What Do We Learn from a Conditional?

Eidon Vale

Abstract

Suppose one learns that a reactor will overheat if it is restarted. Has one thereby acquired evidence that the reactor will not be restarted? A prominent recent account of trivalent conditionals says yes: learning the conditional has the same effect as learning that either there will be no restart or there will be overheating. I argue that this erases the very distinction trivalence was introduced to capture. A conditional tells us how things stand in the cases in which its antecedent obtains; it need not tell us how likely those cases are. When the source of conditional information is neutral about whether the antecedent obtains, learning the conditional should therefore leave the antecedent's probability unchanged. I call this Domain Preservation. I show that the resulting update is uniquely characterized by familiar minimum-change principles and extends from simple conditionals to arbitrary trivalent contents. Actual conditional reports may also convey information about speakers, causes, or the likelihood of their antecedents, so they need not always preserve domain probability. The result is a layered account of conditional learning that respects the difference between a conditional being false and its simply not applying.

1 Introduction

Suppose that a safety engineer tells us:

> If the reactor is restarted tonight, its core will overheat.

We trust the engineer and come to accept what she says. What, exactly, have we learned?

The most natural answer is that we have learned what will happen on the supposition that the reactor is restarted. We may previously have been unsure whether a restart would result in overheating; afterward, we are certain that it would. But we need not have learned anything about whether the reactor will be restarted. Imagine that the decision will be made later by a fair coin toss. Learning what will happen if the coin lands heads gives us no reason to revise our confidence that it will land heads.

This simple thought creates a problem for a natural account of conditional learning. The material conditional corresponding to the engineer's statement is the proposition that either the reactor will not be restarted or it will overheat. To learn that proposition by ordinary Bayesian conditionalization is to eliminate restart-without-overheating possibilities and renormalize over everything that remains. That makes a restart less probable. If the four combinations of restart and overheating initially receive equal probability, conditioning on the material conditional lowers the probability of a restart from one-half to one-third. The engineer's warning thereby becomes evidence about the future coin toss.

The result is not an accidental feature of the numbers. The material conditional can be made true in either of two ways: there may be no restart, or there may be a restart followed by overheating. No-restart possibilities therefore count positively toward the truth of the material proposition. They do not appear to count positively toward the truth of the engineer's warning. Her warning concerns what happens in the event of a restart. A failure to restart the reactor would not confirm it; it would leave the warning untested.

This contrast motivates trivalent theories of conditionals. On a simple version of the view, *if A, then B* is true when A and B are true, false when A is true and B false, and *void* when A is false. The third value is not a third way for the conditional to be correct. It marks cases in which the issue raised by the conditional does not arise.[1] Probabilistically, the conditional's credibility is its proportion of truth among the worlds at which it is not void. For a simple conditional, this yields

$$\Pr(A \to B) = \frac{\Pr(AB)}{\Pr(AB) + \Pr(A\neg B)} = \Pr(B \mid A).$$

The view thus vindicates the familiar claim that confidence in a conditional tracks conditional confidence, while explaining why false-antecedent cases do not count in its favor.[2]

Recent work by Égré et al. (2025, 2026) develops this idea into a compositional semantics and probability theory. Their framework distinguishes truth from voidness when assigning probabilities to conditionals. Surprisingly, however, the distinction disappears when a conditional is learned. On their proposed dynamics, becoming certain of $A \to B$ has the same effect on one's unconditional probabilities as conditioning on the material conditional $\neg A \vee B$. No-A worlds are semantically void, yet for purposes of updating they are treated like worlds that verify the conditional.

[1] The idea has a long history in work on conditional assertion and conditional bets; see Cooper (1968) and Belnap (1970). For a recent discussion of how trivalence supports the probabilistic treatment of conditionals, see Lassiter (2020).

[2] The equation is often called Adams's Thesis, after Adams (1975). Edgington (1995) remains the classic philosophical defense of the connection between conditionals and conditional probability.

The crucial step is not hidden in the formalism. Égré, Rossi, and Sprenger observe that the indicative and material conditionals are equivalent in their logic of certainty: making either certain requires assigning no probability to $A\neg B$. They then reason that, since learning makes its object certain, contents that impose the same certainty condition should produce the same update. But this moves too quickly. A constraint on which posteriors are admissible does not determine which admissible posterior the evidence supports. Learning events can exclude the same possibilities while bearing differently on the possibilities that remain. Sameness of posterior certainty is not yet sameness of evidence.

Trivalence itself tells us where a difference may arise. Although $A \to B$ and $\neg A \vee B$ have the same false cases, they classify $\neg A$-worlds differently. Those worlds verify the material conditional but are void for the indicative. Material conditionalization treats them as positive survivors and consequently makes A less probable. If voidness marks cases outside the issue settled by the conditional, that result requires further justification.

I will argue that the distinction between falsity and voidness should constrain learning whenever the acquisition of conditional information is *domain-neutral*. Let S be the event through which an agent receives the information. The acquisition is domain-neutral when S is equally likely whether the conditional's nonvoid domain is entered or not. This is an ordinary claim about the information-producing process. In the reactor example, it is true when the engineer's access to the consequences of a restart, and her willingness to report them, are independent of the later coin toss. It is false when a speaker's choice to mention an antecedent is itself evidence that the antecedent is likely.

Bayes's theorem then yields the central principle:

> **Domain Preservation.** When conditional information is acquired through a process that is neutral about whether the conditional's nonvoid domain is entered, learning it should preserve the probability of that domain.

For a simple conditional $A \to B$, the domain is just the set of A-worlds. Domain Preservation therefore says that domain-neutral conditional information changes one's view of what happens given A without changing one's confidence in A. Probability moves from restart-without-overheating possibilities to restart-with-overheating possibilities, while the balance between restart and no restart remains fixed.

This formulation avoids defining "pure conditional information" by the posterior it produces. Domain neutrality can be assessed independently by asking how the evidence was generated. A requested engineering forecast may be neutral about whether the contemplated intervention occurs. An unsolicited medical warning may not be neutral about whether its antecedent is a live possibility. A conditional can also interact with causal or explanatory

information that bears on its antecedent. The resulting total updates may differ even when the conditional content is the same.

The proposal accordingly separates two questions. First, how should an agent revise the balance between the conditional's true and false cases? Second, what does the acquisition process indicate about whether those cases arise at all? The trivalent probability of a conditional answers the first question. The probability of its domain answers the second. Conditional content and an information-producing process can jointly constrain both, but they should not be collapsed into the single demand that the conditional become certain.

The resulting update is numerically familiar. For a simple conditional it is Adams conditioning: $P(B \mid A)$ changes while $P(A)$ is preserved. The contribution of this paper is not the rediscovery of that rule. It is an account of why the Adams invariance is privileged by a trivalent semantics, how it extends to any content with a compositionally determined true/false/void partition, and why materialization embodies a rival evidential model rather than following from certainty or minimal change alone.

The formal result supports, but does not replace, that philosophical claim. Domain-preserving revision uniquely minimizes every divergence in a broad class, subject to the new conditional probability and domain neutrality. Materialization uniquely minimizes the same divergences when certainty is the only constraint. Minimum change therefore cannot settle the dispute before we decide what the evidence leaves fixed.

The stakes extend beyond a narrow dispute over indicative semantics. Conditionals communicate policies, dispositions, causal dependencies, predictions, threats, promises, and plans. In each case there is a difference between learning what obtains under a condition and learning whether the condition obtains. Section 2 reconstructs the materialization argument and identifies the gap between certainty equivalence and evidential equivalence. Section 3 derives Domain Preservation from a source-neutrality condition. Section 4 gives the resulting update. Section 5 shows how different acquisition models explain apparently conflicting judgments about conditional learning. Section 6 addresses objections and limits. The formal representation result is stated and proved in the Appendix.

2 The philosophical problem with materialization

The problem can be stated without taking a stand on every detail of a trivalent semantics. Let W be a finite set of possible worlds, and let P be a probability distribution over them. A trivalent content E divides W into three cells:

$$E^{\mathsf{T}}, \qquad E^{\mathsf{F}}, \qquad E^{\mathsf{V}},$$

the worlds at which E is respectively true, false, and void. Write

$$D_E = E^{\mathsf{T}} \cup E^{\mathsf{F}}$$

for the *domain* of E: the worlds at which the issue settled by E arises. Its trivalent probability is

$$\pi_P(E) = \frac{P(E^{\mathsf{T}})}{P(E^{\mathsf{T}}) + P(E^{\mathsf{F}})}, \tag{1}$$

provided $P(D_E) > 0$. This is just ordinary probability measured within the nonvoid domain. If E is bivalent, there are no void worlds and (1) reduces to $P(E)$.

For a simple conditional with bivalent antecedent and consequent,

$$(A \to B)^{\mathsf{T}} = AB, \qquad (A \to B)^{\mathsf{F}} = A\neg B, \qquad (A \to B)^{\mathsf{V}} = \neg A.$$

Equation (1) then gives $\pi_P(A \to B) = P(B \mid A)$. The numerator measures the A-cases in which the conditional succeeds; the denominator measures all the cases in which it is tested. The $\neg A$-cases belong to neither.

The materialization result concerns what happens when an agent learns E. Égré et al. (2026, §6) define trivalent conditionalization by

$$\pi_P^E(H) = \pi_P(E \to H). \tag{2}$$

Their Proposition 5 establishes that formulas equivalent in the certainty-preserving logic induce the same update. The accompanying rationale is that learning E makes E certain; hence, if certainty in E and certainty in E' always coincide, learning either should yield the same posterior.

For a simple conditional, that principle yields materialization. Let

$$M_E = E^{\mathsf{T}} \cup E^{\mathsf{V}} = W \setminus E^{\mathsf{F}}$$

be E's *material shadow*. In the simple case, M_E is $\neg A \vee B$. Since $A \to B$ and $\neg A \vee B$ have the same falsifying worlds, they are equivalent in the certainty logic. Proposition 6 therefore implies that updating on the former has the same effect on the underlying distribution as ordinary conditioning on the latter:

$$P_{\mathsf{Mat}}^E(X) = P(X \mid M_E). \tag{3}$$

It removes the false worlds and preserves all prior probability ratios among the true and void

worlds that remain. Their Updating Theorem then shows that this result also minimizes every f-divergence from the prior subject to making E certain.[3]

The reactor case displays what this means. Let R say that the reactor is restarted and O that it overheats. Suppose initially that the four possibilities are equiprobable. Table 1 compares materialization with the domain-preserving update defended below.

World	Prior	Material	Domain-preserving
RO	1/4	1/3	1/2
$R\neg O$	1/4	0	0
$\neg RO$	1/4	1/3	1/4
$\neg R\neg O$	1/4	1/3	1/4

Table 1: Two ways of becoming certain that $R \to O$.

Both updates make $P(O \mid R) = 1$. They differ in how they get there. Materialization eliminates $R\neg O$ and renormalizes globally, so $P(R)$ falls from 1/2 to 1/3. The domain-preserving update reallocates the probability of $R\neg O$ to RO while leaving the no-restart worlds untouched, so $P(R)$ remains 1/2. The latter update represents the engineer as changing our view of the consequences of a restart, not our view of whether the independent coin will produce one.

The two updates therefore disagree neither about success nor about which propositions become certain. Both make the conditional and its material shadow certain. They disagree about how probability should be distributed among possibilities compatible with those certainties. That is a difference in evidence even when it is not a difference in categorical commitment.

A small example makes the point vivid. Let B say that the butler committed the crime and G that the gardener did. Beginning from a uniform distribution, learning the disjunction $B \vee G$ by Bayesian conditionalization raises $P(B)$ from 1/2 to 2/3. Learning the conditional $\neg B \to G$ by the rule defended here leaves $P(B)$ at 1/2. Both updates eliminate $\neg B\neg G$ and make both sentences certain. But the disjunction counts B-worlds as positive ways for the evidence to be true, whereas those worlds are void for the conditional. The sentences can have the same certainty consequences without presenting the same evidence.

This exposes the gap in the argument from certainty equivalence to update equivalence. Equivalence in the certainty logic determines the same *admissible set*: after either sentence is learned, no probability may remain on $\neg B\neg G$. It does not select the same member of that

[3]These are Definition 8, Propositions 5 and 6, and Theorem 4 of Égré et al. (2026). The Reduction Theorem extends the result to arbitrary compounds by identifying a simple conditional with the same trivalent semantic profile.

set. The disjunction and conditional bear differently on B-worlds, and the two posteriors preserve that difference while agreeing about every resulting certainty.

The point is general. To be certain of a trivalent content E is only to assign no probability to E^{F}. That constraint does not settle the division of probability between E^{T} and E^{V}. Yet the distinction between those cells is exactly what separates an indicative conditional from its material shadow. Proposition 5 effectively coarse-grains the trivalent partition as soon as learning begins: once the false cell is removed, truth and voidness are dynamically indistinguishable. That is an additional principle about evidence, not a consequence of the semantics or of the requirement of success.

Materialization makes the additional principle concrete. Because true and void worlds both survive the elimination of falsity, it preserves their prior probability ratio. Equivalently, it treats the learning signal as present at every true and every void world, and absent at every false world. The next section argues that this is the wrong acquisition model when the source is neutral about whether the conditional's domain is entered. In that case, the signal should discriminate truth from falsity within the domain without favoring the void cases over the domain as a whole.

3 From source neutrality to Domain Preservation

Voidness does not by itself entail any unconditional probabilistic independence. A false antecedent can be correlated with the truth of the consequent, with a speaker's evidence, or with the fact that the conditional was asserted. The relevant independence concerns not the events described by the conditional, but the process through which the conditional information is acquired.

Let S be the event that the agent receives the relevant signal: an engineer's report, the output of a model, an experimental result, or any other vehicle of information. Assume for the moment that both D_E and E^{V} have positive probability. The acquisition process is *domain-neutral* for E just in case

$$P(S \mid D_E) = P(S \mid E^{\mathsf{V}}). \tag{4}$$

This says that receiving the signal is no evidence about whether the issue raised by E arises. It does not say that the signal is uninformative. Within D_E, its likelihood may differ sharply between E^{T} and E^{F}, and hence change $\pi_P(E)$.

Domain Preservation now follows directly. By Bayes's theorem and (4),

$$\begin{aligned} P(D_E \mid S) &= \frac{P(S \mid D_E)P(D_E)}{P(S \mid D_E)P(D_E) + P(S \mid E^{\mathsf{V}})P(E^{\mathsf{V}})} \\ &= P(D_E). \end{aligned} \tag{5}$$

Thus:

> **Domain Preservation.** When conditional information about E is acquired through a domain-neutral process, the posterior should satisfy
>
> $$Q(D_E) = P(D_E).$$

This is not a definition of pure evidence in terms of its posterior. Equation (4) is a claim about the likelihood of an observable signal under two possible conditions. It can be assessed from the design of an experiment, the expertise and incentives of a source, or the causal route by which the information was obtained. Preservation is its Bayesian consequence.

The reactor case supplies a clean instance. The engineer is asked in advance what overheating behavior a restart would produce. Her access to that dependency and her willingness to answer are independent of the coin toss that will later determine whether the reactor is restarted. Her answer can be much more likely when restart leads to overheating than when it does not, while remaining equally likely whether the coin later produces a restart or no restart. The report is informative about the contrast inside the restart domain and neutral about entry into it.

Materialization represents a different signal. For categorical learning it assigns likelihood one to every world in $E^{\mathsf{T}} \cup E^{\mathsf{V}}$ and likelihood zero to every world in E^{F}. If $t = P(E^{\mathsf{T}})$, $f = P(E^{\mathsf{F}})$, and $r = P(D_E) = t + f$, then

$$P(S_{\mathsf{Mat}} \mid D_E) = \frac{t}{r}, \qquad P(S_{\mathsf{Mat}} \mid E^{\mathsf{V}}) = 1.$$

Unless the conditional was already certain within its domain, the material signal is more likely at void worlds than at domain worlds. It is therefore evidence against entering the domain. Materialization does not merely make E certain in a neutral way; it builds a particular answer to the domain question into the learning event.

Why should trivalent semantics privilege the neutral baseline? Because the semantic point of voidness is to distinguish failure from non-application. A conditional bet is lost in E^{F} but called off in E^{V}; a prediction is falsified in the former and untested in the latter. A source that reports on whether the prediction succeeds when tested can discriminate truth from

falsity without thereby predicting whether a test occurs. Trivalence identifies the boundary relative to which source neutrality is to be assessed. Without that semantic structure, the relevant independence condition would have to be supplied independently.

The same thought appears in the Ramsey-test picture. To assess $A \to B$, one asks how credible B is under the supposition that A. One can learn the answer to that hypothetical question without learning how likely one is to enter the supposition. A weather service can report what a venue will do if it rains without issuing a weather forecast. More generally, the practical point of conditional inquiry is often to investigate a dependency while leaving the probability of its antecedent open.

Finally, the argument explains a familiar feature of Adams conditioning. Adams's rule changes $P(B \mid A)$ while preserving $P(A)$. Trivalent semantics identifies A as the nonvoid domain of $A \to B$, while source neutrality explains why its probability is invariant.[4] The present account therefore reconstructs rather than replaces Adams conditioning: it derives the rule's characteristic invariance from a semantic partition plus a claim about the acquisition process.

Not every process is neutral. An unsolicited warning may reveal that a speaker regards its antecedent as likely; information about a dangerous consequence may affect an agent's future choice; and a conditional may interact with causal or explanatory evidence about its antecedent. Such cases require a different likelihood model. This is compatible with evidence that conditional reports can raise, lower, or preserve antecedent probabilities (Douven, 2012; Vandenburgh, 2021), and with experiments showing sensitivity to sources and contexts (Collins et al., 2020). The claim is not that one update fits every utterance. It is that trivalent content identifies a principled domain-neutral baseline, while the acquisition model explains departures from it.

4 Domain-preserving conditionalization

The update itself is simple. Let

$$t = P(E^{\mathsf{T}}), \qquad f = P(E^{\mathsf{F}}), \qquad v = P(E^{\mathsf{V}}), \qquad r = t + f = P(D_E).$$

Suppose the evidence requires the posterior trivalent probability of E to be s. The number s may be one, as when the agent categorically accepts the conditional, or it may represent uncertain testimony. The update gives total probability sr to the true cell, $(1 - s)r$ to the

[4]For the relevant probability kinematics, see Adams (1975) and Bradley (2005); for applications to Judy Benjamin and conditional learning, see Douven and Romeijn (2011) and Douven (2012). Bradley (2000) offers an earlier preservation-based treatment of conditionals.

false cell, and v to the void cell. It preserves prior ratios within each cell.

Definition 1 (Domain-preserving conditionalization)**.** Assume initially that $t, f > 0$. Domain-preserving conditionalization on E to level $s \in [0, 1]$ produces the distribution

$$P_{\mathsf{DP}}^{E,s}(w) = \begin{cases} P(w)\dfrac{sr}{t}, & w \in E^{\mathsf{T}}, \\ P(w)\dfrac{(1-s)r}{f}, & w \in E^{\mathsf{F}}, \\ P(w), & w \in E^{\mathsf{V}}. \end{cases} \tag{6}$$

The three multipliers can be read as relative likelihoods of the learning signal. Their prior-weighted average within the domain is

$$\frac{t}{r}\frac{sr}{t} + \frac{f}{r}\frac{(1-s)r}{f} = 1,$$

which equals the multiplier at every void world. The signal is therefore domain-neutral, while its relative likelihood at true and false worlds is exactly what is needed to produce posterior conditional probability s.

The rule additionally preserves every prior probability ratio within each truth-value cell. It thereby represents a signal that bears on the balance between truth and falsity, is neutral about domain entry, and carries no finer information within the cells. In the reactor example, categorical learning sets the probability of $R\neg O$ to zero and transfers its mass, proportionally, to RO; the no-restart worlds are untouched. The representation result stated and proved in the Appendix shows that this is the unique minimum-change update for every strictly convex g-divergence, including Kullback–Leibler divergence, subject to success and Domain Preservation (Theorem 2).

The result clarifies the disagreement with materialization. Égré et al. (2026) appeal to minimum-divergence results in support of their update. But divergence minimization cannot decide what the evidence requires. If the only constraint is that E become certain, minimizing divergence selects materialization: eliminate E^{F} and rescale all surviving worlds. If Domain Preservation is included, the same class of divergences unanimously selects (6). The disagreement concerns the constraints supplied by conditional content, not the choice of a mathematical measure of change.[5]

Two limiting cases help locate the proposal. First, if E is bivalent, then E^{V} is empty and

[5]This is one instance of a general moral about relative-information methods: the representation of the evidential constraint does substantive work. The Judy Benjamin literature is an especially clear illustration; see van Fraassen (1981); Douven and Romeijn (2011).

$D_E = W$. Categorical domain-preserving conditionalization simply eliminates the false worlds and renormalizes over the true ones. It is ordinary Bayesian conditionalization. The proposal does not replace Bayesian learning for propositional evidence; it extends it to contents whose semantic domain may be smaller than the space of possibilities.

Second, if E is the simple conditional $A \to B$, then (6) is Adams conditioning. Writing $p = P(B \mid A)$, the posterior of any proposition X is

$$\begin{aligned} P_{\mathsf{DP}}^{A\to B,s}(X) = P(X \cap A \cap B)\frac{s}{p} \\ + P(X \cap A \cap \neg B)\frac{1-s}{1-p} + P(X \cap \neg A). \end{aligned} \tag{7}$$

The last term is left alone, so $P(A)$ is preserved while $P(B \mid A)$ becomes s. The numerical rule is familiar. What the present account adds is an explanation of why its invariance is appropriate: the antecedent is the conditional's nonvoid domain.

4.1 The cost of materialization

The difference between the two updates is systematic. For categorical learning, materialization gives

$$P_{\mathsf{Mat}}^{E}(E^{\mathsf{V}}) = \frac{v}{1-f}, \qquad P_{\mathsf{Mat}}^{E}(D_E) = \frac{t}{1-f},$$

whereas domain-preserving conditionalization gives

$$P_{\mathsf{DP}}^{E,1}(E^{\mathsf{V}}) = v, \qquad P_{\mathsf{DP}}^{E,1}(D_E) = t + f.$$

The amount by which materialization lowers the probability of the domain, relative to the domain-preserving update, is

$$P_{\mathsf{DP}}^{E,1}(D_E) - P_{\mathsf{Mat}}^{E}(D_E) = \frac{fv}{1-f}. \tag{8}$$

The updates therefore differ exactly when the prior gives positive probability to both false and void worlds. Those are precisely the cases in which the evidence is genuinely informative and the distinction between falsity and voidness has work to do.

For a simple conditional, let $r = P(A)$ and $p = P(B \mid A)$. Materialization yields

$$P_{\mathsf{Mat}}^{A\to B}(A) = \frac{rp}{1 - r(1-p)} < r \tag{9}$$

whenever $0 < r < 1$ and $p < 1$. It always makes the antecedent less likely. Domain

Preservation leaves its probability at r. Equation (9) is the general form of the antecedent flight seen in the reactor case.

4.2 Beyond simple conditionals

Nothing in (6) depends on the syntax of E. It requires only the partition into true, false, and void worlds. This matters because the semantics of Égré et al. (2026) assigns such a partition compositionally to Boolean compounds and nested conditionals, and proves that each is equivalent to a simple conditional $A_E \to B_E$ for some bivalent A_E and B_E. Domain-preserving learning of the complex formula then changes $P(B_E \mid A_E)$ while preserving $P(A_E)$.

The point is not to add another technical extension for its own sake. Adams conditioning by itself tells us how to revise on a simple $A \to B$, but it does not tell us what counts as the domain of a conjunction or nesting of conditionals. A compositional trivalent semantics does. It identifies which possibilities test the complex content, and the update preserves the total probability of those possibilities.[6]

The resulting dynamics also preserves the ordinary inferential use of a conditional. If an agent first learns $A \to B$ categorically and later learns A, ordinary conditioning on A makes B certain. Conditional information changes how the agent is disposed to revise when its antecedent is learned, without pretending that the antecedent has already become less likely. This is the familiar practical role of a conditional: it stores information for use under a possible future supposition.

5 Conditional content and total evidence

The most serious worry about Domain Preservation is that conditional reports often do change antecedent probabilities. The response should not be to deny the examples. It should be to represent the additional information that distinguishes them.

The separation is transparent if the posterior is described by two coordinates:

$$s = \pi_Q(E), \qquad r' = Q(D_E).$$

The first records confidence in the conditional within its domain. The second records confidence that the domain is entered. Minimum-divergence revision subject to both targets

[6]Under the Cooper semantics, for example, the domain of a conjunction such as $(A \to B) \wedge (C \to D)$ is the set of worlds at which at least one component antecedent obtains. The exact reduction is supplied by the semantics; the updating rule need not inspect the formula's syntax.

gives

$$P^{E;s,r'}(w) = \begin{cases} P(w)\dfrac{sr'}{t}, & w \in E^{\mathsf{T}}, \\ P(w)\dfrac{(1-s)r'}{f}, & w \in E^{\mathsf{F}}, \\ P(w)\dfrac{1-r'}{v}, & w \in E^{\mathsf{V}}. \end{cases} \tag{10}$$

This equation is not intended as a complete psychology of testimony. It makes a conceptual point. Domain-neutral acquisition is the special case $r' = r$. Other evidence may supply a different value for r', or may impose still finer constraints within the three cells.

Materialization is itself one choice of the second coordinate. For categorical evidence it sets $s = 1$ and

$$r'_{\mathsf{Mat}} = \frac{t}{t+v}.$$

It is therefore not a neutral way of making E certain. It couples certainty in the conditional with a particular reduction in domain probability. Equation (10) permits that joint update when the context supports it. What it denies is that every acquisition of the conditional supports it.

Three kinds of case illustrate the distinction. In *domain-neutral* cases, a source is authoritative about what happens given A but not about whether A. A hotel employee can report what the hotel will do if it rains without having any meteorological information; an engineer can report the consequences of a restart without knowing the future coin toss. Here Domain Preservation supplies the total update.

In other cases an utterance provides *pragmatic domain evidence*. A physician who volunteers "If the biopsy is malignant, we will begin chemotherapy immediately" may make malignancy seem more salient or probable. The effect depends on a model of why a cooperative speaker would mention that possibility. The same sentence, offered in response to a narrowly framed question about hospital policy, need not have it. The words may express the same conditional content while the speech acts provide different total evidence.

Finally, some conditionals communicate or interact with *structural information*. If we know that passing an exam would lead to a ski trip and then observe the purchase of ski equipment, the conditional can make passing a better explanation of what we observe (Douven, 2012). If a conditional is learned as part of a causal model, evidence can propagate across its domain boundary (Vandenburgh, 2021). In these cases a posterior that preserves $P(A)$ may be wrong. But the movement in $P(A)$ depends on background information and structural assumptions, not merely on becoming confident in the A–B connection.

Edgington's coin case shows that the different answers correspond to different acquisition models. A coin is equally likely to be double-headed or double-tailed and, independently,

equally likely to be flipped or not. We learn:

If the coin was flipped, it landed heads.

Let D say that the coin is double-headed and F that it was flipped. Initially, the four combinations have probability 1/4. The conditional is true at DF, false at $\neg DF$, and void in the two no-flip cases. Table 2 displays three possible likelihood models for the report. The entries are relative likelihoods, so multiplying a row by a common positive constant makes no difference.

Acquisition model	DF	$\neg DF$	$D\neg F$	$\neg D\neg F$	$P(D \mid S)$
Material shadow	1	0	1	1	2/3
Domain-neutral content	2	0	1	1	3/4
Stable-type report	1	0	1	0	1

Table 2: Three evidential models for Edgington's coin.

The first row represents materialization. The report is present throughout the material shadow and absent at the false world. It is twice as likely when the coin was not flipped as when it was, and so constitutes evidence against F. Global renormalization gives $P(D \mid S) = 2/3$.

The second row represents categorical domain-neutral conditional information. Within the flip domain, the report discriminates perfectly between heads and tails. Its average likelihood conditional on F is nevertheless the same as its likelihood conditional on $\neg F$. It therefore preserves $P(F)$ and, absent further information within the cells, leaves both no-flip worlds at their prior probabilities. The result is $P(D \mid S) = 3/4$.

The third row represents a report grounded in the coin's stable type. An informant who inspected the coin may be disposed to issue the conditional at both DF and $D\neg F$, but not at either double-tailed world. This signal is also neutral about whether a flip occurred: its average likelihood is the same conditional on F and on $\neg F$. Yet it carries additional information *within* the void cell, ruling out $\neg D\neg F$ and making D certain.[7]

Thus 2/3, 3/4, and 1 are not rival attempts to extract a posterior from the bare demand that the conditional become certain. They answer different questions about how the report was generated. Domain Preservation rules out the first model when the source is neutral about the flip; minimum change within the truth-value cells distinguishes the second from the third when no stable-type information is supplied. The hard case therefore supports a layered account rather than an unconditional choice among three numbers.

[7]The example originates with Edgington (1995) and is revisited by Égré et al. (2026). The latter also treats the judgment that $P(D \mid S) = 1$ as depending on information about the source. The present comparison adds that their value 2/3 and the domain-preserving value 3/4 likewise correspond to distinct likelihood models.

This layered account agrees with Eva et al. (2020) that interpreting evidence must be distinguished from updating on an interpretation. It adds that a trivalent content does more than impose the probability-one constraint $P(E^{\mathsf{F}}) = 0$. By distinguishing true from void cases, it supplies an invariance appropriate to a domain-neutral acquisition model. Other contexts may supply a different domain likelihood or finer information within the truth-value cells.

6 Objections and scope

6.1 Does source neutrality merely relocate the assumption?

One might agree that (4) entails Domain Preservation but deny that semantics gives us any reason to impose that condition. The proposal would then replace an unexplained posterior constraint with an unexplained likelihood constraint.

Semantics alone does not entail source neutrality, and the argument does not require that it do so. The information-producing process is partly empirical. What the semantics contributes is an independently motivated partition of worlds into cases in which the conditional succeeds, fails, or does not apply. That partition allows us to ask two questions of a source: whether it discriminates truth from falsity within the domain, and whether it discriminates domain from voidness. In a requested conditional forecast, the first answer may be yes and the second no. In an unsolicited warning, both may be yes. These are different causal hypotheses about a report, and not merely different descriptions of its posterior.

Nor is materialization free of such a hypothesis. Its likelihood is constant over true and void worlds and zero over false worlds. As Section 3 showed, that makes the signal more likely when the conditional is void than when its domain is entered. There may be contexts in which this is the right model, but it is not forced by the fact that the conditional becomes certain. The choice is between acquisition models, not between a model and an assumption-free application of Bayes.

Source neutrality also does less work than the complete update rule. It fixes the posterior probability of the domain but does not determine probability ratios within the true, false, or void cells. The representation theorem preserves those ratios as a separate minimum-change commitment. Edgington's stable-type model illustrates the distinction: a signal can be neutral about whether the coin was flipped while carrying further information within the no-flip cell. Keeping source neutrality and cellwise conservatism separate makes the assumptions of the account visible.

6.2 Is domain-neutral acquisition too artificial?

One might accept this reasoning for carefully controlled examples while doubting that ordinary conditional testimony ever isolates content so neatly. Speakers choose antecedents for reasons; conditionals often convey causal structure; hearers draw on background knowledge. Perhaps the proposed baseline has little application outside the philosopher's laboratory.

But semantic theory routinely isolates literal contribution from pragmatic and source-sensitive effects. The fact that an assertion of p can suggest that the speaker has special evidence, declined to assert something stronger, or regards p as relevant does not make its content explanatorily idle. Conditionals carry especially rich additional information, which makes a baseline more useful rather than less. Without it, every effect of a conditional utterance is attributed to an undifferentiated learning event.

Domain-neutral cases are also common. We solicit forecasts conditional on policy choices, engineering predictions conditional on interventions, legal consequences conditional on rulings, and plans conditional on uncertain contingencies. In doing so, we often deliberately hold open the question whether the condition will obtain. Empirical work showing that source and context alter responses to conditional testimony fits this picture: those factors supplement a stable content rather than showing that there is no content-level update (Collins et al., 2020).

6.3 Limits and alternatives

The proposal is not a complete theory of radical revision. Like Bayesian and Adams conditioning, (6) has support limitations. If $P(D_E) = 0$, the trivalent probability is undefined. If learning requires positive posterior probability in a cell to which the prior assigned zero, a minimum-divergence update absolutely continuous with respect to the prior cannot deliver it. Popper functions, lexicographic probabilities, imaging, or a prior-expansion stage may be required.[8]

Nor does the argument establish that trivalent semantics is superior to every suppositional, dynamic, causal, or imaging theory. Dynamic accounts can model the probability of a conditional in terms of hypothetical update rather than the probability of a three-valued content (Ciardelli and Ommundsen, 2024). Imaging and causal approaches can represent relations among worlds that the true/false/void partition omits. The present argument is directed at a question internal to the trivalent project: once false-antecedent cases are treated as semantically void, should that distinction disappear when the conditional is learned?

The answer is no. Trivalence has the resources to preserve the central suppositional

[8]For imaging approaches to uncertain conditional evidence, see Günther (2018); for the difficulties posed by zero probabilities and radical conditional learning, see Günther and Trpin (2023).

intuition—a conditional changes one's view under its antecedent without automatically changing one's view of the antecedent—within a compositional, truth-conditional framework. Materialization obscures those resources by treating the difference between truth and voidness as dynamically idle.

7 Conclusion

A conditional-learning event can bear on two different questions: what happens within the conditional's domain, and whether that domain is entered. Trivalent semantics represents the distinction. True and false worlds answer the first question; void worlds mark the cases in which it does not arise.

The demand that a learned conditional become certain does not preserve this structure. An indicative conditional and its material shadow exclude the same false worlds, so they impose the same certainty constraint. But they disagree about whether the remaining false-antecedent worlds verify the information or lie outside its domain. Inferring update equivalence from certainty equivalence erases that disagreement rather than showing it to be epistemically irrelevant.

The acquisition process determines whether the distinction matters in a given case. When the learning signal is equally likely whether the domain is entered or not, Bayes's theorem requires Domain Preservation. The resulting minimum-change update alters the balance of truth and falsity while preserving domain probability and prior ratios within each truth-value cell. It reduces to Bayesian conditionalization for bivalent evidence and to Adams conditioning for simple conditionals. A compositional trivalent semantics supplies the domains needed to apply the same principle to complex contents.

Materialization corresponds to a different likelihood model: the signal is present at every true and void world and absent at every false one. Unless the conditional was already certain, that signal is more likely when the conditional is void than when its domain is entered. It is therefore evidence against the domain, not a neutral implementation of conditional certainty.

Actual reports may depart from either simple model. Speaker choice can make an antecedent salient; causal structure can transmit information across the domain boundary; and a source may possess stable-property information that distinguishes worlds within the void cell. These possibilities explain why no single posterior fits every conditional utterance. They do not show that the true/false/void partition is dynamically idle. They show why a theory must represent both conditional content and the process through which it is acquired.

Learning a conditional and learning its material counterpart may therefore produce the same certainties without supplying the same evidence. For conditionals, the difference between

falsity and irrelevance is not merely a difference in truth value. It is a difference in what there is to learn.

A Representation theorem

Let $g : \mathbb{R}_{\geq 0} \to \mathbb{R}$ be differentiable and strictly convex, with $g(1) = 0$, and define

$$D_g(Q, P) = \sum_{w \in W} P(w) g\left(\frac{Q(w)}{P(w)}\right). \tag{11}$$

Kullback–Leibler divergence is the special case $g(x) = x \log x$.

Theorem 2 (Representation)**.** *Let W be finite, let P have full support, and suppose $t, f > 0$. For any $s \in (0, 1)$, $P^{E,s}_{\mathsf{DP}}$ is the unique distribution that minimizes every strictly convex divergence $D_g(Q, P)$ subject to*

$$\pi_Q(E) = s, \tag{12}$$

$$Q(D_E) = P(D_E). \tag{13}$$

The categorical cases are obtained by the corresponding boundary solution whenever the target cell has positive prior probability.

Proof of Theorem 2. The two constraints fix the total posterior mass of all three cells. Since $Q(D_E) = r$ and $\pi_Q(E) = s$,

$$Q(E^{\mathsf{T}}) = sr, \qquad Q(E^{\mathsf{F}}) = (1 - s)r, \qquad Q(E^{\mathsf{V}}) = v.$$

It remains to determine how each fixed mass is distributed within its cell. Consider any cell C with $P(C) > 0$ and required posterior mass $Q(C) = m$. Its contribution to the divergence is

$$P(C) \sum_{w \in C} P(w \mid C) g\left(\frac{m}{P(C)} \frac{Q(w \mid C)}{P(w \mid C)}\right).$$

The $P(\cdot \mid C)$-weighted mean of $Q(w \mid C)/P(w \mid C)$ is one. By strict convexity and Jensen's inequality, the displayed expression is uniquely minimized when

$$\frac{Q(w \mid C)}{P(w \mid C)} = 1$$

for every $w \in C$. Equivalently, $Q(w) = P(w)m/P(C)$ throughout C. Applying this result to the true, false, and void cells yields exactly (6). At a boundary value of s, one of the first two

cells is assigned zero mass and the same argument applies to the remaining nonzero cells. □

References

Adams, Ernest W. 1975. *The Logic of Conditionals*. Dordrecht: Reidel.

Belnap, Nuel D. 1970. Conditional assertion and restricted quantification. *Noûs* 4(1): 1–12.

Bradley, Richard. 2000. A preservation condition for conditionals. *Analysis* 60(3): 219–222.

Bradley, Richard. 2005. Radical probabilism and Bayesian conditioning. *Philosophy of Science* 72(2): 342–364.

Ciardelli, Ivano, and Adrian Ommundsen. 2024. Probabilities of conditionals: Updating Adams. *Noûs* 58(1): 26–53.

Collins, Peter J., Karolina Krzyżanowska, Stephan Hartmann, Gregory Wheeler, and Ulrike Hahn. 2020. Conditionals and testimony. *Cognitive Psychology* 122: 101329.

Cooper, William S. 1968. The propositional logic of ordinary discourse. *Inquiry* 11(1–4): 295–320.

Douven, Igor. 2012. Learning conditional information. *Mind & Language* 27(3): 239–263.

Douven, Igor, and Jan-Willem Romeijn. 2011. A new resolution of the Judy Benjamin problem. *Mind* 120(479): 637–670.

Edgington, Dorothy. 1995. On conditionals. *Mind* 104(414): 235–329.

Égré, Paul, Lorenzo Rossi, and Jan Sprenger. 2025. Certain and uncertain inference with indicative conditionals. *Australasian Journal of Philosophy* 103(3): 569–596.

Égré, Paul, Lorenzo Rossi, and Jan Sprenger. 2026. Probability for trivalent conditionals. *Mind*, in press.

Eva, Benjamin, Stephan Hartmann, and Soroush Rafiee Rad. 2020. Learning from conditionals. *Mind* 129(514): 461–508.

Günther, Mario. 2018. Learning conditional information by Jeffrey imaging on Stalnaker conditionals. *Journal of Philosophical Logic* 47(5): 851–876.

Günther, Mario, and Borut Trpin. 2023. Bayesians still don't learn from conditionals. *Acta Analytica* 38(3): 439–451.

Lassiter, Daniel. 2020. What we can learn from how trivalent conditionals avoid triviality. *Inquiry* 63(9–10): 1087–1114.

van Fraassen, Bas C. 1981. A problem for relative information minimizers in probability kinematics. *British Journal for the Philosophy of Science* 32(4): 375–379.

Vandenburgh, Jonathan. 2021. Conditional learning through causal models. *Synthese* 199: 2415–2437.