This monograph is a collection of conference contributions chosen by the editors who led a three-year project on Evolution, Cooperation, and Rationality. The collected works are held together by a 6-page introduction identifying common strands and differences of positions in the different chapters. Since no two chapters have a common author the chapters do not build on each other. Rather, they offer a variety of perspectives on a number of different aspects of Rationality & Evolution. The monograph does thus not, despite the possibly somewhat misleading title, offer a unified account of Evolution & Rationality. Neither is there a clearly identifiable overarching theme that runs through the entire book.

Who may benefit from engaging with this book: Anyone with some background in rationality and an interest in evolution interested in the broadening of horizons. The chapters are written at a level of abstraction that should make them accessible to a wide audience, including interested students with some previous exposure to the area, i.e. some basic understanding of decision making such as utility functions, pay-off functions and mixed Nash equilibria. After a pre-selection by a convenor the book may very-well serve as the basis for student presentations in master or Ph.D. seminars.

This book is unsuitable for non-initiated readers and anyone looking for a unified account or a systematic treatment of either rationality, evolution, or rationality & evolution.

The chapters vary widely in style and content: From rather technical chapters to thirty-page long philosophical discussions not containing a single formula. Given the wide scope of the book and the limited space available for this review I shall concentrate in the following on two chapters which I consider most worthy of discussion.

An evolutionary perspective on the unification of the behavioral sciences by Herbert Gintis lays out four main types of models that are used to investigate decision making; these are in turn economic, sociological, biologic and psychological models. To combine these four distinct and, according to Gintis, incompatible models the following five conceptual units are introduced: a) gene-culture coevolution, b) the sociopsychological theory of norms, c) game theory, d) the rational actor model and e) complexity theory.

The chapter proceeds by giving overviews of the developments of these units over the last decades and shows how they are interdependent in various ways. Although, there is hardly any genuinely new material nor an agenda for future research, the overview over different aspects it provides is illuminating.

I would thus suggest to a non-expert reader to read this chapter before the other chapters of this monograph. The other chapters may then be (better) understood against the backdrop of the overview provided by Herbert Gintis.

Brighton & Gigerenzer offer an answer to the question: "Are rational actor models "rational" outside small worlds?" Since this chapter is most relevant to philosophically interested readers I shall now solely focus on it.

In this chapter the authors consider Savage's problem (Leonard Jimmie Savage. The Foundations of Statistics. Dover Publications, 1954) for large worlds. That is, how to draw inductive inference in highly complex worlds ripe with uncertainty. Inductive inference is here understood as identifying systematic patterns in observations.

At first, the authors describe how an organism drawing inductive inferences faces a myriad of types of uncertainty. Not only is an inductive organism confronted with the

following basic categories of uncertainty: Stochasticity, underspecification, misspecification and nonstationarity the organism faces uncertainties, which arise from the discrepancy between the environment and abstract properties of the hypothesis space, but the organism also has to deal with real-world uncertainties.

A small world à la Savage satisfies two distinct properties: i) the structure and the properties of a problem are known a priori and ii) the structure of the hypothesis space and prior allow exact calculations of posterior probabilities. Given the uncertainties an organism faces in large worlds (as well as further complications), these two properties cannot be assumed to hold in large worlds. This then raises the question if/how rational actor models à la Savage can be fruitfully applied to large world problems?

Determining an optimal response in a large world is pointless, all one can do is to look for better and better models and better decision heuristics. These heuristics sometimes achieve better results than every fully rational Bayesian method advocated by a small world model of the problem.

Thus, decision heuristics and not rational actor models should be studied to gain a better understanding of inductive inference in large worlds. One key question is which heuristics are robust in the face of uncertainties. A further important question is when and why heuristics are robust. An answer to the latter question may rely in substantial parts on insights drawn from studying rational actor models.

In their concluding remarks the authors argue that humans are astonishingly well-adapted to large uncertain worlds in which they outperform the most advanced man-made machinery. That is we are not biased, error-prone users of dumb heuristics. Thus, understanding inductive inference in humans is a highly intractable empirical question that cannot be understood, let alone fully solved, by the study of rational actors in small worlds. As our understanding of large worlds grows our need for small worlds diminishes.

As an inductive logician, the pragmatic position advocated by Brighton & Gigerenzer appears very sensible to me. If the aim of Inductvie Logic (IL) is to understand, describe, explain, explicate, support or improve an organism's inductive inferences, then any meaningful evaluation of a formal system of IL has ultimately to depend on the real world. Thus, any Pure Inductive Logic (PIL) (Jeff Boris Paris and Alena Vencovskà. Pure Inductive Logic. Cambridge University Press, 2013. forthcoming.) can only serve as a proxy or intermediate towards understanding, describing, explaining, explicating, supporting or improving inductive inferences in the real world.

This view is in-line with the more practical approach Design Science (DS) (A.R. Hevner, S.T. March, J. Park, and S. Ram. Design Science in Information Systems Research. MIS Quarterly, 28(1):75105, 2004) takes, which investigates, among other questions, how to evaluate man-made artefacts, such as a formal system of IL. The mantra in DS is that any developed artefact ought to be useful to some users, these users may for instance be humans performing inductive inferences or scientists studying inductive inferences.

If the main goal of IL is to study how an "ideal" epistemic agent ought to make inductive inferences, then Brighton & Gigerenzer's argument only applies for agents situated in a large world. Given their entirely pragmatic approach, I assume that Brighton & Gigerenzer would express the view that studying such agents only in small worlds without ever drawing any conclusions for large worlds is a futile exercise. Again, I find

myself in agreement with Brighton & Gigerenzer.

On one important point I disagree with the advocated position. I believe that by studying small worlds our understanding of them improves which allows us to fruitfully consider more and more complex small worlds. This brings these models closer to large worlds in relevant aspects. I concur with Kim Sterelny's argument in Chapter 10 that, other things being equal, the increase in complexity buys us improved explanatory and predictive leverage over real-world systems. Concluding, I cannot foresee a day when complex rational actor models go out of fashion.