Finite-Capacity Prospect Theory: State-Dependent Risk Preferences under Resource, Attention, and Boundary-Risk Constraints

A Boundary-Risk Account of Loss Aversion, Reference Dependence, and Probability Weighting

Yining Wu

yining.wu@alumni.upenn.edu

May 2026

Abstract

Prospect Theory identifies three robust regularities in decision-making under risk: loss aversion, reference dependence, and nonlinear probability weighting. These regularities are usually treated as descriptive deviations from expected-utility rationality. This paper proposes a finite-system bridge account: for a bounded agent maintaining identity under resource scarcity, limited attention, limited distinguishability, and finite update capacity, these regularities can be interpreted as state-dependent boundary-risk management strategies rather than fixed irrationality constants. The paper models an economic agent as a finite boundary-maintaining decision system with a vector-valued resource buffer, a perceived distance to constraint threshold, a maintained reference baseline, and a limited probability-distinction budget. Loss aversion arises from asymmetric boundary-risk slopes near the lower resource constraint; reference dependence arises from finite-memory prediction baselines; and probability weighting arises from finite precision allocation over probability and event-consequence space. The model predicts that Prospect Theory parameters are functions of resource buffers, attention capacity, and boundary-risk context: loss aversion should rise under scarcity and downside exposure; reference baselines should adapt more slowly under depletion; and probability weighting should become more distorted for low-capacity agents and boundary-relevant rare events. Numerical simulations illustrate the mechanism, compare static and state-dependent parameter recovery in synthetic data, and show a nudge-bandwidth region in which complex choice architecture interventions fail or backfire under depleted decision capacity. The paper does not claim to derive the empirical coefficient $\lambda \approx 2.25$ exactly, nor to reduce all economic behavior to thermodynamics. Its contribution is narrower: it formulates a testable finite-capacity bridge between behavioral economics, bounded rationality, and finite-system maintenance theory.

Keywords: prospect theory; loss aversion; reference dependence; probability weighting; bounded rationality; scarcity; finite capacity; behavioral economics; decision theory; nudge; resource constraints; active finite distinction systems.

JEL codes: D01, D81, D91, C91, I31, H31.

Scope of the Paper. This paper does not claim that all behavioral-economic phenomena are reducible to physics, nor that the empirical loss-aversion coefficient $\lambda \approx 2.25$ is derived exactly from first principles. It does not replace Prospect Theory, cumulative prospect theory, expected utility theory, bounded rationality, active inference, ecological rationality, or evolutionary psychology. The paper develops a finite-capacity bridge model in which economic agents are treated as bounded decision systems with finite resource buffers, attention, memory, precision, update capacity, and tolerance for boundary risk. The purpose is to generate testable predictions about when and why Prospect Theory parameters should vary across resource states, cognitive load, and boundary-risk contexts.

1 Claim-Status Summary

Table 1: Central claims, epistemic status, and demotion conditions.

Claim	Status	What would weaken or falsify it
Economic agents are finite-capacity decision systems	Operational bridge	Decision behavior invariant to memory, stress, attention, and resource constraints
Loss aversion increases near perceived resource thresholds	Testable bridge hypothesis	λ remains invariant across physiological, financial, cognitive, and downside-protection states
Reference dependence reflects maintained prediction baselines	Computational bridge	Reference points fail to correlate with prediction-error integration or baseline-updating dynamics
Probability weighting reflects finite precision allocation	Testable bridge hypothesis	Weighting curvature remains invariant under cognitive load, stress, event class, and resource state
Rare high-consequence events receive excess salience or precision	Domain bridge	Survival-relevant and neutral rare events show identical weighting after controls
Nudges consume decision bandwidth	Policy bridge	Nudge effectiveness independent of stress, scarcity, fatigue, administrative burden, and cognitive load
Prospect parameters are state-dependent	Central prediction	Stable λ, reference adaptation rate, and weighting curvature across resource regimes and boundary-risk contexts
Reverse stress test	Strong falsification condition	Severe resource depletion changes neither λ, reference rigidity, probability weighting, nor decision noise after adequate power and controls

2 Introduction

Kahneman and Tversky (1979) transformed decision theory by showing that risky choice is not well described by the expected-utility model. In Prospect Theory, outcomes are evaluated relative to a reference point; losses loom larger than gains; and objective probabilities are transformed into nonlinear decision weights. These empirical regularities were later extended in cumulative prospect theory (Tversky and Kahneman, 1992). Decades of behavioral-economics research have confirmed their importance for finance, insurance, labor supply, health, household choice, and public policy (Barberis, 2013). A recent meta-analysis of Prospect Theory parameters reports substantial between-study variation and systematic procedure dependence, suggesting that the field still lacks a satisfactory account of why parameter estimates vary across environments and tasks (Imai et al., 2025).

The success of Prospect Theory creates a deeper question. Are loss aversion, reference dependence, and probability weighting merely descriptive anomalies relative to rational choice? Or do they reveal the structure of decision-making by finite agents under physical, cognitive, social, and informational constraints? Standard explanations appeal to evolutionary adaptation, bounded rationality, affective processing, or ecological rationality (Simon, 1955; Gigerenzer and Todd, 1999). These accounts are valuable, but they often leave the key parameters of Prospect Theory as stable psychological facts rather than state-dependent control variables.

This paper develops a finite-capacity account. An economic agent is not an ideal optimizer with unlimited representation, memory, attention, and energy. It is a finite boundary-maintaining system. It must maintain a viable boundary, allocate limited attention, update a compressed internal state, preserve a resource buffer, and act under uncertainty. In the language of Active Finite Distinction Systems (FDS), an agent has a finite distinction projection and finite distinguishability budget: it cannot represent all outcome states, probabilities, future contingencies, and action consequences with equal precision. When task demand exceeds accessible capacity, the agent must compress, coarse-grain, externalize, prune, relax the task, or fail.

The paper's central thesis is that Prospect Theory parameters are not fixed irrationality constants. They are state-dependent finite-system control variables:

$$(\lambda_t, \alpha_t, \gamma_t) = f(F_t^{\text{eff}}, C_t, B_t), \tag{1}$$

where λ_t is loss aversion, α_t is reference-baseline adaptation, γ_t is probability-weighting curvature, F_t^{eff} is effective resource buffer, C_t is available decision capacity, and B_t is boundary-risk context. The model predicts that scarcity, fatigue, stress, financial precarity, time pressure, administrative burden, and downside exposure should change the measured parameters of risky choice.

The contribution is not to replace Prospect Theory. It is to make Prospect Theory dynamic. The paper provides a state-dependent model, falsifiable predictions, simulation code, synthetic parameter-recovery exercises, and policy implications for nudge design under scarcity.

3 Prospect Theory: Three Regularities

Let x denote an outcome relative to a reference point r. A common value function is

$$v(x) = \begin{cases} x^a, & x \geq 0, \\ -\lambda(-x)^b, & x < 0, \end{cases} \tag{2}$$

where $\lambda > 1$ captures loss aversion (Kahneman and Tversky, 1979). The empirical value $\lambda \approx 2.25$ is often treated as a benchmark, but estimates vary by domain, method, subject pool, and framing.

Reference dependence means that utility is not evaluated over final wealth alone. It is evaluated over changes relative to a reference point. The reference point may depend on status quo, expectations, past outcomes, social comparison, or contractual entitlement.

Probability weighting means that agents transform objective probabilities p into decision weights $w(p)$. A one-parameter inverse-S curve can be written as

$$w(p;\gamma) = \frac{p^\gamma}{p^\gamma + (1-p)^\gamma}, \quad 0 < \gamma \leq 1. \tag{3}$$

When $\gamma < 1$, small probabilities are overweighted and moderate probabilities are compressed. More flexible specifications include the Prelec form (Prelec, 1998).

The finite-capacity account does not deny these descriptive forms. It asks why the parameters should move with resource state, attention capacity, and boundary-risk context.

4 Finite-System Preliminaries for Behavioral Economics

Definition 1 (Finite boundary-maintaining decision system) *An economic agent is modeled as a finite decision system*

$$A_t = (B_t, M_t, R_t, \mathbf{F}_t, C_t, \pi_t, \Psi_t, U_t, \tau_t), \tag{4}$$

where B_t is a maintained boundary or viability condition, M_t is memory, R_t is a reference baseline, $\mathbf{F}_t$ is a vector of resource buffers, C_t is available decision capacity, π_t is a finite distinction projection, Ψ_t is the decision environment, U_t is an action-update policy, and τ_t is the decision horizon.

The agent's finite distinction projection induces a decision-relevant budget

$$N_A(t) = |\operatorname{Im}(\pi_t)|, \qquad C_A(t) = \log_2 N_A(t). \tag{5}$$

This budget is not only memory. It includes attention, channel access, time pressure, numeracy, working memory, emotional regulation, and access to external records or institutional supports.

4.1 Vector resource buffers and substitutability

A key revision relative to earlier thermodynamic language is that the resource buffer is vector-valued:

$$\mathbf{F}_t = (F_t^m, F_t^e, F_t^s, F_t^a, F_t^\tau), \tag{6}$$

where F^m denotes monetary or financial buffer, F^e energetic or metabolic buffer, F^s social or institutional support, F^a attention or cognitive bandwidth, and F^τ time buffer. These resources are not equally substitutable. Money can sometimes compensate for time, food, or security; sleep loss or acute illness may be much less fungible.

We define effective resource buffer by a CES aggregator:

$$F_t^{\text{eff}} = \left(\sum_i \omega_i (F_t^i)^\eta \right)^{1/\eta}, \qquad \sum_i \omega_i = 1. \tag{7}$$

The parameter η controls substitutability. When η is high, resources are more fungible. When $\eta \to -\infty$, the effective buffer approaches the minimum resource dimension. This captures a central prediction: depletion of non-fungible resources such as sleep, health, time, or safety should produce more nonlinear changes in risk preferences than depletion of more substitutable resources.

4.2 Perceived distance to boundary

Economic behavior depends on perceived, not merely objective, distance to constraint. Let F_{crit} denote the boundary-risk threshold and define objective distance

$$d_t = \frac{F_t^{\text{eff}} - F_{\text{crit}}}{F_{\text{scale}}}. \tag{8}$$

The perceived distance is

$$\widehat{d}_t = \theta_t d_t, \tag{9}$$

where θ_t summarizes belief, trauma history, social comparison, uncertainty, information quality, and perceived institutional support. A wealthy agent with low perceived security may have low $\widehat{d}_t$ despite high objective resources. Conversely, a poorer agent with strong social insurance may have a larger perceived buffer than monetary wealth alone implies.

5 Loss Aversion as Boundary-Risk Asymmetry

Let $R(d)$ be a boundary-risk potential, decreasing and convex in perceived distance d:

$$R(d) = \kappa (d + \epsilon)^{-\nu}, \qquad \kappa > 0,\ \epsilon > 0,\ \nu > 0. \tag{10}$$

For a small resource change $\Delta > 0$, the local loss-aversion coefficient induced by boundary-risk asymmetry is

$$\lambda(d, \Delta) = \frac{R(d - \Delta) - R(d)}{R(d) - R(d + \Delta)}. \tag{11}$$

Because R is convex and decreasing near the threshold, downward movements toward the boundary increase risk more than equal upward movements reduce it. Thus $\lambda(d, \Delta) > 1$ generically when the agent is close to the boundary.

Proposition 1 (Boundary-risk loss aversion) *If perceived boundary-risk potential $R(d)$ is decreasing and convex over the relevant neighborhood, then a finite agent evaluating symmetric resource changes around d exhibits local loss aversion $\lambda(d, \Delta) > 1$.*

Interpretation. The empirical coefficient $\lambda \approx 2.25$ is treated here as a calibration target and empirical benchmark, not as an exact first-principles constant. The stronger prediction is state-dependence: λ should increase as $\widehat{d_t}$ falls, and it should increase more sharply when the depleted resource is non-fungible.

Prediction 1 (Non-fungible resource depletion) *Depletion of non-fungible resources such as sleep, health, safety, or time should increase λ more nonlinearly than equivalent monetary depletion when other buffers cannot compensate.*

Prediction 2 (Perceived boundary risk) *After controlling for objective wealth, agents who perceive themselves as closer to a constraint boundary should show higher loss aversion.*

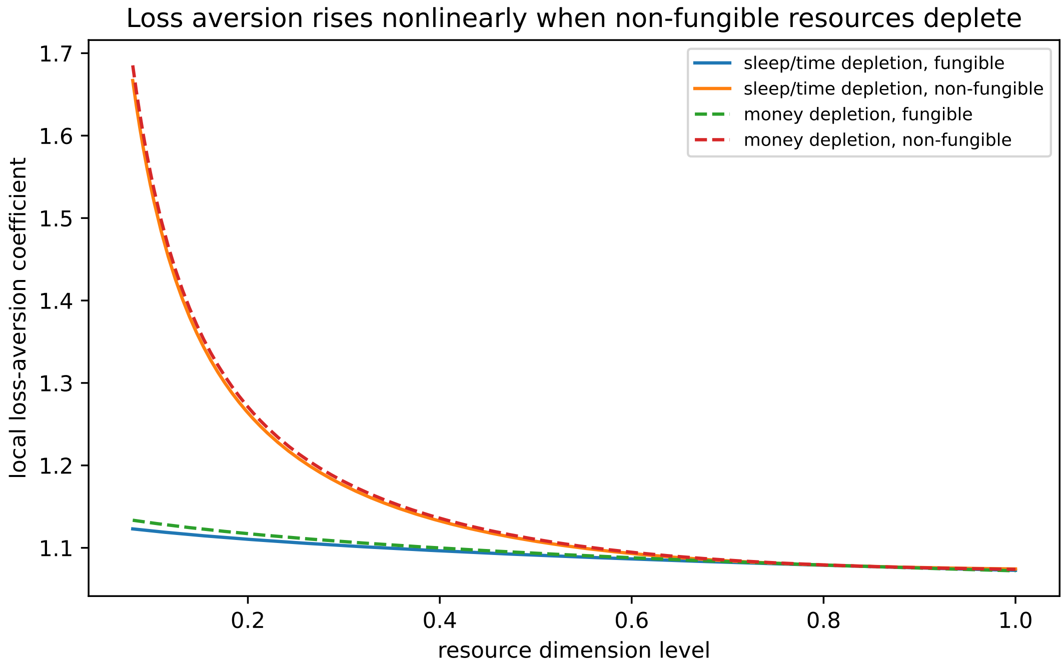

Figure 1: Loss aversion under fungible and non-fungible resource depletion. The simulation uses a CES effective buffer. When resources are weak substitutes, depletion of sleep/time generates sharper nonlinear increases in the local loss-aversion coefficient.

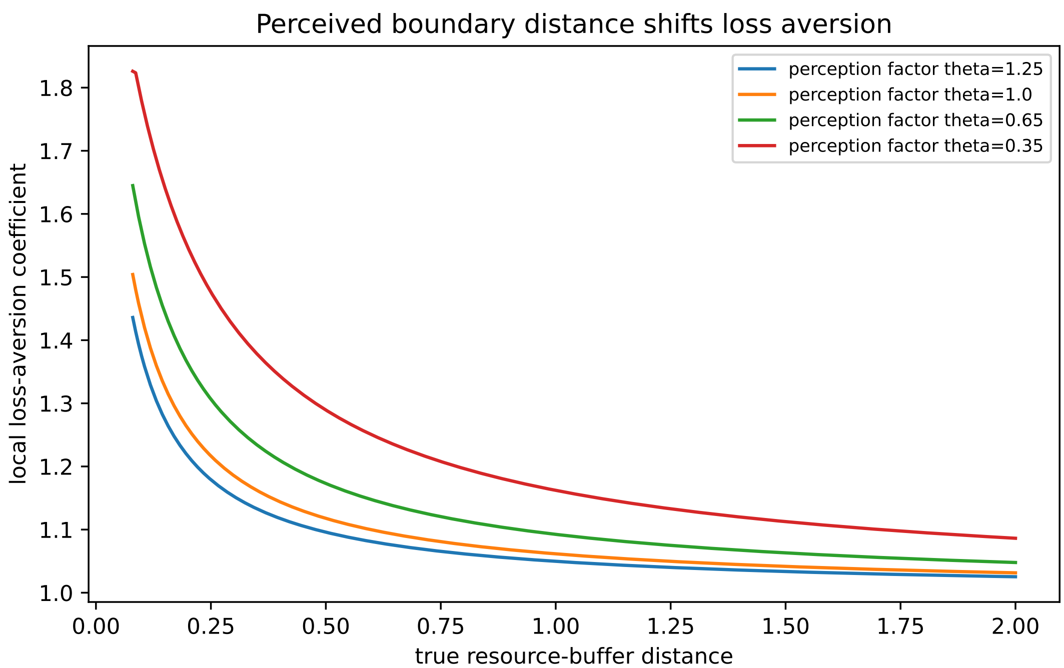

Figure 2: Perceived boundary distance shifts loss aversion. Lower perception factor θ makes the same objective resource buffer appear closer to the constraint boundary, increasing λ.

6 Reference Dependence as Maintained Baseline under Finite Updating

A finite agent does not compute utility over all possible terminal states. It maintains a compressed baseline r_t and evaluates local prediction errors:

$$\mathrm{PE}_t = x_t - r_t. \tag{12}$$

We model baseline updating as

$$r_{t+1} = r_t + \alpha_t(x_t - r_t) - \kappa_t P_t, \tag{13}$$

where α_t is adaptation rate, P_t is pruning or baseline-reset pressure, and κ_t is pruning strength. The finite-capacity hypothesis is that

$$\alpha_t = \alpha_{\min} + (\alpha_{\max} - \alpha_{\min})\sigma(a_0 + a_1 \log(\widehat{d}_t + \epsilon) + a_2 \log(C_t + \epsilon)), \tag{14}$$

where $\sigma(\cdot)$ is the logistic function. When resource buffer and decision capacity are high, reference points update quickly. When they are low, reference points rigidify.

Prediction 3 (Reference rigidity under depletion) *Reference-point adaptation should slow under stress, sleep deprivation, cognitive load, financial scarcity, administrative burden, and time pressure.*

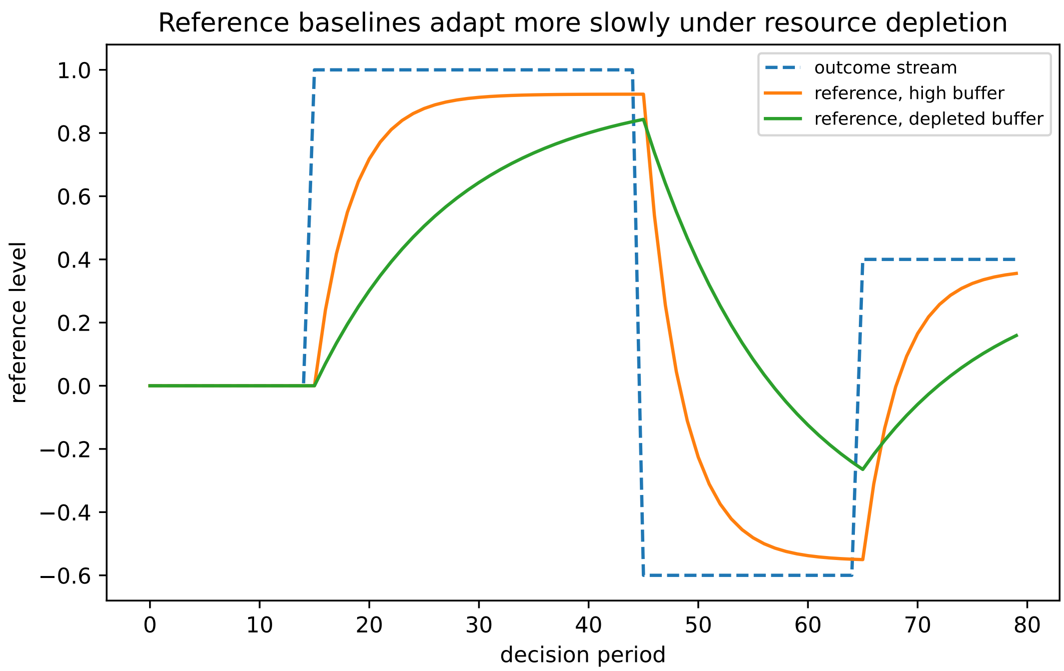

Figure 3: Reference baselines adapt more slowly under depletion. Two agents face the same outcome stream. The high-buffer agent updates quickly; the depleted agent maintains a stale baseline longer.

7 Probability Weighting as Finite Precision Allocation

Finite agents cannot estimate all probabilities with equal precision. Let $\tau(p, E)$ denote precision allocated to probability p for event class E. The event class matters: a one percent probability of a neutral lottery and a one percent probability of job loss, eviction, illness, or predation do not carry the same boundary relevance.

Suppose estimation error has expected cost $C_{\text{err}}(p, E, \tau)$ and precision is constrained:

$$\int_0^1 \tau(p, E)\, dp \le \tau_{\text{total}}(C_t). \tag{15}$$

The allocation problem is

$$\min_{\tau(p,E)} \int_0^1 C_{\text{err}}(p, E, \tau(p, E))\, dp \quad \text{s.t.} \quad \int_0^1 \tau(p, E)\, dp \le \tau_{\text{total}}. \tag{16}$$

Under standard smoothness assumptions, the Euler condition implies that precision is allocated toward regions where marginal error cost is high. A useful reduced-form expression is

$$\tau^*(p, E) \propto \sqrt{\frac{\partial^2 C_{\text{err}}(p, E)}{\partial p^2}}. \tag{17}$$

The decision weight can then be written as

$$w(p, E) = \frac{\int_0^p \tau^*(u, E)\, du}{\int_0^1 \tau^*(u, E)\, du}. \tag{18}$$

In reduced form, this generates event-class-dependent inverse-S weighting:

$$w(p; \gamma_{tE}) = \frac{p^{\gamma_{tE}}}{p^{\gamma_{tE}} + (1-p)^{\gamma_{tE}}}, \tag{19}$$

where lower γ_{tE} represents stronger distortion. We model

$$\gamma_{tE} = \gamma_{\min} + (\gamma_{\max} - \gamma_{\min})\sigma(g_0 + g_1 \log C_t - g_2 B_E - g_3/\widehat{d}_t). \tag{20}$$

Here B_E is boundary relevance of event class E. Low capacity, high boundary relevance, and low perceived distance all reduce γ and amplify tail weighting. This finite-precision interpretation is closely related to efficient-coding accounts of probability weighting, which predict that probability distortions should shift with the distribution of probabilities and prior beliefs encountered by the decision-maker (Frydman and Jin, 2023).

This mechanism explains why insurance and gambling can coexist. In both cases the agent overrepresents tail events. Insurance responds to rare catastrophic losses. Gambling can respond to rare escape opportunities when the current boundary state is poor. The common cause is not irrationality per se; it is finite precision applied to events that can change the boundary regime.

Prediction 4 (Event-class-dependent weighting) *Equally rare events should receive different decision weights when their boundary relevance differs. Health, housing, job loss, safety, and existential financial risks should generate stronger tail overweighting than abstract neutral lotteries.*

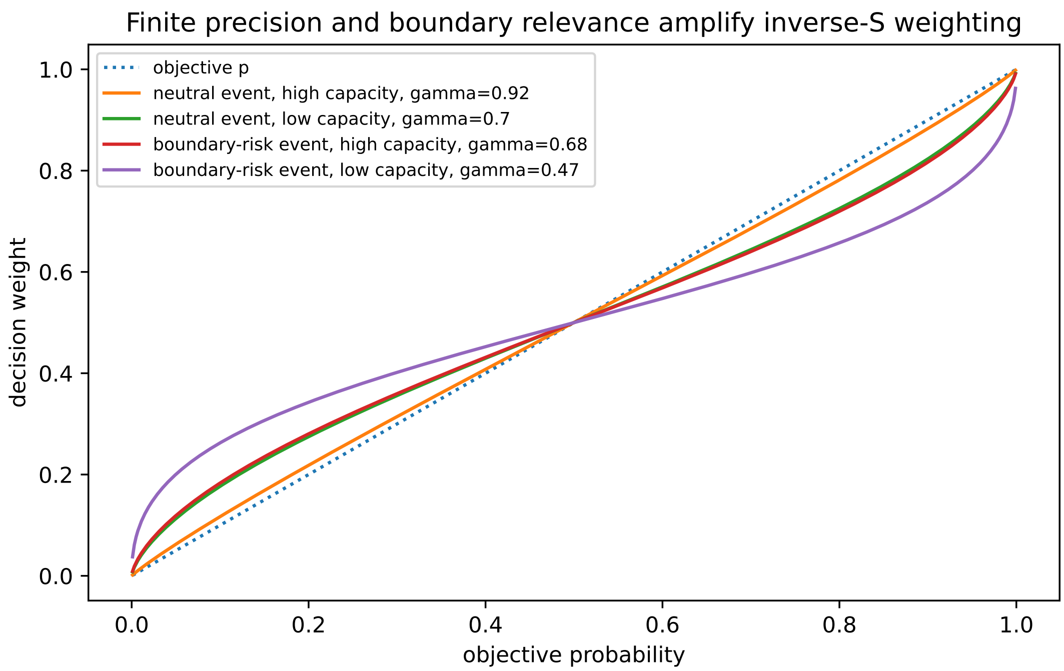

Figure 4: Finite precision and boundary relevance amplify inverse-S probability weighting. Low capacity and high boundary-risk event classes lower γ, increasing tail overrepresentation.

8 A Unified State-Dependent Prospect Model

Let a risky option yield outcomes x_i with probabilities p_i. Standard Prospect Theory evaluates

$$V = \sum_i w(p_i)v(x_i - r). \tag{21}$$

The finite-capacity version is

$$V_t = \sum_i w(p_i, E_i; C_t, \widehat{d}_t, B_{E_i})\, v(x_i - r_t; \lambda_t, a_t, b_t). \tag{22}$$

The key parameters are

$$\lambda_t = 1 + A(\widehat{d}_t + \epsilon)^{-\nu}, \tag{23}$$

$$\alpha_t = \alpha_{\min} + (\alpha_{\max} - \alpha_{\min})\sigma(a_0 + a_1 \log(\widehat{d}_t + \epsilon) + a_2 \log(C_t + \epsilon)), \tag{24}$$

$$\gamma_{tE} = \gamma_{\min} + (\gamma_{\max} - \gamma_{\min})\sigma(g_0 + g_1 \log C_t - g_2 B_E - g_3/\widehat{d}_t). \tag{25}$$

Expected utility is recovered as a limiting case when $F_t^{\text{eff}} \gg F_{\text{crit}}$, C_t is abundant, and boundary risk is low:

$$\lambda_t \to 1, \quad \gamma_{tE} \to 1, \tag{26}$$

while reference baselines adapt rapidly enough that local deviations approximate expected terminal values.

9 Numerical Simulations and Synthetic Identification

This section reports deterministic numerical simulations. They are not fits to real human data. Their purpose is to make the model operational, show qualitative implications, and provide an identification template for future empirical work. All code and CSV outputs are released with this manuscript.

9.1 Synthetic state-dependent parameter recovery

We simulate 1,200 agents with monetary, sleep/time, and attention resources. Effective buffer is generated with a low-substitutability CES aggregator and a random perception factor θ_i. True parameters are generated from the finite-capacity equations above and then measured with noise. We compare a static Prospect Theory model, which uses constant parameters, to a state-dependent finite-capacity model that regresses observed parameters on $\log(\widehat{d_i})$ and resource measures.

Table 2: Synthetic recovery comparison. The finite-capacity model is compared to a static-parameter model on simulated subject-level parameter estimates. Lower RMSE and AIC indicate better recovery.

Parameter	Model	RMSE	AIC	ΔAIC
lambda	static Prospect parameter	0.249	-3338.6	0.0
lambda	finite-capacity state-dependent	0.190	-3970.8	-632.2
gamma	static Prospect parameter	0.119	-5103.4	0.0
gamma	finite-capacity state-dependent	0.041	-7679.8	-2576.5
alpha	static Prospect parameter	0.058	-6825.8	0.0
alpha	finite-capacity state-dependent	0.027	-8648.4	-1822.6

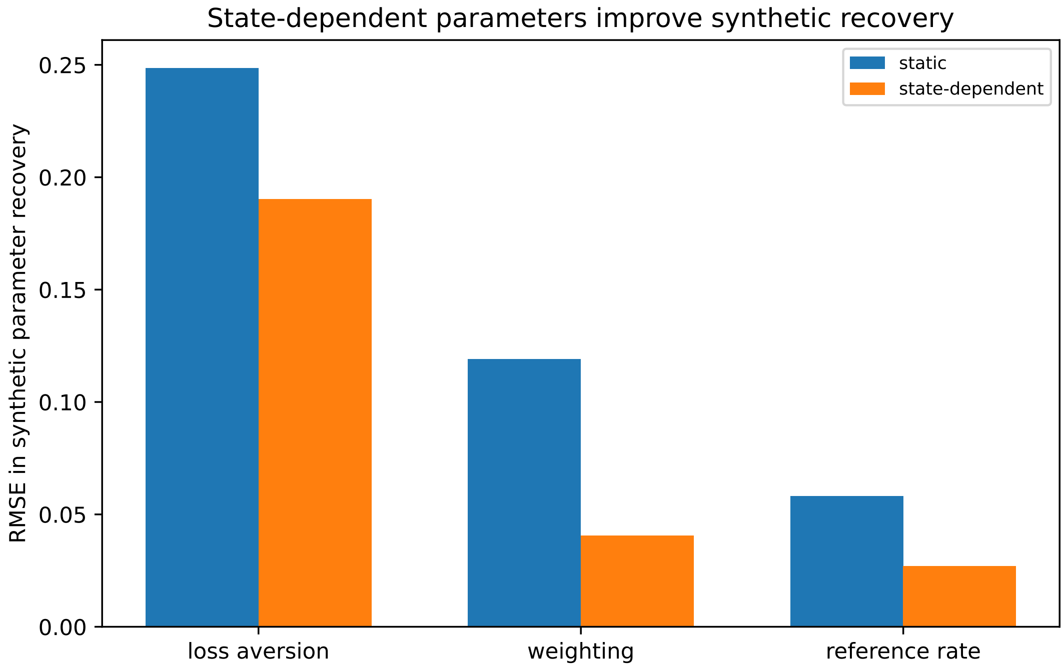

Figure 5: Static versus state-dependent parameter recovery in synthetic data. The finite-capacity specification improves recovery of loss aversion, probability weighting, and reference adaptation in the simulated environment.

9.2 Nudge bandwidth simulation

Choice architecture interventions are often described as low-cost. For a resource-depleted agent, however, a nudge can impose additional distinctions: compare alternatives, verify eligibility, process defaults, update plans, and overcome switching costs. Define nudge bandwidth requirement

$$C_t^{\text{nudge}} = I(\text{attributes}; \text{action}) + C_{\text{switch}} + C_{\text{verification}} + C_{\text{admin}}. \tag{27}$$

A nudge succeeds when available decision capacity exceeds the requirement:

$$C_t^{\text{available}} > C_t^{\text{nudge}}. \tag{28}$$

If $C_t^{\text{available}} < C_t^{\text{nudge}}$, the intervention may fail or backfire through overload.

Figure 6 is the policy-facing summary of the model. It shows that a choice-architecture intervention can be welfare-improving for agents with sufficient resource buffer, neutral for agents near the capacity boundary, and harmful for agents whose effective buffer is already depleted. In the finite-capacity view, a nudge is not a frictionless intervention. It consumes decision bandwidth. When the required nudge bandwidth exceeds the agent's available capacity, the intervention can fail or backfire.

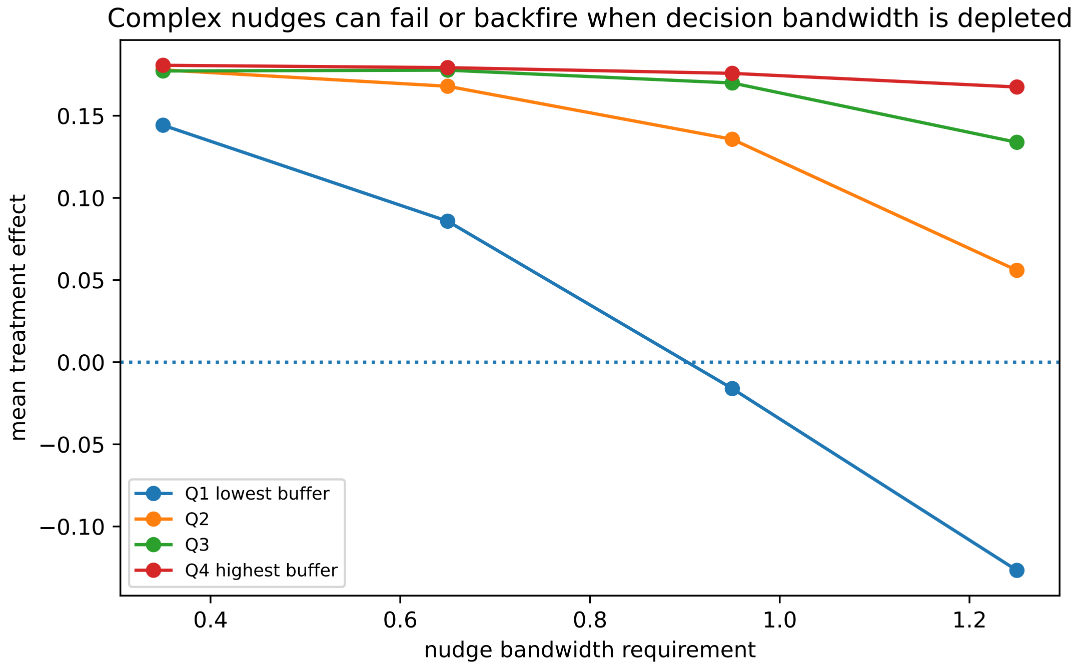

Figure 6: Nudge bandwidth and policy backfire. The simulation shows that a nudge improves outcomes only when the agent has enough available decision capacity to process the intervention. In low-buffer regimes, the same nudge can impose additional cognitive and administrative load, reducing net value. This illustrates the Buffer-First Principle: restore resource and decision bandwidth before imposing complex choice architecture.

10 Empirical Identification Strategy

The model is empirically useful only if F_t^{eff}, C_t, and $\widehat{d}_t$ can be proxied. Candidate measures are summarized in Table 3. Existing scarcity and poverty research already suggests that resource constraints affect cognition and decision quality (Mani et al., 2013; Haushofer and Fehr, 2014; Mullainathan and Shafir, 2013). Stress and affective regulation also affect risky choice (Porcelli and Delgado, 2009; Sokol-Hessner et al., 2009). Recent experiments on cognitive load and acute stress suggest that capacity constraints do not uniformly alter all preferences, but they can impair risky or computationally demanding decisions, especially when working memory, time pressure, or cortisol responses are involved (Dominiak and Duersch, 2024; Doroc, Yadav and Murawski, 2025). The finite-capacity model turns these findings into parameter-level predictions.

A reduced-form estimation strategy is

$$\lambda_i = \lambda_0 + \lambda_1 \widehat{d}_i^{-1} + \lambda_2 \text{NonFungibleDepletion}_i + \lambda_3 \text{DownsideProtection}_i + u_i, \tag{29}$$

$$\alpha_i = \alpha_0 + \alpha_1 \log C_i + \alpha_2 \log(\widehat{d}_i) + u_i, \tag{30}$$

$$\gamma_{iE} = \gamma_0 + \gamma_1 \log C_i - \gamma_2 B_E - \gamma_3 \widehat{d}_i^{-1} + u_{iE}. \tag{31}$$

A horse race against static Prospect Theory tests whether state variables improve out-of-sample prediction of choices or estimated parameters.

Heart-rate variability (HRV) is a particularly useful candidate proxy because it tracks autonomic flexibility and stress-regulation capacity in real time. In the FDS-E1 model, lower HRV should correspond to a smaller effective maintenance buffer or a smaller perceived distance from the boundary-risk threshold, predicting higher loss aversion, slower reference adaptation, stronger probability weighting distortion, and reduced nudge bandwidth. HRV is not the resource buffer itself; it is an autonomic window onto the maintenance state. As a non-invasive, continuous, and physiologically grounded measure, HRV can bridge between laboratory decision tasks and field studies of scarcity, stress, and overload.

Table 3: Behavioral parameters and finite-system interpretations.

Parameter	FDS interpretation	Observable proxies
λ	Boundary-risk slope asymmetry near resource threshold	Financial runway, downside exposure, cortisol, HRV, sleep deprivation, perceived scarcity
α	Reference-baseline update rate	Reversal learning, prediction-error integration, baseline updating after shocks
γ	Probability precision curvature	Probability-weighting estimates, numeracy, cognitive load, event class, tail-event salience
F^{eff}	Effective resource buffer	Income runway, liquid assets, sleep, health, time, social support, institutional insurance
C	Available decision capacity	Working memory, time pressure, attention, administrative burden, information complexity
B_E	Boundary relevance of event class	Health, job loss, eviction, violence, food insecurity, catastrophic financial loss, social exclusion
θ	Perceived distance mapping	Trauma history, uncertainty, belief state, social comparison, perceived institutional reliability

11 Policy Implications: The Buffer-First Principle

The model refines nudge theory (Thaler and Sunstein, 2008). A nudge is not free from the agent's perspective. It can require attention, comprehension, verification, switching, and memory. When an agent is depleted, an intervention that is simple for a high-capacity agent may exceed the low-capacity agent's available bandwidth.

Policy Corollary 1 (Buffer-First Principle) *Before imposing complex choice architecture on resource-depleted agents, policy should restore the resource and bandwidth conditions required for the agent to use the choice architecture.*

This principle favors cash transfers, debt relief, administrative simplification, default enrollment, automatic eligibility, shorter forms, stable scheduling, food security, housing stability, and reduction of time taxes before complex information campaigns. Recent reviews of choice architecture interventions emphasize substantial heterogeneity and the possibility of ineffective or counterproductive nudges; the nudge-bandwidth model provides one mechanism for such heterogeneity (Szaszi et al., 2025). It does not reject nudges. It states that nudges require bandwidth, and bandwidth is state-dependent.

Prediction 5 (Buffer-First Policy Corollary) *If a target population is operating near its boundary-risk threshold, interventions that add informational, administrative, switching, or verification demands should be preceded by buffer-restoring policies such as cash transfers, debt relief, default simplification, administrative burden reduction, food security, sleep protection, or housing stability. Otherwise, the intervention may consume the very decision capacity it requires.*

12 Relation to Existing Theories

Prospect Theory. The model does not replace Prospect Theory. It makes Prospect Theory parameters state-dependent. Because loss aversion and probability weighting are conceptually distinct (Blavatskyy, 2024), the model treats λ and γ as separate state-dependent parameters rather than reducing one to the other.

Expected utility. Expected utility becomes a limiting case of high resource buffer, high capacity, low boundary risk, and stable baselines.

Bounded rationality. FDS gives a maintenance interpretation of bounded rationality: agents economize not only on computation but also on boundary risk and update cost.

Active inference. Reference dependence is naturally related to prediction-error baselines and generative-model updating (Friston, 2010).

Ecological rationality. What looks like irrationality in modern financial markets may be an adaptation to environments in which resource buffers were small and rare boundary-changing events mattered.

13 Limitations and Falsification

This paper has several limitations. First, it does not derive $\lambda \approx 2.25$ exactly. Second, economic losses are not literal thermodynamic losses in any simple sense; they are symbolic claims on future resource access. Third, humans are not assumed to be globally optimal in modern environments.

Fourth, culture, institutions, learning, and market structure are not reduced to biology. Fifth, the model requires empirical proxies for F^{eff}, C, B_E, and θ.

The finite-capacity bridge is strongly weakened if the following reverse stress test fails:

> If severe resource depletion changes neither loss aversion, nor reference rigidity, nor probability weighting, nor decision noise after adequate measurement power and controls, then the finite-capacity bridge account is strongly undermined.

More specific demotion conditions are: (i) λ invariant across resource stress, physiological depletion, downside protection, and financial precarity; (ii) reference-point adaptation invariant across prediction-error integration, memory load, and cognitive load; (iii) probability-weighting curvature invariant across event class, cognitive load, and perceived consequence severity; (iv) nudge effectiveness independent of scarcity, fatigue, and administrative burden.

14 Conclusion

Prospect Theory showed that human beings are not ideal expected-utility maximizers. The finite-system view explains why they should not be. A physically realizable decision-maker has finite buffers, finite memory, finite attention, finite precision, finite update capacity, and finite tolerance for boundary risk. Under these constraints, loss aversion, reference dependence, and probability weighting are not arbitrary defects. They are candidate signatures of finite agents managing resource buffers, prediction baselines, and limited distinguishability under uncertainty.

The main contribution of this paper is a state-dependent bridge:

$$(\lambda_t, \alpha_t, \gamma_t) = f(F_t^{\text{eff}}, C_t, B_t, \theta_t). \tag{32}$$

This formulation suggests that poverty, stress, overload, and administrative burden are not merely psychological states. They are also finite-capacity maintenance states that change the effective parameters of decision-making. A finite agent cannot be nudged, priced, informed, or incentivized as if its decision capacity were infinite. Traditional economic models that ignore system finiteness may be elegant, but they risk describing behavior in a vacuum.

Code and Data Availability

The deterministic Python code, figures, and CSV files used in this manuscript are included in the accompanying replication package. Running `code/generate_results.py` regenerates all figures, CSV outputs, and the model-comparison table.

A Notation

Symbol	Meaning
F_t^i	Resource buffer in dimension i
F_t^{eff}	Effective resource buffer after substitutability aggregation
F_{crit}	Boundary-risk or constraint threshold
d_t	Objective distance from threshold
$\widehat{d}_t$	Perceived distance from threshold
θ_t	Perception or belief mapping from true to perceived distance
C_t	Available decision capacity
N_A	Number of operationally distinguishable decision states
λ_t	Loss-aversion coefficient
α_t	Reference-baseline adaptation rate
γ_{tE}	Probability-weighting curvature for event class E
B_E	Boundary relevance of event class E
C_t^{nudge}	Decision bandwidth required by a nudge

B Simulation Parameters

The simulations use fixed parameter values documented in `code/generate_results.py`. The boundary-risk potential is $R(d) = \kappa(d + \epsilon)^{-\nu}$ with default $\epsilon = 0.08$ and $\nu = 1.15$. CES effective buffers use weights on monetary, sleep/time, and attention resources. Low substitutability is represented by negative CES curvature. Probability weighting uses $w(p; \gamma) = p^\gamma/[p^\gamma + (1 - p)^\gamma]$. The nudge-bandwidth simulation generates a synthetic population with resource-dependent available capacity and nudge-complexity requirements.

C Reproducibility Checklist

1. The script is deterministic and uses a fixed random seed for synthetic population simulations.
2. All figures are generated from the source code.
3. All CSV files are regenerated from the source code.
4. The model-comparison table is generated from the synthetic parameter-recovery exercise.
5. No proprietary or human-subject data are used in this draft.

AI Assistance Disclosure

AI-assisted tools were used to improve readability, wording, and LaTeX formatting. The author reviewed and edited all content and remains responsible for all claims, references, simulations, and conclusions. No AI system is listed as an author.

References

Barberis, N. C. (2013). Thirty years of prospect theory in economics: A review and assessment. *Journal of Economic Perspectives*, 27(1), 173–196.

Friston, K. (2010). The free-energy principle: A unified brain theory? *Nature Reviews Neuroscience*, 11, 127–138.

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

Haushofer, J., and Fehr, E. (2014). On the psychology of poverty. *Science*, 344(6186), 862–867.

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

Mani, A., Mullainathan, S., Shafir, E., and Zhao, J. (2013). Poverty impedes cognitive function. *Science*, 341(6149), 976–980.

Mullainathan, S., and Shafir, E. (2013). *Scarcity: Why Having Too Little Means So Much*. Times Books.

Porcelli, A. J., and Delgado, M. R. (2009). Acute stress modulates risk taking in financial decision making. *Psychological Science*, 20(3), 278–283.

Prelec, D. (1998). The probability weighting function. *Econometrica*, 66(3), 497–527.

Shaffer, F., and Ginsberg, J. P. (2017). An overview of heart rate variability metrics and norms. *Frontiers in Public Health*, 5, 258.

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

Sokol-Hessner, P., Hsu, M., Curley, N. G., Delgado, M. R., Camerer, C. F., and Phelps, E. A. (2009). Thinking like a trader selectively reduces individuals' loss aversion. *Proceedings of the National Academy of Sciences*, 106(13), 5035–5040.

Blavatskyy, P. R. (2024). A behavioral definition of loss aversion. *Economics Letters*, 235, 111555. doi:10.1016/j.econlet.2024.111555.

Dominiak, A., and Duersch, P. (2024). Choice under uncertainty and cognitive load. *Journal of Risk and Uncertainty*, 68, 133–161. doi:10.1007/s11166-024-09426-6.

Doroc, K., Yadav, N., and Murawski, C. (2025). Acute stress impairs decision-making at varying levels of decision complexity. *Communications Psychology*, 3, 179. doi:10.1038/s44271-025-00355-x.

Frydman, C., and Jin, L. J. (2023). On the Source and Instability of Probability Weighting. NBER Working Paper No. 31573, revised April 2025. doi:10.3386/w31573.

Imai, T., Nunnari, S., Wu, J., and Vieider, F. M. (2025). Meta-Analysis of Prospect Theory Parameters. CESifo Working Paper No. 12334.

Szaszi, B., Goldstein, D. G., Soman, D., and Michie, S. (2025). Generalizability of choice architecture interventions. *Nature Reviews Psychology*, 4, 518–529. doi:10.1038/s44159-025-00471-9.

Thaler, R. H., and Sunstein, C. R. (2008). *Nudge: Improving Decisions About Health, Wealth, and Happiness*. Yale University Press.

Tversky, A., and Kahneman, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, 5, 297–323.

Wu, Y. (2026a). Finite distinguishability budgets and maintenance bounds for physical observers. Zenodo. doi:10.5281/zenodo.20234249.

Wu, Y. (2026b). Active finite distinction systems: A formal core for boundary maintenance under finite capacity. Zenodo. doi:10.5281/zenodo.20158923.