


This article was downloaded by: 10.2.97.136

On: 30 Mar 2023

Access details: *subscription number*

Publisher: *Routledge*

Informa Ltd Registered in England and Wales Registered Number: 1072954 Registered office: 5 Howick Place, London SW1P 1WG, UK



## **Routledge Handbook of Bounded Rationality**

Riccardo Viale

### **Rationality without optimality**

Publication details

<https://test.routledgehandbooks.com/doi/10.4324/9781315658353-20>

Henry Brighton

**Published online on: 29 Oct 2020**

**How to cite :-** Henry Brighton. 29 Oct 2020, *Rationality without optimality from:* Routledge Handbook of Bounded Rationality Routledge

Accessed on: 30 Mar 2023

<https://test.routledgehandbooks.com/doi/10.4324/9781315658353-20>

**PLEASE SCROLL DOWN FOR DOCUMENT**

Full terms and conditions of use: <https://test.routledgehandbooks.com/legal-notices/terms>

This Document PDF may be used for research, teaching and private study purposes. Any substantial or systematic reproductions, re-distribution, re-selling, loan or sub-licensing, systematic supply or distribution in any form to anyone is expressly forbidden.

The publisher does not give any warranty express or implied or make any representation that the contents will be complete or accurate or up to date. The publisher shall not be liable for an loss, actions, claims, proceedings, demand or costs or damages whatsoever or howsoever caused arising directly or indirectly in connection with or arising out of the use of this material.

# RATIONALITY WITHOUT OPTIMALITY

## Bounded and ecological rationality from a Marrian perspective

*Henry Brighton*

### **Does rationality imply optimality?**

The idea that rationality implies optimality is so widely assumed as to seem barely worth discussing. Optimal Bayesian decision makers in the cognitive sciences, optimal foragers in biology, and Bayesian maximizers of expected utility in economics are different faces of the same interdisciplinary orthodoxy. Does bounded rationality offer something different? I will argue that ecological rationality, a development of Herbert Simon's (1957) notion of bounded rationality, provides a fully-fledged alternative to the view that rational decisions under environmental uncertainty are optimal probabilistic decisions (Gigerenzer and Selten, 2001; Gigerenzer and Brighton, 2009; Brighton, 2019). For those familiar with the concept of ecological rationality, the study of simple heuristics will likely spring to mind. The study of simple heuristics examines how minimalist decision making algorithms that ignore information can, depending on the structure of the environment, outperform familiar and more sophisticated algorithms that weigh and integrate information, or explicitly optimize a rationally justified criterion (e.g., Gigerenzer et al., 1999; Brighton, 2006; Brighton and Gigerenzer, 2007; Gigerenzer and Brighton, 2009). Rather than focus on these findings, I will scrutinize an argument often made by critics of ecological rationality: Insights arising from the study of ecological rationality are, after the fact, formally reducible to instances of optimal Bayesian inference and, furthermore, require principles of Bayesian inference to fully explain them. The claim is that ecological rationality neither undermines the use of optimality nor challenges orthodox rationality.

I will use Marr's (1982) levels of analysis to undermine this claim, and instead argue that ecological rationality does not imply optimality. Marr's levels organize and relate functional, algorithmic, and implementational theories of information processing systems. In Marr's terms, simple heuristics studied under the rubric of ecological rationality address algorithmic level concerns, where the question is *how* organisms make decisions, not the functional question of *what* constitutes a rational decision and *why*. For some theorists, this elementary distinction between function and mechanism clarifies why ecological rationality is subservient to orthodox rationality (e.g., Chater et al., 2003; Gintis, 2012). My counterargument will be developed in three stages: (1) ecological rationality involves more than a conjecture about simple heuristics, it also rests on a statistical formulation of the problem of inductive inference under unquantifiable

uncertainty that neither assumes nor implies the existence of an optimal solution (Breiman, 2001; Brighton and Gigerenzer, 2007, 2012; Brighton, 2019); (2) this statistical perspective exposes limitations in the ability of optimal functioning to characterize functional responses to ill-posed problems, and undermines the idea that function and mechanism can be studied independently; and (3) ecological rationality is incompatible with and exposes the limitations of Marr's (1982) levels of analysis, but finds support in Marr's (1977) largely overlooked views on the broader functioning of biological information processing systems.

### Marr and Poggio's three levels of analysis

David Marr's book *Vision* (1982) set out an innovative approach to understanding vision as an information processing problem. As well as detailing computational models of visual processing, Marr argued that a deeper understanding of information processing systems in general, not just the visual system, requires three levels of analysis. Originally developed with Tomaso Poggio (Marr and Poggio, 1976; Poggio, 2012), these three levels map onto fundamental distinctions in the study of rationality that will recur throughout this discussion. At the highest level, termed the computational level, the goal is to identify an appropriate formal system, a calculus, used to define the problem and derive its solution. This abstract rendering of the problem details what problem the system attempts to solve and why. For example, a pocket calculator solves the problem of performing arithmetic calculations and, so the story goes, human decision makers solve the problem of maximizing expected utility when making decisions under uncertainty. Calculi used to develop computational level theories include Bayesian statistics (Griffiths and Tenenbaum, 2006), deductive logic (Baggio et al., 2014), and the mathematics of signal processing (Marr and Hildreth, 1980).

Moving down to Marr's second level, the algorithmic level, the goal is to specify the algorithms and data structures needed to compute the solution. For example, if part of the computational level theory of a pocket calculator is to calculate factorials, then at the algorithmic level this calculation, for reasons of efficiency, might be implemented using Stirling's approximation,  $n!$  approximates sign square root  $n! \approx \sqrt{2\pi n} \left(\frac{n}{e}\right)^n$  which incurs a loss in precision but requires fewer operations to compute. Similarly, it has been proposed that simple heuristics might be used to approximate a full Bayesian computation specified at the computational level (Chater et al., 2006; Oaksford and Chater, 2009).

Marr's third and final level of analysis is the implementation level, which considers the constraints of physically instantiating the proposed algorithms and data structures in, say, biological machinery or a conventional digital computer. Combined, Marr's levels consider what problem the system attempts to solve, what the solution to this problem is, how this solution is computed, and how this computation is realized in a physical device. Notice that each Marrian level corresponds to a different level of analysis rather than a different level of organization. Levels of organization in nature typically reflect part-whole relationships such as those that separate proteins from cells and individuals from groups, whereas levels of analysis separate different forms of understanding (Bechtel and Shagrir, 2015; McClamrock, 1991). For Marr, studying an information processing system without a formal understanding of the problem being addressed should be avoided, and he advocated a top-down approach to developing theories of biological information processing systems. This methodology, though, does not exclude a theorist from adopting alternative priorities when working within the Marrian framework. Constraints from implementation and algorithmic levels can, depending on the question, also drive theory development.

Why place Marr's levels at the center of this discussion? To begin with, Marr's levels of analysis are seen as "a core tenet of cognitive science" (Peebles and Cooper, 2015, p. 188), defining the de facto categories used to orientate cognitive models in relation to rationality claims. Bayesian optimality modeling in the cognitive sciences, for instance, seeks computational level theories that formulate the problems faced by the cognitive system and what constitutes a rational solution (Tenenbaum and Griffiths, 2001; Chater et al., 2006; Griffiths et al., 2012). Rarely are these models interpreted as hypotheses about *how* the cognitive system computes. Cognitive process models such as simple heuristics, on the other hand, are categorized as algorithmic level theories. Both kinds of theory can be interpreted as models of human behavior, but only the second kind involves an explicit conjecture about cognitive processing. From this Marrian perspective, bounded rationality is typically seen as an attempt to inform computational level theory development by importing constraints arising from the algorithmic and implementation levels. For example, Anderson's (1990) rational analysis of cognition and the more recent study of resource rationality (Griffiths et al., 2015) use these constraints to revise the problem specification, and consequently, what problem to which the computational level theory specifies an optimal response. Now, on the one hand, if we assume that simple heuristics are to Bayesian rationality what Stirling's approximation is to arithmetic, then as critics as have argued, studying ecological rationality without appeal to Bayesian statistics is like trying to understand how birds fly without a theory of aerodynamics (Chater and Oaksford, 1999) or understanding how people play billiards without a theory of the physics of motion (Gintis, 2012). On the other hand, both Simon (1991) and advocates of ecological rationality have rejected the proposal that the study of bounded and ecological rationality reduces to studying optimality (e.g., Gigerenzer and Selten, 2001; Brighton, 2019). What is the basis of this disagreement, and can Marr's levels be used to clarify and resolve the issue?

### The statistical foundations of ecological rationality

Ecological rationality is, and always has been, more than the study of simple heuristics. The goal is to understand the adaptive fit between organisms and the structure of the environment by examining three interacting components: (1) algorithmic models of how organisms make inductive inferences, with a particular focus on simple heuristics; (2) the properties of natural environments whose probabilistic structure is either uncertain or unknown; and (3) a formulation of the problem of statistical inference that defines and quantifies the meaning of an adaptive fit. To focus this discussion, I will sidestep experimental research examining the use of simple heuristics in humans and other animals, and instead focus on the overarching hypothesis: Simple heuristics that ignore information are a vital part of how organisms successfully cope with the uncertainty of the natural world.

Results supporting this hypothesis are termed less-is-more effects, and they detail how minimalist processing strategies improve the accuracy of decisions relative to more complex and supposedly sophisticated strategies commonly assumed in the cognitive sciences and beyond. For example, when deciding which of two objects in the environment scores higher on some criterion of interest, such as which of two food items has a higher calorific content, a user of the recognition heuristic will choose the recognized object if the other object is unrecognized. A partly ignorant decision maker using this heuristic can make more accurate inferences than a more informed decision maker who weighs and integrates all available information (Goldstein and Gigerenzer, 2002). Similarly, a user of the simple heuristic take-the-best ignores correlations between cues and uses a single cue to make an inference, yet can make more accurate inferences than a decision maker who attempts to model these correlations and weighs and integrates all

available cues (e.g., Gigerenzer and Goldstein, 1996; Czerlinski et al., 1999; Brighton, 2006; Brighton and Gigerenzer, 2007; Gigerenzer and Brighton, 2009; Şimşek and Buckmann, 2015). In both cases, a simple heuristic (component 1 of the interaction above) confers a function by exploiting properties of natural environments (component 2 of the interaction above). Of absolutely crucial importance to the coming discussion is component 3 of the interaction described above: The formulation of the statistical problem facing the organism. This third component also requires theoretical commitments to be made, and it is these commitments that establish, most directly, ecological rationality as a substantive alternative to orthodox rationality.

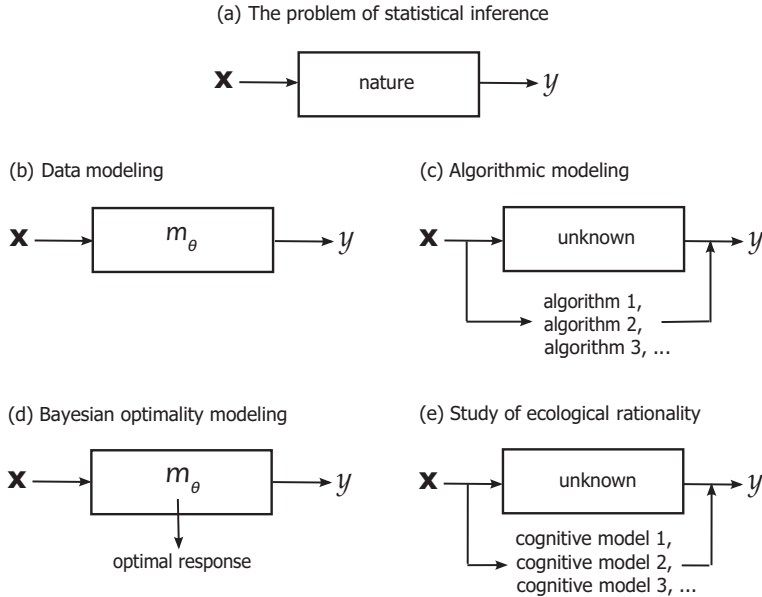
### ***The role of statistics in constructing the rationality problem***

Simple heuristics like take-the-best are supervised learning algorithms. They tackle the problem of making inductive inferences from observations of the environment where the goal is to distinguish systematic patterns that govern the observations from any noisy and accidental patterns. What is a rational solution to this problem? If we possess a complete probabilistic understanding of the environmental processes that generate the observations, then the answer is simple: The optimal Bayesian response relative to a prior and a hypothesis space expressing this probabilistic knowledge. But what if we lack what Binmore (2009) termed this “metaphysical hotline to the truth?” Or, what if the decision maker is not in what Savage (1954), the inventor of Bayesian decision theory, termed a “small world” where identifying a mutually exclusive and exhaustive set of future states of the world, consequences, and the actions that map between them is a realistic possibility? Put simply, how should an optimal response be formulated when partial or complete ignorance precludes probabilistic quantification of the uncertainties we face?

Orthodox responses to these questions cite uninformed priors, second-order probabilities, imprecise probabilities, or some other technique that assumes and proposes how we should quantify the varying degrees of uncertainty we face. Another response is to devise alternatives to Bayesian decision theory but retain the goal of optimality (e.g., Gilboa, 1987; Gilboa and Schmeidler, 1989). A third option is to accept that “it is sometimes *more* rational to admit that one does not have sufficient information for probabilistic beliefs than to pretend that one does” (Gilboa et al., 2012, p. 28). This third option represents the fork in the road where orthodox rationality and ecological rationality part ways. The orthodox view proceeds under the assumption that we can, by hook or by crook, always quantify environmental uncertainty. The study of ecological rationality sets out to examine the implications of unquantifiable uncertainty. The statistician Leo Breiman (2001) characterized two cultures of thought underlying these issues, and they map directly onto the divergent practices of orthodox rationality and ecological rationality. Figure 17.1 depicts these relationships, which are worth considering in detail because they prove critical in constructing the formal nature of the problem.

### ***Orthodox rationality and the statistical culture of data modeling***

If we start with observations, each relating a set of independent variables to a dependent variable, then the environment can be seen as a black box containing data-generating machinery that determines the joint distribution over the inputs to the black box (independent variables) and outputs (the dependent variable) shown in Figure 17.1(a). Much of traditional statistical inquiry, and this includes the practice of formulating optimal probabilistic responses to environmental uncertainty, requires that we make a conjecture about the contents of this black box, depicted in Figure 17.1(b). We might work from the assumption that dependent and independent variables are linearly related, for example. Taking a Bayesian perspective, we might



*Figure 17.1* How should the problem of statistical inference be formulated? We start with observations of an unknown functional relationship between input vectors  $x$  (the independent variables) and output  $y$  (the dependent variable) determined by some aspect of nature, shown in (a). Following Breiman (2001), the statistical culture of data modeling views the conjecture of a stochastic model  $y = m_{\theta}(x)$  of nature's black box as an essential step in statistical inference, shown in (b). Rather than attempt to model the contents of nature's black box, algorithmic modeling is an incremental search for learning algorithms that can, to varying degrees, accurately predict the input-output relationship, shown in (c). Bayesian optimality modeling conducts data modeling in order to define an optimal response (d), while the study of ecological rationality conducts algorithmic modeling and interprets predictive models as potential cognitive models, shown in (e).

Source: Diagrams (a–c) adapted from Breiman (2001).

formulate an hypothesis space, prior distribution, and various parameters that we fit using the available observations. The defining characteristic of this approach is that, at some point, a conjecture is made about the contents of the black box, and it is relative to this conjecture that subsequent claims about rationality and optimality are made. For instance, Anderson's rational analysis of cognition is a methodology for formulating the problems that the environment poses the cognitive system, and using rational solutions to these problems to derive theories of cognitive processing (Anderson, 1990, 1991; Chater and Oaksford, 1999). In both Anderson's methodology and Bayesian optimality modeling, formulating a probabilistic model of the task environment is a key step, illustrated in Figure 17.1(d) (Tenenbaum and Griffiths, 2001; Chater et al., 2006; Griffiths et al., 2012). Breiman (2001) used the term “data modeling” to describe this general approach to statistical inquiry.

### *Ecological rationality and the statistical culture of algorithmic modeling*

An alternative approach to statistical modeling is what Breiman termed algorithmic modeling. When algorithmic modeling, we refrain from making a conjecture about the contents of the black box and instead try to predict its behavior. As shown in Figure 17.1(c), the observations

are used to estimate the predictive accuracy of competing models of inductive inference, which in practice usually means comparing machine learning algorithms using a model selection criterion, such as cross-validation (Stone, 1974), where algorithms are judged on their ability to generalize to unseen cases (Hastie et al., 2001; Bishop, 2006). Crucially, learning algorithms and the probabilistic assumptions they imply tend not to be seen as models of the environment, but rather inductive biases that may explicitly introduce model infidelities in order to reduce a component of prediction error known as variance, described below (such techniques are part of the theory of regularization; see, e.g., Chen and Haykin, 2002). Among the algorithms being considered, one or more algorithm will achieve the greatest reduction in prediction error. Findings like these in no way license an optimality claim. Instead, they provide an indication of the kinds of algorithmic design decisions or statistical assumptions that improve our ability to reduce prediction error, thereby suggesting further algorithms worth evaluating.

This process is fundamentally exploratory, yields a functional understanding of the algorithms being considered, yet in no way invokes the concept of optimality to explain model performance. Specifically, one can seek to improve predictive accuracy relative to existing models without assuming that an optimal model can ever be identified, or, given the complexity and uncertainty of the natural world, meaningfully exists. How, though, is ecological rationality related to the culture of algorithmic modeling? From the outset, key results in the study of ecological rationality have rested almost exclusively on algorithmic modeling studies, where the ability of simple heuristics to reduce prediction error relative to other models has been analyzed in environments for which the optimal response is unknown and arguably unknowable (Czerlinski et al., 1999; Martignon and Laskey, 1999; Brighton, 2006; Gigerenzer and Brighton, 2009). Such studies typically proceed by investigating model performance in new environments, or new heuristics generated by combining a set of basic building blocks that include rules for information search, rules for stopping search, and rules for making decisions (Gigerenzer et al., 1999). Combining these building blocks, which are often inspired by experimental studies of decision making in humans and other animals, is one way of steering the algorithmic modeling process. In short, the study of ecological rationality, schematized in Figure 17.1(e), proceeds by studying algorithmic functioning in environments composed on observations rather than environments rendered as probabilistic models that imply optimal responses. Next I will argue that in theory too, ecological rationality is not in the business of studying optimality.

### ***Bayesian reductionism and the limitations of optimal function***

Despite the preceding points, critics have argued that when a heuristic works, one still needs a rational explanation for why it works, and so the concept of ecological rationality is ultimately subservient to Bayesian rationality. Oaksford and Chater (2009), for example, argue that Bayesian optimality modeling “cannot be replaced by, but seeks to explain, ecological rationality” (p. 110). This view extends to critics of Bayesian optimality modeling, such as Jones and Love (2011), who argue that the two approaches are “highly compatible” because “any inference algorithm implicitly embodies a prior expectation about the environment” (p. 186). Thus, even though ecological rationality does without strong probabilistic conjectures and optimality claims in practice, in theory, these concepts will be required to explain functional success. These arguments, which express a kind of Bayesian reductionism, are worth dissecting in turn.

*Argument 1: Bayesian thinking is required to explain functional success*

When a learning algorithm performs well in a particular environment, incurring lower prediction error than the other algorithms being considered, how can this behavior be explained? Seen through the lens of probabilistic optimality, there is a strong tendency to view this functional success as indicating a close match between the prior distributional assumptions of the algorithm and the probabilistic structure of the environment. An alternative pattern of explanation that has proven critical to understanding the success of simple heuristics is to analyze the prediction error of competing algorithms using the bias/variance decomposition (Gigerenzer and Brighton, 2009). Specifically, we can decompose prediction error into two controllable components, bias and variance, and one uncontrollable component, noise, such that they combine additively as follows (O’Sullivan, 1986; Geman et al., 1992; Hastie et al., 2001; Bishop, 2006):

$$\text{Total error} = (\text{bias})^2 + \text{variance} + \text{noise} \quad (17.1)$$

The bias component reflects the inability of the model to capture what is systematic in the observations. The variance component, in contrast, reflects the sensitivity of the model’s predictions to different observations of the same problem, such as a different sample from the same population. All algorithms face what is termed the bias/variance dilemma because methods for reducing variance tend to increase bias, whereas methods for reducing bias tend to increase variance (Geman et al., 1992). For example, a heuristic may suffer from high bias due to having a single free parameter, but this simplicity can lead to a greater reduction in variance when learning from small samples (Brighton and Gigerenzer, 2015). A competing algorithm may strike a different trade-off between bias and variance, and whether this point proves functional will depend on the sample size and the properties of the environment. Many of the algorithmic tricks used by simple heuristics and other learning algorithms to incur low prediction error are methods for reducing variance. In short, the idea that Bayesian thinking is required to explain the performance of learning algorithms is inaccurate, both in theory and in practice.

*Argument 2: Bayesian explanations of functional success should be preferred*

Even if the principles of Bayesian inference are not required to explain functional success, perhaps, given their elegance and foundations in probability theory, they should be preferred. A counterexample to this intuition is the naïve Bayes classifier, a simple learning algorithm that makes a strong assumption that the features (also termed cues, attributes, or independent variables) are conditionally independent of each other, given the class. Despite this assumption being highly unlikely to be met in practice, the naïve Bayes classifier has a long history of performing surprisingly well, particularly when learning from sparse data. Why is this? Research spanning several decades has repeatedly invoked the concepts of bias and variance to explain why the naïve Bayes classifier performs well relative to other algorithms: Despite introducing a strong bias, the conditional independence assumption is often highly effective at reducing variance (e.g., Domingos and Pazzani, 1997; Friedman, 1997; Hand and Yu, 2001; Ng and Jordan, 2002; Webb et al., 2005). Thus, even when we have a clear probabilistic formulation of a learning algorithm, it can prove more insightful to explain functional success using the bias/variance perspective than a perspective rooted in probabilistic optimality.



*Argument 3: All learning algorithms imply an optimal Bayesian response*

Given any supervised learning algorithm, let us assume that a probabilistic reformulation of this algorithm is always possible, such that the probabilistic conditions under which the algorithm makes Bayes optimal decisions are made explicit. Wouldn't this then prove that ecological rationality reduces to the study of optimal Bayesian inference? No. All this proves is that one aspect of algorithmic functioning is optimal functioning. When algorithmic modeling, learning algorithms can, and typically do, confer a function without being optimal. They also confer a function by virtue of their performance *relative* to other learning algorithms. Indeed, if the optimal response is indeterminable, then relative function is all we have to work with. More generally, the conditions under which an algorithm is optimal and the conditions under which it confers a function outside these optimality conditions cannot be assumed to be directly related. Reductions in prediction arise from a complex interaction between the learning algorithms being considered, the number of observations available, and properties of the environment (e.g., Perlich et al., 2003). And these complexities mean that the optimality conditions of most learning algorithms are simply not known, and attempts to identify them tend to be only partially successful (e.g., Domingos and Pazzani, 1997; Kuncheva, 2006). To summarize, optimality results are

- (1) of limited practical use outside the study of idealized environments for the simple reason that we tend not to know the probabilistic structure of natural environments;
- and (2), limited in their ability to provide a reliable indication of how well an algorithm will perform outside its optimality conditions, and relative to other algorithms.

### **The theory-dependence of Marrian decomposition**

To recap, my argument is that Breiman's (2001) statistical cultures of data modeling and algorithmic modeling differ in how they view and respond to uncertainty, these differences have direct implications for the study of rationality, and crucially, these differences are reflected in the practices of Bayesian optimality modeling and the study of ecological rationality. The fly in the ointment is optimality. Given all this, how can the study of ecological rationality be reconciled with the Marrian perspective? Resolving this issue is of central importance because any incompatibility between ecological rationality and Marr's levels of analysis would leave us in the precarious position of advocating a view on rationality at odds with "the canonical scheme for organizing formal analyses of information processing systems" (Griffiths et al., 2015, p. 217) and "one of the most well-known theoretical constructs of twentieth century cognitive science" (Willems, 2011, p. 1). On the other hand, if ecological rationality proves to be entirely compatible with Marr's levels, then the existence of a computational-level theory would suggest, as is typically assumed, that the whole approach reduces to an optimality claim.

### ***Ecological rationality from a Marrian perspective***

Recall that the objective of a computational level theory is to identify an appropriate formal system, a calculus, used to define the problem and derive a solution. Ultimately the computational level should define the input-output mappings to be computed, perhaps approximately, by an algorithmic level theory. The problem we face now is that the statistical "goal" of algorithmic modeling, and hence ecological rationality, is both exploratory and algorithmic in nature, echoing Selten's (2001) point that "bounded rationality cannot be precisely defined. It is a problem that needs to be explored" (p. 15). In contrast to data modeling, algorithmic

modeling doesn't supply a solution derived using a calculus but instead a relative understanding of the efficacy of competing algorithms. The problem facing the organism remains, only this problem is not well posed in Hadamard's (1902) sense of requiring that a solution exists, is unique, and is stable with respect to small perturbations to the problem definition. Nature doesn't always, or even often, present well-posed problems in this sense, and it is imperative that our conception of rationality should reflect this. Taken in conjunction with Marr's (1977) proposal that a computational level theory should be "expressed independently of the particular way it is computed" (pp. 37–38), the conclusion that ecological rationality is incompatible with the demands of a computational level theory is hard to avoid.

### ***The limits of Marrian decomposition: type-1 and type-2 theories***

This conclusion, that Marr's levels of analysis fall short as a canonical system and imply a limited view of algorithmic functioning, is in fact consistent with Marr's thinking when considered in light of his 1977 article published in *Artificial Intelligence* (Marr, 1977). This article addresses some of the issues discussed above, yet is largely overlooked in contemporary discussions of Marr's work. Here, Marr drew the distinction between biological information processing systems with type-1 theories (those which yield to a computational-level analysis) and those with type-2 theories (those which don't). An archetypal type-1 theory for Marr is the Fourier Transform, a mathematical operation central to edge detection that can be approximated at the algorithmic level in a number of ways (Marr and Hildreth, 1980). Marr's examples of archetypal type-2 theories include protein unfolding and the grammar of natural language. Neither of these problems, Marr speculates, are likely to have useful mathematical abstractions. He warns, for instance, that "an abstract theory of syntax may be an illusion" (p. 42), such that productive study at the computational level is likely to be questionable.

Marr's overarching point is that we should prioritize the study of problems with type-1 theories but not expect that all interesting problems will have them. This suggests three possibilities that need to be considered. First, the theorist seeks a computational level theory, finds one, and then invokes Marr's three levels of analysis to guide further theory development. Second, the theorist seek a computational level theory, fails to find one, and then pursues an alternative form of functionalist inquiry. Third, the theorist fails to adopt a functionalist perspective in the first place, overlooks a potentially insightful computational theory, and misses the opportunity to arrive at the kind of understanding that Marr advocates. Marr (1977) reserves his criticism for this third scenario, which he saw as a "waste of time" because it "amounts to studying a mechanism not a problem" (p. 46). Advocates of ecological rationality find themselves in the second scenario, and set out to examine what implications these problems have for the study of rationality. In summary, the idea that theories of cognitive processing that make functionalist claims, including claims about rationality, must be reconcilable with Marr's three levels of analysis rests on a narrow conception of the kinds of problem that the cognitive system is likely to face. This category of problem was one that Marr undoubtedly saw as a priority, but he also recognized that this category offers only a partial view of the kinds of problem tackled by biological information processing systems.

### **Rationality without optimality**

It is commonly assumed, outside the cognitive sciences as well as within them, that "the only viable formulation of perception, thinking, and action under uncertainty is statistical inference, and the normative way of statistical inference is Bayesian" (Edelman and Shahbazi, 2011, p. 198).

Furthermore, despite the existence of alternative methods for formulating the problem of statistical inference, the corollary that “all such methods would ultimately reduce to Bayesian inference” (p. 198), if true, leaves little room for maneuver in the study of rationality. This is a form of Bayesianism, the view that the uncertainty surrounding the problems faced by the cognitive system is quantifiable, and that rational solutions to these problems are optimal probabilistic solutions. One can dispute this view without questioning Bayesian methods, probability theory, or mathematical statistics. The dispute concerns the application of these methods outside what Savage (1954), the founder of Bayesian decision theory, termed “small worlds.” Savage considered it “utterly ridiculous” (p. 16) to apply his theory in large worlds where uncertainty surrounds which states, actions, and consequences are relevant, and where we suffer from varying degrees of ignorance that undermine our ability to quantify the uncertainties we face.

Deductive logics in their various guises are used to define rationality in domains that demand deductive rather than inductive reasoning, so the expectation that our conception of rationality must adapt to fit the nature of the problem should be both familiar and uncontroversial. The claim here is that ecological rationality is an adaptation to the problem of decision making in large worlds. These are environments that we can observe but our partial ignorance precludes them from being probabilistically quantifiable (Binmore, 2009; Brighton and Gigerenzer, 2012), and a key implication of this claim is that optimality ceases to be a meaningful characteristic of rational decisions under these conditions. Critics, in contrast, have argued that because ecological rationality involves a conjecture about simple heuristics, it necessarily focuses on Marr’s algorithmic rather than the computational level. And given that rationality claims are traditionally made at the computational level, the assumption is that ecological rationality must therefore inherit principles of Bayesian rationality to explain the success of simple heuristics. In response, I have argued that Marr’s levels of analysis are incapable of resolving this debate because the statistical commitments driving the study of ecological rationality make the problem of statistical inference exploratory, algorithmic, and ill-posed. Marr’s levels of analysis, on the other hand, are tailored to well-posed problems with a precise mathematical formulation enabling the derivation of an optimal solution. As Marr (1977) himself noted, not all of the problems faced by biological systems are well-posed in this way.

In many areas of science, optimality modeling is seen as a heuristic of discovery that introduces abstractions that may or may not prove insightful when formulating problems, deriving solutions, and understanding the functioning of organisms and other complex systems (Dupré, 1987; Parker and Maynard Smith, 1990; Schoemaker, 1991). Seen in this way, it becomes imperative to consider the boundary conditions of optimality modeling, and recognize when the abstractions and assumptions needed to formulate optimality claims cease to be insightful and instead provide elegant solutions to problems of a different nature. Is the use of optimality in the study of rationality any different? I have argued that it isn’t, and the basis for my argument is a statistical one rooted in the uncertainty arising from what Simon (1957) referred to as informational bounds, rather than the more commonly assumed computational, cognitive, or biological bounds that motivate the study of bounded and ecological rationality. Seen from this statistical perspective, the study of ecological rationality can cast new light on the limited ability of probabilistic optimality to capture what is functional outside idealized worlds, those where all aspects of uncertainty are probabilistically quantifiable.

## References

- Anderson, J. R. (1990). *The adaptive character of thought*. Lawrence Erlbaum, Hillsdale, NJ.  
Anderson, J. R. (1991). Is human cognition adaptive? *Behavioral and Brain Sciences*, 14, 471–517.

- Baggio, G., van Lambalgen, M., and Hagoort, P. (2014). Logic as Marr's computational level: Four case studies. *Topics in Cognitive Science*, 7(2), 287–298.
- Bechtel, W. and Shagrir, O. (2015). The non-redundant contributions of Marr's three levels of analysis for explaining information-processing mechanisms. *Topics in Cognitive Science*, 7(2), 312–322.
- Binmore, K. (2009). *Rational decisions*. Princeton University Press, Princeton, NJ.
- Bishop, C. M. (2006). *Pattern recognition and machine learning*. Springer, New York.
- Breiman, L. (2001). Statistical modeling: The two cultures. *Statistical Science*, 16, 199–231.
- Brighton, H. (2006). Robust inference with simple cognitive models. In C. Lebiere, and R., Wray (Eds.), *Between a rock and a hard place: Cognitive science principles meet AI-hard problems (AAAI Technical Report SS-02-06)* (pp. 189–211). AAAI Press, Menlo Park, CA.
- Brighton, H. (2019). Beyond quantified ignorance: Rebuilding rationality without the bias bias. Economics Discussion Papers No 2019–25, Kiel Institute for the World Economy.
- Brighton, H. and Gigerenzer, G. (2007). Bayesian brains and cognitive mechanisms: Harmony or dissonance? In N. Chater, and M., Oaksford (Eds.), *The probabilistic mind: Prospects for Bayesian cognitive science* (pp. 189–208). Cambridge University Press, Cambridge.
- Brighton, H., and Gigerenzer, G. (2012). Are rational actor models “rational” outside small worlds? In S. Okasha and K., Binmore (Eds.), *Evolution and rationality: Decisions, co-operation and strategic behaviour* (pp. 84–109). Cambridge University Press, Cambridge.
- Brighton, H. and Gigerenzer, G. (2015). The bias bias. *Journal of Business Research*, 68, 1772–1784.
- Chater, N. and Oaksford, M. (1999). Ten years of the rational analysis of cognition. *Trends in Cognitive Sciences*, 3(2), 57–65.
- Chater, N., Oaksford, M., Nakisa, R., and Redington, M. (2003). Fast, frugal, and rational: How rational norms explain behavior. *Organizational Behavior and Human Decision Processes*, 90, 63–86.
- Chater, N., Tenenbaum, J. B., and Yuille, A. (2006). Probabilistic models of cognition: Conceptual foundations. *Trends in Cognitive Sciences*, 10(7), 287–291.
- Chen, Z. and Haykin, S. (2002). On different facets of regularization theory. *Neural Computation*, 14(12), 2791–2846.
- Czerlinski, J., Gigerenzer, G., and Goldstein, D. G. (1999). How good are simple heuristics? In G. Gigerenzer, P. M. Todd, and the ABC Research Group (Eds.), *Simple heuristics that make us smart* (pp. 119–140). Oxford University Press, New York.
- Domingos, P. and Pazzani, M. (1997). On the optimality of the simple Bayesian classifier under zero-one loss. *Machine Learning*, 29, 103–130.
- Dupré, J., (Ed.) (1987). *The latest on the best: Essays on evolution and optimality*. MIT Press, Cambridge, MA.
- Edelman, S. and Shahbazi, R. (2011). Survival in a world of probable objects: A fundamental reason for Bayesian enlightenment. *Behavioral and Brain Sciences*, 34(4), 197–198.
- Friedman, J. H. (1997). On bias, variance, 0/1-loss, and the curse-of-dimensionality. *Data Mining and Knowledge Discovery*, 1, 55–77.
- Geman, S., Bienenstock, E., and Doursat, R. (1992). Neural networks and the bias/variance dilemma. *Neural Computation*, 4, 1–58.
- Gigerenzer, G. and Brighton, H. (2009). Homo heuristicus: Why biased minds make better inferences. *Topics in Cognitive Science*, 1, 107–143.
- Gigerenzer, G. and Goldstein, D. G. (1996). Reasoning the fast and frugal way: Models of bounded rationality. *Psychological Review*, 103(4), 650–669.
- Gigerenzer, G. and Selten, R. (2001). *Bounded rationality: The adaptive toolbox*. MIT Press, Cambridge, MA.
- Gigerenzer, G., Todd, P. M., and the ABC Research Group (Eds.) (1999). *Simple heuristics that make us smart*. Oxford University Press, New York.
- Gilboa, I. (1987). Expected utility with purely subjective non-additive probabilities. *Journal of Mathematical Economics*, 16, 65–88.
- Gilboa, I., Postlewaite, A., and Schmeidler, D. (2012). Rationality of belief or: Why Savage's axioms are neither necessary or sufficient for rationality. *Synthese*, 187, 11–31.
- Gilboa, I. and Schmeidler, D. (1989). Maxmin expected utility with a non-unique prior. *Journal of Mathematical Economics*, 18, 141–153.
- Gintis, H. (2012). An evolutionary perspective on the unification of the behavioral sciences. In S. Okasha and K., Binmore (Eds.), *Evolution and rationality: Decisions, co-operation and strategic behaviour* (pp. 213–245). Cambridge University Press, Cambridge.
- Goldstein, D. G., and Gigerenzer, G. (2002). Models of ecological rationality: The recognition heuristic. *Psychological Review*, 109, 75–90.

- Griffiths, T. L., Lieder, F., and Goodman, N. D. (2015). Rational use of cognitive resources: Levels of analysis between the computational and the algorithmic. *Topics in Cognitive Science*, 7(2), 217–229.
- Griffiths, T. L. and Tenenbaum, J. B. (2006). Optimal predictions in everyday cognition. *Psychological Science*, 17(9), 767–773.
- Griffiths, T. L., Vul, E., and Sanborn, A. N. (2012). Bridging levels of analysis for probabilistic models of cognition. *Current Directions in Psychological Science*, 21(4), 263–268.
- Hadamard, J. (1902). Sur les problèmes aux dérivées partielles et leur signification physique. *Princeton University Bulletin*, 13, 49–52.
- Hand, D. J. and Yu, K. (2001). Idiot's Bayes: Not so stupid after all? *International Statistical Review*, 69, 385–398.
- Hastie, T., Tibshirani, R., and Friedman, J. (2001). *The elements of statistical learning: Data mining, inference, and prediction*. Springer, New York.
- Jones, M. and Love, B. (2011). Bayesian fundamentalism or enlightenment? On the explanatory status and theoretical contributions of Bayesian models of cognition. *Behavioral and Brain Sciences*, 34, 169–231.
- Kuncheva, L. I. (2006). On the optimality of naïve Bayes with dependent binary features. *Pattern Recognition Letters*, 27, 830–837.
- Marr, D. (1977). Artificial intelligence: A personal view. *Artificial Intelligence*, 9, 37–48.
- Marr, D. (1982). *Vision*. Freeman, San Francisco, CA.
- Marr, D. and Hildreth, E. (1980). Theory of edge detection. *Proceedings of the Royal Society of London B*, 207, 187–217.
- Marr, D. and Poggio, T. (1976). From understanding computation to understanding neural circuitry. *AI Memo* 357.
- Martignon, L. and Laskey, K. B. (1999). Bayesian benchmarks for fast and frugal heuristics. In G. Gigerenzer, P. M. Todd, and the ABC Research Group (Eds.), *Simple heuristics that make us smart* (pp. 169–188). Oxford University Press, New York.
- McClamrock, R. (1991). Marr's three levels: A re-evaluation. *Minds and Machines*, 1(2), 185–196.
- Ng, A. Y. and Jordan, M. I. (2002). On discriminative vs. generative classifiers: A comparison of logistic regression and naïve Bayes. In T., Dietterich, S., Becker, and Z. Ghahramani (Eds.), *Advances in neural information processing systems (NIPS)* 14. MIT Press, Cambridge, MA.
- Oaksford, M. and Chater, N. (2009). The uncertain reasoner: Bayes, logic, and rationality. *Behavioral and Brain Sciences*, 32, 105–120.
- O'Sullivan, F. (1986). A statistical perspective on ill-posed inverse problems. *Statistical Science*, 1, 502–518.
- Parker, G. A. and Maynard Smith, J. (1990). Optimality theory in evolutionary biology. *Nature*, 348, 27–33.
- Peebles, D. and Cooper, R. P. (2015). Thirty years after Marr's vision: Levels of analysis in cognitive science. *Topics in Cognitive Science*, 7(2), 187–190.
- Perlich, C., Provost, F., and Simonoff, J. S. (2003). Tree induction vs. logistic regression: A learning curve analysis. *Journal of Machine Learning Research*, 4, 211–255.
- Poggio, T. (2012). The levels of understanding framework, revised. *Perception*, 41(9), 1017–1023.
- Savage, L. J. (1954). *The foundations of statistics*. Wiley, New York.
- Schoemaker, P. J. H. (1991). The quest for optimality: A positive heuristic of science? *Behavioral and Brain Sciences*, 14, 205–245.
- Selten, R. (2001). What is bounded rationality? In G. Gigerenzer and R. Selten (Eds.), *Bounded rationality: The adaptive toolbox* (pp. 13–36). MIT Press, Cambridge, MA.
- Simon, H. A. (1957). *Models of man*. John Wiley & Sons, New York.
- Simon, H. A. (1991). Cognitive architectures and rational analysis: Comment. In K. Van Lehn (Ed.), *Architectures for intelligence* (pp. 25–39). Erlbaum, Hillsdale, NJ.
- Şimşek, O. and Buckmann, M. (2015). Learning from small samples: An analysis of simple decision heuristics. In D. D. Lee, M. Sugiyama, and R. Garnett (Eds.), *Advances in Neural Information Processing Systems (NIPS)* 28. MIT Press, Cambridge, MA.
- Stone, M. (1974). Cross-validated choice and assessment of statistical predictions. *Journal of the Royal Statistical Society B*, 36, 111–147.
- Tenenbaum, J. B. and Griffiths, T. L. (2001). Some specifics about generalization. *Behavioral and Brain Sciences*, 24, 762–778.
- Webb, G. I., Boughton, J. R., and Wang, Z. (2005). Not so naïve Bayes: Aggregating one dependence estimators. *Machine Learning*, 58, 5–24.
- Willems, R. (2011). Re-appreciating the why of cognition: 35 years after Marr and Poggio. *Frontiers in Psychology*, 2, 244.