

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

E Pluribus Multa In Unum: The Rationality Multiverse

Permalink

<https://escholarship.org/uc/item/5bp0g99w>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 34(34)

ISSN

1069-7977

Authors

Besold, Tarek R.
Kuhnberger, Kai-Uwe

Publication Date

2012

Peer reviewed

E Pluribus Multa In Unum: The Rationality Multiverse

Tarek R. Besold and Kai-Uwe Kühnberger
({tbesold | kkuehnbe}@uni-osnabrueck.de)

Institute of Cognitive Science
University of Osnabrück
49069 Osnabrück, Germany

Abstract

The paper argues for a new view on and an approach to rationality as a concept of study and modeling paradigm of human behavior. After critically reviewing classical (normative) approaches to rationality, decision-making, and rational behavior, we present cornerstones of a positive, integrative, and holistic conception of these cognitive capacities. A discussion of key elements of this new view is given, and possible consequences and implications are considered.

Keywords: Rationality; Human-level intelligence; Subject-centered model; Cognitive capabilities; Integrative model.

Introduction

With Aristotle's famous characterization of man as a rational animal ("zoon logikon") in his *Metaphysics* (Tredennick, 1933-35), and the ascription of a rational principle to the human being in his *Nicomachean Ethics* (Broadie & Rowe, 2002), the idea of seeing rationality and rational behavior as indispensable parts of our humanity has been introduced into human self-conception, starting centuries of inquiry into the nature and properties of this alleged *conditio sine qua non*. Still millennia after Aristotle, Descartes is only one amongst many famous scholars explicitly mentioning the question in his writings: "*But what is a man? Shall I say 'a rational animal'? No; for then I should have to inquire what an animal is, what rationality is, and in this one question would lead me down the slope to other harder ones (...)*" (Cottingham, Stoothoff, & Murdoch, 1984).

Nowadays, with the background of modern sciences, and especially with the advent of cognitive science, new questions and perspectives have been added to the question for the nature of rationality. Instead of merely studying properties and features of its manifestation in humans and their behavior, and measuring those against standards for rationality derived from normative theories, prediction and modeling aspects of theories of rationality are gaining more and more importance. These efforts can take various forms, ranging from formal studies of customer preferences and behavior, over cognitive modeling of decision-making and choice processes, to rational agency projects within AI.

Nonetheless, although this shift of emphasis within the study of rationality is clearly happening, it is a change of priorities within the existing overall setting, but not a revolutionary process overthrowing existing paradigms, or creating new approaches and ideas. In the following, we want to argue that the latter is what would be needed, on the one hand for real progress with respect to the aforementioned modeling and prediction tasks, but on the other hand also for the sake

of a deeper understanding and progress within rationality research itself.

The paper is structured as follows: The next section gives an overview of classical accounts and paradigms within rationality research, together with well-known objections and caveats to these standard frameworks. Then, a review of recent empirical and theoretical findings is presented, indicating a way to a new understanding and modeling of rationality and rational behavior. This new conception, for the time being called "subject-centered rationality", is sketched and elaborated on in a dedicated section. The penultimate section discusses basic features of this new stance in a juxtaposition of (expected) standard objections and reservations with solutions and answers to the former. Concludingly, the paper is put into a bigger context within the respective fields.

Standards in Rationality

In the following, when talking about rationality, we mostly want to refer to the manner in which people derive conclusions when considering things deliberately in the domain of individual problem-solving, also including the conformity of one's beliefs with one's reasons for those beliefs or of one's actions with one's reasons for those actions. In consequence, at least for our considerations, beliefs and knowledge (which are also seen as presuppositions for respective actions), that is the epistemic aspects of rationality, are our main concerns. Within this context, rationality is intrinsically connected to an optimality principle, making a decision rational if it is not just reasoned, but if it is also in certain ways optimal for achieving a goal or solving a problem.

Several centuries of thought and investigation in relation to rationality and its manifestations in humans and their behavior resulted in the formation of mainly four abstract general models (most of which also bring along corresponding normative theories, and even definitions, of rationality):

- Logic-based systems (cf. e.g. (Evans, 2002)): A belief is rational, if there is a logically valid reasoning process to reach this belief relative to available/given background knowledge.
- Probability-based frameworks (cf. e.g. (Griffiths, Kemp, & Tenenbaum, 2008)): A belief is rational, if the expectation value of this belief is maximized relative to given probability distributions of background beliefs.
- Game theory-based models (cf. e.g. (Osborne & Rubinstein, 1994)): A belief is rational, if the expected payoff of

maintaining the belief is maximized relative to other possible beliefs.

- Accounts based on the use of heuristics (cf. e.g. (Gigerenzer, Hertwig, & Pachur, 2011)).

Unfortunately, when comparing the different conceptions, it shows that the frameworks (and also the resulting definitions of rationality) are in many cases almost orthogonal to each other, making them in the best case incommensurable, if not inconsistent or even partly contradictory in their modeling assumptions. Also, in many cases the predictive power of these classical theories of rationality turns out to be rather limited (at least when applied to real-world examples instead of artificially simplified and constructed scenarios), as they have more normative or postdictive-explanatory character.

Even more, although each of the listed accounts has gained merit in modeling certain aspects of human rationality, the generality of each such class of frameworks has at the same time been challenged by psychological experiments or theoretical objections. On the one hand, studies by Byrne on human reasoning with conditionals (Byrne, 1989) indicated severe deviations from classical logic (see below), a finding also supported by human subjects failing at a seemingly simple logical task in the famous Wason-selection task (Wason, 1966). Similarly, Tversky and Kahneman's Linda problem (Tversky & Kahneman, 1983) illustrates a striking violation of the rules of probability theory. On the other hand, game-based frameworks are questionable due to the lack of a unique concept of optimality in game-theory which could possibly support different "rational behaviors" for one and the same situation (just think of the plethora of different equilibrium concepts which have been derived from the original Nash equilibrium, cf. e.g. (Halpern, 2008)). Finally, heuristic approaches to judgment and reasoning (Gigerenzer, 2008) are often seen as approximations to a rational ideal and in some cases could work in practice, but mostly lack formal transparency and explanatory power. Also, from a methodological or philosophical point of view, severe reservations can be put forward: Due to the open nature of the collection of heuristics (the "heuristic toolbox") propagated in most current accounts, the possibility of falsification and refutation of modeling assumptions and theories is not guaranteed (as always another heuristics could be introduced, covering cases previously not accounted for), and a (reasonable) completion of the model can neither be checked for, nor guaranteed at any point.

A More Cognitive Perspective

Already starting out with Simon's seminal work on "*A Behavioral Model of Rational Choice*" (Simon, 1955) more than half a century ago, conceptually different takes on rationality and rational behavior have been introduced into the discussion. Where logic-based, probability-based, and game-based accounts normally do not take into account limitations and cognitive properties of the reasoner (i.e., in our case, the human agent), the last decades have seen a growth of awareness for the importance and indispensability of these factors

in models of (human) rationality. Nonetheless, from our point of view, understanding the full meaning and implications of the commitment to a more "cognitively adequate" theory of rationality still is in an early stage, with researchers in different disciplines, both on the more theoretical and the empirical side, only starting to fully integrate these concepts into their accounts of rationality and rational behavior. In the following, we want to sketch some important developments and insights, preparing the ground for the subsequent presentation of our account of "subject-centered rationality".

Theoretical Considerations

In (Simon, 1955), Simon articulates a simple but groundbreaking insight: (Human) agents are bounded in their resources and computational power. Once accepted, this proposition has far reaching consequences for the entire conceptual endeavor of formalizing and modeling rationality. All of a sudden, internal limitations of the reasoner, like limitations on working memory and computational power, but also external constraints, like limited time for decision-making, or incomplete and possibly also false information, have to be accounted for when creating a framework for rationality. Where this might not necessarily be a problem at first sight, it casts more than just a slight shadow of doubt on some of the fundamental assumptions underlying many "classical" accounts of rationality: It might be the case that reaching a conclusion via *modus ponens* in a logic-based model, once the preconditions are fulfilled, can computationally be realized rather simple - but what if the reasoner has to deal with incomplete information? If expectations have to be maximized in a probability-based model, where actually do the priors come from, once their existence cannot just be imported as given by the modeling assumptions? For many game-theoretic settings, even modern computers have a hard time finding equilibria or optimal solutions to not overly complex problems - how then can a mere human take the corresponding decision within a split second?

Although these issues clearly were identified as urgent questions, subsequent attempts at solutions mostly tried to solve the appearing problems within the original models, instead of putting the foundations to the proof and (possibly) having to construct entirely new models of rationality. Nonetheless, this has changed over the course of the last years. In a reply to Colman's article "*Cooperation, psychological game theory, and limitations of rationality in social interaction*" (Colman, 2003), Kokinov challenges traditional views on rationality (Kokinov, 2003). Taking an initial stance similar to Colman's, that is agreeing that rationality fails as both, descriptive theory of human decision-making and normative theory for good decision-making, Kokinov reaches a different, more radical conclusion than Colman did before. Instead of trying to fix the concept of rationality by redefining it, adding formerly unconsidered criteria for optimization of some kind, he proposes to replace the concept of rationality as a theory in its own right by a multilevel theory based on cognitive processes involved in decision-making. Where

Colman proposes a collection of ad-hoc strategies for explaining the deviations from rationality which people exhibit in their behavior, Kokinov proposes analogy as means of unifying the different, formerly unconnected parts of Colman's attempt at describing the mechanisms of decision-making. In Kokinov's view, the classical concept of utility making has to be rendered as an emergent property, which will emerge in most, but not all, cases, converting rationality itself into an emergent phenomenon, assigning rational rules the status of approximate explanations of human behavior.

Of course, this also defines a rather extreme position. But also scholars in "more standard" disciplines of rationality research, as for example decision theory, have become aware of fundamental problems with the traditional conception of rationality. In (Gilboa, 2010), Gilboa (in accordance with an earlier definition in (Gilboa & Schmeidler, 2001)) defines rationality as a subjective concept: "(...) *a mode of behavior is rational for a given decision maker if, when confronted with the analysis of her behavior, the decision maker does not wish to change it.*". This of course has far-reaching consequences. First of all, rationality becomes subject-centered, in that what is rational might vary with the population in question. But probably the most important implication is the dependence on the individual subject's abilities and limitations. If the decision maker does not understand the analysis, or why her behavior is not judged as rational, she cannot be judged irrational for not complying with the alleged norm of rationality. If limited cognitive capacities do not allow the reasoner to understand the rules he should follow in his reasoning (e.g. the Neumann-Morgenstern theory (von Neumann & Morgenstern, 1944)), but would always take the same decision again, he has to be called rational.

Experimental Evidence

But also on the more empirical side, there is evidence galore that cognitive capacities and limitations have a clear influence on behavior and decision-making. Evidence for a crucial role of analogy as cognitive ability in decision-making can be found in psychological studies on decision-making and choice processes. An overview by Markman and Moreau (2001), based on experiments and observations from psychological studies (amongst others on consumer behavior and political decision-making), reaches the conclusion that there are at least two central ways how analogy-making influences choice processes. Analogies to other domains can provide means of representation for a choice situation, as generally speaking the making of a decision relies on a certain degree of familiarity with the choice setting. In many cases of this kind, analogy plays a crucial role in structuring the representation of the choice situation, and thus may strongly influence the outcome of a decision. Also, structural alignment (a key process of analogy-making) plays a role when comparing the different possible options offered by a decision situation, with new options being learned by comparison to already known ones. An experimental study by Kokinov (2005) demonstrated that people actually do use analogies in the process

of decision-making, with significant benefit already if only one case is found to be analogous to the choice situation under consideration. Furthermore, evidence has been found that there is no significant difference between close and remote analogies in this process, and that people are not limited to rely only on analogous cases from their own experience, but that also cases which were only witnessed passively (e.g., by being a bystander, or learning about a situation from reports in the media) may have beneficial influence.

Another example can be given in form of well-known studies on human decision-making under time pressure, which show a change in the applied inference procedure. In (Rieskamp & Hoffrage, 2008), the authors report that, whilst the best predicting model of human inference for decision making in an unstressed condition was a weighted linear model integrating all available information, when time pressure was induced, best predictions were obtained by using a simple lexicographic heuristic (Fishburn, 1974). When speculating about the precise way in which the induced pressure influences the reasoning process, the presumed change from a more complex strategy using complex relational structures to a simple single-attribute-based procedure, one possible explanation can again be found in research on analogy-making: In (Tohill & Holyoak, 2000), it is reported that anxiety made participants of an analogical-reasoning experiment switch from a preference for complex relational mappings to simple attribute-based mappings. So presupposing that analogy to already familiar situations serves as a basis for the decision-making, this reduction of the complexity of the mapping would be in line with the observed change in strategies.

Cornerstones of "Subject-Centered Rationality"

Continuing and expanding thoughts already started for example in (Besold et al., 2011) and (Besold et al., 2012), in this section, we want to present some key features and cornerstones of a new account and understanding of rationality which we call "subject-centered rationality". What shall be presented is not yet another model of rationality and rational behavior, but an overall view and meta-conception of rationality, defining a supporting and limiting background and context for the construction of new models:

1. Rationality in a human context clearly has to be considered as a subject-centered notion, demanding for the integration of subject-related properties and constraints. This goes in line with Simon's and Gilboa's already discussed positions in that there is no use establishing models and norms which can never be implemented or fulfilled by human reasoners due to limitations of the agent or of its environment. Going back to the very beginnings, taking Aristotle's characterization of humans as rational animals as defining statement, a framework of rationality which cannot be applied at any moment by a human in everyday life has to be considered as limited in its usefulness and adequacy. This means that especially limiting constraints of human cognition, as for

example the computational boundedness of human agents and thus also the computational complexity of the rationality theory at hand, have to be accounted for.

2. The main aspect of theories and models of rationality has to be their use and applicability as positive theories, and not as mere normative or postdictive explanatory accounts. A valuable and adequate account of rationality has to provide a feasible prediction of human rational behavior and decision-making when being provided with the information the human reasoner can access at the moment of reasoning, that is, when possibly only having access to incomplete knowledge, when having to deal with ambiguities and possible false assessments of situations, etc.
3. A feasible theory of rationality does not have to be committed to one single formal modeling paradigm, but instead of being monolithic should pursue a holistic approach. Different formalisms and approaches to modeling should be unified and integrated into one account, providing the amalgamation of the different approaches with the union of their respective advantages and particular strengths, whilst mitigating each others deficiencies and weaknesses. Logic-based formalisms can be used alongside probability-based techniques alongside heuristic elements alongside models of cognitive core capacities alongside game-theoretic means for utility maximization. This integration of means and paradigms is governed by two guiding principles, the appropriateness with respect to the boundedness and the limited resources of the human reasoner, and the respect for and reproduction of human particularities in rational behavior and decision-making.

As should have become obvious from the above listing, we clearly consider rationality as a concept which is not connected to a particular formalism or theoretical modeling paradigm, but instead see it as possibly a plethora of different mechanisms competing, interacting and contributing to what we externally observe as a single capacity. Of course, this brings along challenges and questions, but from our point of view at the same time offers even more chances and opportunities. In the following, we want to give two examples where incorporating different (formal) paradigms into (distinct) traditional contexts has shown to be highly profitable:

With “algorithmic rationality”, Halpern and Pass (2011) proposed a framework including computational costs in otherwise game-theoretical notions of rationality, allowing to directly take into account the complexity dimension of agents’ boundedness when modeling rationality with game-theoretical means.

“Probabilistic dynamic epistemic logic” by Kooi (2003) presents an attempt at amalgamation between probabilistic and logical views of rationality, which can for example be used when treating with aspects of rationality at the intersection between epistemic logic and game theory.

Also, we strongly advocate the integration of factors and mechanisms describing the influence of cognitive capacities

and abilities into models of rationality, possibly offering entirely new perspectives and explanations for classical paradoxes of human rationality, as illustrated by the following examples:

Well-known empirical studies by Byrne (1989) question whether human reasoning can be covered by a classical logic-based framework. Presented with the information given in Table 1 and asked for what can be concluded from this, from 1. 46% of subjects conclude that Marian will not study late in the library, erring with respect to classical logic (as denial of the antecedent does not validate a negation of the consequent). Also, from 2. 96% of subjects conclude that Marian will study late in the library, whilst only 38% of subjects reach the same conclusion from 3.. Thus an introduction of another antecedent (without any indication that the antecedent should not hold) dramatically reduced the number of subjects applying a simple modus ponens in their process of forming a conclusion. Giving the task and the find-

Table 1: Inferences and Conditionals (Byrne, 1989)

1. If Marian has an essay to write, she will study late in the library. She does not have an essay to write.
2. If Marian has an essay to write, she will study late in the library. She has an essay to write.
3. If Marian has an essay to write, she will study late in the library. She has an essay to write. If the library stays open, she will study late in the library.

ings a more cognitive capacities-oriented view than supported by the logic-based framework, the results concerning conclusions drawn by the subjects can for example be explained through analogy. People faced with the information given in 1. will recall similar conversations they had before, using these known situations as basis for their decision on what to conclude. According to Grice (Grice, 1975), in conversations, speakers are supposed to provide the hearer with as much information as is needed for exchanging the necessary information, a rule which also goes in accordance with our everyday observation. Thus, when being given the additional information that “Marian does not have to write an essay.”, the set of candidate situations for establishing an analogy to the present one will be biased towards situations in which this information had an impact on the outcome, resulting in the conclusion that Marian would not study late in the library either. Regarding 2. and 3., a similar conjecture seems likely to hold: By additionally mentioning the library, similar situations in which the library might actually have played a crucial role (e.g., by being closed) will be taken into account as possible base domains of the analogy, causing the change in conclusions made.

Already more than a decade ago, in the field of decision theory and economics Gilboa and Schmeidler (1995) devel-

oped an (at least partly) case-based theory and model for decision-making under uncertainty. In their model, cases are primitive and provide a simple axiomatization of a decision rule that selects an act to be performed based on the act's past performance in similar cases. Each act is evaluated by the sum of the utility levels that resulted from using this act in past cases, where the degree of (dis)similarity between the past cases and the problem at hand is accounted for by weighting the respective utility by the value of a similarity measure between both situations. Remarkably, this formal approach in a natural way gives rise to (amongst others) the notions of satisficing decisions and aspiration levels (cf. also (Gilboa & Schmeidler, 2001) for a detailed account).

Solutions for Problems of a Subject-Centered Notion of Rationality

Of course, a positive subject-centered approach to rationality brings along quite some ground for reservations and skepticism. In the following, we want to address some of the most probable objections.

Doesn't your account collapse into pure subjectivity, making it not usable as basis for a general theory and framework anymore? No, it does not. The idea to overcome the problem of pure subjectivity is to identify central cognitive mechanisms, limitations and properties, common to all humans, and use those as basis for building up the theory. This is not a commitment to any particular modeling paradigm, but rather a strong statement assigning the positive aspects of the model (i.e. the adequate reflection of actual human properties) a higher priority than detail decisions for how to model a particular capacity, or overall methodological modeling consistency.

But then, how do you want to (theoretically) explain and (practically) preserve the individual aspects of your subject-centered modeling notion? In real life, each human comes equipped with different gifts and talents. Distinct cognitive capacities, although present in (almost) every human, are developed to a different extent. This also gives a cognitively-based model of rationality the possibility to account for different individual behaviors and decisions. Although there is an overall unified framework on a general human level, once reliable predictions shall be made on an individual scale, the weights and levels of developments between these different capacities will have to be assessed and adapted.

Your account does not provide any normative power, a key feature one would expect from a theory of rationality? To the contrary, normativity can be introduced on two levels, allowing for a distinction between a more general form of rationality, and a subjective one. Even more, contrary to some of the classical approaches, it is even possible to provide a quantitative account of performance on a normative scale. On a higher level, normativity can be introduced via the question "Given the general cognitive mechanisms and model, how well does the individual perform compared to the general optimal case?", that is, by assessing whether, tak-

ing into account goal-oriented behavior, the chosen way of acting resulted in the best outcome possible for an ideal representative of the species. If this is the case, the respective behavior or decision has to be considered as generally rational. If this is not the case, comparing the quality of the outcome to the performance of other individuals of the species can provide a quantification of the quality of the decision or behavior on a normative scale (e.g. "the subject performed at least as rational as 80% of his species"). On the individual level, normativity can be introduced via the question "Given the individual distribution of properties and limitations of the individual, how well does the subject perform compared to the individually optimal case?", that is, by assessing whether, taking into account goal-oriented behavior, the chosen way of acting resulted in the best outcome possible for the individual (a notion reminiscent of Gilboa's already aforementioned idea of rationality (Gilboa, 2010)). Also here, similar to the more general case, a quantitative aspect can be introduced to judging a behavior or decision rational: Provided that the model has accurately been fitted to the individual, reflecting its cognitive capacities, properties and limitations to a sufficient degree, it can be assessed to what extent the subject made use of its theoretical capacities.

Conclusion

In the present paper, we give an account of basic principles and cornerstones of our positive conception of a theory and framework for rationality and rational behavior, envisioning a modeling paradigm which integrates different perspectives and approaches into a holistic system, giving rise to an integrated multiverse of rationality, replacing the multiple (mostly) mutually exclusive competing universes which there currently are.

From our point of view, our perspective on rationality offers several advantages not only within the field of cognitive science, but also for neighboring disciplines. A positive, predictively usable theory and framework for rationality (moreover if equipped with a quantitatively accessible notion of normativity) would allow for manifold applications, for example within decision theory and psychology (serving as an initial test bed for conjectures and research hypotheses), but also in more technical fields such as human-computer interaction (allowing for more natural and better adapting interfaces between man and machine) or artificial intelligence (greatly contributing to an overall model of human intelligence; cf. e.g. (Besold, 2011) for an AI-centric perspective).

Of course, there still are numerous open questions left for future investigation: How can the given cornerstones, characterizations and properties of "subject-centered rationality" be developed into a completely worked out meta-theory? Which are the key particularities, properties and limitations of human cognition that have to be integrated and accounted for by a subject-centered theory of rationality? How compatible are the already existing "classical" frameworks for rationality with our proposed view? How do our meta-level considera-

tions relate to conceptual work done in other relevant fields modeling (and possibly predicting) human behavior, and thus most likely dealing with similar questions? What would be a promising paradigm for an implementation: a highly modular bottom-up approach starting out by modeling one facet of rationality, consecutively adding more modules later, an entirely hybrid top-down approach, applying an amalgamated broad mixture of formalisms and a very general modeling paradigm from the very beginning, addressing different forms and facets of rationality by specialization within the overall framework, or something in between?

Although these are challenging and demanding questions, and a complete answer to any of those still is far from being visible, we are convinced that each single one of them is worth scientific effort and attention already by itself, in their totality moreover promising key insights into a core concept of human intelligence and cognition.

References

- Besold, T. R. (2011). Rationality in/through/for AI. In J. Romportl, P. Ircing, E. Zackova, R. Schuster, & M. Polak (Eds.), *Proc. of Extended Abstracts of Beyond AI 2011*.
- Besold, T. R., Gust, H., Krumnack, U., Abdel-Fattah, A., Schmidt, M., & Kühnberger, K. (2011, July). An Argument for an Analogical Perspective on Rationality & Decision-Making. In J. van Eijck & R. Verbrugge (Eds.), *Proc. of the Workshop on Reasoning About Other Minds (RAOM-2011)*. CEUR-WS.org, Vol. 751.
- Besold, T. R., Gust, H., Krumnack, U., Schmidt, M., Abdel-Fattah, A., & Kühnberger, K.-U. (2012). Rationality Through Analogy - Towards a Positive Theory and Implementation of Human-Style Rationality. In *Proc. of MATHMOD 12 Vienna*. (to appear)
- Broadie, S., & Rowe, C. (Eds.). (2002). *Aristotle Nicomachean Ethics: Translation, Introduction, and Commentary*. Oxford University Press.
- Byrne, R. (1989). Suppressing valid inferences with conditionals. *Cognition*, 31(1), 61–83.
- Colman, A. M. (2003). Cooperation, psychological game theory, and limitations of rationality in social interaction. *Behavioral and Brain Sciences*, 26(2), 139–198.
- Cottingham, J., Stoothoff, R., & Murdoch, D. (Eds.). (1984). *The Philosophical Writings of Descartes* (Vol. II). Cambridge University Press.
- Evans, J. (2002). Logic and human reasoning: An assessment of the deduction paradigm. *Psychological Bulletin*, 128, 978–996.
- Fishburn, P. (1974). Lexicographic orders, utilities and decision rules: A survey. *Management Science*, 20, 1442–1471.
- Gigerenzer, G. (2008). *Rationality for Mortals: How People Cope with Uncertainty*. Oxford University Press.
- Gigerenzer, G., Hertwig, R., & Pachur, T. (Eds.). (2011). *Heuristics: The Foundation of Adaptive Behavior*. Oxford University Press.
- Gilboa, I. (2010). Questions in Decision Theory. *Annual Reviews in Economics*, 2, 1–19.
- Gilboa, I., & Schmeidler, D. (1995). Case-Based Decision Theory. *The Quarterly Journal of Economics*, 110, 605–639.
- Gilboa, I., & Schmeidler, D. (2001). *A Theory of Case-Based Decisions*. Cambridge University Press.
- Grice, H. P. (1975). Logic and Conversations. In P. Cole & J. L. Morgan (Eds.), *Syntax and Semantics, Vol. 3: Speech Acts* (pp. 41–58). Academic Press.
- Griffiths, T., Kemp, C., & Tenenbaum, J. (2008). Bayesian Models of Cognition. In R. Sun (Ed.), *The Cambridge Handbook of Computational Cognitive Modeling*. Cambridge University Press.
- Halpern, J. Y. (2008). Beyond Nash Equilibrium: Solution Concepts for the 21st Century. In *Proc. of the 27th Annual ACM Symposium on Principles of Distributed Computing*.
- Halpern, J. Y., & Pass, R. (2011, June). Algorithmic Rationality: Adding Cost of Computation to Game Theory. *ACM SIGecom Exchanges*, 10(2), 9–15.
- Kokinov, B. (2003). Analogy in decision-making, social interaction, and emergent rationality. *Behavioral and Brain Sciences*, 26(2), 167–169.
- Kokinov, B. (2005). Can a Single Episode or a Single Story Change our Willingness to Risk? The Role of Analogies in Decision-Making. In B. Kokinov (Ed.), *Advances in Cognitive Economics*. NBU Press.
- Kooi, B. P. (2003). Probabilistic Dynamic Epistemic Logic. *Journal of Logic, Language and Information*, 12, 381–408.
- Markman, A., & Moreau, C. (2001). Analogy and analogical comparison in choice. In D. Gentner, K. Holyoak, & B. Kokinov (Eds.), *The Analogical Mind: Perspectives from Cognitive Science* (pp. 363–399). MIT Press.
- Osborne, M., & Rubinstein, A. (1994). *A Course in Game Theory*. MIT Press.
- Rieskamp, J., & Hoffrage, U. (2008). Inferences under time pressure: How opportunity costs affect strategy selection. *Acta Psychologica*, 127, 258–276.
- Simon, H. A. (1955, February). A Behavioral Model of Rational Choice. *The Quarterly Journal of Economics*, 69(1), 99–118.
- Tohill, J., & Holyoak, K. (2000). The impact of anxiety on analogical reasoning. *Thinking & Reasoning*, 6(1), 27–40.
- Tredennick, H. (Ed.). (1933–35). *Metaphysics*. Harvard University Press.
- Tversky, A., & Kahneman, D. (1983). Extensional versus intuitive reasoning: The conjunction fallacy in probability judgement. *Psychological Review*, 90(4), 293–315.
- von Neumann, J., & Morgenstern, O. (1944). *Theory of Games and Economic Behavior*. Princeton University Press.
- Wason, P. C. (1966). Reasoning. In B. Foss (Ed.), *New Horizons in Psychology*. Penguin.