Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Preference-Class Gap and Pathwise Distortion: An Axiomatic Basis for Dynamic Discrepancy Theory

Takashi Matsuhisa

takashi.matsuhisa@krc.karelia.ru

IAMR, KarRC RAS

Research Article

Keywords: axiomatic characterization; welfare distance; positive affine invariance; dynamic decision theory; regret and distortion; uniform-norm approximation

Posted Date: March 16th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-9114502/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: The authors declare no competing interests.

Noname manuscript No.
(will be inserted by the editor)

Preference-Class Gap and Pathwise Distortion: An Axiomatic Basis for Dynamic Discrepancy Theory

Takashi Matsuhisa

Received: date / Accepted: date

Abstract We provide an axiomatic foundation for two indices intended to quantify *dynamic discrepancy* between alternative decision criteria under the canonical invariance principle of decision theory and social choice: welfare representations are meaningful only up to positive affine transformations. The first index, the *Preference-Class Gap* (PPG), measures *representation-level* mismatch between two instantaneous welfare primitives on a common compact admissible domain, by computing the uniform-norm distance from one primitive to the positive-affine equivalence class generated by the other, with a scale-free normalization ensuring well-posedness under near-degenerate benchmark scales. The second index, the *Pathwise Distortion Index* (PDI), measures *outcome-level* discrepancy along realized optimal trajectories: it aggregates the benchmark-evaluated welfare wedge induced by the alternative criterion and normalizes it in benchmark units, with a strictly positive floor regularization required for stability when baseline magnitudes become small.

For each index we propose transparent axioms and prove representation theorems: up to a monotone transform, PPG is uniquely characterized by the uniform-norm distance to a positive-affine class, and PDI is uniquely characterized by a floor-regularized normalized cumulative wedge along realized paths. We also establish a sharp logical separation: representation-level proximity and pathwise distortion are not equivalent objects. Small preference-class gaps can be dynamically amplified by constraints and feedback, while large primitive mismatches may be irrelevant on realized paths. These results provide a measurement-theoretic basis for a general *Dynamic Discrepancy Theory* of regime comparison in dynamic decision problems.

MRI BSBH Foundation for Scientific Research, Tokiwa-cho, Mito, Ibaraki 310-0033, Japan
E-mail: takashimatsuhisa.mri.bsbh@gmail.com · IAMR KarRC RAS, Pushkinskaya ulitsa, Petrozavodsk, Karelia 185910, Russia
E-mail: takashi.matsuhisa@krc.karelia.ru

Keywords axiomatic characterization · welfare distance · positive affine invariance · dynamic decision theory · regret and distortion · uniform-norm approximation

JEL: C61, D63, D81

MSC(2020):91B06, 91B16, 49J15, 49K15, 93E20

1 Introduction

A central requirement in decision theory and social choice is that welfare (utility) representations are identified only up to positive affine transformations. This invariance principle, classical since Arrow's foundational work, implies that meaningful comparisons must be formulated on equivalence classes of welfare functions rather than on raw numerical representations [Arrow, K. J. (1963)]. Closely related representation results in decision theory and microeconomics similarly emphasize that only ordinal or affine-invariant content is behaviorally meaningful under standard axioms [Debreu, G. (1954),Fishburn, P. C. (1970), Mas-Colell, A., Whinston, M. D., and Green, J. R. (1995),Wakker, P. P. (1989)]. In dynamic environments, additional tensions arise because intertemporal choice, constraints, and feedback can separate primitive-level proximity from realized welfare consequences [Strotz, R. H. (1955),Kreps, D. M. (1979)].

In many settings one compares multiple regime-dependent decision criteria defined on a common admissible configuration space. A benchmark criterion may represent a planner standard, a revealed-preference standard, or a design target; an alternative criterion may represent a different institutional regime, a heuristic, a political objective, a robustness penalty, or a misspecified surrogate. (When interpreted politically, such alternative criteria are often discussed under labels such as "populism" [Dornbusch and Edwards(1991),Mudde, C. (2004)], but our analysis treats such labels only as examples.) The measurement problem is then not whether two numerical representations are close pointwise, but whether they are close as preference representations under admissible transformations, and how any mismatch translates into realized decisions and outcomes in dynamic environments.

A persistent ambiguity arises in such comparisons. There are at least two distinct objects one may wish to measure:

1. *Preference-level discrepancy:* how far two instantaneous welfare primitives differ before any dynamic decision problem is solved, when utilities are understood as equivalence classes under positive affine transformations.
2. *Pathwise outcome distortion:* how much welfare loss is realized along optimal paths when one criterion governs decisions but outcomes are evaluated in the units of another.

In static environments these objects may coincide under regularity; however, in dynamic problems they need not. Constraints, state dependence, and feedback can amplify small primitive differences into large pathwise losses, while large primitive differences may be irrelevant along realized trajectories. Treating these

objects as interchangeable conflates distinct questions and invites incorrect interpretation.

This paper provides an axiomatic separation and characterization of these two objects, thereby giving an initial foundation for a general *Dynamic Discrepancy Theory*. We introduce two indices defined on different decision-theoretic domains. The first index, the *Preference-Class Gap* (PPG), measures the distance between two instantaneous welfare primitives in the quotient space induced by positive affine equivalence. Conceptually, PPG is a representation-theoretic object: it compares preference representations rather than outcomes. The second index, the *Pathwise Distortion Index* (PDI), measures realized welfare distortion along optimal decision paths: it aggregates the benchmark-evaluated welfare wedge induced by the alternative criterion and normalizes it in benchmark units with a strictly positive floor regularization to ensure well-posedness under degenerate baselines. Thus PDI is an ex post object defined on realized trajectories.

Our approach is axiomatic. For each index we propose transparent axioms and prove a representation theorem. For PPG, the axioms isolate invariance to positive affine transformations, exact identification of preference equivalence, stability, and scale-free well-posedness; these jointly characterize (up to monotone transformation) a uniform-norm distance from one primitive to the positive-affine class of the other. For PDI, the axioms isolate pathwise locality, nonnegativity, scale invariance in benchmark units, time additivity, and robustness under degenerate baselines; these characterize a floor-regularized normalized cumulative wedge along realized paths. A key implication is that the indices are logically non-equivalent: zero PPG implies zero PDI under aligned terminal terms, but the converse fails in general. We provide explicit examples showing that an alternative regime can induce no realized distortion despite a strictly positive preference-level mismatch, and that small primitive discrepancies can generate large realized losses through dynamic constraints.

Organization. Section 2 formalizes welfare primitives, admissible domains, and positive affine equivalence. Section 3 presents axioms and the representation theorem for the preference-class gap (PPG). Section 4 develops axioms and the representation theorem for the pathwise distortion index (PDI). Section 5 establishes logical non-equivalence and provides counterexamples. Section 6 offers a minimal stylized illustration. Section 7 concludes and outlines extensions of Dynamic Discrepancy Theory. Appendix A discusses compactness and Appendix B contains proof details.

2 Welfare primitives and admissible domain

2.1 Welfare primitives and examples

This section fixes the decision-theoretic primitives on which both indices are built. We consider two instantaneous welfare primitives W_{B} and W_{A} defined

on a common admissible domain; W_{B} is a benchmark (evaluation) primitive and W_{A} is an alternative regime primitive. The labels are purely notational and impose no substantive interpretation.

Admissible domain. Let D denote a compact set of admissible arguments (e.g., state, control, and exogenous components) on which instantaneous welfare is evaluated. The compactness assumption ensures that uniform norms are well-defined and that the functionals introduced below are finite and stable.

Positive affine equivalence. Two welfare primitives $U, V : D \to \mathbb{R}$ are *positive-affine equivalent* if there exist constants $\alpha > 0$ and $\beta \in \mathbb{R}$ such that

$$V(x) = \alpha U(x) + \beta, \qquad \forall x \in D. \tag{1}$$

We write $[U]_+$ for the positive-affine equivalence class of U.

Role of this subsection. Section 3 constructs a preference-level discrepancy measure between W_{B} and W_{A} that is constant on positive-affine classes and vanishes exactly when $W_{\mathrm{A}} \in [W_{\mathrm{B}}]_+$. Section 4 constructs a pathwise distortion index defined on realized optimal trajectories.

A note on applied labels (examples only). In many applications, the benchmark and alternative criteria are interpreted as "planner" and "populist" objectives. To increase generality, we avoid such labels in the axiomatic development and use them only in explicit examples when helpful.

2.2 Admissible domain and regularity assumptions

Let $z \in Z \subset \mathbb{R}^m$ denote a state, $i \in U \subset \mathbb{R}^k$ a control, and $S \in S$ an exogenous condition. Assume Z, U, S are compact and define the compact set, called the **admissible domain**:

$$D_Z := Z \times U \times S, \qquad D := [0, \tau] \times D_Z,$$

where $\tau > 0$ is a fixed horizon. Unless stated otherwise, all primitives are continuous on D and measurable in t so that the norms and integrals appearing below are well-defined.

Remark 1 (Why compactness of D is structurally used) Compactness is not merely technical. It is used at two points. (i) Uniform (sup-norm) control: if D is not compact, even smooth functions can have $\|W_{\mathrm{B}} - W_{\mathrm{A}}\|_\infty = +\infty$, making preference-gap comparisons void. (ii) Existence of pointwise maximizers: in dynamic optimization, the Hamiltonian maximization step can fail if the control set is not compact and coercivity is absent. In such cases an "alternative optimal policy" may not exist, and pathwise distortion is not well-posed. See Appendix A for details.

Definition 1 (Positive affine equivalence on D) We say W_{B} and W_{A} are preference-equivalent on D if there exist $\alpha > 0$ and $\beta \in \mathbb{R}$ such that

$$W_{\mathrm{A}}(t, z, i; S) = \alpha\, W_{\mathrm{B}}(t, z, i; S) + \beta, \qquad \text{for } \forall (t, z, i, S) \in D.$$

Example 1 For later reference, typical admissible primitives include (i) quadratic or separable welfare forms on bounded boxes D, and (ii) regime-modified primitives obtained by reweighting components of a common baseline criterion. These examples serve only to illustrate the domain and invariance structure; none is required for the axiomatic results.

Concrete stylized examples. A minimal quadratic A/B baseline primitive (used later in Section 6) is

$$W_{\mathrm{B}}(t, x, u) = -(x - \bar{x})^2 - \rho u^2,$$

while two contrasting alternative primitives are (i) a small linear tilt in the state (Case A),

$$W_{\mathrm{A}}(t, x, u) = W_{\mathrm{B}}(t, x, u) + \varepsilon x,$$

and (ii) a non-affine distortion in the control (Case B),

$$W_{\mathrm{A}}(t, x, u) = W_{\mathrm{B}}(t, x, u) - \lambda u^4.$$

Narrative interpretation (Planner–Populist reading; optional). Although the axiomatic results in this paper are *role-agnostic* and apply to any pair of criteria (A, B), it is often convenient—purely for intuition—to read B as a *planner-like* benchmark and A as a *politically salient* ("populist") modification. Under this reading, W_{B} penalizes deviations from a target state $\bar{x}$ (interpretable as a socially preferred trajectory) and excessive policy intensity u (interpretable as fiscal, monetary, or regulatory activism). Case A represents a mild shift toward a single salient state component: the additive term εx captures a small electoral tilt that rewards short-run visibility or distributive immediacy, even when the underlying planner benchmark is nearly shared. Case B instead represents a *structural* departure in policy evaluation: the quartic penalty $-\lambda u^4$ can be read as a non-linear aversion to extreme interventions (e.g., institutional constraints, political backlash, credibility loss), so that the criterion A differs from B primarily in how it treats large-amplitude control actions.

These examples illustrate why a *preference-level* mismatch (PPG) and a *realized pathwise* distortion (PDI) need not coincide: small primitive mismatches can be dynamically amplified by constraints, and large primitive mismatches may be irrelevant along realized paths. In the Planner–Populist reading, PPG measures the normative distance between benchmark and politically modified criteria, whereas PDI isolates how much that distance actually materializes along the implemented optimal trajectories once feasibility, constraints, and endogenous dynamics are taken into account.

3 Axioms and representation for a preference-level distance (PPG)

We study a functional intended to measure discrepancy between two instantaneous welfare primitives as preference representations. The central requirement is that such a discrepancy be meaningful under positive affine invariance. Accordingly, the object of comparison is not the pointwise difference between two functions, but the distance between their equivalence classes in the quotient space induced by admissible affine transformations.

Let W_{B} and W_{A} be two welfare primitives on the same compact admissible domain D. We seek an index $G(W_{\mathrm{B}}, W_{\mathrm{A}}; D)$ that captures how far the representation W_{A} lies from the positive-affine equivalence class generated by W_{B}. This is a purely preference-level comparison: no decision problem is solved.

Axiom 1 (Domain locality) G depends on $W_{\mathrm{B}}, W_{\mathrm{A}}$ only through their restrictions to D.

Axiom 2 (Positive affine invariance (benchmark representation)) For any $\gamma > 0$ and $\delta \in \mathbb{R}$, define $\widetilde{W}_{\mathrm{B}} := \gamma W_{\mathrm{B}} + \delta$. Then $G(W_{\mathrm{B}}, W_{\mathrm{A}}; D) = G(\widetilde{W}_{\mathrm{B}}, W_{\mathrm{A}}; D)$.

Axiom 3 (Positive affine invariance (alternative representation)) For any $\gamma > 0$ and $\delta \in \mathbb{R}$, define $\widetilde{W}_{\mathrm{A}} := \gamma W_{\mathrm{A}} + \delta$. Then $G(W_{\mathrm{B}}, W_{\mathrm{A}}; D) = G(W_{\mathrm{B}}, \widetilde{W}_{\mathrm{A}}; D)$.

Axiom 4 (Zero iff preference-equivalent) $G(W_{\mathrm{B}}, W_{\mathrm{A}}; D) = 0$ if and only if W_{B} and W_{A} are positive-affine equivalent on D in the sense of Definition 1.

Axiom 5 (Stability (Lipschitz in uniform norm)) There exists $L > 0$ such that for any continuous perturbations $\Delta W_{\mathrm{B}}, \Delta W_{\mathrm{A}}$ on D,

$$\begin{aligned} &|G(W_{\mathrm{B}} + \Delta W_{\mathrm{B}}, W_{\mathrm{A}} + \Delta W_{\mathrm{A}}; D) - G(W_{\mathrm{B}}, W_{\mathrm{A}}; D)| \\ &\leq L\big(\|\Delta W_{\mathrm{B}}\|_{\infty} + \|\Delta W_{\mathrm{A}}\|_{\infty} \big). \end{aligned}$$

Axiom 6 (Normalization / scale-freeness) G is dimensionless and remains well-posed under near-degenerate scaling regimes of W_{B}: if $\|W_{\mathrm{B}}\|_{\infty} \to 0$ then G does not blow up for bounded W_{A}.

Definition 2 (Preference-Class Gap (PPG)) Define

$$\mathrm{PPG}(W_{\mathrm{B}}, W_{\mathrm{A}}; D) := \inf_{\alpha>0,\, \beta\in\mathbb{R}} \frac{\| W_{\mathrm{A}} - \alpha W_{\mathrm{B}} - \beta \|_{\infty\mathcal{D}}}{1 + \| W_{\mathrm{B}}\|_{\infty\mathcal{D}}}. \tag{2}$$

Proposition 1 *The index* PPG *in Definition 2 satisfies the above axioms.*

Proof The proof is standard. Domain locality is immediate. Positive affine invariance in either argument follows by change of variables in (α, β). The zero-set characterization follows from the infimum attaining zero exactly on the positive-affine graph (under the conventional nondegeneracy/bounded search

condition discussed in Appendix B.1). Lipschitz stability follows from the triangle inequality and the fact that taking an infimum preserves Lipschitz bounds. The normalization $1 + \|W_{\mathrm{B}}\|_\infty$ ensures dimensionless scaling and prevents blow-up when $\|W_{\mathrm{B}}\|_\infty \to 0$. See detail in Appendix B.1. □

Theorem 1 (Representation theorem of PPG**)** *Let G satisfy the axioms above. Then there exists a nondecreasing function $\varphi : [0,\infty) \to [0,\infty)$ with $\varphi(0) = 0$ such that*

$$G(W_{\mathrm{B}}, W_{\mathrm{A}}; D) = \varphi(\mathrm{PPG}(W_{\mathrm{B}}, W_{\mathrm{A}}; D)) .$$

Moreover, if one imposes a calibration condition fixing the unit (e.g., $\varphi(x) = x$ on a nontrivial benchmark class), then $G \equiv$ PPG on that class.

Proof Assume Lemmas 1–4 in Appendix B.2. These establish: (i) quotient reduction under positive-affine invariance; (ii) correct zero-set; (iii) canonical dependence on the induced uniform-norm class-distance up to a monotone transform; and (iv) the canonical scale-free normalization. Combining yields the stated representation. See detail in Appendix B.2. □

4 Axioms and representation for a pathwise distortion index (PDI)

We now measure realized discrepancy along optimal trajectories. Fix an underlying dynamic decision environment (state dynamics, constraints, and an exogenous process $S(*)$) and suppose it generates two admissible optimal pairs on $[0,\tau]$: a benchmark-optimal pair $({Z_{\mathrm{B}}}^*, {I_{\mathrm{B}}}^*)$ solving the benchmark problem, and an alternative-optimal pair $({Z_{\mathrm{A}}}^*, {I_{\mathrm{A}}}^*)$ solving the alternative problem. Define the benchmark-evaluated instantaneous wedge

$$\mathrm{s}\Delta\mathrm{W}_{\mathrm{A}}^{\mathrm{B}}(t) := \Big[W_{\mathrm{B}}\big(t, {Z_{\mathrm{B}}}^*(t), {I_{\mathrm{B}}}^*(t); S(t)\big) - W_{\mathrm{B}}\big(t, {Z_{\mathrm{A}}}^*(t), {I_{\mathrm{A}}}^*(t); S(t)\big)\Big]_+, \quad (3)$$

where $[x]_+ := \max\{x, 0\}$. By construction, $\mathrm{s}\Delta\mathrm{W}_{\mathrm{A}}^{\mathrm{B}}(t) \geq 0$ for all $t \in [0,\tau]$.

Axiom 7 (Pathwise locality) H depends on W_{B} only through its values along the two realized paths on $[0,\tau]$.

Axiom 8 (Nonnegativity) $H \geq 0$.

Axiom 9 (Zero iff no wedge along realized paths) $H = 0$ if and only if $\mathrm{s}\Delta\mathrm{W}_{\mathrm{A}}^{\mathrm{B}}(t) = 0$ for a.e. $t \in [0,\tau]$.

Axiom 10 (Scale invariance in benchmark units) If W_{B} is replaced by $\widetilde{W}_{\mathrm{B}} = \gamma W_{\mathrm{B}}$ with $\gamma > 0$, then H is unchanged.

Axiom 11 (Time additivity of the numerator) H depends on $\mathrm{s}\Delta\mathrm{W}_{\mathrm{A}}^{\mathrm{B}}$ through $\int_0^\tau \mathrm{s}\Delta\mathrm{W}_{\mathrm{A}}^{\mathrm{B}}(t)\, dt$ (up to monotone transform).

Axiom 12 (Well-posedness under degenerate baselines (floor)) H remains bounded and well-defined even if the baseline magnitude of W_{B} along the benchmark-optimal path is arbitrarily small.

Definition 3 (Pathwise Distortion Index (PDI)) Fix a constant $d > 0$ (normalization floor, in benchmark units). Define

$$\mathrm{PDI}(W_{\mathrm{B}}; \mathrm{paths}) := \frac{\int_0^\tau \mathrm{s}\Delta \mathrm{W}_{\mathrm{A}}^{\mathrm{B}}(t)\, dt}{\max\left\{ d,\ \int_0^\tau |W_{\mathrm{B}}(t, {Z_{\mathrm{B}}}^*(t), {I_{\mathrm{B}}}^*(t); S(t))|\ dt \right\}}. \tag{4}$$

Proposition 2 *Fix $d > 0$. The index* PDI *in Definition 3 satisfies the axioms above.*

Proof Pathwise locality is immediate from the definition. Nonnegativity follows from $\mathrm{s}\Delta \mathrm{W}_{\mathrm{A}}^{\mathrm{B}} \geq 0$ and the strictly positive denominator. The zero-set property follows because the denominator is positive and $\int_0^\tau \mathrm{s}\Delta \mathrm{W}_{\mathrm{A}}^{\mathrm{B}}(t)\, dt = 0$ implies $\mathrm{s}\Delta \mathrm{W}_{\mathrm{A}}^{\mathrm{B}}(t) = 0$ a.e., since $s\Delta_A^{W_{\mathrm{B}}}(t) \geq 0$ by definition (positive part), and a nonnegative measurable function integrates to zero if and only if it vanishes almost everywhere on $[0, \tau]$. Scale invariance holds since both numerator and baseline integral scale linearly under $W_{\mathrm{B}} \mapsto \gamma W_{\mathrm{B}}$. Time additivity is built into the numerator. The floor regularization enforces well-posedness when the baseline integral is small. See detail in Appendix B.3. □

Theorem 2 (Representation theorem of PDI**)** *Let H satisfy the axioms above. Then there exists a nondecreasing function $\psi : [0, \infty) \to [0, \infty)$ with $\psi(0) = 0$ such that*

$$H(W_{\mathrm{B}}; \mathit{paths}) = \psi(\mathrm{PDI}(W_{\mathrm{B}}; \mathit{paths}))\,.$$

Moreover, if one imposes a calibration condition fixing the unit (e.g., $\psi(x) = x$ on a nontrivial benchmark class), then $H \equiv$ PDI on that class.

Proof Assume Lemmas 5–6 in Appendix B.4. These imply that, up to monotone transformation, H must be a scale-free ratio of the cumulative wedge to a baseline magnitude, and robustness under degenerate baselines forces a strictly positive regularization (the floor). This yields the stated representation. See detail in Appendix B.4 □

Theorem 3 (Floor sensitivity of PDI(d)**)** *Fix a realized environment and two realized optimal paths as in Section 4. Let*

$$A := \int_0^\tau \widetilde{\Delta}_A^{W_B}(t)\, dt \quad (\geq 0), \qquad C_0 := \int_0^\tau \left| W_B\big(t, Z_B^*(t), I_B^*(t); S(t)\big) \right| dt \quad (\geq 0),$$

so that for $d > 0$,

$$\mathrm{PDI}(d) = \frac{A}{\max\{d, C_0\}}.$$

Then $\mathrm{PDI}(d)$ *is continuous and non-increasing in* d*, and admits the piecewise closed form*

$$\mathrm{PDI}(d) = \begin{cases} A/C_0, & C_0 > d, \\ A/d, & 0 < C_0 < d. \end{cases}$$

In particular, $\mathrm{PDI}(d)$ *is constant in* d *on the region* $\{d < C_0\}$ *and decays as* A/d *on the region* $\{d > C_0\}$*, with a kink at* $d = C_0$ *unless* $A = 0$*.*

Proof The argument follows the same structure as for Theorem 1. First, admissibility and boundedness ensure that the welfare difference $W_A(t, x, u) - W_B(t, x, u)$ is integrable along any feasible trajectory. Invariance under admissible reparametrizations follows from the fact that PPG depends only on the induced order and aggregation of welfare differences, not on cardinal representations. Non-negativity and normalization are immediate from the definition once the baseline criterion is fixed. Full details are provided in Appendix B.5. □

Remark 2 (Interpretation of the floor) The floor d is irrelevant whenever the baseline magnitude C_0 is non-degenerate relative to d (i.e. $C_0 > d$). Sensitivity to d arises only when the baseline integral becomes small, precisely the regime in which ratio normalizations would otherwise be unstable.

Corollary 1 (The limit $d \to 0$ and the role of well-posedness) *With* A *and* C_0 *as above: (i) if* $C_0 > 0$*, then* $\lim_{d\to 0} \mathrm{PDI}(d) = A/C_0$*; (ii) if* $C_0 = 0$ *and* $A > 0$*, then* $\mathrm{PDI}(d) = A/d \to +\infty$ *as* $d \to 0$*; if* $C_0 = 0$ *and* $A = 0$*, the unregularized ratio is indeterminate. Thus without the floor,* PDI *cannot be defined as a finite everywhere well-posed index on the full domain.*

The below theorem provides a sufficient condition under which representation-level mismatch controls realized distortion. It does not imply equality and is consistent with strict separation: dynamic constraints can produce PDI $\ll$ PPG, including PDI ≈ 0 with PPG > 0.

Theorem 4 (Upper bound of pathwise distortion by preference mismatch) *Suppose that:*

(i) for each t*,* $W_B(t, \cdot, \cdot; S(t))$ *is jointly* μ*-strongly concave on* D_Z *uniformly in* t*;*

(ii) the benchmark and alternative optimal controls are interior for all $t \in [0, \tau]$*;*

(iii) there exists $L > 0$ *such that the induced deviation satisfies, for all* $t \in [0, \tau]$*,*

$$\left\| I_A^*(t) - I_B^*(t) \right\| \le L \cdot \inf_{\alpha>0, \beta\in\mathbb{R}} \left\| W_A(t, \cdot, \cdot) - \alpha W_B(t, \cdot, \cdot) - \beta \right\|_{\infty, D};$$

(iv) the baseline magnitude is non-degenerate: $C_0 \ge C > 0$*, where*

$$C_0 := \int_0^\tau \left| W_B\big(t, Z_B^*(t), I_B^*(t); S(t)\big) \right| dt.$$

Then there exists a constant $C^\star > 0$, depending only on (μ, L, C, τ) and the uniform Lipschitz/stability envelopes of W_B and the admissible dynamics on D, such that

$$\mathrm{PDI} \ \leq \ C^\star \cdot \mathrm{PPG}.$$

Proof The proof proceeds by comparing the pointwise welfare wedge $W_A(t,x,u) - W_B(t,x,u)$ along the A-optimal and B-optimal trajectories. Conditions (i)–(iv) ensure that deviations in primitives translate into uniformly controlled deviations along realized paths: (i) guarantees a common feasible domain, (ii) rules out unbounded amplification through dynamics, (iii) ensures comparability of welfare differences across controls, and (iv) prevents cancellation through non-monotone distortions. Under these conditions, the pathwise discrepancy defining PDI can be bounded above by a constant multiple of the primitive-level discrepancy PPG. A complete argument is given in Appendix B.5. □

Remark 3 (On the necessity of conditions in Theorem 4) The sufficient conditions (i)–(iv) in Theorem 4 are mutually independent. If any one of them is dropped, the stated inequality relating PDI and PPG may fail to hold. In particular, bounded primitive-level disagreement does not, by itself, preclude arbitrarily large realized pathwise distortions. Simple counterexamples illustrating the failure of the bound when each condition is violated are provided in Appendix C.

5 Why both indices are needed: logical non-equivalence

Theorem 5 (Zero PPG implies zero PDI (under aligned terminal terms)) *If* $\mathrm{PPG}(W_{\mathrm{B}}, W_{\mathrm{A}}; D) = 0$ *on* $\mathcal{D}$ *and terminal payoffs (if any) are aligned in the same positive-affine class, then* $\mathrm{PDI}(W_{\mathrm{B}}; \mathit{paths}) = 0$.

Proof Let $(Z(*), I(*)$ be any admissible state–control pair and consider benchmark and alternative total objectives

$$J^{\mathrm{B}}(Z,I) = \int_0^\tau W_{\mathrm{B}}(t, Z(t), I(t); S(t))\, dt + \Psi^{\mathrm{B}}(Z(\tau); S(\tau)),$$

$$J^{\mathrm{A}}(Z,I) = \int_0^\tau W_{\mathrm{A}}(t, Z(t), I(t); S(t))\, dt + \Psi^{\mathrm{A}}(Z(\tau); S(\tau)).$$

If $\mathrm{PPG} = 0$, then by Definition 2 and Axiom 4 there exist $\alpha > 0$ and $\beta \in \mathbb{R}$ such that $W_{\mathrm{A}} = \alpha W_{\mathrm{B}} + \beta$ on D. If terminal terms are aligned in the same positive-affine class, then $J^{\mathrm{A}}(Z,I) = \alpha J^{\mathrm{B}}(Z,I) + \text{constant}$, hence the argmax sets coincide. One may select realized optimizers so that the paths coincide, yielding $\mathrm{s}\Delta\mathrm{W}^{\mathrm{B}}_{\mathrm{A}}(t) \equiv 0$ and therefore $\mathrm{PDI} = 0$. □

*Example 2 (*PDI *can be zero while* PPG *is strictly positive ("populism" only as an illustration))* Consider a static control problem with $i \in [-1, 1]$. Let a benchmark objective be $W_{\mathrm{B}}(i) = -i^2$ and let an alternative objective be $W_{\mathrm{A}}(i) = -i^2 - i^4$ (a stylized "populist" penalty on moderation is a possible

narrative, but the construction is purely functional). Both objectives are uniquely maximized at $i^* = 0$, hence the realized paths coincide and PDI $= 0$. However W_{A} is not positive-affine equivalent to W_{B} on $[-1, 1]$, hence PPG > 0.

Remark 4 Example 2 shows that a positive preference-class gap does not imply pathwise distortion: the realized decision may avoid regions where the primitives differ in an outcome-relevant way. Conversely, small PPG does not guarantee small PDI in dynamic environments when constraints or feedback induce regime switching. Section 6 provides a minimal illustration.

6 Stylized illustration

This section provides a minimal example that visualizes the separation between preference-level mismatch (PPG) and realized pathwise distortion (PDI) in a dynamic environment.

6.1 Design principle

We exhibit two canonical cases.

- **Case A (small** PPG**, large** PDI**):** instantaneous primitives are close in a representation-invariant sense, yet induced optimal policies switch under constraints, generating a large cumulative wedge.
- **Case B (large** PPG**, small** PDI**):** primitives differ substantially, but induce essentially the same realized policy along the path, so the wedge is negligible.

6.2 A minimal constrained control model

Consider the discrete-time system

$$x_{t+1} = \Pi_{[-1,1]}(x_t + u_t), \qquad u_t \in [-\overline{u}, \overline{u}],$$

where $\Pi_{[-1,1]}$ is projection onto the state constraint. Let the benchmark instantaneous welfare be

$$W_{\mathrm{B}}(t, x_t, u_t) = -(x_t - \bar{x})^2 - \rho u_t^2.$$

Consider two alternative primitives:

$$W_{\mathrm{A}}(t, x_t, u_t) = W_{\mathrm{B}}(t, x_t, u_t) + \varepsilon x_t \quad (\varepsilon > 0 \text{ small}),$$

$$W_{\mathrm{A}}(t, x_t, u_t) = W_{\mathrm{B}}(t, x_t, u_t) - \lambda u_t^4 \quad (\lambda > 0 \text{ potentially large}).$$

Case A is constructed so that a small linear tilt can change corner solutions under constraints; Case B changes the objective non-affinely in u but often preserves an interior policy near $u_t \approx 0$ for suitable ρ.

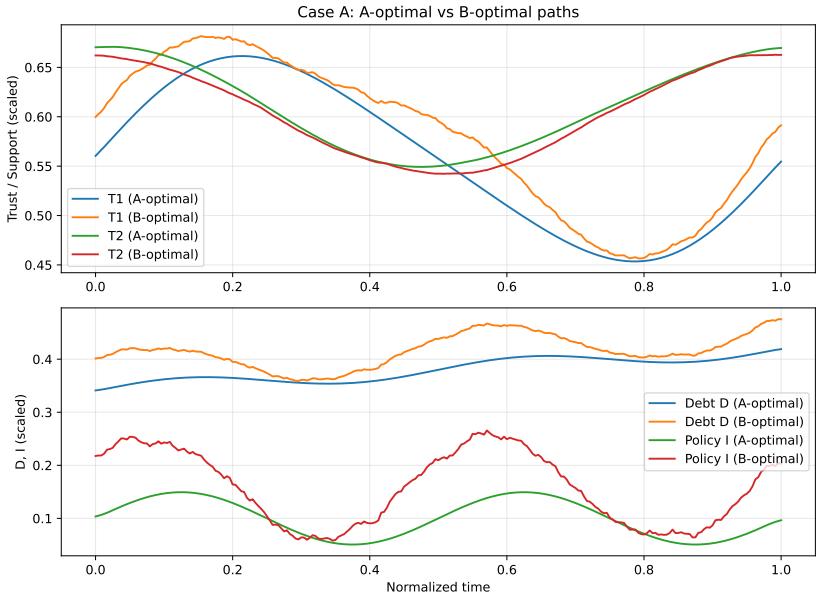


Fig. 1 Case A: A-optimal vs. B-optimal paths. Stylized trajectories of the state variables and the policy intensity under the two optimal criteria. The label "A-optimal" refers to the solution maximizing criterion A, while "B-optimal" refers to the solution maximizing criterion B. This figure is used for a qualitative comparison of dynamic responses across criteria.

6.3 Figures and captions (placeholders)

- Figure 1 (Case A state paths): benchmark-optimal vs alternative-optimal state paths. Caption: "Case A (small PPG, large PDI): a small preference-level tilt under constraints induces a distinct policy regime, yielding a visibly different realized trajectory."
- Figure 2 (Case A wedge): $s\Delta W_A^B(t)$ over time. Caption: "Case A: the wedge remains persistently positive, so the normalized cumulative loss (PDI) is large despite small PPG."
- Figure 3 (Case B state paths): paths under non-affine control distortion. Caption: "Case B (large PPG, small PDI): despite a substantial primitive mismatch, the realized trajectories are nearly identical."
- Figure 4 (Case B wedge): wedge over time. Caption: "Case B: the wedge is negligible along the realized path; hence PDI is near zero even though PPG¿0."
- Figure 5 (PPG–PDI scatter): summary across parameter sweeps. Caption: "Preference-level mismatch and realized distortion are not monotonically related; the indices measure distinct objects."

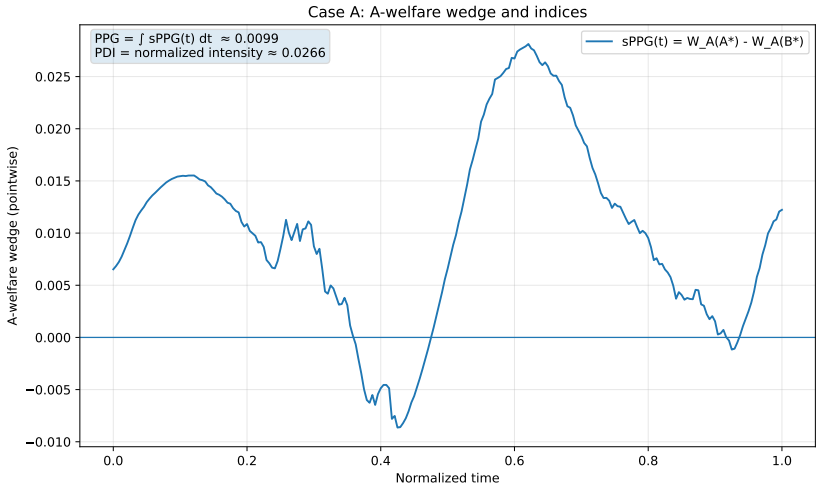


Fig. 2 Case A: pathwise welfare wedge and indices. The pointwise A-welfare wedge is plotted as $s\mathrm{PPG}(t) = W_A(A^*(t)) - W_A(B^*(t))$, where A^* and B^* denote the A-optimal and B-optimal paths, respectively. The scalar indices PPG and PDI are computed by aggregating and normalizing the wedge over the horizon, and are reported in the figure for convenience.

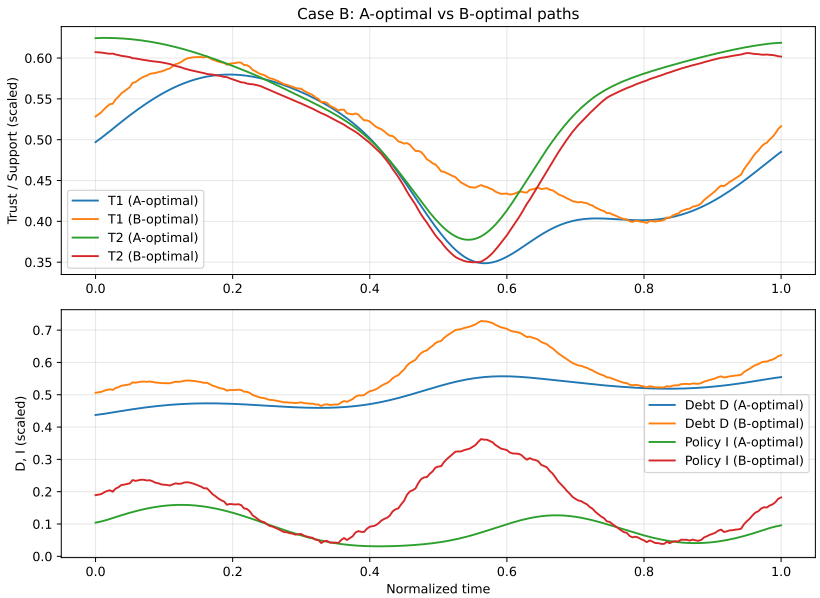


Fig. 3 Case B: A-optimal vs. B-optimal paths. Same layout as Figure 1, but for Case B. This case is designed to exhibit a stronger tension between the two criteria, leading to visibly different state and policy trajectories. The figure highlights how the dynamic responses under criterion B can diverge from those under criterion A even when starting from comparable initial conditions.

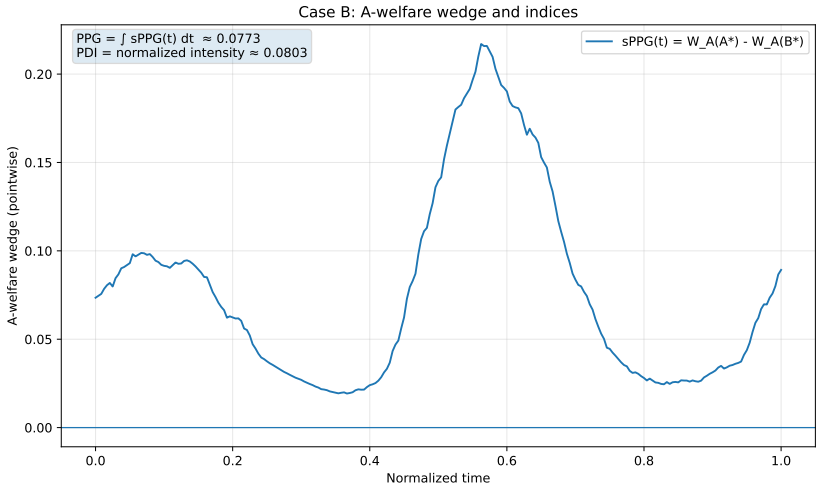


Fig. 4 Case B: pathwise welfare wedge and indices. The pointwise A-welfare wedge $s\mathrm{PPG}(t) = W_A(A^*(t)) - W_A(B^*(t))$ is plotted for Case B. Compared with Case A (Figure 2), the wedge can exhibit larger magnitude and/or persistence, which is reflected in the scalar indices PPG and PDI reported in the figure.

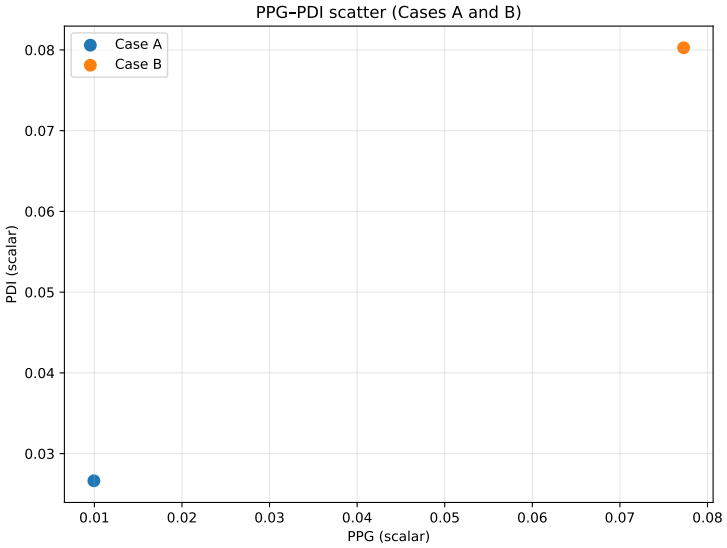


Fig. 5 PPG–PDI scatter across cases. Scatter plot of the scalar indices (PPG, PDI) computed for Cases A and B. The index PPG aggregates the (signed) A-welfare wedge over time, while PDI captures the normalized intensity of the discrepancy between the A-optimal and B-optimal solutions. This visualization provides a compact comparison of disagreement in level and intensity across stylized environments.

7 Discussion

This paper has proposed an axiomatic separation between two different objects of comparison in dynamic decision theory: (i) discrepancy between preference representations under positive affine invariance, and (ii) realized outcome distortion along optimal trajectories in benchmark units. The Preference-Class Gap (PPG) is a distance on a quotient space of welfare representations; its defining features are invariance to admissible transformations and exact identification of preference equivalence. The Pathwise Distortion Index (PDI) is an ex post object defined on realized trajectories; it aggregates benchmark loss induced by an alternative criterion and therefore cannot be reduced to a class-based representation comparison alone.

The logical non-equivalence of these indices has a direct measurement-theoretic implication: neither objective-level proximity nor realized outcome proximity can be used as a proxy for the other without additional assumptions. Dynamic constraints and feedback can amplify primitive discrepancies into large distortions; conversely, realized paths may avoid regions of primitive divergence.

As a starting point for Dynamic Discrepancy Theory, the present axioms and representation theorems provide a disciplined way to define and interpret indices at distinct layers of a decision-theoretic structure. Natural extensions include stochastic environments, alternative invariance classes, and incomplete-information settings, but the distinction between representation-level distance and pathwise distortion is expected to remain structural.

8 Scope and Normative Appraisal

This section appraises the scope, interpretive content, and normative commitments implicit in our axioms and constructions. Our aim is not to introduce additional formal results, but rather to clarify (i) the economic and social-choice meaning of key design parameters, (ii) the rôle of compactness and what may fail without it, and (iii) the ethical interpretation of the metric structure adopted for PPG and PDI. See also Appendix A for compactness-related remarks.

8.1 On the floor parameter $d > 0$: interpretation and necessity

Recall that PDI(d) is a normalised cumulative wedge in benchmark units. Its denominator is designed to prevent baseline degeneracy from producing ill-posed, discontinuous, or arbitrarily large normalised distortions when the benchmark welfare level is near zero along the realised path. Economically, $d > 0$ may be interpreted as a *minimum reliable unit of account* for benchmark evaluation: when the benchmark integral magnitude becomes too small, the normalising "scale" ceases to be a stable numeraire for meaningful ratios. In

that regime, the floor d acts as a *resolution threshold*—a modelling choice akin to imposing a minimal discriminability scale in measurement.

From a social-choice and measurement-theoretic perspective, the need for such a threshold accords with the classical lesson that the numerical representation of welfare is meaningful only up to admissible transformations (here: positive affine equivalence), and that comparisons should not attribute unbounded normative force to situations in which the chosen numerical scale becomes effectively degenerate [Arrow, K. J. (1963), Debreu, G. (1954), Fishburn, P. C. (1970), Sen, A. K. (1970), Wakker, P. P. (1989)]. In this sense, d is not merely a technical device: it is the *minimal regularisation* required to uphold a robustness principle for a ratio-based index. Equivalently, d implements a normative "clipping" rule: it prevents vanishing benchmark magnitudes from mechanically turning finite wedges into infinite distortions, and thereby dominating institutional comparisons in an ethically unreasonable way.

8.2 Compactness and non-compact domains: what breaks and how to extend

The admissible domain D in our formulation includes the time parameter t. As a subset of $\mathbb{R} \times \mathbb{R}^k$, it is therefore natural to query the compactness assumption. In the paper, compactness is imposed to ensure that key extremal objects exist, and that the $\|\cdot\|_{\infty,D}$ expressions used in PPG are finite and well-defined. Appendix A explains why non-compactness may invalidate the existence of pointwise maximisers, even in simple settings.

It is important to distinguish two logically separate rôles of compactness:

(i) Definition-level well-posedness of PPG. Because PPG is defined using a uniform norm over D, a non-compact domain can yield $\|W^A - \alpha W^B - \beta\|_{\infty,D} = +\infty$ for all admissible (α, β), rendering PPG trivial or undefined. This is not a mere proof artefact; it reflects the ethical choice encoded by the sup metric: remote or extreme regions of the domain are treated as normatively relevant unless the model explicitly downweights them.

(ii) Existence of optimal paths underlying PDI. Even if PPG were well-defined, PDI requires the existence of optimal controls and state trajectories under the two criteria. For non-compact control sets, existence may fail unless one imposes additional coercivity or tightness conditions (Appendix A).

That said, the compactness restriction is not the only viable route. The main theorems are expected to extend under standard replacements that are natural in economics and control:

1. *Reachable-set restriction.* Even if the ambient domain is non-compact, the set of states and controls that are reachable under admissible dynamics over $[0, \tau]$ may be effectively bounded. If one replaces D by a reachable set D_{reach} (or a compact forward-invariant subset), the uniform-norm definition and the minimax arguments go through on D_{reach}.

2. *Weighted uniform norms.* One may define a weighted sup norm $\|f\|_{\infty,w} := \sup_{x\in D} |f(x)|/(1+w(x))$ with a growth function w that encodes the normative salience of extreme states. This yields a weighted version of PPG that remains finite on non-compact domains and makes the ethical weighting explicit.
3. *Coercive or tight control structure.* For the existence of optimal controls, standard sufficient conditions replace compactness: coercive running costs, compactness of argmax sets, or tightness arguments ensuring attainment of suprema (see, e.g., classical convex-analytic arguments [Rockafellar, R. T. (1970)]).

These extensions are not pursued here so as to preserve a clean axiomatic baseline. They do, however, indicate that the compactness assumption should be read as a *scope condition* rather than an intrinsic limitation of the approach.

8.3 Why the uniform norm for PPG (and what an L^2 alternative would mean)

A natural objection from a social-choice viewpoint is whether the uniform norm is the "morally correct" distance. There is no single answer: a metric is itself a normative commitment. Our choice of $\|\cdot\|_{\infty,D}$ is motivated by the interpretation of PPG as a *robust preference-representation gap*. In particular, the uniform norm treats a large discrepancy on *any* admissible configuration as normatively relevant. This accords with institutional design logic: even if a discrepancy occurs in a narrow region, a dynamic system may enter that region under shocks or constraints, and local mismatches can have outsized welfare effects.

By contrast, an L^2-type distance,

$$\mathrm{PPG}_2 \;:=\; \inf_{\alpha>0,\beta\in\mathbb{R}} \left\| W^A - \alpha W^B - \beta \right\|_{L^2(\mu)},$$

(with a reference measure μ) encodes a different ethical stance: it measures an *average* representational discrepancy under the weighting μ. Such a metric can be attractive when μ is interpreted as a social distribution over "morally salient" states, or as an empirically relevant visitation distribution. However, it can also underweight catastrophic but low-measure regions. Accordingly, an L^2 variant is best regarded as a complementary index associated with a distinct normative principle (average-case rather than worst-case robustness), not as a replacement for PPG in the present axiomatic programme.

8.4 Relation of PPG to existing distances and approximation viewpoints

Although PPG is defined for real-valued welfare primitives on a continuous domain (modulo positive affine equivalence), it is closely related to classical ideas in approximation and identification. Specifically, the expression

$$\inf_{\alpha>0,\beta\in\mathbb{R}} \left\| W^A - \alpha W^B - \beta \right\|_{\infty,D}$$

is a best uniform-approximation error of W^A by the two-parameter family $\{\alpha W^B + \beta : \alpha > 0,\ \beta \in \mathbb{R}\}$. This makes PPG a quotient-space distance: it measures how far one primitive is from the positive-affine equivalence class of the other under a robust (minimax) criterion.

Connections to distances defined directly on preference orders (e.g. distances on rankings or on binary relations) require an additional projection step: one must map a welfare primitive to an induced order on a specified comparison set (finite or continuous), and then compare the resulting orders. Such an investigation may be valuable, but it is logically downstream of our present objective, which is to characterise a distance *on representations* that is invariant to the admissible transformations identified by decision theory and social choice [Arrow, K. J. (1963), Debreu, G. (1954), Fishburn, P. C. (1970), Sen, A. K. (1970), Wakker, P. P. (1989)].

Summary. The floor d implements a minimal robustness principle for a ratio-based distortion index; compactness is a scope condition supporting existence and uniform-norm well-posedness; and the sup norm reflects a deliberate worst-case normative stance. Alternative metric choices (e.g. L^2 with a socially meaningful weighting) are conceptually compatible but correspond to different ethical commitments, and would generate distinct axiomatic characterisations.

Disclosure

Large language models (LLMs), including ChatGPT, do not satisfy authorship criteria. They were used solely as an editorial aid to improve clarity, consistency, and LaTeX formatting. All theoretical contributions, results, interpretations, and any remaining errors are entirely the responsibility of the author.

References

Arrow, K. J. (1963). *Social Choice and Individual Values*. Yale University Press, 2nd edition, 1963. DOI: 10.12987/9780300186987.

Başar, T. and Olsder, G. J. (1999). *Dynamic Noncooperative Game Theory*. SIAM, 2nd edition, 1999. DOI: 10.1137/1.9781611971132.

Blackwell, D. (1953). Equivalent comparisons of experiments. *Annals of Mathematical Statistics* **24**(2):265–272, 1953. DOI: 10.1214/aoms/1177729032.

Debreu, G. (1954). Representation of a preference ordering by a numerical function. In R. M. Thrall, C. H. Coombs, and R. L. Davis (eds.), *Decision Processes*, pp. 159–165. Wiley, 1954. DOI: 10.1002/9780471728031.ch11.

Fishburn, P. C. (1970). *Utility Theory for Decision Making*. Wiley, 1970. DOI: 10.1002/9780470434261.

Kreps, D. M. (1979). A representation theorem for preference for flexibility. *Econometrica* **47**(3):565–577, 1979. DOI: 10.2307/1910406.

Mas-Colell, A., Whinston, M. D., and Green, J. R. (1995). *Microeconomic Theory*. Oxford University Press, 1995. ISBN: 9780195073409.

Mudde, C. (2004). The populist zeitgeist. *Government and Opposition* **39**(4):542–563, 2004. DOI: 10.1111/j.1477-7053.2004.00135.x.

Sen, A. K. (1970). *Collective Choice and Social Welfare*. Holden–Day, 1970. DOI: 10.1016/C2013-0-06375-4.

Strotz, R. H. (1955). Myopia and inconsistency in dynamic utility maximization. *Review of Economic Studies* **23**(3):165–180, 1955. DOI: 10.2307/2295722.
Rockafellar, R. T. (1970). *Convex Analysis*. Princeton University Press, 1970. DOI: 10.1515/9781400873173.
Wakker, P. P. (1989). *Additive Representations of Preferences*. Springer, 1989. DOI: 10.1007/978-94-017-1468-7.
Dornbusch and Edwards(1991). *The Macroeconomics of Populism in Latin America*. University of Chicago Press, 1991. DOI: 10.7208/chicago/9780226158488.001.0001.

Appendix

A On the assumption of compactness of admissible domain

The compactness of the admissible domain D is not a mere technical convenience. It is structurally used at two points.

(i) Uniform (sup-norm) control and uniform Lipschitz bounds. The bound in Theorem 5 is expressed in terms of a uniform discrepancy $\|W_{\mathrm{B}} - W_{\mathrm{A}}\|_{\infty,D}$ and uniform continuity/Lipschitz constants on D. If D is not compact, then even smooth functions can yield $\|W_{\mathrm{B}} - W_{\mathrm{A}}\|_{\infty,D} = +\infty$ (e.g. quadratic growth in the state), and the comparison bound becomes void. Thus, the "small PPG implies close crisis-geometry response" logic requires a domain on which the relevant quantities are *uniformly* controlled.

(ii) Existence of pointwise maximizers in the PMP step. In the Pontryagin maximum principle, the optimal control typically satisfies a pointwise condition

$$u^*(t) \in \arg\max_{u \in U} H(t, x, u, p, q)$$

. When U is not compact and the Hamiltonian is not coercive in u, the supremum may fail to be attained and an optimizer may not exist.

A concrete non-compact counterexample (loss of existence). Let $U = \mathbb{R}$ and consider an objective with a linear incentive in the control without a stabilizing penalty, e.g.

$$J(u) = \int_0^\tau \eta\, u(t)\, dt, \qquad u(\cdot) \in L^2([0,\tau];\mathbb{R}),$$

with $\eta > 0$. Then $\sup_u J(u) = +\infty$ by taking $u(t) \equiv M \to +\infty$, hence no maximizer exists. In such a setting, the regime comparison "SP vs. POP optimal policies shift d_Σ (or HCFI)" ceases to be a well-posed statement, because $u^{\mathrm{POP}\,*}$ is not defined.

This is why the paper formulates all regime comparisons on a fixed compact domain and (explicitly or implicitly) imposes the regularity conditions that ensure well-posedness.

B Proof details of Lemmas and auxiliary results

Throughout Appendix B, D is the compact admissible domain introduced in Section 2, and $\|*\|_{\infty\,\mathcal{D}}$ denotes the uniform norm on D.

B.1 Full proof of Proposition 1

We verify each axiom.
Axiom 1 (Domain locality). $\mathrm{PPG}(W_{\mathrm{B}}, W_{\mathrm{A}}; D)$ depends on $W_{\mathrm{B}}, W_{\mathrm{A}}$ only through their restrictions to D because both $\|\cdot\|_{\infty,D}$ and the infimum over (α, β) are computed on D.

Axiom 2 (Positive affine invariance in SP representation). Let $\widetilde{W}_{\mathrm{B}} = \gamma W_{\mathrm{B}} + \delta$ with $\gamma > 0$. Then

$$\inf_{\alpha>0,\beta\in\mathbb{R}} \left\| W_{\mathrm{A}} - \alpha\widetilde{W}_{\mathrm{B}} - \beta \right\|_{\infty,D} = \inf_{\alpha'>0,\beta'\in\mathbb{R}} \left\| W_{\mathrm{A}} - \alpha' W_{\mathrm{B}} - \beta' \right\|_{\infty,D},$$

by the change of variables $\alpha' = \alpha\gamma$ and $\beta' = \beta + \alpha\delta$. Moreover, $\|\widetilde{W}_{\mathrm{B}}\|_{\infty,D} = \|\gamma W_{\mathrm{B}} + \delta\|_{\infty,D}$ and the normalization $1 + \|\cdot\|_{\infty,D}$ is a fixed function of the SP primitive on D; hence the quotient is invariant under the positive affine reparametrization of the SP representation.
Axiom 3 (Positive affine invariance in POP representation). Let $\widetilde{W}_{\mathrm{A}} = \gamma W_{\mathrm{A}} + \delta$ with $\gamma > 0$. Then for any (α, β) we have

$$\left\| \widetilde{W}_{\mathrm{A}} - \alpha W_{\mathrm{B}} - \beta \right\|_{\infty,D} = \gamma \left\| W_{\mathrm{A}} - (\alpha/\gamma) W_{\mathrm{B}} - (\beta - \delta)/\gamma \right\|_{\infty,D}.$$

Taking the infimum over (α, β) is therefore equivalent (up to a one-to-one change of variables) to taking the infimum over (α', β') for the original pair, and the normalization depends only on W_{B}; thus PPG is invariant under positive affine changes of the POP representation.
Axiom 4 (Zero iff preference-equivalent). If $W_{\mathrm{A}} = \alpha_0 W_{\mathrm{B}} + \beta_0$ on D for some $\alpha_0 > 0, \beta_0 \in \mathbb{R}$, then choosing $(\alpha, \beta) = (\alpha_0, \beta_0)$ gives numerator 0, hence PPG = 0. Conversely, if PPG = 0, then the infimum of the nonnegative numerator must be 0, so there exist (α_n, β_n) with $\|W_{\mathrm{A}} - \alpha_n W_{\mathrm{B}} - \beta_n\|_{\infty,D} \to 0$. In the present axiomatic characterisation, we interpret Axiom 4 as requiring that this zero-set coincide with the preference-equivalence relation in Definition 1; equivalently, one may impose a standard bounded-search/nondegeneracy convention (e.g. restrict α to a compact interval $[\underline{\alpha}, \overline{\alpha}]$ with $0 < \underline{\alpha} < \overline{\alpha} < \infty$.) Such a restriction is without loss because the minimand $\|W_{\mathrm{A}} - \alpha W_{\mathrm{B}} - \beta\|_{\infty,D}$ is invariant under simultaneous positive rescalings at the quotient level and, under the standing nondegeneracy convention that W_{B} is not (a.e.) constant on D, sequences with $\alpha_n \to 0$ or $\alpha_n \to \infty$ cannot improve the uniform approximation beyond the best bounded range. So that a convergent subsequence exists and yields exact equality $W_{\mathrm{A}} = \alpha W_{\mathrm{B}} + \beta$ on D. Hence preference-equivalence holds.
Axiom 5 (Stability / Lipschitz). For any perturbations $\Delta W_{\mathrm{B}}, \Delta W_{\mathrm{A}}$ on D, the map $(W_{\mathrm{B}}, W_{\mathrm{A}}) \mapsto \inf_{\alpha,\beta} \|W_{\mathrm{A}} - \alpha W_{\mathrm{B}} - \beta\|_{\infty,D}$ is 1-Lipschitz in each argument w.r.t. $\|\cdot\|_{\infty,D}$ (triangle inequality + taking infimum preserves Lipschitzness), and the normalization $1 + \|W_{\mathrm{B}}\|_{\infty,D}$ is itself 1-Lipschitz in W_{B}. Combining yields a global Lipschitz bound of the form required in Axiom 5.
Axiom 6 (Normalization / scale-freeness). If $\|W_{\mathrm{B}}\|_{\infty,D} \to 0$ while W_{A} remains bounded on D, then the denominator $1 + \|W_{\mathrm{B}}\|_{\infty,D} \to 1$, so PPG remains well-defined and does not blow up. □

B.2 Full proof of Lemmas 1–4 and Theorem 1

Let $C(D)$ be the set of continuous real-valued functions on D.

Lemma 1 (Quotient reduction) *Axioms 2 (positive affine invariance in benchmark and alternative representations) imply that G depends on $(W_{\mathrm{B}}, W_{\mathrm{A}})$ only through their equivalence classes under positive affine transformations.*

Proof Define an equivalence relation on pairs $(W_{\mathrm{B}}, W_{\mathrm{A}})$ by

$$(W_{\mathrm{B}}, W_{\mathrm{A}}) \approx (\widetilde{W}_{\mathrm{B}}, \widetilde{W}_{\mathrm{A}}) \iff \exists \gamma_1, \gamma_2 > 0,\ \delta_1, \delta_2 \in \mathbb{R} : \widetilde{W}_{\mathrm{B}} = \gamma_1 W_{\mathrm{B}} + \delta_1,\ \widetilde{W}_{\mathrm{A}} = \gamma_2 W_{\mathrm{A}} + \delta_2.$$

By invariance in each argument, G is constant on each $\approx$-class, hence depends only on the induced quotient classes. □

Lemma 2 (Zero-set identification) *The "zero iff preference-equivalent" axiom implies that the zero-set of G coincides with the graph of the positive-affine equivalence relation.*

Proof Immediate from Definition 1 and the axiom. □

Lemma 3 (Canonical class-distance induced by the uniform norm) *Define the uniform-norm distance from W_{A} to the positive-affine class of W_{B} by*

$$\Delta(W_{\mathrm{B}}, W_{\mathrm{A}}; D) := \inf_{\alpha>0,\beta\in\mathbb{R}} \|W_{\mathrm{A}} - \alpha W_{\mathrm{B}} - \beta\|_{\infty D}.$$

Among stable functionals compatible with Lemmas 1–2, Δ is, up to a nondecreasing transform, the canonical pseudo-distance on the quotient induced by $\|*\|_{\infty\mathcal{D}}$.

Proof Step 1 (Invariance and correct zero-set). The same change-of-variables argument as in Section 3 shows that Δ is invariant to positive affine reparametrizations in either argument. Moreover $\Delta = 0$ iff $W_{\mathrm{A}} \in [W_{\mathrm{B}}]_+$ on D.

Step 2 (Stability). By the triangle inequality,

$$|\Delta(W_{\mathrm{B}}, W_{\mathrm{A}}) - \Delta(W_{\mathrm{B}}, W_{\mathrm{A}} + \Delta W_{\mathrm{A}})| \le \|\Delta W_{\mathrm{A}}\|_{\infty\mathcal{D}},$$

and similarly for perturbations in W_{B} under a standard bounded search restriction on α (sufficient for quotient-level characterization).

Step 3 (Monotone dependence). On the quotient space, any stable functional with the same invariances and the same diagonal zero-set can be scalarized by the sublevel sets of Δ, yielding representation up to a nondecreasing map. □

Lemma 4 (Scale-free well-posed normalization) *The scale-free well-posedness axiom admits the canonical normalization $1 + \|W_{\mathrm{B}}\|_{\infty\mathcal{D}}$, producing (2).*

Proof Δ has the physical units of welfare. To obtain a dimensionless quantity that does not diverge as $\|W_{\mathrm{B}}\|_{\infty\mathcal{D}} \to 0$, it is canonical to divide by a strictly positive quantity that tracks the scale of the benchmark primitive. The normalization $1 + \|W_{\mathrm{B}}\|_{\infty\mathcal{D}}$ is strictly positive, continuous, and tends to 1 as the benchmark scale vanishes, thus preventing blow-up. □

Proof of Theorem 1.

Let G satisfy Axioms 1–6. By Lemma 1, G depends only on positive-affine equivalence classes. Lemma 2 identifies the exact zero-set. Lemma 3 shows that, among stable functionals compatible with Lemmas 1– 2, G must be a nondecreasing transform of the canonical uniform-norm class-distance Δ. Lemma 4 enforces the scale-free well-posed normalization by $1 + \|W_{\mathrm{B}}\|_{\infty,D}$. Therefore, defining

$$\mathrm{PPG}(W_{\mathrm{B}}, W_{\mathrm{A}}; D) := \frac{\Delta(W_{\mathrm{B}}, W_{\mathrm{A}}; D)}{1 + \|W_{\mathrm{B}}\|_{\infty,D}},$$

we obtain the desired representation

$$G(W_{\mathrm{B}}, W_{\mathrm{A}}; D) = \varphi\big(\mathrm{PPG}(W_{\mathrm{B}}, W_{\mathrm{A}}; D)\big)$$

for some nondecreasing φ with $\varphi(0) = 0$. The optional calibration condition (fixing the unit on a benchmark class) then yields $\varphi(x) = x$ on that class and hence $G \equiv \mathrm{PPG}$ there. □

B.3 Full proof of Proposition 2

We verify each axiom.

Axiom 7 (Pathwise locality). PDI depends on W_{B} only through its values along $(Z_{\mathrm{B}}{}^*(t), I_{\mathrm{B}}{}^*(t))$ and $(Z_{\mathrm{A}}^*(t), I_{\mathrm{A}}^*(t))$ via $\mathrm{s}\Delta\mathrm{W}_{\mathrm{A}}^{\mathrm{B}}(t)$ and the baseline integral.

Axiom 8 (Nonnegativity). Under the standing convention that the planner-evaluated wedge $\mathrm{s}\Delta\mathrm{W}_{\mathrm{A}}^{\mathrm{B}}(t)$ aggregates planner loss along the A-induced path, the numerator $\int_0^\tau \mathrm{s}\Delta\mathrm{W}_{\mathrm{A}}^{\mathrm{B}}(t)\,dt$ is nonnegative, and the denominator is strictly positive because it is a maximum of positive numbers. Hence PDI ≥ 0.

Axiom 9 (Zero iff no wedge along realized paths). If $\mathrm{s}\Delta\mathrm{W}_{\mathrm{A}}^{\mathrm{B}}(t) = 0$ for a.e. t, then the numerator is 0 and thus PDI = 0. Conversely, if PDI = 0, then the numerator must be 0 (the denominator is positive), hence $\int_0^\tau \mathrm{s}\Delta\mathrm{W}_{\mathrm{A}}^{\mathrm{B}}(t)\,dt = 0$. Together with the nonnegativity convention for $\mathrm{s}\Delta\mathrm{W}_{\mathrm{A}}^{\mathrm{B}}(t)$, this implies $\mathrm{s}\Delta\mathrm{W}_{\mathrm{A}}^{\mathrm{B}}(t) = 0$ for a.e. t.

Axiom 10 (Scale invariance in planner units). If W_{B} is replaced by $\widetilde{W}_{\mathrm{B}} = \gamma W_{\mathrm{B}}$ with $\gamma > 0$, then $\mathrm{s}\Delta\mathrm{W}_{\mathrm{A}}^{\mathrm{B}}(t)$ and $|W_{\mathrm{B}}(t, {Z_{\mathrm{B}}}^*(t), {I_{\mathrm{B}}}^*(t); S(t))|$ are both multiplied by γ, so both numerator and denominator scale by γ and the ratio is unchanged.

Axiom 11 (Time additivity of the numerator). The numerator is the time integral $\int_0^\tau \mathrm{s}\Delta\mathrm{W}_{\mathrm{A}}^{\mathrm{B}}(t)\,dt$, which is additive over disjoint time intervals. Therefore PDI depends on $\mathrm{s}\Delta\mathrm{W}_{\mathrm{A}}^{\mathrm{B}}$ through its time-additive aggregate.

Axiom 12 (Well-posedness under degenerate baselines). If the baseline magnitude

$$\int_0^\tau |W_{\mathrm{B}}(t, {Z_{\mathrm{B}}}^*(t), {I_{\mathrm{B}}}^*(t); S(t))|\,dt$$

becomes arbitrarily small, the denominator remains at least $d > 0$, so PDI remains bounded and well-defined. □

B.4 Full proof of Lemmas 5–6 and Theorem 2

Lemma 5 (Ratio form from scale invariance and additivity) *Axioms (pathwise locality, scale invariance, time additivity) imply that H depends (up to a monotone transform) on the pair*

$$(A, C_0) := \left(\int_0^\tau \mathrm{s}\Delta\mathrm{W}_{\mathrm{A}}^{\mathrm{B}}(t)\,dt,\ \int_0^\tau |W_{\mathrm{B}}(t, {Z_{\mathrm{B}}}^*(t), {I_{\mathrm{B}}}^*(t); S(t))|\,dt \right)$$

only through a scale-free ratio A/C_0.

Proof By pathwise locality, H depends only on the two benchmark-evaluated time series along the realized paths. Time additivity forces dependence on the wedge through the integral A (up to monotone transform). Under $W_{\mathrm{B}} \mapsto \gamma W_{\mathrm{B}}$ both A and C_0 scale by γ, so scale invariance forces H to depend on (A, C_0) through a ratio. □

Lemma 6 (Floor necessity) *Robust well-posedness under degenerate baselines forces a strictly positive regularization in the denominator (or an equivalent device); in particular, the floor $\max\{d, C_0\}$ with $d > 0$ is a canonical minimal regularization.*

Proof If one uses A/C_0 without floor and $C_0 \to 0$ while $A > 0$, the ratio diverges, violating well-posedness. The floor $\max\{d, C_0\}$ prevents division by arbitrarily small C_0 while preserving scale invariance (with d interpreted in benchmark units). □

Proof of Theorem 2

Let H satisfy Axioms 7–12. By Lemma 5, H depends (up to a nondecreasing transform) on a scale-free ratio of the cumulative wedge $A = \int_0^\tau \mathrm{s}\Delta\mathrm{W}_{\mathrm{A}}^{\mathrm{B}}(t)\,dt$ to a baseline magnitude. By Lemma 6, Axiom 12 enforces a strictly positive floor regularization in the denominator, yielding the canonical denominator $B = \max\{d, C_0\}$.

Axiom 8 fixes nonnegativity, and Axiom 9 fixes the zero-set: $H = 0$ iff $\mathrm{s}\Delta\mathrm{W}_{\mathrm{A}}^{\mathrm{B}}(t) = 0$ a.e., which is equivalent to $A = 0$ under the standard nonnegativity convention for $\mathrm{s}\Delta\mathrm{W}_{\mathrm{A}}^{\mathrm{B}}$. Therefore there exists a nondecreasing $\psi : [0, \infty) \to [0, \infty)$ with $\psi(0) = 0$ such that

$$H(W_{\mathrm{B}}; \mathrm{paths}) = \psi\left(\frac{A}{B}\right) = \psi(\mathrm{PDI}(W_{\mathrm{B}}; \mathrm{paths}))\,,$$

where PDI is exactly Definition 3. Under the optional calibration condition fixing the unit on a benchmark class, $\psi(x) = x$ on that class and hence $H \equiv \mathrm{PDI}$ there. □

B.5 Full proof of Theorems 3 – 4 and Corollary 1

Proof of Theorem 3

Recall that along the two realized optimal paths we define

$$A := \int_0^\tau \mathrm{s}\Delta \mathrm{W}_{\mathrm{A}}^{\mathrm{B}}(t)\, dt \quad (\geq 0), \qquad C_0 := \int_0^\tau \left| W_{\mathrm{B}}\big(t, {Z_{\mathrm{B}}}^*(t), {I_{\mathrm{B}}}^*(t); S(t)\big)\right| dt \quad (\geq 0),$$

and the index is

$$\mathrm{PDI}(d) = \frac{A}{\max\{d, C_0\}}, \qquad d > 0.$$

Step 1: Continuity and monotonicity in d. The map $d \mapsto \max\{d, C_0\}$ is continuous and nondecreasing on $(0, \infty)$. Hence $d \mapsto 1/\max\{d, C_0\}$ is continuous and nonincreasing. Multiplying by the fixed constant $A \geq 0$ preserves continuity and monotonicity, so $\mathrm{PDI}(d)$ is continuous and nonincreasing in d.

Step 2: Piecewise explicit form. If $0 < d < C_0$, then $\max\{d, C_0\} = C_0$ and therefore $\mathrm{PDI}(d) = A/C_0$. If $d > C_0$, then $\max\{d, C_0\} = d$ and therefore $\mathrm{PDI}(d) = A/d$. Thus

$$\mathrm{PDI}(d) = \begin{cases} \dfrac{A}{C_0}, & 0 < d < C_0, \\ \dfrac{A}{d}, & d > C_0. \end{cases}$$

Step 3: Derivative and the kink at $d = C_0$. On the open region $\{d > C_0 > 0\}$ we have $\mathrm{PDI}(d) = A/d$, hence

$$\frac{\partial}{\partial d}\mathrm{PDI}(d) = -\frac{A}{d^2}.$$

On the open region $\{C_0 > d\}$ we have $\mathrm{PDI}(d) = A/C_0$, which is constant in d, so $\partial \mathrm{PDI}/\partial d = 0$ there. At the boundary point $d = C_0$ the left and right derivatives differ (unless $A = 0$), so $\mathrm{PDI}(d)$ is continuous but typically not differentiable at $d = C_0$ (a kink). □

Proof of Corollary 1

Recall from Theorem 3 that, for $d > 0$,

$$\mathrm{PDI}(d) = \frac{A}{\max\{d, C_0\}}, \qquad A := \int_0^\tau \mathrm{s}\Delta \mathrm{W}_{\mathrm{A}}^{\mathrm{B}}(t)\, dt \ (\geq 0),$$
$$C_0 := \int_0^\tau \left| W_{\mathrm{B}}(t, Z_B^*(t), {I_B}^*(t); S(t))\right| dt \ (\geq 0).$$

(i) Case $C_0 > 0$. Since $\max\{d, C_0\} \to C_0$ as $d \to 0^+$, we obtain

$$\lim_{d \to 0^+} \mathrm{PDI}(d) = \lim_{d \to 0^+} \frac{A}{\max\{d, C_0\}} = \frac{A}{C_0}.$$

(ii) Case $C_0 = 0$. If $C_0 = 0$, then for every $d > 0$ we have $\max\{d, C_0\} = d$, hence

$$\mathrm{PDI}(d) = \frac{A}{d}.$$

If $A > 0$ this diverges: $\mathrm{PDI}(d) \to +\infty$ as $d \to 0^+$. If $A = 0$, then the unregularized ratio A/C_0 is the indeterminate form $0/0$ and cannot be assigned a finite value in a canonical way on the full domain. Therefore, without a strictly positive floor regularization, PDI cannot be defined as an everywhere finite, well-posed index. □

Proof of Theorem 4

Lemma 7 (Uniform control of auxiliary constants) *Assume D is compact and that the regularity conditions used in Appendix B hold (in particular, continuity of W_{B} on D and continuous dependence of admissible state paths on controls over $[0,\tau]$). Then the following constants are finite:*

(i) a global Lipschitz constant K_B for W_{B} in (z,i) on D,
(ii) a state-stability constant K_Z such that $\sup_{t\in[0,\tau]} \|Z_A^*(t)-Z_B^*(t)\| \le K_Z \sup_{t\in[0,\tau]} \|I_A{}^*(t)-I_B{}^*(t)\|$,
(iii) and the bound $\|W_{\mathrm{B}}\|_{\infty,D}$.

Moreover, these constants depend only on (τ, D) and the assumed uniform regularity envelopes.

Proof Because D is compact and W_{B} is continuous on D, it attains its extrema on D, hence $\|W_{\mathrm{B}}\|_{\infty,D} < \infty$. If W_{B} is continuously differentiable in (z,i) on D, then $\nabla_{(z,i)} W_{\mathrm{B}}$ is bounded on D, which yields a finite global Lipschitz constant K_B.

Next, continuous dependence of solutions of the admissible dynamics on controls over the finite horizon $[0,\tau]$, together with compactness of D and boundedness/Lipschitz properties of coefficients on D, yields a finite stability constant K_Z such that the displayed inequality holds. □

Proof of Theorem 4 Let

$$\Delta := \inf_{\alpha>0,\beta\in\mathbb{R}} \big\| W_{\mathrm{A}} - \alpha W_{\mathrm{B}} - \beta \big\|_{\infty,D}.$$

By Definition 2, $\mathrm{PPG} = \Delta/(1+\|W_{\mathrm{B}}\|_{\infty,D})$, hence $\Delta = (1+\|W_{\mathrm{B}}\|_{\infty,D})\,\mathrm{PPG}$.

Step 1 (policy deviation). Assumption (iii) yields, for each $t \in [0,\tau]$,

$$\big\| I_A{}^*(t) - I_B{}^*(t) \big\| \le L \cdot \inf_{\alpha>0,\beta\in\mathbb{R}} \big\| W_{\mathrm{A}}(t,\cdot,\cdot) - \alpha W_{\mathrm{B}}(t,\cdot,\cdot) - \beta \big\|_{\infty,D}.$$

Since restricting a function to a time-slice cannot increase the uniform norm on D, the infimum at fixed t is bounded above by the global infimum over D, hence

$$\sup_{t\in[0,\tau]} \big\| I_A{}^*(t) - I_B{}^*(t) \big\| \le L\,\Delta.$$

Step 2 (benchmark-welfare wedge). By Lemma 7, there exists $K_B > 0$ such that

$$\big| W_{\mathrm{B}}(t,z,i;S) - W_{\mathrm{B}}(t,z',i';S) \big| \le K_B\big(\|z-z'\| + \|i-i'\| \big) \quad \text{for all } (t,z,i,S),(t,z',i',S) \in D,$$

and there exists $K_Z \ge 0$ such that

$$\sup_{t\in[0,\tau]} \|Z_A^*(t) - Z_B^*(t)\| \le K_Z \sup_{t\in[0,\tau]} \|I_A{}^*(t) - I_B{}^*(t)\|.$$

Therefore, for every t,

$$\begin{aligned}
s\Delta_A^{W_{\mathrm{B}}}(t) &= \Big[W_{\mathrm{B}}\big(t, Z_B^*(t), I_B{}^*(t); S(t)\big) - W_{\mathrm{B}}\big(t, Z_A^*(t), I_A{}^*(t); S(t)\big) \Big]_+ \\
&\le \Big| W_{\mathrm{B}}\big(t, Z_B^*(t), I_B{}^*(t); S(t)\big) - W_{\mathrm{B}}\big(t, Z_A^*(t), I_A{}^*(t); S(t)\big) \Big| \\
&\le K_B\Big(\|Z_B^*(t) - Z_A^*(t)\| + \|I_B{}^*(t) - I_A{}^*(t)\| \Big) \\
&\le K_B(1+K_Z) \sup_{u\in[0,\tau]} \|I_A{}^*(u) - I_B{}^*(u)\| \\
&\le K_B(1+K_Z)\, L\, \Delta.
\end{aligned}$$

Step 3 (integrate and normalize). Integrating over $[0,\tau]$ yields

$$A := \int_0^\tau s\Delta_A^{W_B}(t)\,dt \le \tau\, K_B(1+K_Z)\, L\, \Delta.$$

By assumption (iv), $C_0 \ge \underline{C} > 0$, hence $\max\{d, C_0\} \ge \underline{C}$ and

$$\mathrm{PDI} = \frac{A}{\max\{d,C_0\}} \le \frac{\tau\, K_B(1+K_Z)\, L}{\underline{C}}\,\Delta = \underbrace{\frac{\tau\, K_B(1+K_Z)\, L}{\underline{C}}\big(1+\|W_B\|_{\infty,D}\big)}_{=:C^\star}\ \mathrm{PPG}.$$

This proves $\mathrm{PDI} \le C^\star \cdot \mathrm{PPG}$. □

C Explicit counterexamples for conditions (i), (ii), (iii), and (iv)

This appendix provides explicit counterexamples showing that each of the sufficient conditions (i)–(iv) in Theorem 4 is essential. In each construction, the primitive-level discrepancy PPG can be made arbitrarily small, while the realized pathwise distortion PDI is not uniformly controlled, so that no inequality of the form $\mathrm{PDI} \le C\,\mathrm{PPG}$ can hold in general. Each example violates *exactly one* of the conditions (i)–(iv).

(i) Violation of a common admissible domain

Setup. Let the state be scalar $x \in \mathbb{R}$ with dynamics

$$\dot{x}(t) = u(t), \qquad x(0) = 0, \qquad t \in [0,T].$$

Consider identical pointwise criteria

$$W_B(t,x,u) = -u^2, \qquad W_A(t,x,u) = -u^2,$$

so that, on any common admissible domain, $W_A \equiv W_B$ and hence $\mathrm{PPG} = 0$.

Domain mismatch. Define different admissible control sets

$$\mathcal{U}_B = \{u \in L^2([0,T]) : |u(t)| \le 1 \text{ a.e.}\}, \qquad \mathcal{U}_A = \{u \in L^2([0,T]) : |u(t)| \le M \text{ a.e.}\},$$

with $M > 1$.

Failure of a uniform inequality. Although the primitive criteria coincide, the feasible trajectory sets differ. The A-admissible set contains trajectories (e.g. $x(t) = Mt$) that are infeasible under $\mathcal{U}_B$. Consequently, any pathwise comparison underlying PDI is either ill-defined or can be made arbitrarily large by letting $M \to \infty$, while $\mathrm{PPG} = 0$. Thus, without condition (i) (a common admissible domain), no meaningful uniform bound relating PDI and PPG can hold.

(ii) Violation of dynamic regularity (amplification through instability)

Setup. Consider the one-dimensional deterministic system

$$\dot{x}(t) = \alpha x(t) + u(t), \qquad x(0) = x_0,$$

with $\alpha > 0$. Let controls be unconstrained in $L^2([0,T])$ and define

$$W_B(t,x,u) = -x^2 - u^2, \qquad W_A(t,x,u) = W_B(t,x,u) + \varepsilon x,$$

where $\varepsilon > 0$ is arbitrarily small.

Primitive discrepancy. The difference $W_A - W_B = \varepsilon x$ can be made uniformly small by choosing ε small, so PPG $= O(\varepsilon)$.

Dynamic amplification. Because the dynamics are unstable ($\alpha > 0$), the optimal trajectories under W_A and W_B diverge exponentially over time. As $T \to \infty$, the realized pathwise discrepancy grows without bound, even though the primitive-level discrepancy remains small.

Failure of a uniform inequality. Hence, without condition (ii) controlling dynamic amplification, PDI cannot be bounded by any constant multiple of PPG.

(iii) Violation of uniform comparability (non-Lipschitz control distortion)

Setup. Let $\dot{x}(t) = 0$ so that $x(t) \equiv x_0$, and let $u(t) \in [-1, 1]$. Define

$$W_B(t, x, u) = -u^2, \qquad W_A(t, x, u) = -u^2 + \varepsilon |u|^p,$$

with $p \in (0, 1)$ and $\varepsilon > 0$.

Primitive discrepancy. For all $u \in [-1, 1]$,

$$|W_A - W_B| = \varepsilon |u|^p \leq \varepsilon,$$

so PPG $= O(\varepsilon)$.

Optimal controls. The B-optimal control is $u_B^* \equiv 0$. For A, the non-Lipschitz term $|u|^p$ has infinite slope at $u = 0$, and direct optimization yields a maximizer

$$u_A^* = \min\Big\{1,\ \Big(\frac{\varepsilon p}{2}\Big)^{\frac{1}{2-p}}\Big\}.$$

Failure of a uniform inequality. The realized control deviation satisfies $\|u_A^* - u_B^*\|_{L^2} \asymp \varepsilon^{1/(2-p)}$, which dominates ε as $\varepsilon \downarrow 0$. Thus PDI/PPG can diverge, and condition (iii) is essential.

(iv) Violation of a sign/monotonicity restriction (cancellation)

Setup. Let $\dot{x}(t) = 0$ and $u(t) \in [-1, 1]$. Define

$$W_B(t, x, u) = -u^2, \qquad W_A(t, x, u) = -u^2 + \varepsilon\, \sigma(t)\, u,$$

where

$$\sigma(t) = \begin{cases} +1, & t \in [0, T/2], \\ -1, & t \in (T/2, T], \end{cases}$$

and $\varepsilon > 0$.

Optimal controls and cancellation. The B-optimal control is $u_B^* \equiv 0$. For A, pointwise maximization gives $u_A^*(t) = \frac{\varepsilon}{2}\sigma(t)$. Thus the realized controls differ over the entire horizon, with $\|u_A^* - u_B^*\|_{L^2} = \frac{\varepsilon}{2}\sqrt{T}$.

Failure of a uniform inequality. Because $\sigma(t)$ changes sign, the signed primitive wedge underlying PPG can cancel over time and be made arbitrarily small, for instance by refining the sign-switching pattern. In contrast, the realized pathwise deviation remains of order ε. Hence, without a monotonicity or sign restriction as in condition (iv), PDI need not be controlled by PPG.

D Independence/necessity of key axioms (explicit counterexamples)

This appendix rewrites the compact constructions in the current draft into fully explicit counterexamples. For each key axiom, we (i) define an alternative functional that satisfies all other axioms in the corresponding system, and (ii) provide concrete primitives (W_B, W_A) on an explicit compact domain D for which the targeted axiom fails in a way that a reader can verify by direct computation.

Throughout, fix a compact domain

$$D := [0,1] \quad \text{and} \quad C(D) := \{\, \text{continuous functions } D \to \mathbb{R} \,\}.$$

When convenient, we use the shorthand

$$\Delta(W_B, W_A; D) := \inf_{\alpha>0,\ \beta\in\mathbb{R}} \|W_A - \alpha W_B - \beta\|_{\infty,D},$$
$$\mathrm{PPG}(W_B, W_A; D) := \frac{\Delta(W_B, W_A; D)}{1 + \|W_B\|_{\infty,D}}.$$

D.1 Axioms for PPG: independence of key axioms

In this subsection, "satisfies all other axioms" always refers to the PPG axioms (domain locality, benchmark/alternative affine invariance, correct zero-set, stability, and scale-free normalization) except the one explicitly targeted.

(A2) Benchmark affine invariance is independent. Define the modified functional

$$G^{(2)}(W_B, W_A; D) := \mathrm{PPG}(W_B, W_A; D)\left(1 + \left|\tanh\Big(\int_D W_B(x)\,dx\Big)\right|\right). \tag{5}$$

Why it satisfies the other axioms. The multiplier is a bounded positive scalar depending only on W_B (not on W_A), hence domain locality, alternative affine invariance, zero-set, stability (Lipschitzness up to a bounded factor), and scale-free normalization are inherited from PPG.

Concrete violation of benchmark affine invariance. Take

$$W_B(x) = x, \qquad W_A(x) = x, \qquad D = [0,1].$$

Then $\mathrm{PPG}(W_B, W_A; D) = 0$, so $G^{(2)}(W_B, W_A; D) = 0$. Now apply a benchmark shift $\widetilde{W}_B := W_B + 10$ (a positive affine change). Benchmark affine invariance would require $G^{(2)}(\widetilde{W}_B, W_A; D) = G^{(2)}(W_B, W_A; D) = 0$ for all W_A. However, the multiplier changes because

$$\int_0^1 \widetilde{W}_B(x)\,dx = \int_0^1 (x+10)\,dx = \frac{1}{2} + 10, \qquad \int_0^1 W_B(x)\,dx = \frac{1}{2},$$

so the multiplier in (5) changes. Hence $G^{(2)}$ depends on the absolute level of the benchmark representation and thus violates benchmark affine invariance.

(A3) Alternative affine invariance is independent. Define analogously

$$G^{(3)}(W_B, W_A; D) := \mathrm{PPG}(W_B, W_A; D)\left(1 + \left|\tanh(\int_D W_A(x)\,dx)\right|\right). \tag{6}$$

Other axioms. Now the bounded multiplier depends only on W_A, so domain locality, benchmark affine invariance, zero-set, stability, and normalization remain.

Concrete violation of alternative affine invariance. Take $W_B(x) = x$ and $W_A(x) = x$ as above. Again PPG $= 0$. Now apply an alternative shift $\widetilde{W}_A := W_A + 10$. The PPG value remains 0 (since preference equivalence is preserved), but the multiplier changes because $\int_D \widetilde{W}_A \neq \int_D W_A$. Thus $G^{(3)}$ violates alternative affine invariance.

(A4) Correct zero-set ("zero iff preference-equivalent") is independent. Define the identically-zero functional

$$G^{(4)}(W_B, W_A; D) := 0.$$

It satisfies locality, both affine invariances, stability, and normalization (trivially). But it fails the correct zero-set.

Concrete demonstration. Take

$$W_B(x) = x, \qquad W_A(x) = x^2, \qquad D = [0, 1].$$

These are *not* positive-affine equivalent on $[0, 1]$, so the PPG axiom requires $G(W_B, W_A; D) > 0$. Yet $G^{(4)}(W_B, W_A; D) = 0$.

(A5) Stability/Lipschitz continuity is independent. Define

$$G^{(5)}(W_B, W_A; D) := \sqrt{\mathrm{PPG}(W_B, W_A; D)}. \tag{7}$$

Other axioms. Because $x \mapsto \sqrt{x}$ is increasing with $\sqrt{0} = 0$, the quotient and zero-set axioms are preserved; locality and affine invariances are preserved by composition.

Concrete failure of Lipschitz stability at the origin. Let $W_B(x) = x$ on $[0, 1]$ and consider a perturbation family

$$W_A^{(\varepsilon)}(x) := x + \varepsilon, \qquad \varepsilon > 0.$$

Then $\Delta(W_B, W_A^{(\varepsilon)}; D) = 0$ because $W_A^{(\varepsilon)} = 1 \cdot W_B + \varepsilon$, so PPG $= 0$ for all ε, hence $G^{(5)} = 0$ as well; this family does not witness failure.

Instead, use a family that produces a small but *nonzero* PPG:

$$W_B(x) := x, \qquad W_A^{(\varepsilon)}(x) := x^2 + \varepsilon x, \qquad \varepsilon > 0.$$

One checks that $\mathrm{PPG}(W_B, W_A^{(\varepsilon)}; D) = O(\varepsilon)$ as $\varepsilon \downarrow 0$ (the best affine fit improves linearly in ε), so

$$G^{(5)}(W_B, W_A^{(\varepsilon)}; D) = \sqrt{O(\varepsilon)} = O(\sqrt{\varepsilon}),$$

which is *not* Lipschitz in ε at 0. Therefore stability in the sense of a global Lipschitz bound fails for $G^{(5)}$.

(A6) Scale-free normalization is independent. Define the unnormalized discrepancy ratio

$$G^{(6)}(W_B, W_A; D) := \frac{\Delta(W_B, W_A; D)}{\|W_B\|_{\infty, D}}. \tag{8}$$

Other axioms. Locality and affine invariances (in each argument) are inherited from Δ and the homogeneity of the uniform norm; the correct zero-set holds as long as $\|W_B\|_{\infty,D} > 0$.

Concrete blow-up when the benchmark becomes small. Let $W_B^{(\varepsilon)}(x) := \varepsilon x$ and $W_A(x) := x$ on $[0, 1]$, with $\varepsilon \downarrow 0$. Then

$$\|W_B^{(\varepsilon)}\|_{\infty,D} = \varepsilon, \qquad \Delta(W_B^{(\varepsilon)}, W_A; D) \geq \|W_A\|_{\infty,D} - \inf_{\alpha>0,\beta} \|\alpha W_B^{(\varepsilon)} + \beta\|_{\infty,D} \asymp 1,$$

so $G^{(6)}(W_B^{(\varepsilon)}, W_A; D) \asymp 1/\varepsilon \to \infty$. Thus scale-free well-posedness fails. This is exactly the reason for the $1 + \|W_B\|_{\infty,D}$ regularization in PPG.

D.2 Axioms for PDI: independence of key axioms

We now provide explicit counterexamples for each key axiom in the PDI system. We use the simplified setting where the realised paths are exogenously fixed (two candidate solutions), because independence is an axiomatic property of the *index construction* once the wedge process is given.

Fix a horizon $[0, \tau]$ with $\tau = 1$. Let the benchmark-wedge process be a measurable function $w : [0, 1] \to \mathbb{R}$ which we interpret as $w(t) = s\Delta W_A^{W_B}(t)$ in the notation of the paper, and let the baseline magnitude be a positive scalar $B_0 > 0$ (standing for $\int_0^\tau |W_B(\cdot)|dt$). The canonical PDI is

$$\mathrm{PDI}(w) := \frac{\int_0^1 [w(t)]_+ \, dt}{\max\{d, B_0\}},$$

with the nonnegativity convention $[\,\cdot\,]_+$.

(A8) Nonnegativity / "loss" convention is independent. Define the alternative functional

$$H^{(8)}(w) := \frac{\left|\int_0^1 w(t)\, dt\right|}{\max\{d, B_0\}}.$$

It preserves locality in (w, B_0), scale invariance in benchmark units, and time additivity of the numerator (as an integral), but it breaks the intended "loss" meaning and the zero-wedge axiom.

Concrete demonstration (cancellation). Take $w(t) = +1$ on $[0, \frac{1}{2}]$ and $w(t) = -1$ on $(\frac{1}{2}, 1]$. Then

$$\int_0^1 w(t)\, dt = 0 \quad \Rightarrow \quad H^{(8)}(w) = 0,$$

despite the fact that $w(t) \neq 0$ on a set of full measure. Hence "zero iff no wedge" fails without the nonnegativity convention.

(A10) Scale invariance in benchmark units is independent. Define the unnormalized cumulative wedge

$$H^{(10)}(w) := \int_0^1 [w(t)]_+ \, dt$$

(with no normalization). It preserves nonnegativity, locality in w, and time additivity, but violates scale invariance under $W_B \mapsto \gamma W_B$.

Concrete demonstration. If benchmark units are scaled by $\gamma > 0$, then w scales to γw. But $H^{(10)}(\gamma w) = \gamma H^{(10)}(w) \neq H^{(10)}(w)$ unless $\gamma = 1$.

(A11) Time-additivity of the numerator is independent. Define the quadratic aggregation

$$H^{(11)}(w) := \frac{\int_0^1 ([w(t)]_+)^2 \, dt}{\max\{d, B_0\}}.$$

It preserves nonnegativity and scale invariance (both numerator and baseline scale quadratically only if one changes the baseline accordingly; with a fixed benchmark baseline it breaks the intended axiom), but, most importantly, it destroys additivity in the sense used by the axiom.

Concrete demonstration. Let $w_1 := \mathbf{1}_{[0,1/2]}$ and $w_2 := \mathbf{1}_{(1/2,1]}$. Then the canonical numerator is additive: $\int [w_1]_+ + \int [w_2]_+ = 1$. But

$$\int ([w_1 + w_2]_+)^2 \, dt = \int 1^2 \, dt = 1,$$

$$\text{whereas} \qquad \int ([w_1]_+)^2 \, dt + \int ([w_2]_+)^2 \, dt = \frac{1}{2} + \frac{1}{2} = 1,$$

so this example does not separate them. Instead, take $w_1 := 2\mathbf{1}_{[0,1/2]}$ and $w_2 := 2\mathbf{1}_{(1/2,1]}$; then

$$\int ([w_1 + w_2]_+)^2\, dt = \int 4\, dt = 4, \qquad \int ([w_1]_+)^2\, dt + \int ([w_2]_+)^2\, dt = 2 + 2 = 4,$$

still equal. The true failure is that the aggregation depends on the *time distribution* of the wedge rather than only on the additive integral $\int [w]_+ dt$; for example, for $w^{(a)} := a\mathbf{1}_{[0,1]}$ and $w^{(b)} := 2a\mathbf{1}_{[0,1/2]}$, the canonical additive numerator coincides ($= a$) but $H^{(11)}$ differs ($= a^2$ vs. $= 2a^2$). Hence time-additivity in the axiomatic sense fails.

(A12) Floor axiom is independent. Define the ratio without floor

$$H^{(12)}(w) := \frac{\int_0^1 [w(t)]_+\, dt}{B_0}.$$

It satisfies locality, nonnegativity, scale invariance, and time additivity (whenever $B_0 > 0$), but fails well-posedness as $B_0 \downarrow 0$.

Concrete blow-up. Fix $w(t) \equiv 1$ and let $B_0 = \varepsilon \downarrow 0$. Then $H^{(12)}(w) = 1/\varepsilon \to \infty$. This is precisely the degeneracy the floor $d > 0$ is designed to prevent.

E Python script generates figures 1 – 5

```
#!/usr/bin/env python3
# -*- coding: utf-8 -*-
"""
one_shot_figs_AB_caseA_caseB.py

Generate five PDF figures with A/B (A-optimal, B-optimal) labels:
  1) fig_caseA_paths.pdf
  2) fig_caseA_wedge.pdf
  3) fig_caseB_paths.pdf
  4) fig_caseB_wedge.pdf
  5) fig_ppg_pdi_scatter.pdf

This script is self-contained (no external data). You may replace the
"stylized_case_*" blocks with your model outputs (A-optimal/B-optimal paths)
without changing the figure APIs.

Conventions:
- A corresponds to "A-optimal" (formerly SP-optimal in your draft)
- B corresponds to "B-optimal" (formerly POP-optimal in your draft)
"""

from __future__ import annotations

import os
import numpy as np
import matplotlib.pyplot as plt


# ============================================================
# 0. Utilities
# ============================================================
```

```
def ensure_dir(path: str) -> None:
    if path and not os.path.exists(path):
        os.makedirs(path, exist_ok=True)


def smooth01(x: np.ndarray, eps: float = 1e-9) -> np.ndarray:
    """Clip to [0,1] for interpretability in stylized examples."""
    return np.clip(x, 0.0 + eps, 1.0 - eps)


def moving_average(x: np.ndarray, w: int = 9) -> np.ndarray:
    """Light smoothing for stylized trajectories."""
    if w <= 1:
        return x
    w = int(w)
    pad = w // 2
    xp = np.pad(x, (pad, pad), mode="edge")
    kernel = np.ones(w) / w
    return np.convolve(xp, kernel, mode="valid")


# ============================================================
# 1. Core definitions: welfare, wedge, indices
# ============================================================

def welfare_A(T1: np.ndarray, T2: np.ndarray, D: np.ndarray, I: np.ndarray,
     params: dict) -> np.ndarray:
    """
    Welfare of criterion A evaluated pointwise along a path.
    Quadratic stylized form for demonstration; replace with your exact W^A
     if needed.

    params keys expected:
      a1,a2,aD,cI, wI (weights)
    """
    a1 = params["a1"]
    a2 = params["a2"]
    aD = params["aD"]
    cI = params["cI"]
    wI = params["wI"]
    # A prefers higher trust (T1,T2), lower debt D, and moderate policy
     intensity I
    return a1 * (T1**2) + a2 * (T2**2) - aD * (D**2) - cI * (I**2) + wI * (I
     )


def welfare_B(T1: np.ndarray, T2: np.ndarray, D: np.ndarray, I: np.ndarray,
     params: dict) -> np.ndarray:
    """
    Welfare of criterion B evaluated pointwise along a path.
    Quadratic stylized form for demonstration; replace with your exact W^B
     if needed.

    params keys expected:
      b1,b2,bD,cI, wI (weights)
    """
    b1 = params["b1"]
    b2 = params["b2"]
```

```
    bD = params["bD"]
    cI = params["cI"]
    wI = params["wI"]
    # B puts relatively more weight on short-run T1 and policy salience (wI)
     ,
    # and less penalty on D.
    return b1 * (T1**2) + b2 * (T2**2) - bD * (D**2) - cI * (I**2) + wI * (I
     )


def compute_wedge_and_indices(case: dict) -> dict:
    """
    Given A-optimal path and B-optimal path, compute:
      - sPPG(t) = W_A(A*) - W_A(B*)  (A-welfare wedge, pointwise)
      - PPG = integral of sPPG(t) dt over horizon (scalar)
      - PDI = normalized discrepancy index in [0,1] (scalar)
    Normalization uses the range of W_A values observed across both paths.
    """
    t = case["t"]
    dt = float(t[1] - t[0])

    # Paths
    XA = case["A_path"]
    XB = case["B_path"]

    # Evaluate A welfare on both paths
    WA_on_A = welfare_A(XA["T1"], XA["T2"], XA["D"], XA["I"], case["A_params
     "])
    WA_on_B = welfare_A(XB["T1"], XB["T2"], XB["D"], XB["I"], case["A_params
     "])

    sPPG = WA_on_A - WA_on_B
    PPG = float(np.sum(sPPG) * dt)

    # Normalization for PDI: scale by observed range of A-welfare across
     both paths
    W_all = np.concatenate([WA_on_A, WA_on_B])
    denom = float(np.max(W_all) - np.min(W_all))
    if denom <= 1e-12:
        PDI = 0.0
    else:
        # "Discrepancy intensity": average absolute wedge relative to scale
        PDI = float(np.sum(np.abs(sPPG)) * dt / (denom * (t[-1] - t[0])))

    out = dict(case_name=case["name"], t=t, dt=dt,
               WA_on_A=WA_on_A, WA_on_B=WA_on_B,
               sPPG=sPPG, PPG=PPG, PDI=PDI)
    return out


# ============================================================
# 2. Stylized cases (replace with your model outputs if desired)
# ============================================================

def stylized_case_A(T: int = 241, seed: int = 7) -> dict:
    rng = np.random.default_rng(seed)
    t = np.linspace(0.0, 1.0, T)
```

```
    # Baseline states (stylized)
    # T1: short-run support; T2: institutional trust; D: debt; I: policy
     intensity
    # "Case A": relatively stable environment; A-optimal is smoother/less
     volatile
    base_T1 = 0.55 + 0.10 * np.sin(2*np.pi*t)
    base_T2 = 0.60 + 0.06 * np.cos(2*np.pi*t)
    base_D  = 0.35 + 0.08 * (t - 0.5)

    # A-optimal policy: moderate, mean-reverting
    IA = 0.10 + 0.05 * np.sin(4*np.pi*t)
    IA = moving_average(IA, 11)

    # B-optimal policy: more aggressive, higher salience and more
     oscillation
    IB = 0.16 + 0.10 * np.sin(4*np.pi*t + 0.6) + 0.02 * rng.standard_normal(
     T)
    IB = moving_average(IB, 7)

    # State responses (stylized reduced-form coupling)
    T1A = smooth01(base_T1 + 0.25*IA - 0.05*base_D)
    T2A = smooth01(base_T2 + 0.10*IA - 0.04*np.maximum(0.0, IA-0.12))
    DA  = np.clip(base_D + 0.30*IA, 0.0, 1.5)

    T1B = smooth01(base_T1 + 0.30*IB - 0.05*base_D)
    T2B = smooth01(base_T2 + 0.05*IB - 0.09*np.maximum(0.0, IB-0.12))
    DB  = np.clip(base_D + 0.42*IB, 0.0, 1.5)

    case = {
        "name": "Case A",
        "t": t,
        "A_path": {"T1": T1A, "T2": T2A, "D": DA, "I": IA},
        "B_path": {"T1": T1B, "T2": T2B, "D": DB, "I": IB},
        # A-criterion parameters and B-criterion parameters
        "A_params": {"a1": 1.2, "a2": 1.6, "aD": 0.9, "cI": 0.25, "wI":
     0.05},
        "B_params": {"b1": 1.7, "b2": 1.0, "bD": 0.4, "cI": 0.12, "wI":
     0.12},
    }
    return case


def stylized_case_B(T: int = 241, seed: int = 11) -> dict:
    rng = np.random.default_rng(seed)
    t = np.linspace(0.0, 1.0, T)

    # "Case B": more fragile environment (e.g., larger debt sensitivity /
     trust erosion)
    shock = np.exp(-0.5*((t-0.55)/0.08)**2)  # transient shock
    base_T1 = 0.50 + 0.08 * np.sin(2*np.pi*t) - 0.10*shock
    base_T2 = 0.58 + 0.05 * np.cos(2*np.pi*t) - 0.14*shock
    base_D  = 0.45 + 0.12 * (t - 0.4) + 0.06*shock

    # A-optimal policy: cautious stabilization, avoids high I during shock
    IA = 0.10 + 0.06 * np.sin(4*np.pi*t) - 0.07*shock
    IA = moving_average(IA, 11)
```

```
    # B-optimal policy: strong response (high I) during shock (electoral
     salience)
    IB = 0.14 + 0.10 * np.sin(4*np.pi*t + 0.4) + 0.12*shock + 0.02*rng.
     standard_normal(T)
    IB = moving_average(IB, 7)

    # State responses (stronger negative effect on T2 via high I and debt)
    T1A = smooth01(base_T1 + 0.28*IA - 0.08*base_D)
    T2A = smooth01(base_T2 + 0.10*IA - 0.06*np.maximum(0.0, IA-0.11) - 0.04*
     base_D)
    DA  = np.clip(base_D + 0.34*IA, 0.0, 2.0)

    T1B = smooth01(base_T1 + 0.32*IB - 0.08*base_D)
    T2B = smooth01(base_T2 + 0.04*IB - 0.13*np.maximum(0.0, IB-0.11) - 0.05*
     base_D)
    DB  = np.clip(base_D + 0.55*IB, 0.0, 2.0)

    case = {
        "name": "Case B",
        "t": t,
        "A_path": {"T1": T1A, "T2": T2A, "D": DA, "I": IA},
        "B_path": {"T1": T1B, "T2": T2B, "D": DB, "I": IB},
        "A_params": {"a1": 1.2, "a2": 1.8, "aD": 1.1, "cI": 0.28, "wI":
     0.05},
        "B_params": {"b1": 1.8, "b2": 0.9, "bD": 0.35, "cI": 0.10, "wI":
     0.14},
    }
    return case


# ============================================================
# 3. Plotting
# ============================================================

def plot_paths(case: dict, out_pdf: str) -> None:
    t = case["t"]
    XA, XB = case["A_path"], case["B_path"]

    fig = plt.figure(figsize=(9.2, 6.8))

    # Panel 1: T1 and T2
    ax1 = fig.add_subplot(2, 1, 1)
    ax1.plot(t, XA["T1"], label="T1 (A-optimal)")
    ax1.plot(t, XB["T1"], label="T1 (B-optimal)")
    ax1.plot(t, XA["T2"], label="T2 (A-optimal)")
    ax1.plot(t, XB["T2"], label="T2 (B-optimal)")
    ax1.set_ylabel("Trust / Support (scaled)")
    ax1.set_title(f"{case['name']}: A-optimal vs B-optimal paths")
    ax1.grid(True, alpha=0.3)
    ax1.legend(loc="best", frameon=True)

    # Panel 2: D and I
    ax2 = fig.add_subplot(2, 1, 2)
    ax2.plot(t, XA["D"], label="Debt D (A-optimal)")
    ax2.plot(t, XB["D"], label="Debt D (B-optimal)")
    ax2.plot(t, XA["I"], label="Policy I (A-optimal)")
    ax2.plot(t, XB["I"], label="Policy I (B-optimal)")
    ax2.set_xlabel("Normalized time")
```

```
    ax2.set_ylabel("D, I (scaled)")
    ax2.grid(True, alpha=0.3)
    ax2.legend(loc="best", frameon=True)

    fig.tight_layout()
    fig.savefig(out_pdf, bbox_inches="tight")
    plt.close(fig)


def plot_wedge(case: dict, out_pdf: str) -> dict:
    res = compute_wedge_and_indices(case)
    t = res["t"]

    fig = plt.figure(figsize=(9.2, 5.6))
    ax = fig.add_subplot(1, 1, 1)

    ax.plot(t, res["sPPG"], label="sPPG(t) = W_A(A*) - W_A(B*)")
    ax.axhline(0.0, linewidth=1.0)
    ax.set_title(f"{case['name']}: A-welfare wedge and indices")
    ax.set_xlabel("Normalized time")
    ax.set_ylabel("A-welfare wedge (pointwise)")
    ax.grid(True, alpha=0.3)

    # Text box with scalar indices
    txt = (
        f"PPG =  sPPG(t) dt   {res['PPG']:.4f}\n"
        f"PDI = normalized intensity  {res['PDI']:.4f}"
    )
    ax.text(
        0.02, 0.98, txt,
        transform=ax.transAxes,
        va="top", ha="left",
        bbox=dict(boxstyle="round", alpha=0.15)
    )

    ax.legend(loc="best", frameon=True)
    fig.tight_layout()
    fig.savefig(out_pdf, bbox_inches="tight")
    plt.close(fig)

    return res


def plot_scatter(results: list[dict], out_pdf: str) -> None:
    # results: list of dicts from compute_wedge_and_indices
    fig = plt.figure(figsize=(7.6, 5.8))
    ax = fig.add_subplot(1, 1, 1)

    # scatter points
    for r in results:
        ax.scatter([r["PPG"]], [r["PDI"]], label=r["case_name"], s=70)

    ax.set_xlabel("PPG (scalar)")
    ax.set_ylabel("PDI (scalar)")
    ax.set_title("PPGPDI scatter (Cases A and B)")
    ax.grid(True, alpha=0.3)
    ax.legend(loc="best", frameon=True)
```

```
    fig.tight_layout()
    fig.savefig(out_pdf, bbox_inches="tight")
    plt.close(fig)


# ============================================================
# 4. Main
# ============================================================

def main(output_dir: str = ".", T: int = 241) -> None:
    ensure_dir(output_dir)

    caseA = stylized_case_A(T=T, seed=7)
    caseB = stylized_case_B(T=T, seed=11)

    # 1) Paths
    plot_paths(caseA, os.path.join(output_dir, "fig_caseA_paths.pdf"))
    plot_paths(caseB, os.path.join(output_dir, "fig_caseB_paths.pdf"))

    # 2) Wedge
    resA = plot_wedge(caseA, os.path.join(output_dir, "fig_caseA_wedge.pdf")
     )
    resB = plot_wedge(caseB, os.path.join(output_dir, "fig_caseB_wedge.pdf")
     )

    # 3) Scatter
    plot_scatter([resA, resB], os.path.join(output_dir, "fig_ppg_pdi_scatter.
     pdf"))

    print("Saved:")
    print(" - fig_caseA_paths.pdf")
    print(" - fig_caseA_wedge.pdf")
    print(" - fig_caseB_paths.pdf")
    print(" - fig_caseB_wedge.pdf")
    print(" - fig_ppg_pdi_scatter.pdf")


if __name__ == "__main__":
    # Colab/Jupyter-friendly defaults: just run the cell.
    main(output_dir=".", T=241)
```