

Intelligent Epistemology

MU and Epistemic Zero

Emad Mostaque

Intelligent Internet

August 2026

Abstract

MU begins with the thinnest fact on which reasoning can stand: consistent inference is possible. The proof takes one line, since for any consistent proposition P , reflexivity gives $P \vdash P$. Everything else turns on the relation the word *from* names: a conclusion follows from its grounds exactly insofar as every answer-relevant dependence is carried by those grounds.

Let $q : \mathcal{X} \rightarrow \mathcal{P}$ identify different presentations of the same problem. Dependence on the problem then has the exact form

$$F = \widehat{F} \circ q.$$

Equivalent descriptions therefore agree, and labels, coordinates, search order, priors, and other working choices contribute to the result precisely when the problem makes them part of its grounds. Surviving possibilities form the full answer, while empty and unattained cases carry their certificates. Together these form one discipline, Epistemic Zero: assume nothing beyond what the constraints demand, and surrender nothing they deliver.

Write the full result of a problem Π as $\text{Answer}(\Pi)$. MU then has four quantitative forms, each unique in its domain. In a determinate Boolean plausibility domain it forces probability up to order-preserving regratuation. Applied to informational change, it forces exact marginal–conditional decomposition and therefore relative entropy up to scale. Applied to starting belief, it forces relative-entropy projection from every permitted reference, with Shannon maximum entropy as the finite counting-symmetry form. Applied through time, it forces retention of every feature compatible with new constraints and therefore KL projection. The four laws are the arithmetic of Epistemic Zero, and beyond them values and costs govern choice while channels carry the connection to the world.

The same method resolves, reconstructs, and locates the classical problems of epistemology and scientific reasoning. Regress, rule-following, indifference, confirmation, induction, Gettier cases, scepticism, disagreement, scientific realism, causal inference, the norms of inquiry, and the value of knowledge become instances of one question: what full answer do their grounds support?

$$\mathbf{MU} : \quad \Pi \longmapsto \text{Answer}(\Pi).$$

Contents

I	MU and the Logic of Inference	3
1	What It Means to Follow From	3
2	The MU Theorem	4
3	MU Generates the Method	4
II	The Ground and Its Limits	9
4	The Anatomy of an Inference Episode	9
5	The Foundation Theorem	10
6	Familiar Forms of Complete Answers	10
III	The Mathematics of Rational Belief	13
7	Three Distinctions	13
8	The Quantitative Laws of MU	14
9	Probability	16
10	Relative Entropy	20
11	Starting Beliefs: Reference Measures and Maximum Entropy	24
12	Updating	28
13	Full Answers Can Be Sets	31
14	From Belief to Choice	33
15	Summary of MU's Quantitative Forms	35
16	Evidence Channels	36
17	Convergence to Truth	39
18	What Each Condition Selects	42
IV	Classical Problems: Resolutions, Reconstructions, and Boundaries	43
19	Foundations, Rules, and Experience	43
20	Belief, Uncertainty, and Action	46
21	Induction, Confirmation, and Explanation	52
22	Knowledge, Testimony, and Scepticism	56
23	What the Arguments Establish	60
V	Science, Computation, and Society	64
24	Criticism, Falsification, and Confirmation	64
25	Underdetermination, Realism, and Theory Change	64
26	Bounded Rationality	66
27	Social Epistemology	67
28	Artificial Intelligence and Causal Models	68
29	A Rational Reconstruction of Scientific Method	70
VI	Norms and the Value of Knowledge	73
30	Belief, Choice, and Pragmatic Reasons	73
31	Norms of Inference	73
32	The Value of Knowledge	74
33	Internalism and Externalism	75
	Conclusion	77
	References	78

Part I

MU and the Logic of Inference

When you know a thing, to hold that you know it; and when you do not know a thing, to allow that you do not know it: this is knowledge.

Confucius, *Analects* 2.17, tr. Legge

1 What It Means to Follow From

1.1 A theory of principle

Every argument carries a promise: the conclusion comes from its grounds. A constructive account would build reasoning from proposed mechanisms. A theory of principle begins instead with a constraint that every genuine inference must satisfy and asks which forms meet it. The inquiry proceeds in that form: state the problem, determine whether an answer exists, classify every answer up to the relevant equivalence, and require a mere redescription to leave the result unchanged. Results are graded by their statement form throughout: theorems, propositions, and corollaries are proved; arguments are defended on stated grounds without claim of proof; remarks locate, delimit, and connect.

1.2 The word *from*

The argument turns on *from*. If an answer is presented as following from its grounds, any feature that changes the answer must either appear among those grounds or remain visible as unresolved freedom. An answer that changes with an unstated label, coordinate, ordering, prior, or preference draws on more than the grounds supplied. The word is the ordinary marker of a language-independent relation, and the commitment attaches to that relation rather than to the vocabulary.

MU is the floor: consistent inference exists. Because MU is about inference, we can then ask what any inference commits itself to when it claims that a conclusion follows from its grounds.

Definition 1.1 (Inferential relation and inference episode). An **inferential relation** is a reflexive, rule-governed relation under which conclusions follow from premises. An **inference episode** is an actual use of such a relation by a person, institution, process, or machine. The formal core concerns the relation; claims about agents concern episodes.

Definition 1.2 (Consistent inference). A consistent inference is a valid premise–conclusion instance with a consistent premise set in the chosen consequence relation.

Definition 1.3 (MU). MU is the principle that consistent inference is possible. Let $CI(I)$ mean that I is a valid inference from a consistent premise set. In symbols,

$$\boxed{\text{MU} \equiv \exists I \text{ CI}(I).}$$

The formula says that at least one such passage exists. The MU Theorem proves that it does. The Dependence Theorem then unfolds what any such passage commits itself to when its conclusion is said to come from its grounds. These are two movements of one principle: existence, then form.

2 The MU Theorem

Theorem 2.1 (Existence of consistent inference). *Every reflexive inferential relation whose language contains at least one consistent proposition contains a consistent inference.*

Proof. Let P be a consistent proposition. Reflexivity gives the valid instance $P \vdash P$. Its premise set is consistent, so it is a consistent inference and MU follows. $\square$

The identity instance establishes existence and is discharged by existential introduction.

Corollary 2.2 (No consistent denial). *In any consequence structure satisfying the hypotheses of Theorem 2.1 whose language contains the proposition $\neg \mathbf{MU}$ that no consistent inference exists, $\neg \mathbf{MU}$ has no consistent assertion: if $\neg \mathbf{MU}$ were consistent, reflexivity would make the identity inference from $\neg \mathbf{MU}$ to $\neg \mathbf{MU}$ a consistent inference, witnessing the existence it denies.*

The proof is the theorem's own: identity supplies the witness and the denial supplies the premise, so the denial exhibits the result's self-referential stability. The remainder is independent of the particular proposition P . The witness leaves behind the relation it exhibits: a conclusion that comes from grounds depends on those grounds.

3 MU Generates the Method

An inference claims that its conclusion follows from its grounds. Once that dependence is made explicit, the answer varies exactly with information the problem retains. Every answer-relevant feature appears in the problem, and a complete report preserves everything the problem supports.

A problem includes the language in which the grounds are expressed, the candidate answers, the rule of consequence, the question being asked, the required form of reply, and what counts as the same answer.

Definition 3.1 (Inference problem). An inference problem is stated by the tuple

$$\Pi = (\mathcal{L}, \mathcal{H}, \mathcal{C}, \vdash, \mathcal{Q}, \mathcal{O}, \sim),$$

where $\mathcal{L}$ is the representation language, $\mathcal{H}$ the candidate domain, $\mathcal{C}$ the premises and constraints, $\vdash$ the consequence relation, $\mathcal{Q}$ the question, $\mathcal{O}$ the required form of answer, and $\sim$ the rule for when two outputs count as the same answer. Equivalent descriptions of the same structure represent the same problem. Throughout the paper, *the problem* means this whole specification.

Every choice carries its grounds. A component of Π may be stipulated, inherited from a domain, or supported by a further argument. Writing it down identifies its role and exposes the grounds that support it. A language, consequence relation, reference measure, output form, or equivalence offered as a conclusion becomes the answer to a higher-order problem and follows the same dependence rule.

A selector may be included in the problem. The point it chooses is then conditional on that selector, whose own standing depends on its grounds.

Definition 3.2 (Presentations and problems). Let $\mathcal{X}$ be the class of working presentations and $\mathcal{P}$ the class of problems. A surjection

$$q : \mathcal{X} \longrightarrow \mathcal{P}$$

identifies presentations of the same problem by forgetting labels, coordinates, enumeration order, temporary algorithms or gauges, and any other detail the problem treats as irrelevant. For each $\Pi \in \mathcal{P}$, let $\mathcal{A}_\Pi$ be its answer space, and let $\mathcal{A} = \bigsqcup_{\Pi \in \mathcal{P}} \mathcal{A}_\Pi$. An answer rule on presentations assigns $F(x) \in \mathcal{A}_{q(x)}$.

A structure-preserving redescription $u : \Pi \rightarrow \Pi'$ carries answers by a map $u_* : \mathcal{A}_\Pi \rightarrow \mathcal{A}_{\Pi'}$.

Theorem 3.1 (Dependence on grounds). *An answer rule $F : \mathcal{X} \rightarrow \mathcal{A}$ depends only on the problem precisely when it is constant on every class of presentations identified by q . Equivalently, there is a unique rule $\widehat{F}$ on problems such that*

$$F = \widehat{F} \circ q.$$

For equivalent descriptions $u : \Pi \rightarrow \Pi'$, the answers must also transport coherently:

$$\widehat{F}(\Pi') = u_* \widehat{F}(\Pi).$$

Proof. If $F = \widehat{F} \circ q$, presentations of the same problem receive the same answer. Conversely, if F is constant on each such class, define $\widehat{F}(\Pi) = F(x)$ for any presentation x with $q(x) = \Pi$. Constancy on these classes makes the definition independent of x , and surjectivity gives existence and uniqueness. The transport equation extends the same requirement from alternative presentations to equivalent descriptions of the problem itself. $\square$




Figure 1: Dependence on grounds. Left: an answer rule depends only on the problem exactly when it factors through the quotient q , so $F = \widehat{F} \circ q$. Right: across an equivalent redescription u , answers transport coherently, $\widehat{F}(\Pi') = u_* \widehat{F}(\Pi)$.

An answer follows from a problem exactly when every distinction that changes the answer survives in the problem.

Put two reasoners before the same problem. A difference in their answer spaces reveals a difference in grounds, representation, or procedure. Naming that difference places it in the problem. Until then, the full answer retains every possibility supported by the shared grounds. Coordinates, labels, file order, algorithms, gauges, and priors acquire inferential force through their explicit role in the problem. Figure 1 displays both conditions.

Definition 3.3 (Complete answer). For a problem Π , let $\text{Answer}(\Pi)$ denote the answers that satisfy it, together with the equivalences it declares. A report is **complete** when it preserves this entire answer space. Completeness is exhaustion relative to the question; computability, decidability, and single-valuedness are separate properties.

Remark 3.1. When equivalences among answers matter, the result retains them together with the answers. Category theorists call the resulting object a groupoid. The ordinary-language rule is simple: keep the answers together with the equivalences the problem gives.

Theorem 3.2 (Full-answer preservation). *Let R be presented as the complete answer to Π . Then R preserves every satisfying answer, every distinction left inequivalent by $\sim$, and every unresolved family. A question, output form, equivalence relation, or selector compresses the result precisely when it makes the corresponding distinction irrelevant. Those data define the richer problem whose complete answer is the compressed result.*

Proof. Definition 3.3 identifies completeness with preservation of the whole answer space. A revised equivalence, point-output protocol, or selector changes that space and thereby defines a richer problem. Theorem 3.1 adds presentation independence: every compression follows from structure retained by the problem itself. $\square$

Corollary 3.3 (Canonical point selection). *A point is canonical exactly when the problem contains structure that distinguishes and fixes it. When a symmetry exchanges satisfying answers while preserving the problem, the invariant result is their orbit.*

Proof. Let a preserve the problem and exchange x and y . Every point rule that follows from the problem is invariant under a . The invariant output therefore retains the exchanged answers together. An order, label, coordinate, or tie-breaker produces a point once it becomes part of the richer problem. $\square$

Remark 3.2 (Two symmetric candidates). Suppose the answers are a and b , and the swap $a \leftrightarrow b$ preserves every feature of the problem. The full answer is the symmetric pair. Enlarging the problem by an order $a \prec b$ supplies the distinction that selects a first.

Corollary 3.4 (Method and full answers). *Whenever an answer is presented as following solely from a problem:*

1. *every answer-relevant fact appears in that problem;*
2. *the answer depends on the problem and transports coherently across equivalent descriptions, which mathematicians call naturality;*
3. *working choices contribute to the result through an explicit role in the problem;*
4. *a complete report preserves every supported answer and every supported distinction.*

Proof. The MU Theorem establishes that consistent inference exists. Theorem 3.1 unfolds the dependence built into inference as defined in definitions 1.1 and 1.2. Definition 3.3 and Theorem 3.2 then give the two directions of completeness. Together they yield the method and the rule for complete answers. $\square$

Corollary 3.5 (Non-Smuggling as fidelity to grounds). *Non-Smuggling has three affirmative forms:*

1. **content fidelity:** *each proposition tracks the support carried by its premises;*

2. **selection fidelity:** each point choice tracks a selector carried by the problem;
3. **calculus fidelity:** each operation tracks the structure supplied by the problem and by the role the operation plays.

One rule for what follows. MU establishes the floor and governs the passage from grounds to answer. Logic, representations, reference measures, values, resources, and channels define the subject under study. Probability, relative entropy, starting beliefs, revision, choice, convergence, surviving families, and impossibility certificates are the corresponding full answers.

In symbols, the problem gives the subject matter and MU returns its complete answer:

$$\boxed{\text{MU} : \quad \Pi \longmapsto \text{Answer}(\Pi).}$$

The structure present fixes the resolution of the result: a point, an orbit, a wider family, or an empty requested class certified by proof.

Definition 3.4 (Four useful shapes of a complete result). After quotienting by the equivalence fixed by the problem, the answer space has one of four useful shapes:

1. one equivalence class with a representative fixed by every symmetry of the problem;
2. one orbit whose symmetry fixes the answer-class while leaving its representatives equivalent;
3. more than one inequivalent answer, retained as the complete family;
4. an empty admissible class or unattained optimum, accompanied by its certificate.

The first two distinguish uniqueness of the answer-class from whether the problem singles out a representative.

3.1 Epistemic Zero

The bilateral rule above is Epistemic Zero:

$$\boxed{\text{add only supported structure} \quad \text{and} \quad \text{preserve every supported distinction}.}$$

In plain words: assume nothing beyond what the constraints demand, and surrender nothing they deliver. MU leaves the subject matter exactly where the problem places it and governs how that subject matter bears on the answer. Its thinness is its reach: every substantive condition stays visible, while one rule governs what follows. Part III states its four arithmetic forms.

3.2 The axiomatic method as a special case

A mathematical problem specifies a language $\mathcal{L}$, an axiom theory T , a consequence relation $\vdash$, a question, a form of answer, and an equivalence. The general rule then becomes ordinary deductive closure.

Theorem 3.6 (Deductive closure). *For the theorem-generation problem determined by T , the complete theorem set is*

$$\boxed{\text{Cn}_{\vdash}(T) = \{\varphi \in \text{Sent}(\mathcal{L}) : T \vdash \varphi\}.}$$

Proof. The closure contains exactly the sentences supported by T and $\vdash$. It therefore preserves the whole theorem output and its exact boundary. $\square$

Remark 3.3 (Inconsistent theories). The displayed closure is syntactic. In an explosive logic, an inconsistent theory has the universal sentence set as its closure, which records trivialisation. The consistent-inference problem treats consistency as an admissibility condition and records its violation explicitly. Paraconsistent consequence relations define a different theorem-generation problem with their own closure.

Part II

The Ground and Its Limits

Give me a place to stand, and I will move the earth.

Attributed to Archimedes

4 The Anatomy of an Inference Episode

An abstract inferential relation needs rules and content. An actual episode also needs something that carries the relation out. Keeping these apart prevents us from confusing a proof with an event of proving.

4.1 L: Logic

Logic is the rule structure that distinguishes valid from invalid transitions and allows inferences to compose. The domain under study fixes the consequence relation.

4.2 C: Content

Content is the subject matter represented in the premises and conclusion, together with the relations that make the former bear on the latter. Symbol manipulation can instantiate a calculus, but interpreting it as inference about a domain requires a semantics.

4.3 A: Agent or Operator

The **agent or operator** is whatever carries out the inferential transition in an episode: a person, institution, proof assistant, algorithm, or machine. The theorem concerns this function and makes no assumption about consciousness.

Theorem 4.1 (The three aspects of an inference episode). *For an episode E presented as drawing conclusions from premises,*

$$E \text{ is an inference episode} \iff L(E) \wedge C(E) \wedge A(E).$$

Here L is the rule that makes the transition valid, C the premises and conclusion, and A the operator that carries the transition out. Within the class of inference episodes, none of the three can be removed while the episode itself is retained.

Proof. If E is an inference episode, “from” supplies L , “premises” and “conclusions” supply C , and “drawing” supplies A . Conversely, when an operator applies a rule to given premises and thereby produces a conclusion, an inference episode occurs. The three apparent separations fail in different ways: $L \wedge C \wedge \neg A$ leaves an abstract inferential relation with no episode; $L \wedge A \wedge \neg C$

leaves rule execution with no specified inferential content; and $C \wedge A \wedge \neg L$ leaves a change of state with no valid transition from premises to conclusion. $\square$

Remark 4.1 (Scope). The theorem is relative to the object being analysed. An abstract inferential relation contains rules and content; an inference episode adds the operator that carries the relation out. Within an episode, Logic, Content, and Agent are inseparable aspects of one event. Definition 1.1 retains abstract consequence relations at the relational level.

5 The Foundation Theorem

The order matters. MU is established first; its application to inference follows afterwards.

Theorem 5.1 (Foundation theorem). *Within the class just defined:*

- (i) **Proof order.** *MU is proved independently of the later method, which is then applied to every subsequent proof, question, and assessment.*
- (ii) **Minimality.** *Any proposition sufficient to ground the existence of consistent inference entails MU; MU is therefore the weakest claim sufficient for that task.*
- (iii) **Self-application.** *MU has a direct proof, and every valid assessment of it is itself an inference governed by the relation MU concerns. Reflexive application follows proof and preserves the proof order.*

Proof. Part (i) follows from the identity proof in Theorem 2.1 and the Dependence Theorem. For part (ii), any proposition sufficient to ground the existence of consistent inference entails that at least one such inference exists, which is exactly MU. For part (iii), MU has already been established; every later assessment is an inference episode and therefore falls under the method of Corollary 3.4. The proof order places every later consequence above the established floor. $\square$

5.1 Criterion and content

The problem of the criterion untangles once method and subject matter occupy their proper roles. The method follows from what it is for an answer to come *from* its inputs: relevant information is visible, equivalent descriptions agree, and a complete report preserves the full supported answer. Premises, models, and channels provide the material being assessed. Experience tests their fit to the world, while the inferential form governs the assessment.

6 Familiar Forms of Complete Answers

Gödel, Turing, Tarski, Fitch, Arrow, Goodman, and No Free Lunch are complete achievements of reason (Gödel, 1931; Turing, 1936; Tarski, 1936; Fitch, 1963; Arrow, 1951; Goodman, 1955; Wolpert and Macready, 1997). Each fixes a domain and classifies exactly what that domain permits. Set beside Shannon's positive classification (Shannon, 1948), they display the principal shapes of a full result: a theorem, a surviving family, a reference-relative answer, or a certified empty class.

Proposition 6.1 (Impossibility completes a classification). *Suppose a problem asks for an object satisfying conditions $C_1, \dots, C_n$. A proof that the admissible class is empty gives the complete classification of that problem. A constructive answer becomes available through a revised condition, a restricted domain, an enlarged output type, or added structure.*

Proof. The impossibility proof classifies every candidate in the defined class as inadmissible. The empty answer space therefore exhausts the requested domain and completes the result. $\square$

Result	What the problem gives	Form of the complete answer
Gödel	A sufficiently expressive, effectively axiomatised formal theory, arithmetisation, derivability conditions, and the relevant consistency hypothesis	An internal proof ceiling: under the usual hypotheses, the system's own consistency lies beyond its theorem closure
Turing	A model of effective computation and the demand for a total procedure deciding whether arbitrary programs halt	The total halting-decider class is empty
Tarski	A sufficiently expressive object language and the demand for an internally definable truth predicate satisfying the full truth scheme	Full truth moves to a metalanguage or a suitably restricted object language
Fitch	Factive knowledge, ordinary modal closure, and the claim that every truth is knowable	Universal knowability collapses to universal knowledge; unknown truths require a weaker knowability claim
Arrow	Individual preference orderings, at least three alternatives, a collective ordering, and the fairness and independence conditions	Every aggregation rule satisfying the full package is dictatorial
Goodman	Evidence, a projective question, and an unresolved predicate or representation language	A family of projective rules remains until representation and invariance structure are supplied
No Free Lunch	An unrestricted problem class with performance averaged uniformly over all objective functions	Uniformly averaged performance is equal across learners; problem structure creates advantage
Algorithmic induction	A coding language or universal reference machine	Universality and simplicity are fixed only relative to that reference, with the given invariance bounds
Shannon	A source, channel, output alphabet, and the continuity and composition conditions	Entropy and channel quantities take their classified information-theoretic form

6.1 Formal ceilings: Gödel, Turing, and Tarski

Gödel's incompleteness theorems are often heard as a warning about the reach of formal reason. They establish exact limits after a rich formal problem has been specified. Their statement and proof already require a consequence relation, consistency conditions, and conclusions that follow from premises. MU is prior in logical role because the incompleteness theorem is itself a valid result from its grounds.

Applied to a Gödelian problem Π_G , MU returns the applicable incompleteness classification. The domain-specific hypotheses belong to Π_G . The complete result is the theory's internal proof closure, while the expressible truths may extend beyond it. Gödel marks a ceiling inside formal reasoning while presupposing the floor on which that ceiling is proved.

Turing and Tarski deepen the same point. Turing identifies the empty class of universal halting deciders within the model of computation. Tarski returns a level distinction: full truth

for a sufficiently expressive object language lives in a metalanguage or a restricted hierarchy. These are complete results about what the given structures permit.

6.2 Knowability and Fitch’s boundary

Fitch’s paradox isolates a different limit. Let K be factive knowledge and suppose every truth is knowable:

$$\varphi \longrightarrow \Diamond K\varphi.$$

If some truth p is unknown, then $p \wedge \neg Kp$ is true. Universal knowability would make it possible to know that conjunction. But knowing it yields both Kp and $K\neg Kp$; factivity of the second conjunct gives $\neg Kp$, contradicting Kp . Hence, under these modal rules,

every truth is knowable $\implies$ every truth is known.
--

Inquiry proceeds under a knowability principle whose strength preserves room for unknown truths (Fitch, 1963). Fitch’s result classifies the unrestricted schema and identifies the exact revision points: the schema, the modal logic, or the conception of knowledge.

6.3 Jointly imposed conditions: Arrow

Arrow’s theorem asks for a social aggregation rule satisfying a package of conditions. The theorem returns an empty non-dictatorial class for the full package. Practical design then proceeds by relaxing a condition, restricting the preference domain, introducing interpersonal structure, or changing the requested output. MU makes that price visible.

6.4 Generalisation needs structure: Goodman and No Free Lunch

Goodman’s riddle (Goodman, 1955) and No Free Lunch meet at one lesson: *successful generalisation requires structure supplied by the problem*. Inductive bias is the structure through which a learner generalises. MU makes that structure explicit, testable, or inherited from the problem. When representation, problem distribution, or invariance remains unresolved, the full answer retains the surviving family.

6.5 Reference-relative universality: Solomonoff and Kolmogorov

Algorithmic induction illustrates the same requirement (Solomonoff, 1964; Kolmogorov, 1965). A universal machine supplies the coding against which simplicity is measured. The invariance theorem bounds variation across suitable machines while preserving finite-scale reference dependence. Thus “universal” means universal relative to a chosen coding scheme. More generally, MU places every limit theorem inside the result supported by its problem: a certified empty class can itself be the complete answer. *The limitation of reasoning is still something known by reasoning.*

Part III

The Mathematics of Rational Belief

A wise man proportions his belief to the evidence.

David Hume, *An Enquiry Concerning Human Understanding*

The method now turns to rational belief: how degrees of support combine, how informational change is measured, how starting beliefs are formed, how beliefs are revised, and how values enter choice. Boolean structure, scalar representation, reference measures, and channel models name the objects under study. MU supplies every law that carries those objects from grounds to answer.

A fully stated domain can admit one mathematical form. Logicians call this *relative categoricity*; the plainer phrase *uniqueness within the domain* captures the result. The domain names the object. MU fixes the form that follows from it.

7 Three Distinctions

Three distinctions organise the argument and its philosophical applications.

7.1 Constraints and conclusions

Definition 7.1 (Constraint and conclusion). A **constraint** is part of what the problem gives. A **conclusion** is what those constraints support under the chosen consequence relation. Non-Smuggling governs the passage from the first to the second.

Constraints include more than observations. They can include logical relations, symmetries, apparatus, form of answer, and equivalence. Treating one of these as hidden background makes the problem appear more determinate than it is; treating a conclusion as an input makes a derivation circular.

7.2 Constraint language

Definition 7.2 (Constraint language $\mathcal{L}$). The constraint language of a quantitative Boolean case may contain:

1. logical and partition relations;
2. expectation or moment constraints for the observables in the problem;
3. interval and inequality constraints;
4. symmetries and invariances;
5. likelihood or channel conditions supplied by the apparatus.

These are the forms used here. Any further constraint must be made explicit if it is to affect the answer.

Remark 7.1 (Constraint and assumptive structure). Every output-affecting assumption belongs among the grounds or follows from them. Coordinates, priors, and orderings acquire inferential force through that explicit role.

7.3 Internal support and connection to truth

Definition 7.3 (Internal and external). The **internal** question is whether an answer follows from the agent's grounds. The **external** question is whether those grounds represent the world and arrive through channels that reliably track the relevant truth.

Internal support and connection to truth are separate coordinates. A false model can support impeccable updating; a lucky guess can be true without justification. The channel and convergence sections classify the conditions under which internal updating tracks the world.

7.4 Constitutive and hypothetical

Definition 7.4 (Constitutive and hypothetical). A condition is **constitutive** when it belongs to what the practice is, for example that an inferential conclusion depend on its grounds. A claim is **hypothetical** when it is one proposition assessed within that practice.

Keeping these levels distinct resolves familiar regress demands. A model's connection between past and future is a substantive question. An audit of all inference is itself an inference and therefore remains inside the same constitutive form.

7.5 A simple diagnostic

The distinctions cut the classical problems at their joints. Induction separates constitutive support from projective hypotheses: the demand for justification outside all inference dissolves, while convergence remains conditional on model and channel assumptions. Goodman separates constraints from representation: the question moves to the chosen predicate language, invariances, and apparatus, and unresolved structure returns a family. Gettier separates internal support from the external route: empirical knowledge joins internal support to a route that reliably tracks truth. Scepticism separates global denial from local channel doubt: global denial of all inference is defeated by the MU Theorem, while local doubt remains calibrated uncertainty about channels and models.

8 The Quantitative Laws of MU

Each quantitative domain names an object: plausibility, informational change, starting belief, revision, choice, or learning from a channel. MU generates the laws of each object by preserving every dependence carried by its grounds, every distinction carried by its state, and every equivalence carried by its descriptions. Four recurring movements do the work: dependence, retention, refinement, and completion. Their arithmetic forms are probability, relative entropy, maximum entropy, and KL projection.

Theorem 8.1 (Four quantitative laws generated by MU). *The finite quantitative domains developed in this Part have four MU laws:*

- (i) *determinate Boolean scalar plausibility has probability as its unique form, up to order-preserving regraduation;*
- (ii) *total informational change has exact marginal–conditional decomposition and relative entropy as its unique form, up to positive scale;*
- (iii) *least-assuming completion from each permitted reference is relative-entropy projection, with Shannon maximum entropy as the finite counting-symmetry form;*
- (iv) *revision preserves every feature of the current state compatible with new constraints and therefore takes the form of KL projection.*

Full-Answer Preservation carries every remaining plurality as its complete family and every empty or unattained case as its certificate.

Proof. The next four sections prove items (i)–(iv). Full-Answer Preservation supplies the concluding sentence. In every case the domain names the represented object and MU supplies the law that fixes its full answer. $\square$

Given a truth-representing channel, prior support, stable sampling, and distinguishability, repeated KL revision concentrates on the represented truth.

$$\boxed{\mu \xrightarrow[c_0]{D_{KL}\text{-projection}} P_0 \quad P_0 \xrightarrow[c_1]{D_{KL}\text{-projection}} P_1}$$

Probability supplies the states; relative entropy supplies the geometry; the two projections generate MaxEnt completion and KL revision.

Definition 8.1 (Point and set-valued answers). A **point answer** occurs when one answer remains after the relevant equivalences and explicit selectors are taken into account. A **set-valued answer** occurs when several inequivalent possibilities remain; the family itself is then the full answer, and a later selector defines a downstream problem.

One principle, four laws, every answer shape. The four laws are one rule seen four times. Probability preserves the structure of determinate Boolean support. Relative entropy accounts for every informational change exactly once. Maximum entropy carries only the concentration purchased by the constraints. KL projection carries every compatible feature of the current state through revision. Each problem resolves into the exact shape supported by its grounds: point, orbit, family, attained optimum, unattained optimum, or certified empty class.

8.1 Representation and hypothesis space

A problem states the representations through which evidence and candidate states are described.

Definition 8.2 (Representation language). The language $\mathcal{L}$ contains the predicates, observables, transformations, and relations admitted by the problem. In empirical work it is constrained by the apparatus and modelling purpose; in mathematics it is supplied by the signature of the theory. Apparatus restricts what can be observed and may leave several representations available.

Definition 8.3 (Hypothesis space). The hypothesis space $\mathcal{H}$ is the candidate class over which the question is asked. Its members are expressible in $\mathcal{L}$, while the problem may select a proper

subset of the formulas expressible there. Representation choice and model-class choice are visible inputs whose consequences remain traceable through the update.

Reasoners work through available representations and instruments. The method applies across substrates, but every application remains situated: changing the language, apparatus, or model class changes the question. Those changes must remain visible so that their consequences can be audited.

Remark 8.1 (Completion over finite and infinite spaces). On a finite hypothesis space, permutation or combinatorial symmetry makes counting the natural reference. On a continuous space, entropy and relative entropy are defined relative to a reference measure, and uniformity follows from invariant length or counting structure. On countably infinite or noncompact spaces, a normalised completion follows from suitable reference, constraint, compactness, and coercivity conditions. An unattained optimum is recorded as the full result for that case.

The probability theorem gives MU's form for determinate scalar Boolean plausibility. Partial comparisons produce the full family of compatible representations.

9 Probability

Probability is MU's unique form for determinate scalar plausibility on a Boolean event structure (Cox, 1946, 1961; Aczél, 1966; Halpern, 1999; Kraft et al., 1959). Here *Boolean* and *scalar* name the object under study, as *group* names an algebraic object. MU supplies every inferential law of its calculus. Probability is Epistemic Zero written as arithmetic: every supported distinction is preserved and every undetermined distinction remains open.

Definition 9.1 (Determinate Boolean plausibility domain). A determinate Boolean plausibility domain consists of:

- (D1) a Boolean algebra of propositions and a total plausibility order with a separable, operationally rich scalar range;
- (D2) a scalar state complete for conjunction and disjoint union: the scalar inputs and Boolean relations exhaust the operation-relevant information;
- (D3) symmetric refinement: rational subdivisions exist, common refinements preserve the represented event, and the realised partial ranges are dense in their target intervals.

A partial comparison produces its full family of compatible orderings or probability representations.

Proposition 9.1 (MU generates the probability calculus). *In every determinate Boolean plausibility domain, MU generates:*

- (C1) **scalar substitution**: equal scalar states are interchangeable, yielding universal scalar operations for conjunction and disjoint union;
- (C2) **retention of support**: each operation is strictly increasing on every nondegenerate argument whose order distinction the scalar state preserves;
- (C3) **refinement coherence**: equivalent refinements and common refinements agree;
- (C4) **Boolean coherence**: associativity, identity, zero, complementation, and distributivity follow from the corresponding logical identities;
- (C5) **continuity**: dense refinement and strict monotonicity make the operations continuous.

Proof. D2 and the Dependence Theorem give C1: every operation represented on scalar states factors through those states. Full-Answer Preservation gives C2 by carrying each represented order distinction into the output. D3 and presentation invariance give C3. Logically identical Boolean expressions represent one proposition and therefore have one plausibility value, giving C4. Finally, D3 gives each monotone one-variable section a dense realised range; density fills every interval and yields continuity. The rectangle squeeze used in Theorem 9.3 then gives joint continuity. $\square$

Theorem 9.2 (MU forces probability). *Every determinate Boolean plausibility domain, together with the calculus generated in Proposition 9.1, admits an order-preserving regraduation p for which*

$$p(A \wedge B \mid C) = p(A \mid C)p(B \mid A \wedge C),$$

$$p(A \vee B \mid C) = p(A \mid C) + p(B \mid C) - p(A \wedge B \mid C),$$

and

$$p(\neg A \mid C) = 1 - p(A \mid C).$$

Probability is therefore MU's unique scalar Boolean calculus, up to order-preserving regraduation. Partial comparison yields the full family of compatible representations; an empty representation class yields its certificate.

9.1 Derivation

The answer depends on the constraint set as a whole. Every finite reordering presents the same problem and therefore gives the same full result. Sequential algorithms may pass through different intermediate states; the stated derivation identifies their relation to the simultaneous problem.

The given structure determines each retained comparison. Where it leaves propositions incomparable, the full answer preserves that incomparability together with every compatible ordering or probability representation. Every operation also transports coherently across equivalent descriptions. Labels, coordinates, and proposition identities acquire relevance through the equivalence structure of the problem.

Theorem 9.3 (Continuity from operational richness). *A strictly increasing scalar composition whose realised partial ranges are dense in their target intervals is continuous.*

Proof. For a monotone one-variable section, every discontinuity is a jump and therefore omits a nonempty open interval from the realised range. Density excludes such a gap, so every section is continuous. To obtain joint continuity at (x, y) , choose $x_- < x < x_+$ and $y_- < y < y_+$ so close that the four one-variable boundary changes are arbitrarily small. Monotonicity gives

$$F(x_-, y_-) \leq F(x', y') \leq F(x_+, y_+)$$

whenever (x', y') lies in the intervening rectangle. Separate continuity makes both bounds converge to $F(x, y)$, and the squeeze proves joint continuity. $\square$

Remark 9.1 (Neighbouring regularity domains). A globally fixed discontinuity or flat region defines a neighbouring domain. Its realised order contains a gap or compresses a distinction represented as relevant. Reproducibility assigns the rule to that domain and preserves the distinction from the probability case.

9.2 Associative composition

Lemma 9.4 (Monotone Cauchy). *Let $h : (0, \infty) \rightarrow \mathbb{R}$ be monotone and satisfy $h(x + y) = h(x) + h(y)$. Then $h(x) = h(1)x$.*

Proof. Additivity gives the result on positive rationals. Rational sequences increasing and decreasing to a real x then squeeze $h(x)$ to $h(1)x$ by monotonicity. $\square$

Theorem 9.5 (Multiplicative representation of conjunction). *Let $I = [b, t]$ and let $F : I^2 \rightarrow I$ be continuous and associative, strictly increasing in each argument whenever the other lies in $(b, t]$, with*

$$F(x, t) = F(t, x) = x, \quad F(x, b) = b.$$

Then there is a continuous strictly increasing $g : I \rightarrow [0, 1]$ with $g(b) = 0$, $g(t) = 1$, and

$$F(x, y) = g^{-1}(g(x)g(y)).$$

Proof. This is the standard strict associative-operation representation; the details below make the order-density step explicit (Aczél, 1966).

Rational powers. Strict monotonicity gives cancellation. Fix $a \in (b, t)$ and define $a^{n+1} = F(a^n, a)$. Since $F(x, a) < F(x, t) = x$ for $x > b$, the sequence (a^n) decreases. Its limit must be b : if it were $\ell > b$, continuity would give $F(\ell, a) = \ell$, contradicting $F(\ell, a) < F(\ell, t) = \ell$.

For each n , the continuous strictly increasing map $y \mapsto y^n$ sends I onto I , so a has a unique n th root $a^{1/n}$. Associativity and cancellation define $a^{m/n} = (a^{1/n})^m$ independently of the chosen representation of m/n and give the usual power laws. If $n > m$, then $a^{1/n} > a^{1/m}$; otherwise raising both sides to the power mn would imply $a^m \leq a^n$, contrary to the strict decrease of integer powers. Thus $a^{1/n}$ increases to a limit $\ell \leq t$. If $\ell < t$, then $\ell^k \downarrow b$, so for some k one has $\ell^k < a$; but $a^{k/n} \rightarrow \ell^k$ and $a^{k/n} > a$ for $n > k$, a contradiction. Hence $a^{1/n} \uparrow t$.

Order density and an additive coordinate. Uniform continuity gives

$$\sup_{z \in I} |F(z, a^{1/n}) - z| \rightarrow 0$$

after any fixed order-preserving coordinate is placed on I . Given $b < x < y < t$, choose n so large that consecutive points in the chain $t, a^{1/n}, a^{2/n}, \dots$ are separated by less than $y - x$. Since the chain descends from t to b , one of its terms lies in (x, y) . The rational powers are therefore order-dense.

Define

$$\lambda(x) = \sup\{r \in \mathbb{Q}_{>0} : a^r \geq x\}.$$

Order density makes $\lambda : (b, t) \rightarrow (0, \infty)$ a continuous strictly decreasing bijection. The power law first gives $\lambda(F(a^r, a^s)) = r + s$ on rational powers. Approximation from above and below, together with continuity of F and λ , extends this identity to all $x, y \in (b, t)$:

$$\lambda(F(x, y)) = \lambda(x) + \lambda(y).$$

Multiplicative change of scale. Putting $g(x) = e^{-\lambda(x)}$ and extending by $g(b) = 0$, $g(t) = 1$ gives a continuous strictly increasing bijection with $g(F(x, y)) = g(x)g(y)$, which is the claimed representation. $\square$

After regraduating by g , conjunction is multiplication. Write the resulting scale as p .

9.3 Disjoint union and complementation

For disjoint propositions, D2 and Proposition 9.1 supply a universal partial operation $x \oplus y = p(A \vee B \mid C)$. Boolean union makes it commutative and associative, with identity 0; C2 and D3 make it continuous and strictly increasing on its domain triangle.

Lemma 9.6 (Additive representation of disjoint union). *There is a continuous strictly increasing $h : [0, 1] \rightarrow [0, 1]$, $h(0) = 0$, $h(1) = 1$, such that*

$$h(x \oplus y) = h(x) + h(y)$$

whenever $x \oplus y$ is defined. Compatibility with Boolean distributivity forces $h(x) = x^m$ for some $m > 0$.

Proof. Rational points. Write $n \odot x$ for the n -fold disjoint sum when it is defined. The symmetric-refinement condition D3 supplies, for every $n \geq 1$, a unique equal part e_n with $n \odot e_n = 1$, and Proposition 9.1 identifies the same rational part across common refinements. For $0 \leq m \leq n$ define

$$u_{m/n} = m \odot e_n.$$

Associativity, commutativity, and common-refinement coherence give

$$u_r \oplus u_s = u_{r+s} \quad (r, s, r+s \in \mathbb{Q} \cap [0, 1]),$$

and strict monotonicity orders the u_r exactly as their rational labels. D3 makes these rational points order-dense in the realised interval.

Additive change of scale. Define

$$h(x) = \sup\{r \in \mathbb{Q} \cap [0, 1] : u_r \leq x\}.$$

Density and strict monotonicity give $h(u_r) = r$ and make h a continuous strictly increasing bijection of $[0, 1]$. If $x \oplus y$ is defined, rational nets approaching x and y from below and above, together with continuity, give

$$h(x \oplus y) = h(x) + h(y).$$

The normalisation $h(0) = 0$, $h(1) = 1$ is built into the construction.

Compatibility with conjunction. Boolean distributivity gives

$$z(x \oplus y) = (zx) \oplus (zy).$$

For fixed z , define

$$T_z(r) = h(z h^{-1}(r)).$$

The displayed identity implies $T_z(r+s) = T_z(r) + T_z(s)$ whenever $r+s \leq 1$. Continuity and Lemma 9.4 give $T_z(r) = c(z)r$. Taking $r = 1$ gives $c(z) = h(z)$, hence

$$h(zx) = h(z)h(x).$$

Every continuous increasing multiplicative map on $[0, 1]$ with $h(1) = 1$ has the form $h(x) = x^m$ for some $m > 0$. $\square$

Set $q = h \circ p = p^m$. Because h is multiplicative, conjunction remains multiplication; because h additivises $\oplus$, disjoint union becomes addition. Since A and $\neg A$ are disjoint and exhaust the Boolean universe,

$$q(A \mid C) + q(\neg A \mid C) = 1.$$

Inclusion–exclusion gives the general sum rule. Renaming q as p completes the proof of Theorem 9.2.

Remark 9.2 (Neighbouring qualitative domains). Finite qualitative systems display the neighbouring forms clearly. Some total qualitative orders lie outside additive probability representation and therefore return an empty representation class: Kraft, Pratt, and Seidenberg exhibit a five-atom order with no agreeing measure, and Scott supplies the necessary and sufficient cancellation conditions (Kraft et al., 1959; Scott, 1964). Scalar completeness, symmetric refinement, substitution, and common-refinement coherence select the probability domain. Partial orders, vector-valued states, and context-sensitive operations receive their own full answers under their own specifications.

Remark 9.3 (Dutch books, frequencies, and credence). When credences are offered as fair betting prices and an opponent may combine bets, violations of the probability laws can produce a sure loss. Dutch-book arguments therefore supply an operational consequence of incoherence under a betting interpretation (Ramsey, 1926). Reference measure, prior, and world model enter through their own parts of the problem. Observed frequencies are outputs of a sampling channel, while credences are support states conditional on a model. Long-run frequencies calibrate beliefs while preserving that distinction.

Remark 9.4 (Accuracy and rival routes). Accuracy-first epistemology derives probabilism from dominance with respect to strictly proper scoring rules, and its programme extends the derivation to updating and chance deference (Joyce, 1998; Pettigrew, 2016). Objective Bayesianism selects the equivocal prior through a maximum-entropy norm (J. Williamson, 2010), and preference-first routes recover probability and utility together from axioms on choice (Savage, 1954). Each is a neighbouring route to overlapping conclusions: the accuracy route begins from a scoring rule and a dominance principle, the preference route from choice axioms over acts, and the present route from the dependence analysis and the domain’s richness conditions. These routes begin from different grounds and converge on overlapping quantitative forms; the convergence strengthens the classification, while whether accuracy dominance and dependence on grounds force one another remains open.

10 Relative Entropy

Relative entropy is MU’s unique form for total informational change between finite probability states (Kullback and Leibler, 1951; Shore and Johnson, 1980; Csiszár, 1975; Uffink, 1995). The same transition admits a flat joint description and a staged marginal–conditional description. MU requires both descriptions to carry one total amount of change. Relative entropy is Epistemic Zero’s ledger: every change is counted once, and every equivalent description balances to the same total.

Definition 10.1 (Informational-change domain). An informational-change domain represents a transition $P \rightarrow Q$ by a nonnegative scalar magnitude with the unchanged transition as its identity.

It admits independent product composition, disjoint branch composition, symmetric refinement, and equivalent flat and hierarchical descriptions. Its realised magnitudes are operationally rich enough for the continuous representation of those compositions.

Lemma 10.1 (Final-state branch weighting). *Let a conditional informational change of magnitude d occur on a final-state branch of probability q . Symmetric refinement, branch additivity, and continuity force that branch to contribute qd to the total change.*

Proof. For rational $q = m/n$, refine the final-state partition into n symmetric copies. Refinement invariance preserves the branch's total change, symmetry gives contribution d/n to each copy, and branch additivity gives md/n to the m copies occupied by the branch. Operational richness makes rational branch weights dense in $[0, 1]$, and continuity extends the formula to every q . $\square$

Final-state weighting fixes the informational direction: $D(Q\|P)$ measures the information carried by the revised state Q relative to P , so conditional changes are averaged with the final marginal Q_X . Reversing the transition exchanges the arguments and the branch weights.

Theorem 10.2 (MU generates exact information accounting). *On every finite informational-change domain, MU generates:*

- (I1) *identity and strict retention: $D(Q\|P) \geq 0$, with zero exactly at $Q = P$;*
- (I2) *invariance under common relabelling and equivalent presentation;*
- (I3) *additive composition of independent components after a positive choice of scale;*
- (I4) *exact marginal–conditional decomposition,*

$$D(Q_{XY}\|P_{XY}) = D(Q_X\|P_X) + \sum_x Q_X(x)D(Q_{Y|x}\|P_{Y|x});$$

- (I5) *contraction under every common stochastic channel;*
- (I6) *continuity on positive simplex interiors.*

Proof. I1 is the scalar identity law for change together with Full-Answer Preservation: every represented change receives positive magnitude. I2 follows from the Dependence Theorem because a common relabelling preserves the represented transition.

For I3, independent component changes combine by a universal scalar operation. Regrouping three components gives associativity; the unchanged component is the identity; Full-Answer Preservation gives strict increase; operational richness gives continuity. Aczél's continuous associative-operation representation, proved by the rational-power construction underlying Theorem 9.5, therefore supplies an order-preserving scale on which independent composition is addition (Aczél, 1966).

For I4, Lemma 10.1 gives the contribution of branch x as $Q_X(x)D(Q_{Y|x}\|P_{Y|x})$. The hierarchical description therefore consists of the marginal component together with the final-state-weighted conditional components. Its flat joint description represents the same transition, so the Dependence Theorem gives the displayed equality.

I5 is Lemma 10.4, derived from I1 and I4. I6 follows on positive interiors from the same dense-refinement argument used for probability. The KL characterisation below fixes the interior form, and Corollary 10.6 supplies its canonical support-boundary extension. $\square$

Proposition 10.3 (Relabeling invariance). *A common relabelling preserves the value of informational change.*

Proof. A common relabelling changes the presentation of the pair while preserving the transition. Theorem 3.1 therefore carries the value unchanged across that redescription. $\square$

Lemma 10.4 (Information monotonicity from exact accounting). *For any finite stochastic kernel K and any pair $Q \ll P$,*

$$\boxed{D(QK \| PK) \leq D(Q \| P)}.$$

Thus coarse-graining preserves or reduces measured change.

Proof. Let

$$Q_{XY}(x, y) = Q_X(x)K(y | x), \quad P_{XY}(x, y) = P_X(x)K(y | x).$$

Decomposition through X gives

$$D(Q_{XY} \| P_{XY}) = D(Q_X \| P_X) + \sum_x Q_X(x) D(K_x \| K_x) = D(Q_X \| P_X).$$

Decomposition through Y gives

$$D(Q_{XY} \| P_{XY}) = D(Q_Y \| P_Y) + \sum_y Q_Y(y) D(Q_{X|y} \| P_{X|y}) \geq D(Q_Y \| P_Y),$$

where zero-weight conditional terms vanish by convention. Since $Q_Y = Q_X K$ and $P_Y = P_X K$, the two evaluations give the result. $\square$

Theorem 10.5 (MU forces relative entropy). *Let $D(Q \| P)$ be the informational-change measure generated by Theorem 10.2. For strictly positive finite probability states,*

$$D(Q \| P) = k D_{KL}(Q \| P) \quad (k > 0).$$

MU applies the resulting divergence to starting belief through least-assuming completion and to revision through retention.

Proof. Uniform divergences. Define $c(n) = D((1, 0, \dots, 0) \| U_n)$. Split $n = mk$ outcomes into m blocks of k : the chain rule gives

$$c(mk) = D((1, 0, \dots, 0)_m \| U_m) + 1 \cdot D((1, 0, \dots, 0)_k \| U_k) = c(m) + c(k),$$

where conditional terms carrying zero weight vanish by the standard convention $0 \cdot D(\cdot) = 0$; the only blocks receiving positive weight are well-defined. Lemma 10.4 also gives monotonicity. Let K be the row-stochastic kernel from $n+1$ outcomes to n that fixes cell 1 and, for each remaining row, has

$$K_{i1} = \frac{1}{n^2}, \quad K_{ij} = \frac{n+1}{n^2} \quad (j \geq 2).$$

A direct calculation gives $U_{n+1}K = U_n$ and $\delta_1 K = \delta_1$. The derived data-processing inequality therefore gives $c(n+1) \geq c(n)$. Moreover $\kappa := c(2) > 0$: otherwise D would vanish at the unequal pair (δ_1, U_2) .

Logarithmic growth. For any n, m , choose k with $2^k \leq n^m < 2^{k+1}$. Additivity and monotonicity give

$$k\kappa \leq m c(n) = c(n^m) < (k+1)\kappa.$$

Hence $c(n)/\log n = \kappa/\log 2$; after absorbing the logarithm base into the constant, write $c(n) = \kappa \log n$.

Refinement and nesting. For rational pairs with a common denominator, uniform refinement leaves D unchanged. If $Q = (a_i/N)_i$ and $\hat{Q}$ spreads the i th final-state block uniformly over a_i cells, with $\hat{P}$ defined analogously, the chain rule gives

$$D(\hat{Q}\|\hat{P}) = D(Q\|P) + \sum_i Q_i D(U_{a_i}\|U_{a_i}) = D(Q\|P).$$

Two further consequences of the chain rule will be used: *padding invariance*, obtained by adjoining cells on which both distributions vanish; and the *nesting formula*

$$D(\bar{U}_s\|\bar{U}_t) = \kappa \log(t/s) \quad (s \mid t),$$

where $\bar{U}_s$ is uniform on s of t cells.

Rational pairs and extension. Now let $Q = (a_i/N)$ and $P = (b_i/N)$ be strictly positive rational distributions. Fix $M \geq 1$ and choose an integer K that is a common multiple of the a_i and satisfies $K \geq \max_i a_i/b_i$. In row i , put $s_i = a_i M$ supported cells under the first argument and $t_i = b_i M K \geq s_i$ under the second, each uniform. Evaluate the divergence of the resulting joint pair in two ways. The chain rule through the row variable gives

$$D(Q\|P) + \sum_i Q_i \kappa \log(t_i/s_i).$$

After a permutation, the first joint is uniform on NM cells nested inside the second, which is uniform on NMK cells; the nesting formula therefore gives $\kappa \log K$. Equating the two evaluations yields

$$D(Q\|P) = \kappa \sum_i Q_i \log(a_i/b_i) = \kappa D_{KL}(Q\|P).$$

This holds for every strictly positive rational pair. Continuity extends the identity to every strictly positive pair. $\square$

Corollary 10.6 (Closed-simplex extension). *The lower-semicontinuous envelope of the positive-simplex formula on a finite closed simplex is the usual extended relative entropy:*

$$D(Q\|P) = \begin{cases} \sum_{i:Q_i>0}^k Q_i \log \frac{Q_i}{P_i}, & Q \ll P, \\ +\infty, & Q \not\ll P. \end{cases}$$

Proof. When $Q \ll P$, restrict to the common support and approach every zero coordinate through strictly positive pairs; continuity gives the displayed finite sum. If $Q_i > 0 = P_i$ for some i , then every strictly positive approximation has a term $Q_i^{(n)} \log(Q_i^{(n)}/P_i^{(n)})$ tending to $+\infty$. The liminf envelope therefore gives the stated extended value. $\square$

Remark 10.1 (Exact accounting selects KL). Exact final-state-weighted decomposition selects KL from the wider family of monotone divergences. A joint transition carries its marginal change plus its conditional changes, each weighted by the revised state. This accounting leaves one logarithmic form; Corollary 10.6 then extends it across the support boundary.

Remark 10.2 (One divergence, two uses). Prior completion and updating use one classified divergence with different reference objects:

$$\mu \xrightarrow{\text{I-projection onto prior constraints}} P_0, \quad P_0 \xrightarrow{\text{I-projection onto new constraints}} P_1.$$

The first reference is a measure permitted by the representation and apparatus; the second is the current state of belief. The same divergence governs both maps. Least-assuming completion and retention carry forward every feature compatible with the relevant constraints in their respective settings.

11 Starting Beliefs: Reference Measures and Maximum Entropy

Starting belief is MU applied before new evidence arrives. The problem identifies every permitted reference structure; MU then selects the feasible states carrying the least informational departure from each reference. Maximum entropy is Epistemic Zero before evidence: the selected state carries the least concentration compatible with the constraints and the reference. Coordinates, symmetries, and apparatus remain visible inside the full answer.

Remark 11.1 (Bayes begins from a starting state). Bayes' rule maps a prior and evidence to a posterior (Bayes, 1763). MU brings the starting measure onto the page before revision begins. Maximum entropy is the least-concentrated completion relative to the reference structure carried by the case (Jaynes, 1957, 2003; Uffink, 1995). A unique reference yields a point completion; several permitted references yield their full family.

Definition 11.1 (Permitted reference measures). For a prior problem Π specified in this way, let $\mathcal{R}_\Pi$ denote the representations permitted by the problem and apparatus descriptions together with the transformations that preserve the statistical question. A positive measure class $[\mu]$ on one presentation is **permitted** when it belongs to an invariant assignment

$$r \longmapsto [\mu_r], \quad f_*[\mu_r] = [\mu_{r'}] \quad \text{for every } f : r \rightarrow r' \text{ in } \mathcal{R}_\Pi.$$

Let $\mathfrak{M}_\Pi$ be the family of such permitted values, taken up to positive scale. Normalised measures give ordinary prior cases. A sigma-finite but improper reference may be retained only when the projection problem and posterior propriety are separately well-defined.

For a probability measure $P \ll \mu$, write

$$D(P||\mu) = \int \log\left(\frac{dP}{d\mu}\right) dP.$$

Scaling μ adds a constant independent of P , so the set of minimisers depends only on the reference class.

Definition 11.2 (Projection relative to a reference). For $\mu \in \mathfrak{M}_\Pi$ and feasible probability set $\mathcal{P}_C$, define

$$\text{Comp}_\Pi(\mu) = \arg \min_{P \in \mathcal{P}_C} D(P||\mu).$$

Let

$$\mathfrak{U}_\Pi = \{\mu \in \mathfrak{M}_\Pi : \text{Comp}_\Pi(\mu) = \emptyset\}$$

record nonattained cases.

Definition 11.3 (Prior answer). The complete prior output is

$$\text{Prior}(\Pi) = (\{(\mu, P) : \mu \in \mathfrak{M}_\Pi, P \in \text{Comp}_\Pi(\mu)\}, \mathfrak{U}_\Pi).$$

The first component records every attained case; the second records every permitted case with an unattained point optimum.

Theorem 11.1 (MU forces least-assuming completion). *For every permitted reference μ , MU selects exactly the feasible state or states carrying the least informational departure from that reference. They are the minimisers of the classified divergence, with the full minimiser set and every attainment boundary preserved.*

Proof. The reference records the comparison structure before the constraints are imposed. Relative entropy orders feasible states by the information they add to that reference. Whenever $D(P\|\mu) > D(Q\|\mu)$, selecting P carries a strictly larger informational departure than selecting Q . MU therefore selects the minimum departure, and Full-Answer Preservation carries every tied or unattained case into the result. $\square$

Theorem 11.2 (Prior completion). *The prior problem determines one complete prior answer.*

- (i) *One permitted reference and one attained minimiser give a single point prior.*
- (ii) *Several permitted references or several minimisers give the complete attained family.*
- (iii) *Every unattained permitted case appears in $\mathfrak{A}_\Pi$ with its attainment boundary.*

A chosen reference or minimiser becomes part of a richer problem and yields that richer problem's point answer.

Proof. Apply Least-Assuming Completion to every member of $\mathfrak{M}_\Pi$ and retain the resulting answer space. The construction respects equivalent descriptions and is unchanged by positive rescaling of the reference. $\square$

11.1 Maximum entropy: the law of least information

Theorem 11.3 (MU forces maximum entropy under finite counting symmetry). *Fix a reference measure $\mu \in \mathfrak{M}_\Pi$ and a nonempty convex feasible set. When the relative-entropy projection is attained, strict convexity makes it the unique least-assuming point completion relative to μ . On a finite atom space with permutation-symmetric counting reference U ,*

$$\arg \min_{P \in \mathcal{P}_c} D_{KL}(P\|U) = \arg \max_{P \in \mathcal{P}_c} H(P),$$

where $H(P) = -\sum_i P_i \log P_i$.

Proof. The first claim is Least-Assuming Completion with the characterised divergence. On a finite N -atom space,

$$D_{KL}(P\|U) = \log N - H(P),$$

so the two optimisation problems have identical minimisers/maximisers. $\square$

Maximum entropy is Epistemic Zero in counting coordinates. It is the feasible state with minimum extra information relative to the uniform counting reference. Multiple permitted references produce the full reference-indexed family, and every attainment boundary remains part of the answer.

Theorem 11.4 (Exponential-family form). *When the reference measure is fixed and the feasible set is determined by attained expectation constraints $\mathbb{E}_P[f_j] = c_j$, an interior MaxEnt solution has density*

$$\frac{dP^*}{d\mu}(x) = \frac{1}{Z(\lambda)} \exp \left(\sum_j \lambda_j f_j(x) \right).$$

Proof. This is the Euler–Lagrange equation for the strictly convex relative-entropy objective subject to affine constraints. Existence and boundary cases require the usual integrability and attainment hypotheses. $\square$

11.2 Complexity-constrained priors and Occam ordering

Theorem 11.5 (Complexity-constrained prior). *Let a discrete case declare a complexity function K , a counting reference on the represented hypotheses, and the constraint $\mathbb{E}[K] \leq c$. When the point completion is attained, it has the form*

$$P_\lambda(h) = \frac{e^{-\lambda K(h)}}{Z(\lambda)}, \quad \lambda \geq 0.$$

If several inequivalent complexity representations remain permitted by the problem, the complete output retains the corresponding family.

Proof. Apply the exponential-family theorem to the moment constraint on K . The final clause is the prior-completion theorem. $\square$

This yields a complexity ordering relative to a chosen code or representation. Kolmogorov complexity, for example, depends on a universal machine that remains part of the specification; another machine defines another case.

11.3 Reference measures

Theorem 11.6 (MU fixes the reference answer). *Problem-preserving transformations determine the reference answer. The permitted family contains every reference class preserved under the allowed changes of representation and apparatus in $\mathcal{R}_\Pi$ (Jeffreys, 1946; Kass and Wasserman, 1996). Singleton cases include:*

- (i) *a compact transitive symmetry with a unique normalised invariant measure;*
- (ii) *a problem that explicitly supplies the Fisher information metric and asks for its induced Riemannian volume class; when that volume is normalisable, its normalisation gives the Jeffreys probability case;*
- (iii) *an affine coordinate and length structure with congruence invariance, giving Lebesgue measure and, on a bounded interval, its normalised uniform case.*

Several inequivalent natural references produce their full family. An empty normalised or usable sigma-finite reference class appears with its certificate.

Proof. In the group-symmetry and affine cases, the allowed morphisms act transitively on the relevant representation and the corresponding invariant measure is unique up to scale. In the Fisher case, the metric is part of the problem and its Riemannian volume form is thereby fixed up to the usual normalisation issue. The invariance condition defining a permitted reference preserves these classes under redescription. Multiple permitted classes give multiple objects of $\mathfrak{M}_\Pi$; an empty permitted class gives the empty reference answer. $\square$

Proposition 11.7 (Noncompact symmetry determines an invariant measure class). *Translation symmetry on $\mathbb{R}$ and scale symmetry on $(0, \infty)$ determine sigma-finite invariant measure classes. Each symmetry also carries a normalisation obstruction: every countably additive invariant Borel probability would have total mass either zero or infinite.*

Proof. For translation invariance, let $a = \mu([0, 1))$. Every interval $[n, n + 1)$ has the same mass a , and these intervals are disjoint and cover $\mathbb{R}$. If $a > 0$, countable additivity makes the total mass infinite; if $a = 0$, it makes the total mass zero. Both contradict $\mu(\mathbb{R}) = 1$.

For scale invariance, let $b = \nu([1, 2))$. Every interval $[2^n, 2^{n+1})$ has the same mass b , and these intervals are disjoint and cover $(0, \infty)$. The same alternatives give total mass either infinite or zero, contradicting normalisation. $\square$

A noncompact symmetry therefore determines a sigma-finite measure class, Lebesgue measure dx under translation or Haar measure dx/x under rescaling, together with a normalisation obstruction. A proper probability emerges when the problem adds location, scale, truncation, likelihood, propriety, or another normalising structure.

Lemma 11.8 (Interval equiplausibility). *Suppose the problem states an affine coordinate $r \in [a, b]$, its length structure, and congruence invariance of subintervals as its reference symmetry. Then the unique normalised reference is Lebesgue measure divided by $b - a$.*

Proof. Congruence invariance gives equal measure to equal-length subintervals. Finite additivity, continuity, and normalisation imply

$$\mu([x, y]) = \frac{y - x}{b - a}.$$

The affine and congruence structure determines the uniform measure; numerical bounds alone leave the reference unresolved. $\square$

11.4 Four worked examples

The four shapes of Definition 3.4 each arise inside this section's machinery, and one short computation exhibits each.

A point. Take the fair-die problem: six outcomes with the full permutation symmetry stated. The symmetry admits no invariant probability other than

$$p_i = \frac{1}{6}, \quad H(p) = \log 6,$$

and among all probabilities the uniform one uniquely maximises entropy relative to counting measure, so the symmetry route and the projection route meet at the same point.

A family. Let the stated symmetries leave two invariant references μ_1 and μ_2 standing, with neither preferred by the problem. Each admits its own projection

$$P_k^* = \arg \min_{P \in \mathcal{C}} D(P \parallel \mu_k),$$

and Theorem 11.1 returns the indexed family $\{(\mu_k, P_k^*)\}$ as the complete answer. Selecting one member would add a preference the problem never supplied.

An improper reference class. Under translation on $\mathbb{R}$ the invariant measure class is Lebesgue, and $\int_{\mathbb{R}} dx = \infty$, so no normalised invariant probability exists. The complete answer records the sigma-finite class together with the normalisation obstruction, and further stated structure then restores propriety.

An unattained optimum. On two atoms with uniform reference, minimise $D(p \parallel \mu)$ over the open case $\mathcal{C} = \{p : p_1 > \frac{1}{2}\}$. The infimum is

$$\inf_{p_1 > \frac{1}{2}} \left[p_1 \log 2p_1 + (1 - p_1) \log 2(1 - p_1) \right] = 0,$$

approached as $p_1 \downarrow \frac{1}{2}$ and attained by no member of $\mathcal{C}$, since the minimising point lies on the excluded boundary. The complete answer is the certificate: the case $\mathcal{C}$, the infimum 0, and the boundary point $(\frac{1}{2}, \frac{1}{2})$.

The examples distinguish a point, a reference family, a sigma-finite invariant class, and an unattained optimum, and each is closed by exhibiting rather than by report.

11.5 Finite typicality and reference dependence

The method of types gives a useful independent motivation in the finite counting case. For N i.i.d. draws, a type class with empirical distribution $\hat{P}$ has cardinality exponential in $NH(\hat{P})$ up to subexponential factors. The MaxEnt type is therefore the largest compatible class relative to the product counting measure. The result explains the typicality of the MaxEnt completion relative to the supplied counting measure.

Remark 11.2 (Uniformity depends on the parameterisation). Competing “uniform” answers reveal competing representation cases. Stating the physical randomisation or symmetry identifies the relevant reference. Several permitted structures produce their complete prior family. Indifference is the quantitative expression of symmetry in a defined case.

Definition 11.4 (Credal set). For a Boolean problem Π , the **credal set** is the set of point-valued probability answers permitted by the problem (Walley, 1991; White, 2005; Kopec and Titelbaum, 2016). It contains every attained point answer and no others. It coincides with the raw feasible simplex only when the problem independently supports every feasible point.

12 Updating

Updating is MU applied through time. The current state carries the information already earned; new constraints purchase a specific departure from it. KL projection is Epistemic Zero through time: each new constraint changes exactly what it touches and carries every compatible feature forward. Relative entropy measures that departure, so the update is the full set of KL projections: one point, a tied family, or an attainment boundary (Csiszár, 1975).

Theorem 12.1 (MU’s retention law). *Let P be the current state and let $\mathcal{C}$ be a new feasible constraint set. Revision carries forward every feature of P compatible with $\mathcal{C}$ and changes exactly what joint satisfaction requires. If the case represents change by a divergence D , the complete update is*

$$\arg \min_{Q \in \mathcal{C}} D(Q \parallel P),$$

with all minimisers retained and every unattained optimum reported.

Proof. The new constraints purchase a precise departure from the current state. The divergence orders feasible revisions by that departure. A revision with greater divergence introduces a larger change than another feasible revision, while the constraints support the smaller departure.

MU therefore selects the minimum, and Full-Answer Preservation returns the entire minimiser set. $\square$

Theorem 12.2 (MU forces KL projection). *On the finite probability case of the relative-entropy theorem, let P be the current state and let $\mathcal{C}$ be a feasible set. Define*

$$\text{Upd}_{\mathcal{C}}(P) = \arg \min_{Q \in \mathcal{C}} D_{KL}(Q \| P).$$

This minimiser set is MU's retention update. An unattained infimum is returned with its certificate. On a nonempty closed convex feasible set containing a point with finite divergence from P , the minimiser exists and is unique. For a credal current state, the construction is applied to each member and every attained result is retained.

Proof. Theorem 10.5 fixes the divergence up to a positive scale, preserving its minimisers. Theorem 12.1 supplies the minimisation direction and the full minimiser set. Finite-space existence and uniqueness follow from Corollary 12.8 below. $\square$

Standard conditioning and Jeffrey conditioning are special feasible sets. If the new constraint is $Q(E) = 1$, the minimiser is

$$Q(H) = P(H | E) = \frac{P(E | H)P(H)}{P(E)}.$$

If the new constraint is $Q(E) = q$, the minimiser preserves the old conditionals within E and $\neg E$ and changes only their mixture weights.

Remark 12.1 (The reporting protocol matters). A statement such as “the outcome lies in this disjunction” may correspond to several reporting, sampling, or censoring protocols. Each protocol defines a likelihood kernel. The channel model must therefore retain every protocol consistent with the evidence and carry each posterior forward.

12.1 When sequential updates agree

The divergence theorem supplies exact decomposition and invariance. The retention rule determines how that divergence is used in revision.

Proposition 12.3 (Conditions for sequential agreement). *For a given prior and constraint set, the retention rule has the following properties whenever the relevant structure is present:*

- (i) *the same specified pair has the same complete posterior set;*
- (ii) *equivalent coordinates give equivalent posterior sets;*
- (iii) *product systems with factorised priors and factorwise constraints update factorwise;*
- (iv) *a constraint confined to one conditional component preserves the unconstrained marginal and other conditionals.*

These clauses state uniqueness, coordinate invariance, system independence, and subset independence for revision.

Proof. The first two clauses are MU generation and invariance under equivalent descriptions. For the third, the product structure and factorwise constraints make the objective a sum of factor objectives. For the fourth, exact accounting separates marginal and conditional change,

and retention leaves every unconstrained term at zero change. Each conclusion therefore follows exactly on the structure named in its antecedent. $\square$

12.2 Finite convex geometry

Lemma 12.4 (Convexity of the finite constraint language). *On a finite outcome space, normalisation, support restrictions, linear probability bounds, and expectation equalities or inequalities define closed convex feasible sets. Finite intersections of such constraints are closed and convex.*

Proof. Each condition is an affine equality, affine inequality, or coordinate face of the simplex. These sets are closed and convex, and finite intersections preserve both properties. $\square$

Proposition 12.5 (Finite-space extension). *The KL projection theorem applies to every nonempty closed convex feasible set in a finite simplex that contains a point of finite divergence from the reference. Infinite and noncompact spaces require an explicit topology, tightness, coercivity, and attainment hypotheses.*

Theorem 12.6 (Pythagorean identity). *Let $\mathcal{C} = \{Q : \mathbb{E}_Q[f_j] = c_j\}$ be a linear family on a finite outcome space, let P be strictly positive, and suppose $\mathcal{C}$ contains a strictly positive point. If Q^* minimises $D_{KL}(\cdot \| P)$ on $\mathcal{C}$, then for every $Q \in \mathcal{C}$,*

$$D_{KL}(Q \| P) = D_{KL}(Q \| Q^*) + D_{KL}(Q^* \| P).$$

Proof. The Lagrange equations give

$$\log \frac{Q_\omega^*}{P_\omega} = \lambda_0 + \sum_j \lambda_j f_j(\omega).$$

Its expectation is therefore the same for every $Q \in \mathcal{C}$. Expanding $D_{KL}(Q \| P)$ through Q^* gives the displayed identity. $\square$

Lemma 12.7 (Pinsker inequality). *For finite distributions,*

$$D_{KL}(P \| Q) \geq 2d_{TV}(P, Q)^2.$$

Proof. Coarse-grain by the indicator of $\{P > Q\}$ and apply Lemma 10.4. The problem reduces to the binary inequality $\text{kl}_2(a, b) \geq 2(a - b)^2$. With b fixed, subtract the right side; the resulting function has value and first derivative zero at $a = b$ and second derivative $1/[a(1 - a)] - 4 \geq 0$. $\square$

Corollary 12.8 (Projection existence, uniqueness, and alternation). *On a finite simplex, the KL projection onto a nonempty closed convex set containing a finite-divergence point exists; it is unique because $D_{KL}(\cdot \| P)$ is strictly convex on the relevant face.*

Let $\mathcal{C}_1$ and $\mathcal{C}_2$ be linear families whose intersection contains a strictly positive point. Alternating KL projections between them converge to the unique projection onto $\mathcal{C}_1 \cap \mathcal{C}_2$.

Proof. Existence follows from compactness of the finite simplex and lower semicontinuity; strict convexity gives uniqueness. For alternating projection, Theorem 12.6 yields, for any $R \in \mathcal{C}_1 \cap \mathcal{C}_2$,

$$D_{KL}(R \| Q_k) = D_{KL}(R \| Q_{k+1}) + D_{KL}(Q_{k+1} \| Q_k).$$

Telescoping shows that the step divergences are summable, so Pinsker gives $d_{TV}(Q_{k+1}, Q_k) \rightarrow 0$. Compactness gives an accumulation point Q_∞ ; even and odd subsequences lie in the two closed families, and their asymptotic equality places Q_∞ in the intersection. Passing to the limit in the telescoped identity gives, for every R in the intersection,

$$D_{KL}(R\|Q_0) = D_{KL}(R\|Q_\infty) + D_{KL}(Q_\infty\|Q_0).$$

Thus Q_∞ is the KL projection of Q_0 onto the intersection. Strict convexity makes it unique, so every accumulation point is the same and the whole sequence converges. $\square$

12.3 Path independence

Theorem 12.9 (Path independence). *For compatible finite linear constraints, the simultaneous KL projection is a function of the joint feasible set. Alternating I-projection converges to it under Corollary 12.8. A single sequential pass reaches the same point when the relevant projection operators commute or preserve one another's constraint families.*

Proof. The joint feasible set $\mathcal{C}_1 \cap \mathcal{C}_2$ is invariant under presentation order, and strict convexity gives one simultaneous projection. Alternating convergence is Corollary 12.8. One-pass agreement follows when the projection operators commute at the starting state or preserve one another's constraint families. $\square$

Commutation supplies one sufficient route to one-pass agreement. Nonlinear and nonconvex constraints call for their own algorithms and convergence theorems, while the simultaneous minimiser set remains the retention target whenever it exists.

12.4 Constraint retraction

Removing a constraint creates a new problem with a different feasible set. The state is recomputed from the retained background and the relevant reference or current state. Learning that a source is unreliable is different again: the signal remains, while the channel kernel or its latent state is revised.

13 Full Answers Can Be Sets

When several inequivalent point completions survive, their credal family is the answer. A point appears exactly when the grounds, an explicit reference, or a downstream selector singles one out.

Remark 13.1 (Lawful refusal carries a complete answer). A question can demand a number while its grounds determine an interval, a family, or an empty value class. The full answer is that family or certificate, together with the exact boundary of what the constraints determine. For a human reasoner this is suspension of judgement. For a machine it is lawful refusal. A system that returns the answer's true shape displays greater rational power than one that compresses every problem into a point.

Remark 13.2 (Credal state, recalled). By Definition 11.4, a credal state is the full set of point-valued answers the problem permits. Plurality may arise from an incomplete plausibility comparison, several permitted reference structures, or several tied minimisers. Attainment data remain a separate part of the answer.

Theorem 13.1 (MU returns the credal family). *When several inequivalent point completions remain, the credal family is the full answer. A selector enriches the problem and yields the corresponding downstream point.*

The result handles underspecified priors by returning the compatible family. Sparse evidence, unresolved representation choices, several natural reference structures, and uncertain channel models can each produce a plural state.

Proposition 13.2 (Full-answer uniqueness with plural members). *Identical problems determine the same complete permission set, which may contain one or many members. Agents sharing that credal family may choose different actions because utilities, losses, and resources belong to the downstream decision problem. Different evidence, representation, reference families, channel models, or computational access define different belief problems. Full identity of every answer-relevant input yields identity of the full permission set.*

The distinction resolves the uniqueness–permissivism tension at the level that matters. The answer space is unique even when it contains several permissible points, and practical choice can add further distinctions without changing the belief state.

13.1 Resolute completion

Definition 13.1 (Resolute completion). A **resolute completion** is a point used because an additional reporting or decision protocol requires one. It is chosen by the retention rule relative to a reference or current state; resoluteness holds the selected point fixed until new information changes the problem.

Theorem 13.3 (Resolute completion). *Let $\mathcal{K}$ be a nonempty credal family in a finite probability simplex. Suppose a protocol additionally requires one point and supplies a reference measure μ with finite divergence to some member of $\mathcal{K}$. Then the resolute answer is*

$$\text{Res}_\mu(\mathcal{K}) = \arg \min_{P \in \mathcal{K}} D_{KL}(P \| \mu),$$

with every minimiser retained. On a closed convex $\mathcal{K}$ with a finite-divergence point, strict convexity gives one resolute point. Each permitted reference yields its own resolute case, and the full answer retains the corresponding family. A one-number protocol becomes determinate when it also supplies a reference or selector.

Proof. The point-output requirement fixes the form of the report. The reference supplies the comparison among feasible points. Theorem 11.1 therefore gives the KL minimisers, and Definition 3.4 retains every tied case. Corollary 12.8 gives existence and uniqueness under the finite convex conditions. Holding the resulting point fixed until the problem changes is the operational content of resoluteness; every later change tracks new information or a newly stated selector. $\square$

Remark 13.3 (Forced bets). A reporting protocol may require one numerical credence while the belief state remains a set. A resolute completion then uses an explicit reference or selection rule. A decision problem may also choose an action directly from the credal set through a loss function or robustness criterion. In both cases, the downstream commitment is a new output layered on the underlying belief state.

Remark 13.4 (Uniformity under reparameterisation). Competing “uniform” answers under interconvertible parameterisations reveal an unresolved reference structure. A physical randomisation protocol, affine length, Haar symmetry, or likelihood geometry fixes a case; several natural structures produce their full prior family. The paradox therefore tests and improves the specification of the problem.

Remark 13.5 (Case dependence across limits). A sequence of finite approximations defines a case when the exhaustion and its convergence are part of the problem. Different exhaustions can have different limits, and some sequences carry an unattained limit. The limiting procedure therefore belongs inside the problem and selects according to that specification.

14 From Belief to Choice

A belief state records what a case treats as possible or probable. Choice adds a value functional and a scale for departing from the reference state. These are further inputs layered on the belief state. Once they are supplied, the relative-entropy geometry above fixes the choice distribution and makes its support boundary explicit.

Theorem 14.1 (Choice from reference, value, and cost). *Let $(\mathcal{X}, \mathcal{F}, \mu)$ be a probability space, let $V : \mathcal{X} \rightarrow \mathbb{R}$ be measurable, and let $\tau > 0$. Suppose*

$$0 < Z_{\mu, V, \tau} = \int_{\mathcal{X}} e^{V/\tau} d\mu < \infty, \quad \int_{\mathcal{X}} |V| e^{V/\tau} d\mu < \infty.$$

Define

$$\frac{d\rho^*}{d\mu} = \frac{e^{V/\tau}}{Z_{\mu, V, \tau}}.$$

For every probability measure $\rho \ll \mu$ with $D_{KL}(\rho||\mu) < \infty$ and $\mathbb{E}_{\rho}[|V|] < \infty$,

$$\mathbb{E}_{\rho}[V] - \tau D_{KL}(\rho||\mu) = \tau \log Z_{\mu, V, \tau} - \tau D_{KL}(\rho||\rho^*).$$

Consequently ρ^ is the unique maximiser of*

$$\rho \longmapsto \mathbb{E}_{\rho}[V] - \tau D_{KL}(\rho||\mu)$$

on that domain. If the problem retains several references, value cases, or costs of departure, its complete output retains the corresponding family of tilted completions.

Proof. The integrability assumptions place ρ^* in the displayed domain. Since its density is strictly positive μ -almost everywhere,

$$\log \frac{d\rho}{d\rho^*} = \log \frac{d\rho}{d\mu} - \frac{V}{\tau} + \log Z_{\mu, V, \tau}.$$

Integration with respect to ρ gives the identity. Nonnegativity of relative entropy gives the upper bound $\tau \log Z_{\mu, V, \tau}$, with equality only when $\rho = \rho^*$ almost everywhere. $\square$

Corollary 14.2 (Zero support stays zero). *The choice from reference, value, and cost is equivalent to its reference measure:*

$$\rho^* \ll \mu, \quad \mu \ll \rho^*.$$

In particular,

$$\mu(A) = 0 \implies \rho^*(A) = 0.$$

A specified hard exclusion can be represented by restricting the reference to an allowed set, equivalently by taking the extended value $V = -\infty$ outside that set. Such a restriction may remove support but can never create support beyond the original reference.

Proof. For finite V , the density $e^{V/\tau}/Z$ is positive and finite μ -almost everywhere, proving equivalence. The restricted form follows by applying the theorem to the conditional reference on the allowed set; viewed in the original space, the resulting measure remains absolutely continuous with respect to μ . $\square$

Values reweight the possibilities carried by the reference. On a topological state space the tilt preserves $\text{supp } \mu$. A truth, action, or population outside the represented measure class enters through an enlarged representation or reference.

Proposition 14.3 (Reference, value, and cost are separate inputs). *The triplet (μ, V, τ) occupies three distinct roles. The reference μ records the comparison structure before value is applied; V records what the decision problem rewards or prefers; τ sets the cost of departing from the reference in units of value. Changing any one changes the specified decision problem. The problem itself supplies V and τ , while MU governs the conclusions that follow from them.*

Definition 14.1 (Credal dominance). Let $\mathcal{K}$ be a nonempty set of probability measures, let $\mathcal{A}$ be a finite action set, and let $u(a, x)$ be bounded. Action b credally dominates a when

$$\mathbb{E}_P[u(b, X)] \geq \mathbb{E}_P[u(a, X)] \quad \text{for every } P \in \mathcal{K},$$

with strict inequality for at least one $P \in \mathcal{K}$.

Theorem 14.4 (Decision under imprecise belief). *Under the preceding conditions, the set*

$$\mathcal{A}^*(\mathcal{K}, u) = \{a \in \mathcal{A} : \text{no } b \in \mathcal{A} \text{ credally dominates } a\}$$

is nonempty. Every action outside $\mathcal{A}^$ is worse than another permitted action under every probability state in $\mathcal{K}$. If $\mathcal{A}^*$ is a singleton, the decision problem fixes one action. If several actions survive, the complete decision answer is that set until an ambiguity rule, reference, loss refinement, or institutional protocol is supplied.*

Proof. Credal dominance is transitive and irreflexive. Every finite strict partial order has at least one maximal element, so $\mathcal{A}^*$ is nonempty. Every a outside $\mathcal{A}^*$ is credally dominated by some b that weakly improves expected utility under every permitted state and strictly improves it under at least one. A unique maximal element dominates every other element along a finite dominance chain. Several maximal elements form the full undominated decision set, and further decision structure selects a point from it. $\square$

Proposition 14.5 (The information-cost path). *Suppose $\mathcal{X}$ is finite. As $\tau \rightarrow \infty$, the tilted state converges to μ . As $\tau \downarrow 0$, its mass concentrates on the maximisers of V inside $\text{supp } \mu$; if several maximisers remain, their limiting weights are proportional to their reference masses.*

Proof. For large τ , $e^{V(x)/\tau} \rightarrow 1$ at every point, so normalisation gives $\rho^* \rightarrow \mu$. For small τ , subtract $M = \max_{\text{supp } \mu} V$ from V . Every nonmaximising term then contains $e^{(V(x)-M)/\tau} \rightarrow 0$, while the maximising terms retain their factors $\mu(x)$. Normalisation gives the stated limit. $\square$

Proposition 14.6 (Score decomposition). *Suppose $\mathcal{X}$ carries a chosen Riemannian metric and volume form, μ has a strictly positive C^1 density, and V is C^1 . Then*

$$\nabla \log \frac{d\rho^*}{d\text{vol}} = \nabla \log \frac{d\mu}{d\text{vol}} + \frac{1}{\tau} \nabla V.$$

Thus the local score of the tilted state is the sum of a reference term and a value term, relative to that geometry.

Remark 14.1 (A reversible realisation). Let $p = d\rho^*/d\text{vol}$. Under the usual smoothness, nonexplosion, confinement, and ergodicity hypotheses, the diffusion with generator

$$\mathcal{L}f = \Delta f + \langle \nabla \log p, \nabla f \rangle$$

is reversible with stationary law ρ^* . In Euclidean coordinates this is the overdamped Langevin equation

$$dX_t = \nabla \log p(X_t) dt + \sqrt{2} dW_t.$$

Within the isotropic reversible diffusion class, this is the standard relaxation. Preconditioned and nonreversible processes can realise the same stationary law in other ways.

15 Summary of MU's Quantitative Forms

The quantitative results form one generated sequence. Scalar plausibility yields probability; probability and hierarchical decomposition yield relative entropy; relative entropy and a permitted reference yield starting belief; relative entropy and retention yield revision; value and cost yield choice. Figure 2 displays the sequence.




Figure 2: One generated sequence. Each arrow names the structure a domain supplies; each box is the unique form MU forces once that structure is present.

One source of quantitative force. The domain identifies the object: Boolean plausibility, informational change, starting belief, revision, or choice. MU generates the laws that preserve

its grounds and equivalences. The corresponding theorem supplies the unique form. A change of logic, representation, reference, value, cost, or output question defines a new problem and receives its own full answer.

The classification here concerns the Boolean scalar case. Orthomodular, intuitionistic, relevant, and paraconsistent structures define neighbouring problems. Projection lattices in Hilbert space provide an important non-Boolean case, where Gleason-type theorems classify probability assignments under their own dimensional, additive, and regularity conditions (Gleason, 1957).

Dempster–Shafer belief functions (Dempster, 1967; Shafer, 1976), possibility measures, and other capacities occupy their corresponding domains. Some represent structured credal families, some encode source-combination assumptions in the channel model, and others live on non-Boolean event structures. MU classifies each through the grounds that define it.

16 Evidence Channels

Evidence from the world is represented by stochastic kernels. This connects the distinction between internal support and route reliability to reliabilist and anti-bootstrapping work (Goldman, 1979; Vogel, 2000; Cohen, 2002). Perception, memory, testimony, instruments, and models differ in their signals, errors, dependencies, incentives, and hidden source conditions, yet once those kernels are specified they share the same update form. A single reliability number is only a special summary. Logical consequence remains an internal relation.

Definition 16.1 (Evidence channel). An **evidence channel** is a tuple

$$C = (\mathcal{H}, \mathcal{S}, \mathcal{E}, K),$$

where $\mathcal{H}$ is the world or hypothesis space, $\mathcal{S}$ the signal space, $\mathcal{E}$ a space of source, apparatus, or context states, and

$$K(ds | h, \eta)$$

is a stochastic kernel. Given a prior $P(dh, d\eta)$ and observed signal s , the posterior is

$$P(dh, d\eta | s) \propto K(ds | h, \eta)P(dh, d\eta).$$

The source state η may encode calibration, competence, incentives, selection effects, common causes, or adversarial control.

A scalar reliability parameter is a summary available only in special cases. For a symmetric binary channel one may write

$$r = \Pr(S = H | H) = \Pr(S \neq H | \neg H),$$

while general cases retain sensitivity, specificity, base rates, and dependence as separate features of the kernel. Figure 3 displays the structure.

Proposition 16.1 (Channel update). *For a prior and stochastic kernel, observing s determines the posterior by Bayes’ rule, subject to whatever point, family, or nonexistence structure is already present in the prior. If several channel models remain possible, each is carried forward.*

Proof. Bayes’ rule is the conditional probability identity inside each probability case. Unresolved channel cases remain part of the problem, so the complete posterior carries forward every still-compatible prior-and-channel case. $\square$




Figure 3: How signals bear on belief. A stochastic kernel carries the evidential relation; a scalar reliability is only a special summary.

16.1 Channel learning, composition, and source types

Definition 16.2 (Reliability terminology). For a general channel, **reliability** denotes the kernel or a chosen functional of it. **Fidelity** denotes the probability of error-free transmission through a particular multi-stage model. **Aggregate reliability** denotes the effective kernel after composition or aggregation. The **posterior** is the epistemic state obtained from the prior and kernel; reliability remains a property of the kernel or a chosen functional of it.

Proposition 16.2 (Learning a channel). *Unknown channel parameters and latent source states are ordinary hypotheses. Calibration data update their joint posterior together with the target-world hypotheses. For repeated data (s_i, y_i) with later-verified outcomes y_i , a specified hierarchical kernel $K(s_i | y_i, \eta)$ yields a posterior over η and a posterior-predictive channel. The learned channel is therefore richer than a raw sample proportion.*

Proof. Place the unknown channel state η in the hypothesis space and apply the channel-update proposition to the joint prior $P(h, \eta)$. Track records, controls, cross-checks, and later-verified outcomes are signals informative about η ; integrating over its posterior gives the predictive kernel. $\square$

Proposition 16.3 (Channel identifiability). *A channel's output distribution identifies truth-conditional reliability only when later-verified outcomes, independent controls, another calibrated route, or sufficient structural restrictions anchor the latent truth model.*

Proof. Observed signals determine their marginal law. Truth-conditional calibration additionally requires a decomposition into latent truths and a conditional kernel. In a binary symmetric channel with truth prevalence q and accuracy r ,

$$\Pr(S = 1) = qr + (1 - q)(1 - r).$$

Many pairs (q, r) yield the same observable value; in particular (q, r) and $(1 - q, 1 - r)$ do. More generally, distinct latent-state/kernel pairs can induce the same signal distribution. Agreement of a source with its own past outputs supplies consistency data; external truth labels come from calibration anchors. $\square$

Bootstrapping and easy-knowledge failures are cases of this non-identifiability. Repeated self-agreement establishes consistency; an independent anchor or identifiable structural model establishes truth calibration.

Remark 16.1 (Higher-order evidence). Learning that one is tired, intoxicated, biased, under adversarial pressure, or using a degraded instrument is evidence about η . The same joint calculus

handles this higher-order evidence, changing both the target belief and the estimated quality of the route by which it was formed.

Proposition 16.4 (Channel composition). *If a chain has conditionally independent kernels K_i with explicitly specified intermediate states, its effective channel is their Markov composition. In the independent error-free-transmission model, path fidelity is $\prod_i r_i$; structured errors may cancel, correlate, or amplify.*

Remark 16.2 (Dependence and agreement). Agreement among several reports becomes strong evidence through independence, calibration, and discriminative likelihood. Correlated sources, common training data, shared sensors, selection effects, and copying may make many reports one effective channel. A joint kernel and an independent truth-tracking anchor reveal the evidential value of the agreement and any surviving observational equivalence.

Perception and instruments

Perception and instruments are kernels from world states to registrations. Their evidential value is fixed by discriminative likelihoods, calibration, and the modelled apparatus state. Illusion, masking, calibration drift, and limited resolution are modelled by the same object.

Memory

Memory is a temporally extended reconstruction channel. Its latent state can include retention, interference, present cues, and rehearsal. Confidence in a memory and reliability of the memory process are distinct variables.

Introspection

The occurrence of a present phenomenal state may be immediate, while its classification, cause, persistence, and report remain fallible. Doxastic and motivational self-knowledge are ordinary noisy channels whose reliability varies with occurrence, classification, attention, and self-deception.

Testimony

A report affects odds through its likelihood ratio

$$\frac{K(\text{report } H \mid H, \eta)}{K(\text{report } H \mid \neg H, \eta)}.$$

Competence and honesty contribute differently to this ratio. Testimonial trust therefore follows from a base-rate and source model, with competence, honesty, incentives, and dependence retained as distinct features.

Proposition 16.5 (Base-rate conditional trust). *When a population model and calibration data make truthful reporting more likely than the same report under falsity, testimony raises the posterior of the reported claim. A likelihood ratio of one leaves the prior unchanged; anti-correlation can make the inverted report informative.*

Adversarial sources

Strategic sources require a kernel conditional on goals, information, incentives, and adaptation to the evaluator. Stationarity and independence receive explicit status inside that model.

A priori consequence remains internal

Logical consequence belongs to the internal inferential relation. Its standing comes from the stated consequence system. If $B \vdash H$, every probability representation respecting that logic assigns $P(H \mid B) = 1$; probability one alone remains weaker than entailment.

16.2 Channels and epistemic closure

Definition 16.3 (Robustness of the route). A true belief formed through channel model C is **route-robust** over a stated relevant set W when, throughout W , the same route under the stipulated source conditions continues to track the truth.

Definition 16.4 (Empirical knowledge proposal). An agent empirically knows H through route C when H is true, the belief is internally supported, and the route satisfies the stated robustness or competence condition. This is a proposal for protection against epistemic luck, an anti-luck account in the standard terminology, as a joint condition on truth, support, and route integrity.

Theorem 16.6 (Internal justification is closed). *If B internally licenses P and the specified consequence relation contains $P \vdash Q$, then B internally licenses Q .*

Proof. Closure is composition in the chosen consequence relation. $\square$

Theorem 16.7 (Route conditions govern transmission). *Route robustness transfers across an entailment exactly when the route also discriminates the alternatives introduced by the entailed proposition.*

Proof. A channel may reliably distinguish ordinary hand-present from ordinary hand-absent situations while failing to distinguish either from an adversarial simulation. The entailment relation is internal; the additional discrimination demand is a property of the kernel that distinguishes the truth. Nothing in logical closure supplies it. $\square$

The relevant-world set is a world-facing input used in the Gettier and scepticism analyses.

17 Convergence to Truth

When the channel is specified, the truth is represented and supported, observations follow the sampling law, and false rivals are distinguishable, the posterior concentrates on the represented truth. These are the standard world-facing conditions of Bayesian consistency (Doob, 1949; Schwartz, 1965; Ghosal and van der Vaart, 2017).

Lemma 17.1 (Law of large numbers, bounded case). *Let $Z_1, Z_2, \dots$ be i.i.d. with $|Z_i| \leq M$ and mean μ . Then $\frac{1}{n} \sum Z_i \rightarrow \mu$ almost surely.*

Proof. For centred bounded variables, Hoeffding's bound

$$\Pr(|\bar{Z}_n - \mu| > \varepsilon) \leq 2e^{-n\varepsilon^2/(2M^2)}$$

follows from Markov's inequality applied to $e^{s\bar{Z}}$ and the elementary estimate $\mathbb{E}e^{sZ} \leq e^{s^2M^2/2}$. The bound is summable in n , so Borel–Cantelli yields almost-sure convergence.

On a finite observation alphabet, likelihoods positive on the support of H^* give bounded per-observation log-likelihood ratios. The convergence proof therefore invokes only this lemma; a

hypothesis assigning zero likelihood to an event with positive true probability is eliminated when that event occurs. $\square$

Theorem 17.2 (Posterior Convergence). *Let observations be i.i.d. on a finite alphabet under the specified channel model, with true per-observation distribution p_0 . Let the prior Π over represented hypothesis distributions give positive mass to every KL-neighbourhood of p_0 . Then, for every $\delta > 0$, the posterior mass of*

$$\{p : d_{TV}(p, p_0) \geq \delta\}$$

tends to zero almost surely, exponentially fast. For a finite identifiable hypothesis class this is equivalent to posterior mass on the true hypothesis converging to one.

Proof. Denominator. Fix $\varepsilon > 0$ and choose a smaller closed KL neighbourhood

$$B = \{p : D_{KL}(p_0 \| p) \leq \varepsilon\}$$

with positive prior mass. On a finite alphabet, data processing to each coordinate implies that B is bounded away from zero on $\text{supp } p_0$; hence $\log(p_i/p_{0,i})$ is uniformly bounded and continuous on B . Since the empirical frequencies $\hat{p}_n$ converge almost surely to p_0 ,

$$\frac{1}{n} \log \prod_{j=1}^n \frac{p(X_j)}{p_0(X_j)} = \sum_i \hat{p}_{n,i} \log \frac{p_i}{p_{0,i}}$$

converges uniformly on B to $-D_{KL}(p_0 \| p)$. Therefore, almost surely and eventually, every $p \in B$ has likelihood ratio at least $e^{-2n\varepsilon}$. The marginal likelihood denominator is consequently at least

$$\Pi(B)e^{-2n\varepsilon}.$$

Numerator. Almost surely, $d_{TV}(\hat{p}_n, p_0) < \delta/2$ eventually. If $d_{TV}(p, p_0) \geq \delta$, then $d_{TV}(p, \hat{p}_n) \geq \delta/2$, and Pinsker gives

$$D_{KL}(\hat{p}_n \| p) \geq \frac{\delta^2}{2}.$$

Meanwhile $D_{KL}(\hat{p}_n \| p_0) \rightarrow 0$. Hence, uniformly over the far set,

$$\frac{1}{n} \log \prod_{j=1}^n \frac{p(X_j)}{p_0(X_j)} = D_{KL}(\hat{p}_n \| p_0) - D_{KL}(\hat{p}_n \| p) \leq -\frac{\delta^2}{4}$$

eventually.

Combine. Choose $\varepsilon < \delta^2/16$. The posterior mass of the far set is bounded by

$$\Pi(B)^{-1} \exp\left(-\frac{n\delta^2}{4} + 2n\varepsilon\right),$$

which tends to zero exponentially. In a finite identifiable hypothesis class, a sufficiently small δ isolates p_0 , so posterior mass on the true hypothesis tends to one. $\square$

Remark 17.1 (Discrete, continuous, and set-valued forms). For a finite identifiable class, posterior consistency is point-mass convergence. For continuous parameter spaces, the correct statement is concentration in neighbourhoods; for infinite-dimensional spaces, additional model-specific regularity may be required. If the prior output is a credal set, convergence is evaluated separately

for the cases in which the truth lies in prior support; the resulting posterior set records the cases for which that condition fails.

Here is the boundary between rational method and empirical success. The internal rule fixes how evidence changes belief; channel, support, realisability, and distinguishability determine whether that process finds the truth. Figure 4 illustrates the dynamics.




Figure 4: Illustrative convergence dynamics in a finite identifiable class. As observations accumulate under the channel, support, sampling, and distinguishability conditions, posterior mass concentrates on the true hypothesis while false alternatives decay. The rate depends on their information separation.

Support mechanism. Posterior consistency requires the represented truth to lie in the support of the completed prior. A full-support reference helps, while the constrained projection determines the final support: boundary constraints or a boundary optimum may assign zero mass to a candidate. Common finite affine cases preserve support when the reference is strictly positive, the feasible family contains an interior point, and the attained exponential-family solution is interior. The convergence theorem applies to the completed cases carrying the truth in support.

Corollary 14.2 gives the corresponding boundary for optimisation under a stated value: the Gibbs tilt reweights possibilities already supported by its immediate reference and preserves its exclusions. Realisability remains a separate requirement for convergence to truth.

Remark 17.2 (Convergence conditions). The remaining conditions are substantive:

- **Realisability** holds when the true process is represented in $\mathcal{H}$. Model enlargement addresses misspecification by bringing richer hypotheses into the problem.
- **Distinguishability** can fail for observationally equivalent hypotheses. If two hypotheses make identical predictions for all possible observations, no amount of evidence can distinguish them. This is a limit of empirical inquiry itself.

The convergence guarantee is therefore conditional: when the truth is represented, supported, stably sampled, and distinguishable, the Bayesian update under the retention rule concentrates on it.

Remark 17.3 (Finite-sample counterparts). The guarantee above is asymptotic and model-relative; statistical learning theory supplies the finite-sample, distribution-free counterpart. Uniform convergence over a hypothesis class holds exactly when the class has bounded capacity, and probably approximately correct bounds convert that capacity into sample sizes (Vapnik and Chervonenkis, 1971; Valiant, 1984). Conformal methods trade the model for exchangeability and return finite-sample coverage from that single stated condition (Vovk et al., 2005). The accounting is the same in each case: a guarantee holds inside the structure a method states,

capacity, exchangeability, or a represented and supported truth, and the stated structure carries the price.

18 What Each Condition Selects

Each condition earns its place by excluding a surviving rival.

Result	Rival surviving weaker conditions	Excluding condition
Probability	finite qualitative orders with no agreeing measure	rich refinement, scalar composition, and the corresponding coherence
Probability	vector-valued or context-sensitive plausibility states	a one-dimensional scalar output and universal composition
Probability	discontinuous associative operations	order-density and retention of distinctions already supported by the input
Relative entropy	Rényi and other monotone divergences	exact conditional chain rule
Prior completion	one chosen reference among several	a reference family fixed by the representation and symmetries
Revision	a classified divergence paired with an arbitrary update policy	the retention rule
Update	one selected posterior from a minimiser family	selection fixed by the problem
Channels	agreement among correlated or strategically dependent sources	a joint stochastic kernel and specified dependence structure
Convergence	rational updating under a misspecified or non-identifiable model	realisability, support, stable sampling, distinguishability

An inquiry may end with one answer, a symmetry class, a surviving family, or a proof that no answer of the requested kind exists.

Part IV

Classical Problems: Resolutions, Reconstructions, and Boundaries

Our knowledge can only be finite, while our ignorance must necessarily be infinite.

Karl Popper, *Conjectures and Refutations*

The classical problems are where the method displays its range. Some receive direct resolutions. Some become full families or impossibility certificates. Others are reconstructed through representation, projectibility, channels, computation, value, or the world itself.

Each argument states the exact target it settles and the boundary it reveals. The map at the end gathers the resolutions, reconstructions, and remaining work.

19 Foundations, Rules, and Experience

19.1 Regress, the Trilemma, and the Problem of the Criterion

The Münchhausen Trilemma classifies inferential support as regress, circle, or unsupported stopping (Albert, 1968). MU opens the prior fourth case: a finite direct existence proof. The identity instance establishes the floor first; the later analysis then unfolds its form.

Proposition 19.1 (Grounding result). *Within the consequence structures covered by the MU Theorem, the grounding of MU is neither an infinite regress, nor a derivation of MU from MU, nor a stipulation accepted without proof: it is a direct existence proof. The reflexive analysis of inference begins only after that result has been established.*

Proof. The proof of MU has one reflexive inference instance, $P \vdash P$, for a consistent P , and concludes existentially that at least one consistent inference exists. The proof terminates. Its premise is an instance of reflexivity from reflexivity and a consistent premise. The chosen P disappears under existential introduction. The subsequent method theorem analyses the kind whose existence has already been proved. $\square$

The problem of the criterion asks whether one must begin with reliable cases and infer a criterion, or begin with a criterion and identify reliable cases (Chisholm, 1973). The apparent circle treats method and content as though they had to ground one another in the same respect.

Criterion and content. An answer that comes from grounds depends exactly on the distinctions those grounds retain. Equivalent descriptions agree, and a complete report preserves every supported possibility. Concrete premises, models, representations, and channels provide the material to which this rule is applied. Cases test how well that material fits the world while the inferential form remains fixed.

The same distinction preserves fallibilism. A criterion can govern what follows from a model while the model, representation, or channel remains revisable. The circle dissolves because method and content occupy different roles; the empirical task of improving the content remains.

19.2 Rules, Reasons, and Carroll's Regress

Lewis Carroll's tortoise accepts the premises of a valid argument while demanding one further premise before accepting the conclusion (Carroll, 1895). Adding that premise simply creates another occasion on which a rule must be used.

Proposition 19.2 (Rule use is irreducible to an added premise). *Let A and $A \rightarrow B$ be premises in a consequence system whose rules license modus ponens. Add the proposition*

$$R_0 = (A \wedge (A \rightarrow B)) \rightarrow B.$$

Deriving B from A , $A \rightarrow B$, and R_0 requires the application of an inferential rule. A further premise R_1 reproduces the same role at the next level, and the construction iterates. Rule application is therefore irreducible to additional propositional content.

Proof. From A and $A \rightarrow B$, one may infer $A \wedge (A \rightarrow B)$ and then use R_0 to infer B . Both steps use rules. If the use of R_0 is itself required to appear as a premise, introduce a proposition R_1 stating that A , $A \rightarrow B$, and R_0 license B . Applying R_1 again requires a rule. Repeating the demand yields $R_2, R_3, \dots$ but never turns a proposition into the act of validly passing from premises to conclusion. The regress therefore arises from asking a premise to perform the role of the consequence relation. $\square$

Why Carroll's regress stops. Rules and premises occupy their proper roles. A consequence relation can itself be challenged, compared, or revised within another inferential setting whose grounds and rules are explicit. This closes Carroll's regress at the level of rule use.

A second rule-following problem concerns continuation. Finite conduct can fit more than one total rule (Kripke, 1982).

Proposition 19.3 (Finite data support a family of continuations). *Let*

$$D = \{(x_1, y_1), \dots, (x_n, y_n)\}$$

be a finite graph on an infinite domain X , and choose $x_ \notin \{x_1, \dots, x_n\}$. If one total function $f : X \rightarrow Y$ extends D and Y contains at least two values, then another total function $g : X \rightarrow Y$ extends D while $g(x_*) \neq f(x_*)$.*

Proof. Choose $y_* \in Y$ with $y_* \neq f(x_*)$. Define $g(x) = f(x)$ for $x \neq x_*$ and $g(x_*) = y_*$. Then f and g agree on every observed pair in D and differ at the unobserved case x_* . Thus the finite record supports at least two total continuations. $\square$

The proposition settles the exact underdetermination present in the finite record. Questions about meaning and normativity then receive their own grounds. A unique continuation emerges from further structure: a representation of the rule, a simplicity or invariance standard, a communal practice, an explicit algorithm, or another selection principle. Once that structure is stated, its consequences can be classified. The original finite record returns the family of extensions.

19.3 Hypothesis Generation and Higher-Order Problems

A theory of inference evaluates candidates that have entered the problem. The invention of a language, representation, or hypothesis is a higher-order act that enlarges the problem. The remaining choice can instead be made the subject of a higher-order problem.

Definition 19.1 (Higher-order problem). Let D be a given family of possible specifications. For each $d \in D$, let Π_d be the corresponding inferential problem. The higher-order problem $\Pi^\uparrow$ asks for the complete family of pairs (d, a) with

$$a \in \text{Answer}(\Pi_d),$$

modulo the equivalences stated across specifications.

Theorem 19.4 (Higher-order answer). *For every given family D ,*

$$\boxed{\text{Answer}(\Pi^\uparrow) = \{(d, a) : d \in D, a \in \text{Answer}(\Pi_d)\} / \approx .}$$

If the higher-order problem gives no criterion selecting one d , no member of D may be promoted merely because it was considered first or encoded more conveniently. A selector, comparison rule, or cost placed on D defines an enlarged problem to which the same method applies.

Proof. Every displayed pair is supported by one stated case and therefore belongs to the higher-order answer. Omitting a supported pair would arbitrarily reduce that answer; including a pair from no stated case would add unsupported content. The stated equivalence $\approx$ removes only differences that the higher-order problem itself treats as irrelevant. A point selection among the remaining cases requires further structure and therefore changes the problem. $\square$

Search and evaluation. Let $G_i : Z_i \rightarrow \mathcal{H}$ be hypothesis generators.

- (i) If $\text{im } G_1 = \text{im } G_2$ and the evidential problem is otherwise the same, the complete evaluation answer is the same; enumeration order affects bounded access while the full answer remains fixed by the candidate image.
- (ii) If the images differ, the generators define different candidate problems. Evidence compares the hypotheses present in each candidate set.
- (iii) A given family of generators is handled by Theorem 19.4; the full family is the higher-order answer, and a comparison rule selects a generator when supplied.

Creativity enlarges the candidate space; evaluation then applies the same inferential discipline to the enlarged problem.

19.4 Experience and the Myth of the Given

Experience is often asked to serve two roles at once: to be an input that can support belief and to carry its own complete interpretation. The first role is possible; the second is too strong (Sellars, 1956; Pryor, 2000).

Proposition 19.5 (A signal gains evidential force through its channel). *Let $S = s$ be an experiential signal and H a hypothesis. Its evidential force is the likelihood ratio*

$$\Lambda(s) = \frac{P(s \mid H, \eta)}{P(s \mid \neg H, \eta)},$$

where η specifies the channel and relevant background. The same signal value s can support H , oppose H , or be neutral under different permitted channel models.

Proof. Choose two positive numbers $a, b \leq 1$ as the conditional probabilities $P(s \mid H, \eta)$ and $P(s \mid \neg H, \eta)$. If $a > b$, then $\Lambda(s) > 1$ and the signal confirms H ; if $a < b$, it disconfirms; if $a = b$, it is neutral. The sensory token alone fixes none of these relations. $\square$

The result locates immediate defeasible support in the standing relation between signal, background, and channel. An experience enters the problem as a present signal and shifts belief through a standing perceptual channel. The classification of the signal, the background model, and the channel remain revisable. Immediate first-order support and second-order certification therefore occupy distinct levels.

20 Belief, Uncertainty, and Action

20.1 Certainty, Belief, and Knowledge

Definition 20.1 (Deductive certainty, graded belief, and empirical knowledge). Let B be a stated background and H a proposition.

- (i) **Deductive certainty** obtains when $B \vdash H$ in the chosen consequence relation.
- (ii) **Graded belief** or *doxa* is the probability or credal support assigned to H in the evidential domain beyond deductive entailment.
- (iii) **Empirical knowledge**, in the route-integrity proposal used below, is a true belief that is internally supported by B and produced by a channel that remains reliable across the stated class of relevant nearby cases.

Probability one is a model-relative graded assignment. Deductive certainty comes from entailment, and empirical knowledge joins support to a truth-connected route.

Three epistemic questions. The analysis separates three questions. What follows is governed by the stated consequence relation and deductive closure. How strongly it is supported is governed by the probability or credal-state case. Whether the support tracks truth is governed by the stochastic channel and the relevant world class. Validity, graded support, and knowledge answer different questions. Much of classical epistemology becomes clearer once the three are no longer asked to do one another's work.

A priori support. A conclusion is a priori relative to a stated inferential problem when its support depends on the meanings, axioms, and consequence rules fixed by the problem. This status is internal to that problem; empirical inquiry assesses how the language, axioms, or logic describe the world.

The distinction preserves both sides of the traditional dispute. A proof can justify its conclusion without an empirical channel, while the choice and interpretation of the formal framework remain open to mathematical and philosophical assessment.

20.2 Indifference, Permissivism, and Lawful Refusal

The indifference paradoxes reveal that “spread belief evenly” is incomplete until the measure, representation, or randomisation protocol is fixed. Consider a cube whose side length x lies in $[0, 1]$. Uniformity in side length gives

$$\Pr(x \leq \tfrac{1}{2}) = \tfrac{1}{2}.$$

Uniformity in face area $a = x^2$ gives $1/4$, and uniformity in volume $v = x^3$ gives $1/8$. The three answers reveal a reference family: each calculation is coherent relative to its description, and the visible constraints preserve all three descriptions.

Bertrand’s chord problem has the same form: distinct physical procedures for drawing a “random chord” induce distinct measures. By contrast, a physically symmetric die includes a permutation symmetry among its six faces; the uniform answer is fixed directly by the permutation symmetry carried by the problem (Jaynes, 2003).

Indifference classification. An indifference problem has one of four complete outcomes:

- (i) a unique reference and hence a unique point, when the stated physics or symmetry fixes one;
- (ii) a family of reference-relative answers, when several inequivalent descriptions or protocols remain permitted;
- (iii) a nonattainment result, when a permitted completion problem has an infimum but no minimiser;
- (iv) no admissible case, when the constraints are inconsistent or supply no permitted reference of the required kind.

Cases (ii)–(iv) become point-valued when an enriched problem supplies the missing reference, attainment, or consistency structure.

The Judy Benjamin problem gives a related lesson. Conditional information described in ordinary language may correspond to different formal constraints depending on the reporting protocol. Once the information received is clear, the update problem is ordinary; before that specification, there is no single update to classify (van Fraassen, 1981).

Proposition 20.1 (Reference-class underdetermination). *Let H be a proposition and let $R_1, \dots, R_m$ be reference classes that are each compatible with the information. If the resulting conditionals $P(H \mid R_i)$ differ and the problem gives no relevance relation selecting one class, then the complete answer is the family itself:*

$$\{P(H \mid R_i) : 1 \leq i \leq m\}.$$

Proof. Each conditional is a valid answer to a different permitted reference-class completion. Full-Answer Preservation therefore retains every completion. The family is the exhaustive answer supported by the problem. $\square$

Self-location and Sleeping Beauty. Self-locating problems require uncentered world states, centered locations within those worlds, and a rule describing how the present observation is selected. In the Sleeping Beauty problem, the familiar one-half and one-third answers arise from different precise treatments of awakening evidence, diachronic updating, and the measure over centered awakening-moments (Elga, 2000). Once a formalisation is fixed, its consequences can be calculated. The remaining question is interpretive: which formalisation best captures the

ordinary-language story? The same issue governs anthropic and observer-selection problems. The reference class and observation protocol belong to the question.

Uniqueness of the full answer. One fully defined problem has one complete answer. That answer may be one state or a set of permissible credal states. Epistemic permissivism and rational uniqueness therefore concern different levels: permissivism concerns which states belong to the answer, while uniqueness concerns the answer as a whole (White, 2005; Kopec and Titelbaum, 2016).

A demand for one number adds a reporting or decision problem to the underlying belief problem. Given a permitted reference μ , resolute completion chooses the retained point described in §13.1; if several references remain, the downstream point remains a family. The lawful refusal is itself a complete result: the stated grounds purchase the family and mark the point-selection boundary.

20.3 The Lottery, the Preface, Moore’s Paradox, and Higher-Order Evidence

A ticket in a million-ticket lottery is overwhelmingly likely to lose. Its loss remains graded belief, while deductive certainty and empirical knowledge follow their own conditions.

Proposition 20.2 (The lottery lesson). *For every $p < 1$ there are coherent cases in which $\Pr(H \mid B) > p$ while $B \not\vdash H$. Consequently every threshold below one remains within graded support. Ordinary empirical knowledge also depends on route integrity and relevant alternatives. Stakes alter action and conversational standards while the first-order probability remains fixed.*

Proof. Choose a fair lottery with sufficiently many tickets and let H be the proposition that a specified ticket loses. Its probability can exceed any fixed $p < 1$, while the background remains compatible with that ticket winning. Hence $B \not\vdash H$. The same construction separates numerical confidence from a route connected to truth. $\square$

The preface paradox reverses the direction. An author may have high confidence in each sentence of a long manuscript while also expecting that at least one sentence is wrong (Makinson, 1965). Probability theory makes this coherent because conjunction aggregates the small risks carried by its members.

The distinction that resolves the paradoxes. The lottery separates high probability from entailment and knowledge. The preface shows how rational graded beliefs combine across a long conjunction. What remains is a practical and communicative question: when do stakes, assertion norms, or institutional rules license acting or speaking as if a proposition were settled?

20.3.1 Moore’s paradox

The proposition $p \wedge \neg Bp$ can be true: it may be raining while an agent fails to believe that it is raining. The oddity appears when the agent sincerely asserts the conjunction.

Moorean incoherence is pragmatic and inferential. Suppose a sincere assertion of p is evidence that the speaker believes p . Then a sincere assertion of

$$p \wedge \neg Bp$$

requires the assertion channel simultaneously to represent the speaker as believing p and as denying that belief. The proposition can be logically consistent while its sincere assertion remains unstable as a truthful self-report.

Semantic possibility and the norms of assertion and self-representation occupy distinct levels (Moore, 1942). The related sentence “ p , but my evidence does not support p ” may be coherent when the first clause is a report of truth learned through another route. It becomes epistemically akratic when the speaker treats the second clause as an authoritative assessment of all current grounds while continuing to endorse the first without further reason.

20.3.2 Higher-order evidence

Evidence about the reliability of one’s own reasoning is ordinary evidence about a channel. Let R be a parameter describing the competence or calibration of the first-order process and let D be a peer report, audit, or error signal. A coherent update has the hierarchical form

$$P(H, R \mid E, D) \propto P(D \mid H, R, E)P(E \mid H, R)P(H, R).$$

The answer is the update determined by the independence, reliability, and information carried by D , spanning the full range from steadfastness to deference.

What higher-order disagreement requires. A reliable independent warning about one’s method is a new constraint and should affect belief. A warning generated from the same evidence or from a poorly calibrated source may add little. Higher-order evidence is evidence about the route by which first-order support was produced, and its force follows from that route model.

20.4 Dutch Books, Objective Chance, and Reflection

These results concern three bridges often conflated with the probability calculus itself: prices for bets, objective chance, and one’s own future credence.

Proposition 20.3 (Finite Dutch-book coherence). *Let a finite partition $E_1, \dots, E_n$ be priced by numbers $p(E_i)$, where a ticket on E_i costs $p(E_i)$ and pays one unit exactly when E_i occurs. If the prices violate normalisation or finite additivity, an opponent can form a finite portfolio with a strictly positive payoff in every state, assuming tickets may be bought or sold at the quoted prices and payoffs combine linearly.*

Proof. If $\sum_i p(E_i) < 1$, buying one ticket on every partition cell costs less than one and pays exactly one in every state. If the sum exceeds one, selling the same portfolio yields a sure gain. For disjoint A, B , the ticket on $A \cup B$ has exactly the same statewise payoff as the sum of tickets on A and B ; unequal prices permit buying the cheaper side and selling the dearer. Prices outside $[0, 1]$ are handled by buying or selling the corresponding ticket against cash. Thus the stated betting protocol enforces the probability identities (Ramsey, 1926; de Finetti, 1937). $\square$

The argument is conditional on the practical interpretation of credences as fair prices, unrestricted combination of bets, and linear valuation. It converges with the representation theorem as an operational consequence of the same probability structure.

Proposition 20.4 (Chance–credence bridge). *Let C be a $[0, 1]$ -valued chance variable for a binary event H . Suppose the model is calibrated in the regular-conditional sense,*

$$\mathbb{E}[1_H \mid C] = C \quad \text{almost surely,}$$

and let E be admissible relative to C , meaning $H \perp E \mid C$. Then

$$\mathbb{E}[1_H \mid C, E] = C \quad \text{almost surely.}$$

In particular, on any atom with $P(C = p, E) > 0$, one has $P(H \mid C = p, E) = p$.

Proof. Conditional independence gives $\mathbb{E}[1_H \mid C, E] = \mathbb{E}[1_H \mid C]$ almost surely, and calibration identifies the latter with C . The atomic statement is the corresponding pointwise conditional-probability case. $\square$

This is the conditional core of the Principal Principle (Lewis, 1980). The chance model and admissibility relation enter as grounds; MU then fixes the chance–credence consequence.

Proposition 20.5 (Reflection under stable updating). *Let $(\mathcal{F}_t)$ be a filtration and let $Q_t = P(H \mid \mathcal{F}_t)$ be the posterior generated by one probability model. Then (Q_t) is a martingale:*

$$\mathbb{E}[Q_t \mid \mathcal{F}_s] = Q_s, \quad s < t.$$

Moreover,

$$\mathbb{E}[1_H \mid \sigma(Q_t)] = Q_t \quad \text{almost surely.}$$

Consequently, whenever $\{Q_t = q\}$ is an atom of positive probability,

$$P(H \mid Q_t = q) = q.$$

Proof. The tower property gives

$$\mathbb{E}[Q_t \mid \mathcal{F}_s] = \mathbb{E}[\mathbb{E}[1_H \mid \mathcal{F}_t] \mid \mathcal{F}_s] = \mathbb{E}[1_H \mid \mathcal{F}_s] = Q_s.$$

Since Q_t is itself $\mathcal{F}_t$ -measurable,

$$\mathbb{E}[1_H \mid \sigma(Q_t)] = \mathbb{E}[\mathbb{E}[1_H \mid \mathcal{F}_t] \mid \sigma(Q_t)] = \mathbb{E}[Q_t \mid \sigma(Q_t)] = Q_t.$$

Restriction to a positive-probability atom $\{Q_t = q\}$ gives the final statement. $\square$

Reflection therefore holds when future credence is a posterior in the same model. Anticipated forgetting, model change, strategic distortion, or future irrationality changes the problem and can break the principle (van Fraassen, 1984).

20.5 Degrees of *doxa*

Graded belief admits continuous variation, while the practical vocabulary attached to that variation is context-sensitive. Terms such as *tentative*, *confident*, and *near-certain* denote context-indexed reporting ranges. Their rational use depends on the decision, loss function, channel, and institutional standard in the problem.

Verbal confidence is context-indexed. Probability or a credal interval is part of the belief state. A verbal confidence label is an additional reporting map. Unless that map is specified by the application, no universal boundary between “tentative” and “confident” is supported.

A numerical state, a linguistic report, and a decision threshold are three distinct forms of answer whose connecting maps belong in the problem.

20.6 The Structure of Knowledge

MU is a logical theorem whose status belongs with deductive certainty. Empirical confidence occupies a separate graded scale.

Proposition 20.6 (The logical stability of MU). *Within every consequence structure satisfying the hypotheses of the MU Theorem, no later empirical constraint can overturn MU without changing that consequence structure or denying one of its stated hypotheses.*

Proof. The MU Theorem holds in every consequence structure satisfying its stated hypotheses. Empirical evidence may change the substantive background, hypothesis space, or channel model while leaving that logical result intact. A purported empirical argument for $\neg$ MU would be assessed by the method MU already generates. $\square$

A three-level hierarchy. It is useful to distinguish:

- (i) the **logical floor**: MU and the minimal consequence structure that instantiates it;
- (ii) **deductive certainty relative to a theory**: $T \vdash H$;
- (iii) **empirical knowledge**: true belief reached through an internally supported and reliable route under a stated relevant-alternatives class.

Only the first is foundational here. The second is theory-relative; the third is fallible and connected to truth through a channel.

Definition 20.2 (Knowledge profile). For a proposition H , write

$$\mathcal{E}(H) = (T, W, R),$$

where T records truth, W internal support from the stated grounds, and R the reliability or safety of the route across the relevant class of cases.

Knowledge has three independent coordinates. Truth, internal support, and route integrity are logically distinct. A lucky guess may have $(T, W, R) = (1, 0, 0)$; an internally impeccable subject in a systematically deceptive environment may have $(0, 1, 0)$; a Gettier case may have $(1, 1, 0)$; a reliable but internally unsupported clairvoyant case may have $(1, 0, 1)$; and an ordinary successful knowledge case may have $(1, 1, 1)$. Hence no analysis based on only one of the three coordinates captures the contrasts that motivate the classical cases.

The search for one scalar essence gives way to a knowledge profile. Ordinary knowledge attributions can impose context-sensitive thresholds or credit conditions on that profile without changing its coordinates. The route-integrity proposal used in this paper takes truth, internal support, and adequate route reliability as jointly required for empirical knowledge. Figure 5 places the cases on the profile.

Remark 20.1 (Compatibility with knowledge-first epistemology). A knowledge-first theorist may take knowledge as primitive and use the profile to study its inferential and channel conditions. The formal result needed here is narrower than a lexical reduction: the three coordinates can vary independently, and no internal completion rule can manufacture external connection to truth.




Figure 5: The knowledge profile $\mathcal{E}(H) = (T, W, R)$. The classical cases occupy distinct vertices, so no single coordinate captures the contrasts; empirical knowledge is the $(1, 1, 1)$ vertex.

20.7 Certainty, Doubt, and Action

Definition 20.3 (Certainty). A claim is **deductively certain relative to B** when $B \vdash H$. An agent is **psychologically certain** when it assigns maximal confidence or reports no doubt. The two notions coincide only when the psychological state matches the deductive relation.

Certainty is local to a stated background. Deductive certainty is stable under additions that preserve the premises and consequence relation, but may disappear when premises are retracted, a model class changes, or the consequence relation itself changes. Psychological certainty has no comparable logical guarantee.

Definition 20.4 (Doubt). Doubt is a non-degenerate state of belief. It may occur within one probability distribution, when live alternatives receive non-extreme support, or across a family of distributions, representations, or channel states when the grounds preserve several points.

When suspension is required. If several epistemically inequivalent states remain, the whole set is the answer until a selection rule enters the problem. Given utilities, Theorem 14.4 removes every credally dominated action and returns the undominated decision set. A unique point requires either one survivor or a separately specified ambiguity rule, reference, policy, or institutional protocol.

Remark 20.2 (Terminology). Throughout the remainder, *deductive certainty* names entailment, *graded belief* names probabilistic or imprecise support, and *empirical knowledge* names the route-integrity account. Ordinary uses of “knows” move among these standards with context; the formal distinctions remain fixed.

21 Induction, Confirmation, and Explanation

21.1 Induction: Grounding and Projectibility

The classical problem asks how observed cases support unobserved ones. Deduction preserves the observed record; projective structure carries the inference forward.

21.1.1 The demand for external justification

Inference supplies its own standpoint. Every reason offered for accepting or rejecting a projective rule is itself an inference governed by a given relation of support. Every proposed tribunal is itself an inference, so inferential assessment remains reflexively inside the practice. Projective rules then earn standing through their explicit models, channels, and performance.

Remark 21.1 (Hume locates the projective task). The future-facing result follows from the projective structure, support, sampling, and distinguishability stated in the problem (Hume, 1739; Carnap, 1950). This account separates two questions: inference uses a relation of support through which evidence bears on conclusions; the continued world-tracking success of a particular projective relation is an empirical matter. Realisability, stability, support, sampling, and distinguishability carry that empirical burden. The regress about inference as such ends, and the scientific problem of projectibility receives its explicit conditions. The constitutive half of the diagnosis descends from Strawson's observation that asking whether induction is rational resembles asking whether the law is legal (Strawson, 1952); the convergence conditions carry what his reassurance could not.

Theorem 21.1 (Conditional convergence of inductive learning). *In the finite-alphabet realisable domain of the Convergence Theorem, prior support, stable sampling, and distinguishability imply posterior concentration on every neighbourhood of the true observational distribution. If realisability fails, updating can compare only the candidates present in the model class and may at best concentrate on a predictively optimal region under additional misspecified-learning conditions.*

Proof. The first statement is the Convergence Theorem of §17. The second records its boundary: updating compares the theories present in the candidate class, while an enlarged class brings newly generated theories into comparison. $\square$

The projective task. With the standpoint secured, the choice and empirical success of projective rules are assessed through representation, hypothesis space, channel structure, support, sampling, and distinguishability.

21.2 Confirmation Paradoxes: Ravens and Old Evidence

21.2.1 The ravens paradox

Let H be the hypothesis that all ravens are black and let E be any observation. Bayes' rule gives

$$\frac{P(H | E)}{P(\neg H | E)} = \frac{P(E | H)}{P(E | \neg H)} \frac{P(H)}{P(\neg H)}.$$

Proposition 21.2 (Likelihood determines confirmation). *Evidence E confirms H exactly when $P(E | H) > P(E | \neg H)$. The likelihood ratio is fixed by the sampling and population model. A nonblack nonraven can therefore confirm, disconfirm, or leave H unchanged across different models.*

Proof. The posterior-odds identity shows that the direction of confirmation is determined by the likelihood ratio. Logical equivalence fixes which worlds satisfy H ; the sampling model fixes how observations arise within those worlds. If E has the same probability under H and $\neg H$, its Bayes factor is one. If the population model makes E more likely under H , it confirms H , often only

minutely when ravens are rare. The opposite inequality disconfirms. Thus the paradox arises from combining a logical equivalence with an unstated confirmation model (Hempel, 1945). $\square$

Proposition 21.3 (Confirmation follows the specified hypothesis). *There are probability models in which evidence E confirms H while disconfirming $H \wedge J$.*

Proof. Let

$$P(H \wedge J) = 0.45, \quad P(H \wedge \neg J) = 0.05, \quad P(\neg H \wedge J) = 0.05, \quad P(\neg H \wedge \neg J) = 0.45,$$

and set

$$P(E \mid H \wedge J) = 0.05, \quad P(E \mid H \wedge \neg J) = 1, \quad P(E \mid \neg H \wedge J) = P(E \mid \neg H \wedge \neg J) = 0.01.$$

Then $P(E \mid H) = 0.145 > P(E \mid \neg H) = 0.01$, so E confirms H . But $P(E) = 0.0775$ and hence

$$P(H \wedge J \mid E) = \frac{0.45 \cdot 0.05}{0.0775} \approx 0.290 < 0.45 = P(H \wedge J).$$

Thus E disconfirms the conjunction. Confirmation tracks the specified likelihood relation and the named hypothesis. $\square$

21.2.2 Old evidence

Suppose evidence E is already known, so the agent's current state assigns $P(E) = 1$. A newly proposed theory H can nevertheless gain support when the agent learns a new model relation M connecting H to E .

Proposition 21.4 (Old evidence is new relational information). *Let P_0 be the agent's current state after learning E , so $P_0(E) = 1$. Learning a new relation M between a theory and that evidence updates the current odds by*

$$\frac{P_1(H)}{P_1(H')} = \frac{P_0(M \mid H, E)}{P_0(M \mid H', E)} \frac{P_0(H)}{P_0(H')},$$

whenever the conditionals are defined. The update acts on the newly learned relation M while the old evidence E remains inside the current state.

Proof. The current probabilities $P_0(H)$ and $P_0(H')$ already incorporate the old evidence. Conditioning that state on the newly learned relation M gives

$$\frac{P_0(H \mid M)}{P_0(H' \mid M)} = \frac{P_0(M \mid H, E)}{P_0(M \mid H', E)} \frac{P_0(H)}{P_0(H')},$$

because E has probability one under P_0 . The new information is the theory–evidence relation represented by M ; the old observation retains its existing contribution (Glymour, 1980; Garber, 1983). $\square$

The two puzzles therefore share one lesson: confirmation depends on a model of how evidence bears on hypotheses. Logic alone neither supplies nor replaces that model.

21.3 Goodman, No Free Lunch, and Projectibility

The green–grue construction demonstrates that data support projection only through a represented language (Goodman, 1955). Complexity is representation-relative: a predicate simple in one language may be elaborate in another.

21.3.1 Representation before projection

Proposition 21.5 (Projectibility depends on representation). *A projective inference becomes well-defined through its representation language, measurement apparatus, invariances, and candidate transformations. Several permitted representations produce their complete family; an explicit complexity measure or apparatus structure selects among them.*

Proof. The predicates used to state the data and future claims belong to the input language. Changing the primitive vocabulary changes description lengths and may change which regularities appear simple. By the selection theorem, a preference among such representations must be supported by apparatus, symmetry, coding, or an explicit model-selection criterion. Otherwise the preference is additional structure. $\square$

Definition 21.1 (Projectibility conditions). For a fixed representation, the projectibility conditions record:

- (i) empirical fit under the channel model;
- (ii) stability under the chosen transformations;
- (iii) complexity relative to the chosen coding or reference structure;
- (iv) predictive performance on held-out or future observations.

The riddle becomes a precise question about representation. Occam ordering is legitimate inside a fixed representation. The evidence determines an ordering inside the represented problem, while the higher-order comparison ranges over the surviving representations.

The No Free Lunch theorems give the computational counterpart: in their finite setting, uniform averaging over an unrestricted problem class equalises the performance of every learner or optimiser (Wolpert and Macready, 1997). Successful generalisation therefore requires problem-relative structure, and successful induction makes its bias explicit as exactly that structure. At deployment scale the riddle recurs as distribution shift: a rule fitted to cases examined before some boundary is projected onto cases beyond it, and the choice among extrapolations again lives in representation and apparatus rather than in the record.

21.4 Abduction, Underdetermination, and Explanation

Inference to the best explanation contains at least two stages. Hypothesis generation proposes a space of candidates; evaluation compares them. The method governs the second stage once the first has supplied a hypothesis space.

Reconstruction of abduction. Within a model class, explanatory considerations enter through distinct parts of the problem: Empirical adequacy enters through the likelihood and channel model; parsimony through the reference or complexity case; unification through shared-parameter or common-cause structure; novel prediction through the likelihood on held-out or future observations; and manipulability through the intervention or causal model. Deutsch’s hard-to-vary criterion names the same discipline informally: an explanation is good when its stated

structure leaves no idle freedom, so that variation breaks contact with the constraints (Deutsch, 2011). The specified components jointly determine comparative explanatory standing. Hypothesis generation and representation choice remain separate tasks.

The reconstruction gives each explanatory virtue a distinct formal place and makes comparative evaluation explicit. It shows where each virtue enters and which part of the judgement remains creative or representation-dependent.

Theorem 21.6 (Observational equivalence). *Let M_1 and M_2 be model bundles, including their auxiliaries and apparatus assumptions. Let an experimental policy choose each experiment a_t from the available class $\mathcal{A}$ as a function of the preceding history h_{t-1} . Suppose that for every permitted history and action,*

$$K(\cdot \mid h_{t-1}, a_t, M_1) = K(\cdot \mid h_{t-1}, a_t, M_2).$$

Then no finite history generated by any such adaptive policy changes the prior odds between M_1 and M_2 .

Proof. For a realised history $(a_t, s_t)_{t=1}^n$, the likelihood ratio factors conditionally as

$$\prod_{t=1}^n \frac{K(s_t \mid h_{t-1}, a_t, M_1)}{K(s_t \mid h_{t-1}, a_t, M_2)} = 1.$$

The policy contributes the same action-selection probability under both models because it is a function of the common observed history. Bayes' rule therefore leaves posterior odds equal to prior odds. Distinguishing the bundles requires an experiment or intervention outside the equivalence class, or additional structural information. $\square$

Duhem–Quine underdetermination appears here in its cleanest form. Data test a model together with auxiliaries, and empirically equivalent bundles remain tied under every test on which they agree (Duhem, 1906; Quine, 1951). An explanation may still be preferred for an explicit simplicity, causal, or unifying reason, and that reason appears explicitly as simplicity, causal, or unifying structure in the comparison.

22 Knowledge, Testimony, and Scepticism

22.1 Gettier, Closure, Context, and Stakes

Gettier cases (Gettier, 1963) show that justified true belief can be true through the wrong connection. A right answer can still be bad inference. The failure is structural: internal support and connection to truth answer different questions.

Internal support and route integrity are independent. A true and internally supported belief becomes empirical knowledge when its channel remains robust across the relevant nearby cases. Route integrity supplies the external coordinate alongside internal support.

Definition 22.1 (Route robustness). A belief-forming route is **robust relative to a chosen world class $\mathcal{W}$** when its kernel continues to discriminate the truth-relevant alternatives throughout $\mathcal{W}$ to the standard required by the application. The choice of $\mathcal{W}$ and the standard are part of the problem.

The knowledge profile completes internal support with route integrity. Internal support and connection to truth answer separate questions. The claim that internal justification and truth are sufficient for knowledge dissolves; empirical knowledge adds a condition protecting the route against epistemic luck.

Knowledge through a robust route. A belief counts as empirical knowledge in this reconstruction when:

- (i) it is true;
- (ii) it is internally supported by the problem;
- (iii) its route is robust in the chosen relevant-alternatives class;
- (iv) the success is attributable to the relevant competence or channel.

Remark 22.1 (Safety, sensitivity, reliability, and ability). Safety, sensitivity, average reliability, and ability impose distinct conditions (Lewis, 1996; Goldman, 1979; Williamson, 2000). The channel formalism provides a common space for stating and comparing them while preserving their distinct conditions.

The two-coordinate analysis also handles standard pressure cases for internalism and externalism. A subject in a perfectly deceptive demon world can be internally indistinguishable from a normally situated counterpart while lacking the same external success. Conversely, a reliable clairvoyant who has no accessible reason to trust the faculty can possess an externally successful route without the internal support required by the knowledge proposal used here (BonJour, 1980). The cases show that neither coordinate replaces the other.

The generality problem for reliability. Suppose one token of belief formation belongs to process descriptions $T_1, \dots, T_m$ with different reliabilities $r_i = P(\text{true output} \mid T_i)$. The problem's relevance rule selects or weights the process descriptions; until then the reliability assessment is the family $\{r_1, \dots, r_m\}$.

Reliabilism inherits the reference-class problem. A channel model resolves it through conditioning variables and a level of description fixed before the verdict.

Deductive closure and route preservation. Suppose $B \vdash H$ and the agent is internally justified in accepting B . The entailment transfers internal support from B to H . Route robustness transfers when the supporting channel also discriminates every alternative encoded by H . Closure of empirical knowledge therefore requires preservation of the relevant channel condition.

Closure remains valid for consequence, while route conditions transfer exactly when the entailed reformulation preserves them.

Context changes the knowledge problem. Let $\mathcal{W}_c$ be the class of alternatives relevant in context c , and let ϵ_c be the required reliability or safety standard. A route may count as knowledge-producing relative to $(\mathcal{W}_c, \epsilon_c)$ and fail relative to a more demanding $(\mathcal{W}_{c'}, \epsilon_{c'})$. The attributions are consistent because the problems differ.

Contextualist and relevant-alternatives insights fit here (DeRose, 1995; Lewis, 1996). Whether ordinary language builds context into the meaning of “knows” is a semantic question. The formal result shows how context can alter the relevant alternatives and standard without changing the first-order evidence.

Stakes change action before they change evidence. Fix an epistemic state P and two decision problems with different loss functions L_1 and L_2 . The optimal actions may differ

although P is unchanged. Practical stakes therefore supply reasons for different actions without automatically supplying evidence for different credences.

A theory of pragmatic encroachment may choose to make knowledge attributions sensitive to actionability or stakes (Hawthorne, 2004; Fantl and McGrath, 2009). The formal distinction remains: evidence determines the epistemic state; losses and utilities determine what to do with it.

22.2 Testimony, Easy Knowledge, and Other Minds

Bootstrapping cases use a source's outputs to assess that source's reliability. The channel model identifies the additional calibration anchor that turns this circle into evidence.

Proposition 22.1 (Channel calibration requires an external anchor). *Signals identify a channel's truth-reliability through independently verified outcomes or sufficiently restrictive structural assumptions. Repeated self-agreement supplies consistency data; truth calibration comes from an independent anchor or identifiable structural model.*

Proof. Proposition 16.3 gives the witness: in a binary symmetric channel, many truth-prevalence and accuracy pairs induce the same signal law, so observed frequencies identify only their equivalence class. Verified outcomes, an independently calibrated channel, or identifying structural restrictions select the accuracy. $\square$

Easy-knowledge and bootstrapping cases fit the same pattern (Cohen, 2002; Vogel, 2000). A source contributes to its own calibration through outcomes anchored by independent controls, verified labels, or identifiable structural constraints.

The result also separates *prima facie support* from *self-certification*. A perceptual experience may give immediate defeasible support for "there is a red object here" without a prior argument that vision is reliable (Pryor, 2000). Second-order calibration of vision then proceeds through independent controls or identifiable structure. First-order entitlement and second-order calibration are different questions (Wright, 2004).

Testimonial support. Let S be a source report and H its content. The evidential force of the report is the Bayes factor

$$\frac{P(S \mid H, \eta)}{P(S \mid \neg H, \eta)},$$

where η records competence, honesty, dependence, incentives, and selection. Testimony supports H when this ratio exceeds one, and the source model determines the direction and magnitude of evidential weight.

Expert testimony may be rationally weighty without the hearer reproducing the proof, because division of cognitive labour is itself a calibrated channel structure (Hardwig, 1985). Dependence among experts matters: ten reports copied from one source form one dependent likelihood structure.

Other minds and phenomenal equivalence. Behaviour, testimony, neural evidence, and interaction can support claims about properties that are stable across the well-supported empirical equivalence class: agency, memory, responsiveness, planning, and other capacities reflected in the available channels. If two theories of consciousness induce the same law for every accessible signal and intervention, Theorem 25.1 shows that no evidence from that class can select between

them. Ordinary knowledge of other minds is therefore compatible with a terminal empirical boundary on phenomenal metaphysics.

Empirical equivalence leaves metaphysical identity open. A new discriminating channel refines the empirical answer; until then the equivalence class is the full result.

22.3 Peer Disagreement, Higher-Order Evidence, and Akrasia

A disagreement report is evidence about another route and, often, about one's own reliability. Its force depends on what the peer knows, how the report was generated, and how independent the routes are.

Proposition 22.2 (Disagreement as evidence). *Let D be a peer's report about H . Conditional on the current evidence E and channel state η ,*

$$\frac{P(H \mid E, D, \eta)}{P(\neg H \mid E, D, \eta)} = \frac{P(D \mid H, E, \eta)}{P(D \mid \neg H, E, \eta)} \frac{P(H \mid E, \eta)}{P(\neg H \mid E, \eta)}.$$

An independent, comparably reliable peer can therefore supply substantial higher-order evidence. A report determined by the same evidence and method may add little beyond revealing a disagreement to be explained.

Proof. Bayes' rule now applies after conditioning on E and η . The likelihood ratio measures the information in the peer's report beyond the evidence already shared. If the report is conditionally independent and discriminating, the ratio can be substantial; if it is fixed by the shared evidence, its likelihood factor is one. $\square$

Conciliation and steadfastness follow from the peer channel, shared evidence, and dependence structure in each case (Kelly, 2005; Christensen, 2007; Elga, 2007). If agents have a common prior and their posteriors become common knowledge, Aumann's theorem forces agreement; absent those conditions, rational disagreement can locate a difference in evidence, representation, reliability, or dependence.

Epistemic akrasia. An agent is epistemically akratic when its first-order attitude conflicts with an endorsed higher-order judgement about what its total evidence supports. The conflict should be represented as disagreement between two channels or levels of the model. If the higher-order judgement is treated as certain and authoritative, retaining the conflicting first-order state violates the retention update. If the higher-order judgement is itself uncertain, the correct output can remain a joint distribution over the claim and the reliability of both assessments.

22.4 Global and Local Doubt

Global and local scepticism. The MU Theorem resolves universal scepticism about inference. Cartesian, brain-in-a-vat, perceptual, testimonial, and scientific doubts then become coherent local questions about a channel, model class, or connection to the external world.

Remark 22.2 (Sceptical hypotheses receive probability through a common model). A demon, simulation, or brain-in-a-vat scenario receives probability through a common hypothesis space, reference or prior structure, and signal laws connecting each hypothesis to the evidence (Stroud,

1968; Wittgenstein, 1969; Moore, 1939). Those choices make complexity and likelihood comparisons explicit. MU turns local scepticism into a model-and-channel problem and returns exactly what that problem supports.

A perceptual route may robustly support “I have hands” across ordinary nearby cases while failing to distinguish ordinary worlds from perfectly adversarial simulations. The larger anti-sceptical claim asks more of the route. Its probability follows once a model class and reference structure define the comparison.

The stable result. MU resolves global scepticism about inference. Local scepticism becomes calibrated inquiry into channels, representations, and models. Moorean common sense rests on robustness across the alternatives relevant to the ordinary claim; metaphysical alternatives belong to the wider problem that introduces them.

23 What the Arguments Establish

23.1 A map of the results

The arguments above do the work. This table gathers their results in one place.

Problem	What follows
Grounding regress and the Münchhausen Trilemma	Resolved at the level claimed. MU has a finite direct proof. The proof begins from reflexivity and a consistent premise, terminates by existential introduction, and precedes the later analysis of inference.
Problem of the criterion	Resolved by role separation. The inferential form comes from what it is for an answer to follow from grounds; premises, representations, and channels supply the material assessed by that form.
Rule-following and Carroll’s regress	Resolved in two layers. Rule application is irreducible to an added premise, and finite behaviour supports a family of continuations. Representation, norm, or selection structure determines a unique continuation.
Hypothesis generation and representation choice	Resolved as a higher-order search problem. Generators supply candidates, evaluation assigns standing, and search order governs bounded access to the resulting family.
The Given and immediate experience	Reconstructed. Experience supplies immediate defeasible support through a standing channel and model that connect signal to claim.
A priori support and analyticity	Located internally. Meanings, axioms, and rules determine a priori support within a formal problem; empirical inquiry assesses the framework’s relation to the world.
Indifference, reference classes, and self-location	Resolved by answer shape. Real symmetry determines a point. Several measures, reference classes, or observation protocols produce their full family. Centered states and sampling rules complete self-locating cases.

Problem	What follows
Uniqueness and permissivism	Reconciled. One fully defined problem has one full answer, and that answer may contain several permissible credal states. Uniqueness governs the family; permissivism describes its members.
Lottery, preface, and threshold puzzles	Resolved by level separation. Probability, entailment, knowledge, conjunction, assertion, and action each occupy their own level and receive the corresponding rule.
Dutch books, objective chance, and reflection	Resolved conditionally. Fair-price protocols generate Dutch-book coherence; calibrated chance with admissible evidence guides credence; stable Bayesian updating generates reflection.
Moore's paradox and epistemic akrasia	Reconstructed through self-representation. The Moorean proposition may be true while its sincere assertion conflicts with the speaker's represented belief. Higher-order evidence enters as evidence about one's own channel.
Fitch's knowability paradox	Resolved as an impossibility result. Under factivity and the stated modal closure principles, universal knowability collapses to universal knowledge. A weaker knowability claim accommodates unknown truths.
Induction	Resolved in two stages. Inference supplies its own standpoint. Projective rules are then compared through explicit models, support, sampling, and distinguishability, yielding a time-indexed point or family.
Ravens and old evidence	Resolved conditionally. Likelihood determines confirmation. Old evidence confirms a new theory when the newly learned item is the theory–evidence connection or model structure.
Goodman's new riddle and No Free Lunch	Resolved at the epistemic level. Representation and problem distribution supply the projective structure; evidence compares the cases that make different predictions.
Abduction and scientific underdetermination	Reconstructed. Fit, parsimony, unification, novel prediction, and intervention occupy distinct parts of the problem. New tests or explicit structure separate observationally equivalent models.
Gettier, closure, context, and stakes	Resolved through the knowledge profile. Internal support and route integrity are independent coordinates that jointly support empirical knowledge. Closure preserves knowledge when route conditions transfer. Context changes relevant alternatives; stakes change action thresholds.
Easy knowledge, testimony, and other minds	Reconstructed through channels. External anchors calibrate sources. Testimony earns weight through calibration and dependence structure, and ordinary knowledge of other minds arises through the same channel framework.

Problem	What follows
Peer disagreement and higher-order evidence	Resolved conditionally. An independent reliable peer contributes a likelihood factor according to shared evidence and dependence. Aumann agreement follows from a common prior and common knowledge.
Global and local scepticism	Resolved by scope. MU settles the possibility of inference. Doubt about a channel, model, or world-connection becomes a coherent local problem, and a common model with reference structure assigns the sceptical hypothesis its probability.
Meno and the swamp-ing problem	Substantially explained. Information has nonnegative expected instrumental value, and knowledge adds a stable, reusable, attributable route to truth beyond lucky arrival. Practical value supplies the remaining evaluative layer.
The epistemic is-ought and doxastic voluntarism	Resolved constitutively. The standards of inference bind outputs offered as inference. Practical reasons govern assertion and action, while evidence governs support.
Practical choice under imprecise belief	Resolved as a set-valued decision problem. Credal dominance removes every dominated action. One survivor gives a point; several survivors give the undominated set until value or protocol refines it.
The value of information and the duty to inquire	Resolved conditionally. Cost-free information has nonnegative expected instrumental value. Under a specified expected-utility objective, inquiry is required exactly when its expected value exceeds its cost.
Duhem–Quine underdetermination	Proved within an experiment class. Equal signal laws preserve posterior odds across every available test. New interventions or assumptions enlarge the discriminating class.
Scientific realism, pessimistic meta-induction, and unconceived alternatives	Bounded by the empirical quotient. Data identify models up to observational equivalence. The class-invariant structure is the maximal empirical content; model enlargement introduces unconceived alternatives.
Theory-ladenness and incommensurability	Resolved through representation and translation. Observation is a signal interpreted through a model and apparatus. A bridge preserving shared predictions makes rival frameworks comparable, and constructing that bridge becomes the next explicit problem.
Causation and correlation	Resolved as an identification problem. Interventions, temporal structure, invariance, or further assumptions orient causal direction beyond the observational distribution.
Logical omniscience and bounded rationality	Separated. Deductive closure is the full normative answer, while a finite agent reaches a computably accessible subset at each time. More computation reveals further members of the closure.

Problem	What follows
Expertise, testimony, and epistemic injustice	Reconstructed through calibrated social channels. Credibility tracks predictive relevance and dependence. Social premises evaluate systematic distortions produced by identity power and institutional structure.
Machine understanding and opacity	Resolved at the epistemic level and bounded at the phenomenal level. Reliability, route integrity, and auditability determine epistemic standing. Consciousness theories identical on every accessible signal form one empirical equivalence class whose remaining difference is metaphysical.

23.2 Five Forms of Further Work

The arguments fall into three broad patterns. Direct resolutions establish inference as its own standpoint, rule application as irreducible, thresholds as context-indexed, and channel calibration as externally anchored. Structural relocations place induction in projectibility and sampling, Goodman in representation, local scepticism in models and channels, and disagreement in evidence, dependence, and reliability. Positive reconstructions supply accounts of abduction, empirical knowledge, testimony, scientific inquiry, and the norms internal to inference.

The remaining work now has a clear form: generate candidates, design discriminating experiments, extend computational access, state practical values, and classify distinctions preserved by current empirical equivalence.

Five forms of remaining work. Across the problems treated here, what remains falls into five forms:

- (i) **candidate generation:** inquiry enlarges the problem with a representation, hypothesis, or rule;
- (ii) **empirical discrimination:** new channels or experiments separate the surviving cases;
- (iii) **computational access:** further computation reaches more of the fixed full answer;
- (iv) **practical value:** value, loss, cost, or an ambiguity rule selects among surviving belief or decision options;
- (v) **empirical or metaphysical equivalence:** inquiry records the equivalence class until a new discriminating structure appears.

The corresponding answers are a higher-order family, a probability or credal state on an empirical quotient, an approximation contract, an undominated decision set or value-relative completion, and a proof of non-identifiability. The earlier theorems show where each case belongs: candidate generation in the higher-order problem, empirical discrimination in the quotient theorem, bounded access in the computational analysis, practical underdetermination in the decision results, and empirical equivalence in the channel model.

The scientific, creative, practical, and metaphysical work remains to be done. The gain is knowing what kind of work it is: generate a candidate, run a discriminating experiment, compute further, state the value at issue, or mark the boundary. Each proceeds under the same rule of inference.

Part V

Science, Computation, and Society

We can only see a short distance ahead, but we can see plenty there that needs to be done.

Alan Turing, 1950

The same distinctions now meet the places where inquiry is actually done: experiments, finite computation, social dependence, and artificial systems.

24 Criticism, Falsification, and Confirmation

A scientific claim gains evidential standing through tests whose possible signals separate the hypotheses at issue.

Severe testing. Let $K(s \mid h, \eta)$ be the experimental channel. A test is severe for alternatives h_0, h_1 to the extent that the resulting signal laws are distinguishable and the apparatus state η is sufficiently controlled. A failed severe prediction can strongly reduce support for a theory; a passed severe prediction can increase support when it was appreciably less expected under its rivals.

Falsification and confirmation are complementary regions of one likelihood comparison. A failed prediction often produces a large likelihood penalty. A successful prediction confirms in proportion to the advantage it held over its competitors.

A failed prediction confronts a specified bundle of theory, auxiliaries, apparatus, and background conditions. Revision localises the change supported by the signal and records every alteration of the problem.

25 Underdetermination, Realism, and Theory Change

The observational-equivalence theorem applies directly to scientific realism. Evidence gathered within an experiment class can select only among models whose signal laws differ in that class.

Theorem 25.1 (Empirical equivalence quotient). *Let $\mathcal{A}$ be a class of possibly adaptive experimental policies, and let P_M^a be the probability law on complete observation histories induced by model M under policy a . Define*

$$M_1 \equiv_{\mathcal{A}} M_2 \iff P_{M_1}^a = P_{M_2}^a \text{ for every } a \in \mathcal{A}.$$

If $M_1 \equiv_{\mathcal{A}} M_2$ and both have positive prior probability, then every observation history with positive shared likelihood leaves their posterior odds unchanged:

$$\frac{P(M_1 \mid h, a)}{P(M_2 \mid h, a)} = \frac{P(M_1)}{P(M_2)}.$$

Thus evidence from $\mathcal{A}$ identifies models at most up to the quotient $\mathcal{M}/\equiv_{\mathcal{A}}$.

Proof. Empirical equivalence gives $P(h \mid M_1, a) = P(h \mid M_2, a)$ for every allowed policy and history. Bayes' rule cancels the common likelihood in the posterior-odds ratio. Adaptivity changes the policy as a function of earlier observations while preserving equality of the resulting complete-history laws. $\square$

Corollary 25.2 (Maximal empirical content). *Every property learned from data generated by $\mathcal{A}$ is constant on each $\equiv_{\mathcal{A}}$ -class. Thus empirically decidable properties are exactly the class-invariant properties. The maximal empirical answer is therefore a probability or credal state on the quotient together with the structure invariant within its classes.*

Proof. Any data-based decision rule has the same distribution under empirically equivalent models. If a property took different values within one class, no sequence of such rules could converge with probability one to the correct, different value under both models, because the rules would have identical distributions in the two cases. $\square$

The consequence is sharp (van Fraassen, 1980). Predictive success supports the structure responsible for successful discrimination. New experiments or further metaphysical premises address differences inside an empirical-equivalence class.

Epistemic structural realism. When changing theories preserve mathematical or causal structure across domains of successful prediction, and that structure is invariant under the relevant empirical equivalence classes, it is the strongest realist content supported by those data. Ontological claims that vary within the same classes outrun the empirical result (Worrall, 1989).

The pessimistic meta-induction treats the historical failure of successful theories as higher-order evidence about the reliability of a theory-selection channel (Laudan, 1981). It calibrates confidence in particular realist claims while preserving the evidential value of predictive and structural success.

Proposition 25.3 (Model enlargement gives new alternatives standing). *Let $\mathcal{M}$ be the considered model class. An alternative M_* outside $\mathcal{M}$ receives no posterior probability until the model class is enlarged to include it. Equivalently, in an enlarged class, assigning $P_0(M_*) = 0$ makes $P_n(M_*) = 0$ after every ordinary finite Bayesian update. High posterior concentration on one member of $\mathcal{M}$ therefore establishes comparative success within the considered class unless realisability or class adequacy is separately supported.*

Proof. Updating within $\mathcal{M}$ redistributes probability only among members of $\mathcal{M}$; a newly generated alternative enters through an enlarged candidate class. In an enlarged class, Bayes' rule multiplies prior mass by a likelihood and renormalises, so zero prior mass remains zero after every finite update. The same support-inheritance boundary appears in the Gibbs completion of §14. Evidence can favour a newly introduced alternative only after the problem has been enlarged and the alternative given nonzero standing. $\square$

That is the force of the problem of unconceived alternatives (Stanford, 2006). Posterior concentration establishes comparative success inside a model class; higher-order search addresses the class's scope.

Theory-ladenness and translation. An observation used in inquiry is a signal together with a representation map that places it in a model. If rival theories use different representation maps,

comparison requires a translation preserving enough shared predictions, measurements, or interventions to define a common problem. A supplied bridge turns the frameworks into competitors inside one well-defined comparison; an absent bridge remains a higher-order translation problem.

Kuhnian incommensurability therefore survives without becoming a blank cheque for relativism: it identifies the work needed to build a bridge between representations, and common empirical consequences and retained experimental successes provide the places from which such comparison can begin (Kuhn, 1962; Longino, 1990).

26 Bounded Rationality

An exact theory gives bounded reasoning a target. Approximation becomes meaningful once the reference, objective, cost of departure, and resource budget are clear.

A claim of approximation requires a target. When a bounded procedure is presented as approximating an exact inferential answer, the claim must state:

- (i) the exact problem and target answer;
- (ii) the computational resource budget;
- (iii) an error, regret, calibration, or convergence guarantee;
- (iv) the conditions under which the guarantee holds.

A practical procedure may also be evaluated directly by a loss, calibration test, or empirical benchmark. A target and performance contract make the approximation assessable.

Remark 26.1 (Computation reveals consequences of represented information). Additional computation searches a defined space more thoroughly, reduces approximation error, and exposes consequences that were hard to find. New information or a new selector changes an under-determined problem; computation improves access to what follows from the grounds already present.

Closure under bounded access. Let $\text{Cn}(T)$ be the complete deductive closure of a theory and let $A_t \subseteq \text{Cn}(T)$ be the consequences accessible to a bounded agent by time t . The normative answer remains $\text{Cn}(T)$, while the agent's current computational state is A_t . Further valid computation enlarges A_t within $\text{Cn}(T)$.

Ideal closure and feasible access are different questions. Finite agents therefore require algorithms, complexity bounds, and explicit uncertainty about consequences still awaiting derivation. Learning a proof or a theory–evidence relation can therefore be new information for the agent even when it was implicit in the ideal closure.

A variational approximation must state its family, objective, and approximation bound; a Monte Carlo method its target measure, transition kernel, and mixing or error control; a heuristic search its objective, allowed moves, stopping rule, and failure modes; a calibration model its event class, scoring rule, and sampling regime; and a decision policy its epistemic state, utilities or losses, and resource constraint. The free-energy principle presents perception and action as variational inference driven by a bound whose divergence term is the KL machinery above (Friston, 2010); in this classification it is a model family with an objective, and its epistemic standing is priced by the channel, reference, and bound it states, like any method on this list.

Exact and approximate outputs play different roles. A surrogate becomes rationally usable when its cost, error, and operating conditions are specified.

27 Social Epistemology

Social learning adds channels, dependence structures, protocols, and access conditions.

Theorem 27.1 (Agreement under common priors). *Agents with a common prior whose posteriors are common knowledge cannot agree to disagree about the probability of the same event.*

The scope is a common prior together with common knowledge of posteriors (Aumann, 1976).

Sources of rational disagreement. Persistent disagreement in epistemic states can arise from differences in:

- (i) evidence or access to evidence;
- (ii) reference families or prior cases;
- (iii) model and representation classes;
- (iv) source dependence and channel calibration;
- (v) computational access or approximation.

Differences in utilities, losses, or action thresholds can instead produce rational disagreement about what to do or report while the underlying credences agree. Before attributing either kind of disagreement to irrationality, the relevant inputs should be separated.

Definition 27.1 (Common knowledge). A proposition is common knowledge in a group when everyone knows it, everyone knows that everyone knows it, and so on through the usual transitive closure. Public availability alone is insufficient unless the protocol and its observability are themselves common knowledge.

Social convergence. Agents who share a realisable model, receive sufficiently informative evidence through appropriately controlled channels, preserve support, and update consistently converge in the regime of the Convergence Theorem. Correlated sources, strategic reporting, model misspecification, and asymmetric access can prevent convergence without violating the internal update rule.

Expertise and dependence. Suppose reports $S_1, \dots, S_n$ are conditionally independent given H and their source states. Their joint likelihood factorises. If the reports instead derive from one common source, common dataset, or coordinated incentive, multiplying them as independent double-counts evidence. Rational deference therefore depends on competence, independence, conflict of interest, and the hearer's access to calibration data.

Credibility and irrelevant social markers. Let Z be a social marker, X the source-relevant evidence, and H the truth of the reported claim. If

$$H \perp Z \mid X,$$

then changing a source's weight solely because of Z makes the answer depend on information with no predictive relevance in the problem. Such weighting fails the dependence test. When Z is predictive only because social structures affect access or treatment, the predictive fact and the moral evaluation of the structure must be kept distinct.

The equation captures one component of testimonial injustice (Fricker, 2007). The moral diagnosis goes further: institutions must ask why a marker became predictive, whose evidence was suppressed, and which remedies justice requires.

Remark 27.1 (Collective reference). In social applications a reference measure may encode a community’s inherited background of salience, normality, and admissibility. Explicit analysis of that reference makes institutional priors and exclusions visible.

Remark 27.2 (Epistemic injustice and scope). The channel formalism represents credibility deficits, testimonial suppression, unequal access, strategic silencing, and institutional dependence (Fricker, 2007). Social and moral premises determine their evaluation and remedy.

28 Artificial Intelligence and Causal Models

28.1 Program priors and universal induction

Exponential program priors. Fix a prefix universal machine U . The associated algorithmic semimeasure weights programs exponentially by length and sums over programs whose outputs have the observed prefix. This is a coherent complexity-sensitive reference case relative to U ; it is a complexity-sensitive prior relative to the machine choice.

Universality is indexed to a reference: the universal machine is part of the problem.

Solomonoff induction therefore illustrates (Solomonoff, 1964; Kolmogorov, 1965) both the strength and the discipline of the method. Once the coding machine is fixed, a powerful result follows. Changing the machine changes the case, and adversarial choices can produce pathological “universal” priors. The reference choice must remain explicit.

28.2 Hutter and AIXI: epistemic state versus policy

The components of AIXI. Fix a prefix universal machine U . AIXI contains at least four logically distinct inputs or operations:

- (i) the algorithmic environment semimeasure ξ_U , which supplies a predictive model over percept–reward histories;
- (ii) the reward channel or utility functional, which supplies the objective;
- (iii) a finite horizon or discount convention, which supplies temporal preference;
- (iv) expectimax optimisation, which supplies the policy rule.

Universal induction constrains the predictive model relative to U ; reward, horizon, and policy remain separately specified.

Prediction and policy are distinct (Hutter, 2005; Leike and Hutter, 2015). Changing the reference universal machine changes the belief case, while the ordinary invariance theorem bounds description-length changes asymptotically. Leike and Hutter’s bad-universal-prior constructions provide a decisive control: adversarial universal machines can produce radically different behaviour, so finite-scale AIXI remains reference-indexed. Skill-acquisition measures of intelligence make the same relativity explicit on the empirical side by defining ability against declared priors and experience (Chollet, 2019).

AIXI is uncomputable. Any approximation therefore adds a resource budget, a restricted model class, a search procedure, and a performance guarantee.

28.3 Machine understanding and opacity

The Chinese Room boundary. A system can be evaluated as a source of information through reliable answers, faithful routes, and calibrated uncertainty while the metaphysics of consciousness remains open. Phenomenal understanding and epistemic reliability are independent dimensions. The Chinese Room therefore marks the boundary between epistemic standing and philosophy of mind (Searle, 1980).

Opacity is an audit cost. Algorithmic opacity can reduce the ability to detect hidden dependence, contamination, brittle shortcuts, and failure under distribution shift. Accuracy and auditability are independent dimensions. Reliability, explanation, controllability, and accountability are distinct requirements (Burrell, 2016). A generative system that reports content its constraints never paid for instantiates smuggling at scale; the Non-Smuggling standard applies to machine reporters unchanged, and the failure called hallucination is its violation under thin constraints.

For machine authority, the distinction matters. A system may be accurate enough to deserve weight while still being too opaque for high-stakes delegation, because action also requires governance, contestability, and responsibility.

28.4 Alignment Depends on the Reference

At a fixed decision state, let $\pi_{\text{ref}}(a)$ be a reference policy, let $Q(a)$ be an action value, and let $\tau > 0$. The choice theorem gives

$$\pi^*(a) = \frac{\pi_{\text{ref}}(a)e^{Q(a)/\tau}}{\sum_b \pi_{\text{ref}}(b)e^{Q(b)/\tau}}.$$

Equivalently, π^* uniquely maximises

$$\mathbb{E}_\pi[Q] - \tau D_{KL}(\pi \| \pi_{\text{ref}}).$$

In a sequential soft-control problem, the formula applies statewise after the corresponding soft Bellman recursion has determined Q . The one-step tilt is one component of the sequential policy construction.

The same variational form appears in KL-regularised reinforcement learning and idealised RLHF objectives (Christiano et al., 2017; Haarnoja et al., 2018; Ziegler et al., 2019). In RLHF, Q is estimated from preference data through a reward-model channel. The Gibbs result fixes the policy tilt conditional on that learned value model. The components remain distinct:

Alignment is reference-based as well as value-relative. Reference support sets the available action class, and a misspecified or truncated reference can exclude relevant actions. The reward ranks that class, and reward misspecification and reward hacking live here. The cost parameter governs departure from the reference: too small a value permits over-optimisation, too large a value leaves the reference dominant. The optimisation protocol determines implementation, and the implemented system is assessed by its approximation guarantee to the ideal variational optimum. Approaches that treat the objective itself as uncertain place the value function inside the hypothesis space, so preference evidence arrives through a channel and the reward becomes a posterior object rather than a given (Russell, 2019). The same variational result applies directly to the design of reasoning systems; broader economic interpretations are separate applications.

28.5 Causal inference

Causal discovery requires more than an observational probability distribution (Pearl, 2000; Hacking, 1983). It adds a causal model class, intervention semantics, and assumptions connecting graph structure to observed independences.

Theorem 28.1 (Causal direction requires identifying structure). *Every nondegenerate correlated bivariate Gaussian distribution admits both a linear Gaussian structural representation*

$$X \longrightarrow Y$$

and a linear Gaussian structural representation

$$Y \longrightarrow X$$

with independent disturbances. Hence causal direction requires identifying structure beyond the observational distribution.

Proof. Let (X, Y) be jointly Gaussian with nonzero covariance. Regress Y on X :

$$Y = aX + \varepsilon_Y, \quad a = \frac{\text{Cov}(X, Y)}{\text{Var}(X)}.$$

For jointly Gaussian variables the residual ε_Y is independent of X . Regressing in the opposite direction gives

$$X = bY + \varepsilon_X, \quad b = \frac{\text{Cov}(X, Y)}{\text{Var}(Y)},$$

with ε_X independent of Y . Both structural equations generate the same joint observational law. Interventions or further restrictions are needed to orient the relation. $\square$

Faithfulness supplies causal identification structure. Faithfulness or stability makes causal structure identifiable by excluding exact cancellations. Its standing comes from the causal problem, empirical performance, and robustness checks.

The method contributes auditability: Markov, faithfulness, causal sufficiency, intervention, temporal, and selection assumptions remain visible throughout the result. Correlation becomes causal information only through a stated identification theorem.

29 A Rational Reconstruction of Scientific Method

Scientific inquiry is a cycle of representation, prediction, channel design, testing, criticism, revision, and replication, set against the familiar problems of falsification, underdetermination, and theory change (Popper, 1934; Duhem, 1906; Kuhn, 1962).

The scientific cycle. A well-specified empirical inquiry contains:

- (1) a representation and candidate model class;
- (2) explicit background assumptions and reference structure;
- (3) hypotheses that expose differing signal laws;
- (4) apparatus and channels whose errors are modelled;
- (5) severe tests and replication;

- (6) updating, model criticism, and revision;
- (7) a record of remaining freedoms, underdetermination, and failure.

The cycle states the inferential structure shared by well-formed empirical inquiry.

Remark 29.1 (Boundary results complete scientific inquiry). Scientific success includes selecting a theory, excluding a model class, identifying a boundary, and retaining an observationally equivalent family. A severe experiment may certify model-class failure, parameter non-identifiability, or an observationally equivalent family. Each is a complete result when it exhausts what the experiment can establish. Scientific completion records exactly that result.

29.1 Popper, Duhem, and Kuhn

Popper's lasting insight is that a theory earns standing by exposing itself to outcomes that would count against it (Popper, 1934). Naive falsificationism goes too far when it treats a failed prediction as testing one isolated sentence. Duhem's point is that a prediction depends on a package of theory, auxiliaries, initial conditions, and apparatus assumptions (Duhem, 1906). Rational criticism records which element of the package changes and what evidence licenses the revision.

Kuhn's account of scientific revolutions adds a second complication: rival theories may organise salience, measurement, and even the space of candidate questions differently (Kuhn, 1962). Rational comparison then requires explicit translations, overlapping predictions, retained successes, and new discriminating tests. Theory change thereby becomes rationally assessable; Kuhn's account marks the creative and representational work that must be supplied before comparison can proceed.

Thus falsification is retained as severe exposure, the Duhem problem is absorbed into a fully specified test package, and Kuhnian theory change is reframed as comparison across changing representations. What remains outside the formal cycle is the invention of the new representation and the social organisation of inquiry.

29.2 Model revision

A new theory is rationally preferred when it improves predictive and explanatory performance under the comparison while retaining the old theory's successful limits where those limits remain tested. "Explaining old success" is a likelihood and representation constraint on the successor.

Changes of representation, salience, instruments, and model class require a specified translation and empirical comparison between the old and new problems.

29.3 Demarcation

Scientific standing depends on a record like the following:

- Are the hypotheses and auxiliaries explicit?
- Do possible observations discriminate among them?
- Are channels calibrated and independently auditable?
- Are failed predictions allowed to revise the theory?
- Are representation and complexity choices disclosed?
- Are results reproducible or at least independently checkable?

Scientific standing grows through exposure to discriminating signals, explicit revision of the problem, and a public record of every change.

29.4 Replication and convergence

Replication tests whether an apparent result survives changes of sample, operator, apparatus, laboratory, and analysis. Independent replication therefore probes hidden channel dependence and common-mode error. Under the conditions of the Convergence Theorem, accumulated informative evidence drives agreement; outside them, replication can reveal exactly which condition failed.

Part VI

Norms and the Value of Knowledge

If any man is able to convince me that I do not think or act right, I will gladly change; for I seek the truth, by which no man was ever injured.

Marcus Aurelius, *Meditations*, tr. Long

Belief and action form successive layers. Once an answer is offered as inference, standards come with the claim: give the grounds, follow them, and align every conclusion and selection with their support. Practical and moral judgement begins one layer later, when agents, values, costs, and relations to others enter (Korsgaard, 1996; Clifford, 1877). This Part keeps the layers apart, then shows where they meet.

30 Belief, Choice, and Pragmatic Reasons

Evidence, assertion, and action are different outputs. A value or payoff supplies a reason to act, while evidence determines the belief state.

Proposition 30.1 (Practical reasons govern action). *Let two agents share the same belief state P but face different utility functions U_1 and U_2 . Their optimal actions may differ while every credence remains the same. Consequently a wager, threat, reward, or high stake can change rational action without changing evidential support.*

Proof. Expected utility maximisation chooses

$$a_i^* \in \arg \max_a \mathbb{E}_P[U_i(a, H)].$$

Changing U_i can change the maximiser while P is held fixed. The decision problem therefore carries the difference, while the evidence remains fixed. $\square$

Pragmatic reasons govern action, and evidential support governs belief. The same distinction clarifies doxastic voluntarism: agents choose investigations, reports, and actions, while the grounds determine support. The ordinary word “knowledge” may remain sensitive to pragmatic context; the distinction between support and stakes stays fixed.

Decision paradoxes such as Newcomb’s problem require the relation among prediction, causation, policy, and utility to be stated. Evidential and causal decision theories formalise different dependence structures, and the supplied structure determines which problem is being solved.

31 Norms of Inference

Definition 31.1 (Three kinds of claim). (i) A **descriptive claim** says how agents in fact reason.

- (ii) A **constitutive claim** states what must be true for an activity to count as the activity under description.
- (iii) A **practical or moral norm** supplies a reason to undertake, continue, or prioritise an activity among competing ends.

Theorem 31.1 (Norms of inference). *An agent or process that presents an output as inference is bound, in that capacity, by the standards of Part I: explicit grounds, valid dependence, invariance under equivalent descriptions, and selection fixed by the problem. When the output is also presented as complete, it must preserve the whole supported answer. These standards are constitutive of outputs presented as inference.*

Proof. The method theorem analyses what it is for an answer to count as following from its grounds. An output satisfies the inferential description when every dependence is carried by its stated grounds. A complete output additionally preserves every supported case and distinction. The standards are therefore internal to the claims being made. $\square$

The inferential standpoint and its practical boundary. A reasoned rejection of every inferential norm is itself offered as an inference and is therefore assessable by inferential standards. This makes the epistemic standpoint inescapable for reasoned evaluation. Practical and moral duties arise from the further ends and constraints governing inquiry, effort, and the place of truth among other values.

If an output is offered as following from reasons, the constitutive norms of inference apply to it. Whether one should enter the inquiry, or rank truth above other goods, is a further practical question. Once a utility and inquiry cost are supplied, Corollary 32.2 gives the exact instrumental answer; categorical moral obligation remains outside that decision problem.

The epistemic is-ought result is constitutive: it moves directly from the identity of inference to the standards of that activity, just as the rules of proof bind a purported proof and preservation laws bind a purported homomorphism.

Scope of the normative result. The Non-Smuggling rule follows from the activity's constitutive standard and requires no additional moral premise once the output is presented as an inference. An independent normative premise is required only for claims about whether an agent ought to enter the inquiry, how much effort it should spend, or how epistemic goods trade against non-epistemic ones.

32 The Value of Knowledge

Knowledge can be valuable in several different senses, beginning with the Meno problem of why knowledge is worth more than true belief (Plato, 1997).

Its instrumental value is that reliable belief improves prediction, coordination, and action. Its value constitutive of inquiry is that support connected to truth is success under the practice's own aim. Its social value is that shared reliable channels support trust and division of cognitive labour. Final or moral value requires a further account of why knowledge is worth having independently of its uses.

Theorem 32.1 (Information weakly increases optimal instrumental value). *Let X be a state, S an optional signal, $\mathcal{A}$ a finite action set, and $u(a, X)$ a bounded utility. Define*

$$V_0 = \max_{a \in \mathcal{A}} \mathbb{E}[u(a, X)]$$

and

$$V_1 = \mathbb{E}_S \left[\max_{a \in \mathcal{A}} \mathbb{E}[u(a, X) \mid S] \right].$$

Then

$$\boxed{V_1 \geq V_0.}$$

Information that is free and optional has nonnegative optimal expected utility.

Proof. Let a_0 maximise the decision without the signal. After observing any $S = s$, the inner maximum is at least the conditional expected utility of retaining a_0 . Taking expectations gives

$$V_1 \geq \mathbb{E}_S \mathbb{E}[u(a_0, X) \mid S] = \mathbb{E}[u(a_0, X)] = V_0.$$

□

Corollary 32.2 (When inquiry is instrumentally required). *Suppose obtaining S costs $c \geq 0$ and the practical objective is expected utility net of inquiry cost. Inquiry is strictly preferred when*

$$V_1 - V_0 > c,$$

rejected when $V_1 - V_0 < c$, and leaves the agent indifferent at equality. The result is exact for the supplied utility, cost, and option to ignore the signal. Categorical moral duties belong to the wider value problem.

32.1 Meno and the swamping problem

A merely true belief and a piece of knowledge may guide the same immediate action. The Meno problem asks why knowledge is worth more. The route-integrity answer is that knowledge is stable under perturbation, reusable across nearby problems, and attributable to a method. A true answer reached accidentally may guide one action; a truth-connected method supports navigation, transfer, and correction.

The swamping problem asks what a reliable route adds once truth is already present. Truth is the successful outcome; a robust route explains why that success recurs, survives relevant counterfactual changes, and can be credited to the agent or channel. This supplies the extra epistemic value, while final moral value belongs to the practical domain.

33 Internalism and Externalism

Internalism and externalism answer different questions (Goldman, 1979; Lewis, 1996; Williamson, 2000).

Internal support and connection to truth. An epistemic assessment has at least two coordinates:

- (i) **internal support**: whether the state is a valid answer to the agent's grounds;
- (ii) **external success**: whether the representation and channels connect those constraints to the world in the required way.

A demon-world subject may be internally impeccable and externally unsuccessful. A lucky guesser may be externally correct and internally unsupported.

Reliabilist, safety, competence, and virtue accounts supplement the internal analysis. Internal support records what the evidence justified; external success records how the route connected that support to the world.

Conclusion

MU is the floor of this account, and the floor is thin by design: consistent inference is possible, and nothing more is claimed at the base. Everything built here is what that one fact, taken seriously inside each stated domain, turns out to contain. An answer offered as following from its grounds must draw on those grounds alone, so every answer-relevant dependence lives in the problem, equivalent descriptions agree, and a complete report keeps the whole supported result. The principle adds nothing to any problem it meets, and by adding nothing it lets each problem show exactly what it holds.

Held to arithmetic, the discipline takes four forms, each unique in its domain. Determinate Boolean plausibility is probability. Total informational change is relative entropy. Least-assuming completion is relative-entropy projection, and under finite counting symmetry it is Shannon's maximum entropy. Revision is KL projection, carrying forward every feature the new constraints permit. Values and costs then govern choice, channels carry belief's connection to the world, and realisability, support, stable sampling, and distinguishability decide when that connection finds the truth. These are laws rather than customs: within their stated domains, no consistent alternative exists.

The classical problems, met one by one, resolve at exactly the level their grounds support. Regress ends at the exhibited floor. Indifference returns a point, a family, or a certified boundary according to the structure present. Induction becomes a set of explicit projective and world-facing conditions with the price of each on the table. Goodman moves projectibility into representation, Gettier joins internal support to route integrity, and scepticism separates the self-defeating universal doubt from the tractable local kind. Scientific realism reaches the quotient its experiments fix, and the limit theorems contribute their ceilings and empty classes as complete answers in their own right. A certified boundary is a finding, and a lawful refusal is an answer with its reasons attached.

What remains is work with a shape: generate a candidate, design a discriminating experiment, compute further, state the value at issue, or record the equivalence that survives. Each task is the same discipline carried into new ground, and the ground itself asks for nothing. It was there before the argument began, and every argument, including any brought against this one, stands on it.

One rule for what follows. Four quantitative laws. Every answer shape.

References

References are grouped by their role in the argument and alphabetized by first author within each group.

Foundations, Logic, and Formal Limits

- H. Albert. *Traktat über kritische Vernunft*. Mohr Siebeck, Tübingen, 1968. English edition: *Treatise on Critical Reason*, Princeton University Press, Princeton, 1985.
- L. Carroll. What the tortoise said to Achilles. *Mind*, 4(14):278–280, 1895.
- R. M. Chisholm. *The Problem of the Criterion*. Marquette University Press, 1973.
- F. B. Fitch. A logical analysis of some value concepts. *Journal of Symbolic Logic*, 28(2):135–142, 1963.
- K. Gödel. Über formal unentscheidbare Sätze der *Principia Mathematica* und verwandter Systeme I. *Monatshefte für Mathematik und Physik*, 38:173–198, 1931.
- S. A. Kripke. *Wittgenstein on Rules and Private Language*. Harvard University Press, 1982.
- B. Stroud. Transcendental arguments. *Journal of Philosophy*, 65(9):241–256, 1968.
- A. Tarski. Der Wahrheitsbegriff in den formalisierten Sprachen. *Studia Philosophica*, 1:261–405, 1936. English translation in *Logic, Semantics, Metamathematics*, 1956.
- A. M. Turing. On computable numbers, with an application to the Entscheidungsproblem. *Proceedings of the London Mathematical Society*, 42:230–265, 1936.

Probability, Information, and Representation

- J. Aczél. *Lectures on Functional Equations and Their Applications*. Academic Press, 1966.
- T. Bayes. An essay towards solving a problem in the doctrine of chances. *Philosophical Transactions of the Royal Society*, 53:370–418, 1763.
- R. Carnap. *Logical Foundations of Probability*. University of Chicago Press, 1950.
- R. T. Cox. Probability, frequency and reasonable expectation. *American Journal of Physics*, 14(1):1–13, 1946.
- R. T. Cox. *The Algebra of Probable Inference*. Johns Hopkins Press, 1961.
- B. de Finetti. La prévision: ses lois logiques, ses sources subjectives. *Annales de l'Institut Henri Poincaré*, 7:1–68, 1937.
- A. M. Gleason. Measures on the closed subspaces of a Hilbert space. *Journal of Mathematics and Mechanics*, 6:885–893, 1957.
- J. Y. Halpern. Cox's theorem revisited. *Journal of Artificial Intelligence Research*, 11:429–435, 1999.
- E. T. Jaynes. Information theory and statistical mechanics. *Physical Review*, 106(4):620–630, 1957.
- E. T. Jaynes. *Probability Theory: The Logic of Science*. Cambridge University Press, 2003.
- J. M. Joyce. A nonpragmatic vindication of probabilism. *Philosophy of Science*, 65(4):575–603, 1998.
- A. N. Kolmogorov. Three approaches to the quantitative definition of information. *Problems of Information Transmission*, 1(1):1–7, 1965.
- C. H. Kraft, J. W. Pratt, and A. Seidenberg. Intuitive probability on finite sets. *Annals of Mathematical Statistics*, 30(2):408–419, 1959.
- R. Pettigrew. *Accuracy and the Laws of Credence*. Oxford University Press, Oxford, 2016.
- F. P. Ramsey. Truth and probability. In R. B. Braithwaite, editor, *The Foundations of Mathematics and Other Logical Essays*. Kegan Paul, 1931. Written 1926.
- L. J. Savage. *The Foundations of Statistics*. Wiley, New York, 1954.
- D. Scott. Measurement structures and linear inequalities. *Journal of Mathematical Psychology*, 1(2):233–247, 1964.
- C. E. Shannon. A mathematical theory of communication. *Bell System Technical Journal*, 27:379–423 and 623–656, 1948.

Priors, Updating, Chance, and Imprecise Belief

- I. Csiszár. I-divergence geometry of probability distributions and minimization problems. *Annals of Probability*, 3(1):146–158, 1975.
- A. P. Dempster. Upper and lower probabilities induced by a multivalued mapping. *Annals of Mathematical Statistics*, 38(2):325–339, 1967.

- J. L. Doob. Application of the theory of martingales. In *Le Calcul des Probabilités et ses Applications*, pages 23–27. CNRS, Paris, 1949.
- S. Ghosal and A. van der Vaart. *Fundamentals of Nonparametric Bayesian Inference*. Cambridge University Press, 2017.
- H. Jeffreys. An invariant form for the prior probability in estimation problems. *Proceedings of the Royal Society A*, 186:453–461, 1946.
- R. E. Kass and L. Wasserman. The selection of prior distributions by formal rules. *Journal of the American Statistical Association*, 91(435):1343–1370, 1996.
- S. Kullback and R. A. Leibler. On information and sufficiency. *Annals of Mathematical Statistics*, 22(1):79–86, 1951.
- D. Lewis. A subjectivist’s guide to objective chance. In R. C. Jeffrey, editor, *Studies in Inductive Logic and Probability*, volume 2, pages 263–293. University of California Press, 1980.
- L. Schwartz. On Bayes procedures. *Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete*, 4:10–26, 1965.
- G. Shafer. *A Mathematical Theory of Evidence*. Princeton University Press, 1976.
- J. E. Shore and R. W. Johnson. Axiomatic derivation of the principle of maximum entropy and the principle of minimum cross-entropy. *IEEE Transactions on Information Theory*, 26(1):26–37, 1980.
- J. Uffink. Can the maximum entropy principle be explained as a consistency requirement? *Studies in History and Philosophy of Modern Physics*, 26(3):223–261, 1995.
- B. C. van Fraassen. A problem for relative information minimizers in probability kinematics. *British Journal for the Philosophy of Science*, 32(4):375–379, 1981.
- P. Walley. *Statistical Reasoning with Imprecise Probabilities*. Chapman and Hall, 1991.
- J. Williamson. *In Defence of Objective Bayesianism*. Oxford University Press, Oxford, 2010.

Knowledge, Scepticism, and Epistemic Luck

- L. Bonjour. Externalist theories of empirical knowledge. *Midwest Studies in Philosophy*, 5:53–73, 1980.
- S. Cohen. Basic knowledge and the problem of easy knowledge. *Philosophy and Phenomenological Research*, 65(2):309–329, 2002.
- K. DeRose. Solving the skeptical problem. *The Philosophical Review*, 104(1):1–52, 1995.
- A. Elga. Self-locating belief and the Sleeping Beauty problem. *Analysis*, 60(2):143–147, 2000.
- E. Gettier. Is justified true belief knowledge? *Analysis*, 23(6):121–123, 1963.
- A. Goldman. What is justified belief? In G. Pappas, editor, *Justification and Knowledge*. Reidel, 1979.
- J. Hawthorne. *Knowledge and Lotteries*. Oxford University Press, 2004.
- D. Lewis. Elusive knowledge. *Australasian Journal of Philosophy*, 74(4):549–567, 1996.
- G. E. Moore. Proof of an external world. *Proceedings of the British Academy*, 25:273–300, 1939.
- J. Pryor. The sceptic and the dogmatist. *Noûs*, 34(4):517–549, 2000.
- W. Sellars. Empiricism and the philosophy of mind. In H. Feigl and M. Scriven, editors, *Minnesota Studies in the Philosophy of Science*, volume 1, pages 253–329. University of Minnesota Press, 1956.
- J. Vogel. Reliabilism leveled. *Journal of Philosophy*, 97(11):602–623, 2000.
- T. Williamson. *Knowledge and Its Limits*. Oxford University Press, 2000.
- L. Wittgenstein. *On Certainty*. Blackwell, 1969.
- C. Wright. Warrant for nothing (and foundations for free)? *Aristotelian Society Supplementary Volume*, 78(1):167–212, 2004.

Rational Belief, Permissivism, and Epistemic Value

- W. K. Clifford. The ethics of belief. *Contemporary Review*, 29:289–309, 1877.
- J. Fantl and M. McGrath. *Knowledge in an Uncertain World*. Oxford University Press, 2009.
- M. Kopec and M. G. Titelbaum. The uniqueness thesis. *Philosophy Compass*, 11(4):189–200, 2016.
- C. M. Korsgaard. *The Sources of Normativity*. Cambridge University Press, 1996.
- D. C. Makinson. The paradox of the preface. *Analysis*, 25(6):205–207, 1965.
- G. E. Moore. A reply to my critics. In P. A. Schilpp, editor, *The Philosophy of G. E. Moore*. Open Court, 1942.

- Plato. *Meno*. In J. M. Cooper, editor, *Plato: Complete Works*. Hackett, 1997.
- B. C. van Fraassen. Belief and the will. *The Journal of Philosophy*, 81(5):235–256, 1984.
- R. White. Epistemic permissiveness. *Philosophical Perspectives*, 19:445–459, 2005.

Induction, Confirmation, and Philosophy of Science

- D. Deutsch. *The Beginning of Infinity: Explanations that Transform the World*. Allen Lane, London, 2011.
- P. Duhem. *La théorie physique: son objet, sa structure*. English translation: *The Aim and Structure of Physical Theory*, Princeton University Press, 1954.
- D. Garber. Old evidence and logical omniscience in Bayesian confirmation theory. In J. Earman, editor, *Testing Scientific Theories*, Minnesota Studies in the Philosophy of Science, volume 10. University of Minnesota Press, 1983.
- C. Glymour. *Theory and Evidence*. Princeton University Press, 1980.
- N. Goodman. *Fact, Fiction, and Forecast*. Harvard University Press, 1955.
- I. Hacking. *Representing and Intervening*. Cambridge University Press, 1983.
- C. G. Hempel. Studies in the logic of confirmation. *Mind*, 54:1–26 and 97–121, 1945.
- D. Hume. *A Treatise of Human Nature*. 1739–1740.
- T. S. Kuhn. *The Structure of Scientific Revolutions*. University of Chicago Press, 1962.
- L. Laudan. A confutation of convergent realism. *Philosophy of Science*, 48(1):19–49, 1981.
- K. Popper. *Logik der Forschung*. English translation: *The Logic of Scientific Discovery*, Hutchinson, 1959.
- W. V. O. Quine. Two dogmas of empiricism. *The Philosophical Review*, 60(1):20–43, 1951.
- P. K. Stanford. *Exceeding Our Grasp: Science, History, and the Problem of Unconceived Alternatives*. Oxford University Press, 2006.
- P. F. Strawson. *Introduction to Logical Theory*. Methuen, London, 1952.
- B. C. van Fraassen. *The Scientific Image*. Oxford University Press, 1980.
- J. Worrall. Structural realism: The best of both worlds? *Dialectica*, 43(1–2):99–124, 1989.

Social Epistemology and Collective Reasoning

- K. J. Arrow. *Social Choice and Individual Values*. Wiley, 1951; second edition, 1963.
- R. J. Aumann. Agreeing to disagree. *Annals of Statistics*, 4(6):1236–1239, 1976.
- D. Christensen. Epistemology of disagreement: The good news. *The Philosophical Review*, 116(2):187–217, 2007.
- A. Elga. Reflection and disagreement. *Noûs*, 41(3):478–502, 2007.
- M. Fricker. *Epistemic Injustice: Power and the Ethics of Knowing*. Oxford University Press, 2007.
- J. Hardwig. Epistemic dependence. *The Journal of Philosophy*, 82(7):335–349, 1985.
- T. Kelly. The epistemic significance of disagreement. In T. S. Gendler and J. Hawthorne, editors, *Oxford Studies in Epistemology*, volume 1, pages 167–196. Oxford University Press, 2005.
- H. E. Longino. *Science as Social Knowledge*. Princeton University Press, 1990.

Computation, Artificial Intelligence, and Causality

- J. Burrell. How the machine “thinks”: Understanding opacity in machine learning algorithms. *Big Data & Society*, 3(1), 2016.
- F. Chollet. On the measure of intelligence. arXiv:1911.01547, 2019.
- P. F. Christiano, J. Leike, T. Brown, M. Martic, S. Legg, and D. Amodei. Deep reinforcement learning from human preferences. In *Advances in Neural Information Processing Systems*, volume 30, 2017.
- K. Friston. The free-energy principle: a unified brain theory? *Nature Reviews Neuroscience*, 11(2):127–138, 2010.
- T. Haarnoja, A. Zhou, P. Abbeel, and S. Levine. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In *Proceedings of ICML*, volume 80, pages 1861–1870, 2018.
- M. Hutter. *Universal Artificial Intelligence: Sequential Decisions Based on Algorithmic Probability*. Springer, 2005.

- J. Leike and M. Hutter. Bad universal priors and notions of optimality. *Proceedings of Machine Learning Research*, 40:1244–1259, 2015.
- J. Pearl. *Causality: Models, Reasoning, and Inference*. Cambridge University Press, 2000.
- S. Russell. *Human Compatible: Artificial Intelligence and the Problem of Control*. Viking, New York, 2019.
- J. R. Searle. Minds, brains, and programs. *Behavioral and Brain Sciences*, 3(3):417–457, 1980.
- R. J. Solomonoff. A formal theory of inductive inference. *Information and Control*, 7:1–22 and 224–254, 1964.
- L. G. Valiant. A theory of the learnable. *Communications of the ACM*, 27(11):1134–1142, 1984.
- V. N. Vapnik and A. Ya. Chervonenkis. On the uniform convergence of relative frequencies of events to their probabilities. *Theory of Probability and Its Applications*, 16(2):264–280, 1971.
- V. Vovk, A. Gammerman, and G. Shafer. *Algorithmic Learning in a Random World*. Springer, New York, 2005.
- D. H. Wolpert and W. G. Macready. No free lunch theorems for optimization. *IEEE Transactions on Evolutionary Computation*, 1(1):67–82, 1997.
- D. M. Ziegler et al. Fine-tuning language models from human preferences. arXiv:1909.08593, 2019.