the probability of a set of events happening (“dying of natural causes”) are subadditive relative to the sum of the estimated probabilities of the set’s constituent elements (Tversky and Koehler, 1994) (see figure 3). When people consider the likelihood of death by natural causes, the probabilities they estimate are lower than the summed probabilities those same people accord to death by cancer, heart disease, or other natural causes. Directing attention to the elements within a set increases estimates of its probability. Belief elicitation formats that focus attention on a particular hypothesis increase the likelihood accorded it (Moore et al., 2015).

Fourth, because focusing attention on a particular hypothesis inflates its subjective probability, doing so with all hypotheses in series can inflate the total. Asking someone to consider the probability of hypotheses one at a time, and treat these as separate estimation problems

⁷This observation explains why we pursued a model of misspecified probability spaces, instead of a model in which agents set some prior probabilities to zero. Although one could generate overprecision with such a model, it lacks the element of finite cognition, and so does not explain why prompting agents to consider a hypothesis would change their beliefs. However, Ortoleva (2012) offers a model with priors over priors, in which agents occasionally reconsider the subset of probabilities they set to zero, sometimes inducing paradigm shifts.

with different consideration sets, then the summed estimates will exceed 1. A famous illustration comes from a study in which basketball fans estimated, one at a time, the probability that each of the eight teams headed into the NCAA quarterfinals would win the tournament; they summed to 240 percent (Fox and Tversky, 1998).

Fifth, by extending our theory to a multivariate environment, it also provides an explanation for the conjunction fallacy, where the probability assigned to the event “ X and Y ” (famously, “Linda is a feminist and bank teller”) can be higher than just X (“Linda is a feminist”) (Tversky and Kahneman, 1983). In short, this can happen because bringing Y into the mix necessitates holding more potential hypotheses in one’s head, leaving open more opportunity to not consider all of them. In this example, if raising both the feminist status and bank teller status leads the agent to forget to consider the non-feminist non-bank-tellers (perhaps the modal category!), this leads to more overprecision when considering X and Y than just X .

Our theory has the clear implication that overprecision will be a property of all finite thinking systems, including animals. Being too sure of their estimates of, say, the value of foraging at a particular site, might lead them to explore too little (Garrett and Daw, 2020; Stephens and Krebs, 1986). Overprecision should also afflict the judgments of artificially intelligent agents, since they too must contend with a complex world using finite computational capacities. Indeed, the problem of overconfident AI has been known for some time (Oelrich et al., 2020).

We close by noting the relevance of an epistemological challenge described elsewhere as “theory uncertainty” or “model uncertainty” (Hansen and Sargent, 2001; Levine et al., 2002). Scholars attempt to understand the world by developing theories and applying models to account for empirical data. But our models will be incorrect or incomplete for reasons that we do not anticipate; if we had anticipated them, we could have accounted for them. In the absence of the ability to consider all possible models, there will always be unknown unknowns that put us at risk of being too sure our theories are correct. Consequently, none of our theories will be as likely to be correct as we think they are, when we are thinking about them. Anticipating this, we ought to adjust our confidence downward, but determining precisely how much depends on us knowing what we do not know.

References

- M. H. Bazerman and D. A. Moore. *Judgment in managerial decision making*. John Wiley & Sons, 2012.
- R. Benabou and J. Tirole. Self-confidence and personal motivation. *Quarterly Journal of Economics*, 117(3):871–915, 2002.
- R. A. C. Croft and M. Dailey. On the measurement of cosmological parameters. *arXiv:1112.3108 [astro-ph, physics:gr-qc]*, Jul 2015.
- C. W. Eriksen and H. W. Hake. Multidimensional stimulus differences and accuracy of discrimination. *Journal of Experimental Psychology*, 50(3):153–160, 1955. ISSN 0022-1015(Print). doi: 10.1037/h0047863.
- C. R. Fox and A. Tversky. A belief-based account of decision under uncertainty. *Management Science*, 44(7):879—895, 1998.
- X. Gabaix. Behavioral inattention. *Handbook of Behavioral Economics*, 2:261–343, 2019.
- N. Garrett and N. D. Daw. Biased belief updating and suboptimal choice in foraging decisions. *Nature Communications*, 11(11):3417, Jul 2020. ISSN 2041-1723. doi: 10.1038/s41467-020-16964-5.
- N. Gennaioli and A. Shleifer. What comes to mind. *The Quarterly journal of economics*, 125(4):1399–1433, 2010.
- I. Hacking. *The Emergence of Probability: A Philosophical Study of Early Ideas about Probability Induction and Statistical Inference*. Cambridge University Press, 1975.
- L. Hansen and T. J. Sargent. Robust control and model uncertainty. *American Economic Review*, 91(2):60–66, 2001. ISSN 0002-8282.
- E. B. Hartman. The influence of practice and pitch-distance between tones on the absolute identification of pitch. *The American Journal of Psychology*, 67(1):1–14, 1954. ISSN 0002-9556. doi: 10.2307/1418067.
- N. Harvey. Confidence in judgment. *Trends in Cognitive Sciences*, 1(2):78–82, 1997.
- M. Henrion and B. Fischhoff. Assessing uncertainty in physical constants. *American Journal of Physics*, 54(9):791–798, 1986.

- R. M. Hogarth and N. Karelaia. Entrepreneurial success and failure: Confidence and fallible judgment. *Organization Science*, 23(6):1733–1747, 2012.
- D. D. P. Johnson and J. H. Fowler. The evolution of overconfidence. *Nature*, 477(7364): 317–320, 2011.
- D. Kahneman. *Thinking fast and slow*. Farrar, Straus and Giroux, 2011.
- D. Kahneman and S. Frederick. Representativeness revisited: Attribute substitution in intuitive judgement. In T. Gilovich, D. Griffin, and D. Kahneman, editors, *Heuristics and Biases: The Psychology of Intuitive Judgement*, pages 49–81. Cambridge University Press, 2002.
- D. Kahneman and A. Tversky. On the psychology of prediction. *Psychological review*, 80(4): 237, 1973.
- E. S. Levine, A. E. Schulz, and M. White. Future galaxy cluster surveys: The effect of theory uncertainty on constraining cosmological parameters. *The Astrophysical Journal*, 577(2): 569, 2002.
- R. López-Pérez, A. Rodríguez-Moral, and M. Vorsatz. Simplified mental representations as a cause of overprecision. *Journal of Behavioral and Experimental Economics*, page 101681, Jun 2021. ISSN 2214-8043. doi: 10.1016/j.socec.2021.101681.
- C. G. Lord, M. R. Lepper, and E. Preston. Considering the opposite: A corrective strategy for social judgment. *Journal of Personality and Social Psychology*, 47(6):1231—1243, 1984.
- B. A. Lyons, J. M. Montgomery, A. M. Guess, B. Nyhan, and J. Reifler. Overconfidence in news judgments is associated with false news susceptibility. *Proceedings of the National Academy of Sciences*, 118(23), Jun 2021. ISSN 0027-8424, 1091-6490. doi: 10.1073/pnas.2019527118. URL <https://www.pnas.org/content/118/23/e2019527118>.
- C. R. M. McKenzie, M. J. Liersch, and I. Yaniv. Overconfidence in interval estimates: What does expertise buy you? *Organizational Behavior & Human Decision Processes*, 107: 179–191, 2008.
- D. A. Moore and A. S. Dev. Individual differences in overconfidence. In V. Zeigler-Hill and T. Shackelford, editors, *Encyclopedia of Personality and Individual Differences*. Springer, 2018.
- D. A. Moore, E. R. Tenney, and U. Haran. Overprecision in judgment. *The Wiley Blackwell handbook of judgment and decision making*, 2:182–209, 2015.

- 248 T. Odean. Are investors reluctant to realize their losses? *The Journal of finance*, 53(5):1775–
249 1798, 1998.
- 250 O. Oelrich, S. Ding, M. Magnusson, A. Vehtari, and M. Villani. When are bayesian model
251 probabilities overconfident? *arXiv:2003.04026*, Mar 2020.
- 252 P. Ortoleva. Modeling the change of paradigm: Non-bayesian reactions to unexpected news.
253 *American Economic Review*, 102(6):2410–2436, 2012.
- 254 P. Ortoleva and E. Snowberg. Overconfidence in political behavior. *American Economic Re-*
255 *view*, 105(2):504–535, 2015.
- 256 H.-H. Por and D. V. Budescu. Eliciting subjective probabilities through pair-wise comparisons.
257 *Journal of Behavioral Decision Making*, 30(2):181–196, 2017.
- 258 L. Rayo and G. S. Becker. Evolutionary efficiency and happiness. *Journal of Political Econ-*
259 *omy*, 115(2):302–337, 2007.
- 260 D. W. Stephens and J. R. Krebs. *Foraging theory*. 1986.
- 261 N. Stewart, G. D. A. Brown, and N. Chater. Absolute identification by relative judgment.
262 *Psychological Review*, 112:881–911, 2005.
- 263 A. Tversky and D. Kahneman. Extensional versus intuitive reasoning: The conjunction fallacy
264 in probability judgment. *Psychological review*, 90(4):293, 1983.
- 265 A. Tversky and D. J. Koehler. Support theory: A nonextensional representation of subjective
266 probability. *Psychological Review*, 101(4):547—567, 1994.
- 267 E. van den Steen. Rational overoptimism (and other biases). *American Economic Review*, 94
268 (4):1141–1151, 2004.
- 269 I. Yaniv. Receiving other people’s advice: Influence and benefit. *Organizational behavior and*
270 *human decision processes*, 93(1):1–13, 2004.

A Appendix

Deriving the main result

The “real” model Take a set of hypotheses $\mathcal{X}$ and possible realizations of evidence $\mathcal{E}$. The full outcome space is every possible combination of hypothesis and evidence $\Omega = \mathcal{X} \times \mathcal{E}$. Let $P(X, E)$ be the objective probability of hypothesis $X \in \mathcal{X}$ and evidence $E \in \mathcal{E}$. For P to be a proper probability distribution, it must be the case that $\sum_{X \in \mathcal{X}} \sum_{E \in \mathcal{E}} P(X, E) = 1$.

Assumptions Suppose an agent forms beliefs that only include a limited *consideration set* of the hypotheses $\hat{\mathcal{X}} \subseteq \mathcal{X}$. The agent may also not include in her model all possible kinds of evidence which could be revealed (and, again, may not even be aware of all possibilities): let $\hat{\mathcal{E}} \subseteq \mathcal{E}$ be the ones considered. All that will matter for what follows is that the evidence which is actually observed is a part of $\hat{\mathcal{E}}$. Formally:

Assumption 1 (Finitude). *The outcomes considered by the agent are $\hat{\Omega} = \hat{\mathcal{X}} \times \hat{\mathcal{E}}$, where $\hat{\mathcal{X}} \subseteq \mathcal{X}$ and $\hat{\mathcal{E}} \subseteq \mathcal{E}$.*

To complete the agent’s model, we define a probability for each outcome considered, i.e., for $\hat{\Omega} = \hat{\mathcal{X}} \times \hat{\mathcal{E}}$. To distinguish them from the “real” probabilities P , we designate the agent’s probabilities for outcomes with Q . How do these Q s relate to the true P s?

It may seem that the simplest assumption would be to set $Q(X, E) = P(X, E)$. But if $\hat{\Omega}$ does not include some outcomes which receive positive probability in the real model, then $\sum_{X \in \hat{\mathcal{X}}} \sum_{E \in \hat{\mathcal{E}}} Q(X, E) < 1$. That is, the agent does not have a proper probability distribution over the outcomes she thinks about. Intuitively, if the agent thinks through her entire model, she may realize that something must be missing because probabilities do not sum to 1. If we allow the agent to have a model which is not only incorrect but incomplete (and don’t mind abusing terminology to call the Q ’s probabilities), then our main result holds.

There are an infinite number of ways that the Q ’s can differ from the P ’s, which may or may not be a proper probability model. The constraint we place on these is that the *relative* probabilities always correct. That is, for any (X, E) and (X', E') such that $P(X', E') > 0$:

$$\frac{Q(X, E)}{Q(X', E')} = \frac{P(X, E)}{P(X', E')}. \quad (7)$$

Rearranging, we can write this:

$$Q(X, E) = \frac{P(X', E')}{Q(X', E')} P(X, E). \quad (8)$$

Since this is true for all (X, E) , we can think of this as taking each $P(X, E)$ and scaling it up

by $\frac{P(X', E')}{Q(X', E')} \equiv \lambda$. In other words, once we “set” $Q(X', E')$ (for any event which gets assigned positive probability in the real model), this uniquely determines how to scale up all of the other Q ’s. Formally:

Assumption 2. *[At least relative probabilities] There exists a constant $\lambda > 0$ such that $Q(X, E) = \lambda P(X, E)$.*

If $\lambda = 1$ then Q s and P s are equal. More generally, if the agent’s “probabilities” differ from the real model, all are multiplied by the same constant, λ .

The last assumption, which does not affect our main result, but does ensure that the agent has a proper probability model, is to set this constant such that $\sum_{X \in \hat{\mathcal{X}}} \sum_{E \in \hat{\mathcal{E}}} Q(X, E) = 1$.

Assumption 3. *[Valid probability model] $\lambda = (\sum_{X \in \hat{\mathcal{X}}} \sum_{E \in \hat{\mathcal{E}}} P(X, E))^{-1}$.*

There are two ways to think about this pair of assumptions. One is that agents think in terms of cardinal probabilities, and set these probabilities in a way that (1) has a consistent relationship with the truth, and (2) get “corrected” if the agent realizes they do not sum to 1. For instance, when invited to estimate the percentage of all deaths attributable to natural causes, if a person fails to consider neonatal fatalities, they would have to notice that there is some probability unaccounted-for and might correct their estimates by scaling them up so that they sum to 1 (Tversky and Koehler, 1994).

But if people estimate deaths relative to those causes they think of, unconsidered hypothesis will go unnoticed. Our formalization can accommodate the possibility that the agent does not think of cardinal probabilities in the first place, but only in relative likelihoods. Comparing two options (X', E') and (X'', E'') could return a ratio $Q(X', E')/Q(X'', E'')$, which by assumption 2 must be equal to the ratio of probabilities from the true model $Pr(X', E')/Pr(X'', E'')$. If the agent compares (X', E') to a set of outcomes Ω'' , then this could return $\frac{Q(X', E')}{\sum_{(X'', E'') \in \Omega''} Q(X'', E'')}$ which will return the proper relative probability. Setting the estimate of each event X, E to the likelihood relative to a consideration set including all of $\hat{\mathcal{X}} \times \hat{\mathcal{E}}$ will ensure that Assumption 3 is met. More generally, if people estimate probabilities only relative to a consideration set, it is impossible to realize that the set is incomplete.

Result Given our motivation, we are primarily concerned with the agent’s beliefs about focal hypotheses given some realization of evidence. Write the belief that someone with the real model would hold as $Pr(X|E)$, and the estimate held by the agent with the limited model as $\hat{P}r(X|E)$. A formal statement of the question of how not considering all possibilities affects beliefs given evidence is then “what is the relationship between $\hat{P}r(X|E)$ and $Pr(X|E)$?”

Our main result presumes that the agent does not make any additional mistakes in how she applies Bayes’ rule; she just applies Bayes’ rules given her incomplete view of the possible

hypotheses encoded in assumption 1-2. This implies that her conditional beliefs are given by:

$$\hat{Pr}(X|E) = \frac{Q(X, E)}{\sum_{X' \in \hat{\mathcal{X}}} Q(X', E)} = \frac{\lambda P(X, E)}{\sum_{X' \in \hat{\mathcal{X}}} \lambda P(X', E)} = \frac{P(X, E)}{\sum_{X' \in \hat{\mathcal{X}}} P(X', E)}. \quad (9)$$

If $\hat{\mathcal{X}} = \mathcal{X}$, i.e., all hypotheses are considered, then the leftmost expression is just what the agent with a full model would get from applying Bayes' rule. However, if $X \subset \mathcal{X}$, then not considering all hypotheses will make the denominator smaller, and hence the relevant belief higher, or, overprecise.

Proposition 1. *Suppose an agent with a probability model as specified in assumption 1 and 2 observes evidence E , and does not consider at least one hypothesis which is possible given E ($Pr(X^O, E) > 0$ for some $X^O \notin \hat{\mathcal{X}}$). Then for all hypotheses which are considered ($X \in \hat{\mathcal{X}}$):*

$$\hat{Pr}(X|E) = \lambda' Pr(X|E) \quad (10)$$

where $\lambda' > 1$.

Proof Set $\lambda' = \frac{\sum_{X \in \mathcal{X}} P(X, E)}{\sum_{X \in \hat{\mathcal{X}}} P(X, E)}$. Equation 10 follows immediately from equation 9. If there is an $X^O \notin \hat{\mathcal{X}}$ such that $Pr(X^O, E) > 0$, then $\lambda' > 1$. ■

Note that the degree to which the agent's conditional beliefs are overprecise is not affected by the fact that she may fail to consider all possible evidence, since the only relevant probabilities are those conditional on E . What leads to overprecision is failing to consider hypotheses which are possible given E .

Relationship to theories of behavioral inattention

Gabaix (2019) offers a unifying framework for theories of behavioral inattention in which an agent solves a utility maximization problem given by

$$\max_a u(a, x, m), \quad (11)$$

where $u(\cdot)$ is a utility function, a is a choice variable (e.g., a purchase choice), x is a vector of J parameters of the problem (e.g., characteristics or prices), and m is a vector of length J where m_j parameterizes the degree of behavioral inattention to parameter x_j . The bounds on m have simple interpretations: if m is the 1-vector then the problem is equivalent to the classical optimization problem, which can be written $\max_a u(a, x)$. Alternatively, if $m_j < 1$ for some j then the agent exhibits behavioral inattention. This is defined with respect to a default value of characteristics, given by x^d . For $m = 0$, the agent solves the problem exactly

at $u(a, x^d)$ irrespective of the true values of x , while intermediate values of m convexify the two bounds. A simple formulation of this is that an inattentive agent maximizes as if $x_j = m_j x + (1 - m_j)x_j^d$.

The extension of this approach to our model is straightforward. Rather than thinking of x as a vector of parameters, think of it as a probability vector. Let $m_j = 1$ if hypothesis j is in the consideration set and $m_j = 0$ otherwise. Then our agent would make a decision as if $x_j = \lambda m_j x + (1 - m_j)0$, where λ is the over-precision factor defined above. This entails two simplifications and one extension to the simple formulation described by Gabaix (2019). The simplifications are that m_j is either 0 or 1, and that the default for each “parameter” is $x_j^d = 0$ (i.e., hypotheses not thought of are assigned probability 0). The complication is that there are “spillovers” in inattention, as the λ factor depends on which hypotheses are considered. I.e., not considering one variable affects the perceived value of the others.

A consequence of this extension is that, in addition to the seventeen behavioral findings rationalized by behavioral inattention in Gabaix (2019), we can add the five additional findings that we rationalize using our model.

Extending to explain the empirical results

In this section, we give formal statements for results that speak to the groups of empirical results referenced in the main text.

The first set of empirical results is that beliefs tend to be overprecise, which is a consequence of proposition 1.

The second set of results is that directing attention to a wider range of possibilities decreases overprecision, or the probabilities assigned to individual events more generally. This follows from considering what happens when the set of considered hypotheses increases. Since this is a comparison of the estimates with different consideration sets, we add an argument for the consideration set in our notation for the estimate of event X given evidence E : $Pr(X|E; \hat{\mathcal{X}})$.

Proposition 2. *Suppose assumptions 1-3 hold. Take any $\hat{\mathcal{X}}^0 \subset \mathcal{X}$, and $X' \in \mathcal{X} \setminus \hat{\mathcal{X}}^0$ such that $P(X', E) > 0$. Let $\hat{\mathcal{X}}^1 = \hat{\mathcal{X}}^0 \cup X'$. Then $Pr(X|E; \hat{\mathcal{X}}^1) < Pr(X|E; \hat{\mathcal{X}}^0)$ for all $X \in \hat{\mathcal{X}}^0$.*

Proof By proposition 1, the estimates given consideration set $\hat{\mathcal{X}}^i$ are:

$$\hat{Pr}(X|E; \hat{\mathcal{X}}^i) = \frac{P(X, E)}{\sum_{X \in \hat{\mathcal{X}}^i} P(X, E)}. \quad (12)$$

Since $\hat{\mathcal{X}}^1 = \hat{\mathcal{X}}^0 \cup X'$, $\sum_{X \in \hat{\mathcal{X}}^0} P(X, E) < \sum_{X \in \hat{\mathcal{X}}^1} P(X, E)$, and so $\hat{Pr}(X|E; \hat{\mathcal{X}}^0) > \hat{Pr}(X|E; \hat{\mathcal{X}}^1)$. ■

The third set of results is that the estimate of a group of events increases when asking subjects to think through individual elements of the group. The corresponding formal result shows how breaking forcing an agent to think through more elements within a group increases the estimate assigned to the group as a whole.

Proposition 3. *Suppose assumptions 1-3 hold, and $\hat{\mathcal{X}} \subset \mathcal{X}$. Let $G \subseteq \mathcal{X}$ be a group of outcomes, and assume not all of these are in $\hat{\mathcal{X}}$: $\hat{\mathcal{X}}^G = \hat{\mathcal{X}} \cap G \subset G$. Let $\hat{\mathcal{X}}' = \hat{\mathcal{X}} \cup X'$ be a new consideration set with an additional element of $G \setminus \hat{\mathcal{X}}$, and $\hat{\mathcal{X}}^{G'} = \hat{\mathcal{X}}^G \cup X'$. Then:*

$$\sum_{X \in \hat{\mathcal{X}}^G} \hat{Pr}(X|E; \hat{\mathcal{X}}) < \sum_{X \in \hat{\mathcal{X}}^{G'}} \hat{Pr}(X|E; \hat{\mathcal{X}}'). \quad (13)$$

Proof The estimate about the probability of being in the initial group $\hat{\mathcal{X}}^G$ is:

$$\sum_{X \in \hat{\mathcal{X}}^G} \hat{Pr}(X|E; \hat{\mathcal{X}}) = \frac{\sum_{X \in \hat{\mathcal{X}}^G} P(X, E)}{\sum_{X \in \hat{\mathcal{X}}} P(X, E)} \quad (14)$$

and the “new” group $\hat{\mathcal{X}}^{G'}$:

$$\sum_{X \in \hat{\mathcal{X}}^{G'}} \hat{Pr}(X|E; \hat{\mathcal{X}}') = \frac{\sum_{X \in \hat{\mathcal{X}}^{G'}} P(X, E)}{\sum_{X \in \hat{\mathcal{X}}'} P(X, E)} = \frac{\sum_{X \in \hat{\mathcal{X}}^G} P(X, E) + P(X', E)}{\sum_{X \in \hat{\mathcal{X}}} P(X, E) + P(X', E)} \quad (15)$$

$$> \frac{\sum_{X \in \hat{\mathcal{X}}^G} P(X, E)}{\sum_{X \in \hat{\mathcal{X}}} P(X, E)} \quad (16)$$

$$= \sum_{X \in \hat{\mathcal{X}}^G} \hat{Pr}(X|E; \hat{\mathcal{X}}). \quad \blacksquare \quad (17)$$

The fourth set of results is that asking when asking an agent to assess all of the individual possibilities, the probabilities sum to more than one. A formal statement of this is:

Proposition 4. *Suppose an agent is asked separately to estimate $\hat{Pr}(X_i|E; \mathcal{X}_i)$ for each $X_i \in \mathcal{X}$, and $\hat{\mathcal{X}}_i$ is the comparison set when forming the belief about X_i . Then if there exists an X^O such that $Pr(X^O, E) > 0$ and X^O is not in all of the $\hat{\mathcal{X}}_i$ ’s, then:*

$$\sum_{X_i \in \mathcal{X}} \hat{Pr}(X_i|E; \mathcal{X}_i) > 1. \quad (18)$$

Proof For the i such that X^O is not in $\hat{\mathcal{X}}_i$, $\hat{Pr}(X_i|E) > Pr(X_i|E)$, and for all other estimates $\hat{Pr}(X|E) \geq Pr(X|E)$, and so

$$\sum_{X_i \in \mathcal{X}} \hat{Pr}(X_i|E; \mathcal{X}_i) > \sum_{X_i \in \mathcal{X}} Pr(X_i|E) = 1. \quad (19)$$

The fifth set of results our approach can explain is the conjunction fallacy (Tversky and Kahneman 1983; see Gennaioli and Shleifer 2010 for a related explanation). Put in our notation, suppose our agent forms a belief about the value of two variables given E , call these X and Y . She commits the conjunction fallacy if her estimate of a value of X and Y co-occurring given E is higher than her estimate of X given E (or Y given E). The intuition for why this might occur is that when considering both X and Y , there are more joint possibilities to think through, and hence more potential to be overprecise.

To tie this to the famous “Linda the bank teller” example (Tversky and Kahneman, 1983), let X_1 mean Linda is a bank teller and X_2 mean not a bank teller, and Y_1 mean active in the feminist movement while Y_2 means not active. The evidence E is the story about her (Linda is outspoken, bright, and politically active). Given this story, $Pr(X_1, Y_2|E)$ is plausibly quite small: it is unlikely that Linda is a non-feminist bank teller. It is also plausible that $Pr(X_2, Y_2|E)$ is reasonably large: there are lots of non-feminist-non-bank-tellers, even among former activists (maybe her job is in another domain of social activism). So this may be exactly the kind of distribution where we would expect to see the conjunction fallacy for “not thinking things through” reasons. To pick some particular numbers, say the truth is:

$$\begin{aligned} Pr(X_1, Y_1|E) &= .09 & Pr(X_1, Y_2|E) &= .01 \\ Pr(X_2, Y_1|E) &= .51 & Pr(X_2, Y_2|E) &= .39. \end{aligned}$$

For each of these $Pr(X, Y, E) = Pr(X, Y|E)Pr(E)$, though as we will see it is not necessary to specify $Pr(E)$. This distribution implies that there are few bank tellers that match E , but among the bank tellers the majority are active in the feminist movement. Among the non-bank-tellers meeting description E , slightly more are active in the feminist movement, but the difference is not as large as among the bank tellers.⁸

Of course, the true $Pr(X_1|E) = .1 > Pr(X_1, Y_1|E)$.

If the agent neglects to think through (X_2, Y_2) when estimating $Pr(X_1, Y_1|E)$ she gets:

$$\hat{Pr}(X_1, Y_1|E) = \frac{.09Pr(E)}{.09Pr(E) + .01Pr(E) + .51Pr(E)} \approx .15 \quad (20)$$

which is higher than $Pr(X_1|E) = .1$, i.e., the conjunction fallacy emerges.

In words, the conjunction fallacy arises when (1) bringing in the extra dimension leads the agent to not consider events that are non-trivially likely, and (2) the true probability of the conjunction is not much lower than the true marginal probability. More generally:

Proposition 5. *Take any true conditional distribution on $X \times Y$, $Pr(X, Y|E)$. Let $\hat{\mathcal{X}}_X$*

⁸Formally, conditional on observing E , bank teller and feminist become negatively correlated.

be the set of possibilities though through when asked to form an estimate $\hat{Pr}(X|E)$ and $\hat{\mathcal{X}}_{XY}$ be the set thought through when asked to form an estimate $\hat{Pr}(X, Y|E)$. Then if (i) $\sum_{(X_i, Y_i) \in \hat{\mathcal{X}}_{XY}} Pr(X_i, Y_i|E) < \sum_{X_i \in \hat{\mathcal{X}}_X} Pr(X_i|E)$ and (ii) $\sum_{Y_j \neq Y_i} Pr(X_i, Y_j|E)$ is sufficiently small, then $\hat{Pr}(X_i, Y_i|E) > \hat{Pr}(X_i|E)$

Proof The two relevant estimates are:

$$\hat{Pr}(X_i, Y_i|E) = \frac{P(X_i, Y_i, E)}{\sum_{(X_i, Y_i) \in \hat{\mathcal{X}}_{XY}} Pr(X_i, Y_i|E)} \quad (21)$$

$$\hat{Pr}(X_i|E) = \frac{P(X_i, E)}{\sum_{X_i \in \hat{\mathcal{X}}_X} Pr(X_i, E)} = \frac{P(X_i, Y_i, E) + \sum_{Y_j \neq Y_i} Pr(X_i, Y_j|E)}{\sum_{X_i \in \hat{\mathcal{X}}_X} Pr(X_i, E)}. \quad (22)$$

Assumption (i) implies the denominator in (21) is strictly larger than the denominator in equation (22). So as $\sum_{Y_j \neq Y_i} Pr(X_i, Y_j|E) \rightarrow 0$ and the numerators approach equality, $\hat{Pr}(X_i, Y_i|E) > \hat{Pr}(X_i|E)$. ■