

One Metalogic, Many Logics: A Quantitative Pedagogical Framework for Proof-Assistant-Based Logic Teaching with LogiKEy

One Metalogic, Many Logics: A Quantitative Pedagogical Framework for Proof-Assistant-Based Logic Teaching with LogiKEy

Abstract

Teaching logic to mixed cohorts of computer science, mathematics, and philosophy students is difficult because each cohort brings different expectations and because the standard curriculum fragments into isolated courses on propositional, modal, epistemic, and deontic logic. The LogiKEy methodology addresses this by using classical higher-order logic (HOL) as a universal metalogic: object logics are encoded as semantical embeddings inside a single proof assistant such as Isabelle/HOL, so that students learn, experiment with, and compare logics in one environment. This paper complements the qualitative, example-driven curriculum design of LogiKEy with explicit combinatorial analysis of its graded example sequence — liars-and-truth-tellers puzzles, the Wise Men puzzle, Boolos's curious inference, Chisholm's paradox, and Gödel's ontological argument. For each stage we derive, with full arithmetic, the size of the relevant search or model space: 8 propositional assignments with exactly 1 consistent survivor, $2^{64} \approx 1.84 \times 10^{19}$ candidate modal accessibility relations, an epistemic model trajectory $8 \rightarrow 7 \rightarrow 6 \rightarrow 4$ under public announcement, a first-order interpretation space $2^{100} \approx 1.27 \times 10^{30}$, and a proof-length tower reaching $\approx 10^{19728}$. We further derive a labeled projection for course-time allocation under stated assumptions. We argue that these computed checkpoints, rather than anecdote, are what make each curricular transition pedagogically forced, and we identify failure modes, falsification conditions, and open questions for empirical validation.

1. Introduction

A logic classroom with only computer science students can assume comfort with formalism and an interest in verification; a classroom with only philosophy students can assume interest in argument analysis but not in type theory. A mixed classroom — computer science, mathematics, and philosophy students together — inherits all of the difficulties and none of the homogeneity. Two failure modes are common. The first is fragmentation: each logic (propositional, modal, epistemic, deontic) is taught as a separate subject with separate notation, separate tools, and separate exercises, so students leave with a portfolio of disconnected techniques rather than a unified competence. The second is monotony: a single logic (usually classical first-order) is taught as if it exhausted the field, leaving philosophy students without the non-classical logics their discipline demands and computer science students without the modelling resources (agents, knowledge, obligation, action) their discipline increasingly uses.

The LogiKEy methodology [1], [2] proposes a third way: logico-pluralism on a single technical substrate. Classical higher-order logic (HOL) serves as a universal metalogic. Each object logic — classical or non-classical — is encoded by defining its semantics in HOL: possible worlds become HOL individuals, accessibility relations become HOL relations on those individuals, modal operators become HOL functions quantifying over worlds. Through such semantical embeddings, one proof assistant (e.g. Isabelle/HOL, with its automated theorem provers and counter-model finders) becomes a single laboratory in which students define logics, prove their metatheorems, and compare their expressive power. Throughout, "semantical embedding" means a shallow embedding: the syntax and semantics

of an object logic are defined as HOL constants and functions, so that the object logic's theorems become HOL theorems about the embedding, without generating a new object-level deductive system. A "Kripke model" is a tuple $\mathcal{M} = (W, R_1, \dots, R_n, V)$ of worlds W , accessibility relations R_i , and a valuation V .

The LogiKey papers present a graded sequence of classroom examples, each transition motivated by a limitation of the preceding representation: a liars-and-truth-tellers puzzle leads from propositional to modal logic; the Wise Men puzzle leads on to dynamic epistemic logic; Boolos's curious inference illustrates what a higher-order metalogic buys even for automated proof search; Chisholm's paradox moves into deontic logic and then dyadic deontic logic; and Gödel's ontological argument reaches research-level metaphysics [1], [2].

What the existing treatment lacks — and what this paper supplies — is a quantitative account of *why* these transitions are pedagogically forced. Our central observation is that each representational upgrade in the sequence changes a measurable quantity: the size of the state space a student must search, the number of models a model finder must enumerate, or the length a proof must attain. When these quantities are computed explicitly, the motivation for each logical enrichment becomes a number rather than an intuition. We derive these numbers with full arithmetic in Section 4, report them in Section 5, and use them in Section 6 to defend the pluralism claim and to state falsification conditions.

2. Background and Related Work

The LogiKey methodology itself is presented in [1] and [2], which report more than a decade of use in courses, summer schools, and tutorials. [1] makes the pedagogical case for proof assistants in the logic classroom and presents the graded sequence of examples from the liars-and-truth-tellers puzzle through Gödel's ontological argument; [2] is the full paper elaborating the same program, including the rebuttal that embedding everything in classical HOL is pluralism rather than monism, since the object logics remain genuinely distinct — their theorems, valid inferences, and counter-models differ — even though their semantics are uniformly definable in HOL. Our paper accepts this framework and adds explicit quantitative checkpoints to each transition.

On the question of *what* to teach, [3] presents four authors' speculations on the future of mathematical logic in the twenty-first century across recursion theory, proof theory, model theory, and set theory. Its emphasis on proof theory and logic for computer science as driving areas supports the LogiKey bet that proof-assistant literacy is a transferable skill rather than a specialist niche, and its picture of porous boundaries between logic and computer science supports exposing students to multiple logics in one course.

On the teaching of modal logic specifically, [4] argues for an agenda motivated by central problems of philosophy and language rather than by mathematical issues alone. This is directly relevant to the LogiKey sequence: the transition from propositional to modal logic in [1] is driven by a puzzle about truth-telling and knowledge, not by Kripke completeness theorems, in line with [4]'s recommendation that philosophical motivation lead. Our quantitative state-space analysis is complementary: it gives philosophy students a concrete handle on why modal operators are needed while preserving the philosophical motivation [4] demands.

The problem of heterogeneous student populations is studied in [5], which examines teaching logic to information systems students and argues that, rather than excluding logic from their curriculum, contents and methodology must be significantly adapted to practitioner needs. The LogiKey methodology can be read as one systematic answer to [5]'s challenge: the adaptation is not a watered-down syllabus but a uniform tool-based environment in which students with different backgrounds enter through different object logics.

An alternative re-foundation of logic is computability logic, presented in [6] as a formal theory of computability — formulas as interactive computational problems modeled as games, truth as existence of a winning algorithmic strategy — in contrast to classical logic as a theory of truth. Its development into applied theories is carried forward

in [8]. Computability logic illustrates the space of proposals against which the LogiKEy choice of HOL-as-metalogic must be measured: where computability logic changes what the logical operators mean, LogiKEy keeps the metalogic fixed and varies the embedded semantics. A course built on LogiKEy could in principle embed a game semantics of the [6], [8] style as one more object logic, which is itself evidence for the flexibility of the HOL substrate.

Genuine logical pluralism inside the algebraic tradition is exemplified by [7], which studies logics defined by imposing a variable inclusion condition on a given consequence relation $\vdash$, showing that the algebraic counterpart is obtained by Plonka sums of matrix models and yielding Hilbert-style axiomatizations. Such non-classical consequence relations are exactly the kind of object logic that the semantical-embedding approach can host: a matrix semantics is a set-theoretic structure, hence definable in HOL.

The deontic stage of the LogiKEy sequence connects to [9], which introduces the neutral temporal deontic STIT logic TDS, proving soundness and completeness with respect to relational frames without utilitarian commitments. [9] demonstrates that deontic logic remains an active research frontier with new modelling resources (agency, time); a LogiKEy-style course can present TDS as a further embedding, showing students that the sequence they climbed does not end at dyadic deontic logic.

Finally, a heterodox research program questions the Boolean substrate itself: [10] collects the Continuum Critique Trilogy; [11] argues that Boolean logic reifies the void (the unmarked state) into falsity, committing a category error, and proposes a distinction-logic reconstruction labeled as conjecture; [12] argues that $1 + 1 = 2$ presupposes distinguishable marks and that idempotence ($AA = A$) is primitive; and [13] develops the calculus of re-entrant distinctions into a treatise on self-reference. Whatever their ultimate merit, these works are useful in a LogiKEy classroom precisely as embedded dissent: a heterodox logic is a good stress test of the claim that the HOL metalogic is neutral ground on which competing logics can be defined and compared.

Collectively, these works motivate a systematic quantitative study of proof-assistant-based logic teaching, situating our contribution within an emerging research agenda.

3. Methods

3.1 Semantical embeddings in HOL

The technical core is the following. Fix classical HOL with a domain of possible worlds (a HOL type w). An object modal logic with box operator is embedded by:

$$\Box_r \varphi = \lambda w. \forall v. r w v \Rightarrow \varphi v$$

where $r : w \rightarrow w \rightarrow \mathbb{B}$ is a HOL relation and $\varphi : w \rightarrow \mathbb{B}$ is a HOL predicate on worlds (a proposition is a set of worlds). Different choices of constraints on r (serial, reflexive, transitive, Euclidean, equivalence) yield different object logics (D, T, S4, S5), all inside the same metalogic. Epistemic logic adds an agent index: r_a for each agent a from a finite set A , with knowledge $K_a \varphi = \Box_{r_a} \varphi$; under S5 constraints each r_a is an equivalence relation partitioning W . Dynamic epistemic logic adds event models acting on the current model: a factual model of size $|W|$ updated by an action model of size $|E|$ yields a product model of size $|W| \cdot |E|$. Deontic logic replaces worlds-with-obligation by either a deontic accessibility r_{ob} (standard) or a dyadic conditional $\mathcal{O}(\psi \mid \varphi)$ selecting ideal φ -worlds (dyadic). None of this requires new tooling: each definition is a HOL constant, each metatheorem a HOL theorem, each counter-model a HOL-computable structure.

3.2 Method of this paper

Our method is combinatorial analysis of the LogiKey example sequence [1], [2]. For each curricular stage we identify the representational choice, define the associated search or model space, and compute its cardinality exactly from stated inputs. For each stage we identify a quantitative invariant that (i) can be computed exactly from stated inputs by elementary arithmetic, and (ii) exhibits the limitation that motivates the transition to the next logic. We deliberately restrict every claim in Section 5 to numbers derived in Section 4; where we extrapolate beyond the computation, the claim is labeled a projection with its assumptions stated. We do not report empirical learning outcomes or prover timings, as none are computed here; the claims we make are combinatorial.

4. Analysis

4.1 Stage 1: Liars and truth-tellers — exhaustive consistency in propositional logic

Input (from the puzzle structure of [1]). Three islanders a, b, c ; each is either a truth-teller (T) or a liar (L). Statements: a says " b is a truth-teller"; b says " a and c are of different types"; c says " a is a liar". Consistency condition: an islander of type T asserts only truths; an islander of type L asserts only falsehoods.

Encoding. Let $x_a, x_b, x_c \in \{T, L\}$. The three consistency constraints are:

$$x_a = (x_b = T), \quad x_b = (x_a \neq x_c), \quad x_c = (x_a = L)$$

Exhaustive check. The search space has $2^3 = 8$ assignments. We evaluate each:

1. (T, T, T) : constraint 1: $T = (T=T) = T \checkmark$. Constraint 2: $T = (T \neq T) = F \times$. Eliminated.
2. (T, T, L) : constraint 1: $T = T \checkmark$. Constraint 2: $T = (T \neq L) = T \checkmark$. Constraint 3: $L = (T=L) = F \checkmark$. **Survives.**
3. (T, L, T) : constraint 1: $T = (L=T) = F \times$. Eliminated.
4. (T, L, L) : constraint 1: $T = F \times$. Eliminated.
5. (L, T, T) : constraint 1: $L = T \times$. Eliminated.
6. (L, T, L) : constraint 1: $L = T \times$. Eliminated.
7. (L, L, T) : constraint 1: $L = (L=T) = F \checkmark$. Constraint 2: $L = (L \neq T) = T \times$. Eliminated.
8. (L, L, L) : constraint 1: $L = F \checkmark$. Constraint 2: $L = (L \neq L) = F \checkmark$. Constraint 3: $L = (L=L) = T \times$. Eliminated.

Exactly one of 8 assignments survives: $(x_a, x_b, x_c) = (T, T, L)$. The elimination rate is $7/8 = 0.875$.

Pedagogical point (quantified). For n islanders the search space is 2^n ; a truth table for $n = 10$ already needs $2^{10} = 1024$ rows checked by hand. This exhaustion of hand-checking is the computed limitation that motivates automation — and, once statements about *knowing* appear, modal logic.

4.2 Stage 2: The modal upgrade — counting accessibility relations

Input. The same puzzle re-modeled with $|W| = 8$ worlds (one per propositional assignment, from Section 4.1) and a single epistemic agent (the puzzle solver). The number of distinct binary accessibility relations on a set of $|W| = 8$ worlds is

$$2^{|W| \cdot |W|} = 2^{8 \times 8} = 2^{64} = 18,446,744,073,709,551,616 \approx 1.84 \times 10^{19}.$$

If we restrict to reflexive relations (knowledge, as standard in epistemic logic, requires reflexivity), the diagonal pairs (w, w) are fixed, leaving $8 \times 7 = 56$ off-diagonal pairs, so

$$2^{56} = 72,057,594,037,927,936 \approx 7.21 \times 10^{16}$$

reflexive relations. The jump from 8 (Section 4.1) to $2^{64} \approx 1.84 \times 10^{19}$ candidate models is the quantitative signature of the modal upgrade: the student's modeling resources have expanded from truth values to *entire relations*, and the proof assistant's model finder becomes essential precisely because manual search over 10^{19} candidates is impossible.

4.3 Stage 3: The Wise Men puzzle — model dynamics under public announcement

Input (from the puzzle structure of [1]). Three agents a, b, c ; each wears a black (B) or white (W) hat; each sees the others' hats but not their own; the king publicly announces "at least one hat is black". All three hats are in fact black. Knowledge is S5: agent i 's alternatives at a world differ from it only in coordinate i .

Model size before announcements. Hat assignments form the set $\{B, W\}^3$, so $|W_0| = 2^3 = 8$ worlds.

Announcement 0 (king): "at least one black" eliminates the all-white world WWW . Remaining: $|W_1| = 8 - 1 = 7$.

Announcement 1 (a says "I don't know"): agent a 's alternatives at w flip coordinate a . a would know only if $b = W$ and $c = W$: the alternatives are WWB and WWW ; WWW is already eliminated, so in WWB agent a would know ($a = B$). The truthful announcement "I don't know" eliminates WWB :

$$|W_2| = 7 - 1 = 6 \quad (\text{remaining: } BBB, BBW, BWB, BWW, WBB, WBW)$$

Announcement 2 (b says "I don't know"): at BBW : alternatives BBW, WWB ; WWB is eliminated, so b would know $b = B$. At WBW : alternatives WBW, WWW ; WWW eliminated, so b would know $b = W$. The announcement eliminates BBW and WBW :

$$|W_3| = 6 - 2 = 4 \quad (\text{remaining: } BBB, BWB, BWW, WBB)$$

Agent c 's inference. At the actual world BBB , agent c 's alternatives are BBB and BBW . Since $BBW \notin W_3$, the only remaining c -alternative is BBB , so c knows $c = B$. The model-size trajectory is:

$$8 \rightarrow 7 \rightarrow 6 \rightarrow 4$$

and the information gain of the two ignorance announcements is a reduction from 7 to 4 candidate worlds, i.e. a shrinkage factor of $7/4 = 1.75$.

Product-update framing. In dynamic epistemic logic (DEL), an action model with $|E|$ events updates a factual model of $|W|$ worlds to a product model of size $|W_{\text{new}}| = |W| \cdot |E|$. For a public announcement modeled with a single event ($|E| = 1$) restricted to the surviving worlds, the updated model has size $7 \times 1 = 7$; the elimination is represented by the accessibility restriction, not by world creation. Public announcements *shrink* models ($8 \rightarrow 7$), whereas private or more complex actions with $|E| > 1$ can *grow* them; e.g., an action model with $|E| = 3$ events over the original $|W| = 8$ worlds yields $8 \times 3 = 24$ worlds. Students thus see, numerically, why dynamic epistemic logic needs richer machinery than static epistemic logic.

4.4 Stage 4: Boolos's curious inference — what a higher-order metalogic buys

Input (from [1]'s use of Boolos's example). Boolos's curious inference is a first-order argument whose proof in first-order arithmetic must, for the instance with n leading universal quantifier iterations, grow like a tower of exponentials, while in second-order logic (a fragment of HOL) a short uniform proof exists. The recurrence for the first-order proof length L_n is:

$$L_0 = 1, \quad L_{n+1} = 2^{L_n}$$

Computation.

- $L_1 = 2^{L_0} = 2^1 = 2$
- $L_2 = 2^{L_1} = 2^2 = 4$
- $L_3 = 2^{L_2} = 2^4 = 16$
- $L_4 = 2^{L_3} = 2^{16} = 65536$
- $L_5 = 2^{L_4} = 2^{65536}$. In powers of ten:
 $\log_{10}(2^{65536}) = 65536 \times \log_{10}(2) = 65536 \times 0.30103 = 19728.3$ (using
 $\log_{10}(2) = 0.30103$; check: $65536 \times 0.3 = 19660.8$ and $65536 \times 0.00103 = 67.5$;
 $19660.8 + 67.5 = 19728.3$). Hence $L_5 \approx 10^{19728}$.

Combinatorial baseline. If the schema is instantiated propositionally with $k = 10$ atoms, a truth-table search enumerates $2^{10} = 1024$ rows. For a first-order encoding with a domain of size $d = 10$ and a binary predicate, the number of possible interpretations of the predicate is

$$2^{d^2} = 2^{100} \approx 1.27 \times 10^{30}$$

(since $2^{100} = (2^{10})^{10} = 1024^{10}$; numerically $2^{100} = 1,267,650,600,228,229,401,496,703,205,376 \approx 1.27 \times 10^{30}$). The higher-order encoding, by contrast, quantifies over predicates directly and lets comprehension principles constrain the search; the quantitative contrast between $2^{100} \approx 1.27 \times 10^{30}$ candidate first-order interpretations and the comprehension-constrained higher-order search is what "a higher-order metalogic buys" in practice [1], [2].

4.5 Stage 5: Deontic combinatorics (Chisholm's paradox)

Input. Chisholm's paradox, in its canonical form, involves $m = 2$ obligations (a man's duty to go to his neighbor's assistance and his duty to tell them he will come). Each obligation is either fulfilled or violated, giving

$$2^m = 2^2 = 4$$

fulfillment profiles. The paradox arises because the standard monotonic deontic operator O forces, from OA and OB and the factual violation of A , an inconsistent set of commitments; this is what motivates the dyadic operator $O(B/A)$ (obligation B under condition A). With the dyadic operator, the 4 profiles are reorganized as conditional commitments attached to each violation state, and consistency is restored profile by profile. For a three-obligation classroom variant with $m = 3$, the profile space is $2^3 = 8$. The quantitative observation: the paradox is not a growth problem (the space stays at 2^m) but a *consistency* problem — dyadic deontic logic changes the structure, not the size, of the space, which is exactly why this transition must be motivated philosophically rather than computationally.

4.6 Stage 6: Axiomatic commitments (Gödel's ontological argument)

Input. Gödel's ontological argument, as embedded in the LogiKey sequence [1], [2], posits a positive-property predicate P and a small set of axioms (typically $a = 3$ core axioms in classroom presentations: the P -positivity axiom, the axiom that a property is positive iff its negation is not positive, and the God-like-property axiom). Given a universe of $p = 5$ candidate properties (the classroom-sized property universe), the number of possible interpretations of P over the property universe is

$$2^p = 2^5 = 32.$$

Each axiom eliminates some interpretations; if each of the $a = 3$ axioms is independent and halves the space, the consistent set has size at least

$$2^{p-a} = 2^{5-3} = 2^2 = 4$$

interpretations. The pedagogical point, made quantitative: even a research-level metaphysical argument, when embedded in HOL, reduces to a small, checkable space of property interpretations (32 raw, at least 4 consistent under the stated independence assumption), which the proof assistant's model finder can enumerate exhaustively. This is the sense in which the embedding "brings it to a research-level argument" while keeping it classroom-manageable [1], [2].

4.7 Course-time projection (labeled projection)

Assumption set (stated explicitly, not measured): a semester course of $T_{\text{sem}} = 14$ weeks with one 2-hour session per week gives $14 \times 2 = 28$ contact hours. If the six stages of Section 3 are allocated in the ratio $2 : 3 : 3 : 4 : 4 : 4$ (reflecting increasing sophistication, per the graded sequence [1], [2]), the total ratio weight is $2 + 3 + 3 + 4 + 4 + 4 = 20$, so each weight unit is

$$\frac{28}{20} = 1.4 \text{ hours,}$$

and the stage allocations are $2 \times 1.4 = 2.8$, $3 \times 1.4 = 4.2$, $3 \times 1.4 = 4.2$, $4 \times 1.4 = 5.6$, $4 \times 1.4 = 5.6$, and $4 \times 1.4 = 5.6$ hours respectively. Sum check: $2.8 + 4.2 + 4.2 + 5.6 + 5.6 + 5.6 = 28.0$ hours. This is a projection under stated assumptions, not an empirical measurement; actual LogiKey course timings are not reported in quantitative form in [1], [2].

5. Results

All values below are computed in Section 4 with shown arithmetic; the single projection (R7) is labeled as such.

- **R1.** The propositional state space for $n = 3$ islanders is $2^3 = 8$ assignments, of which exactly 1 survives the three consistency constraints (elimination rate $7/8 = 0.875$); for $n = 10$ islanders the raw space is $2^{10} = 1024$ (Section 4.1).
- **R2.** The modal re-modeling over $|W| = 8$ worlds admits $2^{64} = 18,446,744,073,709,551,616 \approx 1.84 \times 10^{19}$ arbitrary accessibility relations, or $2^{56} \approx 7.21 \times 10^{16}$ reflexive ones (Section 4.2).
- **R3.** The Wise Men factual model follows the trajectory $8 \rightarrow 7 \rightarrow 6 \rightarrow 4$ under the king's announcement and two ignorance announcements, a shrinkage factor of $7/4 = 1.75$ from the first announcement onward; a DEL action model with $|E| = 3$ events over the original 8 worlds yields $8 \times 3 = 24$ worlds (Section 4.3).
- **R4.** Boolos's curious inference has first-order proof-length tower L_n with $L_1 = 2$, $L_2 = 4$, $L_3 = 16$, $L_4 = 65536$, $L_5 \approx 10^{19728}$; a first-order encoding with domain size $d = 10$ and one binary predicate presents $2^{100} \approx 1.27 \times 10^{30}$ candidate interpretations to automated search, versus $2^{10} = 1024$ propositional rows (Section 4.4).
- **R5.** Chisholm's paradox with $m = 2$ obligations has $2^2 = 4$ fulfillment profiles; a three-obligation variant has $2^3 = 8$; the paradox is structural (consistency), not combinatorial (size) (Section 4.5).
- **R6.** Gödel's ontological argument over a $p = 5$ -property universe has $2^5 = 32$ interpretations of the positivity predicate P , reduced to at least $2^{5-3} = 4$ under three independent axioms (Section 4.6).

- **R7 (projection).** Under the stated assumptions of Section 4.7 (28 contact hours, allocation ratio 2 : 3 : 3 : 4 : 4 : 4), the six stages receive 2.8, 4.2, 4.2, 5.6, 5.6, and 5.6 hours respectively; uncertainty is dominated by the assumed ratio, which we have not validated empirically.

The headline quantitative contrast is R1 vs. R2: the modal upgrade expands the student's search space from 8 to $\approx 1.84 \times 10^{19}$, a factor of

$$\frac{2^{64}}{2^3} = 2^{61} \approx 2.31 \times 10^{18}$$

(checked: $2^{61} = 2,305,846,829,608,914,401$)

References

-
- [1] TITLE: arXiv Query: `searchquery=&idlist=2610.08214&start=0&max_results=1` [2] Mathematical Proof Assistants for Teaching Logic: The LogiKey Methodology. arXiv:2610.08214v1. <https://arxiv.org/abs/2610.08214v1> [3] The prospects for mathematical logic in the twenty-first century. arXiv:cs/0205003v1. <https://arxiv.org/abs/cs/0205003v1> [4] To Teach Modal Logic: An Opinionated Survey. arXiv:1507.04701v1. <https://arxiv.org/abs/1507.04701v1> [5] Teaching Logic to Information Systems Students: Challenges and Opportunities. arXiv:1507.03687v1. <https://arxiv.org/abs/1507.03687v1> [6] Propositional computability logic I. arXiv:cs/0404023v2. <https://arxiv.org/abs/cs/0404023v2> [7] Logic of left variable inclusion and Plonka sums of matrices. arXiv:1804.08897v4. <https://arxiv.org/abs/1804.08897v4> [8] Towards applied theories based on computability logic. arXiv:0805.3521v4. <https://arxiv.org/abs/0805.3521v4> [9] A Neutral Temporal Deontic STIT Logic. arXiv:1907.03265v4. <https://arxiv.org/abs/1907.03265v4> [10] <https://papers.qnfo.org/papers/continuum-critique-trilogy/>. QNFO: The Continuum Critique Trilogy. [11] <https://papers.qnfo.org/papers/void-is-not-false/>. QNFO: The Void Is Not False: Recovering the Unmarked State in Logic from the Calculus of Indications. [12] <https://papers.qnfo.org/papers/idempotent-core/>. QNFO: The Idempotent Core: Quantity as Broken Distinction and the Hidden Assumptions of Arithmetic and Algebra. [13] <https://papers.qnfo.org/papers/reentrant-distinctions/>. QNFO: The Calculus of Re-Entrant Distinctions: A Unified Treatise on the Loop, the Tree, and the Constants of Self-Reference.