LUISS

LUISS GUIDO CARLI

Department of Economics and Finance

Degree Program in Economics and business

BACHELOR'S DEGREE

INNOVATION: is it possible to benefit from the unkown?

Relator	Prof. Vallanti
Correlator	Prof. Egidi
Candidato	Samuele Tavagnacco
Matricola	288101
Cattedra	Macroeconomics
Anno Accademico	2024/2025

Table of Contents

Introduction

Innovation sits at the heart of economic growth and social change, yet it emerges in landscapes that resist complete description. The most valuable breakthroughs often materialise in territories that no agent can map or price beforehand. This thesis poses a pointed question: "How can decision makers navigate benefits and risks that lie beyond the reach of standard forecasting?" The answer pursued here blends insights from complexity science with economics. Innovation is approached as the unfolding of a complex, nonlinear system whose boundaries expand endogenously, what Stuart Kauffman calls the adjacent possible. In such open ended environments, traditional optimisation loses power; decision makers frequently resort to heuristics, concise rules of thumb that trade precision for speed and tractability. Yet these same shortcuts, while sometimes indispensable, can also reinforce biases and lag behind structural change. The analysis therefore treats heuristics as context bound tools, useful but never universally sufficient.

The thesis unfolds over four cumulative chapters. Chapter 1 surveys dynamical, chaotic and complex systems, showing why long range economic prediction is structurally limited. Chapter 2 formalises the adjacent possible through combinatorial models, demonstrating how each act of novelty redefines the future choice set in ways that undermine probabilistic closure. Chapter 3 turns to decision theory, assessing when heuristics such as recognition cues or equal weight rules outperform elaborate forecasts. Chapter 4 scales the discussion to the macro level: using an agent based framework (Dosi, Fagiolo, Roventini), it evaluates whether pairing Schumpeterian innovation incentives with Keynesian demand stabilisers yields more robust outcomes than single target policy rules.

Recognising the limits of foresight neither mandates blind faith in heuristics nor calls for their wholesale rejection. Instead, it argues for disciplined pragmatism: employ simple rules when information is sparse or time is short but subject them to continual testing and stand ready to refine or abandon them as environments evolve. The chapters that follow develop this cautious but constructive stance, outlining a pathway beyond today's optimisation paradigm and illustrating how economic actors can navigate the unknown without assuming that any single method, heuristic or otherwise will suffice in all contexts.

Chapter 1: Dynamical systems

"It is better to be roughly right than precisely wrong."

-John Maynard Keynes

Introduction

Innovation is often associated with novelty, uncertainty, and the exploration of previously unknown possibilities. The central question of this thesis “Is it possible to benefit from the unknown?” motivates an examination of innovation as a dynamic process. In particular, we ask whether innovation can be understood as the evolution of a complex system over time. To address this, we draw on the mathematical theory of dynamical systems and on modern complexity science.

Innovation can be understood as a dynamic process that unfolds and transforms over time, but what does this actually imply? In this chapter, we examine the nature of dynamical systems and identify which type most accurately models the process of innovation. We begin by defining what constitutes a dynamical system, followed by a historical overview of how different forms of these systems have been recognized and interpreted across centuries. The chapter concludes with the introduction of the theory of the adjacent possible, offering a framework through which innovation can be interpreted as a specific type of dynamical system.

1.1 Dynamical systems

1.1.1. Dynamical systems, definition

The study of dynamical systems is a branch of mathematics and in the last centuries the mathematicians tried to understand every detail of these systems. The goal of studying dynamical systems is to understand processes in motion that occur in all branches of science[1]. These studies are fascinating because they do not involve a singular disciplinary subject, they vary from chemistry to biology, from physiscs to social sciences (included economics).

"*Dynamical systems theory (or dynamics) concerns the description and prediction of systems that exhibit complex changing behavior at the macroscopic level, emerging from the collective actions of many interacting components. The word dynamic means changing, and dynamical systems are systems that change over time in some way*."[2]

Examples of dynamical systems are the stock market, the solar system, the weather, the world population, the motion of a single pendulum and many others.

What is the reason for the studying of dynamical systems?

Since they change in time the scientists tries to predict where those systems are heading. Some dynamical systems are predictable others are not. We know that the sun will raise tomorrow, we know also that if we heat a pot of water, after a while it will start to bubble and boil. On the other hand, if we spill milk we can't predict exactly how it will spread and it seems impossible to predict NASDAQ a month from now. A solution can be stating that this unpredictability is the result of too many variables present in the system (for example in the economic system). In some cases this is true but this is not the complete answer.[3]

What are those systems we are talking about and what does it mean when we say that the systems evolve over time?

We can define a system as a collection of interacting agents or entities. "*A dynamical system is a system whose state (and variables) evolve over time, doing so according to some rule. How a system evolves over time depends both on this rule and on its initial conditions, that is, the system's state at some initial time.* "[4]

[1] Robert L. Devaney, *An Introduction to Chaotic Dynamical Systems*, 3rd ed. (Boca Raton, FL: CRC Press, 2022).
[2] Melanie Mitchell, *Complexity: A Guided Tour* (Oxford: Oxford University Press, 2009), [page 15].
[3] Devaney, *An Introduction to Chaotic Dynamical Systems*.
[4] Dean Rickles, Penelope Hawe, and Alan Shiell, "A Simple Guide to Chaos and Complexity," *Journal of Epidemiology and Community Health* 61, no. 11 (2007): 933–937,

This initial condition is represented by the variables that form the system. These variables can be the number of people in a population, the velocity of the wind, the price of a commodity, etc. At any given moment, the specific values of a system's variables define its current state. We can visualize a dynamical system by plotting it in a phase space, where each point represents a possible state of the system, that is, a unique combination of variable values at a certain time. As the system changes over time, it traces a path through this space. This path, known as a trajectory, illustrates how the system evolves from one state to another.

1.1.2. Dynamical systems, linear and nonlinear

Dynamical systems also posses various type of properties. There can be simple or complicated system, linear or non-linear, discrete or continuos, deterministic or indeterministic, in equilibrium or non in equilibrium, chaotic or complex. In this paragraph we will focus on the differences between linear and non-linear systems which is foundamental to understand how a system behaves. Later on in the chapter the study will focus on the main differences between deterministic, chaotic and complex systems.

Let us defyine a linear and non linear system. "*A linear system is one you can understand by understanding its parts individually and then putting them together. ... A nonlinear system is one in which the whole is different from the sum of the parts*."[5] A linear system is governed by a system of linear equation (equation of the type $y = mx + q$); wherease a non linear system is governed by a system of non linear equation (all of the other types: quadratics, logarithmics, trigonometrics, …). To understand better the implications of this difference and how these two type of functions behave it is useful to provide an example.

The example we present here is taken from Melanie Mitchell's "*Complexity: A Guided Tour*". Imagine a population of rabbits that doubles in size each year. Starting with 2 rabbits, the population grows to 4 in the second year, 8 in the third, and so on. Now, consider a variation: we place one rabbit on each of two separate islands and apply the same doubling rule. After one year, each island has 2 rabbits, totaling 4; in the second year, each has 4 rabbits, totaling 8; and this pattern continues. *Table 1* below illustrates the results. The first column lists the years. The second and third columns show the populations on the two separate islands (the split scenario), while the fourth column represents the original case with the entire population on a single island.

[5] Mitchell, *Complexity*, [page 22].

Year	n. of rabbit in island 1	n. of rabbit island 2	N of rabbit in case 1
1	1	1	2
2	2	2	4
3	4	4	8
4	8	8	16

Table 1

The table reveals that dividing the rabbit population at the start, as in the second scenario, does not affect the total number of rabbits by the end of each year. This outcome indicates that the system behaves linearly: we can grasp the overall dynamics by understanding the behavior of each part and then combining them. In other words, the whole is simply the sum of its parts. As illustrated in *Figure 1*, when we plot the population of one generation on the vertical axis and the next generation's population on the horizontal axis, the result is a straight line, an indication of the system's linear structure.

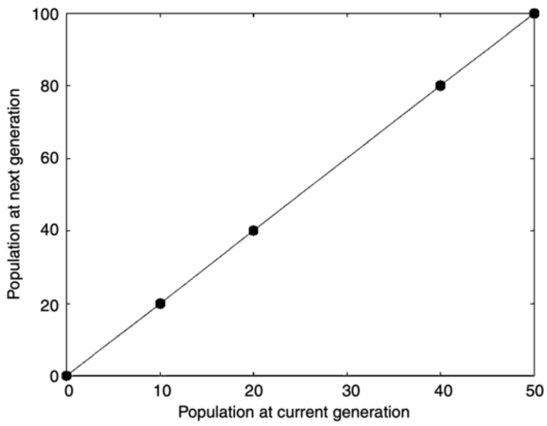

Figure 1[6]

However in the real world the growth is not infinite, so let us put a limit on this duplication of rabbits. To do this we introduce the birth rate, the death rate (the probability an individual will die) and the maximum carrying capacity (the upper limit of the population that the habitat will support).

Suppose we start with:

birth rate=2;

death rate=0.4;

[6] Mitchell, *Complexity*, [page 25].

carrying capacity, k =32.

In biology to compute the population growth from a period to another is used an equation called the logistic equation[7]. The logistic model formula is the following:

$$n_{t+1} = (\text{birthrate} - \text{deathrate}) \cdot \frac{k \cdot n_t - n_t^2}{k}$$

where n_{t+1} is the population at the period t+1. The key question now is what happens when we start with a single population of 20 rabbits compared to two separate populations of 10 rabbits each, one on each island.

In the first scenario, we begin with a single population of 20 rabbits, so $n_t = 20$. Applying the formula, we get:

$$n_{t+1} = (2 - 0.4) \cdot \frac{32 \cdot 20 - 20^2}{32} = 1.6 \cdot 7.5 = 12$$

This result shows that by the end of the first period, the population decreases to 12 rabbits.

In the second scenario, we start with two separate groups of 10 rabbits each, so $n_t = 10$ for both subsets. Applying the same formula to each subset:

$$n_{t+1} = (2 - 0.4) \cdot \frac{32 \cdot 10 - 10^2}{32} = 1.6 \cdot 6.875 = 11$$

Since there are two identical groups, the total population becomes 11+11=22.

This outcome is significantly different from the first case: 12≠22.

"*A nonlinear system is one for which inputs are not proportional to outputs: a small (large) change in some variable or family of variables will not necessarily result in a small (large) change in the system*." [8] *Figure 2* illustrates how this function behaves differently compared to the one presented earlier, highlighting the nonlinear nature of the system.

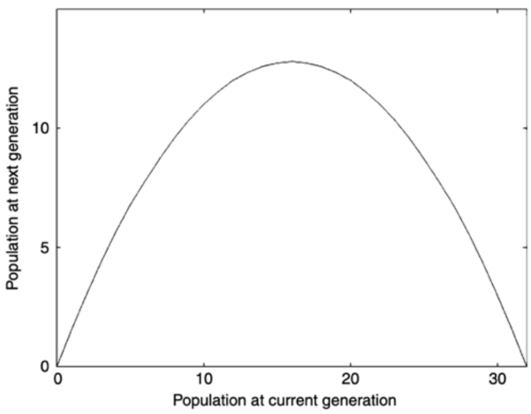

[7] The logistic equation is a model of population growth published by Pierre Verhulst (1845). In this thesis it is not important to understand what is the logistic model but what are the effect of functions of these type, the non linear functions.

[8] Rickles, Hawe, and Shiell, "*A Simple Guide to Chaos and Complexity*," 935.

Figure 2[9]

"*Why are nonlinear systems so much harder to analyze than linear ones? The essential difference is that linear systems can be broken down into parts. Then each part can be solved separately and finally recombined to get the answer. This idea allows a fantastic simplification of complex problems.*"[10]

Understanding the distinction between linear and nonlinear functions, and between linear and nonlinear systems, is essential for grasping why some systems behave predictably while others do not. This distinction serves as a foundation for exploring determinism in complex systems. We will revisit these implications when we examine chaotic systems later on, but first, let us explore the nature of deterministic systems in more detail.

---

9 Mitchell, *Complexity*, [page 26].

10 Steven H. Strogatz, *Nonlinear Dynamics and Chaos: With Applications to Physics, Biology, Chemistry, and Engineering*, 2nd ed. (Boca Raton, FL: CRC Press, 2019), [page 39].

1.2. Deterministic systems

1.2.1. Deterministic systems, history

"Nature and Natures laws lay hid in night: God said, let Newton be! and all was light."

-Alexander Pope, proposed Epitaph for Isaac' Newton, who died in 1727.

Is it possible to predict precisely the future state of a dynamical system? For more than one century the humankind (to be fair just the western society) thought that the answer to this question was yes. During the 18th and 19th centuries western scientists were sure that, given a specific knowledge, it would have been possible to know precisely the past and the future. Isaac Newton constructed the basis for a model in which the universe is seen as a gigantic clock and the nature is governed by immutable and absolute universal physical laws[11]. Nevertheless, it is with Pierre-Simon de Laplace that this vision reaches its most radical expression. In its famous determinism formulation, Laplace ipotize the excistence of an intellect able to know simoultaneously all the law of the nature and all the initial conditions of it. This intellect is capable of predicting, in theory, every past and future event.

" *We may regard the present state of the universe as the effect of its past and the cause of its future. An intellect which at a certain moment would know all forces that set nature in motion, and all positions of all items of which nature is composed, if this intellect were also vast enough to submit these data to analysis, it would embrace in a single formula the movements of the greatest bodies of the universe and those of the tiniest atom; for such an intellect nothing would be uncertain and the future just like the past could be present before its eyes."*[12]

How did we come to this point? How did humanity reach such a level of confidence, perhaps even arrogance, in its understanding of the world?

The theory of dynamical systems has deep historical roots, stretching back to ancient Greek thought. Aristotle was among the first to formulate a structured theory of motion, based in two key principles. First, he proposed that on Earth, an object naturally comes to rest unless a force acts upon it, if a force is applied, it moves in a straight line. In contrast, celestial bodies were believed to move endlessly in perfect circles centered on the Earth. Second, Aristotle asserted that the motion of earthly objects depends on their composition: a rock falls because it is made of the element earth, while smoke rises because it is composed of air. It is worth emphasizing

[11] Isaac Newton, *The Principia: Mathematical Principles of Natural Philosophy*, trans. I. Bernard Cohen and Anne Whitman (Berkeley: University of California Press, 1999).

[12] *Pierre-Simon Laplace, A Philosophical Essay on Probabilities, trans. Frederick Wilson Truscott and Frederick Lincoln Emory (New York: Dover Publications, 1951).*

Aristotle's method of reasoning and his approach to developing theoretical models of motion. He "*was not one to let experimental results get in the way of his theorizing. His scientific method was to let logic and common sense direct theory; the importance of testing the resulting theories by experiments is a more modern notion.* "[13] Aristotele and more in general all the ancient greeks thought was based on the idea that the universe had a rational principle called logos. The human being then could understand this principle through the nature and the philosophy. The Greeks in fact, adapted their theory with what was happening in their surroundings. Aristoteles, Plato and the ancient philosophers conceived the nature as a cosmos, a qualitative order in which the elements of the reality were interconnected following a universal equilibrium. This stability was not just a set of quantitative laws and it was the main way of thinking of western society until XVI century.

After the institution of the scientific revolution and the philosophical contribution of Descartes, this conception of the reality totally changed in the XVI and XVII centuries. The modern science progressively substituted the qualitative vision of the Greeks. The new approach was based on the reduction of nature to mathematical laws .The universe started to be seen as a perfect machine in which universal and deterministic laws were able to predict every effect of an event. This method was based on the principle of cause effect. The classic science affirmed that the nature has no history, it does not evolve: the laws are eternal, immutable and valid everywhere and everytime. [14]

In this perspective, the reality is reducible to a set of mathematical and physical laws that determine the unique evolution of every natural system. The consequence of this statement is the negation of creativity, uncertainty and unpredictability. The universe is conceived as a closed and reversable system where time is only a simple mathematical parameter. This mathematic paradigm influenced the development of classic physic and engineer science, and it produced science model based on linearity. The reality was seen as the sum of the single components and its understanding derived from the analysis of those single components isolated. This model denied and found useless a science whose goal was the investigation of the interaction of the agents.

[13] Mitchell, *Complexity*, [page 17].

[14] Ilya Prigogine and Isabelle Stengers, *Order out of Chaos: Man's New Dialogue with Nature* (New York: Bantam Books, 1984).

1.2.2. Deterministic systems, definition

Let us now examine the mathematical definition and formalization of the determinism paradigm. " *A process is called deterministic if its entire future course and its entire past are uniquely determined by its state at the present time. The set of all states of the process is called the phase space."*[15] This property is called unique evolution and from it we can derive that "*given a state at a specific time there is only one history of transitions consistent with the laws*."[16] In other words, no two different outcomes can arise from identical starting conditions. Deterministic dynamics are typically expressed as equations in state-space form. The system's state is represented by a vector of variables x(t) in an appropriate state space (or phase space), and the time-evolution rule is given by a fixed function or law.

For a continuous-time system, the evolution is often specified by a differential equation:

$\dot{x}(t) = f\big(x(t)\big)$

Where,

$x(t) \in R^n$, is the state vector, all the variables that describe the systems at time t;

$\dot{x}(t)$ is the time derivative of the state;

$f: R^n \to R^n \; is\ a\ deterministic\ function\ (the\ system's\ rule\ or\ law)$;

The equation tells how the system evolves from the initial condition $x(0) = x_0$.Ordinary differential equations can be described within a general theoretical framework using the following system:

$\dot{x_1} = f_1(x_1, \dots, x_n)$

$\vdots$

$\dot{x_n} = f_n(x_1, \dots, x_n)$ [17]

In discrete time, the evolution is given by an iteration:

$x_{n+1} = f(x_n)$

$x_0,\ f(x_0),\ f\big(f(x_0)\big),\ f\Big(f\big(f(x_0)\big)\Big),$

Where n indexes time steps. This recursive map produces a sequence of states x_0, x_1, x_2, determined entirely by the initial value x_0.[18]

[15] Vladimir I. Arnold, *Ordinary Differential Equations*, trans. Roger Cooke (New York: Springer, 1992), [page 13].
[16] https://plato.stanford.edu/entries/chaos/#:~:text=A%20mathematical%20model%20is%20deterministic,if%20it%20exhibits%20unique%20evolution
[17] Strogatz, *Nonlinear Dynamics and Chaos*.
[18] Devaney, *An Introduction to Chaotic Dynamical Systems*.

A contiuous time-system is a system in which the state $x(t)$ changes continuously over time. It is possible to compute the state at any real-valued time $t \in R^n$. In a discrete time system the state x_n changes at discrete time steps (e.g., every second, every iteration). In both cases (continuous or discrete), no external randomness enters: the "rule" f alone drives the dynamics. Determinism thus implies a well-defined trajectory (path in phase space) passing through one state to the next.

It is important to note that deterministic does not necessarily mean simple or predictable in practice. It only guarantees that the outcome is uniquely determined by initial conditions, not that we can easily calculate it.

1.3. Chaotic systems

1.3.1. Chaotic systems, definition

"*Deterministic chaos is the rule, not the exception.*"

- Edward Lorenz

Deterministic systems can nevertheless behave in ways that appear irregular, erratic, or unpredictable over the long term. Chaos refers to such complicated behavior arising within a deterministic framework. In a chaotic dynamical system, the underlying evolution rules are fixed and non-random, yet the system's trajectory shows an aperiodic and seemingly random pattern over time. Crucially, chaotic behavior does not violate determinism: the system still follows one unique trajectory from any given initial state. Chaotic behavior is always deterministic, despite popular confusion of chaos with true randomness. The difference is that chaotic systems are extremely sensitive and complex, making their long-term behavior effectively unpredictable even though it is in principle determined by initial conditions.[19]

Currently, there is no single, universally accepted definition of chaos. However, the scientific community generally agrees on a three-point description, outlined below.

"*Chaos is aperiodic long-term behaviour in a deterministic system that exhibits sensitive dependence on initial conditions*."[20]

In which:

1) Aperiodic long-term behaviour refers to the fact that the system's trajectory does not converge to fixed points or repeat in regular cycles, even as time approaches infinity.
2) Determinism implies that the system operates without any random external inputs; the observed irregularity in behavior stems entirely from its inherent nonlinear structure.
3) Sensitive dependence on initial condition indicates that trajectories starting from nearly identical states diverge rapidly over time, with small initial differences growing exponentially.[21]

Before presenting the mathematical formulation, it is essential to first introduce the concepts of 1) topological transitivity and 2) sensitivity to initial conditions.

1) "$Definition$: $f\colon J \to J$ *is said to be* $topologically\ transitive$ *if, for any pair of open sets* $U, V \subset J$, *there exists* $k > 0$ *such that* $f^k(U) \cap V \neq \emptyset$."[22]

[19] plato.stanford.edu
[20] Strogatz, *Nonlinear Dynamics and Chaos*, [page 798].
[21] Strogatz, *Nonlinear Dynamics and Chaos*.
[22] Devaney, *An Introduction to Chaotic Dynamical Systems*, [page 61].

$f^k(U)$ denotes the k-th iterate of the function f applied to the set U. Specifically:

$$f^k(U) = \underbrace{f(f(\dots f(x)\dots))}_{k \text{ times}}.$$

Basically it means that points from any regions can evventualy reach any other region. In this way every region of the space interact with all of the other region as the time continues.

2) "$Definition$: $f{:}\,J \to J$ has $sensitive\ dependence\ on\ initial\ conditions$ on J if there exists $\delta > 0$ such that, for any $x \in J$ and any neighborhood N of x, there exists $y \in N$ and $n \geq 0$ such that $|f^n(x) - f^n(y)| > \delta$.

Intuitively, a map possesses sensitive dependence on initial conditions if, for each point x in J, there exist points arbitrarily close to x whose orbits eventually separate from the orbit of x by at least δ under iteration of f."[23]

Now it is possible to formulate a mathematical definition of a chaotic system:

$Definition$: Let V be a set. $f: V \to V$ is said to be chaotic on V if f has the following three properties:

1. $periodic\ points\ are\ dense\ in\ V$;
2. $f\ is\ topologically\ transitive$;
3. $f\ has\ sensitive\ dependence\ on\ initial\ conditions$.

The first propertie states that it is possible to find periodic point in every neighborhood. A periodic point is one that returns to its original position after a finite number of iterations of the function. A chaotic system posses 3 important characteristics: unpredictability, indecomposability, and an element of regularity.[24]

1.3.2. Chaotic systems, history

The journey from deterministic predictability to chaotic unpredictability has its roots in the foundational problem of Newtonian dynamics: the n-body problem. Newton was able to solve the 2 bodies problem but, when a third body was added, the conclusions seemed impossible. What does it mean? Newton's laws of motion, when combined with his universal law of gravitation, could precisely describe the motion of two celestial bodies, such as the Earth and the Moon. However, as soon as a third body was introduced, the equations became intractably complex. In 1887, the Swedish King Oscar II announced a mathematical competition to determine the long-term stability of the solar system. The king promised to award a prize to the one who succeeded in solving the problem. Henri Poincaré entered the contest by

---

[23] Devaney, *An Introduction to Chaotic Dynamical Systems*, [page 61].

[24] Devaney, *An Introduction to Chaotic Dynamical Systems*.

attempting to solve the three-body problem using Newton's laws, but he failed. Or better he found that the question was impossible to solve due to approximation problems. He infact stated that:

"*A very small cause which escapes our notice determines a considerable effect that we cannot fail to see, and then we say that that effect is due to chance. If we knew exactly the laws of nature and the situation of the universe at the initial moment, we could predict exactly the situation of that same universe at a succeeding moment. But, even if it were the case that the natural laws had no longer any secret for us, we could still only know the initial situation approximately. If that enabled us to predict the succeeding situation with the same approximation, that is all we require, and we should say that the phenomenon had been predicted, that it is governed by laws. But it is not always so; it may happen that small differences in the initial conditions produce very great ones in the final phenomena A small error in the former will produce an enormous error in the latter. Prediction becomes impossible, and we have the fortuitous phenomenon."*[25]

It is worth highlighting that Poincaré, although he did not formally prove it, recognized the profound effect that even minimal measurement errors can have within nonlinear systems. This insight marked the earliest formal acknowledgment of what we now refer to as chaos, though the term itself was not yet in use. Poincaré's observations were remarkably ahead of his era, and as a result, their significance was not immediately appreciated by the broader mathematical community. For much of the early twentieth century, the investigation of chaotic dynamics remained largely neglected, as attention shifted toward other emerging fields such as topology and quantum mechanics.

In the early 1960s, Edward Lorenz, a meteorologist at MIT, constructed a simplified computational model of the atmosphere. The machine printed numbers that simulated wind and temperature patterns. Although limited in memory and processing power, this toy model had the ability to mimic actual weather behavior. It did not reproduced it exactly, but showed aperiodicity (non-repeating patterns) and unpredictability. The system never returned to the same state, even when inputs seemed almost identical.

Lorenz discovered chaos by accident. During a test, he rounded off input numbers to three decimal places instead of six, expecting identical results. But the simulation quickly diverged into a radically different trajectory. This minor numerical change amplified into a large-scale

[25] Henri Poincaré, *Science and Method*, trans. Francis Maitland (London: Thomas Nelson and Sons, 1914), [page 67].

alteration. It was an early demonstration of what later became known as sensitive dependence on initial conditions.

1.3.3. Chaotic systems, demonstration

To reproduce this effect, we used the classic Lorenz system, a set of three coupled nonlinear differential equations derived from fluid dynamics[26]:

$$\begin{cases} \dot{x} = \sigma(y - x) \\ \dot{y} = x(r - z) - y \\ \dot{z} = xy - bz \end{cases}$$

where the standard parameter values are σ=10, r=28, and b=8/3.

We performed a numerical simulation using Euler's method with two sets of initial conditions:

Case A (3-decimal precision): x =1.235, y =1.346, z =1.457

Case B (6-decimal precision): x =1.234567, y =1.345678, z =1.456789

Euler's method is a simple numerical technique used to approximate solutions to ordinary differential equations (ODEs) by iteratively advancing the solution using the slope at each step:

$x_{n+1} = x_n + h \cdot f(x_n, t_n)$,

where h is the step size and in this case it is h =0.01.

At every iteration step, all calculations were constrained to either three or six decimal digits, just as Lorenz had experienced due to the limited precision of 1960s computing.

The following *Table 2* displays the results at selected iteration steps, along with the Euclidean distance between the states of the two systems:

Step	x_3dec	y_3dec	z_3dec	x_6dec	y_6dec	z_6dec	Distance
1	1,246	1,66	1,435	1,245678	1,659915	1,434555	**0,000556**
2	1,287	1,974	1,417	1,287102	1,974236	1,416977	**0,000258**
5	1,568	2,993	1,4	1,568334	2,993617	1,39919	**0,001072**
20	7,288	15,188	5,103	7,289412	15,19025	5,10427	**0,002941**
50	-4,021	-15,286	33,587	-4,0237	-15,2863	33,58725	**0,002726**
100	-5,303	-7,64	18,389	-5,3009	-7,63854	18,38485	**0,004877**
200	-13,508	-21,229	23,741	-13,4603	-21,209	23,59644	**0,153534**
500	-1,12	0,579	22,632	-0,9949	0,759842	22,6675	**0,222741**

Table 2

Even though both systems use the same mathematical laws and extremely similar initial values, their paths diverge measurably over time. This is not due to randomness, but to the

[26] Once again here it is not important to understand where these equations came from, but what are the implications

amplification of numerical noise inherent to chaotic systems. The divergence begins gradually but persists and grows as the simulation progresses. *Figure 3* provides a visual confirmation of this exponential trend, illustrating the rapid divergence between the two trajectories over time.

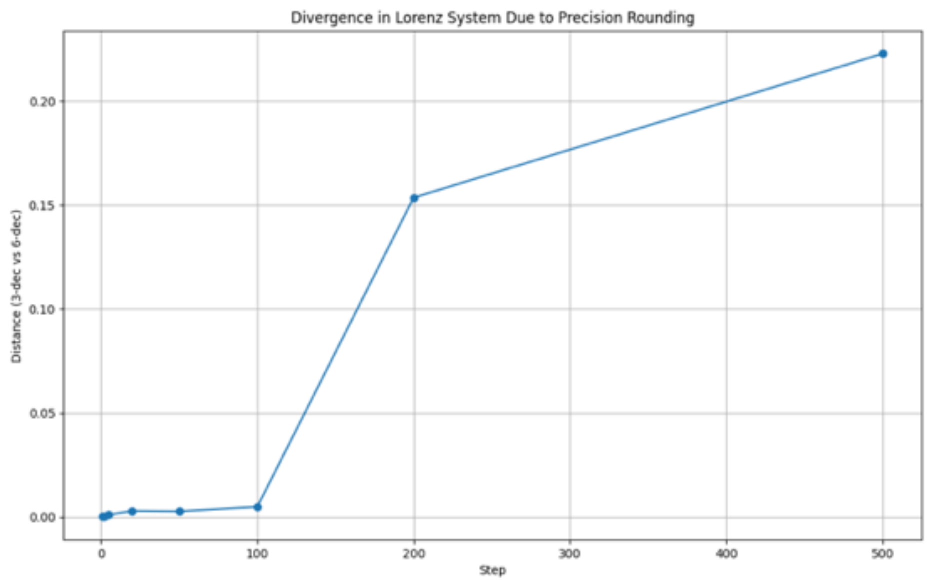

Figure 3

This experiment underscores a fundamental limitation of deterministic modeling in chaotic regimes: perfect predictability is impossible in practice, not because of flaws in the equations, but because no measurement or computation can ever be infinitely precise.

1.3.4. Chaotic systems, implications

The deterministic paradigm assumed that science could build models to fully describe and predict natural phenomena. Chaos theory challenges this idea. It shows that there are structural limits to what science can predict. In non-linear systems, the issue is not a lack of knowledge or measurement errors. The problem lies in the system's inherent instability, which amplifies minimal perturbations beyond any reasonable predictive capacity.[27]

"*In short, the presence of chaos in a system implies that perfect prediction à la Laplace is impossible not only in practice but also in principle, since we can never know x_0 to infinitely many decimal places*."[28]

[27] James Gleick, *Chaos: Making a New Science* (New York: Open Road Integrated Media, 2011).
[28] Mitchell, *Complexity*, [page 33].

The discovery and development of chaos theory have fundamentally transformed science and reshaped humanity's understanding of nature. In the 19th century, such ideas would have seemed implausible, even absurd, to many scientists of the time. Yet, modern research has revealed several groundbreaking insights:

- Apparently random behavior can arise from purely deterministic systems, without any external source of randomness.
- Certain simple, deterministic systems can exhibit behavior that, due to their sensitivity to initial conditions, becomes inherently unpredictable over the long term.
- Despite this unpredictability, chaotic systems often display an underlying structure, a kind of "order within chaos",reflected in universal properties shared across diverse chaotic systems.[29]

This new lens on unpredictability opened the door to a broader scientific revolution: the emergence of complexity theory. If chaos revealed the limits of predictability in simple deterministic systems, complexity expanded that insight to systems composed of many interacting parts, systems where unpredictable, emergent behavior arises not just from sensitivity to initial conditions, but from the structure of interactions themselves. It marked a shift from asking how systems behave to understanding how systems organize, adapt, and evolve.

[29] Mitchell, *Complexity*.

1.4. Complex systems

1.4.1. Complex systems, definition

"*More is different*."

- Philip W. Anderson, nobel laureate in physics 1977.

The studies of complex systems started in the lasts decades of the previous centuries. As for chaos theory, we do not have a globally accepted definition of complex system. Despite its growing influence across disciplines, it is often described as a broad field that concern how independent agents interact and evolve over time in response to each other and the environment. An informal definition can be that: complexity science is "*a field of research that explores how independent agents interact with each other in a variety of ways.*"[30] The most important thing to analyze is that the agents produce collective dynamics that cannot be traced back solely to the characteristics of individual agents. This fact reflects the field's intellectual necessity for interdisciplinarity. Complexity theory draws upon contributions from biology, computer science, systems theory, economics, cognitive science, and social theory, among others. Each discipline brings different models and methods, but all converge on the recognition that real-world systems are not best understood by breaking them down into isolated parts.

This departure from classical reductionism marks a decisive shift in scientific thinking. Classical models, especially the ones derived from Newtonian physics, assumed that if all parts of a system were understood, then it was possible to understand also the collective behaviour of the system. Complexity theory directly challenges this assumption. What it emphasizes is that the whole may be greater than, or even qualitatively different from, the sum of its parts.[31] It focuses in particular on dynamic interactions, non-linear feedback loops, and the emergence of novel structures over time. Instead of offering universal law, complexity offers a language of patterns and change. It helps us understanding how societies transform, how innovation unfolds, and how living systems sustain themselves amid uncertainty. In this light, complexity theory is not merely a scientific theory but a new paradigm for understanding systemic transformation, one that privileges relational thinking, contextuality, and the co-evolution of structure and behavior.

Complexity theory and chaos theory might seem very similar since they share rejection of linear, deterministic frameworks, yet they are conceptually distinct in both scope and

---

[30] John R. Turner and Rose M. Baker, "*Complexity Theory: An Overview with Potential Applications for the Social Sciences,*" *Systems* 7, no. 1 (2019): 1–23.

[31] Turner and Baker, "*Complexity Theory*," 5.

application. As we have seen chaos theory states the unpredictability of a system deriving from sensitive dependance on initial conditions. Such systems are deterministic yet appear random in the long term. Complexity theory, on the other hand, investigates systems composed of multiple interacting components whose collective behavior gives rise to new, emergent pattern. Complex theory concern not only unpredictability but also self-organization, adaptation and evolution over time. Complex systems are often high-dimensional and exhibit behaviors not deducible from their components alone.[32] One of the most important difference between chaotic and complex system is that chaotic systems are typically closed and memoryless; they evolve in ways that are highly sensitive to their initial states but do not learn or adapt. In contrast, complex adaptive systems (CAS), a central concept in complexity theory, are open systems that incorporate feedback, adjust their behavior based on experience, and modify their internal rules in response to environmental changes.[33] Despite these differences, both frameworks share crucial methodological affinities. They emphasize non-linearity and challenges the classical assumption about causality.

1.4.2. Complex systems, CAS (complex adaptive systems)

A central notion in complexity theory is the notion of complex adaptive system (CAS): systems composed of multiple agents that interact with each other and their environment according to local rules, and that are capable of adaptation and learning over time. Differently from static and merely complex system, CAS are defined by their capacity to evolve, aggregate behaviour and anticipation[34]. Each agent within the system follows a limited set of rules, but their collective interactions generate emergent global behaviors that cannot be predicted by analyzing single agents. According to Holland, what distinguishes CAS from other forms of complex systems is their ability to learn and self-organize in ways that increase their fitness in changing environments. These systems evolve by balancing exploitation of known strategies and exploration of new possibilities. In this context, “credit assignment” becomes essential: effective rules or patterns are reinforced while ineffective ones are discarded or recombined. Such mechanisms parallel those found in evolutionary biology and neural learning models.[35] Moreover, CAS are open systems that continually exchange energy, information, or resources with their environment. They do not seek equilibrium, and their structure is often the result of

[32] Rickles, Hawe, and Shiell, “*A Simple Guide to Chaos and Complexity*,” 935.
[33] Turner and Baker, “*Complexity Theory*,” 5.
[34] John H. Holland, “*Complex Adaptive Systems*,” *Daedalus* 121, no. 1 (1992): 17–30.
[35] Holland, “*Complex Adaptive Systems*,” 25.

ongoing, decentralized interactions rather than centralized control. This makes them highly responsive to contextual changes but also difficult to model using classical equations.[36]

Complex Adaptive Systems are not abstract constructs limited to theoretical models, they are present in the natural and social world. One of the most illustrative examples is the immune system. It consists of a decentralized network of cells that continuously adapt to new pathogens. Each immune cell follows simple recognition rules but their interactions result in complex behaviors such as memory, response modulation, and self/non-self differentiation.[37] Similarly, economic markets function as prototypical CAS: they comprise countless agents (consumers, firms, regulators) whose decisions, collectively shape market trends. These systems do not converge toward a single equilibrium but evolve in response to policy shifts, technological changes, and cultural trends. Another example is ecosystems, where species interact in food webs, compete for resources, and evolve. As Holland notes, ecosystems exhibit adaptive behavior over time as populations fluctuate, niches evolve, and environmental pressures reshape the dynamics of interdependence. CAS can also be found in social organizations, such as governments or universities. They operate without a singular controlling entity yet adapt to political, economic, and social demands through distributed decision-making.[38] These organizations frequently reorganize their internal structures in response to environmental changes, demonstrating the core traits of emergence, self-organization, and adaptability. In the digital age, online social networks and digital ecosystems (such as app stores or e-commerce platforms) also are CAS dynamics. Agents engage in real-time feedback cycles that shape content visibility and platform evolution.

These examples highlight the practical relevance of CAS across domains. Their defining features allow them to navigate environments marked by uncertainty, interdependence, and constant change.

1.4.3. Complex systems and information theory

To fully grasp the inner workings of complex adaptive systems, one must consider not only their dynamic structure and emergent behavior, but also how they process, store, and transmit information. Information theory offers a powerful framework for analyzing complexity at a fundamental level, particularly through its capacity to quantify the informational content required to describe or reproduce a system. One of the most profound bridges between complex

[36] Turner and Baker, "*Complexity Theory*," 5.
[37] Holland, "*Complex Adaptive Systems*," 25.
[38] Turner and Baker, "*Complexity Theory*," 5.

systems and information theory emerged in the 1960s, when mathematician Andrey Kolmogorov, along with Gregory Chaitin and (independently) Ray Solomonoff, proposed defining an object's complexity as the length of the shortest computer program that can generate a complete description of that object.[39] This quantity, now known as Kolmogorov complexity or algorithmic information content, measures an object's irreducible information: the more easily a pattern can be compressed into a short description, the "simpler" it is, whereas a pattern that is truly random (lacking any regularity) has no description shorter than a literal transcription of itself and thus possesses maximal complexity[40]. For example, a DNA sequence consisting of a simple repeating motif can be generated by a brief algorithm ("print AC ten times"), yielding low complexity, whereas an equal-length sequence with no apparent pattern would require an algorithm essentially as long as the sequence itself, making it incompressible and highly complex.[41] Kolmogorov first articulated this idea in 1965 (unaware that Solomonoff had already hinted at a similar principle in 1960), and around the same time Chaitin arrived at the same definition from a computational perspective.[42] Their contributions inaugurated the field of algorithmic information theory and inspired further refinement of the concept; over the next decade the initial formulations were improved and their soundness confirmed in practice. Kolmogorov complexity thus provides a universal, machine-independent measure of complexity in abstract systems, marrying the intuitions of complexity science with a rigorous information-theoretic framework. In mathematics, Chaitin used this notion to reveal fundamental limits to knowledge: essentially, no formal axiomatic system of limited size can prove that a given sequence is random if that sequence's Kolmogorov complexity exceeds the information contained in the system's axioms, a result that casts Gödel's incompleteness phenomenon as a natural consequence of information constraints. In the realm of scientific inquiry, Solomonoff's related work showed how algorithmic complexity can guide inductive reasoning: the best theory is the one with the shortest description that reproduces the observed data, an algorithmic form of Occam's razor in which the simplest adequate program (minimum Kolmogorov complexity) is preferred as the explanation of phenomena.[43] Even in the social sciences, this view of complexity has resonated: it implies that truly complex phenomena (whether in biology, society, or economics) carry a high information content and are inherently

---

[39] Mitchell, *Complexity*.
[40] Gregory J. Chaitin, "*Randomness and Mathematical Proof*," Scientific American 232, no. 5 (May 1975): 47–52.
[41] Mitchell, *Complexity*.
[42] Chaitin, "*Randomness and Mathematical Proof*," 49.
[43] Chaitin, "*Randomness and Mathematical Proof*," 49.

unpredictable. Indeed, organizations understood as complex systems have been described as "radically unpredictable," defying any straightforward, compact description or top-down control.[44] In sum, Kolmogorov and Chaitin's work provided a new theoretical lens for complexity, one that equates complexity with information, and in doing so it has enriched our understanding of randomness, structure, and the limits of compressibility in mathematical, scientific, and even social domains.

[44] W. Brian Arthur, "*Complexity and the Economy*" *Science* 284, no. 5411 (1999): 107–109.

1.5 Introduction to adjacent possible theory

As we have seen, Complex Adaptive Systems (CAS) are defined not only by the richness of their internal interactions, but also by their ability to respond to changing environments through continual adaptation. What makes them especially resilient is their capacity to oscillate between two essential dynamics: the efficient use of what is already known (exploitation) and the cautious pursuit of what remains unexplored (exploration). This balance is not a marginal feature, it is fundamental to how these systems evolve, learn, and survive in conditions of uncertainty.[45] It is precisely this dynamic interplay that opens the door to a broader conceptual framework: the theory of the Adjacent Possible.

Stuart Kauffman's notion of the Adjacent Possible offers a convincing way to interpret how complex systems generate novelty. Rather than moving toward chaotic or implausible futures, adaptive systems tend to move incrementally, testing the limits of their current configuration while stretching into new but structurally reachable directions.[46] In other words, at any given moment, a system is surrounded by a range of "next steps" that were not accessible before but become viable as the system evolves. The Adjacent Possible, then, is not just a metaphor for innovation; it is a structural property of complexity itself. As Kauffman explains, the growth of possibility unfolds in tandem with the history of what has already emerged.

This shift in perspective, from solving problems within a known landscape to generating new landscapes entirely, marks a significant extension of complexity theory. It suggests that adaptive systems are not merely reactive or rule following; they are inherently inventive, continually reshaping the profile of their future through the choices they make in the present. The Adjacent Possible thus serves as a conceptual bridge between the descriptive power of complexity science and the creative logic of innovation and transformation.

---

[45] Turner and Baker, "*Complexity Theory*," 5.
[46] Stuart A. Kauffman, *The Origins of Order: Self-Organization and Selection in Evolution* (New York: Oxford University Press, 1993).

Chapter 2: Adjacent possible theory

2.1. Link with Complex adaptive systems

Complex Adaptive Systems (CAS) occupy a central place within the broader field of complexity science due to their unique capacity to generate emergent behaviors through the dynamic interplay of heterogeneous agents, feedback loops and nonlinear interactions. Unlike purely deterministic or chaotic systems, which evolve within fixed and predefined state spaces, CAS are characterized by their ability to evolve with their environment, continuously altering the landscape of possibilities through endogenous novelty production.[47] This ability to transform not only in response to external conditions but also through self-organization and the internal recombination of existing structures leads to a distinctive mode of evolution that cannot be fully captured by traditional models of linear causality or optimization.

It is within this conceptual framework that the Theory of the Adjacent Possible (TAP) becomes particularly salient. Introduced by Stuart Kauffman in the context of evolutionary biology, TAP describes the dynamically expanding space of potential states that become accessible to a system as a direct consequence of the states it has already actualized.[48] At any given moment, a CAS is surrounded not by an infinite array of options but by a set of adjacent possibilities: configurations that were previously inaccessible but become viable due to recent structural transformations within the system.[49] This conception of adjacent possibilities as emerging at the edge of the current state space reframes our understanding of adaptation, it emphasize a process not merely of selection among fixed alternatives but of endogenous expansion into novel, structurally coherent futures.[50]

The relevance of TAP to CAS lies precisely in this capacity to model evolution as an open ended, path dependent process. Each act of exploration or recombination within a CAS does not merely yield a new element, it also alters the configuration of the possibility space itself. As Devereaux, Koppl and Kauffman argue, this leads to a non ergodic movement through an unlistable and generative landscape, where neither the full set of future options nor their associated outcomes can be known ex ante.[51] In this context, the adjacent possible functions as

---

47 Kauffman, *"The Origins of Order."*

48 Andrew Devereaux, Roger Koppl and Stuart Kauffman, *"Creative Evolution in Economics."* Journal of Evolutionary Economics 34 (2024): 489–514.

49 Vittorio Loreto et al., *"Dynamics on Expanding Spaces: Modeling the Emergence of Novelties."* Nature Physics 12, no. 10 (2016): 842–847.

50 Joakim Taalbi, *"Long-Run Patterns in the Discovery of the Adjacent Possible," arXiv* (2023).

51 Devereaux, Koppl and Kauffman, *"Creative Evolution in Economics."*

a generative grammar for novelty, enabling us to understand how systems can evolve in ways that are both constrained by historical precedent and yet radically inventive. The system's trajectory, far from being teleologically determined or statistically predictable, becomes a function of emergent path dependencies, where each realized novelty conditions the set of subsequent possible innovations.[52]

Moreover, empirical and computational modeling work, particularly in social and technological systems, has demonstrated that TAP dynamics are not merely metaphorical. Studies using urn models and combinatorial algorithms have formalized how the appearance of novelty, represented as the addition of new balls or categories, triggers the subsequent emergence of further possibilities in a reinforcing cascade.[53] These dynamics have been shown to explain observed innovation patterns across diverse contexts, such as patent networks, cultural production and social media interactions, where waves of novelty are continuously mixed with the exploitation of past successes.[54] Theoretically, such models support the view that innovation arises not from isolated breakthroughs but from a combinatorial exploration of proximate possibility spaces, aligning with the core logic of CAS as systems that evolve by investigating the edges of their own structure.

In light of the conceptual alignment between Complex Adaptive Systems and the Theory of the Adjacent Possible, this chapter aims to provide a structured and analytical exploration of the Adjacent Possible as a framework for understanding innovation and systemic evolution. It begins by defining the notion of the adjacent possible space, clarifying its theoretical foundations and distinguishing it from adjacent but unrelated concepts in complexity science. The discussion then turns to the temporal evolution of the adjacent possible, structured into four key phases: initial exploration, the emergence of structured complexity, the onset of accelerated combinatorial growth, and the appearance of new systemic constraints that shape future innovation trajectories. The chapter will also delineate the intrinsic characteristics of the adjacent possible space, including its dynamic expansion, path dependence and irreversibility. In addition, it will classify different types of adjacent possible spaces, such as accessible and cognitive adjacent, and it will analyze the implications of each for adaptive processes. Finally, the chapter concludes by examining empirical evidence and statistical regularities observed in

---

[52] Kauffman, "*The Origins of Order"*; Loreto et al., *"Dynamics on Expanding Spaces."*
[53] Loreto et al., "Dynamics on Expanding Spaces."; Francesca Tria et al., "*The Dynamics of Correlated Novelties."* Scientific Reports 4, no. 5890 (2014).
[54] Bernardo Monechi et al., *"Waves of Novelties in the Expansion into the Adjacent Possible." PLOS ONE* 12, no. 6 (2017).

real world systems, such as the emergence of power-law distributions, which support the theoretical predictions of adjacent possible dynamics in domains ranging from technological innovation to social networks and linguistic evolution.

2.2 APT: Definition and concepts

The concept of the Adjacent Possible (APT) emerged in the context of evolutionary biology and complexity theory, it was originally formulated by Stuart A. Kauffman to describe the generative mechanisms underlying biological innovation and the expansion of the biosphere. Dissatisfied with purely Darwinian explanations based solely in selection and variation, Kauffman introduced the idea that the biosphere does not evolve merely through adaptive walks on fitness landscapes, but also by expanding the very configuration space within which these adaptive moves occur.[55] In "*The Origins of Order"*, Kauffman posited that each novel biological entity (whether a molecule, metabolic function, or species) does not just occupy a previously empty niche, it actively reshapes the landscape of what is possible next by creating new affordances for further combinations, adaptations and co evolution.[56] This conceptual shift implied that evolution is not only an exploration within a fixed space of possibilities but also a process that dynamically constructs and enlarges that very space. This perspective resonates across diverse fields including innovation studies, information theory, economics and cultural evolution.[57]

At its core, the Theory of the Adjacent Possible posits that, at any given moment, a system, biological, technological or social, is surrounded by a set of potential states or configurations that are not currently actualized but become reachable due to what already exists.[58] This set of states, the "adjacent possible space", is not infinite but contextually constrained by the components, knowledge and interactions presently available within the system. The adjacent possible thus comprises all the novel configurations that can be generated by recombining, modifying or extending current elements through minimal steps. For example, in molecular evolution, once a new protein folds into a viable structure, it enables new biochemical pathways that were previously inaccessible. In cultural innovation, the invention of the printing press enables the creation of public newspapers, which in turn leads to the emergence of mass media systems. As such, the adjacent possible represents the structured perimeter of innovation: what can be done next, but not what can be done eventually.[59] This concept stands in contrast to the classical optimization frameworks where exploration is often viewed as a search across a static

---

[55] Kauffman, *The Origins of Order*
[56] Cortês et al., "The TAP Equation."
[57] Devereaux, Koppl and Kauffman, "Creative Evolution in Economics."
[58] Lars Björneborn, "Adjacent Possible," in *Springer MRW* (2023).
[59] Taalbi, "Long-Run Patterns."

landscape. Instead, the adjacent possible evolves with each innovation, continuously redrawing the boundaries of feasible change.

2.2.1. Exploration and exploitation

One of the foundational dynamics underlying the theory of the adjacent possible is the interplay between exploitation and exploration, two competing, yet interdependent, modes of systemic behavior that shape the trajectory of innovation and evolution in Complex Adaptive Systems (CAS). Exploitation can be understood as the incremental refinement and repeated use of existing structures, practices or knowledge. It leverages what the system has already actualized, deepening capabilities and improving efficiency by reinforcing known pathways.

As we can see in panel (a) of *Figure 4*, exploitation corresponds to the red arrows operating entirely within the space of already realized nodes (represented in grey), where internal linkages reinforce the current configuration of the system without extending its structural boundaries. In contrast, exploration involves the system's outward movement toward novel configurations that were not previously present but become accessible due to existing structures. Panel (b) of *Figure 4* illustrates this process clearly: the red arrows now extend from the actualized network into a green frontier of new nodes, representing the adjacent possible. These nodes were previously unreachable but now become structurally available due to the specific position and evolution of the current system state.

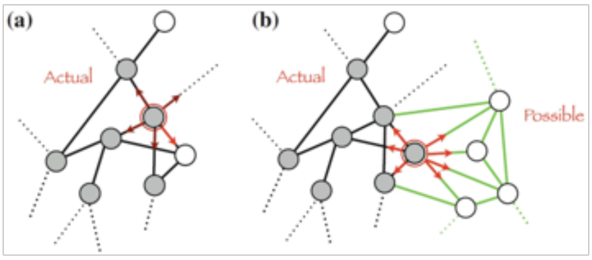

Figure 4, Lars Björneborn, "Adjacent Possible," in *Springer MRW* (2023).

This dichotomy between exploration and exploitation is crucial for adaptive performance: a system that overcommits to exploitation risks entrenchment in suboptimal equilibria, while one that favors unrestricted exploration may suffer from volatility, incoherence or premature abandonment of viable innovations.[60]

[60] Loreto et al., *"Dynamics on Expanding Spaces."*

Within the framework of the adjacent possible, this tension takes on a particularly generative role. The adjacent possible functions as a dynamic, structurally delimited perimeter that expands in tandem with the system's own evolution. It contains all those novel configurations that are reachable through minimal changes, combinations, interactions or recompositions, based on what already exists.[61] What distinguishes the adjacent possible from the merely hypothetical or imaginable is its dependency on actualized history: it is not a universal space of options but a situated, system contingent set of next step innovations. With each realization of novelty the adjacent possible enlarges recursively reshaping itself and opening pathways to further innovations.[62] In this sense, the adjacent possible does not merely describe potential futures, it describes how the very structure of potentiality is transformed by present action. In empirical contexts, such as the evolution of technological platforms or cultural trends, this manifests in the emergence of stable yet open ended innovation paths, where agents continuously negotiate between optimizing known solutions and reaching beyond them into uncertain terrain.[63]

While theoretical models such as the modified Polya's urn have been developed to capture these dynamics formally, their detailed mechanics will be addressed in the following sections. What is crucial to emphasize here, however, is that the adjacent possible is not merely unpredictable, it is, in a fundamental sense, unknowable in advance. As Kauffman famously argued, the growth of the biosphere exemplifies a process in which "*not only do we not know what will happen, we do not even know what can happen.*"[64] This radical unknowability stems from the generative nature of each novel realization, which redefines the structure of the system and produces new affordances that were previously inaccessible and inexpressible.[65] The adjacent possible is therefore not a fixed frontier of well specified alternatives but an emergent and combinatorially expanding field of potentialities. Its boundaries are contingent, recursive and unprestatable. As such, exploration into the adjacent possible implies confronting a shifting horizon: one cannot enumerate the full space of possibilities ex ante because the act of innovation continuously reshapes the possibility space itself. This insight marks a profound epistemological boundary for models of adaptive systems: we are not merely ignorant of future

[61] Stuart A. Kauffman, "*Investigations*" (Oxford: Oxford University Press, 2000); Björneborn, *"Adjacent Possible.*
[62] Devereaux, Koppl and Kauffman, *"Creative Evolution in Economics."*
[63] Elisa Ubaldi et al., *"Emergence and Evolution of Social Networks through Exploration of the Adjacent Possible."* Communications Physics 4, no. 28 (2021).
[64] Kauffman, "*Investigations*".
[65] Devereaux, Koppl and Kauffman, *"Creative Evolution in Economics."*

selections, we are fundamentally incapable of foreseeing the very space in which those selections will unfold.

2.2.2 How the adjacent possible evolves

From an evolutionary perspective, the development of the adjacent possible can be conceptualized as a dynamic, non linear process unfolding across four major phases: initial exploration, the emergence of structured complexity, accelerated combinatorial growth, and eventual systemic constraint. In the initial phase, a system is typically composed of a small number of elements or basic units, which limits the range of available recombinations but also renders the adjacent possible relatively simple and tractable. Because the system's components are few and the interactions among them are minimally constrained, novelty tends to emerge through straightforward combinations of already available elements. As such, this phase is characterized by a relatively linear and incremental innovation trajectory, where each newly realized state adds only modestly to the space of accessible configurations.[66] This early stage reflects the low dimensional geometry of the possibility space, in which the adjacent possible expands gradually and predictably, driven primarily by direct linkages between actualized and nearby unactualized nodes.[67]

As the system evolves, it transitions into a phase marked by the emergence of structured complexity. This is facilitated by the cumulative integration of novel elements and the formation of increasingly sophisticated interactions among them. Here, the adjacent possible expands in a non linear manner: each realized novelty not only enlarges the set of available options but also enhances the system's capacity to combine elements across greater distances in conceptual or physical space.[68] The system begins to exhibit modularity, hierarchical organization and specialized subsystems, each contributing to the diversification of the innovation landscape. In this phase, the adjacent possible is no longer a simple shell around current actualizations but a multi dimensional space whose topography becomes increasingly difficult to map in advance.[69]

The third phase is characterized by a sharp acceleration in the rate of novelty production, a combinatorial explosion that arises once the system has accumulated a sufficiently large and diverse repertoire of elements. At this point, the adjacent possible begins to grow at a super

66 M. Cortês et al., *"The TAP Equation: Evaluating Combinatorial Innovation." arXiv* (2022); Taalbi, *"Long-Run Patterns."*
67 Björneborn, *"Adjacent Possible."*
68 Loreto et al., *"Dynamics on Expanding Spaces."*
69 Devereaux, Koppl and Kauffman, *"Creative Evolution in Economics."*

exponential rate, as each new configuration enables disproportionately more recombinations. This phenomenon is formally captured in the TAP (Theory of the Adjacent Possible) equation, which models the number of new elements generated at time t as a function of all possible combinations of previously existing elements. According to the TAP equation, when the number of combinable components crosses a certain threshold, the system can enter a blow up regime, in which the number of reachable new states increases so rapidly that the system approaches a practical infinity of possibilities in finite time.[70] This phase is not merely theoretical: empirical patterns in technological evolution, patent networks and innovation genealogies frequently exhibit such bursts of radical expansion, confirming the TAP model's predictive power.[71]

However, such explosive growth is neither sustainable nor unbounded. Eventually, the system encounters a series of internal and external constraints that act to modulate and channel the evolution of the adjacent possible. These constraints may take the form of physical resource limitations, cognitive saturation, institutional or regulatory barriers or even path dependencies that restrict viable trajectories. As the complexity of the system increases, so too does the likelihood that new configurations will interfere with or destabilize existing ones, resulting in diminishing returns on novelty and a slowdown in effective exploration.[72] Moreover, feedback mechanisms, both positive and negative, begin to structure which areas of the adjacent possible are accessible, feasible or desirable. This stage does not imply stagnation; rather it introduces a pattern of punctuated equilibrium, in which phases of rapid expansion alternate with periods of consolidation and constraint. In this way, the evolution of the adjacent possible space resembles a dialectic between generativity and limitation, a recursive loop through which systemic innovation is simultaneously enabled and bounded by the history of its own emergence.[73]

---

[70] Cortês et al., *"The TAP Equation."*
[71] Taalbi, *"Long-Run Patterns."*
[72] Kauffman, "*The Origins of Order*."
[73] Björneborn, *"Adjacent Possible"*; Loreto et al., *"Dynamics on Expanding Spaces."*

2.3. Characteristics and types of adjacent possible

A further analytical deepening of the theory of the adjacent possible requires a discussion of its defining characteristics and the various types of adjacency that can structure the evolving possibility space. Among the most distinctive features of the adjacent possible is the ambiguity of causality that arises in systems embedded in creatively evolving environments. In traditionally modeled systems, causality is assumed to be deductive or probabilistically determinable, with outcomes following clearly from antecedents. However, within the adjacent possible framework, this linear causality breaks down. As Devereaux, Koppl and Kauffman argue, agents operating within such systems are not merely observers of a pre existing state space; they are themselves creators of that space through their decisions and interactions.[74] In this epistemological position, causal chains are not only hard to trace, they are being rewritten in real time. When an agent actualizes a novel configuration, that action does not simply select from known options, it restructures the landscape of what is even possible. As a result, actions are not only contingent but generative, making causality in these systems both distributed and temporally non local. This characteristic ambiguity disrupts standard notions of optimization and forecasting, replacing them with path dependence and a radical openness to unforeseen trajectories.[75]

A second critical feature of the adjacent possible is the pluralism of knowledge among agents, which allows for disagreement without contradiction. In creatively evolving systems, no central observer has access to a complete or stable representation of the entire adjacent possible space. Instead, agents operate with locally bounded, historically situated knowledge, leading to a multiplicity of perspectives and strategies. As Devereaux point out, "*all action is entrepreneurial action*", meaning that each decision is based on the actor's particular horizon of knowledge and interpretation of the system's dynamics.[76] This epistemic plurality implies that agents may "agree to disagree" on what is possible, desirable or even real, because the adjacent possible is not a fixed ontological entity but an evolving set of affordances that appears differently depending on one's position within the system. This characteristic challenges classical assumptions of rational consensus and stable expectations, suggesting instead a model of decision making that is speculative, creative and inherently non convergent. Moreover, the process of navigating the adjacent possible is shaped by what Devereaux and colleagues term "local knowledge" which is both the cognitive resource and the generative mechanism through

---

[74] Devereaux, Koppl and Kauffman, *"Creative Evolution in Economics."*
[75] Devereaux, Koppl and Kauffman, *"Creative Evolution in Economics."*
[76] Devereaux, Koppl and Kauffman, *"Creative Evolution in Economics."*

which novelty emerges. This implies that not only is the adjacent possible unknowable in its entirety, but that even its partial contours are subject to disagreement, reinterpretation and invention.[77]

Beyond its defining characteristics, the theory of the adjacent possible also entails a taxonomy of different types of adjacency, which help clarify the structural conditions under which novelty may emerge. One important distinction is between accessible and inaccessible regions of the adjacent possible. Accessibility here refers not to physical reachability per se, but to structural compatibility and developmental readiness. Some portions of the possibility space are adjacent in principle but cannot be accessed in practice without the prior realization of certain precursor states. For example, in the evolution of technology, the invention of the internet made the development of social media platforms structurally possible but these platforms remained inaccessible until key intermediate tools (web browsers, user interfaces) were actualized.[78] This distinction mirrors concepts in systems biology and evolutionary theory where certain mutations or phenotypes become available only after the establishment of specific genetic or environmental conditions.[79] In this sense, accessibility is not static but path dependent: what is reachable now depends on what has already been achieved.

A second typology concerns the difference between physical adjacency and cognitive adjacency. Physical adjacency pertains to material or spatial relationships, such as the proximity of molecules in a chemical reaction network or the compatibility of technologies in an engineering context. In contrast, cognitive adjacency relates to conceptual or perceptual relationships between ideas, theories or problem solving strategies. For example, in the process of scientific discovery, new hypotheses are often generated not by direct observation but by analogical reasoning, by cognitively leaping to a domain that is structurally similar but physically remote.[80] This distinction is crucial because many high impact innovations occur through shifts in cognitive adjacency rather than mere physical recombination. The case of Einstein's theory of relativity, as often cited in historical epistemology, did not arise from a new physical component but from a reconceptualization of time and space, a paradigmatic shift in cognitive adjacency. Importantly, physical and cognitive adjacencies are often connected: a

[77] Devereaux, Koppl and Kauffman, *"Creative Evolution in Economics."*
[78] Taalbi, *"Long-Run Patterns."*
[79] Kauffman, "*The Origins of Order*."
[80] Björneborn, *"Adjacent Possible."*

new material configuration may reveal new conceptual affordances and vice versa, creating a recursive loop between doing and thinking that propels the adjacent possible forward.[81]

In summary, the adjacent possible is not merely an abstract model of future potential; it is a richly structured and deeply contextual framework governed by a set of epistemological, ontological and relational characteristics. Its generative power derives from its ambiguity: of causality, of representation, and of accessibility. The space of the possible is shaped not only by what exists but by how agents interpret, engage with and construct the affordances around them. Whether physically adjacent or cognitively adjacent, accessible now or dependent on prior actualizations, each node in this evolving space redefines the landscape it belongs to. Understanding these features is essential to modeling not only the structure of innovation but the deep unpredictability and creative logic that underlies complex adaptive systems.

[81] Loreto et al., “Dynamics on Expanding Spaces.”

2.4. Empirical evaluation

This section concludes the chapter by examining empirical evidence and statistical regularities observed in real world systems, offering a critical test of the theoretical claims underpinning the adjacent possible framework. Specifically, it presents an empirical evaluation of the Generalized Urn Model with Triggering (GUMT) by applying a set of formally defined metrics: Gini coefficient, Youth coefficient, Recentness and Local Entropy, to diverse datasets drawn from domains such as digital music consumption, social media, collaborative software development, online encyclopedias and literary corpora. These systems, though heterogeneous in form and function, each exhibit innovation dynamics that can be interpreted through the lens of the adjacent possible: as new elements are introduced, they reshape the landscape of what is next possible. By quantitatively assessing how novelty competes with familiarity across time, this empirical analysis not only validates the theoretical predictions of the model but also illustrates how different environments foster or inhibit the expansion into adjacent possible spaces.

2.4.1. GUMT (Generalized Urn Model with Triggering)

The dynamics of novelty emergence in complex adaptive systems can be formally captured using probabilistic reinforcement schemes, among which the Polya urn model is foundational. In its classical formulation the model assumes an urn initially filled with balls of various colors, each color representing a categorical element within a bounded possibility space. At each discrete time step t a ball is drawn at random, its color is noted and it is then returned to the urn along with $\rho \geq 1$ additional balls of the same color. This reinforcement process, where frequent events become increasingly likely results in a power law distribution over frequencies, a statistical signature frequently observed in self organizing systems such as language formation, citation networks and economic growth models.[82] Mathematically, if $X_t \in \mathcal{C}$ denotes the color drawn at time t, and $C_t(c)$ the number of balls of color $c \in \mathcal{C}$ at time t, then the selection probability follows:

$$P(X_t = c) = \frac{C_{t-1}(c)}{\sum_{c' \in \mathcal{C}} C_{t-1}(c')}$$

The key dynamic of the model lies in its update rule, where $C_t(c) = C_{t-1}(c) + \rho$ if the drawn color is c, and remains unchanged otherwise. While this model elegantly formalizes

[82] Strogatz, "*Nonlinear Dynamics and Chaos*."

reinforcement and path dependency, it is inherently limited by the assumption of a fixed possibility space and it cannot accommodate the introduction of genuinely new elements.

To overcome this limitation, the Generalized Urn Model with Triggering (GUMT) introduces two critical extensions: a semantic labeling of elements and a mechanism for the endogenous expansion of the space of possibilities.[83] Each element in the urn, now interpreted not merely as a static type but as a contextually situated entity, is assigned a label $\kappa \in \mathcal{K}$, denoting its semantic or functional category (a music genre, topic, or technological domain). At time step *t* the last drawn element bears label κ and the urn is partitioned accordingly into elements that are semantically related (share the label κ) and those that do not. Let us define the following quantities:

- N_κ: the number of elements in the urn with label κ,
- $N_{\neg\kappa}$: the number of elements in the urn without label κ,
- $\gamma \in [0,1]$: a tunable parameter that controls the system's tendency to favor exploitation of familiar categories.

The weight associated with drawing an element that does not share the label of the last drawn item (a known but semantically unrelated element) is given by the expression

$$\gamma f(N_\kappa, N_{\neg\kappa}),$$

where f is a bounded, monotonic function reflecting the semantic relevance between the current context and the new candidate.[84] Two typical forms for f are employed in the literature:

- A weighted ratio, accounting for exploration bias: $f(N_\kappa, N_{\neg\kappa}) = \frac{N_\kappa}{N_\kappa + \gamma N_{\neg\kappa}}$,
- An unweighted ratio, where all elements contribute equally: $f(N_\kappa, N_{\neg\kappa}) = \frac{N_\kappa}{N_\kappa + N_{\neg\kappa}}$.

To illustrate the model's mechanics, consider a music recommendation platform where each song is categorized by genre. Suppose the user has just listened to a jazz track and the system records this as a draw of an element with label κ =jazz. Let the urn now contain N_κ = 3 jazz tracks and $N_{\neg\kappa}$ = 5 non-jazz tracks (pop, electronic). Using the unweighted choice of f, we have:

$$f(N_\kappa, N_{\neg\kappa}) = \frac{3}{3+5} = 0.375$$

Assuming an exploitation parameter of γ=0.8, the computed weight for drawing a previously encountered but semantically unrelated song becomes:

[83] Vittorio Loreto et al., *"Statistical Physics of Social Dynamics."* Reviews of Modern Physics 81, no. 2 (2009).
[84] Tria et al., *"The Dynamics of Correlated Novelties."*

$$\gamma f(N_{\kappa}, N_{\neg\kappa}) = 0.8 \times 0.375 = 0.3$$

This weight is then compared to those of other competing categories (known and related, or novel and related items) to determine the probability distribution over possible next selections. The function f thus plays a central role in regulating the system's exploration and exploitation trade off. It ensures that semantic similarity is not a fixed attribute but one dynamically inferred from the current composition of the urn, reflecting what has already been discovered and reinforced. The parameter γ, in turn, governs how strongly the system favors previous knowledge, thereby encoding a tunable cognitive inertia.[85]

In this formulation, the GUMT model not only preserves the reinforcement dynamics of the classical Polya process but crucially introduces a mechanism for novelty generation, contextual dependency and semantic proximity. It provides a principled, probabilistic framework through which the adjacent possible unfolds over time, where each new draw potentially reshapes the underlying space of what can be next discovered.[86]

2.4.2. The coefficients

To validate such models and render their dynamics intelligible from an economic perspective, a set of empirical coefficients has been defined to quantify how novelty emerges, competes and persists in real world data. These metrics provide economists with interpretable indicators of innovation cycles, popularity inequality and the temporal structure of adoption patterns to characterize the waves of novelties observed in data.

The Gini coefficient, originally developed to quantify income inequality in economics, can be adapted to measure inequality in the distribution of popularity across elements introduced over time. It provides a critical insight into the reinforcement dynamics underlying systems governed by the adjacent possible.[87] The primary goal of using the Gini coefficient in this context is to assess whether older elements (those introduced early in a sequence) dominate cumulative attention or whether popularity is more evenly distributed across items regardless of their introduction time.[88] The computation begins by sorting all elements (songs, hashtags, or words) in the order of their appearance in the dataset and then calculating their total popularity, defined as the number of times each element occurs in the sequence.

Let us say we have 5 elements (A to E), introduced in this order:

---

[85] Monechi et al., *"Waves of Novelties."*

[86] Kauffman, "*Investigations*."

[87] Corrado Gini, *"Measurement of Inequality of Incomes."* The Economic Journal 31, no. 121 (1921): 124–126.

[88] Monechi et al., *"Waves of Novelties."*

Element	**First Introduced**	**Popularity (Number of Occurrences)**
A	1	40
B	2	30
C	3	20
D	4	7
E	5	3

Table 3

Their cumulative popularity would be 40+30+20+7+3=100.[89] From this, one computes the cumulative share of both the population (ordered by time of introduction) and popularity, as shown in *Table 4*.

Cumulative % of Elements	**Cumulative % of Popularity**
1/5 = 0.20	40/100 = 0.40
2/5 = 0.40	(40+30)/100 = 0.70
3/5 = 0.60	(40+30+20)/100 = 0.90
4/5 = 0.80	(40+30+20+7)/100 = 0.97
5/5 = 1.00	1.00

Table 4

These values are then plotted to construct the Lorenz curve, where the x-axis represents the cumulative share of elements and the y-axis the cumulative share of popularity (*Figure 5*).

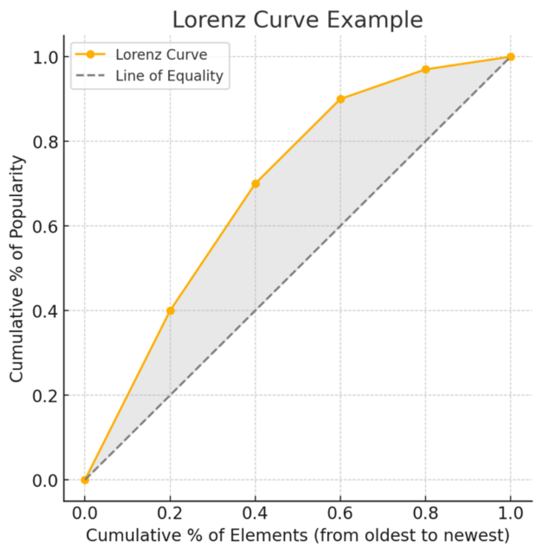

Figure 5

[89] Peter Lambert, *"The Distribution and Redistribution of Income."* 3rd ed. (Manchester: Manchester University Press, 2001), 32–36.

The Lorenz curve is then compared to the 45 degree line of perfect equality, and the Gini coefficient is defined as twice the area between this diagonal and the curve.[90] Mathematically this is expressed as $G=1-2A$, where A is the area under the Lorenz curve normalized to fall within the interval [0,1].[91] A Gini coefficient near 1 indicates that a small fraction of early items capture the vast majority of popularity (a highly unequal distribution), whereas a Gini near 0 implies that attention is more evenly spread regardless of age. In the context of the adjacent possible, this coefficient becomes a tool for understanding systemic inertia: a high G suggests that early discoveries continue to dominate, limiting the emergence of novel items, while a low G reflects a system more open to innovation.[92]

The Youth coefficient (Y) is a temporal indicator designed to measure how quickly a system renews its trending or popular elements, offering a dynamic lens through which one can assess the pace at which novelty emerges in an evolving environment.[93] While the Gini coefficient evaluates inequality in cumulative popularity, the Youth coefficient captures whether newly introduced items are increasingly present in the most recent periods of system activity.[94] This is particularly relevant in systems driven by innovation, where one wishes to understand whether new items such as songs, ideas, or technologies are quickly gaining traction or whether older elements persist over time. To compute Y, the sequence of events is divided into equal-length, non overlapping time windows (say every 100 observations) and in each window the average introduction time of the items that appear is calculated. This creates a sequence of average values, one for each interval.[95] These values are then plotted against the index of each window (first, second, third, etc.), and a linear regression is applied to fit a line through the data. The slope of this line, denoted by λ, indicates the rate at which the system is adopting newer items: a steeper slope suggests a greater presence of recent elements in successive windows.[96] To ensure comparability across datasets with different time window lengths, the coefficient is normalized as $Y = \frac{\lambda}{\Delta\tau}$, where $\Delta\tau$ is the length of the time window. This normalization constrains Y to the interval [0,1] allowing for intuitive interpretation.[97] A value

---

90 Donald B. Rubin, *"The Calculation of Gini Coefficients."* The Review of Economics and Statistics 45, no. 1 (1963): 50–52.
91 Loreto et al., "*Statistical Physics of Social Dynamics*." 602.
92 Monechi et al., "*Waves of Novelties*." Table 1.
93 Tria et al., "*The Dynamics of Correlated Novelties*."
94 Monechi et al., "*Waves of Novelties*."
95 Monechi et al., "*Waves of Novelties*."
96 Peter J. Brockwell and Richard A. Davis, "*Introduction to Time Series and Forecasting"*, 2nd ed. (New York: Springer, 2002).
97 Monechi et al., *"Waves of Novelties."*

of Y close to 1 indicates that every new window contains almost exclusively novel items, suggesting a system characterized by constant innovation and high temporal turnover. By contrast, Y near 0 implies that the composition of active or popular items remains largely unchanged over time, signaling stagnation or inertia.

The Recentness coefficient (R) captures the degree to which the most popular elements in a system at each moment in time are newly introduced.[98] Unlike the Youth coefficient, which examines the average age of all items in use during a given time window, Recentness focuses specifically on the leading item, the single most popular entry in each interval. To compute R, the dataset is divided into equal time windows. In each window the system identifies the most popular item and records the time at which that item was first introduced.[99] The sum of these introduction times forms the numerator. The denominator represents the theoretical maximum this sum could take: it is the sum of the latest possible introduction times had the most popular item in each window always been the most recently introduced.[100] The ratio between the two values yields the Recentness coefficient, which ranges from 0 (purely conservative system) to 1 (system fully favoring the newest items). A simple example illustrates the extremes of the metric. Imagine a system divided into three time windows where, in each, the most popular item is also the newest available one, as seen in *Table 5*.

Window	Most Popular Item	Time of Introduction
1	A	1
2	B	2
3	C	3

Table 5

Numerator (actual sum of intro times): 1 + 2 + 3 = 6;

Denominator (max possible sum): still 1 + 2 + 3 = 6

$R = \frac{6}{6} = 1$

This means the system is highly recent, new items quickly become the most popular.

By contrast, consider a system in which the same item, say item A introduced at time 1, remains the most popular across all three windows (*Table 6*).

[98] Monechi et al., *"Waves of Novelties."*

[99] Monechi et al., *"Waves of Novelties."*

[100] Tria et al., "*The Dynamics of Correlated Novelties."*

Window	Most Popular Item	Time of Introduction
1	A	1
2	A	1
3	A	1

Table 6

Numerator (actual sum): 1 + 1 + 1 = 3;

Denominator (max possible sum): 1 + 2 + 3 = 6,

$R = \frac{3}{6} = 0.5$

This means the system tends to keep old items in the spotlight, even when new ones appear.

In the framework of adjacent possible theory, R offers a lens into how quickly a system capitalizes on new possibilities: a high Recentness score suggests that novelties do not just expand the space of potential choices but are actively elevated to prominence, while a low score reflects structural conservatism, in which innovation is permitted but rarely rewarded with immediate visibility.[101]

The **Local Entropy coefficient**, denoted ⟨h⟩, measures the diversity of attention or popularity within each time window of a dynamic system, providing a statistical lens through which to assess whether attention is monopolized by a few dominant items or more evenly distributed among many.[102] Unlike the Gini or Recentness coefficients, which assess long term accumulation or leading trends, ⟨h⟩ captures the short term competitive structure of the system. To compute it the sequence of events is divided into equally sized time windows. In each window the relative popularity (frequency of appearance) of each item is calculated. These frequencies are then used to compute the Shannon entropy, a measure from information theory that quantifies uncertainty or dispersion in a probability distribution.[103] Formally, for a given window *w*, the entropy is calculated as:

$$h_w = -\sum_{j=1}^{n_w} p_j \log p_j$$

[101] Kauffman, "*Investigations*."
[102] Monechi et al., *"Waves of Novelties."*
[103] Claude E. Shannon, *"A Mathematical Theory of Communication*." Bell System Technical Journal 27, no. 3 (1948): 379–423.

where p_j is the relative frequency of item j in window w, and n_w is the number of distinct items in that window. To allow comparison across windows with different numbers of items, this entropy value is normalized by dividing it by the maximum possible entropy $\log n_w$, yielding a value between 0 and 1.[104] The Local Entropy coefficient ⟨h⟩ is then the average of these normalized entropies across all time windows:

$$\langle h \rangle = \frac{1}{W} \sum_{w=1}^{W} \frac{h_w}{\log n_w}$$

where W is the total number of windows.[105] A value of ⟨h⟩ close to 1 indicates high diversity, meaning that many items share attention within each window, while values near 0 indicate that popularity is concentrated on just one or two items. For example, consider two windows: in the first, if one song accounts for 95% of plays while others are barely played the entropy will be very low. In contrast, if ten songs are all played roughly equally the entropy will be much higher.

2.4.3 Empirical evaluation in different domains

To assess the empirical validity of the adjacent possible theory and the predictive capacity of the Generalized Urn Model with Triggering (GUMT), this section draws upon data and analyses presented in the article "*Waves of Novelties in the Expansion into the Adjacent Possible*" by Monechi, Tria, Ruiz-Serrano, and Loreto. The authors apply the GUMT framework to a diverse set of large scale, real world datasets spanning various domains of cultural and technological activity. By evaluating the behavior of four key statistical indicators: the Gini coefficient, Youth coefficient, Recentness, and Local Entropy. The study examines how novelty emerges and is reinforced across systems such as music streaming, social media, software development, encyclopedic knowledge production, and literary composition. These empirical investigations offer a quantitative test of the adjacent possible hypothesis: that each realized novelty opens up new pathways for further innovation and that the structure of attention and popularity within a system reflects its capacity to explore and populate those new possibilities.

The datasets were drawn from Last.fm, Twitter, GitHub, Wikipedia, and Project Gutenberg, capturing a spectrum from online social behavior to textual production. Last.fm provides a

[104] Peter Harremoës and Flemming Topsøe, *"Maximum Entropy Fundamentals,"* Entropy 3, no. 3 (2001): 191–226.
[105] Monechi et al. *"Waves of Novelties."*

sequence of songs played on a music platform (revealing how new artists or genres enter listeners' repertoires); Twitter offers a stream of hashtags created in tweets (tracking the emergence of new topics in social discourse); GitHub contains the history of software project creations and contributions (reflecting innovation in open-source development); Wikipedia covers the sequential editing or word introduction in a large scale crowd sourced knowledge base; Gutenberg comprises the text of literary works (a proxy for the introduction of new words or concepts in written language over the course of a novel). These five contexts were chosen because they mirror human activities with an element of novelty creation: music consumption, communication trends, software innovation, knowledge accumulation, and literary creativity. The question was whether all these systems, despite their differences, exhibit similar patterns in how novelty competes with the familiar. The empirical results show both commonalities and clear contrasts across the datasets. The results are displayed in *Table 7* below, taken from the article "*Waves of Novelties in the Expansion into the Adjacent Possible"* by Monechi, Tria, Ruiz-Serrano, and Loreto.

System	**G**	**Y**	**R**	**⟨h⟩**
Last.fm	0.491	0.379	0.516	0.982
Twitter	0.405	0.463	0.448	0.961
GitHub	0.706	0.339	0.386	0.907
Wikipedia	0.889	0.035	0.020	0.930
Gutenberg	0.950	0.0103	0.0277	0.909

Table 7

In all cases the *Gini-like coefficient G* was found to be positive ($G > 0$), underscoring a baseline advantage for early entrants (the first elements introduced tend to accumulate more total popularity than those introduced later). However, the magnitude of G varied significantly: in the music (Last.fm), social media (Twitter) and coding (GitHub) data G was moderate (around 0.4-0.7 in value), meaning that while older songs, hashtags, or projects do have an edge, newer ones collectively still command a substantial share of attention. In fact, G in these systems was considerably lower than it would be under a randomized baseline, indicating that real temporal dynamics allow more egalitarian popularity outcomes than a static null model. By contrast, in the textual datasets (Wikipedia and Gutenberg) G was very high (≈0.9), close to the theoretical maximum, implying that the first introduced words or topics utterly dominate the frequency distribution. This resonates with intuition: in a book or encyclopedia core terms introduced early (characters, common concepts) appear repeatedly whereas later introduced words are comparatively rare.

The *Youth coefficient Y* further highlights this difference. For Last.fm, Twitter and GitHub, Y ranged roughly 0.35-0.46, signaling a meaningful pace of renewal: on average, each successive time window contains a fair number of newly introduced popular items (Y well above 0). These systems exhibit a rejuvenation effect: the trending content in, say, each week tends to be newer than that of the previous week, to a much greater degree than one would expect by chance. In the randomized (time shuffled) version of the data, Y fell near zero (since shuffling breaks any temporal novelty pattern), whereas the real data's Y was an order of magnitude higher. For example, Twitter's Y≈0.46 suggests that trending hashtags are frequently fresh ones rather than the same old tags, reflecting the platform's preference for constantly evolving conversations. In GitHub, Y was a bit lower (0.339), consistent with the idea that, while new projects do emerge, developers also continue to star or fork older, established repositories for longer. Meanwhile, Wikipedia and Gutenberg showed negligible Youth coefficients (Y≈0.01-0.04): the content that fills each successive segment of text is almost entirely drawn from the pool of words introduced early on, with virtually no new trending words later. This means a novel or an encyclopedia doesn't keep introducing popular new terms in later chapters, it largely exploits the vocabulary set that was established initially.

The *Recentness R* metric paints a similar picture. In Last.fm and Twitter, R was about 0.45-0.52, meaning that the most popular item in a given interval was as likely to be a recently introduced song or hashtag as it was to be a long established hit. In fact, R≈0.5 in those cases indicates a balance: sometimes an old favorite tops the charts but other times a brand new release (or meme) becomes the most popular of that interval. In GitHub, R was a bit lower (≈0.39), implying that the top project in each time window was more often an older repository (perhaps well known libraries getting more attention), but still sometimes a newcomer would become the most popular. In contrast, Wikipedia and Gutenberg had R near 0 (around 0.02-0.03), meaning that in essentially every segment of text the most frequent word was among the earliest introduced (often a common function word or a main subject word). A new word virtually never overtook the older ones in usage frequency at any point.

Finally, all domains showed high *local entropy ⟨h⟩*, generally between 0.90 and 0.98, indicating that within any given short time frame popularity was not totally concentrated in one item but rather spread across several items. Even in the presence of viral hits, people (or words in a text) still distribute attention to multiple options in parallel. The entropy was slightly lower in GitHub ($\langle h \rangle \approx 0.91$) and the literary texts (≈ 0.91 in Gutenberg) than in music or Twitter ($\langle h \rangle \approx 0.96$-$0.98$), suggesting that occasionally a single repository or a single concept could dominate a time period more strongly in those former cases. But overall, $\langle h \rangle \approx 1$ for both empirical and shuffled data, which implies that each interval typically had a diversity of popular elements rather than a monopoly.

Together, these results illustrate how the expansion into the adjacent possible unfolds differently across socio technical environments. In highly dynamical systems like online music consumption, social media and open-source development, we observe a moderate but palpable level of novelty turnover. New entries consistently arise and capture attention locally (high Y and R relative to baseline), yet older entries are never completely displaced (G remains positive). This corresponds to a balanced exploitation and exploration regime: the community continually explores adjacent possible innovations (new songs, topics or projects) while still remembering and exploiting the successful creations of the past. From an economics perspective, these findings underscore the importance of a mixed strategy in innovation dynamics: systems that foster ongoing novelty (analogous to competitive markets with new entrants) tend to show transient surges of new successes without entirely overturning the old hierarchy, whereas systems that are too exploitative can stagnate with entrenched incumbents. The adjacent possible theory, quantified through GUMT and its empirical validation, suggests that the key to sustained innovation lies in maintaining a balance, allowing the exploration of new possibilities at a reasonable rate while retaining some continuity with the past.

Chapter 3: Heuristics and ecological rationality

3.1. Introduction

The previous chapter explored how innovation expands the adjacent possible, an ever evolving space of new opportunities that cannot be fully mapped in advance. Each novel discovery opens further possibilities, making the future fundamentally open ended. In such an environment of radical uncertainty (in the Knightian sense, where not all outcomes or probabilities are known), traditional models of fully informed optimization lose traction.[106] This raises a set of important questions: What can agents do under uncertainty? How should they act when the very structure of possibilities is uncharted and continually unfolding? And how can innovation be pursued in these conditions, given that no deterministic method or algorithm can reliably chart a course through wholly novel terrains? These questions set the stage for the present chapter.

One plausible answer is that agents rely on heuristics, that is, simplified decision rules or rules of thumb based on experience, which guide action when calculation is impractical. Rather than attempting to compute an optimal strategy (impossible under genuine uncertainty), boundedly rational decision makers fall back on heuristics as practical guides. The notion of bounded rationality, introduced by Herbert Simon, holds that real economic agents face cognitive limits and incomplete information and thus satisfice (seek outcomes that are good enough via heuristics) instead of optimizing in any global sense.[107] Far from being irrational, such heuristics may be adaptive responses to complexity. Indeed, simple heuristics can outperform elaborate models in environments where the future is fundamentally unpredictable and outcomes are not enumerable. Heuristics are seen as part of an adaptive toolbox of strategies that are selected not to maximize a known objective but to fit the specific demands of the environment. Rather than treating heuristics as biased approximations, this approach frames them as necessary and effective tools for coping with uncertainty.[108]

Emerging research in economics similarly suggests that heuristic driven behavior can be not only effective but rational under conditions of deep uncertainty. For example, in an agent based model of a complex evolving economy with technological change, it has been shown that firms using simple fast and frugal rules of thumb can perform as well as, and occasionally better than, those relying on more sophisticated predictors. In such contexts, robust heuristics are not

[106] Frank H. Knight, "*Risk, Uncertainty and Profit.*" (Boston: Houghton Mifflin, 1921).
[107] Herbert A. Simon, "*Administrative Behavior*." 4th ed. (New York: Free Press, 1997), 88–90.
[108] Gerd Gigerenzer and Wolfgang Gaissmaier, *"Heuristic Decision Making."* Annual Review of Psychology 62 (2011).

merely second best approximations but rational responses to ever changing environments.[109] In the absence of a clearly defined optimizing solution, economic agents and organizations rely on such heuristics to experiment, adapt and progressively expand the frontier of the adjacent possible.

In summary, heuristics offer a compelling framework for understanding how boundedly rational agents navigate open ended complex settings. They provide plausible answers to the challenges of uncertainty and innovation when formal optimization fails. This chapter opens with a conceptual definition of heuristics within the context of decision making under uncertainty. It then illustrates their practical applications through selected examples, followed by an analysis of their origins and development.

[109] Giovanni Dosi, *"Rational Heuristics: Expectations and Behaviors in Evolving Economies with Heterogeneous Agents."* Economic Inquiry 58, no. 1 (2020).

3.2. Definition and Concepts

Gerd Gigerenzer's *"Why Heuristics Work"*[110] situates heuristics alongside logic and probability as a fundamental approach to human rationality in the face of uncertainty. Whereas classical logic and probabilistic models strive for exhaustive consistency or optimality, heuristics are characterized by frugality. They deliberately ignore portions of available information and by satisficing rather than optimizing, seeking good enough solutions instead of mathematically optimal ones. Contrary to the common view that heuristics are merely second best shortcuts necessitated by cognitive limitations, Gigerenzer argues that simple heuristics can often rival or even surpass more complex analytic methods in real world decision tasks. He challenges several misconceptions about heuristic reasoning: for example the notion that optimization is always superior to heuristics or that humans use heuristics only because of mental constraints, does not hold universally. In many situations, finding the optimal solution is impossible or impractical, either because the problem is computationally intractable or because attempting to optimize leads to overfitting errors, and in such cases a well chosen heuristic can yield more robust and accurate judgments. This perspective reframes heuristics not as cognitive biases or imperfections but as adaptive strategies shaped by both mind and environment to deal with complexity and uncertainty.

At the heart of Gigerenzer's thesis is the concept of the adaptive toolbox, a Darwinian inspired model of the mind as a collection of specialized cognitive heuristics, their building blocks and the evolved capacities that support them. In this view, human rationality is bounded but not irrational, our minds evolved mental abilities (such as memory, perception of frequencies or recognition) to make decent decisions with limited time and information. Each heuristic in the toolbox is a module tuned to particular types of tasks or environments. For instance, the recognition heuristic capitalizes on the simple fact of whether one recognizes an option or not, leveraging the basic cognitive capacity for recognition memory. Such capacities are not unlimited; indeed, a degree of forgetting or ignorance can improve decision making, an idea captured by the "less is more" effect, where knowing less can lead to better judgments if knowledge beyond a point only adds noise. Gigerenzer builds on Herbert Simon's notion of bounded rationality and the analogy of mind and environment as two blades of scissors: a heuristic's effectiveness comes from the fit between the cognitive strategy and the structure of the environment in which it operates. Rather than viewing departures from logic or probability

[110] Gerd Gigerenzer, *"Why Heuristics Work,"* Perspectives on Psychological Science 3, no. 1 (2008): 20–29.

as errors, this ecological view defines rationality by how well a thinking process succeeds in the world, highlighting that what counts as a rational decision strategy depends on the environment's demands.

3.2.1. Principles Explaining Why Heuristics Work

Gigerenzer and colleagues[111] outline several key principles that explain the power of heuristics in uncertain, complex domains. These principles form a framework for understanding when and why simple rules of thumb can be so effective:

Formal Modeling and Precision: Cognitive heuristics should be described with explicit computational models, not just vague labels like availability or ad hoc dual system accounts. By formalizing a heuristic's process (specifying how information is searched and stopped), researchers can derive clear predictions and rigorously test them. Such models have revealed cases where a simple heuristic predicts outcomes more accurately than do sophisticated multivariate models, undermining the assumption that heuristic reasoning is necessarily inferior. Formal models thus help identify the conditions where more information and computation improve decision accuracy versus where ignoring information is advantageous.

Tractability: Human minds (and even computers) often face decision problems that are computationally intractable, that is, no feasible computation can guarantee the optimal answer. Many real world problems (finding an optimal investment portfolio, planning complex routes, playing games like chess) are NP hard, meaning an exhaustive search of options is practically impossible. Heuristics succeed by simplifying these problems, making decision making tractable. Rather than attempting the impossible task of optimizing over astronomically many possibilities, a heuristic focuses on a manageable subset of cues or alternatives, trading some theoretical accuracy for immense gains in speed and feasibility. In short, we use heuristics not only because our cognitive resources are limited, but because the world's complexity often exceeds the capacities of any optimization approach.

Robustness to Uncertainty: A well designed heuristic tends to be robust, it avoids overfitting noisy data and thus predicts better in new situations. Complex models that finely tune themselves to past observations (for instance, by weighing dozens of factors) can capture idiosyncrasies that won't repeat, thereby mispredicting future cases. In contrast, heuristics deliberately ignore the noise and only use the most important information, thereby generalizing more reliably. Research shows that limiting information and computation, even traits like

[111]Gerd Gigerenzer, *"Why Heuristics Work."*

memory constraints or forgetting, need not be regrettable deficiencies but can actually enhance performance in an uncertain world. By cutting through irrelevant details, heuristics reduce the risk of mistaking random variation for enduring signal.

Evolved Capacities: Heuristics are effective in part because they use natural cognitive abilities that evolution has already optimized. Instead of mirroring statistical equations, heuristics often find solutions that a human brain can execute easily by exploiting innate or learned capacities. For example, the recognition heuristic relies on our ability to quickly recognize familiar names or objects, a capacity that required no formal training. The adaptive toolbox perspective holds that our minds contain many such domain specific skills (from visual perception to social intuition) and heuristics are tailored to take advantage of them. Notably, these capacities themselves might be bounded (memory is fallible and attention is limited) but those bounds can be features, not bugs, when they allow heuristics to ignore low priority information and focus on what matters.

Ecological Fit: The success of a heuristic depends on the match between its structure and the environment's structure, a concept known as ecological rationality. A given heuristic will work well in environments that provide the patterns or cues it needs, and will fail if applied in the wrong context. Gigerenzer emphasizes studying the environment (the second blade of the scissors) alongside the cognitive process: factors like the distribution of information, the cost of time and the typicality of cases determine which heuristic is appropriate. This principle explains why there is no universally best strategy, each heuristic has an intermediate range of problems to which it applies. The mind appears to adapt by selecting heuristics suited to the task at hand, through individual learning, social learning (imitation or instruction) or evolutionary selection of successful rules of thumb. In sum, heuristics work when they exploit the regularities of the environment and understanding those regularities is as important as understanding the cognitive rule itself.

3.3. Examples of heuristics

Heuristics, as practical cognitive strategies for decision making under uncertainty, come in diverse forms tailored to specific contexts.

The *recognition heuristic* is a decision making strategy that exploits a person's ability to recognize one option over another. It is predicated on the idea that when individuals are faced with a choice between two alternatives and only one of them is recognized, they infer that the recognized object has a higher value with respect to a certain criterion. This heuristic is especially effective in environments where recognition correlates with the desired outcome. Formally introduced by Goldstein and Gigerenzer, the recognition heuristic is defined as follows: "*If one of two objects is recognized and the other is not, then infer that the recognized object has the higher value with respect to the criterion*."[112] The cognitive efficiency of the recognition heuristic lies in its frugality, it makes a decision using minimal information, often just a single binary cue: recognition. It avoids the need to recall or evaluate other potentially available data. For example, consider a scenario where an individual is asked: which of two cities, Heidelberg or Erlangen has a larger population. If the individual recognizes Heidelberg but has never heard of Erlangen the heuristic would lead them to infer that Heidelberg is the more populous of the two. This inference is based not on any calculated evaluation, but solely on the recognition of the name Heidelberg. Empirical studies have demonstrated that this simple strategy can lead to highly accurate inferences, especially in domains where recognition is systematically related to the criterion being judged, such as city size, academic citation counts or company revenue.[113]

The *1/N heuristic* embodies a simple allocation strategy in which resources are distributed equally across all available options, without accounting for the specific characteristics or statistical profiles of each alternative. This rule based approach avoids the complexity of optimization and sidesteps the need for precise estimations of probabilities or payoffs, making it useful under conditions of uncertainty and limited information. In practical terms, it has been extensively studied in the domain of financial decision making, particularly in portfolio diversification. For instance, when an investor is faced with five different mutual funds and limited knowledge about their expected returns or covariances, applying the 1/N heuristic would result in allocating an equal 20% share to each fund. Despite its simplicity, research has shown that this heuristic can perform comparably to, and at times even outperform, more

---

[112] Daniel G. Goldstein and Gerd Gigerenzer, *"Models of Ecological Rationality: The Recognition Heuristic."* Psychological Review 109, no. 1 (2002).

[113] Daniel G. Goldstein and Gerd Gigerenzer, *"Models of Ecological Rationality."*

sophisticated optimization based models, especially when those models suffer from estimation errors in small samples or unstable environments. This robustness underscores its ecological rationality: the 1/N heuristic trades precision for reliability in complex real world contexts.[114]

The *fluency heuristic* is a cognitive shortcut that builds on recognition by considering the ease and speed with which information is recalled. When individuals recognize two options, they tend to prefer the one that comes to mind more fluently, assuming it holds greater value or relevance. This heuristic is efficient in environments where fluency correlates with meaningful outcomes, such as frequency or success[115]. For instance, Alter and Oppenheimer demonstrated that stocks with names easy to pronounce ("JAGA") outperformed those with complex names ("XAGY") in early market performance, showing how fluency can guide financial decisions[116]. Similarly, in consumer contexts, products with more readable labels are often judged as more effective[117].

The *take the best heuristic* functions by evaluating options based on a list of cues ordered by predictive validity. It searches through these cues sequentially and bases its decision on the first cue that discriminates between the alternatives. Importantly, once a distinguishing cue is found, the process stops, subsequent cues are ignored. For example, when predicting which of two German cities has a larger population, if one recognizes Hamburg but not Heidelberg, recognition alone might suffice; but if both cities are recognized, one might look at whether either is a state capital (a high validity cue). If only one is, the decision is made without considering further information. This strategy has been shown to match or even outperform more complex compensatory models in domains such as consumer choice and political forecasting.[118]

Tallying does not rank cues or stop after a single discriminating one. Instead, it considers all available cues equally, counting how many cues favor each alternative and choosing the one with the higher tally. For instance, when deciding between two job candidates, if one has

---

114 Victor De Miguel, Lorenzo Garlappi, and Raman Uppal, "*Optimal Versus Naive Diversification: How Inefficient Is the 1/N Portfolio Strategy?"* Review of Financial Studies 22, no. 5 (2009).

115 Gerd Gigerenzer and Peter M. Todd, *Simple "Heuristics That Make Us Smart"* (New York: Oxford University Press, 1999).

116 Adam Alter and Daniel M. Oppenheimer, *"Predicting Short-Term Stock Fluctuations by Using Processing Fluency,"* Proceedings of the National Academy of Sciences 103, no. 24 (2006).

117 Lael J. Schooler and Ralph Hertwig, "*How Forgetting Aids Heuristic Inference."* Psychological Review 112, no. 3 (2005).

118 Gerd Gigerenzer and Daniel Goldstein, "*Reasoning the Fast and Frugal Way: Models of Bounded Rationality."* Psychological Review 103, no. 4 (1996).

experience in leadership, speaks two foreign languages, and holds a relevant degree (three positive cues), while the other has only two of these traits, the first candidate is preferred[119].

The *availability heuristic* operates on the principle that the ease with which instances come to mind serves as a proxy for estimating frequency or probability. When individuals are asked whether tornadoes occur more often in Kansas or Nebraska, they are more likely to choose Kansas, not necessarily because of statistical data, but because media portrayals like The Wizard of Oz make tornadoes in Kansas more memorable. This mental shortcut, while efficient, can lead to systematic biases when vivid or recent events dominate memory recall, even if they are statistically rare.[120]

The *anchoring and adjustment heuristic* refers to the cognitive tendency to rely heavily on an initial reference point (the "anchor") and insufficiently adjust away from it when making quantitative judgments. In one classic experiment, participants asked to estimate the percentage of African countries in the United Nations were influenced by a random number generated by a spinning wheel: those shown a high number guessed higher percentages than those shown a low number.[121] This heuristic explains why sellers might set high initial prices, anchoring buyer expectations, even if the final price is negotiated down. While the initial anchor may be arbitrary or unrelated, its influence on subsequent judgments is substantial and persistent.

Finally, the *satisficing* heuristic, introduced by Herbert Simon, involves setting an aspiration level and selecting the first option that meets or exceeds this threshold, rather than exhaustively searching for the optimal choice.[122] For instance, a person searching for an apartment may decide in advance on a set of acceptable criteria (price under €800, less than 30 minutes from work, natural light) and choose the first listing that satisfies all conditions, rather than evaluating every possible apartment in the market. Satisficing is particularly useful in complex or time constrained environments, where exhaustive comparison is not feasible. This approach reflects bounded rationality in practice: decision makers aim for good enough solutions, conserving cognitive resources while still making effective choices.

---

[119] Konstantinos V. Katsikopoulos, "*Psychological Heuristics for Making Inferences: Definition, Performance, and the Emerging Theory and Practice,"* Decision Analysis 8, no. 1 (2011).

[120] Amos Tversky and Daniel Kahneman, "*Availability: A Heuristic for Judging Frequency and Probability."* Cognitive Psychology 5, no. 2 (1973).

[121] Amos Tversky and Daniel Kahneman, *"Judgment under Uncertainty: Heuristics and Biases." Science* 185, no. 4157 (1974).

[122] Herbert A. Simon, *"Rational Choice and the Structure of the Environment."* Psychological Review 63, no. 2 (1956).

These heuristics demonstrate how individuals make adaptive use of mental shortcuts when confronted with limited time, incomplete information, or cognitive constraints. Although they can lead to errors in certain contexts, they also often yield sufficiently accurate judgments for navigating everyday decisions in a complex world.

3.4. History of ecological rationality

For much of the twentieth century, economic thought was dominated by the neoclassical paradigm, which posited a model of rationality grounded in the figure of Homo economicus: an agent endowed with perfect information, unlimited cognitive resources and stable preferences, capable of calculating and optimizing outcomes to maximize utility in all conceivable situations. This portrayal, while mathematically elegant and normatively coherent, increasingly appeared inadequate in light of empirical evidence from psychology and behavioral studies that revealed systematic deviations from such idealized rationality. Beginning in the 1970s, Daniel Kahneman and Amos Tversky launched a transformative research agenda aimed at empirically documenting the limits of human judgment and decision making under uncertainty. Through a series of influential experiments, they demonstrated that people often rely on intuitive mental shortcuts (heuristics) which frequently led to systematic errors or cognitive biases. These findings culminated in the development of prospect theory, which showed that individuals evaluate outcomes relative to reference points and display loss aversion, contradicting the assumptions of consistent utility maximization.[123] Yet, despite their groundbreaking contributions, Kahneman and Tversky's interpretation of heuristics remained tethered to a deficiency model: heuristics were seen as byproducts of cognitive limitations, fast but flawed approximations of an unattainable rational ideal.[124]

This pessimistic view of heuristics, as sources of bias and irrationality, set the tone for much of behavioral economics for decades, reinforcing the idea that deviation from formal logical or probabilistic norms equated to error. However, a major conceptual shift emerged with the work of Gerd Gigerenzer and colleagues, who proposed an alternative framework known as ecological rationality.[125] In this paradigm, the rationality of a decision strategy is not judged by its internal consistency or alignment with formal models, but by its adaptive success in real world environments. That is, a decision rule is considered rational if it leads to successful outcomes given the structure of the environment in which it operates. This correspondence view of rationality stands in contrast to the coherence view of Kahneman and Tversky, which emphasized consistency with axioms of logic or probability theory. Within this reframed understanding, heuristics are not inferior substitutes for optimization; rather they are functional, evolved strategies that work precisely because they ignore information, reduce computational

---

[123] Daniel Kahneman and Amos Tversky, *"Prospect Theory: An Analysis of Decision under Risk," Econometrica* 47, no. 2 (1979).
[124] Amos Tversky and Daniel Kahneman, *"Judgment under Uncertainty."*
[125] Gerd Gigerenzer, *"Why Heuristics Work."*

demand and exploit statistical regularities in the environment. In conditions of deep uncertainty, where the space of possible outcomes is ill defined and information is costly or incomplete, adding more data or complexity does not necessarily improve decision quality. On the contrary, effective decisions often emerge from simple rules that deliberately ignore some information, leveraging domain specific cues to guide behavior efficiently.

The distinction between logic, probability and heuristics helps further clarify this shift. Logic demands deductive validity and internal consistency; probability theory prescribes the updating of beliefs based on new evidence via Bayesian norms; heuristics, by contrast, are fast and frugal rules that prioritize adaptability over formal rigor.[126] While logic and probability are suited to closed and defined problems, heuristics excel in open, complex and dynamically evolving environments, those that closely resemble the conditions in which humans actually operate. The implication is significant: rather than treating heuristics as cognitive flaws or second best solutions, the ecological rationality approach redefines them as indispensable cognitive tools. These tools are not merely shortcuts, but tailored responses to the structural features of particular tasks or environments. Their effectiveness lies not in approximating formal optimization, but in achieving sufficient accuracy with minimal effort.[127]

This reconceptualization has sparked the development of a science of heuristics, which seeks to empirically identify the environmental conditions under which particular heuristics succeed or fail.[128] This shift moves the field from cataloging biases to understanding the adaptive functions of decision strategies. Moreover, it has had practical applications in areas like medicine, finance and policy design, where simpler heuristics have been shown to outperform complex algorithms in noisy and uncertain settings. For example, in certain diagnostic tasks, physicians using fast and frugal trees made more accurate assessments than those relying on exhaustive checklists or probabilistic scoring systems. Similarly, financial investors relying on recognition based heuristics often outperform those attempting to integrate and weigh all available data, particularly under volatile market conditions.[129] These findings support Gigerenzer's broader thesis: the measure of rationality is not how closely decision makers adhere to normative axioms, but how well their heuristics match the demands of the

---

[126] Gerd Gigerenzer, "The Rationality Wars: Why Heuristics Work,". The Rationality Wars: A Personal Reflection, Max Planck Institute for Human Development, 2021.

[127] Gerd Gigerenzer, *"The Rationality Wars: Why Heuristics Work."*

[128] Konstantinos V. Katsikopoulos, *"Psychological Heuristics for Making Inferences: Definition, Performance, and the Emerging Theory and Practice." Decision Analysis* 8, no. 1 (2011): 10–29.

[129] Gigerenzer, *"Why Heuristics Work."*

environment.[130] In this sense, ecological rationality does not reject the insights of Kahneman and Tversky, but reframes them, acknowledging that while heuristics may sometimes produce biases, they are also the best available tools for navigating an uncertain, information rich world.

[130] Gerd Gigerenzer, Peter M. Todd and the ABC Research Group, "*Simple Heuristics That Make Us Smart*."

3.5. Conclusions on ecological rationality

The evidence assembled in this chapter compels a fundamental re evaluation of what it means to be rational when choices must be made in a world that is noisy and only partially knowable. Under such conditions, the normative benchmark of global optimisation dissolves into an ideal that cannot, even in principle, be realised. What survives is a repertoire of heuristics: fast, frugal rules that deliberately ignore large amounts of information yet succeed by aligning with the broad patterns of their environments. Herbert Simon's classic insight that decision makers satisfice rather than maximise was the first systematic recognition that scarcity of time, attention and computational power makes selective ignorance a necessity.[131] Subsequent experimental and formal work by Gigerenzer and colleagues has shown that these seemingly crude strategies often equal or surpass complex algorithms.[132] Rather than aiming for unattainable perfection, heuristics define a practicable middle ground: they substitute procedural economy for exhaustive calculation, robustness for fragile precision, and adaptive fit for abstract coherence.[133]

Seen through this ecological lens, the long catalogue of biases documented by behavioural research is reconceived not as a series of errors but as evidence of the mind's adaptive toolbox. Each tool is specialised for a recurring class of problems and shaped by evolutionary, developmental, and cultural feedback.[134] In domains where uncertainties are radical and data sparse, the best strategy is often to know what to ignore. The recognition heuristic, the 1/N portfolio rule and the take the best heuristic exemplify how purposeful ignorance can yield outcomes that are good enough in practice and sometimes quantitatively superior to optimisation that underestimates model error. What emerges is a positive programme for the study of rationality: evaluate decision rules by their performance in context, not by their conformance to a context free ideal. This shift has already influenced medicine, legal reasoning and finance, where simpler, transparency friendly heuristics now routinely match or outperform statistical black boxes.[135]

---

[131] Herbert A. Simon, *"Rational Choice and the Structure of the Environment."*

[132] Gerd Gigerenzer, Peter M. Todd, and the ABC Research Group, "*Simple Heuristics That Make Us Smart"* (Oxford: Oxford University Press, 1999).

[133] Lars P. Hansen and Thomas J. Sargent, "*Robustness."* (Princeton: Princeton University Press, 2008), 25–49.

[134] Gerd Gigerenzer, *"Why Heuristics Work."*

[135] Martine M. Isler, Nikolaus Becker, and Gerd Gigerenzer, *"Simple Heuristics in a Complex World: Health Care Decisions." Medical Decision Making* 41, (2021).

Only after we grasp this broader cognitive architecture does it make sense to ask how entire economies, vast ensembles of boundedly rational actors, can move toward resilience and prosperity. The final section of this thesis applies the heuristic perspective to macroeconomic policy, showing how a deliberately simple combination of Schumpeterian innovation incentives and Keynesian demand stabilisers outperforms single objective optimisation in volatile evolutionary environments. By treating that combination of policies as an institutional framework, we illustrate concretely how the principles defended here scale from individual decision making to the broader design of economic systems.

Chapter 4: The Role of Heuristics in Macroeconomics

4.1. Introduction

Heuristics play a quiet but significant role in economic reasoning and policymaking. Everyday economic behaviors follow simple rules: consumers might allocate a fixed percentage of income to savings, firms might set prices by applying a standard markup and investors might follow basic diversification rules (such as the 1/N portfolio heuristic). These practices persist because they are adaptively rational: they work well enough in an uncertain changing world. Indeed, evidence shows that many firms predominantly use cost plus pricing heuristics, setting prices by adding a markup to production cost, rather than continuously resolving profit maximization calculus.[136] Such rules of thumb are not mere anomalies; they reflect the fact that economic agents, constrained by limited information and computational capacity, develop robust strategies that satisficingly achieve their goals.[137] Recent contributions in behavioral and computational economics have begun to formally study these heuristic strategies.[138] Rather than assuming an impossible level of rationality, newer approaches incorporate boundedly rational agents who use simple decision rules. This chapter explores the macroeconomic implications of these ideas. It argues that heuristics are not only descriptively realistic, but often normatively desirable in a complex economy. In what follows, we first examine the economy as a complex evolving system in which agents' bounded rationality is a feature, not a bug. We then critique the standard neoclassical macro assumptions through the lens of ecological rationality, emphasizing how real world decision making relies on heuristics. Next, we contend that in such a complex environment, satisficing heuristics are more appropriate than any single rule optimization strategy. Building on this foundation, the chapter makes the case that a combination of Schumpeterian (innovation driven) and Keynesian (demand management) policies constitutes a robust satisficing macroeconomic strategy, one that may not be optimal in narrow theoretical terms, but which performs well across a range of scenarios. We support this argument with mathematical and empirical evidence, especially drawing on the agent based model of Dosi, Fagiolo, and Roventini (2010) *"Schumpeter meeting Keynes."*[139] Finally, we

---

[136] Alan S. Blinder et al., "*Asking About Prices: A New Approach to Understanding Price Stickiness"* (New York: Russell Sage Foundation, 1998).
[137] Gerd Gigerenzer and Reinhard Selten, eds., "*Bounded Rationality: The Adaptive Toolbox"* (Cambridge, MA: MIT Press, 2002).
[138] Herbert A. Simon, *"Rational Choice and the Structure of the Environment." Psychological Review* 63, no. 2 (1956): 129–138
[139] Giovanni Dosi, Giorgio Fagiolo, and Andrea Roventini, *"Schumpeter Meeting Keynes: A Policy-Friendly Agent-Based Model of Schumpeterian Growth and Fluctuations." Journal of Economic Dynamics and Control* 34, no. 9 (2010).

conclude by proposing an ecologically rational approach to macroeconomic policy design that values robustness and adaptability.

4.2. The Economy as a Complex Evolving System

Modern complexity economics views the macroeconomy as a dynamic complex system, continually evolving and often volatile. The aggregate patterns we observe (growth, cycles, crises) are emergent phenomena arising from these interactions.[140] Crucially, the economy is subject to fundamental uncertainty in the Knightian sense: not all future states of the world or their probabilities are known.[141] It is also driven by incessant innovation, which continually generates novelty.[142] Disequilibrium is the norm in an innovative economy; the system evolves as new technologies and behaviors disrupt old patterns.[143] Because of these factors, the economy does not gravitate toward a single static equilibrium. Instead, it may exhibit multiple possible equilibria or regimes, path dependence and sudden phase transitions (for example, a financial crisis as a turning point). Tiny perturbations or different initial conditions can lead to different outcomes in the long run.[144] In such a world, precise prediction and optimization are difficult.

In complex evolving systems the assumption of unbounded rationality, that agents can compute optimal decisions with full knowledge of the future, no longer holds. Agents face pervasive uncertainty and limitations in information and processing. They are boundedly rational, meaning they must cope with complexity using imperfect mental tools.[145] Instead of solving impossible optimization problems, real economic actors interpret their environment and adapt. They form expectations by extrapolating from past trends or analogies, they experiment and learn over time, and they employ heuristics as cognitive shortcuts. Indeed, to navigate genuine uncertainty, agents make sense of problems by guessing, using past knowledge and experience and by using simple decision making heuristics or rules of thumb.[146] By relying on a heuristic that works in their environment, agents can make decisions with reasonable success without exhaustive calculation. These heuristics might be as simple as basing this year's plans on last

---

[140] Alan Kirman, *"The Economy as an Interactive System."* in *The Economy as an Evolving Complex System II*, eds. W. B. Arthur, S. N. Durlauf, and D. Lane (Reading, MA: Addison-Wesley, 1997), 491–531.

[141] Frank H. Knight, "*Risk, Uncertainty and Profit*".

[142] Joseph A. Schumpeter, *"Capitalism, Socialism and Democracy"* (New York: Harper, 1942).

[143] W. Brian Arthur, *"Complexity and the Economy."*

[144] John Foster, "*From Simplistic to Complex Systems in Economics."* Cambridge Journal of Economics 29, no. 6 (2005): 873–892.

[145] Herbert A. Simon, *"A Behavioral Model of Rational Choice."* Quarterly Journal of Economics 69, no. 1 (1955).

[146] Gerd Gigerenzer, *"Homo Heuristicus: Why Biased Minds Make Better Inferences."* Topics in Cognitive Science 1, no. 1 (2009).

year's outcomes plus an adjustment, a common strategy for firms forecasting demand or individuals managing budgets. As agents use feedback to update their rules, their behavior evolves.[147] The economy thus has an adaptive character: individuals and firms learn from experience, revise their heuristics and occasionally imitate successful strategies of others. This adaptive evolutionary dynamic means the system is never static. It also underscores why a diversity of behaviors (heterogeneity) persists: there is no single objectively optimal strategy in an environment that itself keeps changing. Instead, different heuristics may perform well in different periods.[148] The overall macro dynamics (growth, volatility, etc.) emerge from the collective interaction of these heterogeneous, learning agents. Significantly, this complexity perspective validates the use of heuristics as not only necessary (given human limits) but often appropriate responses to an unpredictable economic context. Agents survive by being ecologically rational, by matching their decision rules to the structure of the environment they face. The economy's complexity, volatility and openness thus go hand in hand with bounded rationality and heuristic decision making by its participants.

[147] Cars Hommes, *"Heterogeneous Agent Models in Economics and Finance."* in Handbook of Computational Economic*s*, eds. K. Judd and L. Tesfatsion, Vol. 2 (Amsterdam: Elsevier, 2006).
[148] Giovanni Dosi, Giorgio Fagiolo, and Andrea Roventini, *"Schumpeter Meeting Keynes."*

4.3. A Critique to Neoclassical Macro Assumptions

Neoclassical macroeconomics rests on a foundation of assumptions that appear increasingly restrictive in light of real world complexity. The representative agent in the dynamic stochastic general equilibrium (DSGE) model is typically depicted as an infinitely rational optimizer, solving intertemporal utility maximization with full information and forming rational expectations. From an ecological rationality standpoint, this portrayal is deeply problematic. Rational choice theory, as conventionally applied, demands context independent adherence to axioms of optimization and consistency.[149] But in the messy reality of the macroeconomy, context is everything. What is rational in a formal model might be practically unattainable for real households who face shocks, incomplete markets and limited foresight. Ecological rationality argues that true rationality lies in fitness to the environment, the effectiveness of a decision rule given the situation.[150] By that measure, the neoclassical agent's behavior is often ill suited for an uncertain world, while simpler heuristic behaviors may perform more succesfully. For example, the common heuristic of setting aside a fixed percentage of income as savings each month might not solve an intertemporal optimization problem, but it yields prudent outcomes in an environment where future earnings are unpredictable. Likewise, firms sticking to a target inventory to sales ratio or a markup pricing rule are using historically learned heuristics that keep operations viable without solving a complex constrained optimization every period. These heuristics are ecologically rational if they exploit stable features of the environment (like habitual customer behavior or cost structures) to make good decisions.[151]

A key critique is that neoclassical macro models ignore the cognitive constraints and adaptive behavior of actual economic agents. The assumption of perfect rationality and its unique equilibrium often imposes a false sense of precision and determinacy on open ended processes. In reality, agents do not have the knowledge to form truly rational expectations about the future path of the economy. Instead, they fall back on adaptive expectations or other forecasting heuristics which can lead to systematic errors or waves of optimism and pessimism (Keynes's

[149] Gilboa, Itzhak, Andrew Postlewaite, and David Schmeidler. *"Rationality of Belief or: Why Savage's Axioms Are Neither Necessary Nor Sufficient for Rationality."* Synthese 187, no. 1 (2012).
[150] Todd, Peter M., Gerd Gigerenzer, and the ABC Research Group. "*Ecological Rationality: Intelligence in the World*." (New York: Oxford University Press, 2012).
[151] Blinder, Alan S., et al. "*Asking About Prices: A New Approach to Understanding Price Stickiness*." (New York: Russell Sage Foundation, 1998).

"animal spirits").[152] Traditional models treat such behavior as exogenous noise or error, rather than a natural outcome of decision making under uncertainty. By contrast, an ecological rationality perspective recognizes these behaviors as meaningful adaptations. They are what one ought to do, not in an idealized world of certainty, but in the actual environment people face. As Gigerenzer and colleagues point out, more is not always better in decision making; more information and computation can lead to poorer outcomes in high uncertainty contexts.[153] Using a heuristic that ignores certain information can actually improve decision outcomes if that information is unreliable or confounding.[154] This insight directly challenges the neoclassical presumption that any departure from full information optimization is inherently problematic.

Another target of critique is the optimization paradigm of neoclassical macroeconomics. In mainstream models, all behavior is derived from a single utility or profit maximizing principle applied uniformly across contexts. This assumes away the diverse heuristics and norms that actual people use. It also assumes agents somehow solve extremely complex problems (like forming expectations over infinite future contingencies) that are well beyond human (or even computer) capacities. The ecological view instead finds rationality in procedures, in the process of reasoning that agents use, not just in outcomes relative to a model.[155] An agent who uses a simple rule that usually yields decent results is procedurally rational, even if an economist's model shows a theoretically superior result was available. The neoclassical paradigm's focus on outcome optimality (often defined narrowly, like maximizing expected utility) misses this procedural adaptiveness. Bounded rationality complements classical rationality by addressing the discrepancy between the assumed perfect rationality of human behavior and the reality of human cognition.[156] In practical terms, boundedly rational agents satisfice and use heuristics because the world is complex and their minds are finite. Ignoring this, leads mainstream macro to fragile conclusions. Indeed, prior to the 2008 financial crisis, DSGE models built on these

---

[152] George A. Akerlof and Robert J. Shiller. "*Animal Spirits: How Human Psychology Drives the Economy, and Why It Matters for Global Capitalism*." (Princeton: Princeton University Press, 2009).
[153] Gigerenzer, Gerd. *"Homo Heuristicus: Why Biased Minds Make Better Inferences."* Topics in Cognitive Science 1, no. 1 (2009)
[154] Gigerenzer, Gerd, and Daniel G. Goldstein. *"Reasoning the Fast and Frugal Way: Models of Bounded Rationality."* Psychological Review 103, no. 4 (1996).
[155] Simon, Herbert A. *"Invariants of Human Behavior."* Annual Review of Psychology 41, no. 1 (1990).
[156] Simon, Herbert A. "*A Behavioral Model of Rational Choice."* Quarterly Journal of Economics 69, no. 1 (1955).

assumptions largely failed to anticipate the collapse or to provide useful guidance, partly because their agents, by assumption, could not even contemplate such out of equilibrium events.[157] Complexity oriented economists note that mainstream equilibrium models offered little policy insight during that crisis. In response, there has been growing interest in alternative approaches that dispense with the usual assumptions of individual optimization and systemic equilibrium. These approaches, including agent based modeling and behavioral macroeconomics, explicitly incorporate heuristics, heterogeneity and out of equilibrium dynamics, providing a richer framework to analyze macroeconomic phenomena.[158]

Empirically, the case against the neoclassical assumptions is supported by observations of actual behavior. Surveys and studies find, for example, that many firms use simple pricing heuristics (like markup pricing or sticky nominal price rules) instead of continuously recalculating profit maximizing prices.[159] In labor markets, wage setting often follows norms and fairness considerations (rules of thumb) rather than clearing the market via auction each period. Households often follow budgeting heuristics or rely on rules like paying oneself first (automatically saving a fixed amount). Financial market actors, notoriously, use heuristic devices (from credit rating shortcuts to trend following) that do not fit the fully rational model but seem indispensable given the complexity of assessing each investment. These patterns underscore that the representative agent with perfect optimization is a fiction. By insisting on it, neoclassical macro misses how real economies operate and why they sometimes malfunction (feedback loops of heuristic driven behavior can generate instability, like speculative bubbles). The ecological rationality critique therefore calls for a more realistic foundation: one acknowledging that agents use heuristics that are often adaptive to their environment, even if they violate abstract rational choice axioms.[160] It is a shift from viewing such behavior as errors to seeing it as an essential feature of a complex economy.

[157] Caballero, Ricardo J. *"Macroeconomics after the Crisis: Time to Deal with the Pretense-of-Knowledge Syndrome."* Journal of Economic Perspectives 24, no. 4 (2010).
[158]Dosi, Fagiolo, and Roventini, "*Schumpeter Meeting Keynes."*
[159] Fabiani, Silvia, et al. "*Pricing Decisions in the Euro Area: How Firms Set Prices and Why."* Oxford Economic Papers 58, no. 3 (2006).
[160] Gigerenzer, Peter M. Todd, and the ABC Research Group. "*Simple Heuristics That Make Us Smart*."

4.4. Satisficing Heuristics vs Single Rule Optimization

Herbert Simon's notion of satisficing provides a crucial bridge between individual heuristics and broader economic outcomes. To satisfice is to aim for an outcome that is adequate on multiple fronts rather than maximize any single objective in a isolation. In a complex economy characterized by multiple goals and trade offs a satisficing approach can be more appropriate than strict optimization. Optimization typically requires reducing a problem to one objective function subject to constraints, yielding one best policy or decision rule. But complex systems defy such simplification, the real economy has many moving parts and actors with different goals. A single rule optimization strategy (for example, a central bank targeting only inflation with a fixed rule or a government single mindedly maximizing GDP growth) risks overlooking important dimensions and can produce fragility. By contrast, satisficing heuristics acknowledge the need to balance and fulfill multiple criteria well enough without assuming a precise trade off can be optimized. This approach is more robust to uncertainty. Rather than betting everything on one model's definition of optimality, satisficing means securing acceptable outcomes under a range of plausible scenarios.[161]

Consider how individuals make complex decisions, like career choices or retirement planning: they rarely perform an exhaustive optimization. Instead, they set aspiration levels (desired salary, location, work life balance) and search until they find an option that meets these criteria.[162]This multi dimensional satisficing is a heuristic approach suited to complex realities. Similarly, firms often set satisfactory thresholds for profits or market share and pursue strategies (marketing, innovation, etc.) aimed at meeting those targets, rather than mathematically maximizing profit each quarter (which is not even computationally tractable given uncertainties). In macroeconomic policy, the analog would be pursuing strategies that ensure key indicators (employment, inflation, growth) remain in acceptable ranges, rather than optimizing one at the expense of others. Satisficing does not mean settling for mediocrity; it means being prudent about the limits of knowledge and the dangers of pushing any single goal to an extreme. In fact, aiming for the theoretical optimum can be risky if the underlying model is wrong or incomplete, a concept known in decision theory as the risk of model uncertainty.

Mathematically, the advantage of heuristics in complex environments can be understood through concepts like the bias/variance tradeoff in predictive modeling. Highly complex optimization strategies may overfit to a specific environment (yielding low bias but high

[161] Herbert A. Simon, "*Models of Bounded Rationality.*" (Cambridge, MA: MIT Press, 1982).
[162] Gerd Gigerenzer and Reinhard Selten, eds., "*Bounded Rationality: The Adaptive Toolbox*."

variance), performing brilliantly under ideal conditions but failing badly when conditions change.[163] A simpler heuristic, while suboptimal in the idealized case, may generalize better across different environments, it has a bit more bias but far less sensitivity to changing circumstances. In other words, heuristics can offer robustness.[164] This is mirrored in macroeconomic contexts: a policy rule or private decision rule that is modest and adaptive may yield more stable outcomes across booms and busts than an aggressive optimization that assumes steady conditions. For instance, a firm that adopts a conservative debt to equity ratio (a heuristic financial policy) may generate lower profits in favorable times due to restrained leverage. However this conservatism enhances its resilience in downturns, helping it avoid bankruptcy, a satisficing approach to survival. At the macro level, a government might maintain buffers (like higher bank capital requirements or fiscal reserves) based on simple heuristics of prudence, sacrificing a bit of growth in the short run but preventing crises, again a satisficing strategy.[165]

The case for satisficing heuristics over single rule optimization is also supported by computational economics experiments. Agent based models, which explicitly simulate heterogeneous agents using simple rules, find that making agents more rational in the classical sense often yields little improvement and can even destabilize outcomes.[166] Dosi et al. (2010) report that in their agent based macro model, increasing firms' computational sophistication (having them use more complex expectation formation rules) did not significantly change average growth rates or the stability of the economy. In their model, firms initially use a naive heuristic, they simply expect demand next period to equal demand in the last period (a basic adaptive expectation). When the authors allowed firms to employ more advanced forecasting or optimization, the macro level results (growth and stability) remained essentially unchanged. This suggests that simple rules were already sufficient to coordinate outcomes; the added complexity did not buy much, and the system's emergent properties were robust to agents' behavioral rules.[167] In fact, the persistence of outcomes hints that the macro dynamics are driven by deeper structural interactions, not by the micro level optimality of agents' decisions.

[163] David J. C. MacKay, "*Information Theory, Inference and Learning Algorithms"* (Cambridge: Cambridge University Press, 2003).
[164] Gerd Gigerenzer, Peter M. Todd and the ABC Research Group, "*Simple Heuristics That Make Us Smart* ".
[165] Nassim Nicholas Taleb, "*The Black Swan: The Impact of the Highly Improbable."*
[166] Cars Hommes, *"The Heterogeneous Expectations Hypothesis: Some Evidence from the Lab."* Journal of Economic Dynamics and Control 35, no. 1 (2011).
[167] Dosi, Fagiolo, and Roventini, *"Schumpeter Meeting Keynes."*

Such findings reinforce the idea that forcing agents to optimize single objectives yields diminishing returns in a complex adaptive system. The economy may self organize in such a way that simple agent heuristics are good enough, or even exactly what is needed, for the system to function.[168] This aligns with Gigerenzer's argument that in an uncertain world, less can be more, and a heuristic matched well to an environment can outperform complex strategies.[169] In sum, satisficing heuristics offer a philosophy of decision making and policy design that prioritizes robust acceptability over fragile optimality.

[168] Leigh Tesfatsion, "*Agent-Based Computational Economics: Modeling Economies as Complex Adaptive Systems." Information Sciences* 149, no. 4 (2003).
[169] Gerd Gigerenzer, *"Homo Heuristicus: Why Biased Minds Make Better Inferences."*

4.5. Combining Schumpeterian and Keynesian Policies

What might a satisficing macroeconomic policy strategy look like? This chapter argues that it resembles a pragmatic combination of Schumpeterian and Keynesian elements, supporting innovation and long run growth on one hand, while stabilizing demand and employment on the other. These two pillars correspond to addressing the supply side and demand side of the economy, respectively. In orthodox theory, one might attempt to derive an optimal policy by, say, maximizing a social welfare function subject to a macroeconomic model's constraints. Such an optimal plan could, for example, prioritize price stability above all or assume that free markets alone will optimally allocate resources for growth. However, in the real economy a policy that satisfices multiple objectives is more robust. Schumpeterian policies refer to measures that promote innovation, technological progress and structural transformation. These include investments in R&D, education and infrastructure, support for entrepreneurship and industrial policies that foster new industries. Such policies drive the economy's evolutionary growth engine, increasing productivity and the frontier of potential output. Keynesian policies refer to demand management tools, such as countercyclical fiscal spending, monetary policy aimed at full employment and automatic stabilizers like unemployment insurance. These policies seek to maintain adequate aggregate demand, mitigate recessions, and avoid prolonged unemployment or excess capacity.[170]

Each of these policy sets addresses a crucial dimension of economic performance: the Schumpeterian aims for high long run growth, while the Keynesian aims for stable short run fluctuations and high employment. In theory, one could try to focus on just one dimension, for example, a singular focus on innovation led growth, assuming the benefits will trickle down eventually; or an exclusive focus on short run stabilization, assuming the long run will take care of itself. However, an ecologically rational view suggests that these objectives are closely connected and that the most robust strategy satisfies both to a reasonable degree rather than maximizing one at the expense of the other.[171] The innovation engine cannot reach its potential if demand is chronically weak, new technologies will not find markets, firms won't invest in R&D if recessions constantly undercut profitability and human capital will deteriorate under prolonged high unemployment. Conversely, stimulating demand without attention to productivity and innovation can lead to overheating or stagnation once the limits of existing

---

[170] Giovanni Dosi, Mauro Napoletano, Andrea Roventini, and Tania Treibich, *"Micro and Macro Policies in an Agent-Based Keynesian Model."* Journal of Economic Behavior & Organization 155 (2018).

[171] Herbert A. Simon, *"A Behavioral Model of Rational Choice."*

capacity are reached. The complementarity between Schumpeterian and Keynesian policies means that doing both yields more than the sum of its parts. Simulations of an economy where these policies work in tandem show long run growth paths characterized by mild fluctuations and acceptable unemployment levels.[172]

Critically, this combined strategy is satisficing rather than optimizing in the sense that it does not seek a theoretical optimum in one dimension. It accepts, for instance, that aggressive innovation policy might need to be tempered by demand concerns, for example by allocating some resources to social safety nets, which may limit short run productivity gains. Conversely, stabilizing the economy might involve tolerating some inefficiencies, such as not fully allowing market forces to eliminate unproductive firms during downturns, because doing so can have long run costs on innovation ecosystems.The goal is to achieve robust prosperity, reasonably high growth with manageable volatility and unemployment, even if neither growth nor stability is pushed to an extreme theoretical maximum. This approach draws inspiration from the real world observation that economies which successfully industrialized and grew rapidly (the East Asian miracles, postwar Western economies) often combined forward looking development policies with active demand management.[173] For example, post World War II America and Europe saw governments investing heavily in science and infrastructure (Schumpeterian) while also institutionalizing Keynesian demand tools (such as built in stabilizers and countercyclical fiscal policy), leading to a period of unprecedented stable growth. Conversely, periods of policy that overly emphasized one side often resulted in either stagnant growth or unstable boom bust cycles.[174]

The agent based computational evidence strongly supports the synergy of these policies. In Dosi, Fagiolo, and Roventini's model of a "*Schumpeter meeting Keynes*" economy, the authors find that Keynesian demand policies are in fact a necessary condition for sustained long run growth in a technologically evolving economy. When the model is run with only Schumpeterian (innovation) forces active but inadequate demand support, the economy can get trapped in a low growth, high unemployment regime. Innovations occur, but their benefits do not translate into output and employment because demand falls short, a result mirroring secular stagnation concerns. Conversely, with robust Keynesian demand management, the economy is

---

172 Dosi, Fagiolo, Roventini, *"Schumpeter Meeting Keynes".*

173 Alice Amsden, "*Asia's Next Giant: South Korea and Late Industrialization."* (Oxford: Oxford University Press, 1989).

174 Barry Eichengreen, "*Globalizing Capital: A History of the International Monetary System."* (Princeton: Princeton University Press, 2008).

delocked from the bad trajectory and can attain a high growth path. The presence of strong aggregate demand allows innovative firms to expand and encourages continual investment, reinforcing growth. Importantly, the model also shows that Keynesian policies dramatically reduce macroeconomic volatility and unemployment.[175] Countercyclical fiscal measures act like an economic parachute during recessions, preventing severe downturns and thereby indirectly sustaining private investment (which might otherwise collapse in a recession). Meanwhile, Schumpeterian policies (like R&D subsidies or higher innovation propensity) also have distinct effects: holding demand policy constant, greater support for innovation tends to raise long term growth potential, but it can come with higher volatility if not counterbalanced.[176] Only when paired with stabilizing demand policies do the fruits of innovation translate into steady growth with mild cycles. The complementary nature of the two policy types means that neither alone is sufficient for the best outcomes: innovation policy alone cannot guarantee low unemployment and smooth cycles, and demand policy alone cannot generate high growth without technological progress. The most resilient macroeconomic performance emerges when both are pursued together.[177]

This combined approach can be seen as a macro level heuristic or rule of thumb: always encourage innovation and always stabilize demand to an extent. It may not satisfy those who seek an elegant one target optimum rule (for instance, a pure inflation targeting central bank or a single minded growth maximization agenda). But it satisfices the multiple goals society cares about, livelihoods, innovations and stability, yielding robustly good outcomes across different scenarios. Even if such a policy mix is not optimal under a particular model's assumptions, it guards against the model being wrong. In a complex economy with uncertainty about correct models, this robustness is immensely valuable. Indeed, Dosi et al.'s model underscores that pushing either policy too far has diminishing returns: once the economy is on a high growth, low volatility path, further increases in fiscal stimulus, for example don't raise growth much more but do slightly, improve stability.[178] There are pragmatic limits and trade offs, so policymakers must use judgment to keep policies in balance, another sign of satisficing behavior.

[175] Dosi, Fagiolo, and Roventini, *"Schumpeter Meeting Keynes."*
[176] Dosi et al., "*Micro and Macro Policies."*
[177] Dosi, Fagiolo, and Roventini, *"Schumpeter Meeting Keynes."*
[178] Dosi, Fagiolo, and Roventini, "*Schumpeter Meeting Keynes."*

4.6. Mathematical and Empirical Support

The arguments above gain firmer ground when we examine formal models and empirical evidence. The agent based model developed by Dosi, Fagiolo, and Roventini (2010)[179] provides a quantitative, mathematical demonstration of how heuristics and policy combinations can shape macro outcomes. In their model, the economy consists of heterogeneous firms in two sectors (capital goods and consumption goods), consumers/workers and a public sector. Crucially, the firms in this model are boundedly rational: they do not solve intertemporal optimization problems. Instead, they follow simple behavioral rules for decisions like production, pricing, investment and R&D. For example, consumption good firms set prices with a fixed markup over costs, a rule grounded in empirical observation of how businesses price. Firms form their expectations of demand in a highly simplified way (one version of the model assumes firms just expect demand to repeat last period's level). Investment decisions are modeled with rules such as: expand production capacity when customer orders exceed current production capability, or invest a certain fraction of profits in R&D each period. Workers/consumers in the model spend their income based on simple consumption rules (with some propensity to consume out of wages and unemployment benefits). There is no omniscient rational planner among these agents, each is following boundedly rational heuristics adapted to its role.

Mathematically, the model is specified by a system of equations describing these behavioral rules and constraints (a production function, an R&D innovation process, accounting identities ensuring consistency). The public sector in the model implements Keynesian policies: it collects taxes and pays unemployment benefits (injecting demand during downturns) and the authors can vary the tax rate or benefit level to simulate more or less aggressive fiscal policy. Schumpeterian policy enters via parameters affecting innovation, such as the fraction of revenues firms devote to R&D or the efficiency of R&D expenditure (these can be interpreted as influenced by policy incentives or technological opportunities). The model is stochastic: innovation success, for instance, is partly random, and firms' productivity evolves stochastically, capturing the uncertainty and heterogeneity of real economies. Rather than solving for a closed form equilibrium, the model is run through computer simulations (Monte Carlo), generating synthetic time series data that can be analyzed.

The findings from this agent based model strongly support the chapter's thesis. First, the model can replicate many stylized facts of macroeconomic data (such as business cycle fluctuations,

179 Dosi, Fagiolo, and Roventini, *"Schumpeter Meeting Keynes."*

firm size distributions, productivity dispersion, etc.) despite its reliance on simple heuristics. This shows that a heuristic driven, out of equilibrium model is capable of producing realistic macro dynamics, a domain once dominated by equilibrium models. Second, the model's policy experiments directly illustrate the interplay of Schumpeterian and Keynesian levers. When the authors simulate a scenario with only innovation drives (high R&D propensity) but minimal fiscal stabilizers, the economy exhibits high output volatility and periods of stagnation or high unemployment. In contrast, a scenario that adds strong Keynesian demand management (higher taxes and benefits that kick in during recessions) moves the economy to a much better trajectory, what they term the good regime. In this regime, output grows at a healthy long run rate and cyclical swings are much gentler, with lower average unemployment.

Notably, the model identifies two distinct regimes of growth depending on policy: one characterized by robust growth with mild cycles and another by either low growth or more severe cycles with high unemployment. These correspond to the presence or absence of the right mix of policies. The matching of high innovation with strong demand support yields the favorable regime, whereas a mismatch (innovation without demand, or demand stimulus without innovation) yields subpar outcomes. This aligns perfectly with our earlier qualitative argument about complementarity.

Quantitatively, the simulations showed that Keynesian policies not only reduce short run volatility and unemployment, but are indeed necessary for maintaining long run growth. Without them, even an innovation rich economy can falter with demand deficiency. On the flip side, the presence of innovation (Schumpeterian forces) is what allows the demand stabilized economy to grow rather than just achieving full employment at static output. Another important quantitative result is that beyond a certain point, increasing the intensity of Keynesian policy (making fiscal stimulus larger or more automatic) yields diminishing returns in terms of growth. After the economy has been lifted to the high growth path, extra stimulus doesn't raise the growth rate further, it does, however, further reduce output volatility and time spent in unemployment, effectively enhancing the stability of the system. This indicates an optimal intermediate range for policy, too little and the economy is unstable, too much and you get stability but no additional growth. Such nonlinear responses are typical in complex systems and warn against extreme policies, reinforcing the idea of satisficing balance. The authors conclude that there is a clear complementarity between Keynesian and Schumpeterian policies in sustaining long run growth paths characterized by mild fluctuations and acceptable unemployment.

Empirical evidence outside of models also supports these claims. Historical data studies (cited by Dosi et al. as well) have found that economies with more aggressive countercyclical policies tend to experience not just shorter recessions but also faster long run growth recoveries, especially in environments with credit constraints or other market imperfections. One empirical study by Aghion and Marinescu (2007) showed that countries with stronger stabilizers had higher average growth, suggesting that preventing deep recessions helps preserve the drivers of growth, a real world validation of the model's insight.[180] Meanwhile, the importance of innovation policy is evident in the divergent productivity and growth trajectories of countries or regions that invest in R&D and human capital versus those that do not. The synergy is perhaps most clearly visible in extraordinary cases like the wartime and post war U.S., where massive demand stimulus (WWII spending, then Cold War investments) coincided with rapid technological innovation (much of it spurred by war and government R&D), yielding the long post war boom with high growth and low unemployment. While such historical episodes are complex, they exemplify how demand and innovation factors working together create robust prosperity.

[180] Philippe Aghion and Ioana Marinescu, *"Cyclical Budgetary Policy and Economic Growth: What Do We Learn from OECD Panel Data?"* NBER Macroeconomics Annual 22 (2007).

4.7. Conclusion on Heuristics in Macroeconomics

The exploration of heuristics in macroeconomics leads to a profound rethink of what rational policy design means in an uncertain complex world. Ecological rationality offers a guiding philosophy: it advises economists and policymakers to design rules and institutions that are well matched to the structure of the economic environment, rather than assuming an idealized world of rational agents and predictable dynamics. An ecologically rational macroeconomic policy would be one that remains effective under the true conditions of the economy, conditions that include fundamental uncertainty, innovation driven change and human cognitive limitations. The chapter has argued that heuristics, once viewed with suspicion, are in fact essential building blocks for such policy. Just as individual agents use heuristics to make reasonable decisions, policymakers can adopt simple robust rules that secure broadly satisfactory outcomes across many scenarios. The recommended combination of Schumpeterian and Keynesian policies represents precisely this type of rule. A straightforward principle of encouraging innovation and maintaining demand can serve as a reliable formula for long run growth with stability, without requiring the solution of an intractable dynamic optimization problem for the entire economy.[181]

In contrast to the neoclassical tradition of seeking an optimal policy (often characterized by complex rules or targets that assume away uncertainty), the ecological rationality approach emphasizes robustness and adaptability. Policymakers should acknowledge what they don't know and cannot know. Rather than placing all faith in a single model's recommendation (which might fail if the model is wrong), it is more reasonable to pursue policies that are justifiable across a range of models and experiences, policies that satisfice multiple goals and guard against extreme outcomes.[182] This perspective resonates with the concept of minimizing regret or pursuing robust control in macroeconomic policy, techniques that explicitly account for model uncertainty and aim for policies that do reasonably well in many potential situations instead of optimizing one forecast.[183] Ecologically rational policy design also values institutional heuristics that have stood the test of time. For instance, automatic stabilizers (like progressive taxes and welfare programs) can be seen as heuristic policies built into fiscal systems. They don't require real time optimization, they simply operate by rule to stabilize demand when income falls and decades of experience show their effectiveness in mitigating

[181]Gigerenzer, "*Simple Heuristics That Make Us Smart*."
[182] Gigerenzer, *"Why Heuristics Work."*
[183] Lars Peter Hansen and Thomas J. Sargent, *"Robust Control and Model Uncertainty." American Economic Review* 91, no. 2 (2001).

the negative effects of downturns.[184] Similarly, innovation policy often operates through simple mechanisms like grants for research, patent systems and education funding. These may not be optimized each year, but as heuristics they continuously foster the ecosystem needed for innovation.[185]

By embracing ecological rationality, macroeconomics can become more relevant and realistic. It moves away from the sterile elegance of fully optimal but fragile policies, toward robustly good policies that acknowledge complexity. The conclusions drawn from our discussion and the Dosi et al. model exemplify this shift.[186] They suggest that the best way to ensure a healthy economy is not to adjust a single lever (like inflation) with unfailing precision. However, to put in place a framework of heuristics that will guide the economy toward desirable outcomes even when surprises happen. This does not mean there is no role for optimization or sophisticated analysis; rather, it means those analyses should be employed in service of designing and improving heuristics and institutions, not as a replacement for them. In practice, an ecologically rational macroeconomic strategy might involve a policymaking process that is experimental and adaptive. This means trying combinations of policies, monitoring outcomes and adjusting heuristics over time based on feedback, much as agents in the economy adapt their rules.[187]

In conclusion, heuristics are not only a cornerstone of individual decision making in economics, but also a cornerstone of wise macroeconomic policy in a complex world. Recognizing the economy as an evolving open system of boundedly rational agents compels us to abandon the mirage of perfect optimization. In its place, we adopt ecological rationality, designing strategies that are simple, sensible and matched to our environment. The marriage of Schumpeter and Keynes in policy is a prime example: it rejects any notion of a perfect, one dimensional solution and instead relies on a heuristic that combines two essential elements to secure a satisficing outcome. Such an approach may not satisfy those who seek the certainty of optimality, but it offers something more valuable in the real world: resilience. As we face future economic challenges, whether technological disruptions, financial cycles or unforeseen shocks, the lesson

---

[184] Alan J. Auerbach and Daniel Feenberg, *"The Significance of Federal Taxes as Automatic Stabilizers" Journal of Economic Perspectives* 14, no. 3 (2000).
[185] Philippe Aghion, Ufuk Akcigit, and Peter Howitt, *"What Do We Learn From Schumpeterian Growth Theory?" Journal of Economic Perspectives* 30, no. 1 (2016).
[186]Dosi, Fagiolo, and Roventini, *"Schumpeter Meeting Keynes."*
[187] Simon, *"Theories of Decision-Making in Economics and Behavioral Science."*

is that a heuristic, adaptive approach grounded in ecological rationality will guide us to more robust and humane outcomes than rigid and outdated attempts at theoretical optimization.[188]

[188] Gerd Gigerenzer, *"Heuristics."* in Heuristics and the Law, ed. Gerd Gigerenzer and Christoph Engel (Cambridge, MA: MIT Press, 2006).

Conclusions

The core insight of this thesis is simple yet far reaching: in an economic world shaped by innovation, complexity and radical uncertainty, the future is not merely unknown, it is unknowable in principle. No amount of data, modelling or analytical power can fully anticipate the emergence of genuinely novel states, technologies or behaviours. This undermines the foundational assumptions of classical optimisation frameworks, which rely on closed, well-defined sets of outcomes and stable preferences. In such conditions, even heuristics, those practical, experience based rules designed to cope with limited information, cannot guarantee success. They may simplify choices and offer robustness in some cases, but they too are fallible. When applied without regard to context, heuristics can mislead, distort judgement, or fail to adapt to changing conditions.

What this thesis offers, then, is not a defence of heuristics as universal solutions, but a broader invitation to rethink how we approach decision making. The chapters have shown that alternatives to dominant optimisation models exist, alternatives that are grounded in ecological rationality, evolutionary learning and the structural properties of the environment itself. Rather than seeking abstract optimality, these approaches prioritise adaptability, feedback and context sensitivity.

The key lesson is that no decision can be judged in isolation from its environment. To act wisely under uncertainty requires first recognising the structure and dynamics of the system in which one is embedded. This means designing strategies, rules and institutions that are not just efficient under known conditions but resilient when those conditions change. By focusing on the interplay between cognition and environment, between rules and contexts, this thesis encourages a shift from rigid precision to informed flexibility.

In a world where surprises are inevitable and prediction has limits, success lies less in mastering the future than in learning how to navigate it. Acknowledging this is not a call to surrender, but an argument for a more grounded, adaptive and realistic economics; one that begins not with perfect foresight, but with an honest reading of the world as it is.

Bibliography

Aghion, Philippe, and Ioana Marinescu. *"Cyclical Budgetary Policy and Economic Growth: What Do We Learn from OECD Panel Data?"* NBER Macroeconomics Annual 22 (2007): 251–278.

Aghion, Philippe, Ufuk Akcigit, and Peter Howitt. *"What Do We Learn From Schumpeterian Growth Theory?"* Journal of Economic Perspectives 30, no. 1 (2016): 3–28.

Akerlof, George A., and Robert J. Shiller. *"Animal Spirits: How Human Psychology Drives the Economy, and Why It Matters for Global Capitalism."* Princeton: Princeton University Press, 2009.

Alter, Adam, and Daniel M. Oppenheimer. *"Predicting Short-Term Stock Fluctuations by Using Processing Fluency."* Proceedings of the National Academy of Sciences 103, no. 24 (2006): 9369–9372.

Amsden, Alice. "*Asia's Next Giant: South Korea and Late Industrialization."* Oxford: Oxford University Press, 1989.

Anderson, P. W., Kenneth J. Arrow, and David Pines, eds. "*The Economy as an Evolving Complex System*". Reading, MA: Addison-Wesley, 1988.

Arthur, W. Brian. "*Increasing Returns and Path Dependence in the Economy.*" Ann Arbor: University of Michigan Press, 1994.

Arthur, W. Brian. "*Complexity and the Economy.*" Science 284, no. 5411 (1999): 107–109.

Auerbach, Alan J., and Daniel Feenberg. *"The Significance of Federal Taxes as Automatic Stabilizers."* Journal of Economic Perspectives 14, no. 3 (2000): 37–56.

Banks, J., J. Brooks, G. Cairns, G. Davis, and P. Stacey. *"On Devaney's Definition of Chaos."* The American Mathematical Monthly. 99, no. 4 (1992): 332–334.

Beinhocker, Eric D. "*The Origin of Wealth: Evolution, Complexity, and the Radical Remaking of Economics.*" Boston: Harvard Business School Press, 2006.

Björneborn, Lars. *"Adjacent Possible."* In *Springer MRW*, 2023.

Blinder, Alan S., et al. "*Asking About Prices: A New Approach to Understanding Price Stickiness."* New York: Russell Sage Foundation, 1998.

Brock, W. A. "*Distinguishing Random and Deterministic Systems." Journal of Economic Theory.* 40, no. 1 (1986): 168–195.

Brockwell, Peter J., and Richard A. Davis. "*Introduction to Time Series and Forecasting*." 2nd ed. New York: Springer, 2002

Byrne, David S. "*Complexity Theory and the Social Sciences: An Introduction*." London: Routledge, 1998.

Caballero, Ricardo J. *“Macroeconomics after the Crisis: Time to Deal with the Pretense-of-Knowledge Syndrome.”* Journal of Economic Perspectives 24, no. 4 (2010): 85–102.

Chaitin, Gregory J. *“Randomness and Mathematical Proof.”* Scientific American 232, no. 5 (1975): 47–52.

Cortês, Marina, Cristina Juárez, Tom F. Varley, and Stuart Kauffman. *“The TAP Equation: Evaluating Combinatorial Innovation.” arXiv*, 2022.

Devaney, Robert L. *”An Introduction to Chaotic Dynamical Systems.”* 3rd ed. Boca Raton, FL: CRC Press, 2022.

Devereaux, Andrew, Roger Koppl and Stuart Kauffman. *“Creative Evolution in Economics.”* Journal of Evolutionary Economics 34 (2024): 489–514.

Dosi, Giovanni. *“Rational Heuristics: Expectations and Behaviors in Evolving Economies with Heterogeneous Agents.”* Economic Inquiry 58, no. 1 (2020): 1–18.

Dosi, Giovanni, Giorgio Fagiolo, and Andrea Roventini. *“Schumpeter Meeting Keynes: A Policy-Friendly Agent-Based Model of Schumpeterian Growth and Fluctuations.”* Journal of Economic Dynamics and Control 34, no. 9 (2010): 1748–1767.

Dosi, Giovanni, Mauro Napoletano, Andrea Roventini, and Tania Treibich. *“Micro and Macro Policies in an Agent-Based Keynesian Model.”* Journal of Economic Behavior & Organization 155 (2018): 35–52.

Eichengreen, Barry. “*Globalizing Capital: A History of the International Monetary System.”* Princeton: Princeton University Press, 2008.

Fabiani, Silvia, et al. *“Pricing Decisions in the Euro Area: How Firms Set Prices and Why.”* Oxford Economic Papers 58, no. 3 (2006): 449–478.

Foster, John. *“From Simplistic to Complex Systems in Economics.”* Cambridge Journal of Economics 29, no. 6 (2005): 873–892.

Gell-Mann, Murray. “*The Quark and the Jaguar: Adventures in the Simple and the Complex.”* New York: W.H. Freeman, 1994.

Gigerenzer, Gerd. *“Heuristics.”* In Heuristics and the Law, edited by Gerd Gigerenzer and Christoph Engel, 17–44. Cambridge, MA: MIT Press, 2006.

Gigerenzer, Gerd. *“Homo Heuristicus: Why Biased Minds Make Better Inferences.”* Topics in Cognitive Science 1, no. 1 (2009): 107–143.

Gigerenzer, Gerd. *“The Rationality Wars: Why Heuristics Work.”* The Rationality Wars: A Personal Reflection, Max Planck Institute for Human Development, 2021.

Gigerenzer, Gerd. *“Why Heuristics Work.”* Perspectives on Psychological Science 3, no. 1 (2008): 20–29.

Gigerenzer, Gerd, and Daniel Goldstein. *"Reasoning the Fast and Frugal Way: Models of Bounded Rationality."* Psychological Review 103, no. 4 (1996): 650–669.

Gigerenzer, Gerd, and Peter M. Todd. "*Simple Heuristics That Make Us Smart."* New York: Oxford University Press, 1999.

Gigerenzer, Gerd, and Reinhard Selten, eds. "*Bounded Rationality: The Adaptive Toolbox."* Cambridge, MA: MIT Press, 2002.

Gigerenzer, Gerd, Wolfgang Gaissmaier. *"Heuristic Decision Making."* Annual Review of Psychology 62 (2011): 451–82.

Gilboa, Itzhak, Andrew Postlewaite, and David Schmeidler. "*Rationality of Belief or: Why Savage's Axioms Are Neither Necessary Nor Sufficient for Rationality."* Synthese 187, no. 1 (2012): 11–31.

Gini, Corrado. *"Measurement of Inequality of Incomes."* The Economic Journal 31, no. 121 (1921): 124–126.

Gleick, James. "*Chaos: Making a New Science*." New York: Open Road Integrated Media, 2011.

Goldstein, Daniel G., and Gerd Gigerenzer. *"Models of Ecological Rationality: The Recognition Heuristic."* Psychological Review 109, no. 1 (2002): 75–90.

Grüne-Yanoff, Till, and Sven Ove Hansson. "*Modeling Rationality, Rationalizing Models."* Minds and Machines 19, no. 3 (2009): 301–316.

Hanel, Rudolf, Stuart A. Kauffman, and Stefan Thurner. *"Towards a Physics of Evolution: Critical Diversity Dynamics at the Edges of Collapse and Bursts of Diversification."* Biological Theory 2, no. 4 (2007): 1–12.

Hansen, Lars Peter, and Thomas J. Sargent. "*Robustness"*. Princeton: Princeton University Press, 2008.

Hansen, Lars Peter, and Thomas J. Sargent. *"Robust Control and Model Uncertainty." American Economic Review* 91, no. 2 (2001): 60–66.

Harremoës, Peter, and Flemming Topsøe. *"Maximum Entropy Fundamentals." Entropy* 3, no. 3 (2001): 191–226.

Holland, John H. *"Complex Adaptive Systems."* Daedalus 121, no. 1 (1992): 17–30.

Hommes, Cars. "*Heterogeneous Agent Models in Economics and Finance."* In Handbook of Computational Economics, edited by K. Judd and L. Tesfatsion, Vol. 2, 1109–1186. Amsterdam: Elsevier, 2006.

Hommes, Cars. *"The Heterogeneous Expectations Hypothesis: Some Evidence from the Lab."* Journal of Economic Dynamics and Control 35, no. 1 (2011): 1–24.

Isler, Martine M., Nikolaus Becker, and Gerd Gigerenzer. *“Simple Heuristics in a Complex World: Health Care Decisions.”* Medical Decision Making 41, no. 1 (2021): 16–27.

Johnson, Steven. *“Where Good Ideas Come From: The Natural History of Innovation.”* New York: Riverhead Books, 2010.

Kauffman, Stuart A. “*Investigations*.” Oxford: Oxford University Press, 2000.

Kahneman, Daniel, and Amos Tversky. *“Prospect Theory: An Analysis of Decision under Risk.”* Econometrica 47, no. 2 (1979): 263–291.

Kauffman, Stuart A. *“The Origins of Order: Self-Organization and Selection in Evolution.”* New York: Oxford University Press, 1993.

Katsikopoulos, Konstantinos V. *“Psychological Heuristics for Making Inferences: Definition, Performance, and the Emerging Theory and Practice.”* Decision Analysis 8, no. 1 (2011): 10–29.

Kirman, Alan. *“The Economy as an Interactive System.”* In The Economy as an Evolving Complex System *II*, edited by W. B. Arthur, S. N. Durlauf, and D. Lane, 491–531. Reading, MA: Addison-Wesley, 1997.

Knight, Frank H. “*Risk, Uncertainty and Profit.”* Boston: Houghton Mifflin, 1921.

Lambert, Peter. *“The Distribution and Redistribution of Income.”* 3rd ed. Manchester: Manchester University Press, 2001.

Laplace, Pierre-Simon. *“A Philosophical Essay on Probabilities.”* Translated by Frederick Wilson Truscott and Frederick Lincoln Emory. New York: Dover Publications, 1951.

Laplace, Pierre-Simon. *“Essai Philosophique sur les Probabilités.”* Paris: Courcier, 1814.

Levin, Simon A. “*Ecosystems and the Biosphere as Complex Adaptive Systems.”* Ecosystems 1, no. 5 (1998): 431–436.

Loreto, Vittorio, Andrea Baronchelli, Alain Barrat, Luca Dall’Asta, and Alessandro Vespignani. *“Statistical Physics of Social Dynamics.”* Reviews of Modern Physics 81, no. 2 (2009): 591–646.

Loreto, Vittorio, Vito D. P. Servedio, Stefano H. Strogatz, and Francesca Tria. *“Dynamics on Expanding Spaces: Modeling the Emergence of Novelties.”* Nature Physics 12, no. 10 (2016): 842–847.

Lorenz, Edward N. *“Deterministic Nonperiodic Flow.”* Journal of the Atmospheric Sciences 20, no. 2 (1963): 130–141.

MacKay, David J. C. “*Information Theory, Inference and Learning Algorithms.”* Cambridge: Cambridge University Press, 2003.

Mitchell, Melanie. *“Complexity: A Guided Tour”*. Oxford: Oxford University Press, 2009.

Monechi, Bernardo, Francesca Tria, Pierluigi Vellucci, and Vittorio Loreto. "*Waves of Novelties in the Expansion into the Adjacent Possible." PLOS ONE* 12, no. 6 (2017).

Morin, Jean Frédéric, Joost Pauwelyn, and James Hollway. *"The Trade Regime as a Complex Adaptive System: Exploration and Exploitation of Environmental Norms in Trade Agreements."* Journal of International Economic Law 20, no. 2 (2017): 365–390.

Newton, Isaac. *"Philosophiæ Naturalis Principia Mathematica."* London: Royal Society, 1687.

Poincaré, Henri. "*Science and Method."* Translated by Francis Maitland. London: Thomas Nelson and Sons, 1914.

Prigogine, Ilya, and Isabelle Stengers. *"Order out of Chaos: Man's New Dialogue with Nature."* New York: Bantam Books, 1984.

Rickles, Dean, Penelope Hawe, and Alan Shiell. *"A Simple Guide to Chaos and Complexity."* Journal of Epidemiology and Community Health 61, no. 11 (2007): 933–937.

Rubin, Donald B. *"The Calculation of Gini Coefficients."* The Review of Economics and Statistics 45, no. 1 (1963): 50–52.

Schooler, Lael J., and Ralph Hertwig. *"How Forgetting Aids Heuristic Inference."* Psychological Review 112, no. 3 (2005): 610–628.

Schuck, Nicolas W., et al. *"Medial Prefrontal Cortex Predicts Internally Driven Strategy Shifts."* Neuron 86, no. 2 (2015): 331–340.

Schumpeter, Joseph A. "*Capitalism, Socialism and Democracy."* New York: Harper, 1942.

Shannon, Claude E. *"A Mathematical Theory of Communication."* Bell System Technical Journal 27, no. 3 (1948): 379–423.

Simon, Herbert A. *"The Sciences of the Artificial."* 3rd ed. Cambridge, MA: MIT Press, 1996.

Simon, Herbert A. *"A Behavioral Model of Rational Choice."* Quarterly Journal of Economics 69, no. 1 (1955): 99–118.

Simon, Herbert A. *"Invariants of Human Behavior."* Annual Review of Psychology 41, no. 1 (1990): 1–19.

Simon, Herbert A. "*Administrative Behavior."* 4th ed. New York: Free Press, 1997.

Simon, Herbert A. "*Models of Bounded Rationality."* Vol. 1. Cambridge, MA: MIT Press, 1982.

Simon, Herbert A. *"Rational Choice and the Structure of the Environment."* Psychological Review 63, no. 2 (1956): 129–138.

Simon, Herbert A. *"Theories of Decision-Making in Economics and Behavioral Science."* American Economic Review 49, no. 3 (1959): 253–283.

Strogatz, Steven H. *“Nonlinear Dynamics and Chaos: With Applications to Physics, Biology, Chemistry, and Engineering.”* 2nd ed. Boca Raton, FL: CRC Press, 2019.

Taalbi, Joakim. *“Long-Run Patterns in the Discovery of the Adjacent Possible.” arXiv*, 2023.

Taleb, Nassim Nicholas. “*The Black Swan: The Impact of the Highly Improbable.”* New York: Random House, 2007.

Tesfatsion, Leigh. *“Agent-Based Computational Economics: Modeling Economies as Complex Adaptive Systems.”* Information Sciences 149, no. 4 (2003): 262–268.

Tria, Francesca, Vittorio Loreto, Vito D. P. Servedio, and Stefano H. Strogatz. *“The Dynamics of Correlated Novelties.” Scientific Reports* 4, no. 5890 (2014).

Todd, Peter M., Gerd Gigerenzer, and the ABC Research Group. “*Ecological Rationality: Intelligence in the World.”* New York: Oxford University Press, 2012.

Turner, John R., and Rose M. Baker. “*Complexity Theory: An Overview with Potential Applications for the Social Sciences.”* Systems 7, no. 1 (2019): 1–23.

Tversky, Amos, and Daniel Kahneman. *“Availability: A Heuristic for Judging Frequency and Probability.”* Cognitive Psychology 5, no. 2 (1973): 207–232.

Tversky, Amos, and Daniel Kahneman. *“Judgment under Uncertainty: Heuristics and Biases.”* Science 185, no. 4157 (1974): 1124–1131.

Ubaldi, Elisa, Alessio Cardillo, Francesco Tria, Vito D. P. Servedio, and Vittorio Loreto. *“Emergence and Evolution of Social Networks through Exploration of the Adjacent Possible.”* Communications Physics 4, no. 28 (2021).

Waldrop, M. Mitchell. *“Complexity: The Emerging Science at the Edge of Order and Chaos.”* New York: Simon & Schuster, 1992.