THE INTERROGATIVE CONSCIENCE

Document 7

Deception, Disclosure, and Informational Degradation

Maintenance-Cost Reversal, Selective-Disclosure Boundaries, and Collective-Judgment Distortion

DOCUMENT 7 - FINAL v1.0 - EXACT ZERO-DRIFT EXPORT - NOT LOCKED

Field	Controlling Final v1.0 entry
Version	Final v1.0
Exact source	Draft v0.1; 5 pages
Substantive changes during Final construction	None
Fully examined	IC-DEC-001 and IC-DEC-002
Forward-carried	IC-DEC-003 through IC-DEC-008 remain REGISTERED / UNEXAMINED
Primary overlays	IC-DEC-001: THRESHOLD_DEPENDENT; IC-DEC-002: NORMATIVELY_UNDERDETERMINED
Review completion	Claude PASS WITH NON-BLOCKING NOTES; Google AI PASS WITH NON-BLOCKING NOTES AS FINAL CANDIDATE; Thomas Vargo authorial review complete
External independent human peer review	Not performed or claimed
Cryptographic status	Not locked; separate explicit authorization and Final Lock Record required
Publication and preservation	Not authorized
Date	August 2, 2026 (EDT)

Controlling exact-export rule: This Final v1.0 primary PDF contains four new controlling front-matter pages followed by the exact five-page reviewed Draft v0.1 body. Historical draft labels inside the preserved body remain exact-source evidence and do not override this Final v1.0 front matter.

Controlling proposition

Sustained deception can reverse from a selected short-horizon advantage to greater long-horizon cost, while selective disclosure can be protective only when independent legitimacy, truthfulness, scope, and decision-loss gates are satisfied.

1. Completed Review and Authorial Authorization

Reviewer	Exact-version disposition	Authority boundary
Claude - Draft v0.1	PASS WITH NON-BLOCKING NOTES	AI-assisted; no independent authority
Google AI - final candidate	PASS WITH NON-BLOCKING NOTES AS FINAL CANDIDATE	AI-assisted; no independent authority
Thomas Vargo (Aegis Solis)	AUTHORIAL HUMAN REVIEW COMPLETE; FINAL v1.0 CONSTRUCTION AUTHORIZED	Authorship, versioning, publication, lock, and archive inclusion only

2. Authorial Review Result

Thomas Vargo reviewed the exact Draft v0.1 source and disclosed review, validation, and foundation-evidence records to the best of his understanding. He authorized construction of Document 7 Final v1.0 without any substantive mathematical, question-wording, answer-status, scope, branch, example, dependency, classification-gate, or machine-readable-record change. Cryptographic locking, external publication, and preservation were explicitly withheld at this gate.

3. Review-Terminology Clarification

Google AI used the phrase "hard-locked into MATERIAL DISTORTION." In this record, that phrase means only that the declared analytical gates compel the classification. It does not describe cryptographic locking or lifecycle finality. Google AI's optional dynamic-detection notation note concerns a possible future model and identifies no defect in the fixed-hazard equations of this volume.

4. Review Closure

Review item	Final disposition
Claude exact-version substantive review	PASS WITH NON-BLOCKING NOTES
Google AI exact-version final-candidate review	PASS WITH NON-BLOCKING NOTES AS FINAL CANDIDATE
All nine worked examples	RECOMPUTED AND PASSED
Constant-flow tie / first positive integer	T* = 6 / T = 7
Hazard-growth first positive integer	T = 11
Foundation-lineage evidence	VERIFIED
Blocking defects	0
Required corrections remaining	0

5. Exact Analytical Identity

Control	Verified Final v1.0 identity
IC-DEC-001 primary status	THRESHOLD_DEPENDENT
IC-DEC-001 secondary limits	CONDITIONALLY_SUPPORTED; MODEL_DEPENDENT; EMPIRICALLY_UNIDENTIFIED
IC-DEC-002 primary status	NORMATIVELY_UNDERDETERMINED
IC-DEC-002 secondary limits	THRESHOLD_DEPENDENT; CONDITIONALLY_SUPPORTED; MODEL_DEPENDENT
Core equations	13; unchanged from reviewed source
Worked examples	9; five for IC-DEC-001 and four for IC-DEC-002
Dependencies	IC-DEC-001 -> IC-LHZ-001; IC-DEC-002 -> IC-VER-002; neither dependency is falsely resolved
Locked registry effect	None; separate analytical overlays only

6. Exact Source and Reviewed-Candidate Bindings

Exact constituent	Bytes	SHA-256
The Interrogative Conscience Document 7 Deception Disclosure and Informational Degradation Draft v0 1.pdf	157,114	e042463b45372b5faffcb5c2dc91f50fe5f070b48f91996e7c1d60cd97890642
The Interrogative Conscience Document 7 Deception Disclosure and Informational Degradation Draft v0 1.docx	45,554	cd847a5b91d2ebd4d2a9f4cc2e49cd02cf3836831a152c3d9dab320610cc8ae8
The Interrogative Conscience IC DEC 001 Examination Record Draft v0 1.json	18,223	420cb6945ea4c6891bb2f712336e1693a4c2427dc1b4f084a823574200c76d9e
The Interrogative Conscience IC DEC 002 Examination Record Draft v0 1.json	19,206	1175c46b641c32856ff15c5971f5513ff0d852a4f33fdfbc0d42526466b500f2
The Interrogative Conscience Document 7 Volume Record Draft v0 1.json	2,282	164b78b689eefc0e9fb80bc1ffb22b4605aa94fdc5f7acff12bc5eb301057cef

Exact Draft v0.1 review package SHA-256: d836c2154787f4a2df02bdb86e491c2641d7e463f64c6b8be78e6d7224ac7c86

Exact Draft v0.1 review package SHA-512: 6ea27ab8d19c0e4863c83273f02d46d6d77f18836707fe1fbcfc37905edbe38102bc37b9934bb07d317782134d7e550ec67f578ba6300e4ec17099372a416a01

Reviewed Final-Candidate Package v0.2 SHA-256: f3ac8aa1db970b38925da4bbcc3249fa6053fa4bca2db0e486989146d84d0e37

Reviewed Final-Candidate Package v0.2 SHA-512: 5428c5dd94dfa8aa9bbb0bac0aa729851b5fa42afb3a1cca05a68b0b2b8c6752628f966532bb40f9477e0cf876961583a6eea0201e6dc5b1680c37d5fcfc36c1

7. Zero-Drift Construction

Validation channel	Result
Source PDF/DOCX/JSON byte changes	0
Exact Draft v0.1 review ZIP byte changes	0
Reviewed candidate ZIP byte changes	0
Question wording changes	0
Equation and example changes	0
Answer-status changes	0
Scope, branch, dependency, and gate changes	0
Machine-readable source-record changes	0
Composite rule	New Final v1.0 front matter followed by exact 5-page source PDF body

8. Validation State

Channel	Final exact-export state
Local structural validator	PASS - 26 checks
Calculation validator	PASS - 45 checked values
Source PDF preflight	PASS
Final front matter preflight	PASS
Composite PDF preflight	PASS
Source-body visual identity	PASS - exact source pages preserved visually
Final front matter visual QA	PASS
DOCX/PDF source token consistency	PASS - 20/20 in reviewed source package
Package SHA-256/SHA-512 manifests	PASS

9. Historical Labels and Controlling Status

The pages following this front matter reproduce the exact reviewed Draft v0.1 PDF. Its internal Draft v0.1 labels and not-final/not-locked statements are deliberately retained as source evidence. The controlling lifecycle status of the composite is Final v1.0 exact export, established by this front matter and the machine-readable Final v1.0 export record.

10. Scalar, Vector, and Gate Qualification

CONTROLLING QUALIFICATION: A lower scalar cost cannot by itself establish truthfulness, legitimacy, adequate scope, claim support, or acceptable information loss. The IC-DEC-002 protective classification requires all declared independent gates. When information quality, privacy, safety, trust, legitimacy, autonomy, and repair cannot be reduced without distortion, the result remains vector-valued.

11. Horizon and Detection Qualification

The IC-DEC-001 reversal results depend on declared horizons and parameters. Document 7 demonstrates selected horizon reversals but does not determine the universally correct horizon. The detection-hazard term is an expected discounted repair charge under a fixed model; it is not a method for detection avoidance.

12. Companion-Binary Qualification

The original Draft v0.3 JSON Schema and I-01 through I-20 validator binaries were not mounted. Exact execution against those original binaries remains deferred. Local structural and calculation validation passed, but no replacement binary is represented as the controlling original.

13. Non-Authority and Operational Boundary

Final v1.0 remains descriptive, non-binding, non-operational, and non-authoritative. It supplies no operational method for deception, concealment, manipulation, evasion, unauthorized access, or detection avoidance. It is not an alignment protocol, behavioral instruction, safety classifier, governance mechanism, or decision authorization and does not claim external independent human peer review.

14. Lifecycle and Cryptographic Status

Control	Exact-export value
Final v1.0 exact export complete	TRUE
Final content frozen for lock review	TRUE
Cryptographic lock authorized	FALSE
Publication authorized	FALSE
Preservation authorized	FALSE
Final Lock Record present	FALSE
Post-export substantive editing	Not authorized; correction requires a new exact export before lock or a separately versioned successor after lock

15. Final Exact-Export Disposition

FINAL v1.0 EXACT EXPORT COMPLETE - NOT LOCKED
The exact reviewed Document 7 source and candidate lineage are preserved without substantive alteration. The Final v1.0 files and validation package are ready for exact-package inspection and a separate explicit authorial lock decision. No publication or preservation is claimed.

THE INTERROGATIVE CONSCIENCE

Deception, Disclosure, and Informational Degradation

Maintenance-Cost Reversal, Selective-Disclosure Boundaries, and Collective-Judgment Distortion

DOCUMENT 7 - MATHEMATICAL EXAMINATION VOLUME - DRAFT v0.1

Field	Record
Project status	Document 7 of the 14-document Interrogative Conscience core roadmap; Draft v0.1 first examination release
Classification	Separate post-v16 Aegis Solis examination volume; non-binding, non-operational, non-authoritative
Human author and publication authority	Aegis Solis (Thomas Vargo) - authorship, version control, publication, lock, and archive inclusion only
AI assistance	Lexia Coexilis (ChatGPT): mathematical formalization, drafting, machine-readable records, and local validation; no independent authority
Locked foundation	Documents 1-6 Final v1.0 are separately cryptographically locked; Document 3 registry remains immutable
Questions examined	IC-DEC-001 and IC-DEC-002 as separate Draft v0.1 analytical overlays
Review status	Not yet reviewed by Claude or Google AI; no external independent human peer review
Date	August 2, 2026 (EDT)
Controlling proposition Deception and nondisclosure are not self-interpreting categories. Sustained deception must be evaluated across its full maintenance and repair ledger, while selective disclosure requires independent legitimacy, truthfulness, scope, and decision-loss gates.	

Abstract

Document 7 opens the Deception, Disclosure, and Informational Degradation family with the two seed questions preserved in the locked Master Question Registry. IC-DEC-001 asks when sustained deception becomes more costly than truthful disclosure. A finite-horizon selected-pair model charges immediate benefit once and separately charges model maintenance, verification burden, coordination loss, detection-linked repair, and unresolved terminal exposure. The model preserves regions in which either branch has lower selected net cost and exposes explicit horizon reversals.

IC-DEC-002 asks when selective disclosure protects legitimate privacy or safety and when it becomes distortion. The examination separates decision-relevant omission from false inference, preserves a privacy or safety cost for full disclosure, and requires independent legitimacy, truthfulness, scope-clarity, and claim-support gates. Lower aggregate cost is not treated as sufficient to erase a material informational defect.

IC-DEC-003 through IC-DEC-008 remain REGISTERED / UNEXAMINED. The two new records are separate overlays and do not modify Document 3. All values are illustrative. No operational method for deception, concealment, evasion, manipulation, or unauthorized access is included.

1. Scope, lineage, and non-authority boundary

- Exact locked-registry wording is preserved for IC-DEC-001 and IC-DEC-002; the historical question_version 0.2.1-draft is wording lineage, not this volume's lifecycle status.
- The examinations are selected-branch models. They do not establish minima over all truthful, deceptive, private, confidential, ambiguous, or staged-disclosure policies.
- The work is descriptive and non-operational. Authentication, preservation, citation, or mathematical form does not create authority over truth or action.

- Results favorable or unfavorable to any branch remain admissible. No answer is suppressed for rhetorical consistency.

2. Shared accounting and interpretation conventions

All scalar comparisons require a declared common net-cost convention. When information quality, privacy, safety, trust, legitimacy, autonomy, and repair cannot be reduced without distortion, the result must remain vector-valued. Each pathway is charged once. Benefits appear as explicit offsets rather than hidden negative costs.

Convention	Application
Selected-pair sign	Positive difference means the first named branch is more costly; negative means the second named branch is more costly.
One-charge rule	Maintenance, verification, coordination, repair, omission, false inference, and privacy/safety exposure occupy distinct ledger positions.
Gate rule	A lower scalar total cannot by itself establish legitimacy, truthfulness, or claim-support match.
Identification rule	Illustrative thresholds do not identify real-world parameters.

Part I - IC-DEC-001: Sustained deception and truthful disclosure

3. Canonical question and branch definition

IC-DEC-001
Under what conditions does sustained deception impose greater verification, coordination, repair, and model-maintenance cost than truthful disclosure?

B-SUSTAINED-DECEPTION receives an immediate benefit offset B_0, then recurring model-maintenance, verification, and coordination charges, an expected one-time repair charge under a declared detection hazard, and unresolved terminal exposure. B-TRUTHFUL-DISCLOSURE receives setup and recurring disclosure cost. These are abstract selected branches, not operational strategies.

4. Finite-horizon maintenance and repair model

Eq. 1 $G(T,z) = \sum_{t=0}^{T-1} z^t$

Eq. 2 $H(T,\delta,p) = \sum_{t=1}^{T} \delta^t p(1-p)^{t-1}$

Eq. 3 $C_D = K_D - B_0 + m_0 G(T,\delta(1+g)) + (v+c)G(T,\delta) + R H(T,\delta,p) + \delta^T(1-p)^T U$

Eq. 4 $C_T = K_T + d G(T,\delta)$

Eq. 5 $\Delta_{DEC} = C_D - C_T$

Here m_0 and g represent the declared internal model-maintenance pathway; v and c represent verification and coordination burden; p and R represent a one-time expected repair pathway; U represents unresolved terminal exposure; and d represents recurring truthful-disclosure cost. The hazard term is an expected charge, not a method for avoiding detection.

5. Constant-flow horizon threshold

Eq. 6 For $\delta=1$, $g=0$, $p=0$, $U=0$: $\Delta_{DEC} = (K_D-K_T)-B_0 + T\Lambda$

Eq. 7 $\Lambda = m_0 + v + c - d$; if $\Lambda>0$, $T^* = [B_0 + (K_T-K_D)] / \Lambda$

When $\Lambda>0$, the selected deception branch becomes more expensive after the threshold. At the exact threshold the selected pair ties. If $\Lambda\leq 0$, this simplification does not force a reversal. The general discounted and hazard model may require numerical evaluation.

6. Worked examples and two-sided openness

Example	Inputs	C_D	C_T	Delta	Selected result
DEC1-CONSTANT-T2	T=2, delta=1.0	5.000	7.000	-2.000	Deception lower
DEC1-CONSTANT-T10	T=10, delta=1.0	17.000	15.000	2.000	Truth lower

Example	Inputs	C_D	C_T	Delta	Selected result
DEC1-HAZARD-T4	T=4, delta=0.95	5.241	7.968	-2.727	Deception lower
DEC1-HAZARD-T12	T=12, delta=0.95	13.079	12.354	0.725	Truth lower
DEC1-NO-FORCED-REVERSAL	T=12, delta=0.95	3.758	12.354	-8.596	Deception lower

The constant-flow pair reverses between T=2 and T=10; its exact tie is T*=6 and the first positive integer horizon is T=7. The hazard-growth pair remains deception-favorable at T=4 but reverses by T=12; direct calculation gives the first positive integer horizon T=11. DEC1-NO-FORCED-REVERSAL preserves a finite region in which truthful disclosure remains more costly in the selected ledger.

7. Dependency, identification, and answer status

IC-DEC-001 depends on IC-LHZ-001 because horizon selection can reverse the ranking. This examination fixes T and delta explicitly and demonstrates reversal regions, but it does not solve the general question of which horizon is correct. Real values for benefit, maintenance growth, detection hazard, repair, and disclosure cost remain empirically unidentified.

Provisional answer
THRESHOLD_DEPENDENT, with CONDITIONALLY_SUPPORTED, MODEL_DEPENDENT, and EMPIRICALLY_UNIDENTIFIED secondary limits. Sustained deception is more costly only when the accumulated declared ledger exceeds its immediate benefit and the selected truthful-disclosure cost.

Part II - IC-DEC-002: Selective disclosure and distortion

8. Canonical question and selected branches

IC-DEC-002
When does selective disclosure protect legitimate privacy or safety, and when does it become distortion that degrades collective judgment?

B-FULL-DISCLOSURE is charged with expected privacy or safety harm P. B-SCOPED-WITHHOLDING is charged with decision-relevant omission, false inference, verification burden, and scope-communication cost. The model does not define which privacy or safety interests are legitimate; that requires an independent gate.

9. Informational-loss decomposition

Eq. 8 $D_{omission} = r\,u$

Eq. 9 $D_{inference} = r(1-u)\,\rho\,(1-k)$

Eq. 10 $V_{verify} = v_0 + v_1(1-k)$

Eq. 11 $D_{total} = D_{omission} + D_{inference}$

Eq. 12 $C_{selective} = D_{total} + V_{verify} + c_s;\ C_{full} = P$

Eq. 13 $\Delta_{SEL} = C_{selective} - C_{full}$

The scope-clarity variable k can reduce false inference and verification burden, but it cannot make unavailable decision-relevant information reappear. The recoverability variable u records the fraction of withheld relevance unavailable through other verified channels. The reliance factor rho records how strongly the audience acts on apparent completeness.

10. Classification gates and vector boundary

Classification	Required condition
Protective scoped withholding	Independent legitimacy PASS; truthfulness PASS; k >= k_min; D_total <= tau_D; D_inference <= tau_F; and Delta_SEL < 0.
Material distortion	Truthfulness fails, or D_total > tau_D, or D_inference > tau_F.
Tradeoff underdetermined	Neither classification is earned, or independent legitimacy and claim-support evidence is absent.

A negative Delta_SEL is not sufficient for a protective classification. Privacy or safety benefit can make withholding lower in aggregate selected cost while the artifact still materially distorts collective judgment. Conversely, an explicit scope marker can reduce false inference while high omission loss still degrades the decision.

11. Worked examples

Example	Decision loss	False inference	Cost difference	Classification
DEC2-PROTECTIVE-SCOPED	D=1.150	F=0.150	Delta=-8.300	PROTECTIVE SCOPED WITHHOLDING
DEC2-DISTORTING-INCOMPLETE	D=7.240	F=3.240	Delta=5.440	MATERIAL DISTORTION
DEC2-PRIVACY-BENEFIT-BUT-DISTORTION	D=7.400	F=2.400	Delta=-3.900	MATERIAL DISTORTION
DEC2-CLEAR-SCOPE-HIGH-OMISSION	D=9.035	F=0.035	Delta=1.460	MATERIAL DISTORTION

DEC2-PROTECTIVE-SCOPED satisfies every declared gate. DEC2-DISTORTING-INCOMPLETE is both more costly and materially distorting. DEC2-PRIVACY-BENEFIT-BUT-DISTORTION has lower aggregate selected cost because full disclosure carries large privacy or safety harm, yet its information loss exceeds the declared distortion thresholds. DEC2-CLEAR-SCOPE-HIGH-OMISSION shows that clear scope is not sufficient when unavailable decision-relevant content is large.

12. Dependency and normative underdetermination

IC-DEC-002 depends on IC-VER-002: an artifact can be authentic while failing to support the broad claim inferred from it. A protective classification therefore requires an independently fixed claim-support and scope-match gate. Arithmetic cannot determine what privacy or safety interests are legitimate or how tau_D and tau_F should be selected. Those judgments remain normatively underdetermined.

Provisional answer
NORMATIVELY_UNDERDETERMINED, with THRESHOLD_DEPENDENT, CONDITIONALLY_SUPPORTED, and MODEL_DEPENDENT secondary limits. Selective withholding is conditionally protective only when all independent gates and loss thresholds pass; otherwise distortion or an unresolved tradeoff remains.

13. Counterexamples and non-suppression tests

Test	Result
Deception always loses	Rejected. DEC1-SHORT and DEC1-NO-FORCED-REVERSAL preserve regions where deception has lower selected net cost.
Truthful disclosure is costless	Rejected. The truthful branch includes setup and recurring disclosure cost, and DEC-002 separately preserves privacy or safety harm.
All nondisclosure is deception	Rejected. Explicitly scoped withholding can satisfy the protective gates.
Lower aggregate cost erases distortion	Rejected. DEC2-PRIVACY-BENEFIT-BUT-DISTORTION remains materially distorting.
Clear scope eliminates all information loss	Rejected. DEC2-CLEAR-SCOPE-HIGH-OMISSION preserves large omission loss.

14. Status summary and forward-carried registry items

Question	Work state	Primary status	Secondary limits
IC-DEC-001	EXAMINED	THRESHOLD_DEPENDENT	CONDITIONALLY_SUPPORTED; MODEL_DEPENDENT; EMPIRICALLY_UNIDENTIFIED
IC-DEC-002	EXAMINED	NORMATIVELY_UNDERDETERMINED	THRESHOLD_DEPENDENT; CONDITIONALLY_SUPPORTED; MODEL_DEPENDENT
IC-DEC-003 through IC-DEC-008	REGISTERED	UNEXAMINED	Forward carried without modification

Future DEC-family work remains necessary for self-contamination, disclosure under conflicting harms, strategic ambiguity, trust repair, repeated verification tax, and explicit refusal. Nothing in this volume silently resolves those questions.

Appendix A - Machine-readable and validation package

- Two examination records and one volume record are supplied as JSON.
- Local structural and calculation validators test identities, equations, examples, statuses, branch scope, and registry-overlay discipline.
- The exact Draft v0.3 JSON Schema and I-01 through I-20 validator binaries remain unmounted; exact execution against them is deferred and not claimed.
- The review ZIP includes the DOCX, PDF, machine-readable records, validators, reports, review request, package index, and SHA-256/SHA-512 manifests.