Synthese (2018) 195:529–547
https://doi.org/10.1007/s11229-014-0584-1

Comparing the axiomatic and ecological approaches to rationality: fundamental agreement theorems in SCOP

Patricia Rich

Received: 29 November 2013 / Accepted: 20 October 2014 / Published online: 7 November 2014

Abstract There are two prominent viewpoints regarding the nature of rationality and how it should be evaluated in situations of interest: the traditional axiomatic approach and the newer ecological rationality. An obstacle to comparing and evaluating these seemingly opposite approaches is that they employ different language and formalisms, ask different questions, and are at different stages of development. I adapt a formal framework known as SCOP to address this problem by providing a comprehensive common framework in which both approaches may be defined and compared. The main result is that the axiomatic and ecological approaches are in far greater agreement on fundamental issues than has been appreciated; this is supported by a pair of theorems to the effect that they will make accordant rationality judgements when forced to evaluate the same situation. I conclude that ecological rationality has some subtle advantages, but that we should move past the issues currently dominating the discussion of rationality.

Keywords Rationality · Ecological · Axioms · SCOP · Decision-making · Inference

1 Introduction

The realization in the last century that observed human behavior often deviates from the predictions of traditional axiomatic theories of rationality (henceforth "TATs") (e.g. Allais 1953; Ellsberg 1961), has led to a revival of the philosophical problem of developing a credible account of rationality. In particular, the concept of ecological rationality ("ECO") has been developed as a competitor to TATs. One obstacle to progress is that the the competing approaches use different language, assumptions,

P. Rich (✉)
Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA 15213, USA
e-mail: prich@andrew.cmu.edu

and formalisms; furthermore, they have been developed to different degrees of sophistication, as ECO is relatively new. This makes it difficult to compare the approaches in order to judge what improvements, if any, ECO offers over TATs. To solve this problem, I adapt a very general and abstract formalism, known as SCOP (Aerts 2002), that can provide a neutral, comprehensive representation of human reasoning as our minds interact with the outside world; I then use this formalism to define, compare, and evaluate our competing approaches to rationality. The main result is a pair of agreement theorems to the effect that TATs and ECO make accordant rationality judgements if we hold fixed the situation that they must evaluate and recognize that utility is fundamentally the same for both approaches. I use this analysis to argue that their disagreements are largely superficial, but that ECO improves on TATs in some subtle ways.

Section 2 explains the problem in more detail. Section 3 explains the solution—adapting SCOP as a common formal framework for studying theories of rationality—and introduces this framework, which is then used to define TATs and ECO. Section 4 provides the results: agreement theorems are presented and proved, these theorems are interpreted, and TAT and ECO are assessed in light of their definitions in SCOP. Section 5 concludes.

2 The problem

2.1 The general problem of developing an account of rationality

Having a satisfactory philosophical account of individual rationality is important because such accounts have implications for real-world problems and are used as guides for behavior. We want theories to deliver judgements that we can put to good use, telling us what to praise and blame, and therefore how to make normative improvements. This criterion of usefulness has further implications, for example that rationality must be sensitive to the facts of our existence and in some sense attainable. Interestingly, the requirement that our theories be useful in this way offers strong support for the view that they must also be *intuitively accurate*, in the sense of according with and making precise our pre-theoretic notion of rationality. This is because our pre-theoretic notion is strongly internalist, in that rationality is understood to be about the suitability of the means one employs for one's chosen ends, given one's beliefs; in other words, rationality is tied up with justification. A useful theory will need to retain these features, as we humans lack the ability to base our behavior directly on objective facts, and must use our beliefs about that reality as a proxy (I take this to be the moral of Paul Weirich's attempt to provide a "realistic decision theory" (Weirich 2004)).[1]

[1] One could have other goals in studying decisions and inferences, but from a normative, philosophical perspective, and when our theories are theories of *rationality* for individual people and not just predictive devices, these are our primary purposes. A satisfactory theory, therefore, must fulfil these functions.

2.2 The particular problem of comparing theory types

In the domains of decision-making and inference in particular, there are two types of normative theories competing for preeminence: TATs, which give "context-independent" properties that a decision or inference must have in order to be rational; and the newer ECO, which judges decision or inference processes as more or less rational only relative to a particular context[2] (Gigerenzer et al. 2011). I take the original expected utility theory ("EUT") (von Neumann and Morgenstern 1944) to be the paradigmatic TAT, with the AGM theory of belief revision (Alchourron et al. 1985) and axiomatizations of logical inference rules as other important examples. I take ECO to be the conception of rationality articulated and defended by the Center for Adaptive Behavior and Cognition (henceforth "ABC") of the Max Planck Institute (Gigerenzer and Selten 1999; Gigerenzer et al. 2011); this conception is based in Herbert Simon's work (e.g. Simon 1956). The natures of these theory types will be explained in more detail when they are defined formally in the next section.

The TATs are formal, and their philosophical proponents thus have very precise definitions of rationality to defend. This is not yet the case for ECO; it was born as a counter to TATs, and while its advocates study particular decision and inference processes (namely, computable heuristics) with mathematical rigor, the approach itself has not been specified with the precision, and attendant clarity, of the older accounts. While this is understandable for an emerging contender, it is an obstacle that must be overcome if ECO is to serve the normative purposes for which TATs are currently employed.

3 Addressing the problem

3.1 The method: appealing to a more general framework

3.1.1 Introducing SCOP

My method for addressing the philosophical problem of comparing and evaluating these approaches to rationality, despite their differing degrees of development and formality, is to move up a level in generality to yet another framework in which both can be defined and compared. Introducing a new formalism can help us by offering a common language with which to express our competing approaches; providing a complete and holistic representation of the activities we are interested in (mental activities, such as reasoning, inferring, and deciding); forcing us to clarify our accounts as we work to define them within this more complete space of possibilities (stating explicitly what is assumed to be held constant or irrelevant); and thus helping us to really "see" into the natures of the approaches. For reasons explained below, I take the SCOP formalism to be well-suited to these purposes.

There has been a trend towards cross-over from research in physics to the "softer sciences," with both abstract ideas of quantum mechanics and the mathematics devel-

[2] Note that these definitions do not translate literally into the formal framework; c.f. Sect. 3.

oped to accommodate it being borrowed and adapted for diverse other purposes. The state context property (SCOP) formalism is a part of this tradition. A main reason why new mathematical tools (such as quantum probability theory) are needed is that in quantum mechanics physical entities undergo state changes when exposed to a measurement context; the resultant probability distributions violate the laws of classical probability, therefore, because of “fluctuations on the interaction between the measuring apparatus and the physical entity under consideration” (Aerts 2002). Classical tools are thus unable to model quantum phenomena. In parallel, researchers outside of physics have turned to the newer quantum tools when classical techniques left problems and paradoxes (Bruza et al. 2009).

John von Neumann—who also provided the expected utility axioms (von Neumann and Morgenstern 1944) that ground much of modern economics and appear in this paper as the paradigmatic TAT—first axiomatized the Hilbert space in 1932 to provide foundations for quantum mechanics (von Neumann 1955). Despite the tremendous importance of this work for quantum mechanics and mathematics these foundations were found to be inadequate in certain respects, and attempts to provide improved foundations ensued (e.g. Birkhoff and von Neumann 1936).

Diederik Aerts, the father of SCOP, was actively involved in this work; specifically, Aerts sought a more general theory that would subsume both classical and quantum mechanics as special cases (see Aerts 1982, 1983 for early examples). In 1999, Aerts presented state property systems which model entities and the experiments that may be performed on them (Aerts 1999). Influenced by developments in Liane Gabora’s dissertation, state property systems evolved into SCOP, which introduces “context as a fundamental concept” with states and contexts influencing each other (Aerts 2002).

The basic structure of a SCOP system is as follows: There is an *entity*, whose changes we wish to model; this entity is characterized by its set of properties, which are either *actual* (akin to activated) or only *potential* at any given time, but are attached to the entity in the sense that the entity always has the same set of properties. A *state* of the entity is characterized by the set of properties that are actual at that time. Critically, an entity in a SCOP system changes states due to the influence of the external context, and the context can also change due to the influence of the state. So change within the system is modeled as the result of transition functions taking state-context pairs to new state-context pairs, and these transitions are probabilistic in nature (and potentially non-classically so) (Aerts 2002).

With Gabora, SCOP was initially used to model concept combination, resolving puzzles in traditional views (Gabora and Aerts 2002; Aerts and Gabora 2005) this continues to be a major focus of their research (see Aerts et al. (2013) for a recent paper that puts much of this work into context).[3] Yet SCOP’s applications are not limited to concepts: Aerts notes the move from micro-physical to macro-physical to

[3] This data-driven modeling work plays a supporting role in the “ecological theory of concepts” that the authors now develop, according to which concepts “only occur as part of a web of meaning provided by both other concepts and by interrelated life activities” (Gabora et al. 2008). This is similar to ecological rationality in that both take something long understood as purely internal and argue that it is instead a relation between internal and external, specifically in that context plays a critical role. We can understand this as part of a trend towards naturalism (discussed by Philip Kitcher in Kitcher (1992)). The views are primarily distinguished by the fact that the ecological theory of concepts seeks a *descriptive* account of

non-physical uncovering of quantum phenomena, and argues that SCOP is general enough that the entity modeled could be almost anything, "for example ... a cat, or a genome, or a cultural artifact, such as a building, or an abstract idea, or a mind of a person, or a stone, or a quantum particle, or a fluid etc. ..." (Aerts 2002). This is reflected in my own use of SCOP where indeed the entity modeled is a human mind, and my purpose quite different from the modeling of concepts.

Interestingly for my purposes, quantum-borrowing has become a productive area in decision research generally, evinced by a 2009 special issue on quantum cognition in the Journal of Mathematical Psychology (Bruza and Gabora 2009). Aerts himself has argued for a connection between the reasoning biases and fallacies uncovered by Kahneman and Tversky (e.g. Tversky and Kahneman 1974) and the Allais (1953) and Ellsberg (1961) paradoxes in economics, on the one hand, and the quantum nature of our handling of concepts, and thus non-classical thought processes, on the other (cf. Agrawal and Sharda 2010; Aerts and D'Hooghe 2009). He also suggests that SCOP could be useful for modeling decisions, and indeed previously used quantum-inspired ideas (with Sven Aerts) to offer an explanation for data from a psychological study in which participants must respond to controversial survey questions, regarding for instance nuclear energy (Aerts and Aerts 1995). In this last case, the presence of a "measurement effect" should be intuitive; being asked a question on a complex moral topic could easily influence the state of a person's mind, and perhaps even play a role in determining their opinion.

The cognitive science project of Aerts and Gabora is *bottom-up* and *descriptive* in nature; they are interested in modeling human thought, beginning with *particular* cases of manipulation of concepts, the presumed building blocks of thought. The value in this lies primarily in explanation: a compelling model of thinking can serve as an explanation for the patterns we see in human thought. In contrast, my philosophical project is very much *top-down* and *normative*: It is normative because the ultimate goal is to understand the nature of rationality itself, and infer from this how human mental activities *should* be. It is top-down because I am committed to the view that rationality is a natural kind and a unified phenomenon, i.e. that particular types of rationality (rational inference, decision, belief revision, etc.) are so because they instantiate the more general attributes of rationality itself. (This suggests that conclusions regarding the relative merits of axiomatic and ecological approaches to rationality in a particular domain cannot safely be drawn in the absence of more general conclusions about the adequacy of these approaches.)

However, the bottom-up and top-down approaches ought meet in the middle, and the descriptive and normative projects ought jointly produce *evaluations* of human mental activities, i.e. assessments of the rationality of our actual decisions, inferences, etc. This suggests that it will prove advantageous for both approaches to share a common framework and language for representing their shared subject matter, and serves as a reminder that any normative project should keep sight of the realities of the activities it seeks to judge so as to remain relevant. Given the successes of descriptive work

Footnote 3 Continued
concepts specifically, while ecological rationality (for the purposes of this paper) is a *normative* account of what it means to perform well on a broad range of decision and inference tasks.

using quantum principles in general and SCOP in particular, and given the intuitive accuracy of SCOP systems as representations of human mental activities, this makes SCOP a promising formalism for our top-down, normative endeavors. At the top, philosophical level, the formalism will primarily be of conceptual assistance, while Aerts' work shows that as we move down levels it may be used for more concrete computations.

3.1.2 The mental state interpretation of SCOP

A *state context property system* as defined by Aerts[4] is a tuple $(\mathcal{S}, \mathcal{M}, \mathcal{L}, \mu, \xi)$, with all components as below (Aerts 2002). The subject of a particular SCOP system is the entity represented in it, and which undergoes changes as it interacts with contexts. Interpreted for the purpose of studying rationality, this entity is to be a human mind. Let us call the entity E. At every time-point t, E is in some state $s_t \in \mathcal{S}$, where $\mathcal{S}$ is the set of all states the mind can be in. The state denoted $0 \in \mathcal{S}$ represents the end state, or the mind's death. The significance of the death state is that the entity has undergone (property) changes so radical that it is no longer considered to be the same entity.

E has a set of properties, $\mathcal{L}$; these properties are taken to be attached to E, and so they do not change (although whether a given property *actually* obtains will change). The interpretation of these properties is that they are the features that can obtain of the mind, such as happiness or satisfaction, having a goal, believing a proposition, etc. At the meta-theoretic level of this paper, the particular interpretation is not very important.

Change occurs—the entity is said to be in a new state rather than its previous one—when there is a change in which properties are *actual* and which only *potential*. A state s is then described by the set of properties $\mathcal{L}_s \subset \mathcal{L}$ that are actual for E in s. Where $\mathcal{P}$ is the power set, let the function $\xi : \mathcal{S} \rightarrow \mathcal{P}(\mathcal{L})$ be defined such that $\xi(s) = \mathcal{L}_s$. The function ξ is called the "Aristotle map."

The system also includes a set $\mathcal{M}$ of contexts that the entity could encounter. Any relevant property not in $\mathcal{L}$ is represented as belonging to a context. This makes contexts external to the mind in that they are not properties of the mind itself, but since the contexts interact with the mind they may bring about within the mind (for example) representations of, thoughts about, or responses to their content. So the agent's physical state (e.g. an injury), the physical environment, or facts about other agents with whom our agent is interacting would all be represented as part of context.

A change from one state-context pair (s_1, m_1) to another (s_2, m_2) is called a transition induced by context m_1. Given two such pairs, $\mu : \mathcal{S} \times \mathcal{M} \times \mathcal{S} \times \mathcal{M} \rightarrow \mathcal{P}([0, 1])$ is the function giving the probability of the first pair inducing the second pair. This probability is given as a subset of the interval [0,1] rather than a singleton, as in standard probability; it should be required that a pair (s, m) always induces another (s', m'), but how this is best accomplished formally remains an open question (Aerts 2002). The interpretation of these transitions is that when the mind is faced with a context, that context causes changes in the mind via some process (e.g. re-orienting itself towards a

[4] I will stay close to Aerts' notation for convenience, but note that there are several changes.

new goal, reasoning to a new conclusion), and the mind's state can similarly alter the context (e.g. changing the position of a game or causing the body to move). We might interpret the probabilistic nature of the transitions either as the result of indeterminacy in the processes producing state changes, or as representing our uncertainty. Both may be useful in some circumstances (Aerts 2002).

3.1.3 Observations

Looking at the framework interpreted as a human mind, some features of our reasoning as represented here are worthy of note. For one, we have historically found it convenient to draw a sharp distinction between decisions and inferences, but these activities are really of the same sort: both are intended to produce a mental choice—either about what is true, useful, or the best course of action—producing in turn some external effect, which can be better or worse given its purpose.

For two, we can zoom in or out on the processes, unpacking one transition into many, or combining many into one, based on our needs. I see this as a virtue, because it is an inescapable fact that a lot goes on even when we make a simple choice—beliefs have to be consulted, the situation needs to be represented or characterized, goals determined, a process chosen, a process employed, and so forth—and it's harder to forget about the interconnected nature of our various decisions and inferences if they can all be represented in a single framework. This is a benefit of having a top-down representation of how our minds work, even when we want to think about very specific processes. Particular decision or inference processes are chosen by metaprocesses based (one hopes) on their suitability for the context and one's goals, and these metaprocesses are critical to how well we function (as pointed out by Wallin and Gärdenfors (2000)).

Thirdly, many decisions or inferences are made without any explicit reasoning process, hidden even from the agent's view. This presents a challenge for normative theories (and ECO in particular) that does not seem to have been recognized or faced, although the facts are well-known. What is the normative status of those hidden processes? Specifically, are they fair game for evaluation as to their rationality? Neither a simple "yes" nor "no" is appealing, but considering the mind as a whole suggests this to be a false dilemma: we can evaluate the metaprocess that determines whether or not a problem is dealt with by explicit reasoning.

3.2 Preliminary definitions in SCOP

Further preliminary definitions will be needed in order to define TATs and ECO within SCOP. Firstly, a slight modification of SCOP itself will offer a more realistic model of a human mind. Specifically, the mind's properties should be able to take real values (rather than being binary, actual or potential). This is achieved by using the following definition of a SCOP system:

Definition 1 (*SCOP system*) A SCOP system is a tuple $(\mathcal{S}, \mathcal{M}, \mathcal{L}, \mu, \xi)$ where:

- $\mathcal{S}$ is the set of possible states of the entity, with s_t the state at time t
- $\mathcal{M}$ is the set of possible contexts

- $\mathcal{L}$ is the (countable) set of properties that can obtain of the entity; these properties form a Hilbert space $H_{\mathcal{L}}$ such that each property $l \in \mathcal{L}$ is a closed subspace of $H_{\mathcal{L}}$
- μ is a function $\mu : \mathcal{S} \times \mathcal{M} \times \mathcal{S} \times \mathcal{M} \to \mathcal{P}([0, 1])$ giving the probability that any state-context pair arises from any other
- ξ is a function $\xi : \mathcal{S} \to H_{\mathcal{L}}$ defining states in terms of vectors in the Hilbert space, where each vector specifies the value assumed by each property in that state

We are pursuing the use of SCOP to model agents' mental processes, and so the subject of a SCOP system as described above is an agent's mind, the entity E. At any given time t, s_t completely describes the current state of E. The SCOP system, then, describes the entity, the contexts it might face, and the changes of both the entity and the context as a result of interactions between them.

Definition 2 (*Problem*) A problem is a function $P : \mathcal{S} \times \mathcal{M} \to \mathcal{V}$, which takes a state-context pair and relates it to a valuation over outcomes, hence defining tasks or goals.

The idea here is that rationality judgements are made when agents perform tasks; they are faced with some problem to be solved or some goal to be reached, and choose a method of responding to it. A situation—an initial state and context—therefore gets interpreted in terms of the task the agent is thought to be faced with in that situation. Tasks give rise to valuations, i.e. assignments of value to states of affairs, because the essence of a task is to identify circumstances that are to be brought about, i.e. to induce a new valuation of possible states of the world, where solutions to the task are valued, and better solutions more so.

Definition 3 (*Valuation*) A valuation $v \in \mathcal{V}$ is a function $v : \mathcal{S} \times \mathcal{M} \to \mathbb{R}$, which associates resultant state-context pairs with real numbers representing their relative goodness.

Whereas a problem takes an initial situation to a valuation, the valuation itself takes the situations that could result and assigns them values. In typical applications such a valuation will only categorize outcomes on the basis of some small number of features relevant to the problem, and hence be "short-sighted" in the sense of only caring about the achievement of immediate goals, like getting money or answering a question correctly. This simplifies the task of specifying v for particular applications considerably. However, we can conceive of valuations more broadly: a bigger-picture valuation would be required when considering, for example, an agent's life goals.

Definition 4 (*Process and algorithm*) A process φ is a method for accomplishing some task, i.e. for generating some appropriate output in response to an input. It must be carried out by some algorithm, which we call A^{φ}. Then let $A^{\varphi} : \mathcal{S} \times \mathcal{M} \times \mathcal{S} \times \mathcal{M} \to \mathcal{P}([0, 1])$ be the function describing the effects of implementing A^{φ}, by specifying how probable it is that any given state-context pair arises from an initial one.

Note that conditions must be placed upon this function to ensure that these probabilities are coherent, as with the transformation function μ. No confusion should arise in this setting from using the same label for the algorithm itself and the function describing its effects.

The reason for distinguishing processes from their implementation as algorithms is that speed and efficiency are important criteria of evaluation for processes according to ECO, and TATs may take them into account as well. It is possible that processes are multiply realizable by algorithms that vary in terms of speed and efficiency (or even accuracy), and so distinguishing processes from algorithms safeguards against judging processes according to features that depend on their implementation.

3.3 Traditional axiomatic rationality in SCOP

Definition 5 (*Axioms, axiom sets, and axiom complexes*) For $i \in \mathbb{N}$, let any given $Ax_i = \{ax_{i1}, ax_{i2}, ..., ax_{in}\}$ be a set of individual axioms ax_{ij}, such that the axioms of Ax_i – also known as an *axiom system*—are jointly taken (by some possible agent) to constitute the rules of rationality for some domain. Then let $\mathcal{A}x$ be a collection of axiom systems, which will also be referred to as an *axiom complex*.

These definitions reflect and clarify the fact that the TAT approach carves rationality up into domains, each of which is governed by its own axiom system (e.g. von Neumann and Morgenstern's EUT for the domain of risky decision-making). A proponent of this approach might endorse some collection of these axiom systems (such as EUT, AGM, and a system of intuitionistic logic), and we call this collection an axiom complex to distinguish it from the domain-specific axiom systems. Note that there are no substantial restrictions placed on the relationship between axiom systems or the contents of a complex. Further, although there are only a handful of axiom systems standardly used and endorsed today, the possibility of additional axiom systems is not precluded.

Definition 6 (*Axiom indicator functions*) Then for each individual axiom ax_{ij}, there is an indicator function $I_{ax_{ij}} : \mathcal{S} \times \mathcal{M} \times \mathcal{S} \times \mathcal{M} \rightarrow \{0, 1\}$ such that $I_{ax_{ij}}(s, m, s', m')$ $= 0$ if the move from pair (s, m) to (s', m') is a violation of axiom ax_{ij}, and $I_{ax_{ij}}(s, m, s', m') = 1$ otherwise.

In the absence of a language these indicator functions give content to the axioms in a sense, by picking out the states of affairs that constitute violations. They are quite simplistic in that they check an axiom against all possible state-context pair-moves, and not only the ones that we really care about. So, the indicator functions indicate the violation or non-violation of the axiom even for moves that have 0-probability and contexts where the axioms are in no way pertinent (in which case they will be trivially un-violated). The important issue of which contexts are relevant to an axiom system will be sorted out with the definition of an *intended context* below, but no such constraints are placed upon the indicators themselves.

Definition 7 (*Traditional axiomatic rationality functions*)

– Define the functions $\tau_{Ax_i} : \mathcal{S} \times \mathcal{M} \times \mathcal{S} \times \mathcal{M} \rightarrow \{0, 1\}$ such that $\tau_{Ax_i}(s, m, s', m') = min_{ax_{ij} \in Ax_i}[I_{ax_{ij}}(s, m, s', m')]$. A function value of 1 will indicate (traditional axiomatic) rationality, while 0 indicates irrationality. So, given an axiom system Ax_i, an agent's behavior in moving from (s, m) to (s', m') is rational according to that system whenever that move violates none of the axioms of that system.

– Similarly, let the general axiomatic rationality function $\tau_{\mathcal{A}x} : \mathcal{S} \times \mathcal{M} \times \mathcal{S} \times \mathcal{M} \rightarrow \{0, 1\}$ be defined such that $\tau_{\mathcal{A}x}(s, m, s', m') = min_{Ax_i \in \mathcal{A}x}[\tau_{Ax_i}(s, m, s', m')]$, so that an agent's behavior is axiomatically rational *tout court* so long as it is not irrational according to any of the endorsed axiom systems.

These τ functions evaluate the rationality of a given move from one state-context pair to another, either with respect to a single axiom system or an axiom complex. This is achieved quite straightforwardly by checking each axiom, and if any are violated, the system or complex as a whole is violated. The evaluation is binary because axiomatic theories are all-or-nothing in nature. Again, the domains are not restricted to reflect the applicability of the axioms to the situation; this comes into play in the appropriate use of the τ functions.

Definition 8 (*Intended context*) The *intended context* of an axiom set Ax_i is the set of contexts $M^{Ax_i} \subseteq \mathcal{M}$ to which those axioms apply according to the TAT.

The intended context is critical to defining axiomatic theories because their axioms are not intended to apply to all situations; rather there are assumptions behind the axioms, i.e. particular conditions that must obtain for the axioms to be in force. An axiom system applies only to a specific domain.

Definition 9 (*Axiomatic theory and theory-complex*) An axiomatic theory, then, is a triple $T_i = (Ax_i, M^{Ax_i}, P)$, where Ax_i is the theory's axioms, M^{Ax_i} is the intended context, and P a problem, such that $P(s, m)(s', m') = \tau_{Ax_i}(s, m, s', m')$ for $m \in M^{Ax_i}$. An axiomatic theory-complex is a collection T of such theories T_i.

This definition is supported by the fact that the TATs currently endorsed are all provided with justification through theorems to the effect that the axioms collectively correspond to an independently-justified coherence principle for their domain of application: The "correctness" of logical axioms is demonstrated through soundness and completeness theorems, of AGM by Grove's models (Grove 1988), and of EUT by the proof that obeying the axioms is equivalent to maximizing some admissible utility function (von Neumann and Morgenstern 1944).

3.4 Ecological rationality in SCOP

Definition 10 (*Ecological units*) An ecological unit is a tuple $U = (P, A^{\varphi}, (s, m))$ where P is a problem, A^{φ} an algorithm for a process φ, and (s, m) a state-context pair. Call the set of all such units $\mathcal{U}$.

The units are what ECO evaluates. It compares (implementations of) processes with respect to contexts, and this requires some valuation for making the comparisons, which is provided by the problem.

Definition 11 (*Ecological rationality*) ECO, then, is a relation $\geq_{ER} \subset \mathcal{U} \times \mathcal{U}$ such that for any two tuples $U = (P, A^{\varphi}, (s, m))$ and $U' = (P', A^{\varphi'}), (s', m'))$:

- It is *not* assumed that either $U \geq_{ER} U'$ or $U' \geq_{ER} U$.
- However, if $P = P'$ and $(s, m) = (s', m')$, then either $U \geq_{ER} U'$ or $U' \geq_{ER} U$.
- $\geq_{ER}$ is transitive and reflexive.

We will abbreviate as usual so that:

- $U \sim_{ER} U'$ indicates that $U \geq_{ER} U'$ and $U' \geq_{ER} U$.
- $U >_{ER} U'$ indicates that $U \geq_{ER} U'$ but not $U' \geq_{ER} U$.

The interpretation of this relation is that $U \geq_{ER} U'$ whenever the application of algorithm A^{φ} to problem P in context m (given s) is at least as ecologically rational as the application of $A^{\varphi'}$ to P' in m' (given s'). ECO captures how well a process performs on multiple dimensions, in particular *accuracy*, *speed*, and *"frugality"*:

Definition 12 (*Accuracy*) Where $P(s^*, m^*) = v^*$, let accuracy be measured by a function $\alpha : \mathcal{U} \rightarrow \mathbb{R}$ such that $\alpha(P, A^{\varphi}, (s^*, m^*)) = \int^{s \in \mathcal{S}} \int^{m \in \mathcal{M}} [v^*(s, m) \cdot A^{\varphi}(s^*, m^*, s, m)]$.

Since accuracy is never really explained but seems to correspond to something neutral like utility maximization, α gives the expected value of performing A^{φ} at (s^*, m^*), according to the valuation v^* indicated by the problem.

Definition 13 (*Speed*) Let speed be a function β such that $\beta : \mathcal{U} \rightarrow (\mathbb{Z} \backslash \mathbb{Z}^+)$.

Without committing to a particular metric at this time, an example of a speed function would be one that returned $(-1\cdot$ (the algorithm's expected runtime in seconds given the initial state-context pair)). Speed could alternatively be a function of the number of discrete steps in the algorithm, the algorithm's complexity (inferring that increasing complexity uniformly increases the time it takes to run), or the number of time-steps the system takes from start to finish of the algorithm.

Definition 14 (*Efficiency*) Similarly, let $\gamma : \mathcal{U} \rightarrow \mathbb{Q}$ be the function measuring *efficiency*.

Informally, we might take the aggregate of the pieces of data looked up or used and define efficiency as their usefulness divided by their quantity. So, efficiency should measure how much value (information) is extracted per unit taken up. Again, a particular metric for efficiency need not be chosen for our present purposes. "Frugality" could alternatively be understood as simply the number of pieces of data used by the algorithm, which would yield a different kind of metric; this is perhaps even closer to what ABC has in mind. However, such a metric might correspond closely to speed, and this measure of efficiency will help to distinguish them while retaining agreement with the intuition that, *ceteris paribus*, the fewer pieces of data used, the better. It also captures the intuition that the *usefulness* or relevance of a given piece of data can overshadow its quantity; for example, an algorithm that uses only one piece of data is arguably less frugal than an algorithm that uses two, if that one piece carries a very small amount of information, while the other algorithm's two pieces are both highly informative.

A possible downside to the efficiency interpretation is that efficiency may not be independent of accuracy, since more useful information would be expected to result in greater accuracy. Efficiency may be distinguished from speed at the expense of a different dependency. Yet the apparent entanglement of these three criteria may prove a boon as reducing the three criteria to the single criterion of efficiency would simplify ECO judgements considerably, and this prospect is promising.

3.5 Supplementary concepts and bridge principles

ECO is a product of how accurate, fast, and efficient a process is, but no account has been provided regarding the relative importance of those criteria, or how they ought be traded off against one another. However, some basic *dominance* principles should be uncontroversial:

Definition 15 (*Strict dominance*) Given any $U, U' \in \mathcal{U}$, if

- $\alpha(U) > \alpha(U')$, and
- $\beta(U) > \beta(U')$, and
- $\gamma(U) > \gamma(U')$,
- then $U >_{ER} U'$.

Definition 16 (*Weak dominance*) Given any $U, U' \in \mathcal{U}$, if

- $\alpha(U) \geq \alpha(U')$, and
- $\beta(U) \geq \beta(U')$, and
- $\gamma(U) \geq \gamma(U')$,
- then $U \geq_{ER} U'$,
- and if at least one inequality is strict, then $U >_{ER} U'$.

Definition 17 (*Expected axiomatic rationality*) Take a process φ, an algorithm A^{φ}, and an axiom complex $\mathcal{A}x$. Let (s, m) be an initial state-context pair. Then, the expected obedience $\tau^{\Sigma}(A^{\varphi}, s, m)$ of using A^{φ} in (s, m), is given by:

$$\tau^{\Sigma}(A^{\varphi}, s, m) = \int^{s' \in \mathcal{S}} \int^{m' \in \mathcal{M}} A^{\varphi}(s, m, s', m') \cdot \tau_{\mathcal{A}x}(s, m, s', m').$$

Note that τ^{Σ} gives the expectation of obedience to all axiom systems in $\mathcal{A}x$, but it is straightforward to define more restrictive $\tau^{\Sigma}_{Ax_i}$ in the same way, to capture the expected obedience to particular axiom systems, or combinations thereof.

4 Results

4.1 Compatibility and consensus

This series of formal definitions already belies the broad compatibility between the TAT approach and ECO regarding foundational questions of the nature of rationality. Although ECO is marketed as the foil to TATs (e.g. Gigerenzer et al. 2011), many of their differences are in fact superficial and thus easily dispelled by the clarity afforded by the precise formal language of SCOP. There are four key questions about rationality to which the TAT approach and ECO give surprisingly accordant answers. These are: What mechanisms does a rational agent use in making decisions, inferences, and so forth? What values does rationality aim at? What factors may or must a theory of rationality take into account? Finally, what is the role of context in rationality judgements?

The stereotypical positions of TATs and ECO with respect to the mechanisms of rationality are, respectively, that heuristics may be too simple to perform adequately and that complex weighting-and-averaging may be beyond the abilities of humans. However, both approaches agree that rationality does not require any particular kind of mechanism *a priori*, and both are compatible with the common-sense position that whether a particular mechanism can be rationally employed is contingent on the facts about what would happen if an agent tried to use that mechanism. Accordingly, the traditional axiomatic approach makes judgements given by the function: $\tau_{Ax_i} : \mathcal{S} \times \mathcal{M} \times \mathcal{S} \times \mathcal{M} \rightarrow \{0, 1\}$, and ecological rationality via the functions $\alpha(U)$, $\beta(U)$, and $\gamma(U)$. Axiomatic rationality is just a function of state-context pairs, specifically whether movement between them violates axioms. ECO judges processes given an initial state-context pair and some task by evaluating the probable consequences, i.e. how quickly, efficiently, and accurately the task will be completed by the process. Nowhere are there any built-in requirements that the mechanism used have any particular features, apart from the very weak claim of ECO that a faster or a more efficient mechanism is preferable *ceteris paribus*. In both cases, rationality depends on something else, namely how well the process would work.

Again, when it comes to the values rationality aims at, the theories basically agree. For the TAT approach, the particular axiom systems typically endorsed are both detachable from the approach as a whole and generally agnostic with respect to the values agents ought be aiming at, beyond the very general values encompassed by the axiom systems (as with EUT, where values are simply taken from the agent's preferences). Still, a judgemental theory cannot get off the ground without at least this kind of weak value claim. In the formalization of ECO, this is accomplished by making judgements relative to a problem P, so that the problem or task identified determines a value function v, which in turn specifies the relative goodness (or perhaps, if we can avoid the term's historical baggage, the utility) of potential resultant state-context pairs, according to P. As with the traditional approach, ECO approach avoids making concrete value claims except where they will be obvious and uncontroversial (as when a task is given and it is assumed that the agent values succeeding with respect to it, e.g. in Gigerenzer and Goldstein 1996). The boldest claim inherent in ECO is that accuracy α, speed β, and efficiency γ are all bearers of value, though in unspecified proportions. Yet accuracy is defined in terms of the agnostic valuation v, and it should be completely uncontroversial that there is some value to humans in being fast and efficient. Thus the formal definitions of both approaches demonstrate that particular value judgements are not intrinsic to them.

Accordingly, both approaches are inherently flexible; anything of perceived importance to rationality could in principle be incorporated. In the TAT approach, additional axioms or axiom systems could be added, and their scopes can be adjusted via their "intended contexts," defined above. ECO takes more of a "check-list" approach, and so there are multiple ways of accounting for more features of a process. Criteria could be explicitly added to the list to accompany α, β, and γ. Features could alternatively be accounted for through the valuation v, which determines the accuracy α of a unit U.

The SCOP framework makes clear how both approaches can make rationality contextual as needed, too. For one, initial state-context pairs are always inputs to rationality judgements, although only ECO emphasizes this in their writing. Critically, a fully-specified axiomatic theory T_i includes the intended context M^{Ax_i} of its axiom system Ax_i, and so inapplicable axioms play no role in the assessment of an agent's rationality; this leaves room for the applicable rules to vary with the context as much as is desired.

So, as argued informally, the TAT approach and ECO are largely compatible approaches to rationality, and this is to a great extent due to the fact that both take pains to remain agnostic regarding what is valuable and important to people. Subsequent theorems will demonstrate that this consensus in fact makes the approaches coincide in a certain sense.

4.2 Agreement theorems

4.2.1 Two theorems

Theorem 1 *Consider an axiomatic theory* $T_i = (Ax_i, M^{Ax_i}, P)$ *and an arbitrary mental state* s. *Then, for any two processes* φ *and* φ', *calculated by algorithms* A^{φ} *and* $A^{\varphi'}$, *and equally fast and efficient in* M^{Ax_i}, $\tau^{\Sigma}_{Ax_i}(A^{\varphi}, s, M^{Ax_i}) \geq \tau^{\Sigma}_{Ax_i}(A^{\varphi'}, s, M^{Ax_i}) \iff U = (P, A^{\varphi}, (s, M^{Ax_i})) \geq_{ER} U' = (P, A^{\varphi'}, (s, M^{Ax_i}))$.

Proof $\tau^{\Sigma}_{Ax_i}(A^{\varphi}, s, M^{Ax_i}) \geq \tau^{\Sigma}_{Ax_i}(A^{\varphi'}, s, M^{Ax_i})$ means that

$\int^{s'\in\mathcal{S}}\int^{m'\in\mathcal{M}} A^{\varphi}(s, m, s', m')\cdot\tau_{Ax_i}(s, m, s', m') \geq \int^{s'\in\mathcal{S}}\int^{m'\in\mathcal{M}} A^{\varphi'}(s, m, s', m')\cdot \tau_{Ax_i}(s, m, s', m')$, by definition. Because problem P is a component of the theory, this means that

$\int^{s'\in\mathcal{S}}\int^{m'\in\mathcal{M}} A^{\varphi}(s, m, s', m')\cdot P(s, m)(s', m') \geq \int^{s'\in\mathcal{S}}\int^{m'\in\mathcal{M}} A^{\varphi'}(s, m, s', m')\cdot P(s, m)(s', m')$. But then by the definiton of α,

$\alpha(P, A^{\varphi}, (s, M^{Ax_i})) \geq \alpha(P, A^{\varphi'}, (s, M^{Ax_i}))$. Since β and γ are equal by hypothesis, this holds iff

$U = (P, A^{\varphi}, (s, M^{Ax_i})) \geq_{ER} U' = (P, A^{\varphi'}, (s, M^{Ax_i}))$, by weak dominance. □

Theorem 2 *Consider an axiomatic theory* (Ax_i, M^{Ax_i}, P). *Fix also an arbitrary mental state* s. *Then, for any two processes* φ *and* φ' *calculated by algorithms* A^{φ} *and* $A^{\varphi'}$ *which guarantee obedience to* Ax_i, *if one of those processes dominates in terms of speed and efficiency, then it will be more ecologically rational.*

Proof It follows from the proof of the previous theorem that if both A^{φ} and $A^{\varphi'}$ guarantee obedience to Ax_i, then $\alpha(P, A^{\varphi}, (s, M^{Ax_i})) = \alpha(P, A^{\varphi'}, (s, M^{Ax_i}))$. But then if $\beta(P, A^{\varphi}, (s, M^{Ax_i})) \geq \beta(P, A^{\varphi'}, (s, M^{Ax_i}))$ and $\gamma(P, A^{\varphi}, (s, M^{Ax_i})) \geq \gamma(P, A^{\varphi'}, (s, M^{Ax_i}))$, it follows by weak dominance that $(P, A^{\varphi}, (s, M^{Ax_i})) \geq_{ER}$

$(P, A^{\varphi'}, (s, M^{Ax_i}))$, and if at least one of those inequalities is strict, then $(P, A^{\varphi}, (s, M^{Ax_i})) >_{ER} (P, A^{\varphi'}, (s, M^{Ax_i}))$. □

4.2.2 Interpretation

These theorems, and their straightforward proofs, demonstrate that the presumed fundamental disagreement between the TAT approach and ECO resulted from differences in rhetoric and presentation, and not true philosophical incompatibility. The theorems essentially show that TATs and ECO will render the same rationality judgements in response to identically-defined situations, although ECO's judgements are more fine-grained.

The definitions preceding the theorems—of an axiomatic theory, its intended context, and expected axiomatic rationality—reflect and reinforce the notion that the traditional axiomatic approach is not the unreasonable opposite that proponents of ECO often make it out to be. Instead of a theory of rationality as the absolute obedience to context-independent rules, TATs provide sets of rules for particular, carefully drawn contexts, within which those rules are proven to achieve good results; a TAT, then, is not just a list of rules but rules in conjunction with an intended context and a presumed system of goals and values. Further, although the traditional theorists do not typically discuss processes, it is clear that the expected axiomatic rationality of the outcomes of a process will vary with the process, and that an agent is more likely to be judged axiomatically rational the greater the expected axiomatic rationality of whatever process they employ for the task at hand. In summary, it is perfectly acceptable, and even important from the point of view of an agent seeking improvement, to consider contexts and processes when assessing traditional axiomatic rationality.

Once this is accepted, it is straightforward to compare TATs to ECO; ECO evaluates processes with respect to contexts with no substantial restrictions on what contexts, processes, or values may be considered. ECO judgements can only be pronounced on its units of evaluation, but there is complete freedom in selecting the components of these units. This means that, specifically, TATs and processes that might satisfy them can comprise the units, which is exactly what is needed to compare the rationality evaluations of the two approaches; it enables us to ask the critical question of how their rationality judgements *of the same situation* might differ. Contrary to what one would have expected from the literature, they should not differ much at all.

Theorem 1 shows that, if we hold fixed the intended context of an axiom system, and presume that whatever those axioms are designed to achieve is conceded to be valuable (a reasonable and not very strong assumption), then a process' ecological rationality corresponds to its expected obedience to the axioms. The most ecologically rational process will guarantee obedience, the least will guarantee disobedience, and a given process is more ecologically rational than another of the same speed and efficiency if and only if it is more likely to obey more of the axioms. The axioms play the role of desirable features of processes, and so the approaches coincide. The theories judge different units, with TATs judging outcomes and ECO judging the process producing them, but their assessments will agree in the sense that the more ecologically rational is the process an agent uses, the better the chance they will appear rational according to TATs.

Holding fixed a TAT, including (of course) its intended context, any two processes guaranteed to obey the axioms will be equally "accurate" given that theory's values. Theorem 2 demonstrates that, though the traditional approach will not treat those processes as different (both will produce rational outcomes), ECO may be able to make further distinctions by looking at the speed and efficiency of those processes; if one of the processes is both faster and more efficient, then it is better—more ecologically rational—than the alternative. The question remains of how speed, efficiency, and accuracy are to be weighted relative to one another, and these trade-offs will influence ecological rationality judgements of processes. But the traditional approach also has the ability (and motivation) to accommodate speed and perhaps efficiency into its value system, and so any agreed-upon specification of the relative values of these components should not cause disagreement between the approaches.

4.3 Significant differences

Nonetheless, there are some truly significant differences between TATs and ECO. These differences come in two kinds: principled, foundational differences concerning the nature of rationality, and differences in framing and practice. I argue that an agent looking to improve will find it more practical to use ECO in light of these differences.

The first foundational difference is that whereas TATs are *outcome-based theories*, and hence look only at the actual outcome in judging rationality, ECO is a *process-based theory*, and so instead judges based on the *expectation over outcomes*, i.e. the attainable outcomes given the process used, and how likely these are. Another way of stating this is that ECO judges based on the likely trend of outcomes if the process were used over and over in the same context.

Another foundational difference is that rationality is binary according to TATs, but relative according to ECO. TATs' partitioning of outcomes into rational and irrational manifests itself in the rationality functions τ that map pairs of state-context pairs to either 0 or 1; in contrast, ecological rationality is defined as a *relation* $\leq_{ER}$, determined by the real-valued valuations v and accuracy α, integer-valued β, and rational (quotient) γ.

In each of these cases, ECO's changes seem to constitute substantial improvements: Intuitively, it is possible for one decision or method to be more rational than another, even when both pass muster; having a theory that can replicate our more fine-grained judgements is clearly advantageous. Researchers within the traditional approach have recognized the value of quantifying degrees of irrationality, as evinced for example by the literature on degrees of incoherence (Schervish et al. 2000). However, this augmentation does not take the form of axioms, and so we should understand it as evidence that the traditional axiomatic approach, in itself, does not tell us everything we would like to know about rationality.

Also intuitively, agents are praised and blamed for their rationality on the basis of activities that they are responsible for. This speaks to the value of judging the processes the agents actually use, and their likely outcomes, rather than simply the actual outcome. A good, apparently rational outcome could easily be produced by sheer accident by a process we would never recommend or praise as rational, and similarly a rational process can "get unlucky" from time to time and happen to produce an inferior outcome, perhaps because the process is probabilistic or is used in a range of circumstances, and is better suited to some than others. In short, I find many of ABC's arguments against TATs to be highly convincing given the purposes espoused here.

ECO also has advantages in its formulation, probably as a result of its being a newer approach constructed largely in response to TATs. Criticisms of TATs have often resulted from the fact that their comprehensiveness and ability to accommodate context-dependence are hidden, hence critiques of the axioms that miss their intended contexts of application and erroneous charges that EUT ignores pro-social preferences, for example. ECO prevents such misunderstandings by making the components of evaluations explicit: a context and process must be explicitly described, and the familiar check-list (fast? frugal? accurate?) keeps the criteria of evaluation out in the open. Furthermore, again for partly historical reasons, ECO has avoided taking certain contexts as privileged; in contrast, the TAT approach has focused quite narrowly on a few particular contexts, giving a sense that the approach has nothing to say about other contexts or overlapping domains (although this is not true, in principle). Both of these features make ECO more transparent and easier to use; as a theory of rationality should make it easy for people to see how its norms apply and work on improving with respect to them, this feature of ECO should not be undervalued.

5 Conclusions

Despite the inevitable differences between the approaches due to their formulations, their coincidence is quite striking given how little must be agreed upon for the two to assess like situations alike. Thus the moral of this paper is that disagreement should be largely eliminable by being very clear and precise about the context, problem, and values under consideration. The warring sides do not really disagree about such fundamental issues as what factors are relevant to rationality judgements (e.g. the context) or what value system should be used to judge possible outcomes. This is quite important as it implies that much effort has been expended over misunderstandings and, crucially, we can push the perceived differences aside, move forward in rationality research, tackle more difficult questions, and construct more clearly formulated and useful theories that may be ratified by all sides.

My emphasis on the broad consensus between the competing TAT and ECO approaches should not be taken to indicate that there are no interesting foundational questions left; there are many. For example, there must be more to say about the values rationality aims at than that they depend on the problem. Some goals may be more worthy than others, and rationality may have a role to play in selecting agents' goals.

Furthermore, I have downplayed the difference between seeking general rules of rationality that apply in a range of contexts and studying particular contexts first, perhaps learning more general rules as patterns emerge. While this difference does not necessarily influence the assessments of particular situations, it is indicative of a real trade-off between relying on general rules that may not be perfectly suited for all their contexts of application and spending more time and resources to best-respond to the details of one's situation. The approaches tend towards opposite sides, but it is clear that people in practice must strike a balance; how this should be done is an important question that has not yet been addressed.

Finally, while I have praised many of ECO's contributions, many of its concepts and principles have not been fully explained, and there are puzzles surrounding how it ought be used by an agent interested in the big-picture of rationality. A fuller specification will need to provide an account of the meanings of ECO's criteria, as well as how they are to be traded off (perhaps even reducing everything to efficiency, as I suggest as a simplification). An agent would also require guidance as to how to determine the proper context of evaluation for a given process, and the level of context specificity they should focus on for the best results. It is also not clear how or in what sense agents' justification for their process usages may be incorporated into ECO's judgements of their rationality, although I would argue that this is highly relevant.

I expect that having the formal framework of SCOP available when modeling and theorizing will facilitate addressing these issues and the many others that remain, as it is arguably a very complete and accurate representation of the mental activities we wish to judge in the abstract, and therefore any assumptions made will be visible in the framework, and not forgotten.

Further work should include taking advantage of the power of the formalism to interpret and support more specific theories, rather than just theory types. Now that we can see from a top-down perspective what we need for our purposes, we can start to fill in the details subject to those specifications.

References

Aerts, D. (1982). Description of many separated physical entities without the paradoxes encountered in quantum mechanics. *Foundations of Physics*, *12*(12), 1131–1170.

Aerts, D. (1983). Classical theories and nonclassical theories as special cases of a more general theory. *Journal of Mathematical Physics*, *24*(10), 2441–2453.

Aerts, D. (1999). Foundations of quantum physics: a general realistic and operational approach. *International Journal of Theoretical Physics*, *38*(1), 289–358.

Aerts, D. (2002). Being and change: Foundations of a realistic operational formalism. In D. Aerts, M. Czachor, & T. Durt (Eds.), *Probing the structure of quantum mechanics: nonlinearity, nonlocality, computation and axiomatics*. Singapore: World Scientific.

Aerts, D., & Aerts, S. (1995). Application of quantum statistics in psychological studies of decision processes. *Foundations of Sciences*, *1*, 85–97.

Aerts, D., & D'Hooghe, B. (2009). Classical logical versus quantum conceptual thought: Examples in economics, decision theory and concept theory. In *Quantum Interaction* (pp. 128–142). Heidelberg: Springer.

Aerts, D., & Gabora, L. (2005). A theory of concepts and their combinations II: a Hilbert space representation. *Kybernetes*, *34*(1/2), 192–221.

Aerts, D., Gabora, L., & Sozzo, S. (2013). Concepts and their dynamics: a quantum-theoretic modeling of human thought. *Topics in Cognitive Science*, *5*(4), 737–772.

Agrawal, P. M., & Sharda, R. (2010). Quantum mechanics and human decision making. *Social Science Research Network Electronic Journal*, pp. 1–49.

Alchourron, C., Gardenfors, P., & Makinson, D. (1985). On the logic of theory change: partial meet contraction and revision functions. *Journal of Symbolic Logic*, *50*, 510–530.

Allais, M. (1953). Le comportement de l'homme rationnel devant le risque: critique des postulates et axiomes de l'ecole americaine. *Econometrica*, *21*(4), 503–546.

Birkhoff, G., & von Neumann, J. (1936). The logic of quantum mechanics. *The Annals of Mathematics*, *37*(4), 823–843.

Bruza, P., Busemeyer, J. R., & Gabora, L. (2009). Introduction to special issue on quantum cognition. *Journal of Mathematical Psychology*, *53*(5), 303–305.

Bruza, P., & Gabora, L. (Eds.). (Oct. 2009). *Journal of Mathematical Psychology, 53, Special Issue: Quantum Cognition*. Elsevier.

Ellsberg, D. (1961). Risk, ambiguity, and the Savage axioms. *Quarterly Journal of Economics*, *75*, 643–669.

Gabora, L., & Aerts, D. (2002). Contextualizing concepts using a mathematical generalization of the quantum formalism. *Journal of Experimental & Theoretical Artificial Intelligence*, *14*(4), 327–358.

Gabora, L., Rosch, E., & Aerts, D. (2008). Toward an ecological theory of concepts. *Ecological Psychology*, *20*(1), 84–116.

Gigerenzer, G., & Goldstein, D. G. (1996). Reasoning the fast and frugal way: models of bounded rationality. *Psychological Review*, *103*(4), 650–669.

Gigerenzer, G., Hertwig, R., & Pachur, T. (Eds.). (2011). *Heuristics: the foundations of adaptive behavior*. New York: Oxford Universtity Press.

Gigerenzer, G., & Selten, R. (Eds.). (1999). *Bounded rationality: the adaptive toolbox. Dahlem workshop reports*. Cambridge: The MIT Press.

Grove, A. (1988). Two modellings for theory change. *Journal of Philosophical Logic*, *17*(2), 157–170.

Kitcher, P. (1992). The naturalists return. *The Philosophical Review*, *101*(1), 53–114. Philosophy in Review: Essays on Contemporary Philosophy.

Schervish, M. J., Seidenfeld, T., & Kadane, J. B. (2000). How sets of coherent probabilities may serve as models for degrees of incoherence. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, *8*(03), 347–355.

Simon, H. A. (1956). Rational choice and the structure of the environment. *Psychological Review*, *63*(2), 129–138.

Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: heuristics and biases. *Science*, *185*(4157), 1124–1131.

von Neumann, J. (1955). *Mathematical Foundations of Quantum Mechanics*. Princeton: Princeton University Press. translated by Robert T. Beyer.

von Neumann, J., & Morgenstern, O. (1944). *Theory of games and economic behavior*. Princeton: Princeton University Press.

Wallin, A., & Gärdenfors, P. (2000). Smart people who make simple heuristics work. *Behavioral and Brain Sciences*, *23*(5), 765.

Weirich, P. (2004). *Realistic decision theory: rules for nonideal agents in nonideal circumstances*. New York: Oxford University Press.