Dependency Mathematics

A Technical Justification of the Corrections Applied When Users Declare Violations of Conditional Independence

This document is the companion to the Numerical Assessment *reference. It presumes familiarity with the Truth Seeker model developed there, particularly the conditional-independence assumption of Section 4.6, which this document is written to address.*

Contents

1 Purpose and Scope

The Truth Seeker assessment described in the *Numerical Assessment* document assumes that the sub-arguments beneath any node are conditionally independent given the truth of that node. When this assumption is violated — because two arguments cite overlapping evidence, or because one argument's truth causally influences another's — the multiplicative aggregation of Bayes factors overstates the strength of the compounded evidence. Users of the application can flag such structural relationships through the *Dependency system*: designated groups of nodes and directional or bidirectional links between them, each carrying a numerical factor. The assessment engine responds by applying corrections that discount overlapping evidence or dampen causally-linked contributions.

This document justifies the mathematics of those corrections. It is written for the same technically qualified reviewer addressed by the *Numerical Assessment* document. The intent is that a reader who accepts the framework of that document should, after reading this one, understand precisely what the Dependency system does, what modelling choices underlie it, and where its corrections rest on principled derivations versus engineered heuristics.

1.1 What this document justifies

The mathematical operations described here are those actually executed by the application's dependency-adjustment engine.[1] Every formula in this document corresponds to code that runs whenever a user has activated dependency calculation on a tree containing at least one dependency group.

The dependency system as a whole comprises three distinct pieces of machinery, all of which this document addresses:

- **A group-level aggregation**, in which each dependency group is assigned a single "group Bayes factor" derived from its member nodes' individual contributions.
- **A correction step**, in which each group's Bayes factor is adjusted based on its declared relationships to other groups. Two correction forms exist, one for mutual (correlational) dependencies and one for causal (*A affects B*) dependencies.
- **A redistribution step**, in which the ratio of adjusted to raw group Bayes factor is turned into a per-node multiplier that the Truth Seeker engine applies to each affected node's odds during its recursive computation.

This document is organised as follows. Section 2 establishes terminology and notation. Section 3 restates the problem the corrections address — what happens under naive multiplication when independence fails — and grounds the remainder of the document in the specific structural failures the system is designed to correct. Section 4 develops the group-level aggregation framework. Section 5 presents the mutual (correlation) correction. Section 6 presents the causal correction. Section 7 covers the structural overlap auto-adjustment; Section 8 the iterative convergence procedure that handles cyclic dependencies. Section 9 develops the per-node multiplier redistribution, and Section 10 works through a complete end-to-end example. Section 11 is the modelling-choice audit; Section 12 the tooltip reference; Section 13 the bibliography.

1.2 What this document defers

The correction formulas of this document are engineering heuristics motivated by information theory and Bayesian intuition, not derivations from a specified joint probability distribution. Section 11 addresses this honestly and states what alternative approaches exist in the literature and why they were not adopted.

The visual boundary rendering of dependency groups on the canvas (convex-hull polygon calculation, label positioning) is not addressed here; it is presentation logic, not part of the assessment mathematics.

[1] Specifically, the pure static methods of `EditorDependencyCalculator` in `iat_editor_ui/lib/state/editor_dependency_calculator.dart`. This is the single, consolidated calculator following the August 2026 refactor; earlier code paths that duplicated this logic in the host application have been removed.

1.3 Structure

Section 2 establishes the terminology, notation, and structural constraints of the dependency system. Section 3 restates the problem the corrections address — what happens under naive multiplication when independence fails — and grounds the remainder of the document in the specific structural failures the system is designed to correct. Section 4 develops the group-level aggregation framework, including the choice of geometric-mean aggregation and the interpretation of the raw versus effective group Bayes factor. Section 5 presents the mutual (correlation) correction in detail, with derivation, worked examples, and limiting-case analysis.

2 Terminology and Conventions

Throughout this document, we adopt the terminology used by the application's user interface. The formal names of code entities are given where relevant, but the primary vocabulary is what a user sees on the screen.

2.1 Dependency system entities

Group (G). A user-defined set of argument nodes drawn from anywhere in the tree. Groups are constructed by lasso-selection of nodes on the canvas: the user draws a boundary around the nodes they wish to associate, and the encircled set becomes the group. Groups are labelled sequentially — "Grp 1", "Grp 2", and so on — with numeric reclamation when a group is deleted. In the mathematics of this document, a group is identified with its membership set $G \subseteq \mathcal{N}$, where $\mathcal{N}$ denotes the set of all nodes in the tree.

The application enforces the following cardinality constraints:

- At most 20 dependency groups per tree.
- At most 100 nodes total tracked by the dependency system.
- Each node may belong to between 1 and 7 groups.

These constraints bound the computational cost of the correction procedure but are not otherwise relevant to the mathematics. A node's membership in multiple groups will introduce a redistribution consideration in Section 9; until then, the reader may assume unique membership without loss of generality.

Link. A directed or bidirectional relationship between two groups, expressing that their evidential contributions are not independent. Each link carries a numerical *factor* and a *dependency type*. Both are user-declared through the Create/Edit Dependency dialog.

Factor (f). A real number in $[-1.0, +1.0]$ expressing the strength and sense of the dependency. The user interface presents the factor as a single input with three named regions of the range:

- $f = +1.0$: *reinforcing.* The two groups' evidences are treated as maximally overlapping (mutual) or as one strongly modulating the other (causal); the correction applies at its full declared strength.
- $f = 0$: *independent.* The declared relationship has no effect; the groups' contributions are treated as though no link had been declared.
- $f = -1.0$: *compensatory.* The relationship inverts: overlap becomes anti-correlation (in the mutual case), and modulation becomes amplification (in the causal case). Section 5.5 and Section 6 describe the resulting behaviour.

The user interface presents the factor field with the helper text "`+1.0 = reinforcing, 0 = independent, -1.0 = compensatory`".

Dependency Type (τ). A categorical choice among three values:

- *Mutual*: bidirectional. Each group's Bayes factor is discounted based on its overlap with the other.
- *A affects B*: unidirectional. Group B's Bayes factor is modulated by group A's truth level.
- *B affects A*: unidirectional, in the reverse direction.

The last of these is offered in the UI for user symmetry but is normalised at save time to a *A affects B*–type link stored on group B pointing to group A. Consequently, only two distinct computational cases exist: *mutual* and *causal* (which we shall refer to uniformly as "*A affects B*" throughout, since that is the persisted form). This document formalises both and does not further discuss the UI-facing third case.

Structural overlap and overlap ratio (r). For two groups G_A, G_B, the *shared* node count is $|G_A \cap G_B|$ and the *union* node count is $|G_A \cup G_B|$. The overlap ratio is:

$$r(G_A, G_B) = \frac{|G_A \cap G_B|}{|G_A \cup G_B|} \tag{1}$$

This is the Jaccard similarity coefficient of the two membership sets. It takes values in $[0, 1]$: zero when the groups are disjoint, one when they are identical. The overlap ratio plays a role in the effective-factor auto-adjustment procedure of Section 7 and in the user-facing overlap warning discussed in Section 3.

2.2 Computed quantities

Raw group Bayes factor ($\mathrm{BF}_G^{\mathrm{raw}}$). A single positive real number summarising a group's evidential contribution, derived from its member nodes' individual Bayes factors. Its computation is developed in Section 4. The word "raw" indicates that no dependency correction has yet been applied.

Effective group Bayes factor ($\mathrm{BF}_G^{\mathrm{eff}}$). The result of applying all incoming dependency corrections to $\mathrm{BF}_G^{\mathrm{raw}}$. In the absence of any declared links involving G, we have $\mathrm{BF}_G^{\mathrm{eff}} = \mathrm{BF}_G^{\mathrm{raw}}$. The effective Bayes factor is the quantity that ultimately feeds back into the Truth Seeker computation through the per-node multiplier redistribution.

Per-node multiplier (m_v). For each node v that participates in at least one dependency group, a real-valued multiplier that will be applied to v's odds during the Truth Seeker recursive computation. Its derivation from the ratio $\mathrm{BF}_G^{\mathrm{eff}}/\mathrm{BF}_G^{\mathrm{raw}}$ is developed in Section 9.

2.3 Notation

Throughout this document, we adopt the following conventions in addition to those inherited from the *Numerical Assessment* document:

- G, G_A, G_B: dependency groups (equivalently, their membership sets).
- $\mathrm{BF}_G^{\mathrm{raw}}, \mathrm{BF}_G^{\mathrm{eff}}$: raw and effective Bayes factors of group G.

- π_v: the Truth Seeker computed index of node v (a probability in $[0, 1]$).
- log: natural logarithm, throughout. This convention makes log-Bayes-factors comparable to information-theoretic weight of evidence in nats. Base-2 conversion is straightforward if bits are required.

3 The Problem Dependencies Correct

The multiplicative aggregation of Bayes factors developed in the *Numerical Assessment* document is a direct consequence of Bayes' theorem under conditional independence. Recall Equation 7 of that document:

$$\mathrm{BF}(D_1, \ldots, D_k) = \prod_{i=1}^{k} \mathrm{BF}(D_i), \qquad \text{under } D_i \perp\!\!\!\perp D_j \mid H \text{ for all } i \neq j. \tag{2}$$

When conditional independence fails, this factorisation is invalid. The joint Bayes factor is no longer the product of the marginals: it is smaller in absolute value (for positively correlated evidence, which double-counts) or, in less common structures, larger. The multiplicative aggregation, applied naively, systematically misstates the strength of the compounded evidence.

3.1 Two structural violations

Two structural patterns are the practical concern in argument mapping.

Common-cause structures. Two or more sub-arguments beneath the same parent draw on overlapping evidence. The classical example is the "multiple witnesses" scenario in which several witnesses corroborate the same claim, but their testimonies trace back to a single conversation, document, or observation. Treated as independent, each testimony would contribute a full Bayes factor; treated correctly, the shared portion of their evidence must be counted only once. Under naive multiplication, five witnesses each contributing $\mathrm{BF} = 4$ would compound to a joint $\mathrm{BF} = 1024$. Under correct treatment reflecting a common source, the joint BF may be closer to 4 or 8 — three orders of magnitude smaller.

Causal chain structures. The truth of one argument is a precondition for another. The second argument cannot legitimately contribute evidence independently of the first: if the first fails, the second's evidential force is not merely reduced but is potentially undefined. Naive multiplication treats them as independent contributors and overstates the compounded evidence. This is a rarer structure in argument trees than common-cause but does occur — for example, a legal argument in which one sub-claim establishes standing and another establishes damages: the damages argument only bears on the case if the standing argument succeeds.

3.2 A worked example of double-counting

Consider a Truth Seeker Central Argument with Prior Probability $P_{\mathrm{CA}} = 0.10$, updated by three Main Arguments M_1, M_2, M_3, each of type *support*, each with Importance $I = 50$ and computed truth level $\pi = 0.80$. Under Balanced Argument Style ($\sigma = 0.67$), each MA contributes:

$$\mathrm{BF}_{\max}(M_i) = 2^{50/19.95} \approx 5.68,$$
$$\mathrm{BF}_{\mathrm{act}}(M_i) = 5.68^{0.80} \approx 4.00.$$

Multiplied against the prior odds $\Omega(P_{\mathrm{CA}}) = 0.111$:

$$\text{posterior odds} = 0.111 \times 4.00^3 \approx 7.11, \qquad \pi_{\mathrm{CA}} \approx 0.877.$$

The tree reports the Central Argument at 87.7% probability. This is the correct answer if the three MAs are genuinely independent lines of evidence.

Suppose, however, that all three MAs are versions of the same underlying claim, drawn from a single source. The joint evidential contribution should be closer to a single Bayes factor of 4 than to $4^3 = 64$. The correct posterior odds would be $0.111 \times 4 = 0.444$, yielding $\pi_{\text{CA}} \approx 0.308$. The naive computation is off by a factor of nearly three in the posterior probability, and by nearly a factor of 16 in the posterior odds.

The Dependency system is the user's mechanism for signalling that the three MAs share evidence, and for arranging that the assessment engine's output moves from the naive 0.877 toward the corrected value. It does not guarantee arrival at exactly 0.308 — as we shall see, the correction is a parameterised heuristic — but it moves the answer in the correct direction with a strength the user controls through the factor parameter.

3.3 The user-facing warning

The application detects one class of likely double-counting proactively. When a user attempts to create or edit a link between two groups whose membership sets have substantial overlap, the following warning is displayed if the declared factor is below a threshold derived from the overlap ratio:

> *Groups share {percent}% of nodes ({shared}/{total}). Factor {factor} may cause double-counting. Recommended minimum: {min}.*

The recommended minimum factor is $0.9 \times r$, where r is the overlap ratio of Equation (1). The threshold is empirical: it reflects the design position that a factor at least 90% of the structural overlap ratio is required to prevent significant double-counting, and any lesser factor is at risk of leaving too much of the shared evidence uncorrected. The same constant 0.9 appears as the auto-adjustment multiplier in the effective-factor computation (Section 7); the two mechanisms are complementary — one warns the user at input time, the other silently boosts the effective correction at compute time — and their consistency is intentional.

The warning is advisory rather than blocking: a user who understands their tree may legitimately declare a lower factor and dismiss the warning. The threshold is not a hard rule of Bayesian reasoning but an engineering judgment about where the risk of significant double-counting begins to matter.

4 The Group-Level Framework

We now develop the machinery by which the application converts individual nodes' Truth Seeker outputs into a single group-level Bayes factor, and by which it distinguishes raw from effective versions of that quantity.

4.1 The raw group Bayes factor

For a dependency group G with member nodes $\{v_1, \ldots, v_k\}$, the *raw group Bayes factor* BF_G^{raw} is defined as the geometric mean of the members' individual Bayes factors:

$$\text{BF}_G^{\text{raw}} = \left(\prod_{i=1}^{k} \text{BF}(v_i) \right)^{1/k} = \exp\left(\frac{1}{k} \sum_{i=1}^{k} \log \text{BF}(v_i) \right). \tag{3}$$

The second form is how the code computes it (as an arithmetic mean in log-space, then exponentiating), for numerical stability. The two forms are algebraically identical.

The special case of an empty group is defined by convention as $\text{BF}_G^{\text{raw}} = 1$, corresponding to no evidential contribution.

4.2 The per-node Bayes factor

The individual node Bayes factor $\mathrm{BF}(v)$ used in Equation (3) is derived from the node's Truth Seeker computed index according to the following procedure. Let $\pi_v = v.\texttt{computedIndex}$ be the node's computed index following the most recent baseline Truth Seeker pass (before any dependency adjustment), and let $C_v = v.\texttt{confidence}$ be its user-declared Confidence.

Two cases arise:

- If $\pi_v = 0$ (the node has not yet been assessed), the fallback uses the node's declared Confidence directly:

$$\mathrm{BF}(v) = \frac{C_v^*}{1 - C_v^*}, \qquad C_v^* = \max(0.01, \min(0.99, C_v)).$$

- Otherwise, the node's computed index is used, with the interpretation depending on the Central Argument's assessment model:
  - *Truth Seeker.* The computed index is already a probability in $[0, 1]$, and is used directly:

$$p_v = \max(0.01, \min(0.99, \pi_v)).$$

  - *Scorekeeper or Gatekeeper.* The computed index is a score in $[-1, +1]$, and is remapped to a probability-like value in $[0, 1]$ by

$$p_v = \max\left(0.01, \min\left(0.99, \frac{\pi_v + 1}{2}\right)\right).$$

  In either sub-case,

$$\mathrm{BF}(v) = \frac{p_v}{1 - p_v}. \tag{4}$$

The branching by assessment model is essential. In Truth Seeker mode, π_v is a probability, and its odds transform $\pi_v/(1 - \pi_v)$ is the appropriate Bayes-factor-like quantity. In WAM modes, $\pi_v \in [-1, +1]$ is a signed score, and interpreting it directly as a probability would place the neutral point at zero rather than at one-half; the remap $(\pi_v + 1)/2$ places the neutral WAM score of zero at the neutral probability of 0.5, at which point the odds are 1 and the node contributes no update.

The clamps to $[0.01, 0.99]$ are defensive: they prevent the odds transformation from diverging at the boundary and mirror the same clamps applied elsewhere in the Truth Seeker engine (see *Numerical Assessment* Section 4.7).

4.3 Why geometric mean

Equation (3) uses the geometric mean of the member Bayes factors. This choice merits explanation, since a Bayesian purist might expect the product instead — after all, independent Bayes factors combine multiplicatively under Bayes' theorem.

The geometric-mean choice is deliberate and rests on three considerations.

The group represents a single conceptual claim, not k pieces of evidence. A dependency group is the user's declaration that a set of nodes should be treated as evidentially connected. The intent of grouping is not to enumerate all evidence at once, but to summarise a cluster of related arguments so that dependency relationships can be applied at the cluster level. If the product were used, the mere act of grouping k nodes would replace their individual contributions with a k-fold compounded contribution — vastly overstating the group's collective strength. The geometric mean instead treats the group as if it were a single "representative" node contributing a typical Bayes factor.

The geometric mean is scale-equivariant. If every member's Bayes factor is multiplied by a constant α, so is the geometric mean. This property is exactly what allows the redistribution step (Section 9) to work correctly: applying a group-level correction ratio $\mathrm{BF}_G^{\mathrm{eff}}/\mathrm{BF}_G^{\mathrm{raw}}$ to each member's odds preserves the group-level relationship. No other symmetric mean (arithmetic, harmonic, quadratic) has this property in the required form. The choice of geometric mean is not merely convenient; it is what makes the per-node redistribution self-consistent.

Log-space aggregation is the natural setting for evidence. Following Good's information-theoretic framing (see *Numerical Assessment* Section 4.2), the weight of evidence is $\log \mathrm{BF}$, measured in nats or bits. The geometric mean of Bayes factors is the arithmetic mean of weights of evidence: it says "on average, each member contributes this much evidence." This is a natural summary statistic for a cluster, more directly interpretable than the product would be.

The reader should not, however, mistake this choice for a Bayesian derivation. The group Bayes factor of Equation (3) is a designed summary statistic, chosen for the properties described above. It does not correspond to any particular joint probabilistic model of the group's members. Section 11 discusses this openly among the modelling limitations of the system.

4.4 Raw versus effective Bayes factor

The raw group Bayes factor $\mathrm{BF}_G^{\mathrm{raw}}$ is the group's contribution under the assumption that its declared relationships to other groups are inactive — i.e., under the same conditional-independence assumption the Truth Seeker engine already makes. The *effective* group Bayes factor $\mathrm{BF}_G^{\mathrm{eff}}$ is the raw factor adjusted for those relationships. In the absence of any incoming links to G, we have $\mathrm{BF}_G^{\mathrm{eff}} = \mathrm{BF}_G^{\mathrm{raw}}$. When one or more incoming links are declared, the corrections of Sections 5 and 6 apply, potentially reducing $\mathrm{BF}_G^{\mathrm{eff}}$ below $\mathrm{BF}_G^{\mathrm{raw}}$ (the normal case, for correlated groups) or, for negative factors and certain causal structures, raising it above.

The ratio $\mathrm{BF}_G^{\mathrm{eff}}/\mathrm{BF}_G^{\mathrm{raw}}$ is what ultimately propagates back into the Truth Seeker computation through the per-node multiplier. That derivation is developed in Section 9; for the remainder of this section and Section 5 we focus on the correction step itself and how $\mathrm{BF}_G^{\mathrm{eff}}$ is computed from $\mathrm{BF}_G^{\mathrm{raw}}$ under the mutual dependency type.

4.5 A worked example of the raw group Bayes factor

To make Equation (3) concrete, we compute $\mathrm{BF}_G^{\mathrm{raw}}$ for several small groups. Assume Truth Seeker mode throughout, so each member's Bayes factor is $\pi_v/(1-\pi_v)$ where π_v is the member's computed index.

Members' computed indices	Per-node Bayes factors	$\mathrm{BF}_G^{\mathrm{raw}}$
$(0.5, 0.5, 0.5)$	$(1.0, 1.0, 1.0)$	1.000
$(0.7, 0.7, 0.7)$	$(2.333, 2.333, 2.333)$	2.333
$(0.8, 0.8, 0.8)$	$(4.000, 4.000, 4.000)$	4.000
$(0.9, 0.9, 0.9)$	$(9.000, 9.000, 9.000)$	9.000
$(0.8, 0.5, 0.6)$	$(4.000, 1.000, 1.500)$	1.817

Three observations. First, for a homogeneous group in which every member has the same computed index, the group Bayes factor equals any single member's Bayes factor. This is the geometric mean's fixed point on identical inputs — exactly the behaviour we want if the group is to represent a "typical member." Second, a group at equipoise ($\pi_v = 0.5$ for every member) yields $\mathrm{BF}_G^{\mathrm{raw}} = 1$: no evidential contribution. Third, in the heterogeneous case, the geometric mean lies between the smallest and largest member Bayes factors, weighted logarithmically. The value 1.817 for $(4.0, 1.0, 1.5)$ is the geometric mean $\sqrt[3]{4 \cdot 1 \cdot 1.5}$ — closer to the smaller members, in the sense that logarithmic averaging is influenced more by low outliers than by high ones.

5 The Mutual (Correlation) Correction

We now develop the first of the two correction forms. The mutual dependency type is applied when the user declares that two groups' evidences are correlated — typically through common-cause structure — and both groups' Bayes factors should be discounted accordingly.

5.1 The formula

Let G_A and G_B be two groups linked by a mutual dependency of factor f. The correction applied to G_A's Bayes factor, given G_B's, is:

$$\mathrm{BF}^{\mathrm{eff}}_{G_A} = \exp\Big(\log \mathrm{BF}^{\mathrm{raw}}_{G_A} - f \cdot \min\big(\log \mathrm{BF}^{\mathrm{raw}}_{G_A}, \log \mathrm{BF}^{\mathrm{raw}}_{G_B}\big)\Big). \tag{5}$$

By symmetry of the mutual type, an equivalent correction is applied simultaneously to G_B. Both groups' effective Bayes factors are computed against each other's raw values in this single formula.

Equation (5) is executed in code by the `applyDependencyFactor` method under the case `DependencyType.mutual`. The natural logarithm is used throughout; the code writes `log` and `exp` without base annotation.

5.2 Information-theoretic interpretation

To understand what Equation (5) does, it is helpful to work in the language of *weight of evidence*, following Good's terminology introduced in the *Numerical Assessment* document. The weight of evidence contributed by a Bayes factor BF is $w(\mathrm{BF}) = \log \mathrm{BF}$, measured in nats. Equation (5) can be rewritten in this notation as:

$$w(\mathrm{BF}^{\mathrm{eff}}_{G_A}) = w(\mathrm{BF}^{\mathrm{raw}}_{G_A}) - f \cdot \min\big(w(\mathrm{BF}^{\mathrm{raw}}_{G_A}), w(\mathrm{BF}^{\mathrm{raw}}_{G_B})\big). \tag{6}$$

That is: the effective weight of evidence contributed by G_A is G_A's raw weight of evidence, less a fraction f of a conservatively-estimated *shared* weight of evidence between G_A and G_B.

The choice of $\min(w_A, w_B)$ as the shared quantity has an information-theoretic justification. If we regard the two groups' evidential contributions as random variables whose informational content is measured in nats, then the mutual information between them — the amount of information that is shared rather than independent — cannot exceed the smaller of the two individual informational contents. This is a standard result: the mutual information $I(X;Y)$ between two random variables satisfies $I(X;Y) \leq \min\big(H(X), H(Y)\big)$, where H denotes entropy. The formula uses this upper bound as a conservative estimate of the shared weight of evidence, and applies a user-controlled fraction f of it as the discount.

We do not claim that the choice of min is derived from a specific mutual-information calculation over a specified joint distribution — no such distribution is available in the argument-mapping setting. Rather, the min is the largest defensible shared quantity: any value larger would risk overcorrecting, since neither group can share more evidence than the smaller of the two contributes in total. The factor f scales this bound continuously from zero (no discount) to one (full discount by the smaller group's evidence).

5.3 Limiting behaviour

Equation (5) has the following limiting behaviour, which the reader may verify by direct substitution:

- **$f = 0$:** $\mathrm{BF}^{\mathrm{eff}}_{G_A} = \mathrm{BF}^{\mathrm{raw}}_{G_A}$. No correction is applied. This is the correct limit for a declared-but-inactive link, and it means users can set up a link at zero factor as a placeholder without changing the numerical assessment.

- **$f = 1$, symmetric case ($\mathrm{BF}^{\mathrm{raw}}_{G_A} = \mathrm{BF}^{\mathrm{raw}}_{G_B}$):** $\mathrm{BF}^{\mathrm{eff}}_{G_A} = 1$. The correction removes the entirety of the group's evidential contribution. In log-space, this corresponds to subtracting a shared weight of evidence exactly equal to the total weight of evidence contributed. In the multiple-witnesses interpretation, this is the case where two groups' evidence is asserted to be fully redundant, and correct treatment reduces the joint contribution of both groups to the contribution of one representative group. Applied symmetrically to both groups' effective Bayes factors, the net effect on the tree is to preserve one full group-worth of evidence rather than two.
- **$f = 1$, asymmetric case ($\mathrm{BF}^{\mathrm{raw}}_{G_A} > \mathrm{BF}^{\mathrm{raw}}_{G_B}$):** The smaller group's weight of evidence is subtracted from both groups' contributions. The larger group retains a residual: $\mathrm{BF}^{\mathrm{eff}}_{G_A} = \mathrm{BF}^{\mathrm{raw}}_{G_A}/\mathrm{BF}^{\mathrm{raw}}_{G_B}$, corresponding to the excess weight of evidence contributed by G_A beyond what is shared with G_B. The smaller group's effective Bayes factor collapses to 1. This is the correct behaviour when the smaller group's evidence is a strict subset of the larger group's, with the larger group also contributing independent information.
- **Continuous behaviour in f:** The formula is monotonically decreasing in f (for positive factors), continuous on $[0, 1]$, and produces effective Bayes factors that are always positive. There are no discontinuities or singular values within the operating range.

5.4 Worked examples

The following table illustrates Equation (5) across representative factor values, for a symmetric pair with $\mathrm{BF}^{\mathrm{raw}}_{G_A} = \mathrm{BF}^{\mathrm{raw}}_{G_B} = 4.0$ (moderate evidence on both sides, $w \approx 1.386$ nats each):

Factor f	Effective BF	Effective w (nats)
0.00	4.000	1.386
0.25	2.828	1.040
0.50	2.000	0.693
0.75	1.414	0.347
1.00	1.000	0.000

The weight-of-evidence column decreases linearly in f, from the full raw value at $f = 0$ down to zero at $f = 1$. The Bayes factor column, being the exponential of the weight, decreases geometrically over the same range, halving with each $\Delta f = 0.5$.

For an asymmetric pair with $\mathrm{BF}^{\mathrm{raw}}_{G_A} = 8.0$ and $\mathrm{BF}^{\mathrm{raw}}_{G_B} = 2.0$ (where G_B contributes the smaller weight of evidence, $w_B \approx 0.693$ nats):

Factor f	Effective BF_A	Effective BF_B
0.00	8.000	2.000
0.50	5.657	1.414
1.00	4.000	1.000

At $f = 1$, the smaller group G_B's effective contribution collapses to unity, and G_A's excess weight of evidence over G_B's ($\log 8 - \log 2 = \log 4$) remains as G_A's effective contribution. This is the limiting case of "the smaller group's evidence is fully shared with the larger."

5.5 Negative factors: the compensatory case

The formula of Equation (5) admits negative factors, and its behaviour under negative f is well-defined and internally consistent, though it corresponds to a rarer structural situation than the positive case. If $f < 0$, the correction becomes:

$$w(\mathrm{BF}^{\mathrm{eff}}_{G_A}) = w(\mathrm{BF}^{\mathrm{raw}}_{G_A}) + |f| \cdot \min\bigl(w(\mathrm{BF}^{\mathrm{raw}}_{G_A}), w(\mathrm{BF}^{\mathrm{raw}}_{G_B})\bigr),$$

that is, the shared weight of evidence is *added* rather than subtracted, amplifying rather than discounting the group's effective contribution. The user-facing interpretation of this case is compensatory: the two groups' evidences are treated as jointly reinforcing beyond their individual contributions, as though one group's presence bolsters the other's credibility.

The negative-factor mutual case is offered by the application for user flexibility rather than out of a strong Bayesian derivation. Its principal legitimate use is to represent situations in which two groups' evidences are structurally anti-correlated — for example, two rival explanations of a phenomenon that cannot both be true, so establishing one weakens confidence in an already-weak alternative. Users who do not have such a structure in mind are advised to leave the factor non-negative. The user interface's helper text "`+1.0 = reinforcing, 0 = independent, -1.0 = compensatory`" names the extreme and lets the user navigate to the appropriate region of the range.

The causal (*A affects B*) correction, developed in Section 6, has a distinct interpretation of the negative-factor case that is generally more useful in practice. Users seeking to represent amplification effects are usually better served by the causal type at a negative factor than by the mutual type at a negative factor.

6 The Causal (*A affects B*) Correction

The second correction form applies when the user declares that one group's evidential contribution modulates another's rather than merely overlapping with it. Structurally, this represents a chain relationship: the truth of the source group's evidence is a precondition, precursor, or gateway for the target group's evidence being counted at full strength.

6.1 The formula

Let G_A be the source group and G_B be the target of a causal (*A affects B*) link with factor f. The correction applied to G_B's Bayes factor, given G_A's, is:

$$\mathrm{BF}^{\mathrm{eff}}_{G_B} = \left(\mathrm{BF}^{\mathrm{raw}}_{G_B}\right)^{1 - f\cdot\pi(G_A)}, \tag{7}$$

where $\pi(G_A)$ is the "truth level" associated with the source group's Bayes factor, defined by the odds-to-probability conversion:

$$\pi(G_A) = \frac{\mathrm{BF}^{\mathrm{eff}}_{G_A}}{1 + \mathrm{BF}^{\mathrm{eff}}_{G_A}}. \tag{8}$$

Equation (7) is executed in code by the `applyDependencyFactor` method under the case `DependencyType.aAffectsB`. Note that the source's *effective* Bayes factor is used in Equation (8), not its raw value; this couples the two groups together and is the reason an iterative procedure (Section 8) is required when links form chains or cycles.

Equation (7) is asymmetric by construction. Only the target group G_B's Bayes factor is modified; G_A's is unaffected by this link. This is the mathematical realisation of the "A affects B, but B does not affect A" semantics that the causal type expresses.

6.2 Interpretation

The dampening exponent $1 - f \cdot \pi(G_A)$ lies in $[1 - f, 1]$ for $\pi(G_A) \in [0, 1]$ and $f \in [0, 1]$. When the source group has high truth level and the factor is high, the exponent is small, and the target's Bayes factor is compressed toward 1: the causal source, being strongly established, has largely explained away the target's evidential force. When the source's truth level is low, the exponent stays close to 1, and the target's Bayes factor is little changed: an unestablished source cannot displace the target's independent contribution.

The multiplicative form $\left(\mathrm{BF}^{\mathrm{raw}}_{G_B}\right)^{1-f\pi(G_A)}$ works cleanly in log-space:

$$w(\mathrm{BF}^{\mathrm{eff}}_{G_B}) = (1 - f \cdot \pi(G_A)) \cdot w(\mathrm{BF}^{\mathrm{raw}}_{G_B}). \tag{9}$$

The target's weight of evidence is scaled by a factor that ranges from 1 (no scaling) at $f = 0$ or $\pi(G_A) = 0$, down to $1 - f$ at $\pi(G_A) = 1$. Fully-established source ($\pi(G_A) = 1$) at maximum factor ($f = 1$) reduces the target's weight of evidence to zero, so its Bayes factor becomes 1.

6.3 Limiting behaviour

Equation (7) has the following limiting behaviour:

- **$f = 0$: $\mathrm{BF}^{\mathrm{eff}}_{G_B} = \mathrm{BF}^{\mathrm{raw}}_{G_B}$.** No correction. Placeholder links at zero factor do not change results.
- **$\pi(G_A) = 0$: $\mathrm{BF}^{\mathrm{eff}}_{G_B} = \mathrm{BF}^{\mathrm{raw}}_{G_B}$.** An unestablished source does not affect the target, regardless of the declared factor. This reflects the modelling intent: causal relationships operate only when the causing group has evidential force.
- **$f = 1$, $\pi(G_A) = 1$: $\mathrm{BF}^{\mathrm{eff}}_{G_B} = 1$.** A fully-established source at maximum factor collapses the target's contribution to unity.
- **Continuous behaviour in both f and $\pi(G_A)$:** No discontinuities; the effective Bayes factor is a smooth function of both.

6.4 Worked examples

For a target group with raw BF = 5.0, the effective Bayes factor as a function of source Bayes factor and declared positive factor:

Source BF	Factor f		
	0.00	0.50	1.00
2.0 ($\pi = 0.667$)	5.000	2.924	1.710
5.0 ($\pi = 0.833$)	5.000	2.557	1.308
10.0 ($\pi = 0.909$)	5.000	2.406	1.158

Reading the middle column: a moderate factor of $f = 0.5$ dampens the target from BF = 5.0 to somewhere between 2.4 and 2.9, depending on how strongly the source is established. The stronger the source, the more thoroughly it explains the target's evidence away.

6.5 Negative factors: the amplification case

The formula of Equation (7) admits negative factors, and its behaviour under negative f is more compelling as a modelling choice than the negative-mutual case discussed in Section 5.5. If $f < 0$ and $\pi(G_A) > 0$, the exponent $1 - f\pi(G_A)$ becomes greater than 1, and the target's Bayes factor is raised to a power greater than 1: it is amplified rather than dampened.

The interpretation is: a well-established source group increases confidence in the target's evidence rather than displacing it. This corresponds to structural situations in which the two groups' evidences are mutually reinforcing — one group's establishment strengthens the interpretive weight of the other. For example, evidence of motive (source group) may amplify the interpretive weight of evidence of opportunity (target group), because the two together suggest deliberate action more strongly than either alone.

Factor f	Effective BF_B (target raw = 5.0, source raw = 5.0)	Behaviour
-1.00	19.118	amplifies
-0.50	9.777	amplifies
0.00	5.000	neutral
$+0.50$	2.557	dampens
$+1.00$	1.308	dampens

The behaviour is smooth and monotonic across the range: the effective Bayes factor increases as factor decreases (with source truth held constant), passing through the raw value at $f = 0$ and continuing to amplification for negative values. This is the compensatory behaviour promised by the user-facing helper text.

7 Structural Overlap Auto-Adjustment

Between the user's declaration of a link's factor and the correction formulas of Sections 5 and 6, an intermediate step transforms the declared factor into an *effective* factor used in the actual correction. This step is the structural overlap auto-adjustment.

7.1 The formula

Let G_A and G_B be two groups linked by any dependency type, and let f_{decl} be the factor declared by the user. The overlap ratio $r = |G_A \cap G_B|/|G_A \cup G_B|$ is computed from the group membership sets (Equation (1)). The effective factor used in the correction formulas is:

$$f_{\mathrm{eff}} = \mathrm{sign}(f_{\mathrm{decl}}) \cdot |f_{\mathrm{decl}}|^{1/(1+r)} . \tag{10}$$

The result is clamped to $[-1, +1]$ before being passed to the correction formulas, though for $f_{\mathrm{decl}} \in [-1, +1]$ this clamp is defensive only. The special cases $r = 0$ (disjoint groups) and $f_{\mathrm{decl}} = 0$ short-circuit to $f_{\mathrm{eff}} = f_{\mathrm{decl}}$.

The code uses a general form with an internal constant `maxBoostStrength = 2.0`: the exponent is $1/(1 + (M - 1)r)$ with $M = 2$. This simplifies to $1/(1 + r)$ as written above. The constant is available for future tuning; at present, no user control over M is exposed.

7.2 Rationale

The auto-adjustment reflects a specific design position: when two groups share substantial structural overlap, the user's declared factor should be treated as stronger than face value, because the shared membership itself constitutes evidence of correlation the user has not explicitly quantified.

Consider two groups with 80% overlap ($r = 0.8$). If the user declares a factor of 0.5, this can be read two ways. Under one reading, the user has assessed the dependency at strength 0.5 and expects that strength to be applied. Under another, the user has assessed the dependency at strength 0.5 *on top of* the correlation already implied by the shared membership, which should push the effective correction higher. The auto-adjustment adopts the latter reading. It transforms $f_{\mathrm{decl}} = 0.5$ at $r = 0.8$ into $f_{\mathrm{eff}} \approx 0.680$, applying a modestly stronger correction than the declared value would suggest under the first reading.

The exponent $1/(1 + r)$ is designed so that the transformation:

- Reduces to the identity at $r = 0$: no adjustment for disjoint groups.

- Pulls $|f_{\mathrm{decl}}|$ toward 1 as $r \to 1$: fully-overlapping groups have their declared factor treated as approaching a full correction.
- Preserves sign: negative factors remain negative, positive factors remain positive.
- Is monotonic in both r and $|f_{\mathrm{decl}}|$: increasing either input increases $|f_{\mathrm{eff}}|$.

The choice of the specific exponential form is engineering: several curves have the same limits and monotonicity, and the one adopted was selected for its numerical smoothness and its ease of tuning through a single constant M. The specific numerical behaviour at intermediate values (see Table 1) is a design consequence, not a derivation.

Table 1: Effective factor as a function of overlap ratio, at declared factor $f_{\mathrm{decl}} = 0.5$, for pairs of groups with 5 members each.

Shared nodes	Union size	Overlap ratio r	Effective factor f_{eff}
0	10	0.000	0.500
1	9	0.111	0.536
2	8	0.250	0.574
3	7	0.429	0.616
4	6	0.667	0.660
5	5	1.000	0.707

7.3 Interaction with the input-time warning

The auto-adjustment operates silently: users do not see the effective factor f_{eff} in the user interface — only the declared factor f_{decl}. This is a deliberate design choice, complemented by the input-time warning discussed in Section 3. Users who would set a factor too low relative to structural overlap are warned at input time; users who accept a moderate factor while overlap is high get a modestly boosted correction through the auto-adjustment. Together, these two mechanisms defend against silent under-correction from opposite directions.

The reader should be aware that this silent transformation is a legitimate subject of critique. It embeds a design opinion — that user-declared factors should be interpreted *additively to structural overlap* rather than *inclusively of it* — into the mathematics without exposing that opinion in the interface. An honest audit is offered in Section 11.

8 Iterative Convergence

The corrections of Sections 5 and 6 operate on the effective, not raw, Bayes factors of the source groups. When a group has multiple incoming links, or when the dependency graph contains cycles, this creates a system of mutually-dependent equations that cannot be solved in a single pass. The application solves this system through fixed-point iteration.

8.1 The iterative procedure

Starting from the raw Bayes factors as an initial guess ($\mathrm{BF}_G^{\mathrm{eff},(0)} = \mathrm{BF}_G^{\mathrm{raw}}$ for all groups), the procedure updates all groups simultaneously at each iteration:

$$\mathrm{BF}_G^{\mathrm{eff},(t+1)} = \left(\mathrm{BF}_G^{\mathrm{eff}} \text{ under all incoming links, applied to } \mathrm{BF}_G^{\mathrm{raw}}\right) \quad \text{using } \mathrm{BF}_H^{\mathrm{eff},(t)} \text{ for source groups } H. \tag{11}$$

Each iteration applies every group's incoming corrections in sequence to its raw Bayes factor, using the previous iteration's effective values as source inputs. The procedure terminates when the maximum change across all groups falls below the convergence threshold:

$$\max_G \left|\mathrm{BF}_G^{\mathrm{eff},(t+1)} - \mathrm{BF}_G^{\mathrm{eff},(t)}\right| < \varepsilon, \qquad \varepsilon = 10^{-3},$$

or when the iteration count reaches its cap of 10, whichever occurs first.

8.2 Why iteration is necessary

The user interface prevents the creation of two-group direct cycles ($G_A \to G_B$ and $G_B \to G_A$ as separate links) but does not, by design, prevent longer cycles. A three-group causal chain $G_A \to G_B \to G_C \to G_A$ is a legitimate structure a user may wish to declare, and the application accepts such declarations. Single-pass computation would give different answers depending on the order in which groups were processed. The fixed-point iteration is order-independent (up to convergence) and is the standard resolution.

Mutual links also require iteration in principle, though their symmetric structure makes single-pass computation give an approximately correct answer. The iteration cleans up the residual asymmetry introduced by processing groups in sequence within each pass.

8.3 Convergence behaviour and failure modes

The iterative procedure has the following practical behaviour, established through numerical experimentation with representative dependency structures:

- **Non-cyclic structures** (trees or DAGs of links) converge in one to two iterations, typically well below the threshold on the second pass. This is the common case.
- **Symmetric mutual chains** of moderate length (three to five groups) converge in five to ten iterations, sometimes approaching the iteration cap. A three-group chain with all-equal raw Bayes factors and factor 0.5 requires the full ten iterations to reach the threshold, with the final change being of order 10^{-3}.
- **Symmetric causal cycles** of three or more groups converge in three to five iterations, more rapidly than the mutual chain case above.
- **Failure to converge within the cap** is possible in principle for pathological structures — for example, a long chain with factors close to 1. When the cap is reached without convergence, the procedure returns the final iteration's values without warning. The reported numbers in such cases should be understood as a heuristic approximation rather than a fixed-point solution.

8.4 Worked example: a three-group symmetric case

Consider three groups A, B, C, all with $\mathrm{BF}^{\mathrm{raw}} = 4.0$, connected by mutual links: $A \leftrightarrow B$ with $f = 0.5$, and $B \leftrightarrow C$ with $f = 0.5$. This is a chain, not a cycle, but B is subject to two incoming corrections and A and C each to one.

Iteration	$\mathrm{BF}^{\mathrm{eff}}_A$	$\mathrm{BF}^{\mathrm{eff}}_B$	$\mathrm{BF}^{\mathrm{eff}}_C$	Max change
init	4.0000	4.0000	4.0000	—
1	2.0000	1.4142	2.0000	2.586
2	3.3636	2.0000	3.3636	1.364
3	2.8284	1.4768	2.8284	0.535
4	3.2915	1.5422	3.2915	0.463
5	3.2210	1.4848	3.2210	0.071
6	3.2826	1.4929	3.2826	0.062
7	3.2737	1.4859	3.2737	0.009
8	3.2815	1.4869	3.2815	0.008
9	3.2804	1.4860	3.2804	0.001
10	3.2814	1.4861	3.2814	0.000972

The procedure converges just at the cap: the maximum change on iteration 10 falls below 10^{-3} with no headroom. This is atypical for well-conditioned problems but not unrepresentative of what a user with substantive mutual structure may see. The final values ($\mathrm{BF}^{\mathrm{eff}}_A = \mathrm{BF}^{\mathrm{eff}}_C \approx 3.28$, $\mathrm{BF}^{\mathrm{eff}}_B \approx 1.49$) are the fixed-point solution up to the threshold: B receives corrections from both sides and is more heavily discounted than A or C, as one would expect from its position in the chain.

For contrast, a fully-cyclic three-group causal case ($A \to B \to C \to A$, all factors 0.5) converges in 5 iterations from the same starting point, terminating with all three groups at $\mathrm{BF}^{\mathrm{eff}} \approx 2.446$. Cyclic causal structures converge faster than mutual chains because the dampening operation is contractive in a stronger sense.

9 Per-Node Multiplier Distribution

The final step of the dependency-adjustment pipeline converts group-level corrections into per-node modifications that the Truth Seeker recursive computation can apply.

9.1 The multiplier formula

For each group G with $\mathrm{BF}^{\mathrm{raw}}_G > 0$, the group-level correction ratio is:

$$m_G = \frac{\mathrm{BF}^{\mathrm{eff}}_G}{\mathrm{BF}^{\mathrm{raw}}_G}. \tag{12}$$

This ratio is applied to each member node of G as a per-node multiplier. When a node v belongs to multiple groups $G_1, \ldots, G_j$, the applied multiplier is the minimum:

$$m_v = \min_{G \ni v} m_G. \tag{13}$$

The minimum is used out of conservatism: for a node participating in multiple corrections, the strongest applicable discount is preferred. This is a design choice that could reasonably have been made differently (see Section 11).

The per-node multiplier is then applied to the node's odds during the tree-service adjustment pass. For a node v with baseline Truth Seeker computed index π_v:

$$\pi_v^{\mathrm{adjusted}} = \frac{m_v \cdot \Omega(\pi_v)}{1 + m_v \cdot \Omega(\pi_v)}, \qquad \Omega(\pi_v) = \frac{\pi_v}{1 - \pi_v}. \tag{14}$$

The odds transformation, multiplication by m_v, and inverse transformation together implement the standard Bayesian update rule: post-adjustment odds equal pre-adjustment odds times the multiplier. This is the same odds-form update the Truth Seeker engine uses throughout, applied here with a multiplier derived from the group-level correction rather than from a child argument's Bayes factor.

9.2 Consistency with the geometric-mean choice

The per-node redistribution of Equation (13) is self-consistent with the geometric-mean group aggregation of Equation (3). Specifically: if every member's individual Bayes factor is multiplied by the same constant m, the geometric mean is also multiplied by m. Applying m_G uniformly to every member of G therefore transforms $\mathrm{BF}_G^{\mathrm{raw}}$ into $m_G \cdot \mathrm{BF}_G^{\mathrm{raw}} = \mathrm{BF}_G^{\mathrm{eff}}$, exactly the intended group-level correction.

This property is the reason Section 4.3 emphasised scale-equivariance as one of the three justifications for the geometric-mean choice. Under any other aggregation — arithmetic mean, harmonic mean, product — the per-node redistribution would not correspond exactly to the group-level correction, and the entire pipeline would have an internal inconsistency. The geometric mean is what makes the framework close on itself.

9.3 Recomputation after adjustment

The application of Equation (14) modifies the computed indices of dependency-participating nodes. Because Truth Seeker propagates computed indices bottom-up, this modification may in principle affect the computed indices of the participating nodes' ancestors. The tree-service adjustment procedure, therefore, performs a two-step pass: first, all affected leaf-like nodes have their computed indices modified by Equation (14); second, all ancestor nodes are recomputed bottom-up through the Truth Seeker recursive procedure, using the modified leaf values.

This second pass is where the dependency system's corrections actually reach the Central Argument's computed value. The multiplier itself is not applied to the CA directly; rather, the CA's computation runs afresh over a tree in which its descendant leaves have been adjusted. The dependency system therefore respects the Truth Seeker aggregation rules throughout, and its output is a Truth Seeker computation over a modified tree, not a bolt-on adjustment to the CA's number.

10 A Complete Worked Example

To make the full pipeline concrete, we work through a small example end-to-end. The setup is chosen to illustrate the double-counting problem from Section 3.1 and the correction the mutual dependency provides.

10.1 Setup

The Central Argument is in Truth Seeker mode with Prior Probability $P_{\mathrm{CA}} = 0.10$ and Argument Style $\sigma = 0.67$ (Balanced). It has three Main Arguments, all of type *support*, all leaves, all with Importance $I = 50$ and Confidence $C = 0.8$. Under the Truth Seeker rules of Equation 5 of the *Numerical Assessment* document, each MA's actual Bayes factor is:

$$\mathrm{BF}_{\mathrm{act}}(M_i) = \left(2^{50/19.95}\right)^{0.8} = 5.681^{0.8} \approx 4.014.$$

10.2 Case 1: Naive (no dependencies)

The naive Truth Seeker computation gives:

$$\begin{aligned}\Omega(P_{\mathrm{CA}}) &= 0.111,\\ \text{posterior odds} &= 0.111 \times 4.014^3 = 7.186,\\ \pi_{\mathrm{CA}} &= 0.8778.\end{aligned}$$

This is the result reported when the user has not declared any dependencies.

10.3 Case 2: Grouping alone, no links

Suppose the user places M_1 and M_2 into a single dependency group G_1, without declaring any link to another group. The dependency-adjustment engine computes $\mathrm{BF}_{G_1}^{\mathrm{raw}}$ from the members' per-node Bayes factors. Each member's computed index is $C = 0.8$, so its per-node Bayes factor (Equation (4)) is $0.8/0.2 = 4.0$. The geometric mean of two identical values is the value itself:

$$\mathrm{BF}_{G_1}^{\mathrm{raw}} = \sqrt{4.0 \cdot 4.0} = 4.0.$$

With no links, $\mathrm{BF}_{G_1}^{\mathrm{eff}} = \mathrm{BF}_{G_1}^{\mathrm{raw}} = 4.0$, the multiplier is $m_{G_1} = 1.0$, and every member node's adjustment is the identity. The CA's computation is identical to Case 1: $\pi_{\mathrm{CA}} = 0.8778$.

This is the expected behaviour and is stated here for completeness. The mere act of grouping nodes does not change the numerical assessment; only declared links between groups do so.

10.4 Case 3: Mutual link between two singleton groups

Now suppose the user recognises that M_1 and M_2 share a common source of evidence and wants to correct for the double-counting. One way to encode this is to place M_1 in G_1 and M_2 in G_2, then declare a mutual link between G_1 and G_2 with factor $f = 0.5$.

Each group has one member, so each group's raw Bayes factor is that member's per-node Bayes factor:

$$\mathrm{BF}_{G_1}^{\mathrm{raw}} = \mathrm{BF}_{G_2}^{\mathrm{raw}} = 4.0.$$

The overlap ratio between G_1 and G_2 is zero (they share no members), so the effective factor equals the declared: $f_{\mathrm{eff}} = 0.5$.

Applying the mutual correction (Equation (5)) to both groups:

$$\begin{aligned}\mathrm{BF}_{G_1}^{\mathrm{eff}} &= \exp\bigl(\log 4 - 0.5 \cdot \min(\log 4, \log 4)\bigr) = \exp(0.693) = 2.000,\\ \mathrm{BF}_{G_2}^{\mathrm{eff}} &= 2.000 \quad \text{(by symmetry).}\end{aligned}$$

The multipliers are:

$$m_{G_1} = m_{G_2} = \frac{2.000}{4.000} = 0.5.$$

Each participating node's odds is multiplied by 0.5. Applied to the leaves' computed indices via Equation (14):

$$\begin{aligned}\pi_{M_1}^{\mathrm{adj}} &= \frac{0.5 \cdot 4.0}{1 + 0.5 \cdot 4.0} = 0.667,\\ \pi_{M_2}^{\mathrm{adj}} &= 0.667 \quad \text{(by symmetry).}\end{aligned}$$

M_3 is not in any group and is not adjusted.

The Truth Seeker engine now recomputes the CA using these modified leaf values. The MAs' new actual Bayes factors are:

$$\begin{aligned} \mathrm{BF}_{\mathrm{act}}(M_1) &= 5.681^{0.667} = 3.184, \\ \mathrm{BF}_{\mathrm{act}}(M_2) &= 3.184, \\ \mathrm{BF}_{\mathrm{act}}(M_3) &= 4.014 \quad \text{(unchanged).} \end{aligned}$$

And the CA's posterior:

$$\begin{aligned} \text{posterior odds} &= 0.111 \times 3.184 \times 3.184 \times 4.014 = 4.521, \\ \pi_{\mathrm{CA}} &= 0.8189. \end{aligned}$$

10.5 Discussion

The user's declaration of a mutual dependency with factor 0.5 has moved the CA's computed value from 0.8778 to 0.8189, a difference of 5.9 percentage points. This is the numerical effect of the correction: not a wholesale reset of the tree's result, but a proportionate discount reflecting the user's assertion that two of the three MAs share evidence.

If the user believed the sharing was more thorough, they could declare a higher factor (say $f = 1.0$), which would produce a stronger discount. If they believed it was minor, they could declare a lower factor ($f = 0.2$), producing a milder discount. If they had not declared the mutual link at all, the tree would report the naive value of 0.8778 and the double-counting would remain uncorrected.

The system does not enforce a particular answer; it makes the correction machinery available and asks the user to declare the strength of the relationships they perceive. The mathematics behind that declaration is what this document has developed.

11 Modelling Choices and Limitations

This section addresses, in a single place, the significant modelling choices made throughout the dependency system and the limitations they impose. Following the structure of the *Numerical Assessment* document's Section 8, the intent is to make explicit what a technically-literate reader could otherwise assemble from close reading of the preceding sections.

11.1 The corrections are heuristics, not derivations

Neither Equation (5) (mutual) nor Equation (7) (causal) is derived from a specified joint probability distribution over the two groups' evidences. Both are engineering formulas motivated by defensible intuitions — the mutual correction by an information-theoretic upper bound on shared information, the causal correction by a modulation semantics in which established causes explain away downstream evidence — but neither is unique to those intuitions. Alternative forms exist in the argument-mapping literature and in the broader Bayesian-network literature. The forms adopted here were chosen for their limiting behaviour, their computational tractability, and their smoothness in the user-controlled parameters.

The reader who is comfortable with this framing should understand the dependency system's output as *Bayes-flavoured corrections applied consistently within a well-specified framework*, not as *the unique correct Bayesian answer under any particular graphical model*. Users seeking the latter should build explicit Bayesian networks rather than argument trees; the two tools serve different purposes.

11.2 The absence of a joint distribution

The most fundamental modelling limitation is that the dependency system does not represent an explicit joint distribution over the truth values of the groups' evidences. A rigorous Bayesian treatment of correlated or causally-linked evidence would specify $P(D_A, D_B \mid H)$ and $P(D_A, D_B \mid \neg H)$ — the joint likelihoods under the hypothesis and its negation — and compute the joint Bayes factor from these. The dependency system does not require the user to specify these joints, and its corrections are functions of user-declared factors rather than of underlying probability values.

This is a legitimate trade-off. Requiring users to specify joint distributions would make the tool inaccessible to non-statisticians and would not, in practice, be more principled: users who cannot honestly assess a factor of 0.5 are unlikely to honestly assess a joint likelihood. The dependency system's honesty is in accepting a user-declared strength as its input, applying that strength through a documented correction formula, and reporting the result. The output is what the formula produces given the user's declared input; users seeking rigour beyond this should recognise that no available UI-driven tool for argument mapping provides fully-derived Bayesian correction for arbitrary dependency structures.

11.3 The silent overlap auto-adjustment

Section 7 describes a transformation of the user's declared factor into an effective factor, based on structural overlap of the groups' membership sets. This transformation happens silently: users see the declared factor in the interface, but the correction actually applied uses the auto-adjusted value.

An argument for this design: users who structure their groups to share nodes are implicitly asserting a correlation stronger than a factor of 0.5 (say) would represent on disjoint groups. The auto-adjustment corrects the interpretation to match the user's structural declaration.

An argument against: silent transformations of user inputs violate the general principle that a system should do what its interface says it does. A user who declares a factor of 0.5 has a reasonable expectation that the correction of Equation (5) is applied with $f = 0.5$, not with a boosted value. The auto-adjustment mechanism, while defensible, is not visible in the interface, and users cannot inspect the effective factor being used.

The application partially compensates for this by displaying the input-time warning (Section 3) when the declared factor is low relative to the overlap ratio. Users who ignore or dismiss the warning receive the auto-boosted correction anyway; users who follow the warning's advice by raising their factor receive a modestly smaller additional boost. Both mechanisms encode the same design intent, but only one is visible.

A stronger design would either expose the effective factor in the interface (adding cognitive load) or remove the silent auto-adjustment (relying entirely on the warning). The current design is a compromise; whether it is the correct one is a legitimate open question.

11.4 Cyclic dependency handling

The application does not detect cycles at input time and relies on the fixed-point iteration procedure of Section 8 to handle them at compute time. The iteration cap of 10 and the threshold of 10^{-3} are engineering constants, not theoretically-motivated bounds. Long chains, cycles with factors close to 1, or asymmetric structures can fail to converge within the cap. When this happens, the final iteration's values are returned without warning to the user.

A stronger design would either detect potentially-non-convergent structures at input time (requiring cycle analysis over the declared dependency graph) or report to the user when convergence has not been reached at compute time. Neither is currently done. Users of pathological structures should be aware that reported numbers in such cases are heuristic approximations.

11.5 Conservative aggregation across multi-group nodes

For nodes participating in multiple groups, Equation (13) applies the minimum multiplier. The rationale is conservative: apply the strongest applicable discount and let the user rely on the correction being at least as aggressive as any individual group would produce. But this is one of several defensible choices. Alternatives include:

- *Geometric mean of multipliers*: aggregate multi-group corrections as their geometric mean, mirroring the group-level aggregation choice. This would produce a milder combined correction.
- *Multiplicative combination*: multiply the multipliers, producing a substantially stronger combined correction (typically an over-correction, since multiple corrections may reflect the same underlying redundancy).
- *First-applied wins*: process groups in some order and apply only the first correction seen. Order-dependent and generally undesirable.

The minimum-multiplier choice is not the unique correct answer; it is a conservative choice that guards against user error at the cost of possibly under-correcting when multiple independent corrections apply. Users who structure their trees so that no node belongs to more than one group avoid this consideration entirely, and this is the recommended usage pattern.

11.6 What the system does not do

For completeness, the dependency system does not attempt:

- **Automatic dependency detection.** The system does not infer dependencies from the tree's structural properties. Users must declare all dependencies explicitly.
- **Transitivity checks.** If G_A affects G_B and G_B affects G_C, no automatic G_A-to-G_C dependency is inferred. Users who want the transitive effect must declare it separately.
- **Interaction with the WAM or Gatekeeper assessment models.** The dependency system is invoked from the tree-service adjustment pass, which currently applies only under the Truth Seeker model. WAM and Gatekeeper computations proceed without dependency correction, on the grounds that WAM's philosophical stance — weighted-average preference aggregation — does not require Bayesian double-counting correction. This is a reasonable design choice but is worth stating.
- **Learning from data.** User-declared factors are not calibrated against outcomes or refined by any statistical procedure. A user who consistently over-declares correlation will produce systematically over-corrected trees.
- **Cross-tree dependency awareness.** Dependencies declared in one tree do not affect any other tree the user may have open. Groups and links are tree-local.

None of these limitations is a bug; each reflects a design boundary. The reader who understands them is in a position to use the tool honestly.

12 Tooltip and UI Text Reference

The following are the authoritative user-facing strings adopted by the application for the dependency system, drawn from the English localisation file. Each string is quoted verbatim.

12.1 Dialog titles and prompts

Create Dependency (dialog title). "Create Dependency"

Edit Dependency (dialog title). "Edit Dependency"

Draw boundary (lasso prompt). "Draw boundary around nodes (tap-drag-release)"

Draw boundary to add (lasso prompt). "Draw boundary around nodes to add (tap-drag-release)"

Select target group. "Select which group to add to"

12.2 Field labels and helpers

Relationship Type. "Relationship Type"

Mutual (dependency type). "Mutual"

A affects B (dependency type). "{a} affects {b}" (with actual group labels substituted for the placeholders)

B affects A (dependency type). "{b} affects {a}"

Factor field label. "Factor (-1.0 to +1.0)"

Factor field helper text. "+1.0 = reinforcing, 0 = independent, -1.0 = compensatory"

12.3 Warnings and errors

Overlap double-count warning. "Groups share {percent}% of nodes ({shared}/{total}). Factor {factor} may cause double-counting. Recommended minimum: {min}"

Factor range error. "Must be between -1.0 and +1.0"

Cannot link to self. "Cannot link a group to itself"

Groups not found. "One or both groups not found"

Link already exists. "Link already exists: {groupA} {desc} {groupB}"

Cannot change to mutual. "Cannot change to Mutual: Separate reverse link exists ({target} → {source}). Delete it first or edit it to Mutual."

Cannot reverse direction. "Cannot reverse direction: Separate link exists ({desc}). Delete it first."

No nodes encircled (lasso error). "No nodes fully encircled. Try again."

12.4 Toggle and status

Calculation on (tooltip). "Dependency Calculation: ON"

Calculation off (tooltip). "Dependency Calculation: OFF"

Calculation enabled (confirmation). "Dependency calculations ENABLED"

Calculation disabled (confirmation). "Dependency calculations DISABLED"

12.5 Terminology reference

UI term	Formal meaning
Group ("Grp 1", "Grp 2", ...)	Membership set $G \subseteq \mathcal{N}$
Link	Directed or bidirectional relationship between two groups
Factor	Real value $f \in [-1.0, +1.0]$
Reinforcing / Independent / Compensatory	$f = +1.0$ / $f = 0$ / $f = -1.0$ semantic regions
Relationship Type	Categorical $\tau \in \{\text{mutual}, \text{aAffectsB}\}$ (persisted form)
Mutual	Bidirectional correlation correction
Affects (arrow)	Unidirectional causal modulation correction
Dependency Calculation	The global toggle enabling correction application
Overlap ratio	Jaccard coefficient of two groups' membership

13 References

The dependency system's mathematics draws on standard results in information theory, Bayesian networks, and the evidence-combination literature.

13.1 Information theory and weight of evidence

- Good, I. J. (1950). *Probability and the Weighing of Evidence*. Charles Griffin. The origin of the "weight of evidence" terminology used throughout this document. The mutual correction of Section 5 works directly in Good's log-Bayes-factor scale.
- Cover, T. M., and Thomas, J. A. (2006). *Elements of Information Theory*, 2nd ed. Wiley. The reference for mutual information and its bounds. The result that $I(X;Y) \leq \min(H(X), H(Y))$ underlies the choice of $\min(\log \text{BF}_A, \log \text{BF}_B)$ as the conservative shared-information bound in the mutual correction formula.

13.2 Bayesian networks and dependency structure

- Pearl, J. (1988). *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann. The reference work on graphical Bayesian models. Chapter 3's discussion of conditional independence is the assumption the dependency system relaxes; Chapter 4's treatment of belief propagation over graphs is the framework from which fixed-point iteration inherits its intuition.
- Pearl, J. (2000). *Causality: Models, Reasoning, and Inference*. Cambridge University Press. The reference for causal graphical models. The causal (*A affects B*) dependency type has structural echoes of Pearl's "do-calculus" interventions, though the specific dampening formula of Equation (7) is not derived from that framework.

13.3 Evidence combination and alternative frameworks

- Shafer, G. (1976). *A Mathematical Theory of Evidence*. Princeton University Press. The Dempster–Shafer framework offers an alternative approach to evidence combination that does not require full specification of prior probabilities and admits explicit representation of ignorance. The mutual correction of Section 5 has structural echoes of Dempster's combination rule, though the two are not equivalent.
- Kass, R. E., and Raftery, A. E. (1995). "Bayes factors." *Journal of the American Statistical Association*, 90(430), 773–795. A comprehensive review of Bayes factors that includes discussion of how correlated evidence should be combined.

13.4 The companion document

- *Numerical Assessment: A Technical Justification of the Mathematics Behind i.a.T.'s Truth Seeker, Scorekeeper, and Gatekeeper Modes*. The companion document, which establishes the Truth Seeker framework this document depends on. Particular references appear throughout to Section 4.2 (calibration of the Argument Style parameter), Section 4.6 (the conditional-independence assumption), and Section 8 (the modelling-audit template followed here).

This document constitutes the complete Dependency Mathematics reference.

This document, together with its companion Numerical Assessment *reference, constitutes the complete mathematical justification for the i.a.T. application's assessment engine.*