

Title	Empiricism, sciences, and engineering: cognitive science as a zone of integration
Authors	Ross, Don
Publication date	2019-03-26
Original Citation	Ross, D. (2019) 'Empiricism, sciences, and engineering: cognitive science as a zone of integration', Cognitive Processing, 20(2), pp. 261-267. doi: 10.1007/s10339-019-00916-z
Type of publication	Article (peer-reviewed)
Link to publisher's version	https://link.springer.com/article/10.1007/s10339-019-00916-z - 10.1007/s10339-019-00916-z
Rights	© Marta Olivetti Belardinelli and Springer-Verlag GmbH Germany, part of Springer Nature 2019. This is a post-peer-review, pre- copyedit version of an article published in Cognitive Processing. The final authenticated version is available online at: http:// dx.doi.org/10.1007/s10339-019-00916-z
Download date	2025-07-30 17:09:49
Item downloaded from	https://hdl.handle.net/10468/7790



University College Cork, Ireland Coláiste na hOllscoile Corcaigh

Empiricism, Sciences, and Engineering: Cognitive Science as a Zone of Integration

(published in Cognitive Processing 20: 261-267)

Don Ross don.ross931@gmail.com

School of Sociology, Philosophy, Criminology, Government, and Politics University College Cork

> School of Economics University of Cape Town

Center for Economic Analysis of Risk Georgia State University

<u>Abstract</u>

An article by Alexandra Kirsch accepted for publication in *Cognitive Processing* occasioned debate among reviewers about broad methodological issues in cognitive science. One of these issues is the proper place of Popperian falsificationism in the interdisciplinary cluster. Another is the tension between abstract models and theories that apply to wide classes of cognitive systems, and models of more restricted scope intended to predict specifically human patterns of thought and behavior. The lead editorial in a Commentary debate invited by the journal's editors considers these issues from the perspective of a pragmatist philosophy of science inspired by Herbert Simon's classic (1969) reflections on the blurring of the distinction between science and engineering in cognitive science.

1. Introduction

Alexandra Kirsch's submission of her article 'A unifying computational model of decision making' (2019) raised interesting policy issues for the editors of *Cognitive Processing.* These were triggered by reports from peer reviewers that diverged sharply in their recommendations, but did not revolve around the usual matters of referee judgment, that is, factual accuracy, and soundness of data interpretation and analysis. Instead, different reviewers defended conflicting philosophical views about what constitutes a genuine contribution to knowledge in cognitive science, and about the proper scope of a journal with the specific mission of advancing knowledge of (mainly) natural information processing. This unusual circumstance offered the editors an interesting opportunity to reflect on the challenges of integrating the disciplinary strands that make up cognitive science. The Editor-in-Chief consequently requested the Handling Editor to initiate a Commentary debate

that begins with the perspective of Kirsch's referees, but will subsequently be open to further contributions from interested readers of the journal.

Methodological disputes are ultimately resolved by the empirical progress of science, not by philosophical arguments. However, reflection on such arguments can efficiently guide research design and coordinate the collective critical interpretation of new findings. Such meta-scientific awareness is often particularly fruitful in scientific research clusters, such as cognitive science, that entangle various disciplinary norms and traditions.

I will not here review or summarise the reviewers' perspectives as stated in their respective reports. Nor will I offer any critical view of the modeling or argumentation in Kirsch's article, the scientific value of which will be measured, in the usual fashion, by the extent of its influence on further research. My purpose is instead to place the divergence of perspectives revealed by the refereeing process in the context of a large relevant literature in the philosophy of science generally, and in the philosophy of cognitive science specifically.

I proceed as follows. Section 2 reflects on a 'Popperian' critical assessment frame suggested by one of Kirsch's reviewers. I will not pronounce for or against the status of this frame as a guide to how science 'should' be done; it is inadvisable for philosophers of science to position themselves as arbiters of 'proper' science, a role which no one has authority to concede. Instead, I will use these thoughts in the shadow of Popper as a basis for discussing, in Section 3, the heterogeneity of conceptions of 'theory' in different rooms of the cognitive science mansion. These differences interact with orthogonal distinctions that sort scientific contributions along a continuum from relatively purely descriptive to primarily practical. This contrast, between 'pure science' and 'engineering', will be the subject of Section 4. Cognitive science incorporates work all along the continuum, even in those parts of it that focus on natural, as opposed to built, cognitive processes. Section 5 concludes.

The discussion to follow was reviewed by the other participants in the present Commentary debate prior to their preparation of their contributions. The Editors of *Cognitive Processing* invite and encourage further opinions from scientific readers, for consideration for publication in forthcoming issues.

2 Falsifiability and empiricism

The most famous philosopher of science *among scientists* is certainly Sir Karl Popper (1902-1994). One reason for this is that Popper was particularly concerned to *demarcate* science from other activities, especially from 'pseudo-sciences' that try to illegitimately borrow the epistemic authority that is institutionally conferred upon scientists. Most scientists are called upon from time to time to indicate at least rough demarcation criteria: students receiving their first scientific training often ask for such markers, and in a world where highly visible and influential political actors constantly incite people to ignore or reject inconvenient scientific findings, a socially responsible scientist must have something to say when asked why her professional

views should be given special weight. Popper's demarcation criterion, though strictly negative, has enjoyed a hugely successful mimetic history. According to it, a necessary condition for a statement, hypothesis, or generalization to be scientific is that it come attached with specified potential empirical observations that, if discovered through a replicable procedure, would demonstrate the statement, hypothesis, or generalization in question to be untrue. This is the requirement that a scientific proposition be *falsifiable* (Popper 1969). Popper does not demand that the scientist currently be practically able to perform a falsification test or position herself for a decisive observation. But her claims must imply a difference to some clearly identified empirically contingent state of affairs.

Few contemporary philosophers of science still baldly pronounce themselves 'Popperians'. I will not review the numerous respects, identified in the philosophical literature, in which Popper's view is too simple, too sweeping, gives too much weight to unrepresentative exemplars from classical physics, and if taken literally would exclude too much activity that is obviously scientific. Interested nonphilosophers can find a good short review of these problems in Catton (2004). At the very least, a viable Popperian philosophy must incorporate responses to Kuhn's (1962) historicism, which was the project launched by Lakatos (1970) and subsequently extended by numerous philosophers of science.

Notwithstanding the complexities involved in what preoccupies philosophers, namely trying to *exactly* state a viable formulation of falsificationism, contemporary scientists often defend and impress upon their students a general 'Popperian attitude' that is unarguably a core part of scientific culture. Setting aside pure mathematics as a domain that receives its own extensive philosophical consideration, theory or speculation that floats free of empirical implications is not scientific. The relationship between a body of theory and its empirical evidence can be complex, non-obvious, or take years to clearly establish; but it is a relationship that no scientist is institutionally allowed to disregard. This principle plays an important role in real disputes over the allocation of scientific resources. For example, critics of string theory as a unifying foundation for physics argue that it is mathematics for which no specific physical interpretation is ever likely to be established on the basis of experiment. Many physicists believe that if *at some point*, string theorists cannot demonstrate that this charge was premature, then physics departments should stop hiring them (Smolin 2006; but see Dawid 2014).

In my view the most serious problem with consistently applying a *generic* Popperian attitude is that the borderline between mathematics and empirical science is both indeterminate and heterogeneous across disciplines. But fuzzy boundaries are still boundaries, as long as one can identify clear exemplars of work on either side of them. When I referee papers in economics I sometimes recommend rejection on grounds that a piece of work, however elegant, is economically empty mathematics, and I do this with confidence despite recognizing that there are other papers about which a similar judgment would be irresolvably debatable.

Let us now descend from the rarified heights of philosophical generalization to more focused concerns around trying to give the Popperian attitude (not the literal Popperian dogma) its due in cognitive science.

3. Theory across disciplines

Cognitive science is an interdisciplinary enterprise. That is, indeed, the entire point of distinguishing it from cognitive *psychology* narrowly conceived. Standard histories of cognitive science present it as a union of elements of psychology, computer science, linguistics, and philosophy of mind that was forged in the 1960s and 1970s as researchers acknowledged deep structural isomorphisms amongst information processing implementation in people, non-human animals, and built computers. Cognitive science arguably began on the basis of a relatively unified theory, the mathematical account of computation (Lewis and Papadimitriou 1981) combined with the then-nascent ontology of effective algorithms for Von Neumann architectures (Knuth 1968). However, by the end of the 1980s rooms were opening in the interdisciplinary mansion that relied on other foundations: parallel distributed processing as a generalization of perceptron architectures (Rummelhart et al 1986), behavior-based robotics (Brooks 2013), swarm intelligence (Bonabeau et al 1999), artificial life (Langton 1989), cognitive neuroscience (Trehub 1991, Churchland and Sejnowski 1992), behavioral economics (Kagel et al 2007), (crossspecies) comparative psychology (Griffin 1984), social-ecological cognition (Byrne and Whiten 1988), developmental cognition (Wellman 1990), and consciousness studies (Baars 1988, Dennett 1991). Contributing disciplines have expanded to include neuroscience, ethology, information science, economics, and sociology.

In this teeming research ecosystem, prospects for a general theory of cognition must now be regarded as exceedingly remote. Most contemporary philosophers of science would urge a relaxed attitude to this fact of life: the history of science is not mainly a story of intertheoretic reduction or unification, occasional triumphs of deep synthesis notwithstanding (Mitchell 2009). Newton's great unification that marked the exemplary scientific achievement for a couple of centuries was in this respect a false portent. In one respect, the picture of science, and cognitive science in particular, as an archipelago instead of a tower makes the Popperian attitude more plausible, since only relatively isolated theories can face the tribunal of experience straightforwardly. A relatively unified edifice would be a densely connected 'web of belief' in which each piece would gain support from the rest and none might generate clear predictions that could be distinctively associated with them (Ouine and Ullian 1970). The Popperian attitude could in such a hypothetical context only be easily maintained with regard to the foundations of the entire multidiscipline, and falsificationism would then amount to trying to provoke Kuhnian revolution. That was certainly not Popper's conception. He would thus likely be encouraged by the state of play in cognitive science, where relatively isolated theories do periodically get rejected based on fresh evidence.

However, cognitive science is characterized by a higher order of complexity than the divergence of foundations on which its proliferation of theories rests. Contributing

disciplines differ from one another in what their practitioners take theory to *mean*, in the operational sense. This, I suggest, is one of two key points in diagnosing the conflicting responses of Kirsch's referees.

All scientists recognize that they must simplify their descriptions of reality in trying to produce generalized knowledge. After all, we have one perfectly accurate model of the world available to us, namely, the world itself; but mute holistic appreciation isn't scientific knowledge. We need to do analysis, and analysis involves idealized abstraction. Broadly speaking, we find two general strategies for going about this. Most disciplines combine the strategies to some extent. But for concreteness of description, and with a view to the specific methodological tensions that characterized the reviewing of Kirsch's paper, I'll illustrate the contrast by comparing psychology and economics as disciplines that are each unusually pure in overwhelmingly emphasizing one abstracting strategy or the other.

Psychologists seek generalizations about processes they typically cannot directly observe, but must infer from measurable effects. Their preferred methodology for trying to achieve stability and convergence in the face of this challenge is to have developed, and to enforce through peer review mechanisms, a close record of relationships between conservatively extended experimental *protocols* and theoretical *constructs* anchored to measurement scoring systems. Most readers of *Cognitive Processing* will be skilled operators and guardians of this record. The approach is derived from the venerable craft of clinical diagnosis, where practitioners have long needed to control subjectivity in conjecturing underlying diseases and syndromes from manifest symptoms. To this the psychologist adds statistical testing of instrument and protocol reliability and validity with respect to construct application. Psychometrics is, to a first approximation, the statistical theory of measures of construct validity.

In this context, psychological *theories* are essentially *hypotheses* about which constructs are implicated in the production of which behaviors. Hypothesis testing very naturally recruits the Popperian attitude as a rationale for practical features of methodology. Identifying a null hypothesis and then designing an experiment that might refute it is falsificationism at work.¹ This encourages psychologists to develop highly trained and sensitive attunement to potential *confounding* causal influences identified *ex ante* during experimental design. In general, psychologists prefer that confounds that cannot be straightforwardly controlled in linear regressions should be shut out of the lab.

This in turn inclines psychologists away from trying to discover 'laws of behavior'. Many psychologists explain their aversion to large-scale generalizations by noting

¹ Philosophers will be quick to point out that it's not quite the real thing, because Popper urged that scientists should attempt to falsify theoretical hypotheses, not null hypotheses. On the other hand, since a null hypothesis is defined as such by reference to a theoretical hypothesis, scientists might reasonably regard this as quibbling over semantics.

that human and animal behavior (and now even behavior of neural networks and multi-agent simulations) is non-deterministic. But various other disciplines (climatology, for example, and molecular genetics) are comfortable with broadsweep generalizations that are stochastic. The work of the philosopher of science Nancy Cartwright points to a deeper explanation of psychologists' scrupulous modesty: to the extent that a lab environment is designed to protect phenomena from the intrusive confounds that contaminate the wild environment, putative laws governing the messy outside seem like 'lies' (Cartwright 1983, 1989, 1999).

Let us now contrast the methodology of psychology with that of economics. To keep the contrast within useful bounds, I will consider only lab-based experimental economics. The language of 'constructs' is foreign to economists; referential terms in their models are taken to be directly isomorphic to real objects and processes. Of course, economists must simplify and idealize causal relationships just like all scientists. Their vehicle for performing this is the *structural model*. Most of the practical art of experimental economics consists in developing tasks for experimental subjects with incentives for behavior provided in such a way that the model can identify predicted changes in behavior that vary with changes in the experimenter-controlled incentives. Apprentice economists learn from experience how to design models with 'the right' degree of parametric structure: too few parameters will mis-specify phenomena, generating 'specification error', and too many will undermine the likelihood of successful identification. Econometrics is the statistical theory of model specification and estimation.

'Theory', for economists, does *not* refer to hypotheses about empirical relationships. In some respects economists think of 'theory' in the way that mathematicians do: theory is simply precise specification. However, in light of their methodology economists use theory at two levels, to construct two different kinds of models. At one level of theory are economic models that specify causal channels (or, in more metaphysically humble language that many economists prefer, channels of 'influence') in the world, and at another level are data models that specify relationships between obervations and inferences that can be made about estimated variable coefficients under different (usually non-linear) regression models. 'Theory', then, refers to formal structural specification of some class of systems that are, relative to some robustness criterion across models (see below), an equivalence class. When one economist doubts that another economist's model is 'economically significant', the former will not complain about a 'rejected hypothesis'; she might instead say that the project she criticizes is an exercise in 'mere theory'.

Because economists estimate structural models, any unobserved influence that is correlated with any variable in the model amounts to a specification error and will generate biased estimation of regression coefficients. Therefore, when economists think of possibly causally relevant factors during experimental design, they ideally set up a new treatment group arm of the experiment where the factor in question can be independently varied. If observation conditions or budget limitations prevent this first-best approach, they will revisit their basic design and their data modeling to ensure that the factor can be identified and its influence estimated. So, whereas psychologists seek to exclude 'confounds' from the lab, economists come up with strategies to bring 'confounds' *into* the lab. The psychologist's lab is a bunker; the economist's is intended to be a microcosm of the world.

Economists are as reluctant as psychologists to speculate about sweeping laws. However, in the case of the former this really does mainly just reflect metaphysical squeamishness about the idea of stochastic causation. (The great philosopher of science C.S Peirce advised against such anxiety; but the idea that a 'real' cause must always exert its effects unless blocked by another identified cause is a bit of folk metaphysics that holds on even among many physicists; see Ladyman and Ross 2013.) Economists show their relative immodesty (by comparison with psychologists) about making inferences from data in the attention they give to estimating cross-experimental robustness of observed effect strengths relative to classes of data models (Neumeyer and Plümper 2017). Their concern with establishing 'external validity' of experiments (Guala 2005) reflects their practice of trying to proxy the world in their labs, rather than trying to eliminate all but a selected aspect of the world from the lab.

Psychologists' emphasis on isolated hypothesis testing gives a different flavor to their empiricism from that of economists. Economists' practice of accepting observed effects because they can successfully specify and identify them, and then accepting them as general causal structures when they can specify, identify and estimate them robustly, is in much greater tension with the Popperian attitude. The leading economic methodologist of the mid-twentieth century, Mark Blaug (see Blaug 1980) was a committed Popperian and caused a generation of economics PhD students to be made to read selections from Popper and Lakatos. It is noteworthy that Blaug ultimately grew deeply dissatisfied with his discipline, and in his final writings (e.g. Blaug 2002) denounced its mainstream practitioners for seldom conducting themselves as Popper counseled. The later Blaug can be interpreted without much strain as wishing that economists did psychology instead. This proves, as it were, the force of the contrast.

Kirsch's interest in specifying broad classes of heuristics that might characterize various cognitive systems, including humans, at an abstract level but perhaps not in specific detail, would not surprise researchers who approach the world in the way that economists do. Because they are not shy about estimating 'externally valid' effects, and because they are motivated to find the broadest range of phenomena that are equivalence classes from the point of view of a model or family of models, it would be natural for an economist to be interested in heuristics that might restrict choice behavior in both people and various kinds of artificial systems, even if processing details varied from instance to instance.

It may seem odd, even to a non-psychologist, to talk about shared heuristics that might yet involve varying processing details. After all, a heuristic *is* a restriction on processing in the first place. But specification can occur at any of a multiple range of scales of exactness. Suppose that, as a heuristic for not (indefinitely) failing to notice e-mails I every day review all messages still in my inbox from one month ago. Suppose that you instead do a daily review of every message from two months ago. Are we using different heuristics or two variants of the same heuristic? Let us change the dimension of imagined processing difference: suppose I scan the set of old messages manually, while you subject exactly the same set to search by a programme you wrote that checks to see whether you acted on them. Are those different heuristics because they implement different algorithms, or the same heuristic because they address the same problem, unreliable attention under pressure from interruptions, by the same general device, a review anchored to the calendar date?

Heuristics have in fact been studied by both psychologists and economists, sometimes acting in interdisciplinary teams (see, e.g., chapters in Gigerenzer and Selten 2001). All are interested in the efficiency of the heuristics they study. Efficiency, after all, is, on some construal or other, what makes a heuristic a heuristic. But psychologists typically concentrate on a heuristic's efficiency in terms of how much computational effort it requires compared with alternatives; whereas the economists are more likely to focus on the ex ante reliability with which rival heuristics get the same job done. These different loci of attention will tend to lead to systematically different levels of specification of equivalence classes. The psychologist will naturally discriminate at a relatively granular scale and seek to test hypotheses that distinguish between variants at this scale, while the economist might aim to model the widest set of implementations that are, within the feasible set given the time or energy budget (i.e., the budget that rules out the first-best solution and motivates the heuristic approach), equally reliable as solutions to a common problem.

4. Science and engineering

In her paper Kirsch cites Herbert Simon's classic book *The Sciences of the Artificial*, originally published in 1969 and revised twice subsequently (1981, 1996). Simon's central question was how there could be a proper *science* of designed systems, as opposed to simply engineering of them. If we think of science as the disinterested examination of the natural world – where by 'natural' people typically mean 'independent of design by people' – then it is indeed puzzling that we might scientifically investigate systems in which the characteristics were chosen by people in the first place.

What the class of designed systems that Simon discusses have in common is that they are intended to optimize, or to reach a performance threshold short of optimization that a designer deems good enough, in conditions that involve adaptive interface with a class of environments. By way of contrast, the problem of building an excellent toaster or car is a relatively pure engineering problem: the only restriction on making it as good as you want is cost, because neither toasters nor cars are very sensitive to varying 'interface conditions'. This is not true of AI systems, or robots, or artificial markets. There are generalizations to be discovered about adaptive interface. Indeed there are *limitless* potential such generalizations because adaptive interfaces are scale sensitive and there are no in-principle limits to the scales at which an inquirer might aim her interest and try to solve her design problem. Discovering these generalizations is a matter for science because scale effects are typically non-linear, and the locations of transition and inflection points need to be found through theoretically disciplined experimentation.

The sciences of the artificial blend into (some) sciences of the natural, because biologists are also in the business of seeking generalizations about adaptive interfaces. It might be said that the systems studied by biologists – and, as a special case, by psychologists – are different because they were not designed. Of course it is important to stress this if by 'designed' we mean 'assembled with a preconceived end in view'. The idea that biological systems were designed in *that* sense is creationism, and anathema to science. But, as the philosopher of cognitive science Daniel Dennett (1987, 1995) has emphasized, natural selection *is* a kind of designer in the very important sense that it conserves better solutions to problems and tends to eliminate inferior solutions to problems. Furthermore, nature's solutions often offer vital clues to engineers because natural selection runs prodigiously long and varied experiments, albeit with only rare instances where treatments and controls are isolated.

The domain of heuristics, Kirsch's topic, is the perfect point of convergence of the sciences of the artificial and the biological (psychological). A solution to a problem is a heuristic just in case it trades off some accuracy in completeness of representation and / or cross-domain reliability for reduced cost (in energy or time) of processing. That is the only kind of solution that natural selection can produce in principle. So all of psychology is a study of heuristics. But I said above that the boundary between the natural and the artificial, in Simon's sense, is not clean or clear. The environment in which the specifically human mind has evolved is a *social* environment, teeming since the great encephalization with engineered artifacts. Humans are niche constructors on an unprecedented scale (Odling-Smee et al 1996; Sterelny 2003), so the dance of adaptation between their minds and their environments is co-evolutionary.

A matter of awkwardness for Kirsch's referees is that her paper is vertiginously poised at what seems to be an unstable or arbitrary point of engagement on this fuzzy borderline between the naturally selected and the engineered. She studies a class of heuristics more general than the human mind, while wanting to make sure that human heuristic processes are a proper sub-class of her target. But her target class cannot be every possible heuristic process, because that would amount to the unfathomable domain of every possible solution procedure that is budgetconstrained.

I doubt that it is possible in principle to pre-isolate the class of processes to which Kirsch's intuitions direct her attention. In Section 3 I contrasted the psychologist's methodology with the economist's. It might seem at first glance that both of these methodologies, in their opposite ways, avoid the arbitrariness that seems to haunt Kirsch's project. The psychologist, it might be said, restricts herself to the *actual* heuristics that natural selection produced in humans. Thus she avoids the problem of an indeterminate boundary in possibility space. The economist, on the other hand, begins from the mathematically identified ideal solution space and systematically introduces costs, according to a rational sequence governed by which solution functions mathematically nest which others. Thus we have a (radically) different procedure that likewise is safe from charges of arbitrariness. Did Kirsch provoke a muddled refereeing process because her work falls into a ravine between wellmotivated disciplinary alternatives?

5 Conclusion

In fact, I think that neither the psychologist's apparent solution to the arbitrariness of the target space, nor the economist's supposed way around it, can actually hold us on the firm dry ground conjured above.

First consider what the psychologist's 'solution' must assume: that we can coherently aim to study the heuristics our brains would have run naturally, by themselves, if they could have undergone development in a psychologist's lab isolated from confounds. This is purely fantastical, meaning not merely practically impossible but not conceivable with the slightest degree of specifiable detail. All people have evolved the heuristics they use while entangled from infancy in a network of artifactual ones. We are, as the cognitive scientist Andy Clark (2003) puts it, 'natural-born cyborgs'.

As for the economist's solution, the ideal decision-theoretic space is usually thought to be the axiomatic expected utility theory of Savage (1954). However, Binmore (2009) has recently forcefully reminded economists of a point Savage made to begin with (and this time, to judge from the literature, the point is getting through). Application of expected utility theory is only 'best' in environments where all ex ante uncertainty can be quantified based on prior knowledge, i.e., by some procedure other than uninformed conjecture. These environments are what Savage called 'small worlds'. Humans use heuristics even in many small worlds, at least when these are novel to them, but the very novelty of such applications suggests that natural selection might not have pre-selected them. In any event, almost all of the decision worlds of interest to cognitive scientists – psychologists, AI researchers, and behavioral economists alike – are large worlds. There is no end to the useful economic modeling that can be applied to such worlds, but from the perspective of being able to elegantly pre-configure the boundaries of this study domain economists are in the same position as Kirsch.

So, I suggest, are we all. What is the domain of 'cognitive processing'? We can't even really say that we know it when we see it, because sometimes it turns up in places where no one would have expected it, e.g. the arms of octopi (Sumbre et al 2001) or 'zero-intelligence' traders in just the right kind of market (Gode and Sunder 1993). There will be other submissions to *Cognitive Processing* that cause referees to

scratch their heads in different ways. That's editorial life in an adaptively open world.

Compliance with ethical standards

The author has no potential conflicts of interest to disclose. The article does not report any results obtained using human participants or animals.

References

Baars, B. (1988). *A Cognitive Theory of Consciousness.* Cambridge University Press.

Binmore, K. (2009). Rational Decisions. Princeton University Press.

Blaug, M. (1980). *The Methodology of Economics*. Cambridge University Press.

Blaug, M. (2002). Ugly currents in modern economics. In U. Mäki, ed., *Fact and Fiction in Economics*, pp. 35-56. Cambridge University Press.

Bonabeau, E., Dorigo, M., and Theraulez, G. (1999). *Swarm Intelligence*. Oxford University Press.

Brooks, R. (2013). Cambrian Intelligence. MIT Press.

Byrne, R., and Whiten, A., eds. (1988). *Machiavellian Intelligence*. Oxford University Press.

Cartwright, N. (1983). *How the Laws of Physics Lie.* Oxford University Press.

Cartwright, N. (1989). *Nature's Capacities and their Measurement*. Oxford University Press.

Cartwright, N. (1999). *The Dappled World*. Cambridge University Press.

Catton, P. (2004). Constructive criticism. In P. Catton and G. Macdonald, eds., *Karl Popper: Critical Appraisals*, pp. 50-77. Routledge.

Churchland, P., and Sejnowski, T. (1992). *The Computational Brain*. MIT Press.

Clark, A. (2003). Natural-Born Cyborgs. Oxford University Press.

Dawid, R. (2014). *String Theory and the Scientific Method.* Cambridge University Press.

Dennett, D. (1987). Evolution, error, and intentionality. In D. Dennett, *The Intentional Stance*, pp. 287-321. MIT Press.

Dennett, D. (1991). Consciousness Explained. Little Brown.

Dennett, D. (1995). *Darwin's Dangerous Idea*. Simon and Schuster.

Gigerenzer, G., and Selten, R., eds. (2001). *Bounded Rationality: The Adaptive Toolbox.* MIT Press.

Gode, D., and Sunder, S. (1993). Allocative efficiency of markets with zerointelligence traders: market as a partial substitute for individual rationality. *Journal of Political Economy* 101: 119-137.

Griffin, D. (1984). *Animal Thinking*. Harvard University Press.

Guala, F. (2005). *The Methodology of Experimental Economics*. Cambridge University Press.

Kagal, J., Battalio, R., and Green, L. (2007). *Economic Choice Theory.* Cambridge University Press.

Kirsch, A. (2019). A unifying computational model of decision making. *Cognitive Processing* [this issue].

Knuth, D. (1968). *The Art of Computer Programming Volume One: Fundamental Algorithms*. Addison-Wesley.

Kuhn, T. (1962). *The Structure of Scientific Revolutions*. University of Chicago Press.

Ladyman, J., and Ross, D. (2013). The world in the data. In D. Ross, J. Ladyman, and H. Kincaid, eds. (2013), *Scientific Metaphysics*, pp. 108-150. Oxford University Press.

Lakatos, I. (1970). Falsification and the methodology of scientific research programmes. In I. Lakatos and A. Musgrave, eds., *Criticism and the Growth of Knowledge*, pp. 91-196. Cambridge University Press.

Langton, C., ed. (1989). *Artificial Life*. Addison-Wesley.

Lewis, H., and Papadimitriou, C. (1981). *Foundations of the Theory of Computation.* Prentice-Hall.

Mitchell, S. (2009). Unsimple Truths. University of Chicago Press.

Neumayer, E., and Plümper, T. (2017). *Robustness Tests for Quantitative Research.* Cambridge University Press.

Odling-Smee, J., Laland, K., and Feldman, M. (1996). Niche construction. *American Naturalist* 147: 641-648.

Popper, K. (1969). Conjectures and Refutations. Routledge and Kegan Paul.

Quine, W.V.O., and Ullian, J. (1970). *The Web of Belief*. Random House.

Rummelhart, D., McClelland, J., and the PDP Working Group (1986). *Parallel Distributed Processing Volume 1: Foundations.* MIT Press.

Savage, L. (1954). *The Foundations of Statistics.* Wiley.

Simon, H. (1969). *The Sciences of the Artificial* (2nd edition 1981; 3rd edition 1996). MIT Press.

Smolin, L. (2006). *The Trouble With Physics.* Houghton Mifflin Harcourt.

Sterelny, K. (2003). *Thought in a Hostile World*. Blackwell.

Sumbre, G., Gutfreund, Y., Fiorito, G., Flash, T., and Hochner, B. (2001). Control of octopus arm extension by a peripheral motor program. *Science* 293: 1845-1848.

Trehub, A. (1991). *The Cognitive Brain.* MIT Press.

Wellman, H. (1990). *The Child's Theory of Mind.* MIT Press.