THE SEARCH FOR
A THEORY OF COGNITION

Early Mechanisms and New Ideas

Edited by

Stefano Franchi and Francesco Bianchini

CONTENTS

Part I: The cybernetic suburb

Part II: AI's peripheries

Part III: Margins of computations

Part IV: At the thresholds of computability

Six

ARTIFICIAL INTELLIGENCE WITH A NATIONAL FACE: AMERICAN AND SOVIET CULTURAL METAPHORS FOR THOUGHT

Slava Gerovitch

In their proposal to the Rockefeller Foundation for the pioneering conference at Dartmouth College that in 1956 founded Artificial Intelligence (AI), John McCarthy, Marvin Minsky, Nathaniel Rochester, and Claude Shannon, wrote that "every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it" (Norman, 2005, p. 710). Decades later, the aspirations of AI remained the same—to grasp the universal principles of thought in order to implement them in a computer. In 1984, Patrick Winston articulated the goals of AI research as follows:

> One central goal of Artificial Intelligence is to make computers more useful. Another central goal is to understand the principles that make intelligence possible .... Artificial Intelligence excites people who want to uncover principles that all intelligent information processors must exploit. (1984, p. 3)

At the same time in the Soviet Union, a budding AI community formulated goals that sounded remarkably similar, "to understand how the human being thinks, what are the mechanisms of thought" (Smirnov, 1975, p. 3). Scientists on both sides of the Iron Curtain understood AI research as a search for fundamental principles of human thinking.

American and Soviet scientists believed that a general, universal, ahistorical mechanism of human thought existed. Yet, as these scientists themselves belonged to different cultures, they had distinct and culturally specific intuitions of human thinking. The "human beings" whom they took as universal categories were, in fact, people who belonged to specific cultures. Their AI models reflected the specificity of their cultures. According to the recent feminist critique of Allen Newell and Herbert Simon's general theory of human problem solving, their "view from nowhere" was in fact "a view from somewhere." A limited set of problem-solving protocols that served as a foundation for the theory represented "the behavior of a few technically educated, young, male, probably middle-class, probably white college students working on a set of rather specialized tasks in a U.S. university in the late 1960s and early 1970s" (Adam, 2005, p. 339).

Commonly accepted patterns of behavior—actions that we perceive as typical and normal—and also different strategies of handling daily situations that we call "common sense" are the foundation of everyday practice in any society. John McCarthy has famously called AI systems "programs with common sense," implying that a fundamental universality of common sense knowledge underlies human thinking (1968). As the anthropologist Clifford Geertz suggested, however, common sense is "historically constructed and … subjected to historically defined standards of judgment. It can … vary dramatically from one people to the next. It is, in short, a cultural system" (1983, p. 76). Geertz warned against "sketching out some logical structure [common sense] always takes, for there is none" (*ibid.*, p. 92), thus unfortunately undermining McCarthy's basic premise.

Everyday practice serves as a mediator for constant exchange of cultural symbols and shapes the cultural vocabulary for any given group. As the anthropologist Hervé Varenne has put it, "American culture is whatever one cannot escape in the United States" (1986, p. 6). Soviet Union had its share of things people could not escape, too. For American urban professionals, cultural influence "spreads through a diffuse network of everyday experiences that range from reading *The New York Times* to using bank cards on automatic teller machines to watching MTV" (Hayles, 1990, p. 4). Soviet intelligentsia's everyday experience looked quite different. They never read *The New York Times*, used automatic teller machines, or watched MTV. They read *Pravda* and underground literature, sat at Party meetings, and stood in lines in food stores. What was typical and normal to them, looked peculiar and exotic to Americans, and vice versa. Yet, even if common sense is not universal, AI models do tell us something—if not about the fundamentals of human thinking in general, then perhaps about specific cultural constructions of common knowledge.

Cultural influence manifests itself not only through typical patterns of behavior and strategies of everyday life, but also through language—via the metaphors by which we live and think, including thinking about thought itself (Lakoff and Johnson, 1980). In this chapter, I will discuss the different cultural metaphors for thought prevalent among American and Soviet intellectuals and explore their connections with specific AI systems. I will maintain that deep cultural factors lay beneath considerable differences in the approaches to AI that American and Soviet scholars developed. While looking for general principles of thinking and behavior, AI specialists implemented in their models their own cultural stereotypes.

1. Different Cultural Metaphors of Freedom: Choice vs. Creativity

The American anthropologist William Beeman has made the following keen observation:

> There are few things all Americans can agree upon. Near the top of the list, along with mother and apple pie, lies "freedom of choice." The power of

> this concept is shown by its ability to be adopted without public comment by groups as diverse as the Advertising Council, pro-abortion forces, and the anti-busing movement. This concept comes close to being sacred for Americans. (Beeman, 1986, p. 59)

Americans face the necessity of choice every day. If we consider such an everyday situation as shopping, the main problem for American customers is making the right (we may say, "healthy") *choice* among an appealing variety of foods and goods. The ability to make the right *choice* is also a crucial part of academic training in the United States. College students choose most of their courses from among a great variety of courses that are offered; routine multiple-choice tests require selecting one right answer among several possibilities. The computer scientist Edward Fredkin in his youth developed "an exact science" of passing such tests. He recalls: "We had a multiple-choice exam, and my bet was that I could pass if I were only shown the answers and none of the questions. Some of the answers went, like, 347, 492, 513, 629—and I did pass the test" (Wright, 1988, p. 15). The AI pioneer Marvin Minsky was also quite good at passing tests, and on this basis he was sent to an experimental school for gifted children (J. Bernstein, 1982, p. 23).

By contrast, most everyday situations in the Soviet Union left the citizen no choice at all. Higher education curriculum prescribed a fixed, pre-determined sequence of courses for every major. The only choice students had was in selecting a preferred athletic activity. Multiple-choice tests were rare. Instead, the student had to spell out all the intermediate steps, and if the algorithm was inefficient (or just different from the textbook's), the grade was lowered, even if the answer was correct. And finally, the Soviet way of shopping posed a different sort of problem for the customer. The problem was not what to choose, but how to find anything at all. With the shortage of many foods and household items, people could obtain sought-after products only via back channels. An ordinary Soviet citizen had to create a unique, long chain of informal social interactions through a network of friends, relatives, friends of relatives, and relatives of friends, so that s/he could find a desired washing machine or a television set at the other end (see Ledeneva, 1998).

The Russian cultural tradition had long treated the notion of choice among preset alternatives with suspicion. While the Americans valued the "freedom of choice," the Russians often spoke of the "burden of choice." The Russian philosopher Nicholas Berdyaev viewed choice as restrictive, instead of liberating:

> Philosophical textbooks generally speak of freedom as identical with "free will," that is to say, as the possibility of choice, of turning right or left .... For me freedom has always meant something quite different. Freedom is first and foremost my independence, determination from within and creative initiative .... The very condition of choice may result in a sense of repression,

> indecisiveness, or even in the complete disappearance of freedom on the part of man. Liberation comes when the choice is made and when I have begun to create. (1965, p. 173).

The Soviet government, even though it forced Berdyaev into exile in 1922, shared his disdain for the freedom of choice and did everything it could to free Soviet citizens from the burden of choice, political or otherwise. Remarkably, the dissidents, for all their opposition to the Soviet regime, felt equally uneasy about the notion of choice. In his 1959 *Free Philosophical Treatise*, published abroad, Aleksandr Esenin-Vol'pin wrote that, "we value freedom conceived as the possibility of choice; but certainly not because we like to choose (the necessity of a choice is sometimes simply horrible and almost always unpleasant!), but because we desire to choose without compulsion." Trying to find a "scientific expression for the concept of freedom," he suggested that "'freedom' cannot always be understood as a possibility of choice" (1961, pp. 135, 137). Facing a narrow spectrum of alternatives, the Russians often perceived it as inherently incomplete. They felt that to widen the spectrum and to create a new, yet unknown solution, was necessary.

Cognitive psychological theories that American and Soviet scholars developed reflected the different cultural values of choice and creativity. The Soviet psychologist Andrei Brushlinskii, for example, rejected the idea that thinking involved a choice among preexisting alternatives. He argued that true thinking must produce a new alternative:

> Actual live thinking, for example, solving a task or a problem, always takes the form of prediction of an initially unknown, future solution. This prediction … makes the act of choosing among alternative solutions unnecessary. The situation of choice among alternatives is characteristic of the formal-logical correlation of the *products* of the mental process of thinking, but it is not characteristic of the *process* itself. (1979, p. 62, emphasis added)

The American cognitive psychologist Jerome Bruner, by contrast, described concept attainment as a process whose every step "can be usually regarded as a choice or decision between alternative steps" (1973, p. 151). Bruner's work showcased the "cognitive revolution" in psychology that we closely associate with the work of the American AI pioneers Herbert Simon and Allen Newell, who placed choice at the heart of their "heuristic search" model of intellectual activity.

AI specialists in the Soviet Union and in the United States sometimes drew on psychological theories, and sometimes psychologists drew on AI models. More habitually, however, AI specialists ignored psychologists' findings, believing that knowledge should flow from AI to psychology, not the other way around.[1] When AI and psychology agreed, this often happened because they both relied on the same cultural stereotypes.

2. Bureaucratic Man: Striving for Control of Social Environment

Herbert Simon, trained at the University of Chicago as a political scientist, early on strove to find the "atomic phenomena of human behavior in a social environment" (1951, quoted in Crowther-Heyck, 2005, p. 5). Explicitly referring to everyday experience, he found this atom of intellectual activity in the act of choice:

> [N]one of us is completely innocent of acquaintance with the gross characteristics of human choice, or of the broad features of the environment in which this choice takes place. I shall feel free to call on this common experience as a source of the hypotheses needed for the theory about the nature of man and his world. (Simon, 1955, p. 100)

Simon drew on a wide array of mathematical theories that offered different formalizations of choice in well-structured environments—such as econometrics, game theory, operations research, utility theory, and the statistical decision theory—disciplines that his biographer Hunter Crowther-Heyck has termed "the sciences of choice." All these theories assumed the act of choice to be free and rational: an individual acted upon his environment, but the environment did not affect the individual's goals or preferences.

Simon also borrowed from another set of disciplines—sociology, social psychology, anthropology, and political science. These "sciences of control," by contrast, emphasized the malleability and docility of an individual that the social environment mold and over which group and societal pressures hold sway. The "administrative man" of the sciences of control appeared utterly incompatible with the "economic man" of the sciences of choice.

Simon's experiences with the International City Managers Association and the Bureau of Public Administration shaped his faith in liberal reform through expert management and led him to believe in the active role of the intellectual in shaping the social environment. Drawing on the sciences of choice and the sciences of control, Simon developed a theory of "bounded rationality." We could solve complex problems by reducing them to a limited set of alternatives and choosing rationally among them. Belonging to an organization limited an individual's choices and thus made rational decision-making possible: "The rational individual is, and must be, an organized and institutionalized individual" (1997, p. 111). Eventually, Simon arrived at a comprehensive "bureaucratic worldview":

> Simon came to define mind and machine, organism and organization, individual and institution, all as highly specialized yet tightly integrated hierarchical systems, each locked in a continual struggle to adapt to its environment as best it could, given its limited powers. To him the mind and the computer were model bureaucracies, and a bureaucracy was a model mind. (Crowther-Heyck, 2005, p. 7)

In his 1956 paper, "Rational Choice and the Structure of the Environment," Simon used the metaphor of a maze to introduce a mathematical model that described how an organism, although based on incomplete information, could, making a sequence of rational choices at branch points, meet a multiplicity of needs (1956, pp. 129–138). This was not merely a convenient description. Extrapolating from his personal experience to the entire humanity, Simon regarded a sequence of rational choices as a "universal" model, a philosophy of life:

> A philosophy of life surely involves a set of principles .... Principles can provide a book of heuristics to guide choice at life's branch points, a thread to keep one on the right path in the maze .... In this chapter, I have been describing my life, and also my personal life philosophy, but I have also been describing the life of Everyperson. (1991, pp. 360, 363)

In the 1950s-1960s, Newell and Simon developed the "heuristic search" approach that quickly became the dominant paradigm for American AI research. According to their model, problem solving activity consisted in finding a path from the initial to the goal state within the *problem space*. This space looked like a branching tree or a labyrinth; at every step of the process, the problem-solver had to choose one of the alternatives—one of the branches that diverged at the point of choice. In the absence of complete information about the labyrinth, or if the labyrinth was too large to make a feasible calculation, Newell and Simon suggested using *heuristics*—rules of thumb—that would help us make the right choice. They believed labyrinth search to be a universal model of intelligence and considered their computer program, the "General Problem Solver," to be a general "theory of human problem-solving" (1963, p. 279).

As Simon and Newell's conceptualization of human behavior grew increasingly formal, the model situations they were drawing on became increasingly circumscribed and regulated: from semi-independent decisions by workers in big organizations to semi-automatic actions of machine-bound operators in air defense control centers to chess players' limited repertoire of permissible moves. In different computer implementations of the heuristic search model—the theorem-proving Logic Theorist, a chess-playing program, and the "universal" General Problem Solver—Newell and Simon tended to focus on situations with complete, unambiguous, computer-friendly descriptions.

Newell and Simon redefined the problem of choice: they no longer spoke of "making decisions," but of "solving problems." If the decision-maker could consider different goals, the problem-solver had to focus on the assigned problem. Decisions turned into "a less contentious, less political, process of allocating 'processor time' to different tasks. Choices were now less decisions about which set of values to accept and more decisions about what set of data to process" (Crowther-Heyck, 2005, p. 214). Politics was reduced to technology: the liberal aspiration to control and purposefully transform the environment turned into a purely technical task of simplifying search in a labyrinth.

When elaborating her cultural "grammar" of American storytelling, the anthropologist Livia Polanyi has emphasized "control" as one of the most important categories of American life. "Proper people" as everyday conversations portray them, are those who "can *control* the world sufficiently to be happy and have power" (1985, p. 140, emphasis in the original). By contrast, in the Soviet case, instead of something you could control, social environment was something that could potentially control you. A Soviet cultural grammar would probably say that "proper people are those who can sufficiently escape control by the world to be happy." In a formalized and abstract form, Soviet AI models translated the everyday struggle for intellectual autonomy of independent-minded intelligentsia.

3. Soviet Controversy over "Thinking Machines"

The idea that computers could perform intellectual tasks stirred a serious controversy in the Soviet Union in the early 1950s. In the paranoid Cold War context, they often viewed scientific and technological innovations coming from the West with great suspicion. In reaction to the popular discussions of "thinking machines" in the West, the Soviet press condemned this idea as a potential technological threat and an ideological subversion. Soviet journalists berated the capitalists for their hidden agenda to substitute a robot for a striking worker and to replace a human pilot who refused to bomb civilians with an "indifferent metallic monster." Soviet philosophers, for their part, attacked the idea of "thinking machines" as "idealistic" (detaching thought from its material basis in the brain) and "mechanistic" (reducing thought to computer operations). Soviet critics lumped all controversial uses of computers under the rubric of "cybernetics" and labeled this field a "reactionary, idealistic pseudo-science." Despite its glaring logical contradictions—they portrayed cybernetics as idealistic and mechanistic, utopian and dystopian, technocratic and pessimistic, a pseudo-science and a dangerous weapon of military aggression—the campaign had a serious impact on Soviet research. As a result of media frenzy, work on "thinking machines" became ideologically unacceptable, and scientific calculations became the limit that early Soviet computer applications had to observe.[2]

But the anti-cybernetics campaign did not dampen the interest of Soviet scientists in computer systems that could perform intellectual tasks. All the first Soviet large electronic digital computers had been installed at defense research institutions, which were relatively immune from ideological pressure and which also gave their employees access to most recent Western publications. Early Soviet champions of cybernetics and AI came largely from these institutions. The mathematician Aleksei Liapunov led the computer programming department at the Division of Applied Mathematics of the Mathematical Institute of the Soviet

Academy of Sciences in Moscow. This Division (after 1966, the Institute of Applied Mathematics) performed calculations for the Soviet nuclear weapons and rocketry programs. These calculations were double-checked against the results that the Computer Center No. 1 of the Ministry of Defense, which was under the supervision of the computer specialist Anatolii Kitov in charge of research and development, obtained. In 1955, taking advantage of the thawing political climate after Stalin's death, Kitov and Liapunov teamed up with the leading mathematician for the nuclear weapons program Sergei Sobolev and published an article in the journal *Problems of Philosophy*, in which they publicly dismissed ideological charges against cybernetics and effectively legitimized research in this field. Trying to calm down the controversy over cybernetics, they avoided such odious terms as "thinking machines" and preferred to speak vaguely of "the theory of automatic high-speed electronic calculating machines as a theory of self-organizing logical processes similar to the processes of human thought" (Sobolev, Kitov, and Liapunov, 1955, p. 136).

As the cybernetics movement grew in strength, bringing under its umbrella all sorts of mathematical models and computer applications in "cybernetic biology," "cybernetic physiology," "cybernetic linguistics," "cybernetic economics," and many other fields, the pendulum of Soviet public attitudes toward "thinking machines" swung in the other direction.[3] The Soviet press began extolling the intellectual abilities of the computer, portraying it as an all-powerful magic tool for solving any problem. Articles entitled "Thinking Machines" and "Bordering on Science Fiction" mushroomed on the pages of newspapers and popular magazines. Journalists quickly dismissed the previous ideological critique by claiming that it applied only to capitalist society:

> If in the capitalist world the introduction of "thinking" machines means the growth of unemployment, exploitation of workers, and fear of the future, in a socialist society, by freeing people from hard, uninteresting work, machines would provide an opportunity to focus on something lofty and joyful—to think, to create, and, in particular, to create new "thinking" machines. (Petrovskii, 1956, pp. 23–24)

Projects exploring intellectual capabilities of computers proliferated. Liapunov set up a machine translation group at the Division of Applied Mathematics, Kitov organized a similar effort at the Computer Center No. 1, Mikhail Bongard of the Institute of Biophysics in Pushchino near Moscow proposed an original approach to pattern recognition, and Nikolai Shanin at the Leningrad (now St. Petersburg) branch of the Mathematical Institute developed a theorem proving program.[4]

Early Soviet research on AI was done largely out of pure enthusiasm, without sustained institutional support. Despite the media hype, the Soviet government showed little interest in "thinking machines." The leaders of the cybernetics movement distanced themselves from AI aspirations, trying to cultivate an image of the computer as an efficient tool and not as an autonomous agent. The chairman of the Cybernetics Council of the Soviet Academy of Sciences, Engineer Admiral Aksel' Berg publicly proclaimed that electronic computers "will be increasingly providing help to man, but will never replace him and *will never think*" (1964, p. 87, emphasis in the original). Computer time remained in short supply, and supervisors did not look favorably on computer programmers' attempts to divert valuable computational resources to investigate problems that aroused their intellectual interests. For example, Aleksandr Kronrod, head of the computational mathematics laboratory at the Institute of Theoretical and Experimental Physics in Moscow, used his control over the Institute's sole M-20 computer to launch work on a chess-playing program. The Institute's administrators viewed spending computer time on games as "overindulgence" (Landis and Yaglom, 2001, p. 1000).

Language reflected the tenuous position of Soviet AI. To stress the metaphorical meaning of the expression "thinking machines," quotation marks always accompanied it. The very term "artificial intelligence" remained controversial, and researchers avoided it. They preferred more neutral-sounding vocabulary, such as "cybernetic psychology," "the study of information processes," or "heuristic programming" (Boiko et al., 1967). In 1964, when the mathematician Dmitrii Pospelov and the psychologist Veniamin Pushkin brought together computer specialists and psychologists interested in AI for a regular colloquium at the Moscow Power Engineering Institute, they named their field "psychonics." The psychonics group directly challenged the Simon-Newell model of thinking and put forward an alternative approach.

4. The Freedom Not To Choose

The term "psychonics" was formed by analogy with bionics. While specialists in bionics hoped to imitate the "design" of living organisms in engineering systems, Pospelov and Pushkin aspired to use psychological knowledge to construct intelligent computers. Pushkin conducted several eye-movement tracking studies of chess players and concluded that each player, rather than searching for the solution in a preset problem space, constructed a different mental model of the position on the board. He asserted that initially we do not structure the human problem space as a tree, and that the process of finding a solution involved creating a new problem space and not, as Newell-Simon's labyrinth model did, "pruning useless branches" (Pushkin, 1971, p. 204).

Soviet AI specialists disliked the labyrinth model not for its inefficiency, but for its departure from their cultural expectations. Even without knowing the conceptual origins of the "General Problem Solver," they associated it with the "bureaucratic apparatus" of labyrinth search. While some followed Newell and Simon's logic and asserted that "the human being thinks by exhaustive search," many others suggested alternative models, for example, thinking as a chain of associations (Kronrod, 2001, pp. 168, 139).

Pushkin and Pospelov conceptualized thinking not as a search, but as a reflection of/on the problem. They argued that we often formulate the descriptions of the current situation in different terms than the description of the goal. In the case of chess, for example, the location of specific pieces on the board describes the initial position while the goal state—a checkmate—requires a higher-level description that involves the inability to move the checked king. The human chess player must be able to go back and forth between low-level and high-level descriptions and build and manipulate different mediating models of the situation. Pushkin and Pospelov argued that situation modeling, not labyrinth search, was the basic intellectual procedure: "Among all the existing words and notions used to describe productive thinking, the most adequate, the most suitable is the Russian word *soobrazhenie* (reflection/imagination) .... The solution *reflects* the situation, based on the *images* or models of its elements" (1972, pp. 140–141, emphasis added).

For Pospelov and Pushkin, humans manifested their creativity in abandoning the old labyrinth, re-conceptualizing the problem, and constructing a new problem space. For example, we could not construct four equilateral triangles out of six matches if we sought the solution on the plane. Constructing a new labyrinth of solutions—in the three-dimensional space—would produce the answer (1972, p. 139).

Where Newell and Simon started with a ready-made structure of the problem, Pushkin and Pospelov suggested that structuring the problem was an essential intellectual step in finding a solution. Building an adequate model of the situation was more important than powerful search algorithms. Pushkin and Pospelov proposed a "semantic language" for formal descriptions of the situation at different levels of generality and developed a system for building relational situation models. They called it *giromat*, borrowing the term from the Polish science fiction writer Stanislaw Lem, who had imagined an intelligent machine capable of adapting its structure to the nature of the problem it was trying to solve.[5] Pospelov and his team implemented this approach, which combined technological and human elements, in computer systems controlling the loading operations in a sea port and other industrial operations.[6]

Pospelov and Pushkin's critique of the labyrinth theory echoed the Soviet cultural perception of choice as a restraint on creativity. In one of his books, Pospelov quoted the Polish dramatist Slawomir Mrozek's mocking exposure of choice as a fetish of freedom. In Mrozek's 1961 play, *Strip-tease*, the following dialog occurs between two people sitting in a room in front of an open door:

> MAN A. What is freedom? It's the possibility of choosing. As long as I am sitting here and know I can leave via that door, I am free. On the other hand,

> the moment I get up and go out, I have made a choice. In other words, I have limited my field of action. I've forfeited my freedom. I become a slave of my going out.
>
> MAN B. But sitting here and not going out, you're also making a choice. You're choosing sitting and not going out.
>
> MAN A. You're wrong. While I'm sitting, I can still go out. But by going out, I will have excluded the possibility of sitting.
> (Mrozek, 1967, p. 55, quoted in Pospelov, 1986, pp. 13–14)

Freedom consists in not making a choice, for any choice is a false choice. Interestingly, Mrozek made *his* choice: he left Poland two years after publishing *Strip-tease* and settled in France. For Eastern bloc intellectuals, the rigidly structured labyrinth of choices that the government offered seemed overly restrictive. Some chose to emigrate; some, like Pospelov and Pushkin, chose to expose the limitations of choice-driven behavior and to create new problem spaces.

5. Intellectuals under an Oppressive Regime: Striving for Autonomy

The historian Oleg Kharkhordin has argued that in the late 1950s-early 1960s the Soviet government changed its tactics of political control of the population. Moving away from the punitive and haphazard terror of the Stalin era, the Khrushchev leadership created a system of "voluntary" mass organizations and entrusted it with crime prevention and workplace discipline enforcement. As a result, a pervasive rational system of preventive control among peers emerged, in which the entire population was enlisted in mutual surveillance. Everyone was watching and being watched at the same time:

> People's patrols … constituted a fearful core of mutual surveillance, which one joined from time to time in order to enforce the standards that one otherwise reluctantly suffered when they were imposed by others. The people's patrols were more meticulous and closer to the people: they were the people policing itself, and thus escape was hardly possible from their omniscient gaze and omnipresent power. (Kharkhordin, 1999, p. 286)

Soviet intellectuals were watched (and watched others) especially closely, and they developed sophisticated strategies of living under the surveillance. Recent studies of Soviet intelligentsia undermine the Cold War stereotypes of the Soviet scientist as blindly supporting or passively resisting government policies (see Gordin, Hall, and Kojevnikov, 2008). A more typical figure would be a physicist working on nuclear weapons during the day and reading underground literature at night (Rassadin, 2004, p. 217). Interested in results, the government allowed the scientists some intellectual license, as long as it was limited to their subject of study. The historian David Holloway has called nuclear weapons laboratories "islands of intellectual autonomy" (Holloway, 1999, p. 175). One theoretical physicist later recalled:

> [P]hysicists constituted a privileged caste, an aristocracy. There were fewer controls on our freedom than on those of any other member of Soviet civil[ian] society. The only laws we felt restricted by were those relating to the conventions of scientific work. Relatively speaking, we were free people. (Mark Azbel, quoted in *ibid.*, p. 187)

Mathematicians and computer specialists working on defense projects enjoyed a similar privileged status. As priests in a temple of an all-powerful goddess, the Computing Machine, they carved out areas of intellectual autonomy in the climate-controlled, limited-access rooms that housed the mammoth-size computers of the first generation. The mathematicians Izrail Gelfand and Mikhail Tsetlin of the defense-research-oriented Division of Applied Mathematics used their portion of intellectual freedom to engage in a study of the central nervous system.

In 1958, Gelfand and Tsetlin organized an informal regular seminar on mathematical models in physiology that began its meetings in the privacy of Gelfand's apartment and later continued for many years at Moscow University (Ivanov, 1998, p. 568). Neurophysiologists traditionally assumed that different nodes within the central nervous system coordinated their activities via a complex system of interconnections. This assumption, however, baffled mathematicians: in a large system, the number of connections would grow so rapidly that any mathematical model would become too complex. Tsetlin and Gelfand, by contrast, proposed a model in which every node regarded the activity of all the other nodes as changes in its environment. They showed that individual nodes did not have to interact directly, but could merely observe external changes and follow a simple adaptive algorithm that minimized their interactions with the environment. This strategy resulted in the purposeful behavior of the system as a whole, if we defined purpose as minimization of the system's interaction with its environment. In this model, purposeful behavior of the whole system did not require great complexity from its subsystems. All individual parts acted quite simply: instead of building complex coordination networks, they tried to avoid interaction. Gelfand and Tsetlin called this adaptive mechanism the "principle of least interaction":

> At each moment, the subsystem solves its own "particular," "personal" problem—namely, it minimizes its interaction with the medium; therefore, the complexity of the subsystem does not depend on the complexity of the entire system .... [O]ur mathematical models allow us (to a certain degree) to imagine the interaction of the nerve centers without considering the complex system of links and the coordination of their activity. (Tsetlin, 1973, pp. 150–152)

In Gelfand and Tsetlin's model, all subsystems acted independently, without direction from a single command center. Given a "payoff function," the individual parts (nervous centers) developed optimal strategies on their own. Their actions affected environmental inputs for all the other parts and led to improvements in individual strategies. Once all the parts had minimized their interactions, they achieved an overall coordination, and this resulted in the optimal strategy for the entire (nervous) system. Tsetlin and Gelfand concluded that the principle of least interaction grounded the purposeful behavior of the central nervous system (*ibid.*, p. 152).

Tsetlin and Gelfand noted that their notion of purposeful behavior as striving for minimal interaction was similar to the British cybernetician W. Ross Ashby's concept of intelligent activity (*ibid.*, p. 150). In the 1950s, Ashby had argued that the machine he designed called the *homeostat*, which did "nothing more than run to a state of equilibrium," could model "many intricate and interesting ways of behaving"(1956, p. 84). By randomly reconfiguring itself, the *homeostat* "learned how to minimize pain in its interactions with the environment" (Pickering, 2002, p. 419). Remarkably, while Tsetlin and Gelfand described minimal interaction in positive terms such as "independence" and "autonomy," British cyberneticians often ridiculed Ashby's idea of thinking as a search for equilibrium; Grey Walter "sarcastically dubbed the homeostat *machina sopora*, the sleep machine" (*ibid.*, p. 417).

The peculiar definition of purposeful behavior as minimization of the system's interaction with its environment resonated with the Soviet intelligentsia's drive to preserve maximum intellectual autonomy. Tsetlin argued that his model of the nervous system had the advantage of non-individualized control: the designer did not need to tell every node in the system what it was supposed to do; the system used its freedom of maneuver to self-organize under most general conditions. At a lecture before the Physiological Society in February 1965, Tsetlin explicitly brought up a comparison of free and forced labor to highlight the advantages of self-organization:

> [T]he work of prisoners is more expensive than that of free men, even though the former are much worse fed and clad, and they work no less. The point is not only that the efficiency of prisoners is lower, but that a prisoner must be fed, clad, and watched by someone else. With a free person the matter is different: … my manager … doesn't have to think when to change my shoes or linen or what to do with my children. (Tsetlin, 1973, p. 125)

The MIT biophysicist Murray Eden once remarked, "One wonders whether it is a reflection of cultural or social differences that Tsetlin chose to study cooperative phenomena in choosing 'expedient' behavior, while American game theory focuses on competition among the players" (1973, p. xi). Tsetlin's model, strictly speaking, was not a mathematical implementation of socialist ideals. It reflected the intelligentsia's peculiar position within the Soviet system, in which "cooperative phenomena" emerged out of individuals' efforts to escape the control that the environment (the state) or other individuals ("people's patrols") exerted. Eden's suggestion of social and cultural roots of different approaches to game theory, however, is worth exploring in greater detail.

6. Individualistic Games of Capitalism: Minimizing Risks in Pursuit of Payoffs

Different cultures cultivate different patterns of individual and collective behavior. Livia Polanyi has contrasted the different conceptualizations of individual rights in American and Italian societies:

> Some time ago, I had a discussion with an Italian friend. We were talking about *rights*. He insisted that *rights* are given to the *individual* by *society* .... The *individual*, in my American scheme, *has rights* by virtue of being an *individual* person .... The distinction is a crucial one. In one case, the Italian, *society* is the core of the world and must be protected. It is also granted the role of actor much more than in the American scheme in which the *individual* is an actor at the center of the world and *society* is, as an American friend put it, "just a bunch of *individuals* anyway." ... [I]t seemed that everything which could be said about *society* was somehow reducible to the relationship of one *individual* to *other individuals*. (1985, pp. 107–108, emphasis in the original)

The concept of social interaction as a competition between self-interested, rationally calculating, yet cautious opponents was a model for the Hungarian-born American mathematician John von Neumann in his 1926 axiomatic formalization of two-person, zero-sum, games that have a finite number of "strategies" (complete plans of the game). He proved the minimax theorem and asserted the existence of an optimal "mixed," or randomized, strategy that would minimize the maximum loss and guarantee winning the "value of the game." He believed that the minimax strategy captured some fundamental aspect of human rationality: "[A]ny event—given the external conditions and the participants in the situation (provided that latter are acting of their own free will)—may be regarded as a game of strategy if one looks at the effect it has on the participants" (John von Neumann, 1928, quoted in Leonard, 1995, p. 735).

Von Neumann's biographer Steve Heims has traced von Neumann's formalism back to his perception of the world as filled with ruthless competitors who viewed all the other players as cunning enemies:

> The harshness of this Hobbesian picture of human behavior is repugnant to many, but von Neumann would much rather err on the side of mistrust and suspicion than be caught in wishful thinking about the nature of people

> and society. His temperament was conditioned by the harsh political realities of his Hungarian experience. The recommended style of "playing the economic game," the emphasis on caution, on calculation of expected consequences, the whole utilitarian emphasis aptly expresses the characteristic ideals of the middle class in capitalist societies. (1980, p. 296)

Von Neumann himself viewed the self-interested player of his game theory not merely as a useful mathematical abstraction, but as true reflection of human nature: "It is just as foolish to complain that people are selfish and treacherous as it is to complain that the magnetic field does not increase unless the electric field has a curl. Both are laws of nature" (quoted in Wigner, 1967, p. 261).

In 1944, von Neumann and his collaborator, the Austrian-born American economist Oskar Morgenstern expanded the original conceptual framework of game theory to treat problems of economics in their book, *Theory of Games and Economic Behavior.* Like von Neumann, Morgenstern sought universal principles of human behavior. In his 1941 essay, "Quantitative Implications of Maxims of Behavior," Morgenstern outlined a set of formal principles that governed behavior based on the Austrian philosophical notion of "subjective rationality" and took into account an individual's knowledge and interpretation of facts (Leonard, 1995, p. 750). The historian Philip Mirowski has suggested that "Morgenstern's individualistic bias derived from his Austrian background."[7]

Von Neumann and Morgenstern explicitly challenged deterministic decision-making that neoclassical economics enshrined and presented the "solution" of an economic game as a probabilistic "stable set" of possible apportionments of payoff among the players. As Mirowski has argued, they treated mixed strategies as "a representation of the stochastic nature of thought itself" and effectively turned minimax strategizing into "the very epitome of the abstract rationality" (1991, p. 237). Mirowski has further suggested that von Neumann and Morgenstern came to believe that game theory could "simulate the behavior of any opponent and therefore serve as a general theory of rationality," and that in their writings, "game theory and artificial intelligence tended to blur together" (1992, pp. 125, 127).

The historian Robert Leonard has described the transition from neoclassical to game-theoretical economics as "part of a general shift in science which involved, broadly speaking, the abandonment of determinism, continuity, calculus, and the metaphor of the 'machine,' to allow for indeterminism, probability, and discontinuous changes of state" (1995, p. 756). Among the indeterminism that von Neumann and Morgenstern celebrated, one thing remained stable though: the rules of the game. The game itself was defined as "simply the totality of the rules which describe it." Every player had to know the rules and could not violate them: "[T]he rules of game " are absolute commands. If they are ever infringed then the whole transaction by definition ceases to be the game described by those rules" (1944, p. 49).

The pay-off matrix that the players knew and that remained fixed throughout the game effectively embodied the rules that specified to the players which combination of strategies would result in which rewards. "It would not surprise me if the greatest contribution of game theory to social science in the end is simply the pay-off matrix itself that, like a truth table or a conservation principle or an equation, provides a way of organizing and structuring problems so that they can be analyzed fruitfully," wrote one sympathetic economist (Thomas Schelling, quoted in Mirowski, 1991, p. 248). Fixing the rules of the game not only made it possible to derive powerful formal results in game theory; it also provided an anchor for the notion of rationality: the world was too complex for deterministic analysis, but it still followed rules, so a stochastically equipped mind could still calculate an optimal set of strategies.

Game theory initially received a lukewarm reception among economists, but as a general model of decision-making it attracted interest of other social scientists. In the words of the historian David Hounshell, in the late 1940s defense analysts at the USA think tank RAND Corporation came to see game theory as "an ideal approach to solving problems of warfare and decision-making" (1997, p. 253). They embraced the notion that formalization as a zero-sum, non-iterative, two-person game was intrinsically akin to the tactical and strategic problems of the Cold War confrontation with the Soviet Union. RAND analysts rushed to apply game theory to a wide range of problems: strategic bombardment systems/target allocation, missile fire scheduling, and psychological warfare (Edwards, 1996, p. 117). It turned out, however, that "non-zero-sum games consisting of multiple moves in which a changing number of players could learn, mix strategies, and pursue non-rational behavior eluded RAND researchers' formalization abilities," and their superiors found the results "quite disappointing" (Hounshell, 1997, p. 254).

Game theory attracted defense analysts not because it demonstrated particular effectiveness, but largely because it fitted their pattern of thought. They did not perceive the constraints of von Neumann and Morgenstern's formalism; on the contrary, they were perfectly satisfied with the simplest case—two-person zero-sum non-iterative games:

> The strategic community had no qualms about the two-person restriction, since they thought in Manichean terms of Blue versus Red, capitalists versus communists, USA versus USSR. But they also had no difficulty with the zero-sum restriction. In deployment or dogfight models, they believed, winning or losing was a simple offsetting payoff structure.... Moreover, the application of games to the face-off of atomic weapons only reinforced

> the two-person zero-sum mindset, including the all-important stipulation that the game could not be repeated. (Mirowski, 1991, pp. 245–246)

Von Neumann and Morgenstern aspired to create an abstract theory of rationality, but what resulted was a snapshot of the Cold War mindset. The historian Paul Edwards has argued that, "as a metaphor, the zero-sum game represented yet another symbolic enclosure," reflecting the "closed world discourse" of defense analysts (1996, p. 118). While economists debated whether game theory truly reflected actual economic behavior, minimax strategizing proved an excellent match for the strategic community's notion of common sense: their vision of the world as an arena of fully computable and manipulable conflicts with known risks and rewards. Defense analysts stripped game theory of its connotations of indeterminacy and reduced it to an impersonal, formalized, mechanized decision-making tool that could guarantee the winning strategy. Game theory, in its emasculated form, helped them "uphold the appearance of technical, calm, measured rationality in the face of bottomless nuclear horror" (Mirowski, 1991, p. 249).

The strategic community asserted that "the significance of game theory as a decision tool is that it eliminates guessing an opponent's intentions" (quoted in *ibid.*, p. 251). While guessing was the opposite of rational problem-solving for USA defense analysts, it was often the only option available to an intelligent decision-maker in the Soviet Union.

7. Collective Games of Socialism: A Play with Changing Rules

Scholars studying Soviet science in the late Stalinist and Khrushchev periods have remarked the ritualistic patterns of behavior typical of the scientific community. Whether scientists engaged in public discussions of the philosophical and ideological meaning of their discipline or tried to jump on the bandwagon of a fashionable intellectual trend, they had to play a game according to the unspoken rules of public behavior of a Soviet scientist (Kojevnikov, 1998; Krementsov, 1997, esp. pp. 239–248; Gerovitch, 2002). Ritual critique of ideological enemies, skillful manipulation with suitable quotes from Marx or Lenin, and ingenious translation of scientific terminology into an ideology-laden language were among the indispensable strategies of Soviet science. The outcome of debates over the validity of scientific theories often depended on the discussants' abilities to play the game. As the historian Alexei Kojevnikov has remarked, "an important thing about these games was that, in theory and often also in practice, their outcomes were not predetermined, but depended upon the play" (1998, p. 28).

The uncertainty over the rewards and punishments for specific strategies complicated the play. Frequent swings in the direction of Stalinist ideological campaigns often left slow thinkers stuck with old, outdated slogans and made them vulnerable to attacks. On the other hand, ideological campaigns had their own inertia and often became self-perpetuating. Some campaigns were not directed from above; they generated spontaneously but were perceived as intentional, as, for example, the anti-cybernetics campaign. Those scientists who could not properly decipher "signals" from above were often perplexed about the rules and direction of the most recent campaign. Stalin's personal attempts to encourage an open debate over scientific issues faced scientists' reluctance to speak out when the official position was not clearly stated (see E. Pollock, 2006).

Mikhail Tsetlin's theory of collective games of automata exemplifies the fundamental uncertainties of Soviet social games. An automaton is a mathematical model of a finite state machine that changes its state according to its transition diagram and the current input. Tsetlin interpreted an automaton as an agent acting in an environment that randomly penalized or rewarded specific behaviors. Unlike the classic von Neumann-type games, Tsetlin studied games in which the automata faced a world filled with uncertainty:

> It should be noted that the automaton games are discussed here from a viewpoint that differs from the one accepted in game theory. Indeed, it is normally assumed in the latter that the game is defined by a system of pay-off functions previously known to the players.... We thought it interesting to consider games played by finite automata having no *a priori* information about the game, and being forced to shape their strategies for each successive replay in the course of the game itself. (1973, p. 6)

In Tsetlin's games, "the players have practically no information about the game. They are ignorant of the number of other players involved, of the situation at any particular moment and even of what kind of game they are actually playing" (Varshavskii and Pospelov, 1988, p. 97). Tsetlin informally compared his model of an agent operating in an environment with unknown and changing rules to a "little animal in the big world" (*ibid.*, p. 132). His friend, neurophysiologist Nicholas Bernstein used a similar metaphor to describe the fundamental uncertainties of intellectual activity: "To use a metaphor, we might say that the organism is constantly playing a game with its environment, a game where the rules are not defined and the moves planned by the opponent are not known" (1967, p. 173).

Tsetlin discovered that in a changing environment, in which the probabilities of penalties and rewards varied over time, automata that did not have too many states were the most successful. In other words, if the rules of the game constantly changed, it was not a good idea for the automaton to remember too much of its history. The more dynamic the environment, the shorter was the optimal depth of the automaton's "memory."

In his study of collective "distribution games," Tsetlin presented a thinly veiled commentary on the economic strategies of individuals under socialism. First, he considered a game in which a group of automata competed for resources (rewards or payoffs) by choosing different strategies. He designed automata that were completely unaware of the relative strengths of different strategies, but would eventually settle on the optimal strategy by reacting to rewards from their environment. Tsetlin showed, however, that they could increase the average gain if the automata played a game-theory version of socialism—a game with a "common fund" in which all gains and losses of individual automata were added up and then equally shared. The drawback was that the common fund camouflaged the link between individual contribution and reward and placed greater demands on the memory capacity of individual automata. We could "reap the benefits of a common fund procedure starting from a certain level of complexity" of automata memory, he concluded. "If the memory capacity is below this threshold, the introduction of a common fund reduces the average gain" (Varshavskii and Pospelov, 1988, p. 100).

In informal discussions, Tsetlin mockingly translated this rule into the clichéd parlance of the Soviet ideological discourse as "the negative effect of wage-leveling in the conditions of the low consciousness of workers" (*ibid.*, p. 101). The Soviet press often blamed the low quality of consumer products on workers' "low level of consciousness." Soviet propaganda routinely called on the workers to raise their consciousness and to work harder for a common fund. Tsetlin provided a mathematical formalization of this ideological dogma, calculating the precise memory capacity ("consciousness level") that we needed to find an optimal strategy in a game with a common fund.

Tsetlin's analysis implied that different types of games (or economic systems), to achieve the maximum payoff, required different mechanisms of decision-making. In a regular distribution game, the achievement of an overall optimum required each player to pursue individual gains. In a redistribution ("common fund") game, local level decision-making had to be far more complex in order to achieve an overall optimum. One imaginative interpretation of this result was that "capitalism is more profitable when the management system is simple and socialism is more profitable when the management system is elaborate" (*ibid.*, p. 102). Game theory appeared to suggest that, given roughly equal memory capacity of economic agents, socialism required greater social controls, while *laissez-faire* capitalism could work if individual agents were left to their devices. One model explicitly introduced a "manager" into the collective of automata, whose goal was to maximize the total gain. By punishing low-earning automata, the manager could eventually prompt them to occupy states that would be less favorable to them individually but result in greater total gain (Gaaze-Rapoport and Pospelov, 1987, p. 247). Tsetlin's followers tried to mitigate the discouraging implications of their modeling by quoting Lenin's dictum, "socialism equals accounting," but the broader social meaning of restrictions that the demands of a "common fund" placed on individuals were not lost on anyone with a sufficiently high "level of consciousness."

Tsetlin's colleagues turned his result into a fundamental principle of human thinking and behavior. Viktor Varshavskii and Dmitrii Pospelov interpreted memory capacity as a general measure of intellectual ability (1988, chap. 3). They correlated our "intellectual level" with the ability to find an optimal strategy in a game in which our immediate actions, produced at a higher level of organization, had no explicit tie to gains and losses. The writings of Soviet AI specialists paradoxically combined a thinly veiled critique of socialist redistribution and a peculiar definition of intellect as the ability to find an optimal strategy of living under socialism.

The notion of a game with unknown or changing rules was quite familiar to liberal intelligentsia. They played a cat-and-mouse game with the Soviet government, constantly challenging the fuzzy boundaries of permissible discourse. While Soviet laws ostensibly proclaimed many democratic freedoms, the actual practice was to suppress any significant dissent by placing it under the vague rubric of "anti-Soviet activity." The leading dissident, mathematician, and logician Aleksandr Esenin-Vol'pin searched not only for a "scientific expression for the concept of freedom," but also for strict definitions of all legal concepts. He exposed gross logical contradictions in the way the Soviet authorities applied liberal (on paper) Soviet laws to suppress dissent. He gained notoriety in the eyes of Soviet officials by insisting that the government literally obey its laws (see Nathans, 2007). The fellow dissident Vladimir Bukovsky even compared him to a computer that could not tolerate any ambiguity or inconsistency (1999, p. 417). Yet the Soviet government had no intention to part with its privilege to arbitrarily change and reinterpret the legal rules. In March of 1968, Soviet officials forcibly placed Esenin-Vol'pin in a mental institution, giving a whole new meaning to Simon's notion of the rational being as an "institutionalized individual."

Ninety-nine Soviet mathematicians signed an open letter protesting the action against Esenin-Vol'pin and asking for his release, thus abandoning the most expedient strategy of behavior under socialism—to minimize our interactions with the political environment—and entering in a direct confrontation with the authorities. This flagrant violation of the "principle of least interaction" had serious consequences for the signatories (Feinberg, 1996, p. 679). Many mathematicians lost their jobs, including Aleksandr Kronrod, who supported the work on a chess-playing program at the Institute of Theoretical and Experimental Physics (*Khronika tekushchikh sobytii* 1968). Georgii Adel'son-Velskii, a key member of the chess group, resigned in protest (Dobruskin and Fomenko, 2008). The entire chess group moved to another institute, where they eventually developed Kaissa, the first world champion among computer programs (Donskoi, 1999). Ironically, Kaissa's success was due to a quite efficient search algorithm that was based on an original pruning technique.

8. Rats and Butterflies

Two metaphors capture the crucial differences in the cultural stereotypes of thought and behavior reflected in AI systems that the Soviet Union and the United States implemented. Life as a maze—a labyrinth in which we must find the right path—became the central metaphor for American AI. For Soviet AI specialists, the central metaphor for decision-making was not the search in a fixed labyrinth, but the flight of a butterfly, charting its flight trajectory through random streams of air.

The metaphor of a labyrinth evoked the behaviorist pattern of B.F. Skinner's experiments on rats running T-shaped mazes and the popular American cultural image of the "rat race." Marvin Minsky started his research in the early 1950s by building a "basically Skinnerian" electronic machine that simulated rats learning to run mazes (J. Bernstein, 1982, p. 37). Allen Newell and Herbert Simon "internalized the major lessons" of behaviorism (1963, p. 280), and made heuristic search the centerpiece of their theory of human thinking and behavior. Their model of problem-solving drew on Simon's study of administrative behavior, which, in turn, took rats running mazes as a paradigmatic case: "A simplified model of human decision-making is provided by the behavior of a white rat when he is confronted, in the psychological laboratory, with a maze, one path of which leads to food" (1991, pp. 85–86). Simon insisted that the limited knowledge and intellectual capacities of a rat reflected the constraints on human rationality better than the assumption of divine omniscience and perfect rationality: "we need a less God-like and more rat-like chooser" (1954, quoted in Crowther-Heyck, 2005, p. 6)

The Soviet AI researchers Viktor Varshavskii and Dmitrii Pospelov drew on a different image—that of a butterfly. They described a system that simulated the behavior of a moth hunted by a bat. When the bat was too close and the moth could not fly away, the moth started dashing around in a chaotic flight:

> The chaotic flight is a series of passive falls with folded wings, sharp turns, loops and dives. In other words, the moth follows a trajectory which makes it more difficult for the bat to predict its location from one moment to the next. We should mention that in experiments the chaotic flight strategy saved the moth's life 70 percent of the time. (1988, p. 77)

A butterfly fluttering in a chaotic current of life and trying to escape a predator—this image was all too familiar to the Soviet scientists who were trying to preserve their intellectual autonomy.

American and Soviet AI specialists sought general principles—universal, timeless mechanisms of thinking and behavior. Their generalizations, however, relied on culturally conditioned cases. The examples that American and Soviet scientists had at their disposal, were, in fact, culturally specific patterns of social organization and decision-making. When trying to grasp universality, AI models manifested just the opposite—the specificity of cultural patterns.

Without knowing it, science often speaks with a national accent. Cultural symbolic systems can manifest themselves in scientific ideas as clearly as in literature or art. In their simulations of human thinking, AI systems truly reflect both mechanisms of reason and patterns of irrationality, individual creativity and social stereotyping, human nature and human culture.

NOTES

1. Newell and Simon, for example, prophesied in 1958 that "within ten years theories in psychology will take the form of computer programs, or of qualitative statements about the characteristics of computer programs" (1958, pp. 7–8). In 1970, Allen Newell described AI as "theoretical psychology," whose role was to generate problems for experimental psychologists to study (pp.~363–400).

2. On the Soviet anti-cybernetics campaign, see Gerovitch, 2002, chap. 3.

3. On the Soviet cybernetics movement, see *ibid.*, chaps. 4–6.

4. For a brief historical overview of Soviet AI research, see Pospelov, 2003, pp. 407–430.

5. Lem described the *giromat* in his 1955 novel *The Magellanic Cloud* (still untranslated into English).

6. For a historical overview, see Pospelov, 1986, pp. 254–258.

7. Mirowski, 1992, p. 137. On Morgenstern, see Leonard, 2004.