Why Inefficiency Survives: The Off-Book Matrix of Rationality

Ryusuke K. Fukumizu
Tokyo Cognitive Lab, Tokyo, Japan

Abstract

Why does "inefficiency" survive? We propose that apparent inefficiency is not a cognitive defect but an artifact of evaluative mismatch. The Off-Book framework identifies two survival requirements unmeasured in laboratory ledgers: biological survival (avoiding ruin) and social survival (avoiding isolation). Our core insight: the subject's condition determines T (trial capacity), and T determines whether expected value becomes a valid reference. Empirical results ($N = 100$; loss aversion $d = 0.99$, sunk cost $d = 0.47$, interaction $p = .004$) confirm that rationality judgments shift when survival risks arise. We argue for a transition from Universal to Relative Rationality—where the scientist's $N = 100$ ensemble and the subject's $N = 1$ intuition both derive from the same principle. Under the same ecology, subjects may hold contradictory rationality due to the attributes they carry in their position. The bias is not in the subject—it is in the ruler.

Keywords: Ergodicity; Non-Fungible Capital; Cognitive Bias; Ecological Rationality; Kahneman-Gigerenzer Debate; Off-Book Capital; Survival Strategy

Introduction

We write down what matters. But what changes us most often stays off-book. Despite decades of research on cognitive biases and the limits of human rationality, these biases remain robust. Simon (1955) showed that rationality operates under cognitive bounds—computation, information, attention. Subsequent research has extended this framework in several directions: cognitive bounds (Kahneman and Tversky, 1979), environmental bounds (Gigerenzer and Todd, 1999; Hertwig et al., 2019), and computational bounds (Lieder and Griffiths, 2020).

This paper extends the analysis of boundedness into an existential register. Rationality must also operate under survival-related constraints—otherwise the computing subject itself may disappear (what do you compute if you have no computer?). This survival orientation aligns with Kenrick et al. (2009)'s Deep Rationality, though they locate it in evolutionarily inherited motives. We share the same theme but complement it with a more computable framework focused on what lies between the cognitive subject and the environment in the present—the joint: the survival condition defined by finite opportunities and the distance to possible ruin in the future. From this perspective, while Taleb conceptualized the risk one holds in their position as "skin in the game" (Taleb, 2018), in the cognitive-scientific context this means that even under the same environment, what is prioritized, what is discounted, and what is ignored differs according to what the subject cannot afford to fail at and what they can. Our Off-Book is the framework that captures the weighting—under-considered in laboratory ledgers—computed by the degree of survival risk each subject carries, using ergodicity and fungibility as rulers. It is the joint between the two blades of Simon's scissors, where each cognitive subject translates the environment into the language of distance to their risks.

Building on this framework, we take the Kahneman-Gigerenzer debate as representative poles of the 30-year disagreement on human rationality and offer a unified account by identifying the survival conditions under which the same behavior splits into "bias" or "adaptation."

Off-Book 1 (Biological Survival) concerns the gap between mathematical abstraction (statistics, expected value, probability) and biological reality (adaptation, survival). Kahneman labeled systematic deviations as "bias" (Tversky and Kahneman, 1974); Gigerenzer and Todd (1999) reframed them as "ecological adaptation." Ergodicity was first formalized by physicists in the late 19th century and recently reintroduced to decision theory and economics by Peters (2019). It provides a clear explanation distinguishing time-averages (individual gain/loss) from ensemble averages (expected value in group statistics). In reality, individual lives are often **non-ergodic**: what works for the group does not guarantee long-term survival for the individual. The risk to avoid is **Ruin**.

Off-Book 2 (Social Survival) concerns value that we need as social animals. Unlike bank balances that are fungible, social capital is often **non-fungible** (irreplaceable); once trust or a relationship is lost, it cannot be easily repurchased. This has been recognized in sociology (Bourdieu, 1986; Coleman, 1988) and is common sense in management (Drucker, 1954; Collins, 2001), but remains underspecified in cognitive science. In reality, social lives are often non-fungible: preserving the social status quo is a rational insurance premium for social survival. The risk to avoid is **Isolation**.

The laboratory's justice is Efficiency. The organism's justice is Viability. The social animal's justice is Safetynet—the bonds that catch you when you fall.

Humans must survive both as biological individuals (avoiding Ruin) and as social animals (avoiding Isolation), while maximizing fungible resources. This dual survival requirement helps explain why "inefficient" behavior persists. **Figure 1** presents the Off-Book Matrix, crossing two survival

In G. Boleda, I. Dautriche, M. Maldonado, C. Saldana and T. Torrent (Eds.), *Proceedings of the 48th Annual Conference of the Cognitive Science Society*

requirements (biological, social) with two ledger conditions (on-book, off-book).

Hypothesis: Under laboratory conditions, where ergodicity and fungibility are assumed, behavior may be labeled as irrational, whereas under real-world existential constraints it may be adaptive.

The contributions of this paper are threefold:

1. **Theoretical**: We identify two boundary conditions—non-ergodicity and non-fungibility—as existential mediators of the Kahneman-Gigerenzer disagreement.
2. **Empirical**: We demonstrate that rationality judgments flip when survival requirements are made salient ($N = 100$; loss aversion $d = 0.99$, sunk cost $d = 0.47$, interaction $p = .004$).
3. **Conceptual**: We introduce the Off-Book framework as a translator of the environment for a subject in decision-making.

Theoretical Framework

Off-Book 1: Non-Ergodicity

Worldview: Math (statistics, modeling, probability) vs Biology (adaptation, environment)
Objective function: Biological Survival
On-Book: Expected value, Ensemble average
Off-Book: Ruin avoidance, Time average

Why $N = 1$ Matters. Traditional experiments in cognitive science evaluate rationality by averaging across many participants (and the more, the better), implicitly assuming ergodicity—which is to imply what is optimal for the group is also optimal for each individual over time. However, in non-ergodic environments, this assumption no longer holds. Peters (2019) noted that humans live **sequentially through time**, not in parallel across a group. Time-averages diverge from ensemble averages, and expected-value maximization can lead to ruin. That is to say, there is no "real statistically average guy"—and it is very hard to become one.

Suppose a coin flip where heads increases the bet by 50% and tails decreases it by -40%. The mathematical expected value becomes 5%—a reasonably favorable bet (for reference, given that the yield of generally safe investments is around 3–8% per year). Suppose: start at \$100, exit below \$30.

Mathematically, expected value holds only when you run a sufficient number of trials—which, strictly speaking, means splitting yourself into 100 parallel copies playing at once. But in reality, you live through time, and only once.

If 100 people play once (ensemble): 50 win (\$150), 50 lose (\$60). The average converges to \$105. No one exits, since each member played only once. By contrast, if 1 person plays six rounds (time-series): win twice, lose four times—23% probability. Trajectory: \$100 $\rightarrow$ \$150 $\rightarrow$ \$225 $\rightarrow$ \$135 $\rightarrow$ \$81 $\rightarrow$ \$49 $\rightarrow$ \$29. One must **exit**. Same game, same odds, same expected value—the ensemble player sees 5% while the time-series player is eliminated. This is non-ergodicity: what works for the group does not guarantee survival for the individual. **Table 1** shows the divergence is structural.

$T \rightarrow N$'s Validity. Consider a casino offering positive expected value (as you know, this is not usually the case). "Not going is irrational," says the maximizer. "Stay away from danger and save money," says the survivor. The crucial insight: **the subject's condition determines T (how many trials one can endure), and T determines whether N (expected value, probability) becomes a valid reference at all.** If one can run 100 rounds, expected value will converge to a 5% gain. But we are often in a position where we cannot afford enough trials for expected value to emerge in the first place.

This is the **relativity of rationality**. Under the same environment, the same expected value, and the same rules, the same game splits into ergodic or non-ergodic depending on who plays—and this flips the direction of decision-making entirely. If you have enough trials (T)—say, you can afford to bet \$10k—expected value converges, because you can afford the 100 rounds. If not, survival risk may emerge before the math has a chance to converge to its average. A decision becomes rational or irrational depending on the person's trial capacity (T).

Unifying Mechanism: Ergodicity. What looks like "Loss Aversion" on the On-Book ledger of laboratories and data often turns out to be, in actual real-life cases (Off-Book), "Ruin Avoidance"—a legitimate rationality not for winning the game, but for staying alive. When T is large, the environment is ergodic, and Kahneman's framework applies. When T is small, the environment is non-ergodic, and Gigerenzer's framework applies. Whether one should 'take the bet' or 'walk away' depends not on whose theory one prefers, or who is more rational, but truly on the subject's position that translates what the environment means to them.

Off-Book 2: Non-Fungibility

Worldview: Measurable (money) vs Sociology/Management (relationships, social capital)
Objective function: Social Survival
On-Book: Fungible resources (money, Key Performance Indicators)
Off-Book: Non-fungible capital (trust, reputation, relationships, belonging)

What the Ledger Misses. Bank balances are measurable. If reduced, they might be recovered—that is, they are reversible. But trust with an important partner may not be recoverable. Like broken glass that cannot be reassembled, it is irreversible. Salaries can be compared. But the damage of being labeled "the office pariah" erodes the ledger indirectly yet inexorably. Investment returns can be calculated. But whether you have friends to rely on when you lose your job—that stays off the books.

Consider the same decision—"should I stay or should I go?"—at different life stages. The trigger is identical: two serious arguments in a row. The question is the same: cut your losses, or stay?

Young and dating: Changing partners may be rational. The relationship is short, the shared history shallow. A compara-

Figure 1: The Off-Book Matrix. Laboratory ledgers record only the On-Book columns (expected value, fungible capital). Off-Book columns (ruin avoidance, isolation avoidance) remain unmeasured yet determine whether "inefficient" behavior is rational. Loss aversion = Ruin avoidance; Sunk cost = Isolation avoidance.

Table 1: Ergodic vs Non-Ergodic Evaluation.

Concept	Lens	Math Phenomenon	Book Classification	Adaptive Strategy
Math	Ergodic	Ensemble Average	On-Book (visible gains)	Optimization, Maximization
Biology	Non-Ergodic	Time Average	Off-Book (invisible ruin risk)	Robustness, Survival

ble relationship can be built with someone else. The loss is relatively fungible—not just because time is abundant, but because **the social support required for survival can likely be obtained elsewhere—through parents, friends, or a new partner**. In principle, every relationship is non-fungible—each person is unique. Each person is unique; what we assess here is not the person but the availability of social support essential for survival.

Married for 20 years: The same two arguments, the same frustration. But the accumulated trust and mutual support are not so easily rebuilt. Twenty years of shared history, mutual adaptation, and accumulated trust cannot be simply replicated with a new partner. The loss is acutely non-fungible.

Why does this matter beyond sentiment? Social capital is not merely emotional—it is a **buffer against biological ruin**. Consider: you fall seriously ill. In the first scenario, you recover in a hospital bed, alone—but young, with time to rebuild networks. In the second scenario, you lie in that same bed having just severed your deepest connection. The absence is not just emotional; it is structural. Who advocates for your care? Who checks that the diagnosis is correct? Who is there when you come home? Social survival is not parallel to biological survival—**it is a leading indicator of it.**

The sunk cost "fallacy" looks different through this lens. On-Book, "clinging to past investments (the effort spent maintaining the relationship) is irrational." But Off-Book, this is not fallacy but "protection of social value that is likely irreplaceable"—a survival strategy in the social ledger. **Table 2** shows that this divergence parallels Off-Book 1.

$T \to N$'s Validity, Revisited. In Off-Book 1, we showed that the same environment splits into ergodic or non-ergodic depending on who is playing. A parallel structure emerges in Off-Book 2: the same loss splits into fungible or non-fungible depending on if it is recoverable. The subject's social capital determines whether a loss is recoverable, and recoverability determines whether "cut your losses" is valid advice at all. The same event, the same logic, but the fungibility of what is lost depends on T—is it recoverable, if so how many times.

The same logic applies to reputation. A scandal in youth is recoverable—"youthful indiscretion." The same act in later years is "a stain on one's legacy"—not because the act differs, but because **the recoverability is unlikely. Non-fungible becomes increasingly non-fungible as T shrinks.**

"Cut your losses" is likely to be valid advice only when you have enough possibilities to rebuild the non-fungibles you cut.

Unifying Mechanism: Fungibility. What appears as

Table 2: Fungible vs Non-Fungible Evaluation.

Concept	Lens	Math Phenomenon	Book Classification	Adaptive Strategy
Measurable	Fungible	Replaceable loss	On-Book (visible cost)	Optimization, Cut losses
Sociology	Non-Fungible	Irreplaceable loss	Off-Book (invisible trust risk)	Social Survival

"Sunk Cost Fallacy" On-Book is "Isolation Avoidance" Off-Book—a legitimate rationality for social survival, not staying irrational. When social capital is abundant and time is plentiful, cutting losses may be rational—relationships could be rebuilt. When social capital is scarce, time is limited, or the relationship is central to one's quality of life, persisting is rational—isolation threatens survival. Whether one should "cut losses" or "persist" depends not on which advice sounds wiser, but on how close the subject stands to social isolation—and the diminished quality of life it entails.

Social survival is not merely a parallel concern to biological survival—it is a **precursor**. When quality of life radically drops, when a company's reputation collapses, when nobody in the community knows your name—these failures in social survival raise the risk of the first existential requirement: biological survival. One may live shorter under chronic stress. One may lose the job needed to make rent next month. One may not survive a natural disaster without a network to call for help. **Off-Book 2 is the leading indicator of Off-Book 1.**

Conversely, when non-fungible social capital is intact, it functions as a **safety net**—the bonds that catch you when you fall. We know this intuitively. This is precisely why we continue to pay the sunk cost of maintaining relationships, even when the ledger says it is irrational. We are not clinging to past investments; we are investing in future survival. Isolation avoidance is not a fallacy—it is insurance.

The Infrastructure of Belonging. Ultimately, society is an infrastructure for individual survival. We pay seemingly irrational costs—maintaining costly relationships, contributing to social welfare, adhering to local norms—not because we expect an immediate return, but because these are the subscription fees for the ultimate insurance.

When we instinctively seek to "get along" with neighbors or fulfill social duties, we are not being "irrational." We are probably performing the most fundamental biological calculation: ensuring that when our fungible resources run low, or we face unprecedented misfortune or accident, and we stand before the absorbing barrier of ruin, the social network remains intact to catch us. The "wasteful" social cost of today is the structural resilience of tomorrow. But this is not a guarantee—we remain fragile to uncertainty. The safety net may hold, or it may not. What the Off-Book framework captures is not certainty but the rationality of hedging against isolation that may lead to ruin.

> "In the moment of crisis, you don't live on your balance sheet; you live on your infrastructure of belonging."

Empirical Study

Method

Participants. One hundred participants ($N = 100$) were recruited via Prolific. All provided informed consent.

Design. We employed a 2 (bias type: loss aversion vs. sunk cost) $\times$ 2 (context: lab-like vs. life-like) within-subjects design. Presentation order of the four scenarios was randomized across participants.

Materials. Four scenarios presented identical "biased" behaviors across different contexts. To test Split 1 (ergodicity), we manipulated the loss aversion scenarios: In the *ergodic condition*, Mr. Brown declined a 5% asset investment with positive expected value (+50%/−10%, $p = .50$ each; expected value = +20%). In the *non-ergodic condition*, Mr. Smith was encouraged by a friend to invest 50% of his retirement savings in a new business venture. The investment had a 50% probability of success: if successful, the invested amount would double (+100%), whereas if unsuccessful, he would lose 50% of his investment (expected value = +25%). Given that rebuilding assets at his age would be realistically difficult, he placed greater weight on potential losses and chose not to invest.

To test Split 2 (fungibility), we manipulated the sunk cost scenarios: In the *fungible condition*, Mr. Johnson held losing stocks despite available alternatives. In the *non-fungible condition*, Mr. Williams maintained a strained 10-year friendship where the shared community made alternatives difficult. Participants rated whether each decision was rational (1 = *Extremely irrational*, 7 = *Extremely rational*).

Results

A 2×2 repeated-measures ANOVA revealed significant main effects of bias type, $F1,99 = 22.75$, $p < .001$, $\eta_g^2 = .07$, and context, $F1,99 = 93.10$, $p < .001$, $\eta_g^2 = .17$. Critically, the interaction was significant, $F1,99 = 8.64$, $p = .004$, $\eta_g^2 = .02$, indicating that the context effect differed by bias type (**Table 3**).

Both effects were significant and in the predicted direction. The context effect was larger for loss aversion ($d = 0.99$) than for sunk cost ($d = 0.47$). This asymmetry may reflect that ergodicity—the distinction between survival and ruin—is more intuitively accessible than fungibility—the distinction between replaceable and irreplaceable capital.

These results support H1: identical behaviors are judged more rational in life-like contexts. The large effect for loss aversion ($d = 0.99$) suggests laypeople intuitively recognize the rationality of caution when ruin is irreversible. The

Table 3: Rationality judgments by condition (N = 100).

Bias Type	Lab-like	Life-like	Δ	$t(99)$	p	d
Loss Aversion (Ergodicity)	4.70 (1.69)	6.40 (1.22)	+1.70	9.88	<.001	0.99
Sunk Cost (Fungibility)	4.28 (1.61)	5.23 (1.31)	+0.95	4.71	<.001	0.47

Note. Standard deviations in parentheses. Lab-like = ergodic/fungible; Life-like = non-ergodic/non-fungible.

medium effect for sunk cost (d = 0.47) indicates that non-fungibility of social capital also shapes rationality judgments, though less dramatically. The significant interaction confirms that rationality judgments are not absolute properties of behaviors, but context-dependent evaluations sensitive to the structural features—ergodicity and fungibility—of decision environments.

Related Work

The need for off-book thinking was already present in the early days of cognitive science. Neisser (1976) warned early that laboratory research captures only a limited portion of real-world cognition. Anderson (1990), Hertwig et al. (2019), and Lieder and Griffiths (2020) responded from computational and environmental perspectives, while Weber and Johnson (2009), Cushman (2013), Schwartz (2004), and Iyengar and Lepper (2000) highlighted the limits of cognition from diverse angles, including affect, morality, and choice overload.

This study introduces two existential constraints. The first is **ergodicity**—a concept reconnected to decision theory by Peters (2019). The second is **fungibility**—drawing on the social capital literature of Bourdieu (1986) and Coleman (1988), and the off-ledger value traditions of Drucker (1954), Collins (2001), and Weick (1995). This study integrates these streams and formalizes the dual split.

Furthermore, the lineage of complexity science and systems theory resonates with this study. Lorenz's (1963) butterfly effect—where minute differences in initial conditions produce divergent outcomes—can be read as a caution against averaging across individuals. Ashby's (1956) law of requisite variety directly addresses this, arguing that adaptive systems must match environmental complexity. Drawing on this and related systems-theoretic perspectives (Wiener, 1961; von Bertalanffy, 1968), we examine how cognitive agents confront complex reality.

Discussion

The Answer to "Why Does Inefficiency Survive?" We return to the title of this study: Why does inefficiency survive? The answer is now clear—it survives because inefficiency is efficiency for survival. Behaviors judged "inefficient" by one ruler are rational strategies when measured by the ruler of biological or social survival. The laboratory often measures deviation from expected value; the organism measures distance from ruin; the social animal measures distance from isolation. The same behavior becomes either "bias" or "adaptation" depending on one's distance from risk and the degree of skin in the game—conditions that dynamically change through time. Our empirical results confirm this: the same participants, significantly affected by these existential constraints, flipped their judgments. They shifted their rationality judgments based on whether ruin was possible (d = 0.99) and whether the loss was recoverable (d = 0.47). Inefficiency, it turns out, is in the eye of the measurer.

Why Loss Aversion Showed the Larger Effect. The most striking finding of this study is that the context effect for loss aversion (d = 0.99) substantially exceeded that for sunk cost sensitivity (d = 0.47). Why? The scenario involving "investing 50% of retirement savings" clearly evoked the possibility of ruin. Participants judged that "even with positive expected value, it is not rational because if you lose, there is no recovery"—Off-Book 1 was cognitively salient. In contrast, the non-fungibility of relationships (Off-Book 2) may not have been as clearly recognized as ruin risk—or participants may have optimistically assumed that relationships, once severed, could be rebuilt.

In either case, both effects were significant and the interaction was confirmed ($F_{1,99} = 8.64$, $p = .004$). This supports H1.

What This Means. We do not measure survival directly. Instead, we measure whether laypeople intuitively recognize survival-relevant contexts as warranting different rationality standards. The large effect for loss aversion suggests that ruin-possibility is cognitively salient. The moderate effect for sunk cost suggests that recoverability is recognized but less vivid. This aligns with our theoretical framework: social survival is not a direct threat, but an indirect factor that ultimately affects biological survival. If ordinary people adjust their rationality judgments based on these structural features, off-book reasoning is already operative in everyday cognition.

Implications for Contemporary Society. Although speculative, our framework offers a lens for reinterpreting several phenomena in contemporary society.

Rising divorce rates, short-term relationships, organizational fluidity, short-term investment horizons—these are often narrated as progress toward "individual and organizational optimization, rationalization, and efficiency." Yet from the perspective of this study, an alternative interpretation emerges: these may be consequences of treating inherently **non-fungible** social capital—relationships, reputation, organizational trust—as **fungible** (replaceable, rebuildable). Simultaneously, they may result from treating inherently **non-ergodic** life decisions as repeatable trials.

Sunk cost sensitivity has functioned as a stabilizing mechanism that avoids irreversible loss. Uniformly eliminating it may increase short-term flexibility while eroding trust capital

in the long run. Similarly, treating loss aversion as a bias to be overcome may increase short-term efficiency while reducing sensitivity to catastrophic risk. The modern optimization imperative—cut your losses, maximize expected value, stay flexible—is sound advice in ergodic, fungible environments. In non-ergodic, non-fungible life, it may be a recipe for ruin and isolation.

Limitations and Future Directions

This study opens four research agendas.

Beyond Loss Aversion and Sunk Cost. This study focused primarily on two biases: loss aversion and sunk cost. However, other cognitive biases and characteristics may also be explicable through the Off-Book Matrix framework. This remains to be tested.

Relative Rationality. This direction points toward a shift from Universal Rationality to Relative Rationality. Building on Simon (1955) and Gigerenzer and Todd (1999), we extend the analysis: rationality depends not only on cognition and environment but on where the subject stands in conjunction to both. This is not relativism in the sense of "anything goes." It is the recognition that the ruler must fit the measured—and that one ruler cannot fit all lives.

Off-Book, through its axes of ergodicity and fungibility, identifies the subject's specific position—the junction between the internal (cognition) and the external (environment)—functioning as a translator of environmental conditions into survival-relevant information. Even under the same environment, if the translation axes differ, entirely contradictory rationalities can hold simultaneously across different subjects; and within a single subject, one can hold contradictory rationalities when the two axes of off-book shift.

$T \to N$'s Validity. The key empirical program is to identify the thresholds: under what conditions does expected value become a meaningful reference for a given subject? When does an environment shift from ergodic to non-ergodic? When does a loss shift from fungible to non-fungible? The dynamics of off-to-on conversion—how non-fungible capital translates into fungible outcomes over time—also deserve fuller treatment.

Under what conditions does N (ensemble expected value) become valid? Suppose a single investment is \$1,000 with an expected return of +5%. With \$100,000, the outcome should converge to the expected value over $N = 100$ trials. But two hidden conditions must hold: **uniform allocation** and **regular timing**.

If you invest \$10,000 on the first trial, the remaining T is constrained to 90, destroying the validity of $N = 100$. If you stop investing for three months when things go poorly, this too destroys $N = 100$. Mathematically, the ensemble holds; operationally, it is perpetually threatened.

This structure extends beyond investment.

Language learning—30 minutes a day for 365 days should lead to acquisition. But "I'm tired today" leads to a week off, "I'm not improving" leads to two months off, and resumption feels like starting over. Streak functions in language apps are dollar-cost averaging for regular timing.

Exercise—three times a week for six months should improve fitness. But no visible results after two weeks triggers "I'm not suited for this," and one injury leads to three months off, then permanent dropout. Gym membership fees are mechanisms that sustain T regardless of emotion.

In short, the agent that destroys the ensemble is neither the market nor probability—it is **human emotion, fear, and bias**. 401(k) automatic payroll deductions and dollar-cost averaging are defense mechanisms against this humanity—iron rules that enforce uniform allocation and regular timing.

Dual Survival Requirement. This study treated Off-Book 1 (biological survival) and Off-Book 2 (social survival) in parallel. Although our empirical results and theoretical rationale support this connection, it is not yet firmly established. Humans must survive both as biological individuals and as social animals—but these requirements may conflict, and their interaction warrants further investigation.

This study has several limitations.

First, we rely on rationality judgments rather than behavioral measures. Participants evaluated hypothetical others; whether they would act according to these judgments in their own lives remains untested.

Second, while participants shifted their judgments based on context, we cannot isolate the causal mechanism. The effect may stem from intuitive recognition of survival relevance, or from other features of the scenarios such as emotional salience or narrative concreteness.

Third, our sample is limited to Prolific participants from predominantly Western, educated backgrounds. Cultural variation—particularly in the valuation of social capital—may moderate these effects. The effect of Off-Book 2 may differ between individualistic and collectivistic cultures.

Fourth, each condition was instantiated by a single scenario, selected to exhaust the 2×2 structure of the Off-Book Matrix. Future work should employ multiple scenarios per condition to separate structural effects from scenario-specific confounds.

Conclusion

Humans must survive both as **biological individuals** (avoiding Ruin) and as **social animals** (avoiding Isolation). A brilliant strategy that leads to bankruptcy is not brilliant. A profitable decision that leaves you friendless is not profitable.

Our empirical results are modest—read scenarios, make judgments. Yet this simplicity opens onto a larger question. The results ($d = 0.99$ for loss aversion, $d = 0.47$ for sunk cost sensitivity) confirm what laypeople already seem to know: rationality is not one thing. It is a matter of relativity, and it diverges based on the subject's interpretation of the environment.

Why does inefficiency survive? The Off-Book Matrix flips inefficiency and efficiency into one another.

"The bias is not in the subject—it is in the ruler."

Generative AI Usage Statement

We used AI (Claude) assistance for figure design, translation, LaTeX formatting, and language editing during the preparation of this manuscript. All intellectual contributions, arguments, and conclusions are the author's own.

Acknowledgements

The author is grateful to the members of Tokyo Cognitive Lab and all who supported this work. The author also thanks the anonymous reviewers, whose comments substantially improved the manuscript.

References

Anderson, J. R. (1990). *The Adaptive Character of Thought*. Lawrence Erlbaum Associates.

Ashby, W. R. (1956). *An Introduction to Cybernetics*. Chapman & Hall, London.

Bourdieu, P. (1986). The forms of capital. In Richardson, J., editor, *Handbook of Theory and Research for the Sociology of Education*, pages 241–258. Greenwood Press.

Coleman, J. S. (1988). Social capital in the creation of human capital. *American Journal of Sociology*, 94:S95–S120.

Collins, J. (2001). *Good to Great: Why Some Companies Make the Leap and Others Don't*. HarperBusiness.

Cushman, F. (2013). Action, outcome, and value: A dual-system framework for morality. *Personality and Social Psychology Review*, 17(3):273–292.

Drucker, P. F. (1954). *The Practice of Management*. Harper & Row.

Gigerenzer, G. and Todd, P. M. (1999). *Simple Heuristics That Make Us Smart*. Oxford University Press.

Hertwig, R., Pleskac, T. J., Pachur, T., and the Center for Adaptive Rationality (2019). *Taming Uncertainty*. MIT Press.

Iyengar, S. S. and Lepper, M. R. (2000). When choice is demotivating: Can one desire too much of a good thing? *Journal of Personality and Social Psychology*, 79(6):995–1006.

Kahneman, D. and Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2):263–291.

Kenrick, D. T., Griskevicius, V., Sundie, J. M., Li, N. P., Li, Y. J., and Neuberg, S. L. (2009). Deep rationality: The evolutionary economics of decision making. *Social Cognition*, 27(5):764–785.

Lieder, F. and Griffiths, T. L. (2020). Resource-rational analysis: Understanding human cognition as the optimal use of limited computational resources. *Behavioral and Brain Sciences*, 43:e1.

Lorenz, E. N. (1963). Deterministic nonperiodic flow. *Journal of the Atmospheric Sciences*, 20(2):130–141.

Neisser, U. (1976). *Cognition and Reality: Principles and Implications of Cognitive Psychology*. W. H. Freeman.

Peters, O. (2019). The ergodicity problem in economics. *Nature Physics*, 15(12):1216–1221.

Schwartz, B. (2004). *The Paradox of Choice: Why More Is Less*. Ecco.

Simon, H. A. (1955). A behavioral model of rational choice. *Quarterly Journal of Economics*, 69(1):99–118.

Taleb, N. N. (2018). *Skin in the Game: Hidden Asymmetries in Daily Life*. Random House.

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

von Bertalanffy, L. (1968). *General System Theory: Foundations, Development, Applications*. George Braziller, New York.

Weber, E. U. and Johnson, E. J. (2009). Mindful judgment and decision making. *Annual Review of Psychology*, 60:53–85.

Weick, K. E. (1995). *Sensemaking in Organizations*. Sage, Thousand Oaks, CA.

Wiener, N. (1961). *Cybernetics: Or Control and Communication in the Animal and the Machine*. MIT Press, Cambridge, MA.