

Enactive Problem Solving and Chatbot Architectures

RICCARDO VIALE

Herbert Simon Society and University of Milano Bicocca, Milano, Italy

SHAUN GALLAGHER

University of Memphis, Memphis, USA

In this paper we review different conceptions of rationality as they apply to decision making and problem solving. We defend the notion of enactive problem solving (EPS) and we consider how it applies to the use of therapeutic chatbots. We start by recounting the development of different models of rationality, from expected value maximization, to subjective expected utility theory, and of bounded rationality from prospect theory, to simple adaptive heuristics. We then focus on problem solving. We show how the centre of gravity of problem solving shifts from the model of computational and cognitive processing, to an action-oriented pragmatic one in EPS, that is, to the possible actions that the body-environment interaction allows. We highlight a notion of the agent's relational autonomy as part of EPS and show that this can be a problem in some chatbot designs in the therapeutic context. EPS suggests that simple early versions of therapeutic chatbots like ELIZA may be more beneficial than more recent sophisticated models involving "Belief-Desire-Intention" (BDI) architecture characterized by predetermined trajectories. The system formed between the agent plus the chatbot tool facilitates a form of enactive problem solving, reducing rigidity on the side of the agent.

Keywords: Bounded rationality; enactive problem solving; chatbot; AI.

1. The repair programme about the limits of rationality

Questions about the limits of rationality are part of a long tradition that includes the 18th century critique of the theory of expected value maximization. The concept of expected value maximization, within the

context of decision theory, is credited to Blaise Pascal and Pierre de Fermat in the 17th century. They recognized that by calculating the expected value of different outcomes, a decision maker could choose the rational action that offered the highest potential payoff. Since then it has been a history of discovering empirical and logical anomalies in the theory. Gigerenzer (2008: 90), for example, has described the historical progression – from expected value maximization (as a standard of rationality) to expected utility theory and then on to prospect theory – as a “repair program” aimed at resuscitating the mathematical operation of weighted integration, based on the definition of mathematical expectation, as a psychological theory of rationality. Although expected value maximization was once regarded as a proper standard of rationality, it ran aground on the St. Petersburg Paradox, and Daniel Bernoulli began the repair program by transforming the outcomes associated with lotteries using a logarithmic utility of money function. This modification survived and grew as expected utility theory which took root in 20th century neoclassical economics. Then came Allais’ Paradox, which damaged expected utility theory’s ability to explain observed behavior, and a new repair appeared in the form of prospect theory, which introduced more transformations with additional parameters to square the basic operation of probability-weighted averaging with observed choices over lotteries.

2. *Rationality as SEU decision making*

The empirical study of human rationality from the post-war period to date has mainly developed by looking at a normative model of decision making. In particular Subjective Expected Utility (SEU) decision making, which stems from the subjective expected utility theory of Savage (1950) that itself extended the results of Von Neumann and Morgenstern (1944).¹

In decision theory, the von Neumann–Morgenstern utility theorem² shows that under certain axioms of rational behaviour, such as completeness and transitivity, a decision maker faced with risky (probabilistic) outcomes of different choices will behave as if he or she is maximizing the expected value of some function defined over the potential outcomes at some specified point in the future. The theory recommends which option rational individuals should choose in a complex situation, based on their risk appetite and preferences. The theory of subjective expected utility combines two concepts: first, a personal utility func-

¹ The way in which this escalation developed is discussed in detail in Mousavi and Tideman (2019).

² Von Neumann and Morgenstern never intended axiomatic rationality to describe what humans and other animals do or what they should do. Rather, their intention was to prove that if an individual satisfies the set of axioms, then their choice can be represented by a utility function.

tion, and second a personal probability distribution (usually based on Bayesian probability theory).³

In other words (Viale 2024), if decision makers want to reach a goal, for example enrolment at university, they face a limited number of alternatives (such as different universities and different courses), each of which has a probable outcome (such as cost, teaching value, difficulties of exams, type of professional training, consequences on the potential job i.e., on their salary, on the location of the employer, and so on) that has consequences related to the utility function. The concepts used to define the decision are therefore information about the world, the risk related to outcomes and consequences; preferences over alternatives; the relative utilities on the consequences; and, finally, the computation to maximize the subjective expected utility. Even if in formal decision theory no explicit reference is made to the actual empirical mental and psychological characteristics of the decision maker, in fact these can be mapped onto psychological processes such as external perceptual incoming inputs or internal mnemonic inputs, in mental representations of the states of the world on the basis of information, in hedonic evaluations⁴ of the states of the world, and in deductive and probabilistic computation on the possible decisions to be implemented on the basis of hedonic evaluations.

3. SEU driven psychology of decision and bounded rationality

On this view the cognitive psychology of decision making precisely reflects the conceptual structure of formal decision theory. In relation to this structure and the normative component derived from it, empirical research in the cognitive psychology of decision making has been developing since the 1950s. Weiss and Shanteau (2021), highlight that in the 1950s Ward Edwards, the founder of the psychology of decision making, began to carry out laboratory experiments to unravel the way in which people actually decide. His experiments, which became the reference

³ This theoretical model has been known for its clear and elegant structure and it's considered by some researchers one of "the most brilliant axiomatic theor[ies] of utility ever developed". Assuming the probability of an event, Savage defines it in terms of preferences over acts. Savage used the states (which are not under one's control) to calculate the probability of an event. On the other hand, he used utility and intrinsic preferences to predict the event's outcome. Savage assumed that each act and state are enough to uniquely determine an outcome. However, this assumption breaks down in cases where the individual lacks sufficient information about the event. In reality Savage explicitly limited the theory to small worlds, that is, situations in which the exhaustive and mutually exclusive set of future states S and their consequences C are known.

⁴ The hedonic approach to economic assessment can be used for evaluating the economic value of goods. The hedonic approach is based on the assumption that goods can be considered aggregates of different attributes, some of which, as they cannot be sold separately, do not have an individual price.

of subsequent generations and in particular of Daniel Kahneman and Amos Tversky's Heuristics and Biases program, have two fundamental characteristics: firstly, the provisions of the SEU are set as a normative reference, and the experimental work has the aim of evaluating when and how the human decision maker deviates from the requirements of the SEU. Ultimately, the aim is to discover the irrational performances in the decision. Secondly, the experiments are not carried out in the real decision-making contexts of everyday life, but in an abstract one of games, gambling, bets and lotteries. In these abstract experimental situations, characterized by risk, the informative characteristics typical of the real environment – such as uncertainty, complexity, poor definition of data, instability of phenomena, dynamic and interactive change with the decision maker, and so on – are entirely absent. In fact, the gambling paradigm was considered the indispensable “fruit fly” of research on decision making (Lopes 1983) in those years, and experimental activity was confined to abstract situations characterized by risk, that is, by a defined and closed set of alternatives with outputs measurable at a probabilistic level and assessable at the level of relative utility. The influence of economics and its methodological imprint linked to physics is evident. Abstraction and simplification are the methodological objectives that allow independent variables to be kept under control and disturbance phenomena to be reduced (Blaug 1980; Viale 1997, 2012, 2013, forthcoming). However, it should not be overlooked that this may lead to the description of psychological abstractions that do not correspond to what happens in the individual's action processes in the reality of daily and professional life.

The empirical and deductive anomalies of the theory of rationality have informed the development of different concepts of bounded rationality. The limitations were found in the suboptimality of the information collection, in suboptimal computational ability to process data, and in the adaptive interaction between mind and environment. Robert Aumann (1962; 1997) advanced five arguments for bounded rationality that represent the different concepts. To summarize his remarks: (1) Even in very simple decision problems, most economic agents are not (deliberate) maximizers; (2) Even if economic agents aspired to pick a maximal element from a choice set, performing such maximizations are typically difficult and most people are unable to do so in practice; (3) Experiments indicate that people fail to satisfy the basic assumptions of rational decision theory; (4) Experiments indicate that the conclusions of rational analysis (broadly construed to include rational decision theory) do not match observed behavior; (5) Some of the conclusions of rational analysis do not agree with a reasonable normative standard.

Accordingly, for a majority of economic researchers across disciplines, bounded rationality is identified with some form of optimization problem under constraints. On the contrary, Gerd Gigerenzer is among the most prominent and vocal critics of the roles that optimization

methods and logical consistency play in commonplace normative standards for human rationality (Gigerenzer and Brighton 2009), especially the role those standards play in Kahneman and Tversky's biases and heuristics program (Kahneman and Tversky 1996; Gigerenzer 1996).

4. *Problem solving action*

Herbert Simon (1986) emphasizes the importance of problem solving and differentiates it from decision making, which he considers a phase downstream of the former. In dealing with a task, humans have to frame problems, set goals and develop alternatives. Evaluations and judgments about the future effects of the choice are the final stages of the cognitive activity. In fact, Simon's research in AI, economic and organizational theory is almost entirely dedicated to problem solving that seems to absorb the evaluation and judgment phase. On the one hand, action, to the extent that it has a role in this process, seems to arise at the end of the problem solving without interruption.

On the other hand, his approach to problem solving highlights the influence of American pragmatism, and in particular of John Dewey (1910) and Charles Sanders Peirce (1931) and William James (1890), on his work. The centre of gravity of the rationality of the action lies in the ability to adapt. And the centre of gravity of adaptation is not so much in the internal environment of the actor, that is, in his or her cognitive characteristics, as in the pragmatic external environment. Simon and Newell write: "For a system to be adaptive means that it is capable of grappling with whatever task environment confronts it. Hence, to the extent that a system is adaptive, its behaviour is determined by the demands of the task environment rather than by its own internal characteristics. Only when the environment stresses its capacities along some dimension – presses its performance to the limit – do we discover what those capabilities and limits are, and are we able to measure some of their parameters" (Simon and Newell 1971: 149).

The metaphor of the ant on the beach (Simon 1981) is illuminating: imagine an ant walking on a beach. Now let's say you wanted to understand why the ant is walking in the particular path that it is. In Simon's parable, you cannot understand the ant's behaviour just by looking at the ant: "Viewed as a geometric figure, the ant's path is irregular, complex, hard to describe. But its complexity is really a complexity in the surface of the beach, not a complexity in the ant." (Simon, 1981 (1988): 80). In other words, to predict the path of the ant, we have to consider the effects of the beach – the context that the ant is operating in. The message is clear: we cannot study what individuals want, need or value detached from the context of the environment that they are in. That environment shapes and influences their behaviour. In this example, the procedural rationality of the ant (finding a suitable behaviour on the beach) allows its substantial rationality (the adaptivity to the irregularity of the beach).

From this metaphor Simon derives a philosophical principle very much in tune with wide embodied cognition: “A man considered as a system capable of having a behaviour is very simple. The apparent complexity of his behaviour over time is largely a reflection of the complexity of the environment in which he finds himself” (Simon 1981 (1988): 81) The behaviour adapts to external purposes and reveals those characteristics of the system that limit its adaptation. Changes in the environment alter the ability to interact with it.

On this interpretation, the correspondence between action and solution of a problem conceptually bypasses the analytic phase of the decision and limits the role of symbolic representation (Viale 2024). The decision-making model based on SEU theory does not correspond to the empirical reality of individual action. In solving any problem, whether opening a door, running to catch a falling ball, replacing a car tyre, calculating for a financial investment, solving tests and puzzles or negotiating with a competitor, the search for the solution corresponds to acting, to a recursive feedback process leading up to the final action.

5. Enactive Problem Solving (EPS)

From the postwar period to today, the study of reasoning, judgment, and decision processes has largely been conducted within the classic cognitive model, often referred to as “Information Processing Psychology”. It has three characteristics: first, thought is computation and occurs as computation. Every mental activity is performed by algorithms similar to the computer’s machine language. Cognition derives from computational procedures that are carried out on abstract symbolic structures.

Cognitivism has portrayed the mind that thinks and decides as if it were in a vat, separated from the body and the environment. The mind is “disembodied” from the body that carries it and “detached” from the environment in which it interacts. The new perspective of 4E (embodied, embedded, extended, and enactive) cognition reveals instead a cognition integrated with the body through action and shaped by the environment with which the body interacts and in which it is located. This bodily interaction with the environment shapes and models the cognitive activity. Enactive approaches emphasize the idea that the body is dynamically coupled to the environment in important ways (Thompson 2007; Di Paolo 2005); they point not only to sensorimotor contingencies (where specific kinds of movement change perceptual input) (O’Regan and Noë 2001), but also to bodily affectivity and emotion (Colombetti 2013) as playing a nonrepresentational role in cognition. Embedded and enactive approaches emphasize action affordances that are body- and skill-relative (Chemero 2009). More generally, most theorists of embodied cognition hold that these ideas help to shift the ground away from orthodox, purely computational cognitive science, which clearly

informs the cognitive psychology of decision making. In this respect, it's not just the internal processes of the mind or brain, but the brain-body-environment system that is the unit of explanation.

For EPS action is part of the cognitive process (Viale, Gallagher, and Gallese 2023); deliberation, judgment and evaluation are intrinsic aspects of an action process that may involve the agent/problem solver moving around to get a better look, or actively engaging with others to gain perspective. Problem solving is a dynamic process based on pragmatic, recursive actions and on positive or negative feedback from the environment and its affordances. New possibilities, and new requirements emerge, not simply from the mind of the problem solver, but are often generated from the social interactions and worldly practices that characterize this situation. These are the factors that feed back into the problem-solving process. The enactive emphasis on flexibility in everyday problem solving emphasizes embodied and social practices – actual doings and material engagements that typically involve others in either direct or indirect ways – rather than internal deliberations, models, representations or rules. Deliberation is not ruled out, but it often occurs in consultation with others as problematic situations develop and as we work on a task, with the working component entering into deliberative practices. Models are often embodied in external drawings or diagrams or material arrangements that we try out as we try to solve the problem. Cognitive processes are extended by the tools and technologies we use. This conception of EPS delivers the action-oriented detail on the kind of adaptivity that Newell and Simon were seeking in the concept of bounded rationality – what they referred to as the “grappling with whatever task environment confronts” the agent (Newell and Simon 1971: 149), where ‘grappling’ means, for EPS, embodied hands-on activities that together with neural and extra-neural processes do the work of cognition.

One important principle central to the enactive view of problem solving is the agent's preservation of autonomy. Even when the agent encounters environmental constraints, problem solving continues only if the agent can improvise or make a move to resolve or make use of that constraint. Autonomy on this model is relational (Gallagher 2020; MacKenzie and Stoljar 2000) rather than implying absolute control. Nonetheless, some degree of autonomy is required for the process to continue.

The centre of gravity of problem solving is therefore no longer located in the computational and cognitive part, but it shifts to the action-oriented pragmatic one, that is, to the possible actions that the body-environment interaction allows (Viale 2024). This position, which places the constraints of the rational activity of choice and decision not so much in the computational possibilities of the human mind as in the mind-body-environment interaction, represents a further development of Herbert Simon's theory of Bounded Rationality (Viale, Gallagher,

and Gallese 2023). The environment cannot be analysed only as a structure of the task through its computational variables⁵. The physical and social environment also generates sensory and motor constraints that influence reasoning and action. And in determining a choice, possible or simulated bodily actions have an influence in shaping the range of possible options and the value attributed to them.

Relevant to the idea of problem solving, there is general agreement that the environment scaffolds our cognitive processes, and that our engagement with the environmental structure, and environmental features, including external props and devices, can shift cognitive load. Already, within the scope of Simon's own work it's clear that only through the enactive interaction between problem solver and environmental affordances is it possible to construct a solution.

For the idea of enactive problem solving, however, it is important to emphasize two things. First, the relational nature of affordances (Viale, Gallagher, and Gallese 2023). It is not just the environment that constrains behaviour; it is also the body's morphology and motor possibilities, and the agent's past experience and skill level that will define what counts as an affordance. The way in which the body couples (or can couple) to the environment, will delineate the set of possibilities or solutions available to the agent. Likewise, affordances can also be limited by an agent's affective processes, emotional states, and moods. It is sometimes not just what "I can" do (given my skill level and what the environment affords), but what "I feel like (or don't feel like)" doing (given my emotional state).

Second, as the pragmatists pointed out, the environment is not just the physical surroundings; it's also social and cultural and characterized by normative structures (Viale, Gallagher, and Gallese 2023). As Gibson (1977; 1979) indicated, affordances can be social. Enactive problem solving also highlights the important role of social and intersubjective interactions. Again, it's not only what "I can" do, but also what "I can't" (or "I ought not") do given normative or institutional constraints, as well as cultural factors that have to do with, for example, gender and race. These are larger issues that range from understanding how dyadic interactions shape our developing skills, to how institutional factors can either enable or constrain our social interactions.

6. *BDI Artificial Intelligence*

The first phase of AI can be described as the growth of so-called Simulative AI involving amongst other things descriptive protocols for com-

⁵ For example, this includes the characteristics of the structure of the environment introduced by Gigerenzer and colleagues (Gigerenzer and Gassmaier 2011) such as uncertainty, redundancy, variability, number of alternatives and sample size. These characteristics derive from symbolically deconstructed empirical phenomena that are manipulated as cues with statistical meaning (tallied, weighted, sequenced and ordered).

puter simulation of human problem solving and thinking, based on the development of information processing psychology. It found its success in the development of expert systems aimed at problem solving and the ex-post simulation of scientific discoveries. It failed though in machine translation, image and speech processing and artificial movement.

One of the greatest successes of the first phase was the “logic theorist” and the general problem solver (GPS) computer program developed by Herbert Simon, Alan Newell and Cliff Shaw in 1956. The GPS was a computer program intended to work as a universal problem solver machine. In contrast to the former Logic Theorist project, the GPS worked with means–ends analysis. The first-generation AI programmes were psychology-based simulations of how humans reason and solve problems. They relied on a Cartesian dualism between mind and body (software and hardware), were fully transparent and predictable. They relied on a model of mind taken from folk psychology (Stich 1985) the belief–desire–intention triad (BDI). This model expressed by information processing psychology and ‘good-old-fashioned’ cognitivism is present in both SEU decision making and in computational models of problem solving.

The BDI model has been used to program intelligent agents. Superficially characterized by the implementation of an agent’s beliefs, desires and intentions, it actually uses these concepts to solve a particular problem in agent programming. The BDI model allows for the representation of the characteristics and methods of achieving a goal in a system built according to the paradigms of software agents. Achieving the goal is what the agent works for. This is identified by a progression of its internal states that tends to implement the task in a stable manner by determining the actions that the agent is able to undertake. The simplest implementation of this concept is an algorithm that returns a value indicating success if the goal is achieved. However, more complex goals can be divided into different categories:

- achievement goals (long-term goals)
- satisfaction goals (recurring goals, such as gathering resources)
- preservation goals (preservation of life and property)
- delta goals (for example, changes in state).

In the BDI approach, intentions and motivational aspects are also included in the achievement of the goal.

These architectures are based on the concept of practical reasoning (Bratman 1999) which allows the analysis of the reasoning process performed by people when they aim to satisfy their expectations in the real world. The aim is to create an artificial system that simulates such behavior. A fundamental moment of practical reasoning is the choice of one’s own objectives among those that are believed to be achievable. This choice process presupposes knowledge of both the options available and one’s own desires, aspects involved in the traditional concept of autonomy (Frankfurt 1988). The decisions taken have different de-

degrees of complexity and importance: fundamental choices can be made, such as deciding to exploit a totally new and unexpected opportunity, or simply deciding the best way to do something. In the theory of practical reasoning this step is described as the transformation of the chosen options into intentions, which have the characteristic of persisting as long as our beliefs make them reasonable. Practical reasoning is composed of two main processes:

- the decision process between a set of different perspectives
- the process for achieving the proposed condition.

In the process described, proactive behavior (persistence of intentions) and reactive behavior (abandonment of intentions because they are no longer convenient and adoption of a new goal) can be clearly identified. Reactive behavior allows us to verify our choices in light of the latest information, instead of blindly focusing on what we have previously decided. The fundamental characteristic of BDI architectures is to use data structures that correspond to the beliefs, desires and intentions referred to in the theory of practical reasoning; decision-making processes are implemented by functions that act on these structures.

7. Chatbot architecture: BDI versus EPS

We consider here the use of chatbots in the therapeutic context, motivated in some cases by problems with intersubjective relations where patients mistrust the therapist and resist therapy, leading to high dropout rates (Martino et al. 2012). AI-supported interventions (in the form of a chatbot) are proposed to address these issues (Szalai 2021). The specific proposal: AI programs in natural language processing “could help patients re-author their self-narratives into more coherent and meaningful sequences, in which they view themselves more like agents and less subject to external control, as well as assign more consistent roles to others” (Szalai 2021). The aim, in part, is to enhance the autonomy of the patient.

The chatbot is a natural language processing-based AI device that is designed to mimic human interaction in order to supplement sessions with the therapist. The system is meant to facilitate the patient’s reflective processes leading to the reframing of self-narrative. The benefit of such a system is that it could reduce motivation for the patient’s antagonistic responses to the therapist. Szalai notes that patients are more willing to reveal information to “virtual humans” than to the human therapist. The use of AI in this context is not meant to solve all problems; nor is it meant to replace the therapist, and it cannot be the complete therapeutic process.

One of the early first attempts at creating a natural language program that could run a simple psychotherapeutic encounter was ELIZA, a rule-based chat system that would basically turn anything you said into a question that was repeated back to you, with the idea of getting

you to elucidate. If all goes well, you seemingly enact or talk your way to a cure. ELIZA doesn't diagnose anything. Rather it simply emulates a Carl Rogers style of talk therapy that mirrors whatever the analysand says (Bassett 2019). It runs a rule that basically asks you to elucidate, with the goal of "self-actualization." Although ELIZA seems somewhat primitive compared to current AI deep-learning chatbots, there are anecdotal reports that subjects appreciated the encounter. As Caroline Bassett puts it, "if humans found ELIZA useful perhaps it was as a mirror, a listening surface which enabled forms of self-examination, self expression, or self re-narrativization. If users found something revealing in their interactions with ELIZA then that something was their own: ELIZA never did, and does not now, deliver injunctions, suggestions—or nudges; and has no program to promulgate" (2019: 809). ELIZA itself doesn't have a goal, and there is no trajectory involved; it's not committed to anything. As such, it grants a high degree of autonomy to the person using it. We suggest that this is a model that is consistent with enactive problem solving. The therapeutic solution emerges from the patient's interaction with the technology in a way that preserves the patient's autonomy.

In contrast to ELIZA, current chatbots can be designed with a goal or trajectory built in. For example, a BDI architecture tends to rigidly stay on course, or "overcommit" to a predetermined goal (Wallis 2022) – this kind of inflexibility could lead patients to an impersonally predetermined self-construction – a situation that can rob the patient of autonomy.⁶ The rigidity noted by Wallis is not easily fixed and may be persistent in such systems, introducing issues of sustainability. Artificial systems may fail or may work differently than expected in changing circumstances (so-called counterperformativity (Bamford and MacKenzie 2018)). Circumstances (including human inclinations) frequently change.

An alternative approach, proposed by Wallis, drawing on enactivist ideas, is to design a system to directly recognize a trajectory (a goal or intention) in the patient's action by recognizing when its own behavior could be complementary, thus forming a shared trajectory. The task is to get the AI system or artificial agent to recognize what the patient's action or expression affords, i.e., to register it in terms of its own possible response with a high probability weighting for appropriateness. Wallis admits, however, that "implementing this on a computer is beyond us at the moment..." (2022: 9). This, however, is not just a pragmatic issue – perhaps to be overcome in the fast pace of AI

⁶ Chatbots offer a form of cognitive scaffolding as a structure meant to support or enable the user's problem solving. One important consideration here is whether these trajectories are transparent to the user. If the user is unaware of how the chatbot is designed for a particular outcome, and how it may be shaping the user's thoughts or experiences, this threatens the user's autonomy. This is sometimes referred to as 'hostile scaffolding' (Timms and Spurrett 2023)

progress – or a theoretical problem that requires more research – it’s a tricky kind of hermeneutical problem. If one allows the AI system to be guided by the patient, the shared trajectory may simply become a reinforced loop back into the patient’s disorder. On this score, however, the same worry applies to ELIZA. This, we note, is a common worry in the therapeutic context – how to allow some significant degree of autonomy (or agency) to the patient, but at the same time effect a solution that moves the patient to a new perspective. In this respect, the solution may depend on the intervention (a meta-guidance) by a seemingly irreplaceable human psychotherapist who is attuned to counter problems caused by structural underdetermination (as in ELIZA) or overdetermination (as in BDI).

References

- Aumann, R. J. 1962. “Utility Theory without the Completeness Axiom.” *Econometrica* 30 (3): 445–462.
- Aumann, R. J. 1997. “Rationality and Bounded Rationality.” *Games and Economic Behavior* 21 (1–2): 2–17.
- Bamford, A. and D. MacKenzie. 2018. “Counterperformativity.” *New Left Review* 113: 97–121.
- Bassett, C. 2019. “The Computational Therapeutic: Exploring Weizenbaum’s ELIZA as a History of the Present.” *AI & Society* 34 (4): 803–812.
- Blaug, M. 1980. *The Methodology of Economics*. Cambridge: Cambridge University Press.
- Bratman, M. E. 1999. *Intention, Plans, and Practical Reason*. Stanford: CSLI Publications.
- Chemero, A. 2009. *Radical Embodied Cognitive Science*. Cambridge: MIT Press.
- Colombetti, G. 2013. *The Feeling Body: Affective Science Meets the Enactive Mind*. Cambridge: MIT Press.
- Dewey, J. 1910. *How We Think*. Boston: D. C. Heath.
- Di Paolo, E. A. 2005. “Autopoiesis, Adaptivity, Teleology, Agency.” *Phenomenology and the Cognitive Sciences* 4 (4): 429–452.
- Frankfurt, H. G. 1988. *The Importance of What We Care About*. Cambridge: Cambridge University Press.
- Gallagher, S. 2020. *Action and Interaction*. Oxford: Oxford University Press.
- Gibson, J. J. 1977. “The Theory of Affordances.” In R. Shaw and J. Bransford (eds.), *Perceiving, Acting, and Knowing*. Hillsdale: Erlbaum, 67–82.
- Gibson, J. J. 1979. *The Ecological Approach to Visual Perception*. Boston: Houghton Mifflin.
- Gigerenzer, G. 1996. “On Narrow Norms and Vague Heuristics: A Reply to Kahneman and Tversky.” *Psychological Review* 103 (3): 592–596.
- Gigerenzer, G. 2008. *Rationality for Mortals: How People Cope with Uncertainty*. Oxford: Oxford University Press.
- Gigerenzer, G. and H. Brighton. 2009. “Homo Heuristicus: Why Biased Minds Make Better Inferences.” *Topics in Cognitive Science* 1 (1): 107–143.

- Gigerenzer, G. and W. Gaissmaier. 2011. "Heuristic Decision Making." *Annual Review of Psychology* 62: 451–482.
- Kahneman, D. and A. Tversky. 1996. "On the Reality of Cognitive Illusions." *Psychological Review* 103 (3): 582–591.
- Lopes, L. L. 1983. "Some Thoughts on the Psychological Concept of Risk." *Journal of Experimental Psychology: Human Perception and Performance* 9: 137–144.
- James, W. 1890. *The Principles of Psychology*. New York: Dover.
- Mackenzie, C. and N. Stoljar (eds.). 2000. *Relational Autonomy: Feminist Perspectives on Autonomy, Agency, and the Social Self*. Oxford: Oxford University Press.
- Martino, F., M. Menchetti, E. Pozzi and D. Berardi. 2012. "Predictors of Dropout among Personality Disorders in a Specialist Outpatients Psychosocial Treatment." *Psychiatry and Clinical Neurosciences* 66 (3): 180–186.
- Mousavi, S. and N. Tideman. 2021. "Beyond Economists' Armchair: The Rise of Procedural Economics." In R. Viale (ed.), *Routledge Handbook of Bounded Rationality*. London: Routledge.
- Peirce, C. S. 1931. *Collected Papers of Charles Sanders Peirce*. Cambridge: Harvard University Press.
- Savage, L. J. 1954. *The Foundations of Statistics*. 2nd ed. New York: Dover.
- Simon, H. A. 1981. *The Sciences of the Artificial*. Cambridge: MIT Press.
- Simon, H. A. 1986. *Decision Making and Problem Solving*. Washington, DC: National Academy Press.
- Simon, H. A. and A. Newell. 1971. "Human Problem Solving: The State of the Theory in 1970." *American Psychologist* 26 (2): 145–159.
- Stich, S. 1985. *From Folk Psychology to Cognitive Science: The Case Against Belief*. Cambridge, MA: MIT Press.
- Szalai, J. 2021. "The Potential Use of Artificial Intelligence in the Therapy of Borderline Personality Disorder." *Journal of Evaluation in Clinical Practice* 27 (3): 491–496.
- Thompson, E. 2007. *Mind in Life: Biology, Phenomenology, and the Sciences of Mind*. Cambridge, MA: Harvard University Press.
- Timms, R. and D. Spurrett. 2023. "Hostile Scaffolding." *Philosophical Papers* 52 (1): 53–82.
- Von Neumann, J. and O. Morgenstern. 1944. *Theory of Games and Economic Behavior*. Princeton, NJ: Princeton University Press.
- Viale, R. 1997. *Cognitive Economics*. Milan: Università Bocconi.
- Viale, R. 2012. *Methodological Cognitivism: Mind, Rationality and Society*. Heidelberg: Springer.
- Viale, R. 2013. *Methodological Cognitivism: Cognition, Science and Innovation*. Heidelberg: Springer.
- Viale, R. 2024. "Enactive Problem Solving: An Alternative to Decision Making." In G. Gigerenzer et al. (eds.), *Companion to Herbert Simon*. Cheltenham: Elgar.
- Viale, R. Forthcoming. "What Is Cognitive Economics?" In R. Viale (ed.), *Handbook of Cognitive Economics*. Cheltenham: Elgar.
- Viale, R., S. Gallagher and V. Gallese. 2023. "Bounded Rationality, Enactive Problem Solving, and the Neuroscience of Social Interaction." *Frontiers in Psychology* 14: 1152866.

- Wallis, P. 2022. “An Enactivist Account of Mind Reading in Natural Language Understanding.” *Multimodal Technologies and Interaction* 6 (5): 32.
- Weiss, D. J. and J. Shanteau. 2021. “The Futility of Decision Making Research.” *Studies in History and Philosophy of Science* 90: 10–14.