

Happiness Is All You Need

Arjun Nitin Barge*
Independent researcher
arjun.n.barge@gmail.com

26 September 2026

Abstract

“It can scarcely be denied that the supreme goal of all theory is to make the irreducible basic elements as simple and as few as possible without having to surrender the adequate representation of a single datum of experience.” —Albert Einstein [1].

I formulate happiness as a protected decision objective: maximize the worst expected standardized wellbeing surplus, then aggregate utility, within declared individual and resource constraints. I combine normalized observation updates with a finite safety algorithm that chooses one action for all states consistent with the available information. I prove conditional floor protection, bound sensitivity to model error, and state when a change of representation preserves outcomes and risk. Exact examples distinguish protection, repair and infeasibility; they also expose policies that improve an aggregate while worsening an individual’s risk. I examine 2,116 World Happiness Report records to test the measurement step: national margins can be computed, but an outcome-containing residual cannot validate a forecast, and a national mean cannot identify individual protection. My contribution is this specified composition of established methods, its conditional proofs and its explicit failure tests. Application requires justified individual measures, a causal action model and prospective evaluation; the aggregate data do not establish an intervention benefit.

1 From a reference point to a protected baseline

I begin with a stated outcome and a baseline. My aim is to include the last person in line in the definition of success: an improvement in the average should not silently expose an individual to greater risk. This requires three separate decisions: what I measure, which actions preserve the declared protection, and how I rank admissible policies. A boundary is useful when it protects a justified floor; it is restrictive when it removes a feasible improvement without serving that protection.

I use happiness as the purpose of this decision rule. The title expresses that priority. Measurement, a causal model of actions, feasible resources and participation determine what the rule can accomplish. I do not infer them from the name of the objective.

I combine a finite probabilistic update, a two-stage robust welfare criterion and an observation-aware protection algorithm. The components draw on quantum instruments and completely positive maps [2, 3], partially observed control [4, 5], feedback [6] and multiple-prior decisions [7, 8]. My contribution is their explicit composition and its tests; the underlying methods remain credited to their published foundations.

*Codex was used for research, analysis code, mathematical checks, drafting and editing.

The argument proceeds from objective to protection, then to information loss and bounded model error. Exact examples test both successful decisions and failures of plausible alternatives. The World Happiness Report then supplies an observed aggregate measure with which to examine the measurement step. The final sections identify the individual and causal information still required for application.

I distinguish proof under stated assumptions, calculation under stipulated dynamics, and description of observed data. Each supports a different claim. No participant received an intervention selected by this criterion. Figures retain parameters, units and labels so that the role of each quantity is visible without relying on color.




Figure 1: My proposed decision loop, using the state and instrument in Eqs. (1)–(3), objectives in Eqs. (4)–(6), and risk constraint (7). The baselines b_i , positive comparison scales s_i , outcome definitions and admissible actions are declared inputs. Spectral colors distinguish mathematical roles. I rank admissible policies before acting. The instrument then predicts observation branches; the next observed record selects the conditional update and begins the next cycle. The diagram specifies a decision process; its arrows do not represent physical forces.

2 Finite model and decision criterion

2.1 State, observation and causal ordering

Let X , A and O be finite classical state, action and observation sets. Let $\mathbf{H}$ be a finite-dimensional Hilbert space. At decision step n , write the classical–quantum state as

$$\Omega_n(x) \succeq 0, \quad \sum_{x \in X} \text{Tr } \Omega_n(x) = 1. \quad (1)$$

For positive $p_n(x) = \text{Tr } \Omega_n(x)$, this is $\Omega_n(x) = p_n(x)\rho_n(x)$. It retains classical–quantum correlations. It does not identify a human belief with a physical quantum state.

For each old state x and action a , take completely positive maps $\mathcal{J}_{n,o,x' \leftarrow x}^a$ whose sum over (o, x') is trace preserving. The unnormalized branch, its probability and its conditional state are

$$\tilde{\Omega}_{n+1}^o(x') = \sum_x \mathcal{J}_{n,o,x' \leftarrow x}^{a_n}[\Omega_n(x)], \quad (2)$$

$$q_n(o) = \sum_{x'} \text{Tr } \tilde{\Omega}_{n+1}^o(x'), \quad \Omega_{n+1}(x') = \tilde{\Omega}_{n+1}^o(x')/q_n(o). \quad (3)$$

The division applies only when $q_n(o) > 0$. The history $\mathcal{H}_n$ contains available observations, past actions, timestamps and delays. Policies are causal: $a_n = \pi_n(\mathcal{H}_n)$. The instrument specifies the actual ordering of dynamics and measurement. A zero-probability record under every retained model rejects that model set; it is not repaired by dividing by zero.

Proposition 1 (Finite-history normalization). *Under the preceding assumptions, every positive-probability branch has a positive, normalized conditional state. Every causal policy induces a normalized probability law on finite observation histories.*

Proof. Complete positivity preserves positivity term by term in (2). Summing branch traces gives

$$\sum_{o,x'} \text{Tr} \tilde{\Omega}_{n+1}^o(x') = \sum_x \text{Tr} \left[\sum_{o,x'} \mathcal{J}_{n,o,x' \leftarrow x}^{a_n} [\Omega_n(x)] \right] = \sum_x \text{Tr} \Omega_n(x) = 1.$$

Normalizing a positive-probability branch preserves positivity and gives unit total trace. Induction, conditioning each action on its available history, proves the history statement. $\square$

For $\dim \mathbf{H} = 1$, taking $\mathcal{J}_{o,x' \leftarrow x}^a(r) = T_a(x' | x)M(o | x')r$ recovers classical transition followed by Bayesian observation. Thus the construction includes classical controlled stochastic models. The protection algorithm and allocation examples below use this classical specialization. The more general instrument notation supplies no empirical claim of a quantum advantage.

2.2 Baseline priority and individual protection

Fix a decision history, a finite population $i = 1, \dots, N$ with $N \geq 1$, and a finite horizon $K \geq 1$. For person i and future step $k = 1, \dots, K$, let $h_{i,n+k}$ be an independently defined outcome, $b_{i,n+k}$ its justified baseline and $s_i > 0$ a comparison scale with the same units as $h_i - b_i$. They are inputs, not inferred universal constants. A nonempty set $\mathcal{P}_n$ contains retained causal models conditioned on that same history. Models assigning the discrete history zero probability are omitted explicitly. Outcomes and utilities are assumed uniformly bounded over retained models and admissible policies.

Define

$$\eta_n(\pi) = \min_{i,k} \inf_{P \in \mathcal{P}_n} \mathbb{E}_P^\pi \left[\frac{h_{i,n+k} - b_{i,n+k}}{s_i} \middle| \mathcal{H}_n \right], \quad (4)$$

$$J_n(\pi) = \inf_{P \in \mathcal{P}_n} \mathbb{E}_P^\pi \left[\sum_{k=1}^K \gamma_k \sum_i w_i u_i(h_{i,n+k}) \middle| \mathcal{H}_n \right], \quad (5)$$

$$\pi_n^* \in \text{lexargmax}_{\pi \in \Pi_{\text{adm}}} (\eta_n(\pi), J_n(\pi)). \quad (6)$$

Weights w_i, γ_k are nonnegative; utilities u_i are declared nondecreasing functions. Equation (6) first maximizes the worst expected standardized surplus, then maximizes aggregate utility among first-stage maximizers. This is a two-stage lexicographic criterion, not full leximin ordering of all individuals. Changing levels to logarithms changes the objective and is not a neutral numerical transformation.

Admissibility must encode resources, physical restrictions, participation, consent and any required risk protection. In particular, for $A_i = \{\exists k : h_{i,n+k} < b_{i,n+k}\}$, a specified protection rule may require

$$\sup_{P \in \mathcal{P}_n} [\mathbb{P}_P^\pi(A_i | \mathcal{H}_n) - \mathbb{P}_P^{\pi_0}(A_i | \mathcal{H}_n)] \leq 0 \quad \text{for every } i. \quad (7)$$

The reference policy π_0 is a future continuation from the same conditioned history with the same evaluation model, horizon and stated resource comparison. Equation (7) requires risk not to increase relative to the reference; it does not require that risk to be zero. A joint chance constraint $\sup_P \mathbb{P}_P^\pi(\bigcup_i A_i | \mathcal{H}_n) \leq \alpha$, with $0 \leq \alpha \leq 1$, is another possible restriction. Neither follows from improved expected welfare. A guarantee based on an omitted person or a mismeasured need does not protect that person.

If agents can deviate, a feasible centralized allocation also needs participation and incentive conditions under the specified information structure. Maximizing (6) alone does not prove equilibrium cooperation. Historical reciprocity models provide conditional mechanisms, rather than universal cooperation guarantees [10].

Proposition 2 (Existence and representation limits). *For a finite nonempty admissible class of deterministic history policies and finite robust objective values, (6) has a solution. An injective recoding of histories, with its inverse retained and all policies and models transported consistently, preserves the attainable objective and risk values. A positive affine change $h'_i = a_i h_i + c_i$, $b'_i = a_i b_i + c_i$, $s'_i = a_i s_i$ preserves η .*

Proof. A real-valued function on a finite nonempty set attains its maximum. Restrict to the maximizers of η and maximize J there. Under injective recoding, decoding gives a bijection of original and recoded policies with identical history probabilities and outcomes. Finally $(h'_i - b'_i)/s'_i = (h_i - b_i)/s_i$ for $a_i > 0$. Preserving J as well requires the utilities to be transported consistently. $\square$

An empty feasible set must be reported. Existence does not imply efficient optimization or exclude better randomized policies. Enlarging the conditional model set at fixed history cannot improve any fixed policy's worst-case surplus. Repeated reoptimization is not automatically dynamically consistent: general robust Bellman claims require additional conditions such as appropriate rectangularity [8]. The finite search below makes no unrestricted robust Bellman claim.

3 Protection with incomplete observations

3.1 An executable protection rule under incomplete observation

My objective is to improve the worst expected standardized surplus while retaining each person's declared floor. I therefore calculate which actions preserve protection, then rank policies restricted to those actions by (η, J) . Observation, perturbation and resource restrictions enter this calculation explicitly. The construction below is a finite classical specialization of the preceding update; it is not a discretization theorem for a continuous quantum state.

Let X, A, O be finite, with $X \neq \emptyset$. Let $A(x) \subseteq A$ be the physically and socially admissible actions at state x . For $a \in A(x)$, let $R(x, a) \subseteq X$ be a nonempty set of possible next states. An adversary may choose any permitted successor at each step, after my action. A deterministic sensor returns $g(x) \in O$. A noisy sensor with specified finite-state transition support can be included by augmenting the state with its realized output and defining the joint transition support explicitly. The current information set $B \neq \emptyset$ contains every state compatible with the available history and model support. Any observation available before the first action has already been incorporated into B ; the next sensor output arrives after the transition. I use a fixed safety set, with all outcome and floor variables included in the state when needed. I may choose only an action in

$$A(B) = \bigcap_{x \in B} A(x).$$

After action a and next observation o , the information update is

$$\Phi(B, a, o) = \{x' \in X : g(x') = o, \exists x \in B \text{ with } x' \in R(x, a)\}. \quad (8)$$

An observation with an empty update is impossible under this support model and is omitted from the recursion. Let $\mathcal{C} \subseteq X$ encode the declared physical, resource and participation restrictions, and

define $S = \{x \in \mathcal{C} : h_i(x) \geq b_i \ \forall i\}$. On the finite collection of nonempty subsets of X , put

$$\begin{aligned}\mathcal{V}_0 &= \{B : \emptyset \neq B \subseteq S\}. \\ \mathcal{V}_{j+1} &= \{B \in \mathcal{V}_0 : \exists a \in A(B), \forall o \text{ with } \Phi(B, a, o) \neq \emptyset, \\ &\quad \Phi(B, a, o) \in \mathcal{V}_j\}.\end{aligned}\tag{9}$$

This is a finite safety-game construction, grounded in minimax reachability and decision making with incomplete information [11, 12]. It implements the requirement that one action must protect all currently indistinguishable possibilities.

Proposition 3 (Protection, information and welfare selection). *Under these assumptions, $B \in \mathcal{V}_j$ if and only if a deterministic observation-adapted policy keeps the current state and the next j states in S for every initial state in B and every permitted successor sequence. The descending recursion stabilizes after at most $2^{|X|} - 1$ strict deletions of sets. Its fixed point $\mathcal{V}_*$ consists exactly of information sets admitting indefinite protection. At each $B \in \mathcal{V}_*$, define*

$$A_*(B) = \{a \in A(B) : \Phi(B, a, o) \in \mathcal{V}_* \text{ for every nonempty update}\}.$$

This action set is nonempty. Every policy that always selects from $A_*(B)$ preserves every floor. Consequently, any nonempty finite class of such policies has a lexicographic maximizer of (η, J) on a finite evaluation horizon; its individual and joint floor-failure probabilities are zero, and $\eta \geq 0$, under every probability law supported on the stated transitions and initial information set.

Proof. For $j = 0$, protection is equivalent to $B \subseteq S$. Assume the characterization at j . An observation-adapted policy must first choose the same admissible action for every $x \in B$. For each possible next observation, the compatible successor states are exactly $\Phi(B, a, o)$. The induction hypothesis gives a protective continuation precisely when each nonempty successor belongs to $\mathcal{V}_j$. The continuations may depend on the observed o , so their existence is also sufficient. This proves the recursion and the finite-horizon assertion.

The predecessor operator is monotone. Since $\mathcal{V}_1 \subseteq \mathcal{V}_0$, induction gives $\mathcal{V}_{j+1} \subseteq \mathcal{V}_j$. There are at most $2^{|X|} - 1$ nonempty information sets, so a fixed point is reached after at most that many strict deletion rounds. Each retained set has a witnessing action, and choosing one whenever that set occurs keeps all future information sets in $\mathcal{V}_*$. Conversely, an indefinitely protective policy supplies a protective prefix for every finite j , so its initial information set belongs to every $\mathcal{V}_j$ and hence the fixed point. A protected path has nonnegative standardized surplus at every included time. Supported probability laws therefore give zero failure probability and nonnegative expected surplus. Finite nonempty policy selection then follows by two successive maximizations, first of η and then of J . $\square$

A finite-horizon policy selected by this rule has a protective continuation: after the evaluated horizon I may apply a fixed witnessing action at each retained information set. Recomputing welfare preferences does not remove protection provided every later action remains in A_* . This safety statement needs no claim of dynamically consistent robust utility optimization. A union of transition supports from several probability models can admit switching paths absent from any single model. The resulting support test is exact for its declared adversarial model and can be conservative for the original family of fixed probability laws.

Algorithm and cost. I enumerate the nonempty information sets, remove those containing an unsafe state, and repeatedly delete a set with no common admissible action whose possible observation updates are all retained. I save the witnessing actions when deletion stops, reject an initial set outside $\mathcal{V}_*$, and rank candidate protected policies by (η, J) . The information state can have $2^{|X|} - 1$ values. Termination is a computability result; it does not promise efficient scaling. An empty protected set is a reportable result, not a reason to silently lower anyone's floor.

Why individual state tests are insufficient. Consider safe states L, R and an unsafe absorbing state $\perp$. Action a keeps L safe but sends R to $\perp$; action b keeps R safe but sends L to $\perp$. Each exactly observed safe state has a protective action. If L and R are initially indistinguishable, neither action protects $B = \{L, R\}$. A sensor that identifies the initial state makes protection feasible. This is a mathematically precise role for better observation; an invertible relabeling of the same observations cannot provide it.

4 Representation, information and changes of scale

4.1 Phase coordinates and observation checks

For a declared phase in radians, the six coordinates $(\sin \theta, \cos \theta, \tan \theta, \cot \theta, \sec \theta, \csc \theta)$ describe a one-dimensional curve wherever every entry is defined. Reciprocal coordinates introduce singularities. The bounded pair $(\sin \theta, \cos \theta)$ identifies phase modulo 2π ; an unwrapped trajectory additionally requires its winding count. Numerical use of a reciprocal chart requires a domain bounded away from its zero denominator. For my decision rule, an encoding must retain the distinctions used by the policy and observation likelihood. Section 6.4 tests this directly with an invertible encoding of four observations.

Identical marginals likewise do not identify a joint law: two fair bits can be perfectly correlated or anticorrelated, giving XOR identically zero or one. For the deterministic doubling map $x \mapsto 2x \bmod 1$ with uniform initial phase, binary half-interval observations have probability 2^{-m} for each length- m word. Exact interval checks cover the 510 words of lengths one through eight. The symbolic finite-record law equals that of independent fair bits. This illustrates a specific identifiability limit, not the indistinguishability of every chaotic process and every stochastic model. Chaos must be established for a specified dynamics [13].

4.2 Information has a declared probability space

For finite classical variables X, O , I use Shannon entropy in nats [14],

$$H(X) = -\sum_x p(x) \ln p(x), \quad I(X; O) = \sum_{x,o} p(x, o) \ln \frac{p(x, o)}{p(x)p(o)}. \quad (10)$$

Zero-probability summands are zero. These quantities require a specified joint law. They measure uncertainty in that declared law; rescaling arbitrary scores to sum to one does not recover an unobserved population distribution.

Proposition 4 (Information preserved and information lost). *For a deterministic map $Y = c(X)$ of finite variables,*

$$I(X; O) = I(Y; O) + I(X; O \mid Y) \geq I(Y; O). \quad (11)$$

A bijective relabeling preserves mutual information. Also $\mathbb{E}_O[H(X \mid O = o)] \leq H(X)$ for the same joint law.

Proof. Because Y is determined by X , the chain rule for the finite sums gives $I(X, Y; O) = I(X; O) = I(Y; O) + I(X; O | Y)$. The conditional term is an average of nonnegative relative entropies, which are nonnegative by the log-sum inequality. If c is bijective, X is determined by Y and the conditional term is zero. Finally $H(X) - H(X | O) = I(X; O) \geq 0$. $\square$

The last inequality is an average over observations, not a promise that each observation reduces uncertainty. It also does not compare different times of an evolving system without specifying the dynamics. More information can improve the available decision set; information alone does not set wellbeing, preferences or feasible resources.

4.3 When a change of scale preserves the decision problem

I connect different descriptions through a specified map, not through matching vocabulary. Let $c : X \rightarrow Y$ be a surjective map between finite state spaces. Suppose the controlled transition kernel $T_a(x' | x)$ obeys the action-wise block condition

$$\sum_{x': c(x')=y'} T_a(x' | x) = \bar{T}_a(y' | c(x)) \quad \text{for all } x, y' \text{ and admissible } a. \quad (12)$$

Require the initial aggregate law $\bar{\mu}(y) = \sum_{x: c(x)=y} \mu(x)$. Also require admissible actions, every person's outcome and floor, and every scored utility to be functions of $c(x)$, with the same values, comparison scales, person weights and time weights in both descriptions. These are substantive assumptions; an average that erases an individual deficit does not meet them. The condition is the controlled finite version of exact Markov aggregation [15].

Proposition 5 (Exact preservation under declared aggregation). *Under these conditions, every policy adapted to the aggregate history has exactly the same aggregate trajectory law, welfare scores and individual failure probabilities when run on the fine model or the aggregate model. Hence the attainable (η, J) and risk values agree for the same class of aggregate-history policies, provided the two retained families use a common index set and the block and initial-law conditions hold for every corresponding pair.*

Proof. Given any fine state in block y and a fixed action a , Eq. (12) gives the same distribution of the next block. Conditional averaging over the unobserved fine states therefore leaves that distribution unchanged. Induction on the aggregate history shows equality of all aggregate path probabilities under the same policy. Every stated score and failure event factors through that path, so its distribution, expectation and robust evaluation are equal. $\square$

This proves preservation for the stated policy class; it does not claim that arbitrary compression preserves all fine-state strategies. When the block condition fails, a memoryless aggregate kernel can give a false trajectory law. A higher-level model must then retain additional state, memory or a quantified approximation error. These conditions give my cross-scale proposal a falsifiable mathematical meaning. Each application must supply its own state, dynamics and measured outcomes.

Compression also requires an information test. For deterministic transitions $A \rightarrow A$, $B \rightarrow C$, $C \rightarrow C$, merging A and B into one observed class loses predictive state information. Starting half in A and half in B , a fitted one-step two-class transition predicts class probabilities $(1/4, 3/4)$ at step two, while the true probabilities are $(1/2, 1/2)$. A reduced memoryless state is not automatically closed; projection methods explicitly motivate memory terms [16].

5 Perturbations, coupled feedback and equilibrium

5.1 How much model error can a protection claim tolerate?

An exact support guarantee has a sharp limitation: a previously excluded harmful transition can invalidate it. For a separate quantitative result in the finite classical specialization, fix the same causal policy and initial law in a nominal model $\hat{P}$ and an actual model P . Assume that, at every common history and action, their joint next-step laws for the full trajectory variables, including latent state and observation, differ in total variation by at most $\epsilon \in [0, 1]$. Here total variation is $\|p - q\|_{\text{TV}} = \frac{1}{2} \sum_z |p(z) - q(z)|$ on the common finite next-step space. The bound applies at every full history shared by the two models, including latent states; the causal policy still uses only its observed history. A bound only on observed marginals is insufficient. Define

$$d_K = 1 - (1 - \epsilon)^K \leq \min\{1, K\epsilon\}. \quad (13)$$

Proposition 6 (Finite-horizon model-error budget). *For any event E determined by the first K transitions,*

$$|P^\pi(E) - \hat{P}^\pi(E)| \leq d_K.$$

For any path score Z lying in an interval of width W , the expectation discrepancy is at most Wd_K . For families, use a common nonempty index set Θ with $\mathcal{P} = \{P_\theta : \theta \in \Theta\}$ and $\hat{\mathcal{P}} = \{\hat{P}_\theta : \theta \in \Theta\}$. Suppose each pair has the same initial law, the stepwise bound holds uniformly in θ , history, action and policy, and all pairs use the same fixed outcome functions, floors, scales and admissible policy class. For each person and time, suppose the standardized surplus lies in the same bounded interval across both model families and all admissible policies, with every such interval having width at most W_g . Then $|\eta(\pi) - \hat{\eta}(\pi)| \leq W_g d_K$. Over the same finite admissible policy class, a nominal first-stage maximizer $\hat{\pi}$ satisfies

$$\eta(\hat{\pi}) \geq \max_{\pi} \eta(\pi) - 2W_g d_K.$$

In particular, a nominal failure probability at most $\alpha - d_K$ implies an actual failure probability at most α , whenever $\alpha \geq d_K$.

Proof. While the coupled histories agree, the common policy selects the same action. A maximal coupling of the two finite next-step laws agrees with probability at least $1 - \epsilon$. Inductively, the probability that every transition agrees is at least $(1 - \epsilon)^K$. An event's indicators can differ only on the complementary set, proving the probability bound. On that set two versions of a score in an interval of width W differ by at most W , giving the expectation bound. Uniform expectation errors are preserved under taking minima and infima. If π^* maximizes the actual first-stage objective, then

$$\eta(\hat{\pi}) \geq \hat{\eta}(\hat{\pi}) - W_g d_K \geq \hat{\eta}(\pi^*) - W_g d_K \geq \eta(\pi^*) - 2W_g d_K.$$

The failure budget follows from the probability inequality. $\square$

For example, $\epsilon = 0.01$ over ten steps gives $d_{10} = 0.095617925$. A zero-failure nominal model alone can then support only the stated upper bound of about 9.56% under this error allowance. These numbers are sensitivity calculations, not an estimated error rate for the happiness data. The second-stage utility selection can change discontinuously near a tie in η ; the bound does not establish a stable lexicographic ranking or a bound on second-stage utility loss.

5.2 What I mean by a happy equilibrium

For an autonomous deterministic closure with policy π , an equilibrium is a state x^* satisfying $F_\pi(x^*) = x^*$. Calling it baseline-respecting additionally requires $h_i(x^*) \geq b_i$ and the physical, resource and participation constraints for every included person. Existence and stability require their own assumptions; neither follows from the word “happiness.”

For a time-homogeneous Markov closure, the corresponding object is an invariant law $\nu = \nu K_\pi$, not the elimination of fluctuations. Its individual failure probabilities still need to satisfy the declared risk constraints. These definitions place equilibrium within the existing model; the population measurements below do not identify such a controlled equilibrium.

5.3 A protected equilibrium under coupled perturbations

The finite safety algorithm addresses allowed trajectories. A complementary sufficient condition makes the role of coupling and perturbation size explicit. Let F be a self-map of a nonempty complete product metric space Z , and suppose its component distances satisfy $d(Fz, Fz') \leq Gd(z, z')$ for a nonnegative matrix G with $\rho(G) < 1$. The convergent nonnegative series $(I - G)^{-1} = \sum_{m=0}^{\infty} G^m$ gives $w = (I - G)^{-1}\mathbf{1} > 0$. Since $Gw = w - \mathbf{1}$, the factor $q = \max_j (Gw)_j / w_j$ is less than one. The weighted metric $d_w(z, z') = \max_j d_j(z_j, z'_j) / w_j$ therefore makes F a contraction. Banach’s theorem yields a unique nominal fixed point z^* [17]. The metric scales and outcome units are declared inputs. This subsection’s ϵ bounds state-space distance; it is separate from the dimensionless probability error in Proposition 6.

Proposition 7 (A baseline-preserving perturbation tube). *Choose a radius $r > 0$ such that the closed ball $D = \{z : d_w(z, z^*) \leq r\}$ satisfies the declared resource and participation restrictions. Suppose every allowed perturbed update $F_\xi : Z \rightarrow Z$ obeys $d_w(F_\xi z, Fz) \leq \epsilon$ for all $z \in D$, where $0 \leq \epsilon \leq (1 - q)r$. If each outcome is L_i -Lipschitz in d_w on D and $h_i(z^*) - b_i \geq L_i r$, then every trajectory starting in D remains in D and satisfies every baseline, for every allowed perturbation sequence.*

Proof. For $z \in D$, the triangle inequality and $Fz^* = z^*$ give

$$d_w(F_\xi z, z^*) \leq \epsilon + q d_w(z, z^*) \leq \epsilon + qr \leq r.$$

Thus D is invariant. Induction and Lipschitz continuity give $h_i(z_n) \geq h_i(z^*) - L_i r \geq b_i$ at every step. $\square$

This is a sufficient condition, not a necessary one. It requires a verified bound for the connected system, a nominal fixed point with enough margin, and bounded perturbations. The gain of an isolated component does not establish these hypotheses for the connected system. The result does not eliminate fluctuations or claim that an arbitrary chaotic system is contractive.

5.4 Coupling and incentive counterexamples

If every map has the same contraction bound $q < 1$ in the same metric, distances between trajectories driven by the same sequence of maps are at most q times their preceding values at each step. The bounded-perturbation assumption of Proposition 7 alone does not impose this stronger condition or require convergence to a constant equilibrium. Two scalar components with isolated gain 0.5 can form the coupled matrix $\begin{pmatrix} .5 & .6 \\ .6 & .5 \end{pmatrix}$, whose eigenvalue 1.1 is unstable. Stable subsystem checks cannot certify the connected loop. Nor does stability imply a truthful model, resource feasibility or a desirable fixed point.

I test cooperation rather than assume that transparency ensures it. In a one-shot prisoner's dilemma with payoffs (3, 3) for mutual cooperation, (0, 5) or (5, 0) for unilateral defection, and (1, 1) for mutual defection, defection strictly dominates for each player even when all payoffs are public. Information and incentive design therefore need separate tests. The welfare criterion is a declared normative choice, not a theorem that all agents will adopt it.

6 Exact tests of the decision rule

6.1 Protection, repair and infeasibility

For an exactly observed finite state x , let

$$S = \{x \in \mathcal{C} : h_i(x) \geq b_i \text{ for every included person } i\}$$

where $\mathcal{C}$ contains the states satisfying the declared physical and participation restrictions. Suppose I choose $a \in A(x)$ before a disturbance $w \in W(x, a)$ and the next state is $F(x, a, w)$. Every admissible action has a nonempty disturbance set $W(x, a)$. The disturbance may take any permitted value at each step. I calculate

$$V_0 = S, \quad V_{k+1} = \{x \in S : \exists a \in A(x) \forall w \in W(x, a), F(x, a, w) \in V_k\}. \quad (14)$$

These are standard worst-case reachability and invariance constructions [11, 18]. By induction, $x \in V_k$ exactly when a state-feedback policy can keep the current state and the next k states in S against every permitted disturbance sequence. The descending sequence stabilizes because the state space is finite. Its fixed point V_∞ is the largest robust controlled-invariant subset of S ; a witnessing action at every retained state preserves S indefinitely in this model. The quantifier order matters: observing the disturbance before choosing the action defines a different problem.

This is the singleton-information specialization of Proposition 3. I use it here because every resource buffer is observed exactly.

I stress-test this distinction with two synthetic integer resource buffers $x_i \in \{0, 1, 2, 3\}$, floors $b_i = 1$, replenishments $u_i \geq 0$ with $u_1 + u_2 \leq r$, and a loss $w \in \{(1, 0), (0, 1)\}$. All actions are integer valued. I use

$$x_i^+ = \min\{3, \max\{0, x_i + u_i - w_i\}\}.$$

The clipping rule is part of this toy resource model: demand exceeding an empty buffer remains unmet; it does not imply that a human need has been satisfied. Among the nine states meeting both floors, the invariant sets contain 0, 8 and 9 states for budgets $r = 0, 1, 2$. At $r = 1$, the corner (1, 1) cannot be protected against every next loss, although the eight other safe states can. I verify the result by enumerating all $3^9 = 19,683$ stationary deterministic policies for that case. If joint loss (1, 1) is added to the allowed disturbances, the $r = 1$ invariant set becomes empty. Protection depends on the uncertainty set as well as the objective.

Repair is a different question from uninterrupted protection. Starting from $R_0 = V_\infty$, the recursion

$$R_{j+1} = R_j \cup \{x : \exists a \in A(x) \forall w \in W(x, a), F(x, a, w) \in R_j\}$$

identifies states from which I can guarantee entry into the protected set within $j + 1$ steps, while permitting interim deficits. Under the single-loss model, $r = 2$ gives layer sizes 9, 13, 15, 16: every state can be repaired within three steps. The repair trajectory must not be reported as having protected the original floor at every time. Additional state restrictions on the intermediate trajectory require their own constrained reachability calculation.




Figure 2: My exact buffer model: $x_i \in \{0, 1, 2, 3\}$, floors $b_1 = b_2 = 1$, integer replenishment $u_1 + u_2 \leq r$, and losses $w \in \{(1, 0), (0, 1)\}$ after the action. The plotted states already meet both floors; their future protection is computed by (14). Both axes show arbitrary buffer units; the vertical axis is person 2. Green circles are starting states from which an action policy keeps both people at or above the floor for every allowed disturbance sequence. Orange crosses cannot be protected indefinitely under that budget. With joint losses added, budget 1 loses all eight protected states; the assumed disturbance set matters.

Expected surplus alone is insufficient. At a one-step floor $(1, 1)$ with unit scales, a certain outcome $(1, 1)$ has $\eta = 0$. A lottery returning $(3, 3)$ with probability 0.9 and $(0, 0)$ with probability 0.1 has $\eta = 1.7$ and a 10% joint floor-failure probability. Without a risk restriction, my first-stage objective selects the lottery. Under a zero-failure restriction, it selects the certain feasible outcome. In the buffer model, independently assigning probability $1/2$ to each single loss gives a maximum achievable survival probability 2^{-K} at $(1, 1)$ when $r = 1$; exact closed-form checks cover all 432 budget, state and horizon combinations for $K = 0, \dots, 8$. A probability statement relies on this stated loss law; the invariant-set result does not.

A stable state can still be undesirable: $x_{n+1} = x_n/2$ converges to zero, below a positive floor. Conversely, a restriction that removes useful actions can shrink the protected set. With fixed dynamics, safety set and disturbances, enlarging $A(x)$ cannot shrink V_∞ , while enlarging the disturbance sets cannot enlarge it. I therefore distinguish a justified protective boundary from a restriction that blocks a feasible improvement. Consent, participation and meaningful exit must also be specified. A mathematical floor cannot substitute for them.

6.2 Discrete transfers and protected budgets

I can represent gifts restricted to multiples of a chosen amount as a constraint on actions. Let $N \geq 1$ recipients have known monetary shortfalls $r_i \geq 0$, all in the same currency unit. Assume one immediate transfer, no transaction costs and no uncertainty in receipt. A transfer g_i covers shortfall i exactly when $g_i \geq r_i$. Let $\mathcal{B} \geq 0$ be the available budget after protecting the provider's own required resources. For a denomination $d > 0$, require $g_i = dk_i$ with $k_i \in \mathbb{Z}_{\geq 0}$ and $\sum_i g_i \leq \mathcal{B}$. Before imposing that budget limit, the minimum total transfer covering every shortfall is

$$C_d = d \sum_{i=1}^N \left\lceil \frac{r_i}{d} \right\rceil, \quad \sum_{i=1}^N r_i \leq C_d < \sum_{i=1}^N r_i + Nd. \quad (15)$$

To verify this, each feasible integer k_i must be at least $\lceil r_i/d \rceil$. Choosing exactly those integers attains the stated cost. The bounds follow by summing $r_i \leq d \lceil r_i/d \rceil < r_i + d$. Therefore every

shortfall can be covered if and only if $C_d \leq \mathcal{B}$. This is an elementary integer-allocation result, not a new welfare law.

For $d = 18$, three shortfalls (10, 10, 10) and budget 36, unrestricted transfers cover all three for 30, whereas the denomination rule requires 54. If the three shortfalls instead equal 18, the minimum cost is 54 with or without that rule. Costs need not increase with every increase in denomination: a shortfall of 19 costs 36 at $d = 18$ but 20 at $d = 20$. The result depends on the shortfalls, budget and denomination together. Multiplying all monetary amounts by the same positive conversion factor leaves feasibility unchanged.

In a common finite action model with unchanged dynamics, observations and uncertainty, removing actions that violate a denomination rule cannot enlarge the protected set: every surviving protected policy was already available. For nonempty finite policy classes, the best pair (η, J) cannot improve in lexicographic order merely by restricting that class. I therefore place a denomination in admissibility and keep the welfare objective unchanged. A monetary shortfall is not a calibrated happiness deficit. Any effect of a giving convention on participation, trust or wellbeing requires measured outcomes and a separately justified transition model; the arithmetic above supplies no such effect.

6.3 Exact optimization of the protected objective

I now optimize Eqs. (4)–(6) themselves in a finite partially observed example. This calculation is separate from the four-step controller below, which optimizes only a one-step surrogate. Let $X = \{L, R, H, D\}$ have outcomes

$$h(L) = (2, 1), \quad h(R) = (1, 2), \quad h(H) = (2, 2), \quad h(D) = (0, 0), \quad b = (1, 1), \quad s = (1, 1).$$

The initial state is L or R . I retain both initial probabilities $P(L) = 1/4$ and $P(R) = 3/4$ as model uncertainty. The sensor gives the same observation for L, R , but identifies H and D . The horizon is two transitions. Utilities are $u_i(h_i) = h_i$, person weights are $(w_1, w_2) = (3, 1)$, and both time weights are 1. Unequal weights are explicitly chosen to test the second-stage tie-break; they are not estimated ethical constants. All four actions are admissible at every state before the protection restriction is applied. The actions have these transition laws, written as distributions over (L, R, H, D) :

Action	From L	From R	From H
a	(1, 0, 0, 0)	(0, 1/4, 3/4, 0)	(0, 0, 1, 0)
b	(1/4, 0, 3/4, 0)	(0, 1, 0, 0)	(0, 0, 1, 0)
c	(1/4, 1/4, 1/2, 0)	(1/4, 1/4, 1/2, 0)	(0, 0, 1, 0)
d	(0, 0, 9/10, 1/10)	(0, 0, 9/10, 1/10)	(0, 0, 9/10, 1/10)

State D is absorbing under every action. All dynamics are stipulated; the numbers are not measured human responses.

A deterministic two-step observation policy specifies its first action and one later action for each of the three possible observations, giving $4^4 = 256$ descriptions. Exact enumeration finds 108 descriptions with zero floor-failure probability in both retained models. Twelve equivalent descriptions maximize the protected objective:

$$(\eta, J) = \left(\frac{3}{4}, \frac{233}{16} \right) = (0.75, 14.5625), \quad \mathbb{P}(\text{any floor failure}) = 0.$$

They first choose c , choose a on the merged L/R observation, and choose any of a, b, c at H ; the action assigned to the unreachable D observation is immaterial. This policy uses the observations and has a safe continuation after the evaluated horizon.

For comparison, optimizing the same expected-welfare criterion over all 256 descriptions admits a risky first action d , giving

$$(\eta, J) = \left(\frac{4}{5}, \frac{72}{5}\right) = (0.8, 14.4), \quad \mathbb{P}(\text{any floor failure}) = \frac{1}{10}.$$

The larger expected-surplus minimum makes it win the first stage despite its risk. This is why my protection rule is part of admissibility, not an implication of the objective. The lexicographic rule also is not a complete ranking of every person’s entire surplus trajectory: its first stage maximizes the smallest expected surplus, and its declared second stage resolves the remaining ties.

6.4 A surrogate that fails individual protection

Two binary states indicate adequate (1) or inadequate (0) outcomes. They start independently with probability 1/2 of adequacy. At each of four steps, each current state is observed correctly with probability 4/5, independently conditional on state. One person can receive one resource unit. Given the current state, next-step adequacy probabilities are

Current state	Helped	Unhelped
Inadequate	0.80	0.10
Adequate	0.95	0.85

These beneficial transition effects are assumptions of a synthetic system. Conditional transitions are independent across people. The controller observes the current state noisily, updates its belief, chooses an action, then evaluates the next state. The final next state is evaluated without an additional controller observation.

The tested Bayesian controller maximizes the minimum next-step adequacy probability, then their sum; exact ties alternate by step. It is a one-step surrogate, not the full-horizon robust optimum in (6). A deliberately misspecified controller assumes observation accuracy 1/5 while the true generator remains 4/5.

Exact rational enumeration covers all $4^4 = 256$ terminal observation records per policy and 1,700 decision nodes across five policies. A separate backward evaluation reproduces every reported outcome. Adequacy totals count the eight post-action person-steps; failure means any inadequate post-action state, excluding the initial state. All active policies spend four units; the no-action policy spends zero.

Table 1: Exact synthetic outcomes, rounded to six decimals. Lower risk is better.

Policy	Total adequacy	Person 0 failure	Either-person failure
No action	3.610156	0.708291	0.914906
Blind alternating	6.342469	0.399422	0.781159
Correct Bayesian model	6.521994	0.421755	0.734955
Wrong observation model	5.666258	0.457918	0.827362
Correct, angular encoding	6.521994	0.421755	0.734955

The correct controller improves total adequacy and joint risk relative to alternating support, but raises person 0’s failure probability by 0.022333, or about 2.23 percentage points. Constraint (7), with blind alternation as π_0 , would reject it. The example establishes a specific failure of the one-step surrogate; it does not prove that every solution of (6) has that failure.

Encoding the four observations as $(\cos(\pi o/2), \sin(\pi o/2))$, $o = 0, 1, 2, 3$, and decoding before likelihood evaluation gives exactly the same beliefs, actions and outcomes. This tests Proposition 2 directly. A reversible angle map adds no information. The error from the wrong likelihood shows that correct normalization alone does not make beliefs accurate.




Figure 3: I can improve the aggregate while worsening an individual’s risk in the exact four-step example. Relative to alternating support, the correct one-step controller moves right and up: more adequacy, but more risk for person 0. The angular encoding coincides exactly with the correct-model point. The observation accuracy is $4/5$, the initial adequacy probabilities are $1/2$, and the helped/unhelped transition probabilities are stated in Section 6.4. These are stipulated transition dynamics, not patient outcomes.

6.5 Adversarial coverage

The current exact-arithmetic checks compare 648 finite safety games with independent policy-graph enumeration, 1,024 model-error cases with full trajectory probabilities, and 60 fine/aggregate trajectory comparisons. Every value for the 256-policy objective example also agrees with a separate matrix calculation. Perturbation-tube tests include equality at the allowed radius, an excessive disturbance and an insufficient baseline margin. These finite tests help detect implementation and transcription errors; the propositions rely on their displayed proofs. Assumptions omitted in the counterexamples are restored explicitly in the theorem statements.

7 Observed wellbeing and the baseline vector

7.1 What the aggregate observations contain

I use the World Happiness Report 2026 Figure 2.1 workbook to examine how an observed aggregate measure can be expressed relative to a declared baseline. The workbook contains 2,116 rows, 168 country or territory labels and 14 year labels spanning 2011–2025; 2013 is absent. These are published ranking windows, not an annual individual panel. The 2025 label denotes the 2023–2025 window. The earliest 2011 label corresponds to the first ranking’s 2005–2011 period and must not be treated as another three-year window [19, 20].

The latest window has 147 means and 147 reported marginal 95% intervals. Only 145 rows have a complete six-factor decomposition and residual. Tajikistan lacks the freedom contribution and residual; the State of Palestine lacks the healthy-life-expectancy contribution and residual. I preserve these missing values. An older 2018 block duplicates rank 120 for Gambia and Sri Lanka and omits rank 130. The original 2019 appendix gives Sri Lanka rank 130 [21]; I flag the workbook inconsistency rather than silently repair it. The latest rank labels pass the same check.

For country c and stated window τ , I retain the observation tuple

$$z_{c,\tau} = (\hat{h}_{c,\tau}, \ell_{c,\tau}, u_{c,\tau}, d_{1,c,\tau}, \dots, d_{6,c,\tau}, r_{c,\tau}, \mathbf{m}_{c,\tau}, \tau). \quad (16)$$

Here $\hat{h}$ is the observed mean on the 0–10 Cantril Ladder, $[\ell, u]$ is its reported interval, d_j are the published explanatory contributions, r is “Dystopia + residual,” and $\mathbf{m}$ records missing entries. I treat reported Ladder points as the declared numerical diagnostic; they do not establish an invariant interpersonal utility scale [19]. The contributions are model-based decompositions, not six separately measured individual needs or probabilities. This tuple specifies the measured input to a prospective observation model. To use it in the finite construction of (2), I would also need to declare a finite observation encoding and calibrate its likelihood. Neither that likelihood nor a causal response to action is estimated here. The national margins below are descriptive diagnostics, not fitted policy values.

7.2 Baseline sensitivity

For a declared national diagnostic floor b_c and a positive scale s_c in Ladder points, define

$$\hat{g}_c = \frac{\hat{h}_c - b_c}{s_c}, \quad [g_c^-, g_c^+] = \left[\frac{\ell_c - b_c}{s_c}, \frac{u_c - b_c}{s_c} \right]. \quad (17)$$

This is the observable counterpart of a standardized surplus, not an estimate of every individual’s surplus in (4). An interval whose lower limit meets a floor concerns the national mean under the survey’s assumptions. It is neither simultaneous 95% protection across countries nor a guarantee about the least-happy person. Dystopia is a regression reference; I do not adopt it as an ethical minimum.

Table 2: Latest-window sensitivity to explicitly hypothetical floors, $s_c = 1$ Ladder point. Intervals are classified by $\ell_c \geq b_c$ (at/above), $u_c < b_c$ (below), or $\ell_c < b_c \leq u_c$ (unresolved). Floors 4, 5 and 6 are diagnostic choices, not validated policy thresholds.

Floor	Mean at/above	Interval at/above	Below	Unresolved
4	134	132	9	6
5	101	99	44	4
6	71	60	71	16

At floor 5, two of the 101 mean-based passes have intervals crossing the floor. Scaling both the observations and the floor from 0–10 to 0–1 leaves these classifications unchanged when s_c is scaled consistently. Replacing s_c with a sampling-uncertainty width changes the ordering and purpose of the objective: measurement precision is not the same as need. In the latest ordered list, 143 of 146 adjacent pairs have overlapping marginal intervals. Overlap is a descriptive warning about reading too much into neighboring ranks; it is not a formal equality test.




Figure 4: Observed national life evaluation and a declared baseline. All 147 observations are retained. The figure is a comparison of means and uncertainty, not a population distribution or a causal response to intervention.

7.3 An accounting fit is not a forecast

The decomposition provides a direct test for outcome leakage. Summing the six contributions and the residual reconstructs the 145 complete means with RMSE 0.000871 Ladder points and maximum absolute difference 0.002. The residual already contains the unexplained part of the target. This is an accounting identity up to published precision, so using it as a predictor would leak the outcome. The resulting $R^2 = 0.999999416$ is not forecasting evidence. I do not fit a new forecast to the overlapping ranking windows.

8 From aggregate evidence to individual protection

8.1 A mean does not identify the person at risk

The preceding measurements cannot identify individual protection. A synthetic population in which everyone scores 7 and one in which 30% score 0 and 70% score 10 both have mean 7. At floor 5, their failure fractions are 0 and 0.3. Even exact knowledge of the mean would not distinguish them. This shows why Proposition 5 requires individual outcomes and floor events to be retained by the declared map. A national mean alone does not meet that requirement.

The counterexample connects the observed-data analysis to the policy tests. In Section 6.4, improved total adequacy concealed an increase in one person’s risk. Here the same observed mean conceals different individual failure fractions. In both cases, I must retain the individual outcome and its declared floor rather than substitute an aggregate score. A national confidence interval quantifies uncertainty about the mean; it does not supply the missing individual outcomes.

8.2 What an application must specify

Application of Eqs. (4)–(6) requires a defined population, individual outcomes and baselines, comparison scales and feasible actions. Protection defined relative to another policy additionally requires a reference continuation from the same history. The observation model must state who is

measured, who can be missed, and how observations relate to the latent state. Missing participants cannot be protected merely by assigning a better score to those observed.

The transition model must then describe the consequences of actions under stated causal assumptions [9]. Associations between national characteristics and life evaluation do not identify those consequences. The World Happiness Report workbook supplies neither individual action histories nor a calibrated response to this proposed policy. It therefore cannot determine the policy’s individual failure probabilities or validate the uniform model-error bound in Proposition 6.

A prospective evaluation would have to specify the outcome, horizon, reference policy and individual-risk endpoints before choosing a favorable result. It would also need to test observation coverage, action feasibility and sensitivity to model misspecification. These are requirements for evaluating the criterion in an application; no such intervention is evaluated here.

9 Scope of the guarantees and remaining limits

The mathematical guarantee is conditional on the declared model. For a finite support model with known observation rules, the protection recursion returns exactly the information sets from which one policy can preserve every floor against all permitted successor choices. Optimization over its saved actions then retains that protection. If the current information set is excluded, the algorithm reports infeasibility under those assumptions. Enlarging resources or changing the permitted actions is a different decision problem, and repair may allow deficits before protection is reached.

The finite information-set construction can require exponentially many sets in the number of latent states. An omitted transition can invalidate exact support protection. The model-error result instead requires its displayed uniform bound on joint transition and observation laws. The aggregation result requires matched initial laws, compatible actions and retention of individual outcomes. The perturbation result requires a verified contraction, bounded disturbances and a sufficient outcome margin. None of those conditions follows from a successful numerical example alone.

The observation analysis establishes descriptive national margins and exposes outcome leakage and loss of individual information. It establishes no causal effect of the proposed criterion. Baselines, scales, utilities and participation remain choices requiring justification with the affected people; the mathematics does not select a universal interpersonal measure. My next task is to specify and evaluate an individual observation-and-action model while retaining these limits explicitly.

10 Conclusion

I place individual protection before the ranking of feasible improvements. The criterion maximizes the worst expected standardized surplus, then aggregate utility; a separate admissibility rule enforces the declared protection. The finite algorithm identifies when indefinite floor protection is possible under its declared finite support model and incomplete observation. The proofs state which changes of representation preserve it and how model error can weaken it.

The tests complete that argument: expected improvement can conceal risk, a useful aggregate can lose an individual deficit, and an accounting fit can resemble prediction. The observed happiness data make the measurement limit concrete. The result is a specified decision architecture with conditional guarantees and explicit failure cases. Applying it to people requires the individual and causal evidence identified above. My aim remains constructive improvement with the last person in line included in the outcome being protected.

Closing reflection

“rhymes,primes all are sublime, leading love of iconoclasm to the top of the dharmic and karmic lines.”

—Arjun Nitin Barge

References

- [1] A. Einstein, *On the Method of Theoretical Physics: The Herbert Spencer Lecture, delivered at Oxford 10 June 1933* (Clarendon Press, Oxford, 1933), p. 9. Original scan: <https://dspace.bcu-iasi.ro/handle/123456789/14725>. Reprinted in *Philosophy of Science* **1**(2), 163–169 (1934), doi:10.1086/286316.
- [2] E. B. Davies and J. T. Lewis, “An operational approach to quantum probability,” *Communications in Mathematical Physics* **17**, 239–260 (1970). doi:10.1007/BF01647093.
- [3] M.-D. Choi, “Completely positive linear maps on complex matrices,” *Linear Algebra and its Applications* **10**, 285–290 (1975). doi:10.1016/0024-3795(75)90075-0.
- [4] R. D. Smallwood and E. J. Sondik, “The Optimal Control of Partially Observable Markov Processes over a Finite Horizon,” *Operations Research* **21**(5), 1071–1088 (1973). doi:10.1287/opre.21.5.1071.
- [5] J. Barry, D. T. Barry and S. Aaronson, “Quantum POMDPs,” *Physical Review A* **90**, 032311 (2014). doi:10.1103/PhysRevA.90.032311.
- [6] A. C. Doherty, S. Habib, K. Jacobs, H. Mabuchi and S. M. Tan, “Quantum feedback control and classical control theory,” *Physical Review A* **62**, 012105 (2000). doi:10.1103/PhysRevA.62.012105.
- [7] I. Gilboa and D. Schmeidler, “Maxmin expected utility with non-unique prior,” *Journal of Mathematical Economics* **18**(2), 141–153 (1989). doi:10.1016/0304-4068(89)90018-9.
- [8] G. N. Iyengar, “Robust Dynamic Programming,” *Mathematics of Operations Research* **30**(2), 257–280 (2005). doi:10.1287/moor.1040.0129.
- [9] J. Pearl, “Causal diagrams for empirical research,” *Biometrika* **82**(4), 669–688 (1995). doi:10.1093/biomet/82.4.669.
- [10] R. Axelrod and W. D. Hamilton, “The Evolution of Cooperation,” *Science* **211**(4489), 1390–1396 (1981). doi:10.1126/science.7466396.
- [11] D. P. Bertsekas and I. B. Rhodes, “On the minimax reachability of target sets and target tubes,” *Automatica* **7**(2), 233–247 (1971). doi:10.1016/0005-1098(71)90066-5.
- [12] J. H. Reif, “The complexity of two-player games of incomplete information,” *Journal of Computer and System Sciences* **29**(2), 274–301 (1984). doi:10.1016/0022-0000(84)90034-5.
- [13] E. N. Lorenz, “Deterministic Nonperiodic Flow,” *Journal of the Atmospheric Sciences* **20**, 130–141 (1963). [https://doi.org/10.1175/1520-0469\(1963\)020<0130:DNF>2.0.CO;2](https://doi.org/10.1175/1520-0469(1963)020<0130:DNF>2.0.CO;2).

- [14] C. E. Shannon, “A Mathematical Theory of Communication,” *Bell System Technical Journal* **27**(3), 379–423 (1948). doi:10.1002/j.1538-7305.1948.tb01338.x.
- [15] C. J. Burke and M. Rosenblatt, “A Markovian Function of a Markov Chain,” *The Annals of Mathematical Statistics* **29**(4), 1112–1122 (1958). doi:10.1214/aoms/1177706444.
- [16] H. Mori, “Transport, Collective Motion, and Brownian Motion,” *Progress of Theoretical Physics* **33**(3), 423–455 (1965). doi:10.1143/PTP.33.423.
- [17] S. Banach, “Sur les opérations dans les ensembles abstraits et leur application aux équations intégrales,” *Fundamenta Mathematicae* **3**, 133–181 (1922). doi:10.4064/fm-3-1-133-181.
- [18] F. Blanchini, “Set invariance in control,” *Automatica* **35**(11), 1747–1767 (1999). doi:10.1016/S0005-1098(99)00113-2.
- [19] J. F. Helliwell, L. B. Akinin, H. Huang, M. Rojas, S. Wang, V. Guerra and A. Danyluk, “International evidence on happiness and social media,” chapter 2 in *World Happiness Report 2026* (Wellbeing Research Centre, University of Oxford, 2026). doi:10.18724/whr-db0f-2b39. Figure 2.1 data and definitions: <https://www.worldhappiness.report/data-sharing/>. The workbook used here is WHR26_Data_Figure_2.1.xlsx; its observed transcription defect is described in Section 7.
- [20] J. F. Helliwell, R. Layard and J. Sachs, “Setting the Stage,” chapter 1 in *World Happiness Report 2015* (Sustainable Development Solutions Network, 2015). <https://files.worldhappiness.report/WHR15.pdf>. Describes the first ranking’s 2005–2011 observation period.
- [21] J. F. Helliwell, H. Huang and S. Wang, *Statistical Appendix 1 for Chapter 2, World Happiness Report 2019*, Figure 12, p. 25 (2019). https://files.worldhappiness.report/WHR19_Statistical_Appendix_01.pdf.