

## Reasoning with heuristics

Brett Karlan

University of Pittsburgh

*Penultimate draft of a paper to appear in Ratio. Please cite the published version whenever possible.*

Which rules should guide our reasoning? Human reasoners often use reasoning shortcuts, called *heuristics*, which function well in some contexts but lack the universality of reasoning rules like deductive implication or inference to the best explanation. Does it follow that human reasoning is hopelessly irrational? I argue: no. Heuristic reasoning often represents human reasoners reaching a *local rational maximum*, reasoning more accurately than if they try to implement more “ideal” rules of reasoning. I argue this is a genuine rational achievement. Our ideal rational advisors would advise us to reason with heuristic rules, not more complicated ideal rules. I argue we do not need a radical new account of epistemic norms to make sense of the success of heuristic reasoning.

### 1. Introduction<sup>1</sup>

How rational is typical human reasoning? Starting in the 1970s with work on biases in psychology (Kahneman & Tversky 1972), a popular answer became: shockingly irrational. Human reasoners use faulty reasoning rules, called *heuristics*, that led to irrational patterns of reasoning. Human reasoners using these rules were described as “systematically irrational” (Armstrong & Minderman 2018, pp. 5607), and the outlook for human rationality was thought to be “bleak” (Nisbett & Borgida 1975, pp. 935).

Perhaps predictably, other authors pushed back against this view of human rationality. In cognitive science, advocates of bounded (e.g. Simon 1997), ecological (e.g. Gigerenzer & Selten 2002), or resource (e.g. Lieder & Griffiths 2020) rationality all argued that human beings only look irrational on an implausibly idealized view of rationality. Norms of deductive inference and probabilistic coherence “promote standards of rationality which it is simply impossible to live up to” (Ryseiw 2008, pp. 1164), and “people can, and often should, use very reliable [heuristics] that ignore lots of evidence and do not properly integrate the evidence they

---

<sup>1</sup> I wish to extend my heartfelt thanks to Thomas Kelly, Grace Helton, Gideon Rosen, Kyle Landrum, Chris Register, Snow Zhang, Daniel Kranzelbinder, Joseph Moore, Eleanor Gordon-Smith, Haley Brennan, Thomas Lambert, Noah Apthorpe, Ari Seff, an anonymous reviewer, and participants in the Princeton Cognitive Science Colloquium for comments and discussions on earlier versions of this paper.

do consider” (Bishop 2006, pp. 217). Some epistemologists have made related arguments, arguing that certain so-called irrational “biases” in fact represent perfectly rational responses to evidence.<sup>2</sup>

In this paper, I approach the debate from a slightly different angle. I want to describe what kind of rational accomplishment heuristic reasoning amounts to, if it is an accomplishment. I argue that reasoning with heuristics often allows reasoners to reach a *local rational maximum* in rationality space, reasoning to more accurate conclusions than if we tried to use “ideal” rules of reasoning. Utilizing the idea of an ideal normative advisor (Railton 1986; Smith 1995), I suggest that our advisors would recommend that creatures like us reason with heuristics. Instead of proposing a new theory of rationality for making sense of heuristic reasoning, as many authors discussed below do, I aim to show how the latest cognitive science research in heuristic reasoning can be easily and fruitfully integrated into the kind of accuracy-oriented epistemology many already adopt.

## 2. Heuristic and ideal reasoning rules

As I will use the term, *reasoning* is a rule-based transition between premise-states (beliefs, credences, and other propositionally-structured mental states) and a conclusion-state.<sup>3</sup> The debate about human reasoning concerns the rationality of different rules for transitioning from premises-states to conclusion-states. In particular, it is a debate about whether heuristic reasoning rules are rationally defective when compared to what we might call *ideal* rules of reasoning. It is difficult to give necessary and sufficient conditions for an ideal rule of reasoning,

---

<sup>2</sup> Examples include belief polarization (Kelly 2008) and hindsight bias (Hedden 2019).

<sup>3</sup> This is consistent with the way many use the term (e.g. Siegel 2019; Richard 2019), but it is not wholly uncontroversial. A strongly intellectualist account of reasoning, which holds the Taking condition (Boghossian 2014) to be something the subject must actively (and perhaps consciously) employ, will not view the transitions I discuss here as instances of reasoning proper. This is fine: advocates of such a view should replace “reasoning” with “rule-based transitions between premise- and conclusion-states” in what follows. These transitions are epistemically interesting in their own right.

but there are some canonical examples. Deductive implication is an important ideal rule of reasoning (e.g. Dogramaci 2015). Probabilistic induction (e.g. Earman 1992) and inference to the best explanation (e.g. Lipton 2005) seem to fit the bill as well.

Against this backdrop, a heuristic reasoning rule can be seen as a kind of reasoning *shortcut*. An ideal reasoning strategy involves considering all of the relevant evidence available, and operating on it with deductive and inductive reasoning rules, to come to your final conclusion. A heuristic reasoning rule, in contrast, recommends that you come to a conclusion without considering all relevant evidence, or by using a non-standard method of transitioning between premise-states and conclusion-states. It might recommend you ignore certain bits of potentially relevant evidence: if premises P1 - P3 support some conclusion, you might ignore P4 - P5 in further deliberation, since this might only confuse or mislead. It might also recommend attribute substitution (Kahneman & Frederick 2002): if you cannot reason through problem A, you might switch to an easier problem B, which is related to but distinct from A.

My argument that heuristic reasoning often represents a local rational maximum for limited reasoners is simple:

- (1) Heuristic reasoning often maximizes accuracy for limited reasoners, especially in information-rich environments.
- (2) Maximizing accuracy is the overriding rational good in these reasoning contexts; so
- (3) Heuristic reasoning often maximizes rationality for limited reasoners.

I first defend the claim that heuristic reasoning is often more accurate than attempts to implement “ideal” reasoning rules. I then defend the philosophical picture of non-ideal reasoning that results.

### **3. The accuracy of heuristic reasoning: some examples**

Perhaps the ideal reasoner could get by only using ideal reasoning rules.<sup>4</sup> But human beings are limited reasoners. Our choice is not between ideal rules ideally implemented and heuristic rules; it is between ideal rules non-ideally implemented and heuristic rules. There is good empirical evidence that heuristic reasoning often leads reasoners to more accurate conclusions than trying (and failing) to implement more ideal rules of reasoning.

A particularly simple heuristic, one that has occupied much recent discussion, is the *recognition heuristic* (Gigerenzer & Goldstein 2011). The recognition heuristic is deployed in situations where one must make a choice between two or more options and one is able to recognize one of the options but not the others. Suppose one must determine whether Munich or Dortmund is a larger city, and one has no knowledge of demographics nor access to Google. Often, subjects will (correctly) come to believe Munich to be the larger city purely on the basis of recognition. In forced-choice situations, the recognition heuristic offers a strategy for coming to an answer when one has very little access to other information that might be relevant.

That humans use the recognition heuristic has been widely demonstrated (Scheibehenne & Bröder 2007; Gaissmaier & Marewski 2010). It seems to outperform other reasoning strategies, at least for problems like city recognition (Borges et al. 1999; Alfano & Skorburg 2018). But I think its usefulness has been slightly mischaracterized. Recognition primarily represents a strategy of accurate *guessing*, rather than belief formation proper. In the Munich/Dortmund example, for instance, it is crucial that those using the recognition heuristic have no other options available to them to seek out more accurate information. If they did, it would be rational for them to gather more information, rather than using recognition to guess. It furthermore seems plausible that the standards for rational guessing are more permissive than those for rational belief. If I have no conclusive evidence in favor of one answer to some question over

---

<sup>4</sup> Though see Egan (2008).

another, but I must make a selection regardless, using a process that gives me a better than 50% chance of being right will be rational for me. In this sense, the recognition heuristic represents a rational response for limited reasoners to take towards (very) limited information. But this is not the same thing as a rule for accurate belief formation in the presence of robust evidence.

Some have argued this is mistaken. They believe the recognition heuristic also produces rational *beliefs* in certain contexts. The most well-discussed instance of this concerns the stock market. Recognition-based heuristics for choosing which stocks to invest in have been claimed to outperform expert money managers operating with ideal reasoning rules (Borges et al. 1999). But it is controversial whether such effects replicate.<sup>5</sup> Furthermore, overreliance on the recognition heuristic has been shown to rationally imperil all but the most well-informed reasoners, leading to silencing of members of groups that are less likely to be recognized by dominant group members (Alfano & Skorborg 2018). Given these limitations, I think it best to view the recognition heuristic as primarily a strategy for rational guessing. Such strategies are important: in situations where we have to guess, we should do so in ways that are rational. But it is a mistake to infer that we use the recognition heuristic primarily as a method of (non-guessing) belief formation. The next two heuristics discussed represent reasoning strategies for rational belief formation, not merely guessing.

A more complicated set of heuristic reasoning rules are found in the suite of *one-reason* heuristics. The Take the Best (TTB) heuristic (Gigerenzer et al. 1991) is the classic example. When trying to decide which of several options is more choiceworthy, TTB recommends looking for a single salient value that differentiates between two options. If you are trying to

---

<sup>5</sup> The initial results from Borges et al. (1999) failed to replicate in Boyd (2001). Ortmann et al. (2008) criticized this failed replication for sampling only college students, rather than pedestrians on the street, which they argued watered down the efficacy of the recognition heuristic. Another failed replication followed in Andersson & Rakow (2007), however, which avoided some of the errors of Boyd (2001), though the initial results continue to be cited by proponents of the recognition heuristic (e.g. Gigerenzer & Goldstein 2011).

decide which of three government policies to recommend adopting, for instance, you should systematically go through each option and attempt to find a single good reason for preferring one policy to another. Once that reason has been found, you eliminate the losing policy from consideration, even if there might be other reasons that speak in favor of the losing policy. The TTB heuristic recommends comparing each policy in a pairwise comparison, looking for a single reason to prefer policy A to policy B, then policy A to policy C, and so on. Once a single salient reason is found, the heuristic recommends terminating reasoning, choosing the policy that wins out. In our example, this will involve two rounds: perhaps policy A will improve the lives of the worst off better than will policy B, and perhaps policy C will be more cost-efficient than policy A. Using TTB, we come to prefer policy C, even though there might be other (slightly less salient) dimensions on which other policies are to be preferred.

TTB reasoning is surprisingly accurate. The heuristic outperforms six competitor methods for predicting the features of various recognized and unrecognized cities, no matter how many of the cities the subjects recognize (Gigerenzer & Goldstein 1999). It also makes more successful medical diagnoses in both simulated and experimental contexts when compared to information-rich reasoning rules (Hoffrage et al. 2005). There is a plausible story for why this is so. Using more informationally sparse modes of reasoning prevents limited reasoners from becoming overwhelmed or confused by (potentially relevant) information. This is not just practically a good thing. It is also epistemically desirable, since it improves the accuracy of our deliberations. There have also been good reasons, however, for expanding and modifying the model since it was first introduced. For one thing, it often doesn't seem that we are consciously or deliberately reasoning in this way (Graefe & Armstrong 2012). It is also not clear that all subjects use TTB in all contexts (see especially Bröder 2000). Luckily, several more recent models have given a broader scope to the basic idea of one-reason reasoning.

Gabaix (2014) presents a particularly plausible version of the heuristic, called the *sparse max model* of reasoning.<sup>6</sup> In this model, a massive space of possible considerations in favor of different conclusions is paired down by a relevance condition that isolates only a few important aspects of a reasoning problem. The model then crafts a “simplified model of the world” (Lieder & Griffiths 2020, pp. 7) on the basis of this small subset of relevant variables. The reasoning problem is solved within this simplified picture. Instead of one-reason reasoning, this is *few-reason* reasoning: we might compare policies A, B, and C on the basis of four relevant variables, for instance, and only make our decisions based on those. But this model proceeds in the spirit of the one-reason heuristic, pairing down information to avoid confusing the agent and maximizing accuracy through “rational inattention” (Simms 2006). Its results are impressive: the model predicts the actual reasoning of limited reasoners better than any model that tries to instantiate expected utility calculations (Gabaix et al. 2006; Gabaix 2014). An extended, sequential version of the model can explain many behavioral economic results (Gabaix 2017).

The sparse max model can reproduce the successes of TTB reasoning while also allowing us to make slightly more informed decisions. As long as we focus on only those variables that are relevant without overwhelming ourselves with too much detail, utilizing a few variables (rather than merely one) in a reasoning problem should be beneficial. Recent work in cognitive science has coalesced around this idea, which Lieder & Griffiths (2020) call “resource rational” heuristical reasoning. Human beings tend to make optimal use of the limited resources they have, concluding reasoning before they start to get diminishing returns from more information. Resource rational reasoning has been found to be more accurate than informationally-intensive (but non-ideally implemented) “ideal” reasoning strategies in areas as widespread as perceptual

---

<sup>6</sup> There have also been formal epistemic models developed that extend the TTB heuristic; see Arló-Costa & Pederson (2013).

decision-making (Bogacz et al. 2006), betting behavior (Lieder et al. 2017), and action planning (Callaway et al. 2018). I think the best case for the accuracy of heuristic reasoning lies with resource rational heuristics such as the sparse max model. Given how widespread their use is in human reasoning, arguing for the rationality of these models is a good way of arguing for the rationality of human heuristic reasoning more generally.

#### 4. Accuracy and practicality

Many other heuristic rules of reasoning are often deployed by human beings,<sup>7</sup> but focusing on resource rational heuristics will be sufficient to make my point. Each strategy involves *off-loading* a certain amount of informational processing onto the environment. The use of salient single (sets of) reasons in resource rational reasoning depends on a connection between the reasoning problem and the salience of certain reasons. If we were to remove subjects from informational environments in which these correlations hold, their accuracy would suffer as a result. We should therefore not think that heuristic reasoning is *modally robust* in the same way that ideal reasoning is. This is one of the ways heuristic reasoning is not rationally optimal. This is consistent with it occupying a local rational maximum, however.

So far, I have pointed to a number of situations where utilizing heuristic reasoning rules maximizes accuracy over non-ideal attempts to implement ideal reasoning rules. How should we defend premise (2) of the argument, connecting accuracy and rationality? For some, this step is trivial, since accuracy is thought to be the only (or the trumping) epistemic value (e.g. Pettigrew 2016). But our defense of the rationality of heuristic reasoning can be broader than this. Suppose one thought that practical reasons partially determine what it is epistemically rational for an agent to believe (Rinard 2017). It seems to me that practical reasons will either

---

<sup>7</sup> See Gigerenzer & Todd (1999), for instance.



bolster, or at least will not go against, accuracy reasons in the cases I have in mind. The practical ease of using heuristics will itself contribute to their rationality in some cases. It is much easier for us to consider four salient reasons for preferring a policy than all possible reasons that might exist. In other cases, the practical ease of the two methods will be on a par: we often decide between (relatively practical) non-ideal implementations of ideal reasoning rules and (relatively practical) heuristic reasoning rules. At the very least, practical reasons should not trump accuracy considerations in these cases.

This allows us to see what I mean when I say that an ideal rational advisor would tell us, in many cases, to reason with heuristics. An ideal advisor should not just tell you what she would do in a particular circumstance. Instead, she should consider what she knows about you, your circumstances, and what you can accomplish utilizing different reasoning strategies. From this position, she can see that you will do better (form more accurate beliefs) if you utilize simpler heuristic rules of reasoning than if you overburden yourself trying to implement information-intensive reasoning rules. It is, in other words, a rational good to avoid clutter and confusion in your reasoning (Friedman 2018). Given these facts, a rational advisor who wanted you to draw the most accurate conclusions you could would tell you, in cases like the ones we have been discussing, to reason with heuristics.

## **5. Philosophical upshots**

The use of heuristic reasoning rules represents a local rational maximum for human reasoners in many cases. I think this is a genuine rational achievement. But one might think this is damning with faint praise. Doing as well as you can relative to some norm does not entail that you have accomplished much of anything. Due to my cognitive limitations, I might not be able to tell I am giving you gasoline to drink rather than gin. I might have done as well as I

could trying to serve you a safe drink. I might thereby have an excuse for my behavior. But there is no sense in which I have achieved anything. I do not think heuristic reasoning is like this, however. In addition to maximizing our rationality relative to our limitations, I have stressed that heuristic reasoning also does a *good job* of getting us to accurate conclusions in the environments in which they function well. We are able to solve a whole host of reasoning problems in ways that are more than sufficient for our purposes, even if they might be made more accurate by ideal beings. The optimality results for resource rational heuristics, in particular, show the genuine success of heuristic reasoning. In addition to being *subjectively* rational to use, heuristic reasoning reaches a relatively high *objective* threshold of accuracy, and it is in this sense that it represents a rational achievement. Of course, they could be more accurate. Heuristic reasoning rules are not rationally perfect. But saying this should not lead us to belittle the rational achievement of their use.

What arises is, I think, a very attractive picture of the nature and evaluation of limited thinking, one that meshes well with traditional concerns about accuracy and belief formation in epistemology. Ideal norms of reasoning set the bar by which we can understand the rationally optimal response to a reasoning problem. In the cases we have been discussing, this roughly translates to the most accurate reasoning rule, though there are practical considerations to discuss as well. We then evaluate limited reasoners on the basis of how accurate their reasoning is, relative to this ideal.<sup>8</sup> The empirical literature shows that heuristics, particularly resource rational heuristics like the sparse max model, outrank other rules, including our attempts to use the methods of the ideal reasoner. From this, I conclude that heuristic reasoning often represents a local rational maximum for limited reasoners, getting them closer to the rational ideal than other tools we have in our cognitive toolkit.

---

<sup>8</sup> This framework is worked out in detail for credence and Bayesian rationality in Staffel (2020).

For some, this picture must seem much too conservative. Many have interpreted heuristic reasoning results as calling for a *norm revision* away from (epistemic) rationality norms and towards some other norms of evaluation. They argue “rationality is a tool for helping organisms to reach their real-world goals, not necessarily to conform to... epistemic norms” (Chase et al. 1998, pp. 207), and that “the laws of logic and probability theory are thus insufficient to give a definition of rationality relevant to any real intelligent agent, because the cost of computation has to be taken into account” (Lieder et al. 2017, pp. 323). On this way of thinking about limited reasoners, traditional epistemic norms are to blame for our confusion about how to make sense of heuristics. Only once we jettison these outdated and unrealistic norms of rational evaluation will we understand the proper function of reasoning.

This view has proponents within the virtue epistemology tradition as well. Fairweather & Montemayor (2017) argue that a conception of ecological rationality tied to the epistemic virtues can help save reliabilism about justification from various situationist challenges in psychology. Some go further: Morton (2010, 2012) uses facts about heuristic reasoning to argue for a complete *norm elimination* for norms of rationality. Morton thinks that “there is no coherent concept of rationality” (Morton 2012, pp. 3), and that we should instead talk about specific intellectual virtues that can be realized by limited agents. Morton bases this argument on the idea that heuristic and other kinds of limited reasoning have no real unity to their method, only working in certain specific and limited contexts.<sup>9</sup> Because there is no good general story to tell about why certain heuristics work in some contexts but not others, norms of rationality make for bad advice for limited agents: “the optimum patterns for agents of some degree of cognitive

---

<sup>9</sup> Alfano & Skorburg (2018), for instance, report that the recognition heuristic tends to work best when comparing two cities in the same country (e.g. New York City vs. Dallas), but not when comparing cities internationally (e.g. Dallas vs. Stockholm). Once American participants had to start considering countries in parts of the world they were less familiar with, their performance (predictably) declined rapidly.

boundedness could only be part of the content of expectations and awareness of agents whose cognitive powers exceed that bound” (Morton 2010, pp. 16).<sup>10</sup>

One could respond to the particularities of each of these arguments. For instance, one might argue that resource rational analysis is exactly the delineation of proper reasoning under bounds that Morton claims we cannot find. But I instead want to make two broader points. First, it would be a mistake to get carried away with the rational appeal of heuristic reasoning. At many points in my argument, I have relied on certain *central cases* where heuristic reasoning performs well. These cases have an important normalcy component built into them. Heuristic reasoning works well only in informational environments similar to the ones we tend to encounter in everyday life. When we are presented with reasoning problems that require us to think in ways not immediately encoded in the normal environment, the accuracy of heuristic reasoning tends to be significantly reduced.<sup>11</sup> Ideal rules of reasoning, in contrast, have a modally robust truth-tracking profile. Competently following the rules of deductive implication will invariably leave the subject with true conclusion-states, given true premise-states. There are two aspects of ideal reasoning that outperform heuristic reasoning, when properly applied: ideal reasoning rules are more likely to be accurate within a particular reasoning problem, and across different possible reasoning environments.

---

<sup>10</sup> Sometimes the debate is had in terms of *approximation* (as in Staffel 2020): does heuristic reasoning approximate good reasoning? I think it does. Virtue epistemologists are skeptical; Morton (2012) goes so far as to label this kind of thinking the “approximation fallacy” (pp. 7; see also Fairweather & Montemayor 2014). But there are two things we could mean by approximation. Heuristic reasoning approaches ideal reasoning in *outcome*: heuristic reasoning is approximately accurate for reasoners like us. But it does not approximate ideal reasoning in *method*, since it often looks nothing like deduction or induction. The disconnect between method and outcome for non-ideal agents is precisely the area in which resource rational heuristics operate well.

<sup>11</sup> The classic example is the case of Linda the bank teller (Kahneman & Tversky 1972). Many of the irrationality results in the biases and heuristics literature are generated in just this way, presenting problems that remove the subject from their normal informational environment, and exposing how defective their responses are outside of that context.

Second, I think the “traditional” picture of epistemic rationality is so central to our (limited) lives that calls for radical norm revision would require truly overwhelming arguments in their favor. Whether or not a person’s belief is rationally justified seems to be crucially connected with our *moral* practices of praising and blaming others (cf. Wallace 1994). Whether or not you have a good excuse for your behavior turns on whether you justifiably believed you were giving me gin instead of gasoline. To jettison norms of rational evaluation, or to radically alter them, would be to jettison or radically alter our norms of moral evaluation as well. This is a revisionary proposal indeed. We should also want our theory of rationality to tell us both what is rational *and* what is irrational in our thinking. A theory that says there is no general answer to the analysis of good and bad reasoning, or that says humans are maximally optimal in their thinking relative to some bounded norm, cannot adequately make sense of the patterns of both good *and bad* reasoning, and what makes them good or bad. In contrast, the view I defend here has a ready answer to these questions. Our reasoning reaches a local rational maximum, but not a global rational maximum, due to the limitations of our thinking. My view captures both what goes right and what goes wrong in our thinking in a way the norm revision and elimination proposals do not.

We might view the proposal I am putting forward here as an application of the general theory of second best to non-ideal epistemology (Lipsey & Lancaster 1956). In Wiens (2020)’s formulation, the general theory of second best states that, although a best option might have features  $F_1$ - $F_n$  that make that option the best option, a second-best option will lack most (or, in some cases, all) of  $F_1$ - $F_n$ . The rational best practices for reasoning might look nothing like rules that it would be good for us limited, second-best reasoners to adopt. Even if an ideal reasoner might systematically consider all relevant information and perform all valid deductive and inductive inferences on such information, these facts tell us little about what we should do with

our limited reasoning capacities. Heuristic rules of reasoning look radically different than their ideal counterparts. On the general theory of second best, this is to be expected.

In summary, and to return to our motivating question, just how rationally admirable is typical human reasoning? I have argued: about as rational as it can be, given our cognitive limitations. Many cognitive scientists have been attracted to this idea (Gershman et al. 2015; Lieder & Griffiths 2020). The notion of a local rational maximum gives a way of expressing the significant cognitive achievement of human heuristic reasoning without giving up on norms of evaluation that are crucial to our epistemic (and moral) lives. We can paint a realistic picture of the accuracy of human heuristic reasoning, while still being able to note its limitations and inaccuracies. We are neither epistemic oracles, nor are we hopelessly irrational in our thinking. The only surprise, perhaps, is that “traditional” accuracy-oriented epistemology has ample tools to make sense of this.

## References

- Alfano, M. & Skorburg, J. (2018). Extended knowledge, the recognition heuristic, and epistemic injustice. In Pritchard, Kallestrup, Palermos & Carter (eds.), *Extended Knowledge*. Oxford University Press. pp. 239-256.
- Andersson, P., & Rakow, T. (2007). Now you see it now you don't: The effectiveness of the recognition heuristic for selecting stock. *Judgment and Decision Making*, (2), 29-39.
- Arló-Costa, Horacio & Pedersen, Arthur Paul (2013). Fast and frugal heuristics: rationality and the limits of naturalism. *Synthese* 190 (5):831-850.
- Armstrong, S., & Mindermann, S. (2018). Occam's razor is insufficient to infer the preferences of irrational agents. In *Advances in Neural Information Processing Systems* (pp. 5598-5609).
- Bishop, M. A. (2006). Fast and frugal heuristics. *Philosophy Compass* 1 (2):201–223.
- Bogacz, R., Brown, E., Moehlis, J., Holmes, P. & Cohen, J. (2006) The physics of optimal decision making: A formal analysis of models of performance in two-alternative forced-choice tasks. *Psychological Review* 113(4):700–65.
- Boghossian, P. (2014). What is inference? *Philosophical Studies* 169 (1):1-18.
- Borges, B., Goldstein, D. G., Ortmann, A., & Gigerenzer, G. (1999). Can ignorance beat the stock market. *Simple heuristics that make us smart*, 59-72.
- Boyd, M. (2001). On ignorance, intuition, and investing: A bear market test of the recognition heuristic. *The Journal of Psychology and Financial Markets*, 2(3), 150-156.
- Bröder, A. (2000). Assessing the empirical validity of the "Take-the-best" heuristic as a model of human probabilistic inference. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 26(5), 1332.

- Callaway, F., Lieder, F., Das, P., Gul, S., Krueger, P. M. & Griffiths, T. L. (2018) A resource-rational analysis of human planning. In: *Proceedings from 40th Annual Conference of the Cognitive Science Society*. Cognitive Science Society.
- Chase, V. M., Hertwig, R., & Gigerenzer, G. (1998). Visions of rationality. *Trends in cognitive sciences*, 2(6), 206-214.
- Dogramaci, S. (2015). Communist Conventions for Deductive Reasoning. *Noûs* 49 (4):776-799.
- Earman, J. (1992). *Bayes or Bust? A Critical Examination of Bayesian Confirmation Theory*. MIT Press.
- Egan, A. (2008). Seeing and believing: perception, belief formation and the divided mind. *Philosophical Studies* 140 (1):47 - 63.
- Fairweather, A., & Montemayor, C. (2014). Inferential abilities and common epistemic goods. In *Virtue Epistemology Naturalized* (pp. 123-139). Springer.
- (2017). *Knowledge, dexterity, and attention: A theory of epistemic agency*. Cambridge University Press.
- Friedman, J. (2018). Junk Beliefs and Interest-Driven Epistemology. *Philosophy and Phenomenological Research* 97 (3):568-583.
- Gabaix, X. (2014). A sparsity-based model of bounded rationality. *The Quarterly Journal of Economics*, 129(4), 1661-1710.
- (2017) Behavioral inattention. *NBER Working Paper No. 24096*. National Bureau of Economic Research.
- Gabaix, X., Laibson, D., Moloche, G. & Weinberg, S. (2006) Costly information acquisition: Experimental analysis of a boundedly rational model. *American Economic Review* 96(4):1043–68.
- Gaissmaier, W., & Marewski, J. N. (2011). Forecasting elections with mere recognition from small, lousy samples: A comparison of collective recognition, wisdom of crowds, and representative polls. *Judgment and Decision Making*, 6(1), 73-88.
- Gershman, S. J., Horvitz, E. J., & Tenenbaum, J. B. (2015). Computational rationality: A converging paradigm for intelligence in brains, minds, and machines. *Science*, 349(6245), 273-278.
- Gigerenzer, G. & Goldstein, D. G. (2011). The recognition heuristic: A decade of research. *Judgment and Decision Making*, 6(1), 100-121.
- Gigerenzer, G., Hoffrage, U., & Kleinbölting, H. (1991). Probabilistic mental models: a Brunswikian theory of confidence. *Psychological review*, 98(4), 506.
- Gigerenzer, G., & Selten, R. (Eds.). (2002). *Bounded rationality: The adaptive toolbox*. MIT press.
- Graefe, A., & Armstrong, J. S. (2012). Predicting elections from the most important issue: A test of the take-the-best heuristic. *Journal of Behavioral Decision Making*, 25(1), 41-48.
- Hedden, B. (2019). Hindsight bias is not a bias. *Analysis*, 79(1), 43-52.
- Hoffrage, U., Garcia-Retamero, R., & Czienskowski, U. (2005). The robustness of the Take The Best Configural heuristic in linearly and nonlinearly separable environments. In *Proceedings of the 27th annual conference of the cognitive science society* (pp. 971-976). Erlbaum Mahwah, NJ.
- Kahneman, D., & Frederick, S. (2002). Representativeness revisited: Attribute substitution in intuitive judgment. *Heuristics and biases: The psychology of intuitive judgment*, 49-81.
- Kahneman, D., & Tversky, A. (1972). Subjective probability: A judgment of representativeness. *Cognitive psychology*, 3(3), 430-454.
- Kelly, T. (2008). Disagreement, dogmatism, and belief polarization. *The Journal of Philosophy*, 105(10), 611-633.
- Lieder, F., & Griffiths, T. L. (2017). Strategy selection as rational metareasoning. *Psychological*

- Review*, 124(6), 762.
- (2020). Resource-rational analysis: understanding human cognition as the optimal use of limited computational resources. *Behavioral and Brain Sciences*, 43, E1.
- Lieder, F., Krueger, P. M. & Griffiths, T. L. (2017) An automatic method for discovering rational heuristics for risky choice. In: *Proceedings from the 39th annual conference of the Cognitive Science Society* (London, UK), pp. 2567–72. Cognitive Science Society.
- Lipsey, R. G., & Lancaster, K. (1956). The general theory of second best. *The review of economic studies*, 24(1), 11-32.
- Lipton, P. (2004). *Inference to the best explanation*. Taylor & Francis.
- Morton, A. (2010). Human bounds: rationality for our species. *Synthese*, 176(1), 5-21.
- (2012). *Bounded thinking: Intellectual virtues for limited agents*. OUP Oxford.
- Nisbett, R. E., & Borgida, E. (1975). Attribution and the psychology of prediction. *Journal of Personality and Social Psychology*, 32(5), 932.
- Ortmann, A., Gigerenzer, G., Borges, B., & Goldstein, D. G. (2008). The recognition heuristic: a fast and frugal way to investment choice? *Handbook of experimental economics results*, 1, 993-1003.
- Pettigrew, R. (2016). *Accuracy and the Laws of Credence*. Oxford University Press UK.
- Railton, P. (1986). Moral realism. *The Philosophical Review*, 95(2), 163-207.
- Richard, M. (2019). Is reasoning a form of agency? In Jackson, B. B. & Jackson, M. B. (eds.), *Reasoning: New Essays on Theoretical and Practical Thinking*. Oxford University Press. pp. 91-100.
- Rinard, S. (2017). No Exception for Belief. *Philosophy and Phenomenological Research* 94 (1):121-143.
- Ross, Lewis (forthcoming). Rehabilitating Statistical Evidence. *Philosophy and Phenomenological Research*.
- Rysiew, P. (2008). Rationality disputes – psychology and epistemology. *Philosophy Compass* 3 (6):1153-1176.
- Scheibehenne, B., & Bröder, A. (2007). Predicting Wimbledon 2005 tennis results by mere player name recognition. *International Journal of Forecasting*, 23(3), 415-426.
- Siegel, S. (2019). Inference Without Reckoning. In Jackson, B. B. & Jackson, M. B. (eds.), *Reasoning: New Essays on Theoretical and Practical Thinking*. Oxford University Press. pp. 15-31.
- Sims, C. A. (2006) Rational inattention: Beyond the linear-quadratic case. *American Economic Review* 96(2):158–63.
- Simon, H. A. (1997). *Models of bounded rationality: Empirically grounded economic reason (Vol. 3)*. MIT press.
- Smith, M. A. (1995). *The moral problem*. Blackwell: Wiley.
- Staffel, J. (2020). *Unsettled thoughts: A theory of degrees of rationality*. Oxford University Press.
- Wallace, R. J. (1994). *Responsibility and the moral sentiments*. Harvard University Press.
- Wiens, David (2020). The General Theory of Second Best Is More General Than You Think. *Philosophers' Imprint* 20 (5):1-26.