A Unified Treatment of Reachability and Indistinguishability Properties: First-Order Logic with Overwhelming Truth

Gergei Bana
Keio University
Tokyo, Japan
gebana@gmail.com

Mitsuhiro Okada
Keio University
Tokyo, Japan
mitsu@abelard.flet.keio.ac.jp

***Abstract*—In the formal verification of complexity-theoretic properties of cryptography, researchers have traditionally attempted to capture "overwhelming truth" (satisfaction with all but negligible probability) via satisfaction on individual traces of probabilistic execution. However, this approach introduces significant complexity when quantification is present: satisfaction of existential quantification over traces often produces witnesses—such as nonce-guessing oracles—that, without further constraints, may not correspond to meaningful global objects like PPT algorithms respecting causal structure. This discrepancy creates significant obstacles when attempting to combine trace properties, such as reachability, with properties that are defined on non-negligible sets such as computability, or global properties, such as algorithmic indistinguishability.**

We resolve this by shifting from defining local satisfaction as satisfaction on individual traces to a semantics based on ever-decreasing non-negligible sets. We demonstrate that the logical key to this unification lies in first-order modal logic S4 with *non-negligible sets* as possible worlds, rather than the propositional S5 fragment with *traces* as possible worlds suggested in previous investigations by the Squirrel Prover team. By introducing a PPT computational first-order S4 Kripke semantics and adopting Fitting's embedding for trace properties, we provide a unified quantified treatment of overwhelming truth for trace properties and indistinguishability, together with a first-order calculus, BC^+, whose sole modal operator is the overwhelming-truth bracket $[\cdot]$. The calculus is sound and, on its fragment, derives exactly the theorems of first-order S4 with persistence; computational completeness holds on the fragment; cut elimination holds propositionally, while at first order every cut is confined to an optimal Barcan shape.

We show that Fitting's embedding naturally accommodates the higher-order quantification used in the Squirrel prover by interpreting function types as sorts in a many-sorted first-order logic; this reduces the need for the specialized `const` predicate and its associated structural restrictions. Finally, using our findings, we present a hybrid semantics for CryptoVampire that eliminates the need for bounded Skolemization.

***Index Terms*—cryptographic protocols, formal verification, computational soundness, modal logic, first-order logic**

I. Introduction

A. Combining Local and Global Security Properties

Modern cryptography is built on probability theory and computational complexity. Security is not absolute but probabilistic—a scheme is secure if no efficient adversary can break it except with negligible probability, where "efficient" means probabilistic polynomial-time (PPT) in a security parameter η (essentially key length), and "negligible" means decreasing faster than any inverse polynomial in η. This framework expresses security properties as asymptotic statements. A central notion is *computational indistinguishability*: two families of probability distributions are indistinguishable if no PPT algorithm can tell them apart with a probability that is more than negligibly better than picking one randomly. Encryption security, for instance, requires that ciphertexts of different plaintexts be computationally indistinguishable. Formal verification of cryptographic protocols should therefore account for this probabilistic, complexity-theoretic foundation—a significant challenge when the goal is to reason symbolically, abstracting away the underlying probability distributions and security parameter.

The landscape of formal protocol verification ranges from purely symbolic methods (e.g., the Dolev-Yao model [22] underlying ProVerif [17]), which ignore computational details entirely, to fully computational game-based approaches (e.g., EasyCrypt [14], CryptoVerif [18]). The Computationally Complete Symbolic Attacker (CCSA) technique [11] (also called BC logic) occupies a middle ground: it uses symbolic terms with computational interpretation, abstracting away probabilities and the security parameter, axiomatizing cryptographic assumptions in first-order logic. In this work we focus on the CCSA technique and the provers based on it Squirrel Prover [5] and CryptoVampire [26], but the results presented here have wider relevance.

An execution trace of a probabilistic algorithm is what we get when we fix the random flow; essentially an element of the underlying probability space. A fundamental challenge when the underlying probabilistic structure and security parameter are abstracted out is combining distinct types of security properties: *local* properties, i.e., trace properties that hold with overwhelming probability (all but negligible), evaluated along individual execution traces; properties requiring non-negligible sets to be meaningful such as PPT computability; and *global* properties such as computational indistinguishability, which are statements about the entire probability space. Local prop-

erties include authentication and reachability conditions. PPT computability only makes sense on non-negligible sets, since on negligible sets such as individual traces, anything can be computed. Indistinguishability is inherently a relation between entire probability distributions that cannot be localized to individual traces. The difficulty lies in providing unified syntax and semantics where all these properties can be reasoned about coherently, especially with first-order quantification—but without explicit probabilities.

To illustrate what can go wrong when quantifying over traces, consider the "nonce-guessing problem" [26]: the formula $\exists x.\,(n = x)$, where n is an honestly generated nonce and x ranges over some infinite domain probabilistically independent of n, is satisfied on every trace because some value always equals the nonce. But this satisfaction does not yield a single PPT algorithm that computes the witness—the Skolem function becomes a non-computable oracle "guessing" the nonce on each trace.

Furthermore, a unified treatment is desirable to combine local and global reasoning without switching between them. Consider an example from [2]: let $[\phi]$ denote that trace property ϕ holds overwhelmingly. Suppose we show $[\phi_1] \vdash [\psi]$ and $[\phi_2] \vdash [\psi]$ through cryptographic arguments using indistinguishability. We want to conclude $[\phi_1 \vee \phi_2] \vdash [\psi]$, but this is non-obvious since cryptographic inference on individual traces is impossible, and $[\phi_1 \vee \phi_2]$ does not imply $[\phi_1] \vee [\phi_2]$. Nevertheless, this move is sound if things are defined correctly.

B. Partial Solutions in Prior Work

We briefly survey how research groups linked to the CCSA technique have offered partial solutions. The original CCSA for *indistinguishability* [11] avoided local quantification entirely: the *global domain* consists of PPT algorithms, satisfaction is over the whole probability space, and trace properties are expressed via $\mathsf{EQ}(_,_)$ and $\mathsf{if}\ _\ \mathsf{then}\ _\ \mathsf{else}\ _$, with overwhelming satisfaction meaning indistinguishability from ***true***. While quantification cannot be embedded within trace formulas this way (only prenex), for bounded sessions this suffices. The authors in [7] gave a complete axiomatization for such propositional reasoning over traces.

The Squirrel Prover [3], [5], built on indistinguishability CCSA, was the first to pursue the unification we also aim at here: its meta-logic places trace properties and indistinguishability in a single logic where they can interact. It allows inductive reasoning over sessions, making quantified trace reasoning unavoidable. A local formula ϕ is interpreted trace-wise, and $[\phi]$ means ϕ is satisfied overwhelmingly. Furthermore, indistinguishability is defined globally. This creates a two-tiered system: second-order local formulas with bit-string domains for trace properties, and multi-sorted first-order global formulas where the domain consists of *arbitrary random variables*, including those not respecting causal structure. The latter is necessary because trace-wise witnesses do not correspond to say PPT algorithms globally as we have seen with nonce-guessing. To handle issues rising from having arbitrary random variables, they introduce the $\mathtt{const}(x)$ predicate expressing good behavior across traces, with associated structural restrictions.

CryptoVampire [26] instead accepts satisfiability of $\exists x.\,(n = x)$ and employs bounded Skolemization to disallow constructing global witnesses from infinitely many local ones. These approaches are effective but add significant complexity.

Recently, Antoine and Baelde [2] abstracted Squirrel's propositional treatment into modal logic S5[1], with traces as possible worlds and $\Box$ expressing overwhelming satisfaction. They established propositional completeness but needed hypersequent calculus since cut-elimination fails in S5. They identify two significant open problems: extending to first-order quantification, and unifying with indistinguishability reasoning.

C. Our Approach: Overwhelming Truth through Fitting's Embedding

We resolve both open problems by returning to the original "twisted" first-order semantics of Bana and Comon [9] for *reachability properties*, combining it with their later indistinguishability CCSA under first-order S4—not S5. The aim is thus the one the Squirrel team pioneered—a single logic for reachability and indistinguishability—but reached as a genuinely *one-tier* system rather than a two-tiered one. The key insight from [12] is that Bana and Comon's semantics is best understood through Fitting's embedding [23] of classical first-order logic into S4, with a standard Kripke semantics where possible worlds are *non-negligible sets of traces* rather than individual traces, accessibility is the subset relation (S' accessible from S iff $S' \subseteq S$), and there is a single domain of PPT algorithms. An atomic trace property is satisfied on set S if satisfied on all traces in S. Fitting proved that if ϕ^* denotes the embedding (placing $\Box\Diamond$ before each subformula, which we call the Fitting-twist), then $\Phi \vdash_{\text{FOL}} \phi$ iff $\Phi^* \vdash_{\text{FOS4}} \phi^*$. First-order inference is thus sound and complete for Fitting-twisted formulas, and hence for Bana and Comon's non-standard semantics.

Note that the BC twisted semantics was only a treatment of reachability properties; S4 with the Fitting-twist appeared only in the semantics, with pure first-order syntax. On unquantified formulas, the BC twisted semantics captures exactly overwhelming satisfaction: since $\Box\Diamond$ means "for all non-negligible sets there is a smaller non-negligible set," negligible sets of non-satisfaction can remain. Existential quantification $\Box\Diamond\exists x.\,\phi^*$ means: for every non-negligible set, there is a smaller one on which a witness exists satisfying ϕ overwhelmingly on that set. Note that in this case *ϕ being satisfied locally, does not mean it is satisfied trace-wise, it means that every non-negligible set contains a smaller one where it is satisfied*. The witness is local (need only satisfy the property on an arbitrarily small non-negligible set) but belongs to the global domain. This eliminates nonce-guessing because PPT algorithms cannot guess nonces with non-negligible probability. The approach was also generalized to situations requiring

[1]For very short introduction on modal logic and S4, S5 we refer the reader to Section II-A.

satisfaction with sufficient certainty [13], where overwhelming satisfaction is a special case.

Since satisfaction of local formulas is defined on non-negligible sets, combining local and global formulas under S4 poses no difficulty. A local formula ϕ is also global with usual Tarski semantics, while $[\phi] \equiv \phi^*$ means ϕ holds overwhelmingly—not standard Tarski semantics but following first-order inference under the bracket $[_]$. Indistinguishability is defined globally. All is unified in a single tier: one language and one satisfaction relation, with no mixed rules articulating separate local and global calculi.

S4 provides technical advantages: propositionally it admits *cut-elimination* in standard sequent calculus, while labelled and prefixed systems are cut-free for quantified S4 with constant domains [25].

D. Contributions

1) **Overwhelming Semantics over Non-negligible Sets** (Section III-B)**:** For unquantified trace formulas, on the whole set of execution traces the BC semantics matches [2] (and Squirrel). We prove that with arbitrary random variables as domain, the BC first-order semantics gives exactly Squirrel's.
2) **Computationally Sound Calculus** BC^+ (Section V)**:** We define a first-order computationally sound calculus for the first-order S4 fragment in which the only modal operator is $[\cdot]$ standing for Fitting-twist: classical logic extended by necessitation (N), distribution (K), seriality (D), the Barcan formula (BF), and a flattening schema (Flat).
3) **Exactness** (Theorem 5.4)**:** On the fragment, BC^+ derives exactly what $\mathsf{FOS4}$ with Barcan formulas derives; passing to the full modal system gains no theorem expressible in the bracket language.
4) **Computational completeness** (Theorem 6.2)**:** Under the explicit model-class assumptions of the theorem, every computationally valid fragment entailment is derivable in BC^+ with persistence.
5) **Cut elimination and confinement** (Theorem 6.3)**:** The calculus admits full cut elimination propositionally; at first order every cut is confined to an optimal Barcan shape, and boxed sequents collapse to classical first-order reasoning on their matrices.
6) **Splitting and Patching with Fitting Magic** (Section VII)**:** Worlds being non-negligible sets, $[\psi]$ makes sense on each of them, not only globally; case splits proved world-wise from a global hypothesis therefore patch together, and $[\phi_1] \vdash [\psi]$ with $[\phi_2] \vdash [\psi]$ yields $[\phi_1 \vee \phi_2] \vdash [\psi]$ by one application of Fitting's theorem joining cryptographic and trace-wise arguments.
7) **Unified local/global reasoning** (Sections VII and VIII)**:** We combine reachability, PPT computability, and global indistinguishability assumptions in one tier, derive bridge rules between the original CCSA treatments, and validate the splitting principle above—realizing in one tier the unification that Squirrel's meta-logic [3], [5] pioneered with two.
8) **Higher-Order Features** (Sections IX and X)**:** Fitting's embedding extends naturally to many-sorted first-order logic, accommodating Squirrel's type system. Function-type sorts as constant domains bridge Squirrel's types and our modal semantics.
9) **Simplification of Squirrel** (Section X)**:** Our unified model is put to practical use here. By redefining $[\phi]$ as BC twisted interpretation and defining global domains for types, Squirrel's inference rules retain soundness while simplifying: we eliminate the $\mathtt{const}(x)$ predicate—"constant-ness" becomes a fundamental property of rigid designators. From the few rules of BC^+, all non-cryptographic rules of Squirrel seem derivable.
10) **Hybrid Semantics for CryptoVampire** (Section XI)**:** A purely first-order twisted hybrid semantics retains CryptoVampire's symbolic term algebra while replacing trace-by-trace evaluation. This eliminates nonce-guessing semantically, removing bounded Skolemization without new syntax.

In other words, simply redefining overwhelming satisfaction of trace formulas simplifies both Squirrel and CryptoVampire for free, without added syntax. We note that the full treatment of indistinguishability axioms within the unified framework is not completed in this paper and is deferred to future work; however, we only lay the semantic foundations.

E. Related work

The verification landscape and the CCSA line are surveyed in Sections I and II above; here we record the neighbouring mechanized frameworks and recent CCSA developments. EasyCrypt [14] and CryptoVerif [18], [20] represent the state of the art in mechanized game-based proofs. EasyCrypt is an interactive theorem prover using probabilistic relational Hoare logic, while CryptoVerif automates proofs via sequences of games. CryptHOL [15] provides a framework in Isabelle/HOL; SSProve [1] offers a foundational Coq framework using state-separating proof methodology. Recent work has extended CryptoVerif to post-quantum security [19]. On the CCSA side, recent developments include PQ-Squirrel [21], SMT-based automation [4], a Coq formalization [8] applied to FOO voting privacy, and Squirrel's inductive proof for FOO [6].

II. Preliminaries

A. Modal Logic Background

We provide a brief overview of modal logic concepts used in this paper. For comprehensive treatments, see [16], [24].

a) Propositional modal logic.: Modal logic extends classical propositional logic with operators $\Box$ (necessity) and $\Diamond$ (possibility), where $\Diamond\phi \equiv \neg\Box\neg\phi$. The semantics uses *Kripke frames*: a set W of *worlds* with an *accessibility relation* $R \subseteq W \times W$ (each element of W represents a possible state of affairs—in temporal logic, accessible worlds are future moments; in probabilistic settings, accessible worlds can represent refinements carrying additional information, i.e., subsets of a sample space). A *Kripke model* $\mathcal{M}$ adds satisfaction to propositions at each world: $\mathcal{M}, w \models P$. Satisfaction then of a

modal formula at a possible world is defined recursively the usual way for $\neg, \wedge, \vee, \rightarrow$, and for modal operators[2]:

- $\mathcal{M}, w \models \Box\phi$ iff for all w' with wRw', $\mathcal{M}, w' \models \phi$
- $\mathcal{M}, w \models \Diamond\phi$ iff for some w' with wRw', $\mathcal{M}, w' \models \phi$

b) S4 and S5.: Different axiom systems correspond to different constraints on R:

- **K**: The minimal normal modal logic. Axiom (K): $\Box(\phi \rightarrow \psi) \rightarrow (\Box\phi \rightarrow \Box\psi)$ (if $\phi \rightarrow \psi$ holds at all accessible worlds then if ϕ holds at all accessible worlds, so does ψ).
- **T**: K + Reflexivity (wRw) axiom (T): $\Box\phi \rightarrow \phi$ (if ϕ holds at all accessible worlds, then ϕ holds at the current world).
- **S4**: T + Transitivity ($wRw' \wedge w'Rw'' \Rightarrow wRw''$) axiom (4): $\Box\phi \rightarrow \Box\Box\phi$ (if at all accessible worlds ϕ holds, then from all accessible worlds, at all further accessible worlds ϕ holds).
- **S5**: S4 + Symmetry ($wRw' \Rightarrow w'Rw$), equivalently T + Euclidean ($wRw' \wedge wRw'' \Rightarrow w'Rw''$) axiom (5): $\Diamond\phi \rightarrow \Box\Diamond\phi$ (if there is an accessible world where ϕ holds, then from all accessible worlds, there is a further accessible world where ϕ holds).

S5's equivalence relation makes all worlds mutually accessible from any world, collapsing modal distinctions: $\Box\Diamond\phi \leftrightarrow \Diamond\phi$ and $\Diamond\Box\phi \leftrightarrow \Box\phi$ in S5, but not in S4.

c) First-order modal logic.: Extending to first-order adds variables, function symbols (which in cryptographic applications represent operations such as encryption, hashing, or pairing), predicate symbols (such as equality or indistinguishability), and quantification over a domain D. A key design choice is whether the domain varies across worlds or remains fixed. We are only interested in *constant domain* semantics, where all worlds share the same domain D and quantification is interpreted at each world over this common domain. This validates the *Barcan formula* $\forall x.\Box\phi \rightarrow \Box\forall x.\phi$ and its converse $\Box\forall x.\phi \rightarrow \forall x.\Box\phi$. With constant domains, quantifiers commute with $\Box$.

d) Fitting's embedding.: Fitting [23] showed that classical first-order logic can be faithfully embedded into first-order S4 by placing $\Box\Diamond$ before every subformula of a formula ϕ, producing what we call the *Fitting twist* ϕ^*. His theorem states:

$$\Gamma \vdash_{\text{FOL}} \phi \quad \text{iff} \quad \Gamma^* \vdash_{\text{FOS4}} \phi^*$$

where $\Gamma^* = \{\gamma^* \mid \gamma \in \Gamma\}$. The significance is that any S4 Kripke model gives rise to a model of classical first-order logic when formulas are read through the twist—even though satisfaction in the Kripke model is non-Tarski: negation of a formula is not the complement (on the set of models) of its satisfaction, disjunction is not the union, and existential quantification does not reduce to pointwise witnessing. Classical first-order inference nevertheless remains sound and complete for Fitting-twisted formulas. This is the key mechanism underlying the BC semantics developed in Section III-B.

[2]We write $\rightarrow$ and $\leftrightarrow$ for implication and biconditional in the object language, and $\Rightarrow$, $\Leftrightarrow$ for implication and equivalence in the metalanguage.

B. The CCSA Technique

As mentioned in the Introduction, CCSA sits between the Dolev–Yao models and the game-based computational treatments. As in Dolev–Yao, the messages exchanged with the attacker are explicit terms. Unlike Dolev–Yao, the attacker's computations are not built from a fixed list of operations it *can* perform—encryption, decryption, pairing—but are left as abstract function symbols $g_1, g_2, \ldots$ with no prescribed interpretation, standing for arbitrary efficient computations. What constrains them are first-order axioms saying what the attacker *cannot* do: these express the security properties of the cryptographic primitives and the underlying hardness assumptions. An attack, accordingly, is not exhibited but excluded: a protocol is secure if the security goal—also formalized in this language—follows from the axioms, and a countermodel means only that the goal does not follow from the axioms assumed, not necessarily—but in practice mostly—that a concrete adversary has been found.

The Computationally Complete Symbolic Attacker (CCSA) technique provides a framework for reasoning about cryptographic protocols using symbolic terms with computational interpretations, axiomatizing security properties in first-order logic rather than proving them via direct game reductions. Terms are interpreted as PPT algorithms operating on random tapes; the details are presented in Section III-B.

CCSA was developed in two stages addressing different classes of security properties. The original *reachability* variant [9], [10] targets trace properties such as authentication or secrecy. It introduced a novel non-Tarski first-order computational semantics capturing overwhelming satisfaction, where quantification ranges over a global domain of PPT algorithms and existential witnesses need only work on non-negligible sets of traces, not on individual ones. This semantics was later understood [12] to be an instance of Fitting's embedding of classical first-order logic into S4.

The subsequent *indistinguishability* variant [11] introduced a predicate $\vec{t} \sim \vec{u}$ for computational indistinguishability of term sequences. Since indistinguishability is inherently a global property—a statement about entire probability distributions—this variant uses standard Tarski semantics with satisfaction defined over the whole probability space. Trace properties are expressed indirectly: "ϕ holds overwhelmingly" is encoded as $b_\phi \sim \textbf{\textit{true}}$, where b_ϕ is a corresponding Boolean term. A complete axiomatization for propositional reasoning was given in [7].

Two provers implement CCSA variants. The *Squirrel Prover* [3], [5], built on the indistinguishability variant, extends it with inductive reasoning over protocol sessions. Squirrel reintroduced explicit formulas for trace properties alongside global indistinguishability formulas, but did not utilize the Fitting-based semantics from the reachability work. Instead, trace formulas are evaluated trace-by-trace, forcing a two-tiered system with different quantification domains at

each level and requiring the $\texttt{const}(x)$ predicate to control quantification across traces. *CryptoVampire* [26] automates trace property proofs by encoding them for saturation-based first-order theorem provers, but its trace-by-trace semantics requires bounded Skolemization to prevent nonce-guessing witnesses.

Recently, Antoine and Baelde [2] abstracted the propositional fragment of overwhelming truth into modal logic S5, but identified first-order extension and unification with indistinguishability as open problems. In this paper, we show that returning to the Fitting-based semantics of the original reachability work resolves both problems and simplifies both tools.

III. First-Order Overwhelming Truth for Local Properties

In this section, we first summarize Antoine and Baelde's propositional logic of overwhelming truth for local (trace) formulas based on S5 [2], then present the alternative BC semantics based on Fitting's embedding into S4 [10], [12], [13]. We re-establish that the two approaches coincide for quantifier-free formulas, while the S4-based approach extends naturally to first-order quantification, and show that for arbitrary random variables the BC semantics is the same as Squirrel's.

A. Antoine and Baelde's Overwhelming Truth in S5

Antoine and Baelde [2] define a modal logic of overwhelming truth where formulas are interpreted over *cryptographic structures*. A cryptographic structure $\mathcal{C}$ consists of an ensemble of sample spaces $\Omega^{\mathcal{C}} = \{\Omega^{\mathcal{C}}_\eta\}_{\eta\in\mathbb{N}}$, one for each security parameter η, together with an interpretation of each propositional variable p from some prefixed set of propositional cariables as an ensemble of functions $p^{\mathcal{C}} = \{p^{\mathcal{C}}_\eta : \Omega^{\mathcal{C}}_\eta \to \{0,1\}\}_{\eta\in\mathbb{N}}$. Each $\Omega^{\mathcal{C}}_\eta$ carries a probability measure μ_η—in the intended cryptographic reading, normally the uniform distribution on random tapes—and every probability below is taken with respect to it. A propositional variable p stands for a concrete trace property, for instance "the signature received by the responder verifies under the initiator's public key": on each fixed stream of the random tape this is either true or false, so it determines a function $\Omega^{\mathcal{C}}_\eta \to \{0,1\}$.

a) Syntax.: The syntax is standard propositional modal logic:

$$\phi ::= p \mid \bot \mid \phi \Rightarrow \psi \mid \Box\phi$$

where p ranges over propositional variables, and $\Diamond\phi \equiv \neg\Box\neg\phi$.

b) Semantics.: The interpretation of each formula ϕ in structure $\mathcal{C}$ is an ensemble of functions $[\![\phi]\!]^{\mathcal{C}} = \{[\![\phi]\!]^{\mathcal{C}}_\eta : \Omega^{\mathcal{C}}_\eta \to \{0,1\}\}_{\eta\in\mathbb{N}}$, defined inductively:

- $[\![p]\!]^{\mathcal{C}} = p^{\mathcal{C}}$
- $[\![\bot]\!]^{\mathcal{C}}_\eta(\rho) = 0$ for all η and $\rho \in \Omega^{\mathcal{C}}_\eta$
- $[\![\phi \Rightarrow \psi]\!]^{\mathcal{C}}_\eta(\rho) = 1$ iff $[\![\phi]\!]^{\mathcal{C}}_\eta(\rho) = 0$ or $[\![\psi]\!]^{\mathcal{C}}_\eta(\rho) = 1$
- $[\![\Box\phi]\!]^{\mathcal{C}}_\eta(\rho) = 1$ iff the function $\eta' \mapsto \Pr_{\rho'\in\Omega^{\mathcal{C}}_{\eta'}}[[\![\phi]\!]^{\mathcal{C}}_{\eta'}(\rho') = 1]$ is *overwhelming*

A function $f : \mathbb{N} \to [0,1]$ is *overwhelming* if $1 - f(\eta)$ is negligible, i.e., for all polynomials q, there exists N such that for all $\eta > N$, $1 - f(\eta) < 1/q(\eta)$. Thus $[\![\phi]\!]^{\mathcal{C}}$ is overwhelming when the function $\eta \mapsto \Pr_{\rho\in\Omega^{\mathcal{C}}_\eta}[[\![\phi]\!]^{\mathcal{C}}_\eta(\rho) = 1]$ is overwhelming.

Note that the clause for $\Box$ is stated pointwise but its right-hand side mentions neither η nor ρ: the condition is on the ensemble as a whole. Hence $[\![\Box\phi]\!]^{\mathcal{C}}_\eta(\rho)$ is a constant function, taking the same value at every security parameter and every tape. This is what makes $\Box$ express a global, asymptotic statement rather than a trace-local one. Note also that this is not a Kripke semantics.

A formula ϕ is *valid* in the logic of overwhelming truth if $[\![\phi]\!]^{\mathcal{C}}$ is overwhelming for all cryptographic structures $\mathcal{C}$.

c) Equivalence with S5.: Antoine and Baelde establish that the modal logic of overwhelming truth coincides exactly with S5. The soundness direction (S5 theorems are valid for overwhelming truth) is relatively straightforward: one verifies that axioms T ($\Box p \Rightarrow p$), 4 ($\Box p \Rightarrow \Box\Box p$), and 5 ($\Diamond p \Rightarrow \Box\Diamond p$) hold under the overwhelming truth interpretation. For completeness, they show that any finite clique counter-model of an S5 formula can be transformed into a cryptographic structure witnessing non-validity.

d) Determined formulas.: A key notion in their work is that of *determined formulas*. A formula ϕ is determined if $\Box\phi \vee \Box\neg\phi$ (equivalently, $\Diamond\phi \to \Box\phi$) is valid—that is, ϕ is either overwhelmingly true or overwhelmingly false, with no middle ground. Determined formulas are those whose Boolean value does not change (overwhelmingly) from trace to trace, and so are constants as discussed in the introduction. Because of the issues mentioned in the introduction they make heavy use of this notion.

Since S5 lacks cut-elimination, Antoine and Baelde adapt Poggiolesi's hypersequent calculus [28], which is perhaps the main contribution, but which significantly complicates the proof system. They identify two open problems: extending to first-order quantification and connecting to the indistinguishability predicate of CCSA: "More generally, it would naturally be desirable to extend our results to richer fragments of the logic of overwhelming truth: one possibility is to extend our language of formulas to the first-order case; it is also possible to stick to the propositional case but incorporate the computational indistinguishability predicate of the CCSA logics. Both directions seem very challenging." The latter seems fundamentally difficult in their framework since indistinguishability cannot be defined on individual traces, yet their semantics has no global possible world.

B. First-Order Overwhelming Truth via Fitting's Twist

We now first present the original BC "twisted" semantics for overwhelming truth based on Fitting's embedding of classical first-order logic into modal logic S4 [23]. This approach, developed in [10], [12], [13], works naturally for first-order logic hence preserving cut-elimination. The original semantics also contained the PPT computability predicate $\triangleright$, which is not a trace property, perhaps suggesting the Fitting embedding was created only for this predicate. However, the embedding works

for any semantics targeting "enough certainty" as demonstrated in [13]. We focus on trace properties until Section VIII, where $\triangleright$ returns in the hiding rule. We do not present security axioms here, we explain that at the indistinguishability CCSA.

a) Syntax.: Bana and Comon build terms from variables, names $n \in \mathcal{N}$ (honestly generated nonces) and function symbols $f \in \mathcal{F}$ (PPT computable functions), together with adversarial function symbols $g \in \mathcal{G}$ representing arbitrary PPT computations available to the attacker. Atomic formulas $P(\vec{t})$ express *trace properties*: predicates P on lists of terms $\vec{t}$ whose truth can be checked on a single random tape, given the bit-string values of their arguments. The primary example is equality $t_1 = t_2$, but any tape-checkable relation qualifies—ordering $t_1 < t_2$ on timestamps or indices and so on. Bana and Comon also used $\vec{t} \triangleright t'$ as PPT computability of t' from $\vec{t}$ since although this cannot be satisfied trace-wise, it can be on arbitrary small non-negligible ensembles, and for them that is what local means, not strictly trace properties.

Formulas are built using standard connectives (modal operators appear only in the semantics, not the syntax):

$$\phi ::= P(\vec{t}) \mid \bot \mid \phi \wedge \psi \mid \phi \vee \psi \mid \phi \to \psi \mid \neg\phi \mid \forall x.\phi \mid \exists x.\phi$$

b) Computational model.: We define a computational model $\mathcal{M}^{\mathcal{C}}$ following [12] with slight modifications. Fix an ensemble of sample spaces $\Omega^{\mathcal{C}} = \{\Omega^{\mathcal{C}}_\eta\}_{\eta\in\mathbb{N}}$, where each $\Omega^{\mathcal{C}}_\eta$ is a standard probability space of *pairs* of tapes $\rho = (\rho_h, \rho_a)$, representing honest and adversarial randomness respectively; the standard choice is $2^{\mathbb{N}} \times 2^{\mathbb{N}}$ with the fair-coin product measure μ_η and its Borel σ-algebra $\mathcal{A}_\eta$, a PPT algorithm reading only polynomially many tape positions; "random" throughout means distributed according to μ_η, and all probabilities are taken with respect to it. We refer to a single $\rho \in \Omega^{\mathcal{C}}_\eta$ as a *trace*.

c) Notation for ensembles of sets.: Throughout, sets of traces are η-indexed families rather than plain sets. An *ensemble* (but we usually just say set) S is a family $S = \{S_\eta\}_{\eta\in\mathbb{N}}$ of measurable sets, $S_\eta \in \mathcal{A}_\eta$. For ensembles S, S' we write $S' \subseteq S$ to mean $S'_\eta \subseteq S_\eta$ for every η. The ensemble S is *negligible* if the function $\eta \mapsto \mu_\eta(S_\eta)$ is negligible, and *non-negligible* otherwise—that is, if there is a polynomial q such that $\mu_\eta(S_\eta) > 1/q(\eta)$ for infinitely many η.

The computational model $\mathcal{M}^{\mathcal{C}}$ consists of:

1) **Domain** D: A set of PPT algorithms operating on random tapes from $\Omega^{\mathcal{C}}$. The technique works for other domains (e.g., quantum computable functions); the key point is that the domain does not depend on worlds. Each $\mathfrak{a} \in D$ takes security parameter 1^η and random tape $\rho \in \Omega^{\mathcal{C}}_\eta$, outputting a bit string. We write $\mathfrak{a}_{\eta,\rho}$ for this output.
2) **Variable assignments**: A valuation $\nu : \mathrm{Var} \to D$ maps variables to domain elements. We write $\nu[x \mapsto \mathfrak{a}]$ for the assignment mapping x to $\mathfrak{a}$ and agreeing with ν elsewhere.
3) **Interpretation of function symbols**: Each n-ary function symbol f is interpreted as a PPT function $f^{\mathcal{M}^{\mathcal{C}}}$ operating on random tapes and *bit-string outputs*. The interpretation $[\![t]\!]^\nu_{\eta,\rho}$ is defined inductively:
   - $[\![x]\!]^\nu_{\eta,\rho} = (\nu(x))_{\eta,\rho}$
   - $[\![f(t_1,\ldots,t_n)]\!]^\nu_{\eta,\rho} = f^{\mathcal{M}^{\mathcal{C}}}_\eta([\![t_1]\!]^\nu_{\eta,\rho},\ldots,[\![t_n]\!]^\nu_{\eta,\rho})$

   Functions apply to *outputs* of subterms, not algorithms themselves, ensuring locality.
4) **Interpretation of adversarial function symbols**: Each n-ary $g \in \mathcal{G}$ is interpreted as a PPT algorithm $g^{\mathcal{M}^{\mathcal{C}}}$ which, unlike the honest symbols, may also read the adversarial tape:
$$[\![g(t_1,\ldots,t_n)]\!]^\nu_{\eta,\rho} = g^{\mathcal{M}^{\mathcal{C}}}_\eta([\![t_1]\!]^\nu_{\eta,\rho},\ldots,[\![t_n]\!]^\nu_{\eta,\rho};\rho_a)$$
5) **Interpretation of predicate symbols**: Each n-ary predicate P (trace property) is interpreted as $P^{\mathcal{M}^{\mathcal{C}}}$ returning 0 or 1 on bit-string outputs:
$$[\![P(\vec{t})]\!]^\nu_{\eta,\rho} = P^{\mathcal{M}^{\mathcal{C}}}_{\eta,\rho}([\![t_1]\!]^\nu_{\eta,\rho},\ldots,[\![t_n]\!]^\nu_{\eta,\rho})$$
   To be a trace property, it depends only on ρ and the argument outputs, with measurable truth sets for fixed arguments. We can allow non-PPT computability for generality, but we can just as well restrict to it.
6) **Names**: $\mathcal{N}$ is the countable set of *names*, the constant symbols denoting randomness drawn honestly—nonces, keys, encryption randomness. Each name $n \in \mathcal{N}$ is interpreted as independent uniform random sampling of $p(\eta)$ bits from honest randomness ρ_h, where p is a fixed polynomial. Distinct names read disjoint slices of ρ_h, so they are probabilistically independent of one another and of anything the adversary computes without them. This will be important for CryptoVampire.

d) Kripke semantics over non-negligible sets.:

Definition 3.1 (Computational Kripke Model): Given a computational model $\mathcal{M}^{\mathcal{C}}$, the associated *computational Kripke model* is the tuple $\mathcal{M}^{\mathcal{KC}} = (W, R, D, \mathcal{I})$ where:

- W is the set of non-negligible measurable ensembles $S = \{S_\eta\}_\eta$, where $S_\eta \in \mathcal{A}_\eta$ and the function $\eta \mapsto \mu_\eta(S_\eta)$ is non-negligible.
- R is the reverse inclusion relation: SRS' iff $S' \subseteq S$
- D is the domain of PPT algorithms, fixed across all worlds
- $\mathcal{I}$ is the interpretation inherited from $\mathcal{M}^{\mathcal{C}}$: for each function symbol f and predicate symbol P, $\mathcal{I}$ assigns the interpretations $f^{\mathcal{M}^{\mathcal{C}}}$ and $P^{\mathcal{M}^{\mathcal{C}}}$ respectively

The accessibility relation is reverse inclusion: SRS' iff $S' \subseteq S$, in the sense of the ensemble notation above. The domain D is fixed across all worlds—the same set of PPT algorithms is quantified over at every world—enabling the Barcan formulas. Equality is a distinguished predicate among the trace properties, read as a congruence: its intended interpretation, literal equality of the bit-string values, is the primary example, but the definition does not hardwire it (cf. Remark 7).

For an atomic formula $P(\vec{t})$ and variable assignment $\nu : \mathrm{Var} \to D$, satisfaction at world S means $P(\vec{t})$ holds throughout

S:

$$\mathcal{M}^{\mathcal{KC}}, S, \nu \models_{S4} P(\vec{t}) \iff \text{for all } \eta \text{ and all } \rho \in S_\eta,\ [\![P(\vec{t})]\!]^\nu_{\eta,\rho} = 1$$

Logical connectives and quantifiers are interpreted in the standard Tarskian way relative to each world:

- $\mathcal{M}^{\mathcal{KC}}, S, \nu \models_{S4} \phi \wedge \psi$ iff $\mathcal{M}^{\mathcal{KC}}, S, \nu \models_{S4} \phi$ and $\mathcal{M}^{\mathcal{KC}}, S, \nu \models_{S4} \psi$
- $\mathcal{M}^{\mathcal{KC}}, S, \nu \models_{S4} \phi \vee \psi$ iff $\mathcal{M}^{\mathcal{KC}}, S, \nu \models_{S4} \phi$ or $\mathcal{M}^{\mathcal{KC}}, S, \nu \models_{S4} \psi$
- $\mathcal{M}^{\mathcal{KC}}, S, \nu \models_{S4} \neg\phi$ iff $\mathcal{M}^{\mathcal{KC}}, S, \nu \not\models_{S4} \phi$
- $\mathcal{M}^{\mathcal{KC}}, S, \nu \models_{S4} \forall x.\phi$ iff for all $\mathfrak{a} \in D$, $\mathcal{M}^{\mathcal{KC}}, S, \nu[x \mapsto \mathfrak{a}] \models_{S4} \phi$
- $\mathcal{M}^{\mathcal{KC}}, S, \nu \models_{S4} \exists x.\phi$ iff there exists $\mathfrak{a} \in D$ such that $\mathcal{M}^{\mathcal{KC}}, S, \nu[x \mapsto \mathfrak{a}] \models_{S4} \phi$

Modal operators were not included in the local formula syntax, but are needed for S4 semantics. This is a transitional step to recover Bana and Comon's non-standard semantics:

- $\mathcal{M}^{\mathcal{KC}}, S, \nu \models_{S4} \Box\phi$ iff for all non-negligible $S' \subseteq S$, $\mathcal{M}^{\mathcal{KC}}, S', \nu \models_{S4} \phi$
- $\mathcal{M}^{\mathcal{KC}}, S, \nu \models_{S4} \Diamond\phi$ iff for some non-negligible $S' \subseteq S$, $\mathcal{M}^{\mathcal{KC}}, S', \nu \models_{S4} \phi$

It is natural to ask whether requiring ϕ at all non-negligible $S' \subseteq S$ is really stronger than requiring it at S alone. For *atomic* formulas it is not since for us now they are trace properties: if $P(\vec{t})$ holds at every trace in S then it holds at every trace of any subset, so $P(\vec{t}) \to \Box P(\vec{t})$ is valid. For compound formulas the two differ: the clause for $\neg P(\vec{t})$ at S merely says $P(\vec{t})$ fails at S, which is compatible with $P(\vec{t})$ holding at some smaller S' and failing on its complement, but in this case $\Box\neg P(\vec{t})$ fails.

e) Fitting's embedding.: Fitting's embedding [23] transforms any first-order formula ϕ into an S4 formula ϕ^* by placing $\Box\Diamond$ before every subformula. By S4 equivalences and Barcan formulas, many occurrences simplify:

$$\begin{aligned}
P(\vec{t})^* &= \Box\Diamond P(\vec{t}) \\
\bot^* &= \bot \\
(\phi_1 \wedge \phi_2)^* &= \phi_1^* \wedge \phi_2^* \\
(\phi_1 \vee \phi_2)^* &= \Box\Diamond(\phi_1^* \vee \phi_2^*) \\
(\phi_1 \to \phi_2)^* &= \Box(\phi_1^* \to \phi_2^*) \\
(\neg\phi)^* &= \Box\neg\phi^* \\
(\forall x.\phi)^* &= \forall x.\phi^* \\
(\exists x.\phi)^* &= \Box\Diamond\exists x.\phi^*
\end{aligned}$$

We call this the *Fitting twist*. Each simplification is an FOS4-equivalence; they are derived in Appendix B, in the more general setting where the arguments may themselves already be twisted, so that the two presentations agree under nesting as well. For a set Γ, let Γ^* denote the Fitting-twisted formulas.

Fitting's theorem states that $\Gamma \vdash_{\text{FOL}} \phi$ if and only if $\Gamma^* \vdash_{\text{FOS4}} \phi^*$.

Only two of the ϕ^* clauses carry modal content. Taking $\neg, \wedge, \forall$ as primitive and reading $\vee, \to, \exists$ as the usual abbreviations, the remaining clauses are forced if we want first-order soundness.

f) BC overwhelming semantics via Fitting.: Bana and Comon used the term "computational" rather than "overwhelming" semantics. A local formula ϕ holds *computationally* on non-negligible set S iff its Fitting twist ϕ^* is S4-satisfied on S:

$$\mathcal{M}^{\mathcal{C}}, S, \nu \models_{BC} \phi \iff \mathcal{M}^{\mathcal{KC}}, S, \nu \models_{FOS4} \phi^*$$

Using Fitting is convenient; the final semantics can be defined without modal operators as Bana and Comon did. The *computational satisfaction* relation $\mathcal{M}^{\mathcal{C}}, S, \nu \models_{BC} \phi$ turns out to be recursively:

Definition 3.2 (BC Twisted Computational Semantics): Let $\mathcal{M}^{\mathcal{C}}$ be a computational model, S a non-negligible set, and $\nu : \text{Var} \to D$ a valuation.

Atomic formulas:

- $\mathcal{M}^{\mathcal{C}}, S, \nu \models_{BC} P(\vec{t})$ iff for all non-negligible $S' \subseteq S$, there exists a non-negligible $S'' \subseteq S'$ such that for all η and all $\rho \in S''_\eta$, $[\![P(\vec{t})]\!]^\nu_{\eta,\rho} = 1$. As we will see this is the same as that $1 - \Pr_\rho(\rho \in S : [\![P(\vec{t})]\!]^\nu_{\eta,\rho} = 1)$ is negligible in η.

Propositional connectives:

- $\mathcal{M}^{\mathcal{C}}, S, \nu \models_{BC} \bot$ never holds
- $\mathcal{M}^{\mathcal{C}}, S, \nu \models_{BC} \phi \wedge \psi$ iff $\mathcal{M}^{\mathcal{C}}, S, \nu \models_{BC} \phi$ and $\mathcal{M}^{\mathcal{C}}, S, \nu \models_{BC} \psi$
- $\mathcal{M}^{\mathcal{C}}, S, \nu \models_{BC} \phi \vee \psi$ iff for all non-negligible $S' \subseteq S$, there exists a non-negligible $S'' \subseteq S'$ such that $\mathcal{M}^{\mathcal{C}}, S'', \nu \models_{BC} \phi$ or $\mathcal{M}^{\mathcal{C}}, S'', \nu \models_{BC} \psi$
- $\mathcal{M}^{\mathcal{C}}, S, \nu \models_{BC} \neg\phi$ iff for all non-negligible $S' \subseteq S$, $\mathcal{M}^{\mathcal{C}}, S', \nu \not\models_{BC} \phi$
- $\mathcal{M}^{\mathcal{C}}, S, \nu \models_{BC} \phi \to \psi$ iff for all non-negligible $S' \subseteq S$, if $\mathcal{M}^{\mathcal{C}}, S', \nu \models_{BC} \phi$ then $\mathcal{M}^{\mathcal{C}}, S', \nu \models_{BC} \psi$

Quantifiers:

- $\mathcal{M}^{\mathcal{C}}, S, \nu \models_{BC} \forall x.\phi$ iff for all $\mathfrak{a} \in D$, $\mathcal{M}^{\mathcal{C}}, S, \nu[x \mapsto \mathfrak{a}] \models_{BC} \phi$
- $\mathcal{M}^{\mathcal{C}}, S, \nu \models_{BC} \exists x.\phi$ iff for all non-negligible $S' \subseteq S$, there exist a non-negligible $S'' \subseteq S'$ and $\mathfrak{a} \in D$ such that $\mathcal{M}^{\mathcal{C}}, S'', \nu[x \mapsto \mathfrak{a}] \models_{BC} \phi$

A formula ϕ is *computationally valid* (written $\mathcal{M}^{\mathcal{C}} \models_{BC} \phi$) iff $\mathcal{M}^{\mathcal{C}}, \Omega^{\mathcal{C}}, \nu \models_{BC} \phi$ for all valuations ν (which implies $\mathcal{M}^{\mathcal{C}}, S, \nu \models_{BC} \phi$ for all non-negligible S).

The "for all non-negligible S', there exists non-negligible S''" pattern corresponds exactly to $\Box\Diamond$ in S4 Kripke semantics, and ensures overwhelming truth by allowing failure on negligible sets.

Theorem 3.3: If $\Gamma \vdash_{\text{FOL}} \phi$ (i.e., ϕ is derivable from Γ in classical first-order logic), then for any computational model $\mathcal{M}^{\mathcal{C}}$:

$$\mathcal{M}^{\mathcal{C}} \models_{BC} \Gamma \implies \mathcal{M}^{\mathcal{C}} \models_{BC} \phi$$

This soundness theorem is highly non-trivial. Bana and Comon proved it directly in [9], but it easily follows from the non-trivial Fitting's theorem and S4 soundness for Kripke semantics [12].

Remark 1 (Completeness of first-order inference): The Fitting-based semantics is also complete. Any consistent set Γ has a standard first-order model $\mathcal{M}$ by classical completeness. This model can be viewed as a degenerate Kripke structure where every non-negligible set is a singleton (the deterministic case), where the Fitting twist becomes trivial and computational satisfaction coincides with standard first-order satisfaction.

Remark 2 (Completeness does not mean real attacks): This completeness does not imply that symbolic attacks yield computational ones. At this point there are no cryptographic elements; it is just a trivial observation. Cryptography is added via cryptographic axioms, and whether the failure of a proof provides a real attack depends on the completeness of those axioms.

C. The BC Semantics and Overwhelming Truth

We first re-establish that the S4-based Fitting-twisted BC semantics and S5-based overwhelming truth of Antoine-Baelde [2] coincide for quantifier-free formulas, and extend to first-order formulas over random variable domains. Intuitively, the $\Box\Diamond$ pattern means that instead of considering satisfaction on traces, we can zoom onto traces by taking any small non-negligible sets: for any non-negligible set there is a smaller non-negligible set where the property in question holds. Hence if a property holds this way, then there cannot be a non-negligible set where the negation holds, as on that set there would be no smaller non-negligible set where it holds (note the negation also has a box outside); therefore it holds overwhelmingly. This result is from [10] though not spelled out rigorously there.

Proposition 3.4 (Equivalence for Quantifier-Free Formulas): Let ϕ be a quantifier-free formula. Then ϕ holds computationally (i.e., $\mathcal{M}^{\mathcal{C}} \models_{BC} \phi$) if and only if ϕ holds with overwhelming probability in the sense of [2].

Proof:[Proof sketch] By structural induction on ϕ. The base case shows that $\Box\Diamond P$ holds in the Kripke model iff P holds except on a negligible set. The inductive cases for conjunction and negation follow from properties of negligible sets and the Fitting twist. See Appendix A for the full proof. □

Under Squirrel's finite-discrete tape convention, the equivalence extends to first-order formulas when the domain consists of all random variables, as in Squirrel's semantics [5].

Since $\Omega^{\mathcal{C}}$ is an η-indexed family rather than a single space, we first make precise what a random variable over it is. A *random variable* with values in a set V is a family $\mathfrak{a} = \{\mathfrak{a}_\eta\}_{\eta\in\mathbb{N}}$ of functions $\mathfrak{a}_\eta : \Omega^{\mathcal{C}}_\eta \to V$, each measurable with respect to μ_η. We write $\mathfrak{a}_{\eta,\rho}$ for $\mathfrak{a}_\eta(\rho)$. A PPT algorithm is the special case in which the family is required to be polynomial-time computable.

Definition 3.5 (Random Variable Domain): A *random variable domain* for a computational model $\mathcal{M}^{\mathcal{C}}$ is a domain D consisting of all random variables $\mathfrak{a} = \{\mathfrak{a}_\eta : \Omega^{\mathcal{C}}_\eta \to V\}_{\eta\in\mathbb{N}}$ in the above sense, for some value space V.

Assumption 1 (Measurability of Trace Sets): For any first-order formula ϕ and valuation ν, and for every η, the set $S^{\eta}_{\phi,\nu} = \{\rho \in \Omega^{\mathcal{C}}_\eta \mid \phi \text{ holds at } \rho \text{ under } \nu\}$ is measurable.

Squirrel settles this issue by bounding tape length: its tapes are finite bit strings, so the sample space is finite, carries the discrete σ-algebra, and every function out of it is measurable [5]. Accordingly, Theorem 3.7 is stated under this finite-discrete convention, so trace-wise witness selection is measurable without invoking a separate measurable-selection theorem. This restriction concerns only the comparison with trace-by-trace random-variable semantics; the Fitting-twisted semantics over a PPT domain does not require it.

Definition 3.6 (Trace-Level Satisfaction and Overwhelming Trace Truth): For $\rho \in \Omega^{\mathcal{C}}_\eta$, define $\mathcal{M}^{\mathcal{C}}, \eta, \rho, \nu \models_{\mathrm{tr}} \phi$ by the ordinary Tarskian clauses at that trace: atoms are evaluated by $[\![P(\vec{t})]\!]^{\nu}_{\eta,\rho}$, the Boolean connectives are pointwise, and quantifiers range over the fixed domain D. A formula ϕ holds *trace-by-trace with overwhelming probability* under ν if there is a negligible function ϵ such that

$$\Pr_{\rho\leftarrow\mu_\eta}\left[\mathcal{M}^{\mathcal{C}}, \eta, \rho, \nu \models_{\mathrm{tr}} \phi\right] \geq 1 - \epsilon(\eta)$$

for every η.

Theorem 3.7 (Equivalence for Random Variable Domains): Let $\mathcal{M}^{\mathcal{C}}$ be a computational model with random-variable domain D and suppose that every $\Omega^{\mathcal{C}}_\eta$ is finite and carries the discrete σ-algebra. For any first-order formula ϕ and valuation ν:

$$\mathcal{M}^{\mathcal{C}}, \Omega^{\mathcal{C}}, \nu \models_{BC} \phi \iff$$
ϕ holds trace-by-trace with overwhelming probability

Proof:[Proof sketch] By structural induction. When $\exists x.\psi$ holds trace-by-trace, we construct a witness random variable $\mathfrak{a}^*$ by choosing a witnessing value separately at each trace. Each section of this selection is measurable because the tape space is finite and discrete, hence $\mathfrak{a}^* \in D$; the converse is by contradiction. See Appendix A. □

Remark 3: This means the BC twisted semantics provides exactly the extension of Antoine and Baelde's S5 semantics to quantification that they seek.

IV. CCSA Technique Based on Indistinguishability

The previous section only concerned local properties, modalities could be hidden in the semantics. Now we consider computational indistinguishability, which cannot be localized, and explain the core of the CCSA Indistinguishability technique on which Squirrel and CryptoVampire are based.

The CCSA technique based on indistinguishability was introduced by Bana and Comon-Lundh [11] as a follow-up to their reachability work [10]. This version represented the real breakthrough: a computationally sound symbolic technique that is syntactically simple, with standard game-based security notions easy to translate. We summarize following [7]; see that paper for complete details. Both Squirrel [3], [5] and CryptoVampire [26] are based on this version, though with

different design choices for quantification and trace properties. We recall its syntax, semantics and axioms because the next section interprets exactly these formulas over possible worlds; note that trace properties enter this variant only encoded as Boolean terms, $b_\phi \sim \textbf{\textit{true}}$.

A. Syntax and Semantics

The syntax is first-order logic with a distinguished predicate $\sim$ for computational indistinguishability. Terms are those of Section III-B—built from names $\mathcal{N}$, honest symbols $\mathcal{F}$ and adversarial symbols $\mathcal{G}$—refined with sorts $\mathcal{S}$ including bool and $\mathsf{message}$ (names have type $\mathsf{message}$); among the honest symbols we distinguish $\mathbf{0}$ (empty message), ***true***, ***false***, $\mathsf{EQ}(_,_)$ for equality testing, $\mathsf{if}\ _\ \mathsf{then}\ _\ \mathsf{else}\ _$ for conditionals, and cryptographic operations (encryption, decryption, signatures).

The key atomic formula $t_1, \ldots, t_n \sim u_1, \ldots, u_n$ asserts computational indistinguishability of two term sequences. The abbreviation $x = y$ stands for $\mathsf{EQ}(x, y) \sim \textbf{\textit{true}}$.

a) Global formulas.: Global formulas are built inductively from indistinguishability atoms $\vec{t} \sim \vec{u}$, Boolean connectives, and quantifiers $\forall x.\phi$, $\exists x.\phi$.

Terms are interpreted exactly as in the computational model of Section III-B (honest symbols act on outputs only, adversarial symbols also read ρ_a); this variant differs only at the formula level. Only the distinguished symbols need their interpretations fixed: $\mathsf{EQ}(t, u)$ evaluates to ***true*** at (η, ρ) iff $[\![t]\!]^\nu_{\eta,\rho} = [\![u]\!]^\nu_{\eta,\rho}$, and $\mathsf{if}\ b\ \mathsf{then}\ u\ \mathsf{else}\ v$ evaluates to $[\![u]\!]^\nu_{\eta,\rho}$ if $[\![b]\!]^\nu_{\eta,\rho} = \textbf{\textit{true}}$, and to $[\![v]\!]^\nu_{\eta,\rho}$ otherwise.

Given valuation ν mapping variables to D, the indistinguishability predicate is: $\mathcal{M}^c, \nu \models \vec{t} \sim \vec{u}$ iff for any PPT distinguisher $\mathcal{A}$,

$$\left| \Pr_\rho[\mathcal{A}([\![\vec{t}]\!]^\nu_{\eta,\rho}; \rho_a) = 1] - \Pr_\rho[\mathcal{A}([\![\vec{u}]\!]^\nu_{\eta,\rho}; \rho_a) = 1] \right|$$

is negligible in η. Unlike reachability CCSA, connectives and quantifiers are interpreted the standard Tarski way—over the whole space, with no twist.

B. Axiomatization

The technique employs two categories of axioms. *Core axioms* formalize basic properties of indistinguishability and conditionals:

- REFL: $t \sim t$
- SYM: $\vec{t} \sim \vec{u} \to \vec{u} \sim \vec{t}$
- TRANS: $\vec{t} \sim \vec{u} \wedge \vec{u} \sim \vec{v} \to \vec{t} \sim \vec{v}$
- FUNCAPP: $\vec{t} \sim \vec{u} \to f(\vec{t}) \sim f(\vec{u})$
- IFTRUE: $b \sim \textbf{\textit{true}} \to \mathsf{if}\ b\ \mathsf{then}\ u\ \mathsf{else}\ v\ \sim u$
- IFFALSE: $b \sim \textbf{\textit{false}} \to \mathsf{if}\ b\ \mathsf{then}\ u\ \mathsf{else}\ v\ \sim v$

A complete axiomatization for $\mathsf{if}_\mathsf{then}_\mathsf{else}_$ with booleans was given in [7].

A second category formalizes hardness assumptions and the security of cryptographic primitives; [7] covers DDH, IND-CPA, IND-CCA1, IND-CCA2, and UF-CMA unforgeability. These axioms are formalized in [7] such that recursively defined constraints on contexts represent execution traces of the games corresponding to the security notion, and the indistinguishability of these execution traces are stated.

a) Example: IND-CPA and IND-CCA axioms.: To illustrate the cryptographic axiom formulation, consider the encryption security axioms. Here $\{m\}^{n_2}_{\mathsf{pk}(n)}$ encrypts m under the public key generated from name n with randomness name n_2, and L returns message length. All of IND-CPA, IND-CCA1, and IND-CCA2 have the same conclusion (one can replace the encrypted text if their lengths are the same):

$$\vec{t}\left[\mathsf{if}\ \mathsf{EQ}(\mathsf{L}(m), \mathsf{L}(m'))\ \mathsf{then}\ \{m\}^{n_2}_{\mathsf{pk}(n)}\ \mathsf{else}\ u''\right]$$
$$\sim \vec{u}\left[\mathsf{if}\ \mathsf{EQ}(\mathsf{L}(m), \mathsf{L}(m'))\ \mathsf{then}\ \{m'\}^{n_3}_{\mathsf{pk}(n)}\ \mathsf{else}\ u''\right]$$

This is an axiom schema; the three notions differ in where the decryption key $\mathsf{sk}(n)$ may occur in $\vec{t}$ and $\vec{u}$: nowhere (CPA), only in decryptions computed before the challenge (CCA1), or anywhere except on the challenge ciphertext itself (CCA2). Strict recursive definitions were given in [7].

In CCSA, one proves indistinguishability of two *executions* rather than two games as in game reductions. Attacker computations are abstract function symbols $f_1, f_2, \ldots$ acting on previous challenger outputs. For protocols Π_l and Π_r whose equivalence expresses the security property, executions unfold as

$$\Phi(\Pi_{l/r})_{i+1} = g_{l/r}(f_{i+1}(\Phi(\Pi_{l/r})_i))$$

where f_{i+1} represents any PPT attacker computation (not a concrete operation like decryption as in Dolev-Yao), and $g_{l/r}$ represents honest branching using $\mathsf{if}\ b\ \mathsf{then}\ u\ \mathsf{else}\ v$ and EQ. Attacker computations are identical on both sides. The goal is proving $\Phi(\Pi_l) \sim \Phi(\Pi_r)$ using the axioms, with branching handled by axioms like $(b, u \sim b', u') \to (b, t \sim b', t') \to (b, \mathsf{if}\ b\ \mathsf{then}\ u\ \mathsf{else}\ t\ \sim b', \mathsf{if}\ b'\ \mathsf{then}\ u'\ \mathsf{else}\ t')$.

The main soundness theorem [11]: if the computationally sound axioms imply $\Phi(\Pi_l) \sim \Phi(\Pi_r)$, then Π_l and Π_r are computationally indistinguishable under those assumptions, for any bounded number of sessions.

As mentioned earlier, a potential attack exists, if from such axioms, the security property of the protocol cannot be verified.

V. First-order Overwhelming Truth combined with Global Formulas

Now we are ready to combine local and global properties and in this section we lay the fundamentals, on which later sections are built. We have two options for unification: go for the full first-order S4 with $\Box$, or just use the modality $[\Phi]$, which in our case stands for the Fitting-twist. Since to our knowledge no work on protocol verification makes use of the full S4, we stick to $[\Phi]$ in this work, and we allow this modality to be applied to any formula; in other words, any property can be localized. We also provide an inference system for this fragment that is sound for the fragment and hence for the computational Kripke-semantics, and derives exactly the fragment theorems of first-order S4 (Theorem 5.4), hence one

has the option to do proofs within the fragment or go out to the full S4.

a) **Syntax:** We consider the following grammar:

$$\Phi ::= P(\vec{t}) \mid \bot \mid \Phi \wedge \Phi \mid \Phi \vee \Phi \mid \neg\Phi \mid \Phi \to \Phi \mid \forall x.\Phi \mid \exists x.\Phi \mid [\Phi]$$

where $[\,\cdot\,]$ is an abbreviation for the Fitting-twist of Section III-B, treated as a primitive operator.

Remark 4 (Two presentations of the nested twist): We allow iteration of $[\,\cdot\,]$. At every bracket one may use either of the equivalent presentations from Section III-B: the *uniform* presentation, which places $\Box\Diamond$ before every subformula, or the simplified presentation displayed there. Thus, for example,

$$[\Phi \to \Psi] \leftrightarrow \Box\big([\Phi]) \to [\Psi]\big).$$

The equivalence remains valid in the presence of nested brackets; Appendix B proves this for each connective and the universal quantifier. Hence the grammar above defines a fragment of first-order S4 independently of which presentation is used.

Throughout the section, Φ, Ψ range over arbitrary formulas and ϕ, ψ over modality-free ones. We thus keep $[\,\cdot\,]$ as the only modality in the syntax, $\Box$ and $\Diamond$ appearing in the semantics alone. The fragment sits inside the full modal language as the image of the twist.

b) **Computational Semantics:** The computational Kripke-semantics defined as in Section III-B, except that predicates do not have to be defined trace-wise. Note, this way formulas in $[_]$ obtain a twist in the semantics.

Let $\mathsf{FOS4}$ denote first-order S4 with the constant domains: to the S4 axioms of Section II-A, the classical first-order axioms, modus ponens and necessitation, it adds the Barcan formula $\forall x.\Box\Phi \to \Box\forall x.\Phi$ (its converse being derivable).

Over this language, we define BC as classical first-order logic together with

(N) from $\vdash_{\mathrm{FOL}} \phi$ infer $\vdash [\phi]$, for modality-free ϕ;
(K) $[\Phi \to \Psi] \to ([\Phi] \to [\Psi])$;
(D) $[\bot] \to \bot$;
(BF) $\forall x.[\Phi] \to [\forall x.\Phi]$;

and BC^+ as BC extended by

(Flat) $[C([\Phi_1], \ldots, [\Phi_k])] \leftrightarrow [C(\Phi_1, \ldots, \Phi_k)]$, for every classical context C.

Definition 5.1 (Local predicates): A predicate P is *local* if the persistence schema

$$\text{(P)} \qquad P(\vec{x}) \to [P(\vec{x})]$$

is sound for it: its truth at a world carries to every non-negligible subset. A signature is *local* if all its predicates are.

We write $\mathsf{FOS4}_P$ for the same system with the persistence schema: $P(\vec{t}) \to \Box P(\vec{t})$ for when all predicates are local.

Trace properties are local, their truth at S being truth at every trace of S, and they remain so under all first-order connectives, evaluated trace-wise. Locality is however the weaker notion: the computability predicate $\vec{t} \triangleright u$ of Section III-B—some single PPT algorithm computes $[\![u]\!]$ from $[\![\vec{t}\,]\!]$ at every trace of the world—is not a trace property, yet it is local, one algorithm witnessing every subset. By contrast the indistinguishability predicate $\sim$ of Section IV is *global*: it typically holds on the whole space while failing on subworlds, so (P) is unsound for it. A formula is local when it is built from local predicates without $[\,\cdot\,]$.

Theorem 5.2 (Soundness for the fragment): Let Γ be a set and Φ a formula of the fragment. If $\Gamma \vdash_{\mathsf{BC}^+} \Phi$ then $\Gamma \vdash_{\mathsf{FOS4}} \Phi$, both sides read through the Fitting twist. Over a local signature the same holds for the persistent systems: if $\Gamma \cup \mathsf{P} \vdash_{\mathsf{BC}^+} \Phi$ then $\Gamma \vdash_{\mathsf{FOS4}_P} \Phi$, where P collects the instances of (P).

Proof:

- Fitting's theorem gives (N)
- (K) follows from $[\Phi \to \Psi] \equiv \Box[\Phi \to \Psi]$ from which $[\Phi \to \Psi]$ follows by Reflexivity of S4.
- For (D): $[\bot] \equiv \bot$ by the simplified definition of Fitting-embedding.
- For (BF): $\forall x.[\Phi] \to [\forall x.\Phi]$ again by the simplified definition of the Fitting-twist.
- For (Flat), recall the uniform form of Fitting's embedding: up to $\mathsf{FOS4}$-equivalence, the twist of a formula is obtained by inserting $\Box\Diamond$ at every subformula. Applying the twist to an already twisted formula merely inserts a second $\Box\Diamond$ at each subformula, and these repetitions disappear by the S4 identity

  $$\Box\Diamond\Box\Diamond\chi \leftrightarrow \Box\Diamond\chi.$$

  Both halves are short chains in $\mathsf{FOS4}$ as defined above. Write (Mon$\Box$) for the derived rule taking $\vdash \alpha \to \beta$ to $\vdash \Box\alpha \to \Box\beta$—necessitation followed by K—and (Mon$\Diamond$) for its dual. From left to right: T gives $\Box\Diamond\chi \to \Diamond\chi$, so (Mon$\Diamond$) gives $\Diamond\Box\Diamond\chi \to \Diamond\Diamond\chi$, which axiom 4 in dual form contracts to $\Diamond\chi$; one application of (Mon$\Box$) yields $\Box\Diamond\Box\Diamond\chi \to \Box\Diamond\chi$. From right to left: axiom 4 gives $\Box\Diamond\chi \to \Box\Box\Diamond\chi$. The dual of T gives $\theta \to \Diamond\theta$, and with $\theta = \Box\Diamond\chi$ we obtain $\Box\Diamond\chi \to \Diamond\Box\Diamond\chi$, and so by (Mon$\Box$), $\Box\Box\Diamond\chi \to \Box\Diamond\Box\Diamond\chi$. Hence $\Box\Diamond\chi \to \Diamond\Box\Diamond\chi$. Thus an outer twist absorbs every inner bracket, which is exactly (Flat).
- Finally the instances of (P) translate to $P(\vec{t}) \to \Box\Diamond P(\vec{t})$, derivable in $\mathsf{FOS4}_P$ from the persistence axiom $P(\vec{t}) \to \Box P(\vec{t})$. The dual of T gives $P(\vec{t}) \to \Diamond P(\vec{t})$, and (Mon$\Box$) turns this into $\Box P(\vec{t}) \to \Box\Diamond P(\vec{t})$; composing with the persistence axiom gives the instance.

□

Corollary 5.3 (Computational soundness): If $\Gamma \vdash_{\mathsf{BC}^+} \Phi$ then Φ holds at every world of every computational Kripke model at which all of Γ holds; for trace properties $P(\vec{t})$, (P) holds.

Proof: The accessibility relation of Definition 3.1 is reverse inclusion, hence reflexive and transitive, and the domain D of PPT algorithms is the same at every world. So $\mathcal{M}^{\mathcal{KC}}$ is a constant-domain S4 Kripke model, a class for which $\mathsf{FOS4}$—S4 together with the classical first-order axioms, modus ponens, necessitation and the Barcan formula—is sound. Theorem 5.2 places every BC^+-derivation inside $\mathsf{FOS4}$, and

the first claim follows. For a trace property $P(\vec{t}) \to \Box P(\vec{t})$ of the persistence axiom of $\mathsf{FOS4}_P$ is valid in the model precisely because such a P is local (Definition 5.1): its truth at a world carries to every non-negligible subset, and those are exactly the accessible worlds. □

Theorem 5.4 (Exactness over FOS4*):* For sets Γ and sentences Φ of the fragment,

$$\Gamma \vdash_{\mathsf{BC}^+} \Phi \quad \text{if and only if} \quad \Gamma \vdash_{\mathsf{FOS4}} \Phi.$$

That is, every modal theorem expressible in the fragment is provable by the principles above, so nothing is lost by reasoning inside it. The prover can choose the more convenient way to work with.

The soundness part is Theorem 5.2; the completeness part is in Appendix C. The intuition is the following: The countermodel is built as a depth-one fan. Suppose that $\Gamma \nvdash_{\mathsf{BC}^+} \Phi$. After using (FLAT) to remove nested brackets, extend $\Gamma \cup \{\neg\Phi\}$ to a maximal Henkin-complete theory T (where Henkin-complete means that whenever T contains an existential sentence $\exists x\, \psi(x)$, it also contains a witnessing instance $\psi(t)$ for some closed term t), and separate the information in T into two parts: B, the bare atoms that should hold at the root, and $F = \{\phi : [\phi] \in T\}$, the modality-free theory that should hold under brackets. The axioms of BC^+ ensure that F is a consistent first-order theory with witnesses in the same term domain. For every modality-free sentence $\phi \notin F$, choose a classical term model M_ϕ that satisfies all of F but falsifies ϕ, and place all these models as terminal successors of a new root w_0, whose atomic valuation is exactly B. At an endpoint, the twist reduces to ordinary classical truth, while at the root, $[\phi]$ holds exactly when every endpoint satisfies ϕ. Thus every member of F is true under brackets at the root, because all endpoints satisfy F, and every sentence outside F is false under brackets there, because its designated endpoint falsifies it. The root therefore has exactly the bare atomic theory B and exactly the bracketed theory F; since all formulas have been flattened, a truth induction shows that it realizes all of T, and hence satisfies Γ while refuting Φ. In the persistent version, the added instances $P \to [P]$ give $B \subseteq F$, ensuring that every atom true at the root is also true at every endpoint.

Then, over a local signature, $\Gamma \cup \mathsf{P} \vdash_{\mathsf{BC}^+} \Phi$ if and only if $\Gamma \vdash_{\mathsf{FOS4}_P} \Phi$, with P the definitional instances of (P).

VI. Completeness for Computational Models and Cut Confinement

Completeness and cut-elimination are not our major aim in this paper; we briefly state here the results we have so far and their surrounding discussion, with the proofs in Appendices D, E and F.

A. Completeness

Soundness (Corollary 5.3) says that what BC^+ derives holds in every computational Kripke model. Completeness asks the converse, and the question must be posed with care: relative to which calculus, and over which models? Following Antoine–Baelde's propositional completeness [2], whose quantified generalization this section provides, we ask it for computational models of *trace* predicates—pointwise trace properties, the paradigm local predicates of Definition 5.1, for which the persistence instances $P \to [P]$ are sound by definition and indeed valid in every model of the class, pointwise atoms persisting to subsets. A calculus complete for this class must therefore derive the (P)-instances, which plain FOS4 does not: the target is BC^+ with the definitional (P)-instances adjoined, equivalently $\mathsf{FOS4}_P$ (Section V). What happens for predicates beyond this class we leave open: the non-local $\sim$, but also the computability predicate $\triangleright$, which is local yet no trace property—the theorem's boundary is pointwise-ness, not locality, so overwhelming computability is not characterized by it either. Even so scoped, the question cannot be asked for the full modal language: on computational models it outruns $\mathsf{FOS4}_P$ already propositionally.

Proposition 6.1 (Incompleteness, already propositional): There are propositional schemas valid in every computational Kripke model but underivable in $\mathsf{FOS4}_P$; a witness, for atomic p, q, r:

$$\Theta := \Box(p \to (q \vee r)) \to \Box(p \to q) \vee \Box(p \to r).$$

Proof: Validity. Let S be a world satisfying the hypothesis and A the set of traces where p holds. If $S \cap A$ is negligible, every sub-world satisfying the pointwise atom p would be a non-negligible subset of it, so both boxed disjuncts hold vacuously. Otherwise $S \cap A$ is itself a world, and it is the largest sub-world of S satisfying p; the hypothesis applied there gives q or r, say q, at every trace of $S \cap A$. Then $\Box(p \to q)$ holds at S: any sub-world satisfying p is contained in $S \cap A$, where q holds pointwise. *Underivability.* The persistent constant-domain tree with a root satisfying no atom and two reflexive successors with valuations $\{p, q\}$ and $\{p, r\}$ is a model of all $\mathsf{FOS4}_P$ axioms and refutes Θ at the root: the hypothesis holds at all three worlds, while each boxed disjunct fails at one successor. □

Nor does adding Θ close the gap: worlds are closed under unions and under intersecting a world with the trace set of an atom, and this structure validates more than any evident finite extension proves. We therefore leave the full modal language unaxiomatized.

The fragment of Section V is the repair; the semantic and proof-theoretic statements follow. Throughout the rest of this section and Appendix D, the signature is assumed countable and effectively presented; the encodings of the realization live on this assumption.

Theorem 6.2 (Computational completeness): Over an effectively presented countable signature, for sets Γ and sentences Φ of the fragment: if Φ holds under Γ in every computational Kripke model of Definition 3.1—a class in which every predicate is a trace property, hence persistent, taken with the relaxation of Remark 7: equality read as a congruence, and names as constants with arbitrary PPT interpretation—then Φ is derivable in BC^+ from Γ and the instances of (P), all of which are definitional over this class. Together with soundness

and Theorem 5.4 this makes three conditions coincide on the fragment: derivability from Γ with the (P)-instances in BC^+, derivability from Γ in $\mathsf{FOS4}_P$, and validity under Γ in these models.

B. Cut Elimination

For $\mathsf{FOS4}_P$ itself the proof theory is largely settled in the literature. Propositional S4 has cut-free Gentzen-style calculi, and prefixed and labelled systems extend cut-freeness to quantified S4 with constant domains and to frame conditions like persistence [24], [25], [27]. What resists is a cut-free *unlabelled* sequent calculus in the presence of the Barcan formula: a constant domain is not a local condition on a sequent, and ordinary quantifier rules cannot express it. The question for us is what survives in BC^+, whose language is narrower—the only modality is $[\,\cdot\,]$, and (Flat) removes nesting—so that one might hope the obstruction disappears with the nesting. It does not, but it localizes completely: cuts cannot be eliminated, yet every one of them can be confined to a single shape.

The natural calculus for BC^+ is classical LK (Gentzen's classical sequent calculus) together with a single interface rule: from $\vec{P}, m(\Phi_1), \ldots, m(\Phi_k) \vdash_{\mathrm{FOL}} m(\Psi)$ infer $\vec{P}, [\Phi_1], \ldots, [\Phi_k] \vdash [\Psi]$, where $m(\cdot)$ erases the brackets of its argument. For example, $P \vdash [P]$ is one such inference.

Theorem 6.3 (Cut confinement and partial cut elimination): In the unlabelled calculus above (made precise in Appendix E), all cuts confine to formulas of the form $\forall x.[\Theta]$—the *Barcan cuts* of that appendix—optimally so; cut elimination holds on sequents without outer quantifiers, with a disjunction property; and boxed sequents collapse to classical first-order logic on matrices.

Full cut elimination is expected in a labelled calculus, and the route is known: labelled and prefixed systems are cut-free for propositional S4 [27] and for quantified S4 with constant domains [25], and hereditary-atom conditions such as $P \to [P]$ are frame conditions of the same geometric kind; assembling the combined calculus for BC^+ is left to future work.

Proposition 6.4 (The propositional case): Theorems 6.2 and 6.3 hold propositionally with elementary proofs: for fragment sentences without quantifiers, inside or outside the brackets, derivability in BC^+ from the definitional instances of (P) coincides with computational validity, and the calculus admits full cut elimination, with the disjunction property: $\vec{P}, [\phi_1], \ldots, [\phi_k] \vdash \vec{Q}, [\psi_1], \ldots, [\psi_l]$ is derivable iff some Q_i is among $\vec{P}$, or the matrices $\vec{P} \cup \{m(\phi_1), \ldots, m(\phi_k)\}$ are inconsistent or derive some $m(\psi_j)$.

Appendix F has the short proofs.

VII. Advantages of the S4-Based Approach

The Fitting embedding into S4 offers several advantages over S5 for reasoning about overwhelming truth. We highlight key benefits here; further advantages appear when combining with indistinguishability-based CCSA in Section IV.

a) **Combining trace properties with indistinguishability.:** Moving to S4 allows combining overwhelming truth with indistinguishability-based CCSA. Since possible worlds are non-negligible sets (not individual traces), we can define computational indistinguishability on each world using standard definitions. But the real strength, as we will see with Squirrel, is that we can argue with formulas of the form $\Theta \vdash \Box([\phi] \to [\psi])$ (which is syntactically the same as $\Theta \vdash [\phi \to \psi]$)) where we can join trace-wise first-order inference and non-negligible set-wise cryptographic arguments. The former works because of Fitting, the latter because the satisfaction is by non-negligible sets.

b) **First-order extension.:** The Fitting-embedding reduces to $(\forall x.\phi)^* = \forall x.\phi^*$ and $(\exists x.\phi)^* = \Box\Diamond\exists x.\phi^*$. The existential case deserves attention. The formula $\Box\Diamond\exists x.\phi^*$ means: for any non-negligible set S', there exists a non-negligible subset $S'' \subseteq S'$ and a witness $d \in D$ such that $\phi^*[d/x]$ holds throughout S''. The witness is *local*—it need only work on the smaller set S'', not all of S'. However, S'' must still be non-negligible; the witness cannot be defined only on a negligible set.

This locality of witnesses resolves the Skolemization problem in trace-by-trace semantics. In trace-by-trace approaches (like CryptoVampire's original semantics), different random tapes may require different witnesses, and when combined, these form a non-computable function. In the Fitting semantics, each witness must work on a non-negligible set, ensuring Skolem functions remain meaningful and computable.

c) **Rigid designators.:** Rigid designators (terms denoting the same object across all worlds) arise naturally from the constant domain assumption. No separate "constant predicate" is needed; special global constants can be introduced via appropriate typing.

d) **Splitting and patching with Fitting magic.:** Returning to the Introduction's disjunction problem: suppose we have a proof system like Squirrel's but the semantics of overwhelming probability is defined with the Fitting-twist, so $[\phi] \equiv \phi^*$ ($[\phi]$ asserts overwhelming truth of the trace formula ϕ). For trace properties ϕ_1, ϕ_2 and ψ, suppose cryptographic arguments using indistinguishability, from a set Γ of not necessarily twisted formulas such as global indistinguishability, establish $[\phi_1] \to [\psi]$ and $[\phi_2] \to [\psi]$ not just on the full probability space, but on every non-negligible subset. That is, we have

$$\Gamma \vdash \Box([\phi_1] \to [\psi]) \quad \text{and} \quad \Gamma \vdash \Box([\phi_2] \to [\psi]),$$

which is $\Gamma \vdash [\phi_1 \to \psi]$ and $\Gamma \vdash [\phi_2 \to \psi]$ by the Fitting-twist. We take these boxed judgments as the given; how they arise from cryptographic assumptions is discussed in Sections IX and X. In our S4 setting this corresponds to $\Gamma \vdash \Box(\phi_1^* \to \psi^*)$ and $\Gamma \vdash \Box(\phi_2^* \to \psi^*)$, which is the same as $\Gamma \vdash (\phi_1 \to \psi)^*$ and $\Gamma \vdash (\phi_2 \to \psi)^*$, so $\Gamma \vdash (\phi_1 \to \psi)^* \wedge (\phi_2 \to \psi)^* \equiv ((\phi_1 \to \psi) \wedge (\phi_2 \to \psi))^*$. Since by Fitting, we can use first-order inference inside the star, we have $\Gamma \vdash (\phi_1 \vee \phi_2 \to \psi)^*$,

which is the same as $\Gamma \vdash \Box((\phi_1 \vee \phi_2)^* \to \psi^*)$, that is,

$$\Gamma \vdash \Box([\phi_1 \vee \phi_2] \to [\psi]),$$

and axiom T discharges the outer box. So from the boxed implications we can derive $[\phi_1 \vee \phi_2] \to [\psi]$. The magic works because the twisted semantics has a box outside: $\Box([\phi] \to [\psi])$ says not only that if ϕ holds overwhelmingly on the whole space then so does ψ, but that the same is true on every non-negligible set. Such situations arise, for example, when protocols branch—some keys generated on one side, others on another, with the same authentication result but for different reasons. One can do the proofs separately. Note also that since $(\phi \to \psi)^* \equiv \Box(\phi^* \to \psi^*)$ is a syntactic identity, the above derivation only needs the star, not the box so it can be formulated in the BC^+ system.

e) **Cut-elimination.:** Propositionally, S4 admits cut-elimination in standard sequent calculus, whereas S5 requires hypersequent calculus—and the axiomatization of Antoine–Baelde is propositional, so this comparison applies. At first order with constant domains, the Barcan formula blocks cut-elimination in unlabelled calculi for S4 and S5 alike; labelled and prefixed calculi are known to be cut-free for quantified S4 with constant domains [25], [27], and hereditary atoms $P \to \Box P$ are frame conditions of the same geometric kind.

VIII. Combining Indistinguishability and Overwhelming Truth through S4

We now link the reachability CCSA's Fitting-based semantics with indistinguishability CCSA's Tarskian semantics. While indistinguishability is primarily for equivalence properties, trace properties can also be expressed—indeed, much of typical proofs involves trace property arguments. A trace property "ϕ holds overwhelmingly" encodes as $b_\phi \sim \textbf{\textit{true}}$ where b_ϕ is the corresponding Boolean term. The $\mathsf{if}\ b\ \mathsf{then}\ u\ \mathsf{else}\ v$ construct represents both protocol branching and proof case analysis, connecting indistinguishability with overwhelming truth via Fitting embedding.

A. If-Then-Else in the Indistinguishability Framework

The key to connecting the two CCSA approaches is the if-then-else construct. In indistinguishability CCSA, $\mathsf{if}\ b\ \mathsf{then}\ u\ \mathsf{else}\ v$ expresses either an honest agent's execution step or a proof case distinction. Combining such terms yields Boolean operations: $\mathsf{if}\ b\ \mathsf{then}\ (\mathsf{if}\ b'\ \mathsf{then}\ u\ \mathsf{else}\ v\ )\ \mathsf{else}\ v$ equals $\mathsf{if}\ b \wedge b'\ \mathsf{then}\ u\ \mathsf{else}\ v$, and $\mathsf{if}\ b\ \mathsf{then}\ u\ \mathsf{else}\ v$ equals $\mathsf{if}\ \neg b\ \mathsf{then}\ v\ \mathsf{else}\ u$. When u, v have type bool, we obtain a propositional calculus over traces. As shown in [7], this calculus has a complete axiomatization[3], and for any b of type bool: $b \sim \textbf{\textit{true}} \leftrightarrow [b = \textbf{\textit{true}}]$. This calculus and this equivalence make the translation below invertible; they reappear as Squirrel's reasoning on local formulas (Section X) and as CryptoVampire's EQ (Section XI).

[3]Note that in [7] $=$ meant overwhelming equality, which now is expressed with the twist.

B. Translation Between If-Then-Else and Local Formulas

Recall the local formulas of Section III-B. Any local predicate P can be viewed as a function symbol of type bool (equality corresponds to EQ). To any local formula ϕ we assign a Boolean term b_ϕ:

- $b_{P(\vec{t})} \equiv P(\vec{t})$
- $b_{\neg\phi} \equiv \mathsf{if}\ b_\phi\ \mathsf{then}\ \textbf{\textit{false}}\ \mathsf{else}\ \textbf{\textit{true}}$
- $b_{\phi\wedge\psi} \equiv \mathsf{if}\ b_\phi\ \mathsf{then}\ b_\psi\ \mathsf{else}\ \textbf{\textit{false}}$
- $b_{\phi\vee\psi} \equiv \mathsf{if}\ b_\phi\ \mathsf{then}\ \textbf{\textit{true}}\ \mathsf{else}\ b_\psi$
- $b_{\phi\to\psi} \equiv \mathsf{if}\ b_\phi\ \mathsf{then}\ b_\psi\ \mathsf{else}\ \textbf{\textit{true}}$

The inverse translation is also possible. For any quantifier-free local formula ϕ:

$$[\phi] \leftrightarrow [b_\phi = \textbf{\textit{true}}]$$

since both express overwhelming truth (recall $[\phi] \equiv \phi^*$, and $b \sim \textbf{\textit{true}} \leftrightarrow (b = \textbf{\textit{true}})^*$ by the previous subsection). By Theorem 3.7, this extends to quantification over arbitrary random variables. The consequence: any Boolean term translates to an unquantified formula whose twisted semantics corresponds exactly to that term being overwhelmingly equal to ***true***. This bridge is used by the inference rules below, and it is how Squirrel's local judgments become the $[\phi]$ of Section X.

C. Unification in S4 Through Fitting's Embedding

At this point it is entirely clear how to unify the two methods. They are unified under the S4 semantics where local quantified formulas get translated by Fitting, while formulas containing the equivalence predicate stay the same. The domain of interpretation for both of them is still D regardless whether we consider local or global formulas. We only have to define the semantics of $\vec{u} \sim \vec{t}$ over arbitrary non-negligible set, which is easy: $\mathcal{M}^{\mathcal{KC}}, S, \nu \models \vec{t} \sim \vec{u}$ iff

$$\left| \Pr_\rho[\mathcal{A}(\llbracket \vec{t} \rrbracket^\nu_{\eta,\rho}; \rho_a) = 1 \mid \rho \in S] - \Pr_\rho[\mathcal{A}(\llbracket \vec{u} \rrbracket^\nu_{\eta,\rho}; \rho_a) = 1 \mid \rho \in S] \right|$$

is negligible.

a) Unified syntax.: No new language is needed: we take that of Section V, its predicates now including the global atom $\vec{t} \sim \vec{u}$ alongside the local ones. Every local formula is thereby also a global one, and $[\phi]$ abbreviates ϕ^* as before. The local/global distinction is the semantic one of Definition 5.1, whose examples already place trace properties and $\triangleright$ on the local side and $\sim$ on the global one.

Semantics: that of the computational Kripke model $\mathcal{M}^{\mathcal{KC}}$ (Definition 3.1), extended by the clause for $\sim$ given above; any further predicate may be interpreted by an arbitrary relation indexed by the worlds.

D. Cryptographic Inference Rules and Soundness

We do not attempt to give a full axiomatization here as we will show in Section X how to import Squirrel's rules to this system. Instead we just introduce a bridge between the two levels, and rules that Squirrel cannot state, or can state only through its helper predicates. To the rules of BC^+ with the (P)-instances adjoined—equivalently, by the persistent form of Theorem 5.4, those of $\mathsf{FOS4}_P$—we add the following.

- For ϕ without quantification and only trace predicates, $[\phi] \leftrightarrow b_\phi \sim \textbf{\textit{true}}$ as explained before. This plays the role of moving between local and global formulas.
- $[\vec{x} \sim \vec{y}] \rightarrow [\vec{x} = \vec{y}]$, the twist extending to the $\sim$ atom by the atomic clause, $[\vec{x} \sim \vec{y}] \equiv \Box\Diamond(\vec{x} \sim \vec{y})$: overwhelming indistinguishability implies overwhelming equality. It cannot be stated in Squirrel, brackets on global formulas being unavailable there. For soundness we show the contrapositive: if $[\vec{x} = \vec{y}]$ fails at a world S, then so does $[\vec{x} \sim \vec{y}]$. Intuitively, on a suitable sub-world a single fixed bit position tells the two interpretations apart at every trace, and there indistinguishability fails on every smaller world. Assuming that $(\vec{x} = \vec{y})^*$ fails, unfolding the twist, some non-negligible $S' \subseteq S$ has only a negligible set of traces on which $[\![\vec{x}]\!] = [\![\vec{y}]\!]$; removing them leaves a non-negligible T with $[\![\vec{x}]\!] \neq [\![\vec{y}]\!]$ at every trace. Pad both interpretations to a common polynomial length $\ell(\eta)$. At each trace of T they differ at some position; taking the first such, the pair formed by this position and the bit of $[\![\vec{x}]\!]$ there takes at most $2\ell(\eta)$ values, so by pigeonhole some fixed position $i(\eta)$ and bit value—say 0—occur on a non-negligible $S^* \subseteq T$. Thus at every trace of S^*, the bit of $[\![\vec{x}]\!]$ at position $i(\eta)$ is 0 and that of $[\![\vec{y}]\!]$ is 1. The distinguisher with $i(\eta)$ hardcoded as nonuniform advice, outputting the bit of its input at position $i(\eta)$, therefore returns 0 on samples of $\vec{x}$ and 1 on samples of $\vec{y}$ at every trace of S^*, and so has advantage 1 conditioned on any non-negligible $T' \subseteq S^*$. Hence $\Diamond(\vec{x} \sim \vec{y})$ fails at S^*, and $(\vec{x} \sim \vec{y})^* = \Box\Diamond(\vec{x} \sim \vec{y})$ fails at S.
- **The hiding rule.** Recall the computability predicate $\vec{t} \triangleright u$ of Section III-B: some single PPT algorithm computes $[\![u]\!]$ from $[\![\vec{t}]\!]$ at every trace of the world—persistent, hence local, yet no trace property, as discussed in Section V. For distinct names n, n' and a tuple of terms $\vec{t}[n]$ in which n' does not occur (the occurrences of n are unconstrained),

$$\vec{t}[n],\, n \;\sim\; \vec{t}[n],\, n' \;\rightarrow\; [\neg\, \vec{t}[n] \triangleright n]\,:$$

if $\vec{t}[n]$ hides n, then n cannot be computed from it on any non-negligible set. Soundness: it suffices that every PPT g succeeds only on a negligible set of traces, for then $\triangleright$ fails at every world and the twisted conclusion, $\Box\neg\Box\Diamond(\vec{t}[n] \triangleright n)$, holds: the distinguisher that runs g on the tuple's first component—on the shared adversarial tape—and tests equality with the second succeeds with g's probability on $([\![\vec{t}[n]]\!], [\![n]\!])$, but at most with probability $2^{-l(\eta)}$ on $([\![\vec{t}[n]]\!], [\![n']\!])$, the fresh $[\![n']\!]$ being independent of its view by disjointness of tape slices; non-negligible success would break the premise. The conclusion's non-negligible-set reading has no counterpart in Squirrel; what is expressible there, through the helper `adv` [5], is the world-free $\tilde{\forall} g.\ \mathtt{adv}(g) \tilde{\Rightarrow} [g(\vec{t}[n]) \neq n]$, equivalent to it by a theorem of our semantics, not theirs.

With $\vec{t}$ not containing n, the hiding rule's premise is the freshness axiom of indistinguishability CCSA and its conclusion the impossibility of nonce guessing [12] that motivated our domain: over arbitrary random variables the rule is unsound—n itself extracts n—while over attacker computations it is sound.

IX. Global and Local Formulas in Squirrel Prover

The Squirrel prover [3] initially defined a meta-logic atop CCSA, then refined it in [5] to align more closely with [7], particularly in how security axioms are formalized using recursive functions for game execution.

Nevertheless, crucial differences remain. One is that Squirrel defines local formula satisfaction over traces, causing the witness and disjunction problems of the Introduction and forcing a global domain of arbitrary random variables—including non-PPT ones and those violating causal structure (depending on the future).

Another difference is that Squirrel allows quantification over term contexts (functions). For example, $\forall f : \mathsf{message} \rightarrow \mathsf{message}.\ \phi(f)$ where f ranges over functions from messages to messages. On each trace ρ, such f determines a concrete bitstring-to-bitstring function, so local quantification over f means: for every choice of such function on the current trace. When f's target type is bool, this becomes quantification over local predicates, making local formulas second-order over traces.

A. The Two-Tiered Structure: Syntax and Semantics

The key distinction in Squirrel is between two formula levels with fundamentally different quantification domains.

a) Local formulas: are Boolean *terms* in disguise—terms of the $\mathsf{if_then_else_}$ calculus mentioned in Section VIII. A predicate application is a bool-typed term, and a connective is its if-then-else combination, so a quantifier-free local formula ϕ *is* the Boolean term b_ϕ of that section, evaluating to 1 on exactly the traces where ϕ holds. What Squirrel adds is quantification, and it stays inside the term language: for each type τ it introduces *function symbols* $\forall_\tau, \exists_\tau : (\tau \rightarrow \mathsf{bool}) \rightarrow \mathsf{bool}$, and $\forall x.\phi$ abbreviates the term $\forall_\tau(\lambda x.\phi)$—again a Boolean term—whose trace-wise evaluation is

$$[\![\forall_\tau(\lambda x.\phi)]\!]^\nu_{\eta,\rho} = 1 \text{ iff } [\![\phi]\!]^{\nu[x\mapsto a]}_{\eta,\rho} = 1 \text{ for all } a \in [\![\tau]\!]_\eta$$

A quantified local formula is thus still a term, returning 0 or 1 on each trace separately. When evaluating $\forall x.\phi$ on trace ρ, variable x ranges over values on that specific trace. Since one can quantify over bool-typed functions, local formulas have second-order expressive power over individual traces.

b) Global formulas: are first-order formulas over the domain of *random variables*—functions from traces to values :

$$\begin{aligned} F \;::=\; & [\phi] \mid \vec{t} \sim \vec{u} \mid \mathtt{const}(t) \mid \mathtt{adv}(t) \mid F \tilde{\wedge} F \\ & \mid F \tilde{\vee} F \mid F \tilde{\Rightarrow} F \mid \tilde{\forall}(x{:}\tau).F \mid \tilde{\exists}(x{:}\tau).F, \end{aligned}$$

with `const` and `adv` discussed in Section IX-C. Global quantifiers $\tilde{\forall} x.F$ range over *arbitrary* random variables that

can take different values on different traces. The overwhelming truth $[\phi]$ means $\Pr_\rho[[\![\phi]\!]_{\eta,\rho} = \textbf{\textit{true}}] \geq 1 - \text{negl}(\eta)$. The rest follows indistinguishability CCSA semantics except the domain is arbitrary random variables, not just PPT algorithms.

c) Sequents.: The proof system of [5] manipulates two kinds of judgments, mirroring the two levels. A *global sequent* $E; \Theta \vdash F$ consists of a typing environment E, a set Θ of global formulas as hypotheses, and a global formula F as goal; it is valid iff $\tilde{\bigwedge}\Theta \tilde{\Rightarrow} F$ is satisfied by all models of E. A *local sequent* $E; \Theta; \Gamma \vdash \phi$ adds a set Γ of *local* hypotheses and a local goal ϕ; it is valid iff

$$\tilde{\bigwedge}\Theta \tilde{\Rightarrow} [\bigwedge \Gamma \Rightarrow \phi]$$

is satisfied by all models of E. The local hypotheses and the local goal live *inside* the bracket: a local sequent asserts the overwhelming truth of a single implication, with only the global hypotheses standing outside. E records locally and globally quantified variables alike, which [5] justifies by its Proposition 2, the quantifier commutation $[\forall(x{:}\tau).\phi] \tilde{\Leftrightarrow} \tilde{\forall}(x{:}\tau).[\phi]$. Both features are put to work in Section X.

B. Higher-Order Features and Many-Sorted Logic

There are two aspects of higher order in Squirrel. One is at the meta-level: encode axiom schema of CCSA using inductive constructions. This does not affect our analysis. The other is at the object level; we restrict our attention to the latter in this paper.

Squirrel's logic is "higher-order" at the object level because quantification over function-typed variables is allowed, making it second-order over individual traces. However, as the syntax of [5] shows for *global* formulas—which matter for us as we only quantify over those—Squirrel's quantification is many-sorted first-order.

At the global level, quantification is over random variables (functions from traces to values). Writing $\tilde{\forall} f : (\mathsf{message} \to \mathsf{message}).\, F$, the variable f ranges over random variables of function type. This is *many-sorted first-order logic*: types generate sorts (base types plus function spaces), application is a family of function symbols, and quantifiers bind variables over sort elements. Section X starts from this observation.

C. Case Analysis and the `const` Predicate

As explained in [2], authentication properties take forms like $[\forall\tau.\text{happens}(\tau) \Rightarrow \text{condition@}\tau \Rightarrow \exists\tau'.\tau' < \tau \wedge \text{event@}\tau']$, expressing that with overwhelming probability, if a condition is checked at any time point, some event happened earlier. Such formulas require global (not trace-wise) arguments. For a protocol with only actions A and B, this reduces to proving

$$\tilde{\forall}(\tau : \mathtt{timestamp}).[\tau = A \vee \tau = B] \tilde{\Rightarrow} [\phi]$$

The problem: when Squirrel quantification is global, witnesses can vary uncontrollably across traces. Antoine and Baelde work around this by assuming τ is *constant*:

$$\tilde{\forall}(\tau : \mathtt{timestamp}).\mathtt{const}(\tau) \tilde{\Rightarrow} [\tau = A \vee \tau = B] \tilde{\Rightarrow} [\phi]$$

D. Squirrel's Machinery for the `const` Predicate

The $\mathtt{const}(t)$ predicate prevents arguments from varying badly across traces. Since trace-by-trace quantification forces arbitrary random variables globally, $\mathtt{const}(t)$ ensures good behavior. It asserts t evaluates to the same value for all ρ, at each η, serving two purposes:

1) *Constant bitstrings*: A timestamp variable x is constant if it denotes the same timestamp across all traces.
2) *Constant handles*: $\mathtt{const}(k)$ for a key name k means k corresponds to a single key generation across uses, not independent samplings per trace.

Squirrel provides inference rules for constancy, including G.CONST:APP (application preserves constancy) and L.R-∀-CONST (universal introduction with $\mathtt{const}(x)$ hypothesis for finite fixed types). Proposition 4 of [5] establishes that for finite fixed types:

$$\mathcal{M} \models [\forall(x:\tau).\phi] \quad \text{iff} \quad \mathcal{M} \models \tilde{\forall}(x:\tau).\mathtt{const}(x) \tilde{\Rightarrow} [\phi]$$

Section X shows that under the twisted semantics this constancy machinery becomes unnecessary: the Barcan content of the twist plays the role of `const`.

X. How Fitting's Embedding Helps Squirrel

In this section we show how we ca apply our BC^+ fragment in which Squirrel inference rules remain sound but can be simplified: the `const` machinery can be removed entirely. Moreover it seems to us that all of Squirrel's non-cryptographic rules can be derived from our BC^+ rules.

a) **Many-sorted logic and "higher-order" features.:** For soundness, quantification over functions can be treated as many-sorted first-order quantification. For example, $[\forall f : (\text{message} \to \text{message}).\forall x : \text{message}.f(f(x)) = f(x) \to f(x) = x]$ is many-sorted FOL with sorts `message` and $(\text{message} \to \text{message})$. The Fitting twist gives:

$$\forall f.\forall x.\Box((f(f(x)) = f(x))^* \to (f(x) = x)^*)$$

All benefits of Fitting's approach apply directly to Squirrel's full term language without separate machinery for "higher-order" features. In particular, in the many-sorted setting the computability predicate becomes expressible through Squirrel's `adv` [5], which holds of a value when some PPT algorithm computes it from the adversarial tape: the $[\neg\, \vec{t} \triangleright u]$ of the hiding rule (Section VIII) is $[\neg\exists g.\ \mathtt{adv}(g) \wedge g(\vec{t}) = u]$, the quantifier ranging over function-typed elements. Quantifying instead over a sort of adversarial *symbols* would be weaker: a signature's interpretation need not exhaust, or be closed under composition of, the attacker computations.

b) **Encoding Squirrel's rules.:** By switching to $[\phi] \equiv \phi^*$, Squirrel's inference rules remain sound. Recall from Section IX that the semantics of a local sequent $E; \Theta; \Gamma \vdash \phi$ is $\tilde{\forall} E.\ (\tilde{\bigwedge}\Theta) \tilde{\Rightarrow} [(\bigwedge \Gamma) \Rightarrow \phi]$. Under the twisted reading the bracketed implication unfolds, by the $\to$- and $\wedge$-clauses of the twist, to the same string as a boxed sequent:

$$[\bigwedge \Gamma \to \phi] \;\equiv\; \Box([\bigwedge \Gamma] \to [\phi]),$$

so for us a local sequent *is* a boxed judgment. Once we switch to the non-negligible-set interpretation, we can therefore write it as

$$E;\ \Theta;\ [\Gamma]\ \vdash\ [\phi]$$

meaning that when Θ holds globally, then on any non-negligible set, the overwhelming truth of Γ implies the overwhelming truth of ϕ—here $[\Gamma]$ abbreviates $\{[\gamma] : \gamma \in \Gamma\}$, and the variables of E are quantified over the domain. This is how we can unify the trace-wise first-order arguments via Fitting, and the cryptographic ones, satisfaction being on non-negligible sets.

This can now be said precisely, checked against the full rule inventory of [5]. The translation into the fragment is uniform, and keeps the typing: a typed variable of E becomes a quantified domain variable of the same type, `adv` among the types; the global connectives and quantifiers $\tilde{\wedge}, \tilde{\vee}, \tilde{\Rightarrow}, \tilde{\forall}, \tilde{\exists}$ become their classical counterparts; $[\phi]$ is kept; `const` atoms become $\top$; $\sim$-atoms stay the same. Writing F° for the result, a global sequent $E; \Theta \vdash F$ reads $\forall \vec{x}_E.(\bigwedge \Theta^\circ \to F^\circ)$, and a local sequent $E; \Theta; \Gamma \vdash \phi$ reads $\forall \vec{x}_E.(\bigwedge \Theta^\circ \to [\bigwedge \Gamma \to \phi])$: fragment formulas when Θ is $\sim$-free, formulas of the unified language of Section VIII otherwise.

c) **L.REWRITE-EQUIV rule.:** One rule we found that does not translate in general:

$$\frac{E; \Theta; \Gamma_1 \vdash \phi_1 \quad E; \Theta \vdash (\Gamma_0 \to \phi_0) \sim (\Gamma_1 \to \phi_1)}{E; \Theta; \Gamma_0 \vdash \phi_0}$$

With $\sim$ in the unified language the rule is statable, and its quantifier-free instances are derivable: the bridge $b_\phi \sim \textbf{\textit{true}} \leftrightarrow [\phi]$ turns the local sequents into $\sim$-atoms, and symmetry and transitivity of $\sim$—sound, and standard axioms of indistinguishability CCSA [11]—transfer one to the other. The general form, however, is unstatable: lacking trace-wise quantification, we cannot encode a quantified $\Gamma \to \phi$ as a term, unless its quantifiers range over finite domains and expand into disjunctions; we are not sure where this would be important.

d) **Derivability of Squirrel's non-cryptographic rules.:** It seems to us that, under this translation, every non-cryptographic, $\sim$-free rule of the Squirrel proof system (Figure 7 of [5])—the `const`-bearing ones included—is derivable in BC^+. We have checked the inventory rule by rule, and state the outcome as a checked claim rather than a theorem, since it rests on our reading and translation of their rules. The verification: a local sequent $E; \Theta; \Gamma \vdash \phi$ is the judgment $\Theta \vdash [\bigwedge \Gamma \to \phi]$ under the outer quantifiers, so the Fitting collapse—(N), (K) and the conjunction principle—reduces any rule whose sequents share Θ to a single classical entailment between matrices, which for each of Squirrel's local rules is exactly its own classical content: the propositional rules, L.AXIOM, L.CUT-LOC and L.REWRITE directly, the last by equality as congruence; L.R-$\forall$, L.L-$\forall$, L.R-$\exists$ and L.L-$\exists$ through the Barcan equivalence $[\forall x.\phi] \leftrightarrow \forall x.[\phi]$, that is (BF) with its derivable converse—their own soundness proofs invoke Proposition 2 of [5] at exactly this point, and for us that proposition *is* (BF). The rules manipulating Θ, and all global-sequent rules, are outer classical logic. Of the mixed rules, G.BYLOC is an identity; L.BYGLOB is weakening under the box, $[\phi] \to [\bigwedge \Gamma \to \phi]$; and L.LOCALISE, moving $[\phi]$ from Θ to ϕ in Γ, is one application of the collapse to $(\bigwedge \Gamma \wedge \phi \to \psi), \phi \vdash_{\mathrm{FOL}} \bigwedge \Gamma \to \psi$. L.R-$\forall$-CONST dissolves into L.R-$\forall$, its side conditions carrying no content here; and G.REWRITE is congruence inside the brackets—in the presence of bare atoms it must stay there, an overwhelming equality not licensing pointwise rewriting.

This makes sense since our fragment is *closed*: any rule whose translation is a theorem-schema of $\mathsf{FOS4}_P$—in particular any computationally sound non-cryptographic rule is derivable in BC^+ by Theorem 5.4.

e) **Case study (splitting and patching).:** Consider Squirrel's case rule, written as boxed judgments:

$$\frac{E; \Theta; [\Gamma], [\phi_1] \vdash [\psi] \quad E; \Theta; [\Gamma], [\phi_2] \vdash [\psi]}{E; \Theta; [\Gamma], [\phi_1 \vee \phi_2] \vdash [\psi]}$$

for arbitrary global hypotheses Θ: this is the splitting-and-patching derivation of Section VII verbatim, its persistence requirement met automatically by the bracketed hypotheses. Equivalently, inside the box it is first-order disjunction elimination, the formula $[\bigwedge \Gamma \wedge \phi_1 \to \psi] \wedge [\bigwedge \Gamma \wedge \phi_2 \to \psi] \to [\bigwedge \Gamma \wedge (\phi_1 \vee \phi_2) \to \psi]$ being valid by Fitting for us, while trace-by-trace for Squirrel. The real difference is not the shape of this rule, but what may prove its premises: under the twisted reading a local sequent asserts $\Box([\bigwedge \Gamma] \to [\psi])$, a statement about non-negligible sets, so it may be established by cryptographic arguments that are not trace-wise valid. Squirrel's local calculus, by contrast, is sound trace by trace, and genuinely probabilistic facts must be routed through the global tier and its `const` machinery.

f) **Case analysis without the `const` restriction.:** Returning to Section IX-C's case distinction, it is clear what happens: prove the $\tau = A$ and $\tau = B$ branches as local sequents—by cryptographic arguments if need be—and patch them with the above rule.

g) **Existential witnesses.:** Fitting's $(\exists x.\phi)^* = \Box\Diamond\exists x.\phi^*$ does *not* require a single witness working across all traces. It means: for all non-negligible S', there exists non-negligible $S'' \subseteq S'$ with $\exists x.\phi^*$ holding throughout S''. Witnesses can vary between non-negligible sets while remaining constant within each.

h) **Working with the existential quantifier.:** Consider the authentication goal of Section IX-C: with $\mathrm{cond@}\tau$ the event that the message–tag pair received at τ passes MAC verification, and $\mathrm{event@}\tau'$ that the honest agent produced that tag at τ', the local sequent

$$E; \Theta; [\Gamma], [\mathrm{cond@}\tau]\ \vdash\ [\exists \tau'.\, \tau' < \tau \wedge \mathrm{event@}\tau']$$

asserts $\Box([\bigwedge \Gamma \wedge \mathrm{cond@}\tau] \to [\exists \tau' \ldots])$, with Γ collecting the execution facts. Without going into detail, the existence of τ' is delivered by the unforgeability of a signature (there had to be a τ' where the signature was created), and the steps are perfectly sound on non-negligible sets: within each, there has to be another—however small, but still non-negligible—

on which a *single* τ' works, as there are only finitely many previous steps. As noted in the previous paragraph, such a τ' does not have to exist globally, as the bracket puts a $\Box\Diamond$ in front of the existential quantifier.

i) **Universal quantification.:** In Squirrel, relating local and global quantifiers requires (Proposition 4 of [5]) for finite fixed types: $\mathcal{M} \models [\forall x.\phi] \iff \mathcal{M} \models \tilde{\forall} x.\mathtt{const}(x) \Rightarrow [\phi]$. With Fitting's embedding, this always holds: $\forall x \phi^* \leftrightarrow (\forall x \phi)^*$, i.e., $\forall x [\phi] \leftrightarrow [\forall x \phi]$ for any type—constant, PPT, or arbitrary random variable. In our semantics there are no separate local and global quantification as everything is global.

j) **Determined formulas.:** Squirrel needs determined formulas ($[\phi] \vee [\neg\phi]$ holds) with additional inference rules to handle case distinctions. With our semantics, determined formulas remain available if needed as $[\phi] \vee [\neg\phi]$.

k) **Eliminating the `const` restriction.:** The entire `const` infrastructure—predicate, type labels {`fixed`, `finite`}, and restricted inference rules—becomes unnecessary. Consider key Squirrel rules requiring constancy; note that in each, the local formula sits bracketed in the *global* zone—exactly the position to which our boxed judgments give a local reading:

G.R-LOC:$\rightarrow$: $\frac{E;\Theta,[\phi]\vdash[\psi] \quad E;\Theta\vdash\mathtt{const}(\phi)}{E;\Theta\vdash[\phi\Rightarrow\psi]}$

L.R-$\forall$-CONST: $\frac{\{\mathtt{fixed},\mathtt{finite}\}\subseteq\mathtt{label}(\tau) \quad E,x{:}\tau;\Theta,\mathtt{const}(x);\Gamma\vdash\phi}{E;\Theta;\Gamma\vdash\forall(x{:}\tau).\phi}$

G.L-LOC:$\vee$: $\frac{E;\Theta,[\phi]\vdash F \quad E;\Theta,[\psi]\vdash F \quad E;\Theta\vdash\mathtt{const}(\phi)\bar{\vee}\mathtt{const}(\psi)}{E;\Theta,[\phi\vee\psi]\vdash F}$

In Fitting's framework the `const` side conditions dissolve. The key is to write these axioms with a ";" in front of the bracketed formulas, which we can do at least when F is also bracketed. In that case the axioms hold on any non-negligible set: the whole space, and arbitrarily small ones. Read so, they all follow from Fitting's theorem and the splitting-and-patching derivation above. The one genuinely global instance is case analysis under an arbitrary conclusion F, where it can be stated for determined formulas—and for those we do not need `const`, as we described above.

l) **One tier, not two.:** Squirrel's two calculi are articulated by mixed rules, chiefly G.BYLOC: from $E;\Theta;\emptyset \vdash \phi$ conclude $E;\Theta \vdash [\phi]$ [5]. Under the twist this rule is not simplified but dissolved: a local sequent *is* the global judgment, so there is nothing left to articulate. The corresponding bridge is the one non-invertible rule of the calculus of [2] and the delicate point of its completeness proof; in the one-tier reading no such commitment point exists. More generally, the one-tier reading merges rules the two-tier system must state twice: Squirrel's global and local axiom rules, together with the mixed $E;\Theta,[\phi];\Gamma\vdash\phi$, are the single identity axiom; the local $\Rightarrow$-introduction and the global G.R-LOC:$\rightarrow$ become one rule, the `const` side condition of the latter carrying no content here; local and global $\forall$-introductions are identified by (BF).

m) **Cut-elimination.:** In the two-tier setting the main obstacle to cut-elimination is the committing bridge G.BYLOC, which, according to the above, is not needed any more.

XI. TWISTED SEMANTICS FOR CRYPTOVAMPIRE'S EVALUATED LOGIC

CryptoVampire uses bounded Skolemization to avoid nonce-guessing oracle functions. We show this issue does not emerge with BC twisted semantics: guessing nonces on non-negligible sets (not traces) is impossible. We also show how CryptoVampire's hybrid model works with twisted semantics. Importantly, CryptoVampire's syntax stays first-order—no S4 or BC$^+$ is needed. Eliminating bounded Skolemization comes for free.

CryptoVampire [26] automates proofs in the BC Logic framework by encoding into first-order logic for saturation-based provers. It focuses on *trace properties* (not equivalence), so indistinguishability $\sim$ is not involved—making it a natural match for BC's original reachability framework (Section III-B).

CryptoVampire's *Evaluated Logic* separates symbolic from computational reasoning through an evaluation predicate. However, the trace-by-trace computational semantics requires complex bounded Skolem machinery. We propose retaining CryptoVampire's symbolic structure while replacing trace-by-trace semantics with BC twisted semantics.

A. CryptoVampire's Evaluated Logic

CryptoVampire distinguishes *symbolic* predicates (depending only on term structure) from *evaluated* predicates (depending on computational execution). Since subterms cannot be checked on bitstrings, it pairs the computational model with a symbolic one over ground terms.

a) Syntax.: Symbolic terms: $t ::= x \mid n[\vec{\imath}] \mid f[\vec{\imath}](\vec{t}) \mid \mathsf{input}(T) \mid \ldots$ where $n[\vec{\imath}]$ are indexed nonces and $\mathsf{input}(T)$ is attacker input at step T. *Formulas* combine symbolic predicates ($t = t'$ for syntactic identity, $t \sqsubseteq t'$ for subterm relation) with evaluation predicate $|t|$ for Boolean terms, plus standard connectives.

b) The EQ function symbol.: The notation $|t| = |t'|$ abbreviates $|\mathsf{EQ}(t,t')|$, where EQ produces a Boolean term evaluating to 1 iff t and t' compute to the same bitstring—the same Boolean equality as in Section IV.

c) The Symbolic Model.: A symbolic model M interprets predicates over ground terms independently of tapes: $M(=)(t,t')$ holds iff terms are syntactically identical; $M(\sqsubseteq)(t,t')$ holds iff t is a subterm of t'.

d) CryptoVampire's Trace-by-Trace Semantics.: The original semantics defines a cryptographic model $\mathcal{L} = (\mathcal{M}, \mathbb{T}, M)$. The interpretation is:

$$\begin{aligned} \mathcal{L}^\eta(t \sqsubseteq t')(\rho) &:= M(\sqsubseteq)(t,t') && \text{(symbolic)} \\ \mathcal{L}^\eta(t = t')(\rho) &:= M(=)(t,t') && \text{(symbolic)} \\ \mathcal{L}^\eta(|t|)(\rho) &:= \min(1, [\![[t]^{\mathbb{T}}]\!]_{\eta,\rho}) && \text{(evaluated)} \end{aligned}$$

Cryptographic satisfiability requires $\Pr_\rho(\mathcal{L}^\eta(\phi)(\rho) \neq 1) \leq \mathrm{negl}(\eta)$. As [26] notes, $|t|$ is "defined *trace-wise*"—creating the issue we address.

B. The Nonce Guessing Problem

The trace-by-trace semantics creates issues with existential quantification over terms. CryptoVampire's Theorem 2 (No Guessing) states that for any ground term m:

$$|n[\vec{i}]| = |m| \Rightarrow n[\vec{i}] \sqsubseteq m$$

This axiom is cryptographically valid: if the ground term m evaluates to the same bitstring as the nonce $n[\vec{i}]$, then m must contain $n[\vec{i}]$ syntactically since guessing a random nonce is negligible from terms that are not created using the nonce. Note that in their semantics the domain is not arbitrary PPT algorithms, but only those that are created from closed terms.

However, Example 16 of [26] shows the universally quantified version fails:

$$\forall x.\, (|n[\vec{i}]| = |x| \Rightarrow n[\vec{i}] \sqsubseteq x)$$

is invalid in trace-by-trace semantics. The problem is as follows. In trace-by-trace semantics, evaluating $\forall x.\phi$ on random tape ρ requires ϕ to hold for all ground terms x. On one fixed string ρ_1 on the tape, some term t_1 might happen to match the nonce value; on another fixed string ρ_2 on the tape (two disjoint sections of the honest tape), a different term t_2 matches. When the formula is Skolemized to prove the negation, we get a Skolem function sk such that $|n[\vec{i}]| = |\mathsf{sk}|$ holds on every random tape—but sk is a non-computable oracle that "guesses" the nonce value differently on each tape.

Example 17 of [26] makes this concrete: assuming the term algebra contains 0 and successor $S(\cdot)$, for any random tape $\rho = (\rho_h, \rho_a)$, the nonce $n[\vec{i}]$ (which is a slice of the honest tape ρ_h) evaluates to some natural number N, and the term $S^N(0)$ evaluates to the same value. The Skolem witness can pick this term differently for each ρ, violating the no-guessing principle.

CryptoVampire addresses this through *bounded Skolem normal forms* (Definition 18), restricting Skolemization to ensure witnesses come from a finite syntactic domain. This is effective but adds considerable complexity to the metatheory.

C. The BC Twisted Semantics for CryptoVampire

Since CryptoVampire handles only trace properties (not equivalence), the original BC reachability semantics (Section III-B) applies directly. We propose keeping CryptoVampire's symbolic structure while replacing the trace-by-trace semantics with the BC twisted semantics. To emphasize that CryptoVampire's language remains purely first-order, we define the semantics *directly*, without modal operators.

Definition 11.1 (BC Twisted Semantics for CryptoVampire): Let $\mathcal{M}^{\mathcal{C}}$ be a computational model, M a symbolic model, S a non-negligible subset of traces, and $\nu : \mathrm{Var} \to D$ a valuation into the domain D of ground terms.

Symbolic predicates (trace-independent, interpreted by M):

- $(\mathcal{M}^{\mathcal{C}}, M), S, \nu \models_{BC} (t = t')$ iff $M(=)(\nu(t), \nu(t'))$
- $(\mathcal{M}^{\mathcal{C}}, M), S, \nu \models_{BC} (t \sqsubseteq t')$ iff $M(\sqsubseteq)(\nu(t), \nu(t'))$

These hold independently of S—if they hold at S, they hold at any $S' \subseteq S$ so we do not need the box diamond pattern.

Evaluation predicate (trace property, requires overwhelming satisfaction):

- $(\mathcal{M}^{\mathcal{C}}, M), S, \nu \models_{BC} |t|$ iff for all non-negligible $S' \subseteq S$, there exists a non-negligible $S'' \subseteq S'$ such that for all η and all $\rho \in S''_\eta$, $[\![[\nu(t)]^{\mathbb{T}}]\!]_{\eta,\rho} = 1$. Equivalently, the traces of S on which $\nu(t)$ does not evaluate to 1 form a negligible set.

In other words, $|t|$ is satisfied iff t evaluates to 1 at all but a negligible set of traces of S.

Propositional connectives, quantification and validity work exactly as in Definition 3.2.

Observe that the symbolic predicates ($=$, $\sqsubseteq$) are *persistent* (*local*, in the terminology of Section VIII)—they depend only on term structure, not on traces—so the "for all S', exists S''" pattern is equivalent with not having them there. Indeed all of CryptoVampire's predicates are local, the evaluation predicate being a trace property; this is why no modality is ever needed in the syntax. Only the evaluation predicate $|t|$ and the logical connectives that can "hide" it (disjunction, negation, implication, existential) require the characteristic BC twisted pattern. Here too , if $\Gamma \vdash_{\mathrm{FOL}} \phi$ in classical first-order logic, then:

$$(\mathcal{M}^{\mathcal{C}}, M) \models_{BC} \Gamma \implies (\mathcal{M}^{\mathcal{C}}, M) \models_{BC} \phi$$

which follows from the general soundness of the BC twisted semantics (Section III-B). The key point is that CryptoVampire can use standard first-order theorem provers without modification—the non-standard semantics ensures soundness automatically.

D. Resolution of the Nonce Guessing Problem

Consider the No Guessing axiom $\forall x.\, (|n[\vec{i}]| = |x| \Rightarrow n[\vec{i}] \sqsubseteq x)$ under the BC twisted semantics (Definition 11.1).

Proposition 11.2: The formula $\forall x.\, (|n[\vec{i}]| = |x| \Rightarrow n[\vec{i}] \sqsubseteq x)$ is valid in the BC twisted semantics.

Proof: By Definition 11.1, validity means: for all ground terms $t \in D$ and all non-negligible sets S,

$$(\mathcal{M}^{\mathcal{C}}, M), S \models_{BC} |n[\vec{i}]| = |t| \Rightarrow n[\vec{i}] \sqsubseteq t$$

By the clause for implication, this means: for all non-negligible $S' \subseteq S$, if $(\mathcal{M}^{\mathcal{C}}, M), S' \models_{BC} |n[\vec{i}]| = |t|$, then $(\mathcal{M}^{\mathcal{C}}, M), S' \models_{BC} n[\vec{i}] \sqsubseteq t$.

If $n[\vec{i}] \sqsubseteq t$, the consequent holds outright, independently of S'. Now suppose $n[\vec{i}] \not\sqsubseteq t$, i.e., the nonce does not appear syntactically in t. Then t's value is computed from tape bits disjoint from the nonce's slice. By the clause for the evaluation predicate, $|\mathsf{EQ}(n[\vec{i}], t)|$ at S' would in particular yield a non-negligible $S'' \subseteq S'$ on which EQ evaluates to 1 at every trace. But $S'' \subseteq \{\rho : [\![t]\!]_{\eta,\rho} = [\![n[\vec{i}]]\!]_{\eta,\rho}\}$, and by independence this match set has measure at most $2^{-\lambda(\eta)}$, with λ the nonce length—negligible. A non-negligible subset of a negligible set is impossible.

Therefore, if $n[\vec{i}] \not\sqsubseteq t$, then $(\mathcal{M}^{\mathcal{C}}, M), S' \not\models_{BC} |n[\vec{i}]| = |t|$ for any non-negligible S', making the implication true. □

This is the *locality of existential witnesses* of Section VII: an existential under the twist needs one witness working on a whole non-negligible set, never one witness per trace. Trace by trace, the negated axiom $\exists x.\,(|n[\vec{i}]| = |x| \wedge n[\vec{i}] \not\sqsubseteq x)$ is satisfiable, because the witness may change from tape to tape—this is exactly the Skolem oracle of Example 16. Under the twisted semantics, the witness must be a *single* ground term matching the nonce on a whole non-negligible set, which the proof above rules out.

XII. Conclusion

We presented a one-tier framework for cryptographic protocol reasoning combining local and global properies: Bana–Comon's twisted semantics, read through Fitting's embedding into first-order S4 with non-negligible sets as worlds, unifies trace properties, PPT computability, and indistinguishability in a single language with a single satisfaction relation. On the fragment where modalities occur as $[\cdot]$ (The Fitting-twist extended to nesting), the calculus BC^+ with the persistence instances is sound and derives exactly the theorems of $\mathsf{FOS4}_P$; we prove computational completeness; cut elimination holds propositionally and at first order all cuts confine to a single Barcan shape. For quantifier-free formulas the semantics coincides with Antoine–Baelde's, and over random-variable domains under Squirrel's finite-discrete tape convention with Squirrel's semantics, resolving the two open problems of [2].

For the tools, the framework simplifies for free: Squirrel's local sequents become boxed judgments, the `const` machinery dissolves into the persistence that brackets carry, and its non-cryptographic rules can be derived; CryptoVampire keeps its first-order syntax while bounded Skolemization becomes unnecessary.

Limitations and future work: completeness is relative to congruence equality, constructor names, and measurable predicates—the fixed cryptographic name interpretation and literal equality remain open, as does an axiomatization covering $\sim$, which would subsume a quantified version of [2]; full first-order cut elimination is expected only in labelled calculi, where it is known for quantified S4 with constant domains [25]; and a proof-assistant implementation remains future work.

Acknowledgment on the Usage of AI

Anthropic's Claude Fable 5 and Open AI's ChatGPT 5.6 Pro were used to gather known results in the literature, check the correctness of our proofs and help with clear presentation.

References

[1] C. Abate, P. G. Haselwarter, E. Rivas, A. Van Muylder, T. Winterhalter, C. Hriţcu, K. Maillard, and B. Spitters. SSProve: A foundational framework for modular cryptographic proofs in Coq. In *34th IEEE Computer Security Foundations Symposium (CSF)*, pages 608–626. IEEE, 2021. Distinguished Paper Award.

[2] T. Antoine and D. Baelde. Propositional logics of overwhelming truth. In *32nd EACSL Annual Conference on Computer Science Logic (CSL 2025)*, LIPIcs. Schloss Dagstuhl, 2025.

[3] D. Baelde, S. Delaune, C. Jacomme, A. Koutsos, and S. Moreau. An interactive prover for protocol verification in the computational model. In *42nd IEEE Symposium on Security and Privacy (SP)*, pages 537–554. IEEE, 2021.

[4] D. Baelde, S. Delaune, and S. Riou. SMT-based automation for overwhelming truth. In *38th IEEE Computer Security Foundations Symposium (CSF)*. IEEE, 2025.

[5] D. Baelde, A. Koutsos, and J. Lallemand. A higher-order indistinguishability logic for cryptographic reasoning. In *38th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*, pages 1–13, 2023.

[6] David Baelde, Adrien Koutsos, and Justine Sauvage. Leveraging cryptographic simulator synthesis for formally verifying the FOO e-voting protocol. In *35th USENIX Security Symposium (USENIX Security 2026)*, Baltimore, USA, 2026. hal-05453231.

[7] G. Bana, R. Chadha, A. K. Eeralla, and M. Okada. Verification methods for the computationally complete symbolic attacker based on indistinguishability. *ACM Transactions on Computational Logic*, 20(1):1–43, 2019.

[8] G. Bana, R. Chadha, A. K. Eeralla, and Q. Zhang. Mechanized analysis of vote privacy using computationally complete symbolic attacker. Coq development, github.com/gebana4/CCSA_FOO_Verification_Public, 2023.

[9] G. Bana and H. Comon-Lundh. Towards unconditional soundness: Computationally complete symbolic attacker. Available at IACR ePrint Archive, Report 2012/019.

[10] G. Bana and H. Comon-Lundh. Towards unconditional soundness: Computationally complete symbolic attacker. In *Proceedings of the 1st International Conference on Principals of Security and Trust (POST'12)*, LNCS, pages 189–208. Springer, 2012.

[11] G. Bana and H. Comon-Lundh. A computationally complete symbolic attacker for equivalence properties. In *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security (CCS '14)*, pages 609–620. ACM, 2014.

[12] G. Bana, K. Hasebe, and M. Okada. Computationally complete symbolic attacker and key exchange. In *Proceedings of the 20th ACM SIGSAC Conference on Computer and Communications Security (CCS '13)*, pages 1231–1246. ACM, 2013.

[13] G. Bana and M. Okada. Semantics for "Enough-Certainty" and Fitting's Embedding of Classical Logic in S4. In *25th EACSL Annual Conference on Computer Science Logic (CSL 2016)*, volume 62 of *LIPIcs*, pages 34:1–34:18. Schloss Dagstuhl, 2016.

[14] G. Barthe, B. Grégoire, S. Heraud, and S. Zanella-Béguelin. Computer-aided security proofs for the working cryptographer. In *Advances in Cryptology – CRYPTO 2011*, volume 6841 of *LNCS*, pages 71–90. Springer, 2011.

[15] D. Basin, A. Lochbihler, and S. R. Sefidgar. CryptHOL: Game-based proofs in higher-order logic. *Journal of Cryptology*, 33(2):494–566, 2020.

[16] P. Blackburn, M. de Rijke, and Y. Venema. *Modal Logic*, volume 53 of *Cambridge Tracts in Theoretical Computer Science*. Cambridge University Press, 2001.

[17] B. Blanchet. An automatic security protocol verifier based on resolution theorem proving (invited tutorial). In *CADE'05*, Tallinn, Estonia, July 2005.

[18] B. Blanchet. A computationally sound mechanized prover for security protocols. *IEEE Transactions on Dependable and Secure Computing*, 5(4):193–207, 2008.

[19] B. Blanchet and C. Jacomme. Post-quantum sound CryptoVerif and verification of hybrid TLS and SSH key-exchanges. In *37th IEEE Computer Security Foundations Symposium (CSF)*, pages 543–556. IEEE, 2024.

[20] B. Blanchet and D. Pointcheval. Automated security proofs with sequences of games. In *Advances in Cryptology – CRYPTO 2006*, volume 4117 of *LNCS*, pages 537–554. Springer, 2006.

[21] C. Cremers, C. Fontaine, and C. Jacomme. A logic and an interactive prover for the computational post-quantum security of protocols. In *IEEE Symposium on Security and Privacy (S&P)*, pages 125–141. IEEE, 2022.

[22] D. Dolev and A. Yao. On the security of public-key protocols. In *Proc. 22-nd Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 350–357, 1981.

[23] M. Fitting. An embedding of classical logic in s4. *The Journal of Symbolic Logic*, 35(4):529–534, 1970.
[24] M. Fitting. *Proof Methods for Modal and Intuitionistic Logics*, volume 169 of *Synthese Library*. D. Reidel Publishing Company, 1983.
[25] M. Fitting and R. L. Mendelsohn. *First-Order Modal Logic*, volume 277 of *Synthese Library*. Kluwer Academic Publishers, 1998.
[26] S. Jeanteur, L. Kovács, M. Maffei, and M. Rawson. CryptoVampire: Automated reasoning for the complete symbolic attacker cryptographic model. In *2024 IEEE Symposium on Security and Privacy (SP)*, pages 3165–3183, San Francisco, CA, USA, 2024. IEEE.
[27] S. Negri. Proof analysis in modal logic. *Journal of Philosophical Logic*, 34(5–6):507–544, 2005.
[28] F. Poggiolesi. A cut-free simple sequent calculus for modal logic S5. *The Review of Symbolic Logic*, 1(1):3–15, 2008.

APPENDIX

A. Equivalence Proofs

We provide the full proofs of the equivalence results of Section III-B.

Proposition 3.4.*[Equivalence for Quantifier-Free Formulas] Let ϕ be a quantifier-free formula. Then ϕ holds computationally (i.e., $\mathcal{M}^{\mathcal{C}} \models_{BC} \phi$) if and only if ϕ holds with overwhelming probability in the sense of [2].*

Proof:

We prove the proposition as the case $S = \Omega^{\mathcal{C}}$ of a statement about all worlds; the relativized form is what the induction needs, as the case for $\neg$ appeals to the hypothesis at other worlds.

Since ϕ is quantifier-free it has a pointwise value $[\![\phi]\!]^{\nu}_{\eta,\rho} \in \{0,1\}$ at each trace. For a world S let $T_S(\phi)$ and $F_S(\phi)$ be the ensembles of traces of S at which ϕ is true, respectively false.

LEMMA. *For every quantifier-free ϕ, valuation ν and world S: $\mathcal{M}^{\mathcal{C}}, S, \nu \models_{BC} \phi$ iff $F_S(\phi)$ is negligible.*

We use two standard facts about negligible ensembles: (i) removing a negligible ensemble from a non-negligible one leaves a non-negligible one; (ii) a union of two negligible ensembles is negligible, so if $A \cup B$ is non-negligible then A or B is.

Proof. By induction on ϕ, recalling that $\models_{BC} \phi$ abbreviates $\models_{S4} \phi^*$.

Atoms. $\phi = P(\vec{t})$ and $\phi^* = \Box\Diamond P(\vec{t})$. We show that $\mathcal{M}^{\mathcal{KC}}, S, \nu \models_{S4} \Box\Diamond P$ if and only if P holds with overwhelming probability. Write $F = F_S(P(\vec{t}))$. If $\mathcal{M}^{\mathcal{KC}}, S, \nu \models_{S4} \Box\Diamond P$ but F were non-negligible, then since $\Box\Diamond P(\vec{t})$ holds at S, and $F \subseteq S$, there is a non-negligible $F' \subseteq F$ with $P(\vec{t})$ true throughout F', contradicting that every trace of F falsifies it. Conversely, if F is negligible then for non-negligible $S' \subseteq S$ the set $S' \setminus F$ is non-negligible by (i) and satisfies $P(\vec{t})$ throughout, so $\Diamond P(\vec{t})$ holds at S'; as S' was arbitrary, $\Box\Diamond P(\vec{t})$ holds at S.

Falsum. $\bot^* = \bot$ holds at no world and $F_S(\bot) = S$ is non-negligible.

Conjunction. $(\phi_1 \wedge \phi_2)^* = \phi_1^* \wedge \phi_2^*$, so this holds at S iff ϕ_1^* and ϕ_2^* both do, and by the induction hypothesis iff $F_S(\phi_1)$ and $F_S(\phi_2)$ are both negligible. Pointwise, a trace falsifies $\phi_1 \wedge \phi_2$ exactly when it falsifies one of the conjuncts, so $F_S(\phi_1 \wedge \phi_2) = F_S(\phi_1) \cup F_S(\phi_2)$. If both are negligible so is the union, by (ii); conversely each is contained in the union, and a subensemble of a negligible ensemble is negligible. Hence $F_S(\phi_1 \wedge \phi_2)$ is negligible iff both $F_S(\phi_i)$ are, as required.

Negation. $(\neg\psi)^* = \Box\neg\psi^*$ and $F_S(\neg\psi) = T_S(\psi)$. If $T_S(\psi)$ were non-negligible it would be a world $\subseteq S$ with $F_{T_S(\psi)}(\psi) = \emptyset$, so the hypothesis *at that world* gives ψ^* there, contradicting $\Box\neg\psi^*$ at S. Conversely, if $T_S(\psi)$ is negligible then for non-negligible $S' \subseteq S$ the set $F_{S'}(\psi) = S' \setminus T_S(\psi)$ is non-negligible by (i), so ψ^* fails at S' by hypothesis; hence $\Box\neg\psi^*$ at S.

Remaining connectives. These reduce to $\wedge$ and $\neg$. If ϕ, ψ are classically equivalent their pointwise values agree, so $F_S(\phi) = F_S(\psi)$; and Fitting's theorem in both directions gives $\phi^* \dashv\vdash_{\mathrm{S4}} \psi^*$, which by soundness of S4 for computational Kripke models makes ϕ^* and ψ^* true at the same worlds. The lemma therefore transfers along $\phi_1 \vee \phi_2 \equiv \neg(\neg\phi_1 \wedge \neg\phi_2)$ and $\phi_1 \to \phi_2 \equiv \neg(\phi_1 \wedge \neg\phi_2)$.

Proposition 3.4 is the case $S = \Omega^{\mathcal{C}}$. □

Theorem 3.7.*[Equivalence for Random Variable Domains] Let $\mathcal{M}^{\mathcal{C}}$ be a computational model with random-variable domain D and suppose that every $\Omega^{\mathcal{C}}_{\eta}$ is finite and carries the discrete σ-algebra. For any first-order formula ϕ and valuation ν:*

$$\mathcal{M}^{\mathcal{C}}, \Omega^{\mathcal{C}}, \nu \models_{BC} \phi \iff$$

ϕ holds trace-by-trace with overwhelming probability

Proof: By definition, $\mathcal{M}^{\mathcal{C}}, \Omega^{\mathcal{C}}, \nu \models_{BC} \phi$ iff $\mathcal{M}^{\mathcal{KC}}, \Omega^{\mathcal{C}}, \nu \models_{S4} \phi^*$. We proceed by structural induction on ϕ. As in Proposition 3.4, the induction statement is relativized to worlds: for every non-negligible S, ψ^* holds at S iff the set of traces of S falsifying ψ trace-by-trace is negligible; the theorem is the case $S = \Omega^{\mathcal{C}}$, and we display the arguments at $\Omega^{\mathcal{C}}$, the relativized forms being identical. Throughout, every tape space is finite and discrete, so all trace sets and all sectionwise witness selections below are measurable.

Base case (atomic formulas): Identical to Proposition 3.4.

Connectives ($\neg$, $\wedge$, $\vee$, $\to$): Identical to Proposition 3.4.

Existential quantification: Let $\phi = \exists x.\psi$. Then $(\exists x.\psi)^* = \Box\Diamond\exists x.\psi^*$.

($\Leftarrow$) For each η let

$$S_{\mathrm{good},\eta} := \{\rho \in \Omega^{\mathcal{C}}_{\eta} : \mathcal{M}^{\mathcal{C}}, \eta, \rho, \nu \models_{\mathrm{tr}} \exists x.\psi\}.$$

The complement of the ensemble S_{good} is negligible. For every η and $\rho \in S_{\mathrm{good},\eta}$ choose a witness $\mathfrak{b}^{\eta,\rho} \in D$ and put $v_{\eta,\rho} := \mathfrak{b}^{\eta,\rho}_{\eta,\rho}$. Define $\mathfrak{a}^* = \{\mathfrak{a}^*_{\eta}\}_{\eta}$ by $\mathfrak{a}^*_{\eta}(\rho) = v_{\eta,\rho}$ on $S_{\mathrm{good},\eta}$ and by an arbitrary fixed value elsewhere. Each $\mathfrak{a}^*_{\eta}$ is measurable because $\Omega^{\mathcal{C}}_{\eta}$ is finite and discrete, hence $\mathfrak{a}^* \in D$. At a trace $\rho \in S_{\mathrm{good},\eta}$, the variables $\mathfrak{a}^*$ and $\mathfrak{b}^{\eta,\rho}$ have the same value, so locality of term and predicate interpretation gives $\mathcal{M}^{\mathcal{C}}, \eta, \rho, \nu[x \mapsto \mathfrak{a}^*] \models_{\mathrm{tr}} \psi$. Now let S be any non-negligible world and put $S' = S \cap S_{\mathrm{good}}$. Removing the negligible complement of S_{good} leaves S' non-negligible. The trace-level falsity set of ψ under $\nu[x \mapsto \mathfrak{a}^*]$ is empty inside S', so the relativized induction hypothesis gives $\mathcal{M}^{\mathcal{KC}}, S', \nu[x \mapsto \mathfrak{a}^*] \models_{S4} \psi^*$. Thus $\exists x.\psi^*$ holds at S'. Since

S was arbitrary, $\Box\Diamond\exists x.\psi^*$ holds at the full space.

($\Rightarrow$) Suppose $\Box\Diamond\exists x.\psi^*$ holds but the trace-level existential is not overwhelmingly true. Its falsity ensemble

$$N_\eta := \{\rho \in \Omega_\eta^{\mathcal{C}} : \mathcal{M}^{\mathcal{C}}, \eta, \rho, \nu \not\models_{\mathrm{tr}} \exists x.\psi\}$$

is then non-negligible and hence is a world. From $\Box\Diamond\exists x.\psi^*$ there are a non-negligible $S \subseteq N$ and $\mathfrak{a} \in D$ such that $\mathcal{M}^{\mathcal{KC}}, S, \nu[x \mapsto \mathfrak{a}] \models_{S4} \psi^*$. By the relativized induction hypothesis, the traces of S at which ψ fails under this valuation form a negligible ensemble. But by the definition of N, ψ fails at every trace of S for every $\mathfrak{a} \in D$, a contradiction.

Universal quantification: Follows from the existential case using the equivalence $\forall x.\psi \equiv \neg\exists x.\neg\psi$: since $\vdash_{\mathrm{FOL}} \forall x.\psi \leftrightarrow \neg\exists x.\neg\psi$, Fitting's theorem gives the equivalence of the starred forms at every world, and the negation case of the lemma applies to quantified formulas unchanged. □

B. Equivalent Presentations of the Nested Fitting Twist

We verify Remark 4: the simplified clauses of Section III-B remain derivable when their arguments themselves contain brackets. Consequently the two presentations of $[\,\cdot\,]$—the *uniform* one, placing $\Box\Diamond$ before every subformula with a bracketed subformula counted as an atom, and the simplified one—define the same fragment up to FOS4-equivalence. Everything is derived inside FOS4 as defined in Section V, from the S4 axioms, necessitation and the Barcan formula; no semantic argument and no induction on formulas is involved.

1) Derived rules: Four consequences of K and necessitation are used throughout, together with T, 4 and BF. (Mon$\Box$): from $\vdash \alpha \to \beta$ infer $\vdash \Box\alpha \to \Box\beta$, by necessitating the premise and applying K. (Mon$\Diamond$): from $\vdash \alpha \to \beta$ infer $\vdash \Diamond\alpha \to \Diamond\beta$, by applying (Mon$\Box$) to $\neg\beta \to \neg\alpha$ and contraposing. ($\Box\wedge$): $\vdash \Box(\alpha\wedge\beta) \leftrightarrow \Box\alpha\wedge\Box\beta$, left to right by (Mon$\Box$) twice and right to left by (Mon$\Box$) on $\alpha \to (\beta \to \alpha \wedge \beta)$ with K. ($\Box\Diamond$): $\vdash \Box\alpha \wedge \Diamond\beta \to \Diamond(\alpha \wedge \beta)$, since (Mon$\Box$) on the same tautology gives $\Box\alpha \to \Box(\beta \to \alpha \wedge \beta)$ while $\Box(\beta \to \gamma) \to (\Diamond\beta \to \Diamond\gamma)$ is a theorem of K. We write T also in its dual form $\alpha \to \Diamond\alpha$ and 4 in its dual form $\Diamond\Diamond\alpha \to \Diamond\alpha$, and use replacement of provable equivalents, admissible in any normal modal logic.

2) Shape, and three facts: Every output of either presentation of the twist has the form $\Box\Diamond\theta$: each clause produces an outermost $\Box\Diamond$, as one sees by reading the clauses off, and a bracketed argument is by definition already of this form. Write α, β for formulas of this shape; they are the twists $[\Phi], [\Psi]$ of fragment formulas. Three facts are used repeatedly.

(F1) Persistence: $\alpha \to \Box\alpha$. With $\alpha = \Box\gamma$ this is axiom 4.

(F2) Stability: $\Box\Diamond\alpha \to \alpha$. With $\alpha = \Box\Diamond\theta$ this is the left-to-right half of the identity $\Box\Diamond\Box\Diamond\theta \leftrightarrow \Box\Diamond\theta$ derived in the proof of Theorem 5.2.

(F3) $\Box\gamma \to \Box\Diamond\gamma$*, for arbitrary* γ. Apply (Mon$\Box$) to the dual of T, $\gamma \to \Diamond\gamma$. This is the derivation already used for the (P)-instances.

3) The clauses: Atoms, disjunction, existential quantifier. The two presentations give literally the same formula—$\Box\Diamond P(\vec{t})$, $\Box\Diamond([\Phi] \vee [\Psi])$, $\Box\Diamond\exists x.[\Phi]$—so there is nothing to derive.

Falsum: $[\bot] \leftrightarrow \bot$. Necessitation gives $\Box\top$, so $\Diamond\bot \leftrightarrow \bot$; replacement turns $\Box\Diamond\bot$ into $\Box\bot$, and T gives $\Box\bot \to \bot$. The converse is ex falso.

Conjunction: $[\Phi\wedge\Psi] \leftrightarrow [\Phi]\wedge[\Psi]$*, that is* $\Box\Diamond(\alpha\wedge\beta) \leftrightarrow \alpha\wedge\beta$. Left to right: (Mon$\Diamond$) and then (Mon$\Box$) applied to $\alpha \wedge \beta \to \alpha$ give $\Box\Diamond(\alpha \wedge \beta) \to \Box\Diamond\alpha$, and (F2) gives α; symmetrically β. Right to left: (F1) gives $\alpha \wedge \beta \to \Box\alpha \wedge \Box\beta$, ($\Box\wedge$) gives $\Box(\alpha \wedge \beta)$, and (F3) gives $\Box\Diamond(\alpha \wedge \beta)$.

Negation: $[\neg\Phi] \leftrightarrow \Box\neg[\Phi]$*, that is* $\Box\Diamond\neg\alpha \leftrightarrow \Box\neg\alpha$. Right to left is (F3) with $\gamma = \neg\alpha$. Left to right: contraposing (F1) gives $\neg\Box\alpha \to \neg\alpha$, that is $\Diamond\neg\alpha \to \neg\alpha$, and (Mon$\Box$) gives $\Box\Diamond\neg\alpha \to \Box\neg\alpha$.

Implication: $[\Phi \to \Psi] \leftrightarrow \Box([\Phi] \to [\Psi])$*, that is* $\Box\Diamond(\alpha \to \beta) \leftrightarrow \Box(\alpha \to \beta)$. Right to left is (F3). Left to right, we first derive $\Box\Diamond(\alpha \to \beta) \to (\alpha \to \beta)$. By (F1) and 4, $\alpha \to \Box\Box\alpha$. By ($\Box\wedge$),

$$\Box\Box\alpha \wedge \Box\Diamond(\alpha \to \beta) \to \Box\big(\Box\alpha \wedge \Diamond(\alpha \to \beta)\big).$$

Inside the box, ($\Box\Diamond$) gives $\Diamond(\alpha \wedge (\alpha \to \beta))$ and (Mon$\Diamond$) applied to $\alpha \wedge (\alpha \to \beta) \to \beta$ gives $\Diamond\beta$; so (Mon$\Box$) on that composite yields $\Box\Diamond\beta$, and (F2) yields β. Assembling, $\alpha \wedge \Box\Diamond(\alpha \to \beta) \to \beta$, which is the claim. Now 4 gives $\Box\Diamond(\alpha \to \beta) \to \Box\Box\Diamond(\alpha \to \beta)$, and (Mon$\Box$) applied to the implication just derived carries this to $\Box(\alpha \to \beta)$. Note where the outer box is needed: $\Diamond(\alpha \to \beta) \to (\alpha \to \beta)$ is *not* derivable, β being stable but not reflected downwards.

Universal quantifier: $[\forall x.\Phi] \leftrightarrow \forall x.[\Phi]$*, that is* $\Box\Diamond\forall x.\alpha \leftrightarrow \forall x.\alpha$. Left to right: (Mon$\Diamond$) and (Mon$\Box$) applied to the instantiation $\forall x.\alpha \to \alpha$ give $\Box\Diamond\forall x.\alpha \to \Box\Diamond\alpha$, and (F2) gives α; since x is not free in $\Box\Diamond\forall x.\alpha$, generalization gives $\Box\Diamond\forall x.\alpha \to \forall x.\alpha$. Right to left: generalizing (F1) and distributing gives $\forall x.\alpha \to \forall x.\Box\alpha$, the Barcan formula gives $\Box\forall x.\alpha$, and (F3) gives $\Box\Diamond\forall x.\alpha$. This is the only clause using BF, and it uses only BF: the converse Barcan formula is not needed. With varying domains the step fails, and with it the clause.

Brackets: $[[\Phi]] \leftrightarrow [\Phi]$*, that is* $\Box\Diamond\alpha \leftrightarrow \alpha$. Left to right is (F2); right to left is (F1) followed by (F3). This is the case that makes the uniform presentation well behaved on nested brackets: an argument that is already a twist absorbs the $\Box\Diamond$ it receives as an atom.

4) Assembly: The clauses above are the only content of the equivalence. To pass from them to the statement that the two presentations of a whole formula agree, replace provably equivalent subformulas node by node, which is legitimate in any normal modal logic and, under quantifiers, by generalization; the recursion follows the formula tree and carries no hypothesis of its own.

C. Exactness of BC^+ *over* FOS4

We prove Theorem 5.4. Only soundness of FOS4 for constant-domain S4 Kripke models is used, and no measure-

theoretic structure enters, the models built here being abstract. Throughout, equality is read as a congruence predicate, as in Definition 3.1. Write $m(\Phi)$, the *matrix* of Φ, for the result of erasing every bracket.

Lemma A.1 (Flat normal form): $\vdash_{\mathsf{BC}^+} [\Phi] \leftrightarrow [m(\Phi)]$ for every fragment formula Φ.

Proof: Induction on the number of brackets: one application of (Flat) to the whole formula, with C its outermost classical context, removes the brackets around the maximal boxed subformulas and exposes the next layer, so no congruence under $[\,\cdot\,]$ is ever needed; the replacement of each maximal boxed subformula by its flat equivalent then happens inside the outer *classical* context, ordinary replacement of provable equivalents; $[[\phi]] \leftrightarrow [\phi]$ is the simplest instance. Flattening is therefore a normal form rather than a restriction. □

1) The inner theory: Let $\Gamma \cup \{\neg\Phi\}$ be BC^+-consistent. By Lemma A.1 we may assume every formula in sight is *flat*—a classical combination of bare atoms and boxed modality-free formulas—harmless on both sides, (Flat) being an $\mathsf{FOS4}$ theorem-schema (Theorem 5.2). Adjoin countably many fresh constants. For the persistent version, extend the schema (P) simultaneously to all atomic instances of the Henkin language Σ^+, including instances containing the new constants. Extend to a set T that is maximal BC^+-consistent and Henkin-complete: for each $\exists x.\Psi \in T$ some constant c has $\Psi[c/x] \in T$. Put

$$B = \{P(\vec{t}) \in T : P \text{ atomic}\}, \qquad F = \{\phi : [\phi] \in T\},$$

so B collects the bare atoms of T and F the matrices of its boxed members.

Lemma A.2: F is consistent and closed under first-order consequence; and it contains B whenever the instances of (P) are among the hypotheses.

Proof: Closure. Suppose $F \vdash_{\mathrm{FOL}} \psi$. By compactness there are $\phi_1, \dots, \phi_n \in F$ with $\vdash_{\mathrm{FOL}} (\phi_1 \wedge \cdots \wedge \phi_n) \to \psi$. By (N), $\vdash_{\mathsf{BC}^+} [(\phi_1 \wedge \cdots \wedge \phi_n) \to \psi]$, and by (K), $[\phi_1 \wedge \cdots \wedge \phi_n] \to [\psi]$. The derived conjunction principle turns the antecedent into $[\phi_1] \wedge \cdots \wedge [\phi_n]$, all of whose conjuncts lie in T; maximality gives $[\psi] \in T$, that is $\psi \in F$. *Consistency.* If $F \vdash_{\mathrm{FOL}} \bot$ then $\bot \in F$ by closure, so $[\bot] \in T$; but then (D) gives $\bot \in T$, contradicting consistency. *Inclusion.* If the instances of (P) are among the hypotheses they lie in T; so $P(\vec{t}) \in B$ gives $P(\vec{t}) \in T$, and modus ponens under maximality gives $[P(\vec{t})] \in T$, i.e. $P(\vec{t}) \in F$. □

Lemma A.3 (Inner saturation): F has the $\forall$-property: if $\phi[c/x] \in F$ for every constant c, then $\forall x.\phi \in F$. Consequently every consistent $F \cup \{\delta\}$, δ modality-free, extends to a maximal theory with the $\exists$-witness property *in the same language*, and so has a term model with domain the closed terms.

Proof: If $\forall x.\phi \notin F$ then $[\forall x.\phi] \notin T$; by (BF) and maximality, $\forall x.[\phi] \notin T$ and hence $\exists x.\neg[\phi] \in T$, and Henkin-completeness of T gives a constant c with $\neg[\phi[c/x]] \in T$, so $\phi[c/x] \notin F$. For the consequence—a derivation being finite, both directions descend to the countable subsignature generated by the formulas involved together with countably many names, and ascend by monotonicity, so assuming the signature countable is no loss—enumerate the modality-free *sentences* and run the witnessed completion of $F \cup \{\delta\}$ without adding constants, taking closed terms as witnesses—the names supply infinitely many. Suppose $\exists x.\theta$ is consistent with the current stage but no closed term witnesses it, and let δ' be the conjunction of the finitely many sentences adjoined so far: closed, so $x \notin \mathrm{FV}(\delta')$. For every constant c, $F \cup \{\delta', \theta[c/x]\}$ is inconsistent, so the deduction theorem gives $\delta' \to \neg\theta[c/x] \in F$; the $\forall$-property yields $\forall x.(\delta' \to \neg\theta) \in F$, hence $\delta' \to \forall x.\neg\theta \in F$ as x is not free in δ'—contradicting the consistency of $F \cup \{\delta', \exists x.\theta\}$. No fresh constants are needed, and none are available: the constant domain forces the completion to stay in the language, and (BF) is the axiom that corresponds to it. In the resulting term model, function symbols are interpreted syntactically, predicates by membership in the completed theory, and equality by the theory's congruence, per the convention above. □

2) Fan models:

Definition A.4 (Fan model): Let F be consistent and deductively closed, B a set of atoms, and $\{M_j\}_{j \in J}$ a nonempty family of classical structures, all with domain D the closed terms, each satisfying F. The *fan model* $\mathcal{F}$ has worlds $\{w_0\} \cup \{w_j\}_{j \in J}$, constant domain D, accessibility the reflexive–transitive closure of $w_0 R\, w_j$ $(j \in J)$, and valuation

$$V(w_0) = B, \qquad V(w_j) = \{P(\vec{t}) : M_j \models P(\vec{t})\}.$$

The frame is reflexive and transitive, and each w_j is an endpoint: the only world accessible from w_j is w_j. The domain is constant by construction, so $\mathcal{F}$ is a model of all $\mathsf{FOS4}$ axioms; when moreover $B \subseteq F$, the valuation is persistent—$V(w_0) \subseteq V(w_j)$, every M_j satisfying F—and $\mathcal{F}$ models $\mathsf{FOS4}_P$ as well.

All structures in play share the domain D of closed terms, and every element of D names itself; satisfaction under an assignment therefore coincides with satisfaction of the corresponding closed-substitution instance. We state the next two lemmas for sentences and apply them to closed instances, the inductions running through instances: each $\forall$-case reduces to the instances $\phi[t/x]$, $t \in D$, of lower complexity.

Lemma A.5 (Endpoints are classical): For every modality-free sentence ϕ and every j: $\mathcal{F}, w_j \models \phi^*$ iff $M_j \models \phi$.

Proof: Induction on ϕ, which by Section III-B may be taken built from $\bot, \wedge, \neg, \forall$ alone. Since w_j accesses only itself, $\Box$ and $\Diamond$ are vacuous there, so each clause of the twist reduces to its classical counterpart: $\Box\Diamond P$ becomes P and $\Box\neg\psi^*$ becomes $\neg\psi^*$, while $\bot$, $\wedge$ and $\forall$ are transparent already; the quantifier case uses that the domain of $\mathcal{F}$ is that of M_j, reducing to closed instances as above. □

Lemma A.6 (The root sees the family): For every modality-free sentence ϕ: $\mathcal{F}, w_0 \models \phi^*$ iff $M_j \models \phi$ for every $j \in J$.

Proof: Induction on ϕ, again over the primitives $\bot, \wedge, \neg, \forall$ of Section III-B; write $(\dagger)$ for "$M_j \models \phi$ for all j". The

worlds accessible from w_0 are w_0 itself and the endpoints. Nonemptiness of J is used throughout.

Atoms. $w_0 \models \Box\Diamond P$ requires a P-world above each accessible world. Above w_j only w_j is available, so every $M_j \models P$ is necessary; and it is sufficient, since it also supplies a witness above w_0.

Falsum, conjunction and universal quantification. These clauses of the twist are transparent, and $(\dagger)$ commutes with them; for $\forall$ the constant domain is again the common domain of the M_j, the clause reducing on both sides to the closed instances $\phi[t/x]$, $t \in D$, covered by the induction hypothesis, and for $\bot$ both sides fail, the right one because $J \neq \emptyset$.

Negation. $(\neg\phi)^* = \Box\neg\phi^*$. If $(\dagger)$ holds for $\neg\phi$, then $\neg\phi^*$ holds at each endpoint by Lemma A.5, and at w_0 because by the induction hypothesis ϕ^* would require every $M_j \models \phi$, false since $J \neq \emptyset$. Conversely $\Box\neg\phi^*$ forces $\neg\phi^*$ at every endpoint, hence $M_j \models \neg\phi$ for all j.

The derived connectives $\vee$, $\rightarrow$, $\exists$ (Section III-B) need no separate treatment. $\square$

3) Proof of Theorem 5.4: Proof: The forward direction is Theorem 5.2. For the converse, assume $\Gamma \nvdash_{\mathsf{BC}^+} \Phi$ and construct T, B, F as above; by Lemma A.2 the hypotheses on F are met. Enumerate the modality-free sentences outside F as $\phi_1, \phi_2, \ldots$; for each, $F \cup \{\neg\phi_i\}$ is consistent by deductive closure, so by Lemma A.3 it has a same-language witnessed term model M_i with domain the closed terms. Adding one term model M_0 of F itself gives a nonempty family, and every $M_j \models F$. Let $\mathcal{F}$ be the resulting fan model.

Atomic truth. For modality-free ϕ, Lemma A.6 gives $\mathcal{F}, w_0 \models [\phi]$ iff every $M_j \models \phi$. If $\phi \in F$ this holds; if $\phi \notin F$ then $M_i \models \neg\phi$ for the index i with $\phi = \phi_i$. So $\mathcal{F}, w_0 \models [\phi]$ iff $\phi \in F$ iff $[\phi] \in T$. For a bare atom, $\mathcal{F}, w_0 \models P(\vec{t})$ iff $P(\vec{t}) \in V(w_0) = B$ iff $P(\vec{t}) \in T$.

Truth lemma. By induction on flat sentences Ψ, $\mathcal{F}, w_0 \models \Psi$ iff $\Psi \in T$. The atomic cases—bare atoms and boxed matrices—are the previous paragraph. The Boolean cases use maximal consistency of T. For $\exists x.\Psi$, Henkin-completeness supplies a witnessing constant, which is an element of D; conversely a witness $t \in D$ gives $\Psi[t/x] \in T$ by the induction hypothesis, whence $\exists x.\Psi \in T$. The case of $\forall$ is dual, using that every element of D is a closed term and hence named.

Since $\Gamma \subseteq T$ and $\neg\Phi \in T$, the model $\mathcal{F}$ satisfies Γ and refutes Φ at w_0. As $\mathcal{F}$ is a constant-domain S4 model, soundness of $\mathsf{FOS4}$ yields $\Gamma \nvdash_{\mathsf{FOS4}} \Phi$. For the local-signature form, adjoin the (P)-instances to Γ: Lemma A.2 then gives $B \subseteq F$, the fan is persistent, and the same argument yields underivability in $\mathsf{FOS4}_P$, the adjoined instances being $\mathsf{FOS4}_P$-theorems: $\Gamma \cup \mathsf{P} \vdash_{\mathsf{BC}^+} \Phi$ iff $\Gamma \vdash_{\mathsf{FOS4}_P} \Phi$, with P the instances $P(\vec{t}) \rightarrow [P(\vec{t})]$. $\square$

Remark 5: The construction uses only depth-one models: and all the modal content sits in the contrast between the root and its endpoints. This is why no measure-theoretic input is needed. Turning $\mathcal{F}$ into a *computational* model—replacing the endpoints by non-negligible sets and the root by the whole space—is exactly the content of Theorem 6.2, and it is there that a negligible but nonempty set of exceptional traces must be inserted to separate a bare atom P from $[P]$.

D. Computational Completeness (Theorem 6.2)

Together with the exactness construction of Appendix C, the following realization proves Theorem 6.2 over its stated class.

Definition A.7 (Term fans): A *term fan* over Σ^+ is a nonempty family $\{M_i\}_{i\in I}$ of classical structures with common universe the closed Σ^+-terms and with every constant and function symbol interpreted syntactically in the same way in every component. An *enriched term fan* additionally contains a structure N with the same universe and the same constant and function interpretations, subject to

$$\mathrm{Th}_{\mathrm{at}}(N) \subseteq \bigcap_{i\in I} \mathrm{Th}_{\mathrm{at}}(M_i).$$

Theorem A.8 (Fan realization): Every countable term fan is realized at the boxed layer by a computational model in the relaxed class of Remark 7: there is such a computational model under which, at the full sample space, exactly the same boxed sentences $[\phi]$ hold—the components realized on a banded partition with the transfer holding against *all* non-negligible sets. If moreover the fan is enriched—$(N; \{M_i\}_i)$ with $\mathrm{Th}_{\mathrm{at}}(N) \subseteq \bigcap_i \mathrm{Th}_{\mathrm{at}}(M_i)$—then a null modification realizes the bare layer as well, and all flat fragment sentences evaluate as at the root (Theorem A.14).

1) Proof of Theorem A.8: Fix a term fan $\mathfrak{F} = \{\mathcal{M}_i\}_{i\in I}$ over the closed terms of Σ^+. Since duplicating a component changes no satisfaction value (the clause for $[\phi]$ quantifies over all components), we may assume $I = \mathbb{N}$, repeating a component if necessary. Throughout, the counter-model is built over the standard atomless sample space: $\Omega_\eta^{\mathcal{C}} = 2^{\mathbb{N}} \times 2^{\mathbb{N}}$ with the product Borel σ-algebra and the fair-coin product measure, PPT algorithms reading only polynomially many bits of their tapes. The class of Definition 3.1 may contain other sample spaces; the construction chooses this member.

a) Step 1: the banded partition.:

Lemma A.9 (Banded partition): There is a family $\{R_i\}_{i\in\mathbb{N}}$ of ensembles such that:

(P1) for every η, $\{R_{i,\eta}\}_{i\in\mathbb{N}}$ is a measurable partition of $\Omega_\eta^{\mathcal{C}}$;
(P2) $\mu_\eta(R_{i,\eta}) = \eta^{-i}$ for $i \geq 1$ and $\eta \geq 3$, and $\mu_\eta(R_{0,\eta}) \geq \frac{1}{2}$ for $\eta \geq 3$; in particular every R_i is non-negligible;
(P3) for every $c \geq 0$ and $\eta \geq 3$: $\sum_{i>c} \mu_\eta(R_{i,\eta}) \leq 2\eta^{-(c+1)}$.

Proof: For $\eta \leq 2$ put $R_{0,\eta} := \Omega_\eta^{\mathcal{C}}$ and $R_{i,\eta} := \emptyset$ for $i \geq 1$. For $\eta \geq 3$, note $\sum_{i\geq 1} \eta^{-i} = \frac{1}{\eta-1}$, and set $m_0(\eta) := 1 - \frac{1}{\eta-1} \geq \frac{1}{2}$ and $m_i(\eta) := \eta^{-i}$ for $i \geq 1$, so that $\sum_{i\geq 0} m_i(\eta) = 1$. Reading the honest tape as the binary expansion of a real $x(\rho_h) \in [0,1]$ pushes the fair-coin measure forward to Lebesgue measure, so with $t_0 := 0$ and $t_{i+1} := t_i + m_i(\eta)$ the sets

$$R_{0,\eta} := \{\rho : x(\rho_h) \in [t_0, t_1) \cup \{1\}\}$$

$$R_{i,\eta} := \{\rho : x(\rho_h) \in [t_i, t_{i+1})\} \quad (i \geq 1)$$

are Borel measurable (x is continuous; the null fiber $x^{-1}(\{1\})$ is assigned to R_0 so that the sets partition $\Omega^{\mathcal{C}}_\eta$ exactly), and have the prescribed measures. This gives (P1) and (P2); each R_i with $i \geq 1$ is non-negligible since $\eta^{-i} > 1/(2\eta^i)$ for every $\eta \geq 3$. For (P3), $\sum_{i>c} \eta^{-i} = \eta^{-(c+1)} \cdot \frac{\eta}{\eta-1} \leq 2\eta^{-(c+1)}$ for $\eta \geq 2$, and for $c = 0$ the sum omits R_0. □

Lemma A.10 (Banding): For a non-negligible ensemble S put

$$I(S) := \{i \in \mathbb{N} : S \cap R_i \text{ is non-negligible}\}.$$

Then:

(B1) $I(S) \neq \emptyset$;
(B2) if $S' \subseteq S$ then $I(S') \subseteq I(S)$;
(B3) if $i \in I(S)$ then $I(S \cap R_i) = \{i\}$;
(B4) $I(\Omega^{\mathcal{C}}) = \mathbb{N}$.

Proof: For (B1) we prove the contrapositive of the stronger covering property: if $S \cap R_i$ is negligible for *every* i, then S is negligible. (Worlds are measurable, non-negligibility referring to their measure, so all sets below are measurable.) Fix $c \in \mathbb{N}$. For every $\eta \geq 3$, countable additivity over the partition (P1) gives

$$\mu_\eta(S_\eta) \leq \sum_{i=0}^{c} \mu_\eta(S_\eta \cap R_{i,\eta}) + \sum_{i>c} \mu_\eta(R_{i,\eta}).$$

The first sum is a sum of $c+1$ negligible functions, hence negligible, so it is at most $\eta^{-(c+1)}$ for all large η; the second is at most $2\eta^{-(c+1)}$ by (P3). Hence $\mu_\eta(S_\eta) \leq 3\eta^{-(c+1)} \leq \eta^{-c}$ for all large η. As c was arbitrary, for every polynomial q we may choose c exceeding its degree and conclude $\mu_\eta(S_\eta) < 1/q(\eta)$ for all large η: S is negligible. (B2) is immediate from $S' \cap R_i \subseteq S \cap R_i$. For (B3): $i \in I(S \cap R_i)$ since $(S \cap R_i) \cap R_i = S \cap R_i$ is non-negligible, and for $j \neq i$ the ensemble $(S \cap R_i) \cap R_j$ is empty at every η by (P1), hence negligible. (B4) holds because every R_i is non-negligible. □

Remark 6 (No fixed-measure splitting): The growth of the polynomial degrees in (P2) is not an artifact: with band measures independent of η, such as the 2^{-i} of a prefix-free splitting, the diagonal ensemble $S_\eta := \bigcup_{i \geq \log_2 \eta} R_{i,\eta}$ has measure about $2/\eta$—non-negligible—while meeting every fixed band in an eventually empty, hence negligible, ensemble; so (B1) fails and no transfer lemma is possible over such a splitting.

b) Step 2: the realized model.: Fix an effective presentation of the countable signature and an injective encoding enc of the closed Σ^+-terms as bit strings, computable in polynomial time together with a *total* polynomial-time decoder $\mathrm{dec} : \{0,1\}^* \to \{\text{closed terms}\} \cup \{\bot\}$ returning $\bot$ on non-codes, the range of enc disjoint from $\{0\}$. Define the computational model $\mathcal{M}^{\mathcal{C}}$ as follows.

Domain. For each closed term t let $\mathfrak{a}_t$ be the constant algorithm that on input 1^η and any tape outputs $\mathrm{enc}(t)$; a fixed output makes it trivially PPT. Put $D := \{\mathfrak{a}_t : t \text{ closed}\}$; the map $t \mapsto \mathfrak{a}_t$ is a bijection from the universe of $\mathfrak{F}$ to D, and we identify valuations into D with valuations $\bar{\nu}$ into the closed terms through it; the statement of Theorem A.8 is read under this identification.

Function symbols. Every function symbol f of Σ^+—honest, adversarial, name or added constant alike—is interpreted as the constructor

$$f^{\mathcal{M}^{\mathcal{C}}}_\eta(s_1, \ldots, s_n) := \begin{cases} \mathrm{enc}(f(t_1, \ldots, t_n)) & \text{if } s_j = \mathrm{enc}(t_j) \text{ for all } j, \\ 0 & \text{otherwise,} \end{cases}$$

a polynomial-time function of its arguments, independent of η and of the tapes. (Adversarial symbols *may* read the adversarial tape; they are not obliged to. Names are discussed in Remark 7 below.)

Predicates. For an n-ary predicate P, define the trace property

$$P^{\mathcal{M}^{\mathcal{C}}}_{\eta,\rho}(s_1, \ldots, s_n) := 1 \iff \rho \in R_{i,\eta} \text{ and } \mathcal{M}_i \models P(t_1, \ldots, t_n),$$

where i is the unique band with $\rho \in R_{i,\eta}$ and $s_j = \mathrm{enc}(t_j)$ for closed t_j (value 0 if some s_j is not an encoding). Its value depends only on ρ and the argument strings, so it is a trace property; for fixed closed arguments its trace set is the countable union of measurable bands

$$A(P(\vec{t})) := \bigcup \{R_i : \mathcal{M}_i \models P(\vec{t})\},$$

measurable but not necessarily PPT-decidable, which the definition of computational models permits.

Thus $\mathcal{M}^{\mathcal{C}}$ is a computational model in the relaxed class of Remark 7—names as constants, congruence equality, measurable predicates—and its Kripke model $\mathcal{M}^{\mathcal{KC}}$ has as worlds *all* non-negligible ensembles—the full frame, not a copy of some abstract frame inside it.

Lemma A.11 (Denotation): For every Σ^+-term t, valuation ν into D, and all η, ρ: $[\![t]\!]^\nu_{\eta,\rho} = \mathrm{enc}(t\bar{\nu})$, where $t\bar{\nu}$ is the closed term obtained by substituting $\bar{\nu}(x)$ for each variable x of t. In particular closed terms denote their own encodings, independently of the trace.

Proof: Induction on t: variables by the definition of $\mathfrak{a}_{\bar{\nu}(x)}$, composite terms because the constructor interpretation acts on encodings exactly as f acts on terms. □

c) Step 3: the transfer lemma.: This is the heart of the realization: satisfaction at an *arbitrary* world is controlled by the bands the world meets non-negligibly. The $\Box\Diamond$ pattern of the clauses in Definition 3.2 is what makes the two directions meet.

Lemma A.12 (Transfer): For every modality-free formula ϕ, every valuation ν into D, and every non-negligible $S \subseteq \Omega^{\mathcal{C}}$:

$$\mathcal{M}^{\mathcal{C}}, S, \nu \models_{BC} \phi \iff \mathcal{M}_i, \bar{\nu} \models \phi \text{ for every } i \in I(S).$$

Proof: Induction on the structure of ϕ, all valuations at once. We use (B1)–(B3) of Lemma A.10 throughout, and for the quantifier cases the classical substitution lemma in the $\mathcal{M}_i$, whose universe is the closed terms, each denoting itself.

Atoms, $\phi = P(\vec{t})$. By Lemma A.11, $[\![P(\vec{t})]\!]^\nu_{\eta,\rho} = 1$ iff $\rho \in A(P(\vec{t\bar{\nu}}))_\eta$. ($\Leftarrow$) Assume $\mathcal{M}_i \models P(\vec{t\bar{\nu}})$ for every $i \in I(S)$, and let $S' \subseteq S$ be non-negligible. Pick $i \in I(S')$ by (B1); then

$i \in I(S)$ by (B2), so $R_i \subseteq A(P(\vec{t}\bar{\nu}))$. The ensemble $S'' := S' \cap R_i$ is non-negligible and satisfies the pointwise condition at *every* η and $\rho \in S''_\eta$; so the atomic clause holds at S. ($\Rightarrow$) Let $i \in I(S)$ and apply the atomic clause to $S' := S \cap R_i$: it yields a non-negligible $S'' \subseteq S'$ on which $P(\vec{t})$ holds pointwise everywhere. A non-negligible ensemble is nonempty at some η; pick $\rho \in S''_\eta$. Then $\rho \in R_{i,\eta} \cap A(P(\vec{t}\bar{\nu}))_\eta$, and since the bands partition $\Omega^{\mathcal{C}}_\eta$, the only band contributing to $A(P(\vec{t}\bar{\nu}))_\eta$ that contains ρ is $R_{i,\eta}$; hence $\mathcal{M}_i \models P(\vec{t}\bar{\nu})$.

$\bot$: both sides fail—the left by the clause for $\bot$, the right because $I(S) \neq \emptyset$ by (B1).

$\wedge$: both sides distribute over conjunction.

$\neg$, $\phi = \neg\psi$. ($\Leftarrow$) Assume $\mathcal{M}_i \not\models \psi$ for every $i \in I(S)$, and let $S' \subseteq S$ be non-negligible. If $S' \models_{BC} \psi$, the induction hypothesis at S' gives $\mathcal{M}_j \models \psi$ for every $j \in I(S')$, a nonempty subset of $I(S)$ by (B1)–(B2)—contradiction. So no non-negligible subset of S satisfies ψ, which is the clause for $\neg\psi$. ($\Rightarrow$) Let $i \in I(S)$. Then $S \cap R_i$ is a non-negligible subset of S, so $S \cap R_i \not\models_{BC} \psi$ by the clause; by the induction hypothesis at $S \cap R_i$, whose band set is $\{i\}$ by (B3), this means exactly $\mathcal{M}_i \not\models \psi$.

$\vee$, $\phi = \psi \vee \chi$. ($\Leftarrow$) Given non-negligible $S' \subseteq S$, pick $i \in I(S') \subseteq I(S)$; say $\mathcal{M}_i \models \psi$. With $S'' := S' \cap R_i$ we have $I(S'') = \{i\}$ by (B3), so the induction hypothesis gives $S'' \models_{BC} \psi$; the clause for $\vee$ holds at S. ($\Rightarrow$) Let $i \in I(S)$ and apply the clause to $S' := S \cap R_i$: some non-negligible $S'' \subseteq S'$ satisfies ψ or χ. By (B1)–(B3), $I(S'') = \{i\}$, so the induction hypothesis gives $\mathcal{M}_i \models \psi$ or $\mathcal{M}_i \models \chi$.

$\rightarrow$, $\phi = \psi \rightarrow \chi$. ($\Leftarrow$) Assume $\mathcal{M}_i \models \psi \rightarrow \chi$ for every $i \in I(S)$, and let $S' \subseteq S$ be non-negligible with $S' \models_{BC} \psi$. By the induction hypothesis ($\Rightarrow$-direction) at S', $\mathcal{M}_j \models \psi$ for every $j \in I(S')$; as $I(S') \subseteq I(S)$, also $\mathcal{M}_j \models \chi$ for every $j \in I(S')$; by the induction hypothesis ($\Leftarrow$-direction) at S', $S' \models_{BC} \chi$. This is the clause for $\rightarrow$. ($\Rightarrow$) Let $i \in I(S)$ and suppose $\mathcal{M}_i \models \psi$. By the induction hypothesis at $S \cap R_i$ (band set $\{i\}$), $S \cap R_i \models_{BC} \psi$; the clause applied to $S' := S \cap R_i$ gives $S \cap R_i \models_{BC} \chi$; the induction hypothesis again gives $\mathcal{M}_i \models \chi$.

$\forall$, $\phi = \forall x.\psi$. Using that every element of D is $\mathfrak{a}_t$ for a unique closed t:

$$\begin{aligned} S, \nu \models_{BC} \forall x.\psi &\iff \forall t:\ S, \nu[x \mapsto \mathfrak{a}_t] \models_{BC} \psi \\ &\iff \forall t\ \forall i \in I(S):\ \mathcal{M}_i, \bar{\nu}[x \mapsto t] \models \psi \\ &\iff \forall i \in I(S):\ \mathcal{M}_i, \bar{\nu} \models \forall x.\psi, \end{aligned}$$

the middle step by the induction hypothesis and the last because the universe of each $\mathcal{M}_i$ is the closed terms.

$\exists$, $\phi = \exists x.\psi$. ($\Leftarrow$) Given non-negligible $S' \subseteq S$, pick $i \in I(S') \subseteq I(S)$. Since $\mathcal{M}_i, \bar{\nu} \models \exists x.\psi$ and the universe is the closed terms, there is a closed t with $\mathcal{M}_i, \bar{\nu}[x \mapsto t] \models \psi$. With $S'' := S' \cap R_i$, $I(S'') = \{i\}$, the induction hypothesis gives $S'', \nu[x \mapsto \mathfrak{a}_t] \models_{BC} \psi$; the clause for $\exists$ holds at S, the witness $\mathfrak{a}_t$ varying with the band met—precisely what the clause permits. ($\Rightarrow$) Let $i \in I(S)$ and apply the clause to $S' := S \cap R_i$: there are a non-negligible $S'' \subseteq S'$ and $\mathfrak{a}_t \in D$ with $S'', \nu[x \mapsto \mathfrak{a}_t] \models_{BC} \psi$. As $I(S'') = \{i\}$, the induction hypothesis gives $\mathcal{M}_i, \bar{\nu}[x \mapsto t] \models \psi$, so $\mathcal{M}_i, \bar{\nu} \models \exists x.\psi$. $\square$

d) Step 4: conclusion.: Proof of Theorem A.8. Proof: The model $\mathcal{M}^{\mathcal{C}}$ of Step 2, with the bands of Step 1, has all the listed properties, and the displayed equivalence is Lemma A.12. For the last claim: since $I(\Omega^{\mathcal{C}}) = \mathbb{N}$ by (B4), the case $S = \Omega^{\mathcal{C}}$ of Lemma A.12 reads

$$\mathcal{M}^{\mathcal{C}}, \nu \models [\phi] \iff \mathcal{M}_i, \bar{\nu} \models \phi \text{ for every } i \iff \mathfrak{F}, \bar{\nu} \models [\phi],$$

so the two models agree on the boxed atoms of the fragment. Fragment formulas all of whose atoms are boxed then agree by induction on the outer structure: the outer connectives are Tarskian on both sides, and the outer quantifiers range over D and over the closed terms respectively, matched by the bijection $t \mapsto \mathfrak{a}_t$. This completes the component layer of Theorem A.8; the bare layer is proved in Section D2. $\square$

Remark 7 (The model class): Three conventions of the construction deserve to be made explicit. (i) *Names and adversarial symbols* are realized as constructors: legitimate outright for adversarial symbols, which may but need not read the adversarial tape; for names, however, Section III-B fixes the interpretation as independent uniform sampling, which the realization does not respect. Stated precisely, Theorem 6.2 is completeness with respect to computational models in which names are interpreted like any other constants; their probabilistic reading, like all cryptographic content, enters through axioms, in accordance with the remark that completeness does not mean real attacks. We do not claim completeness relative to the fixed name interpretation. (ii) *Equality* is realized as an ordinary band-wise trace property—a congruence on each band whenever the equality axioms are among the hypotheses—rather than as literal bit-string equality, which the definition permits, equality being the primary *example* of a trace property rather than a hardwired one. Literal equality would force quotienting the domain by provable equality, and canonical representatives of those classes are in general not computable; completeness for the class with literal equality we leave open. (iii) The predicates constructed are measurable but need not be PPT-decidable, which the definition of computational models explicitly allows; only the domain and the function symbols are PPT. Finally, with constants and witnesses added per sort, everything applies verbatim to many-sorted signatures, including Squirrel's function-space sorts.

2) Proof of Theorem A.14: Let $(D, \mathcal{I}; N, \{M_i\})$ be a countable enriched term fan satisfying the atomic inclusion of Definition A.7. (In the paper's notation, N is the root of the depth-one Kripke fan: $\mathrm{Th}_{\mathrm{at}}(N) = B$, the bare atoms of the Henkin theory.) Run the banded construction above for $\{M_i\}$ unchanged, obtaining a computational model whose atomic trace sets are band-wise unions $A(P(\vec{t})) = \bigcup\{R_i : M_i \models P(\vec{t})\}$.

a) The null copy of the root.: Let $e_\eta := (0^\omega, 0^\omega) \in R_{0,\eta}$ (indeed $x(0^\omega) = 0 \in [t_0, t_1)$) and $E := \{e_\eta\}_\eta$, a single trace per parameter, of measure zero. Redefine the *entire* atomic

structure at e_η to be that of N: for every predicate P and closed terms $\vec{t}$, the modified model satisfies $P(\vec{t})$ at e_η iff $N \models P(\vec{t})$, and at every other trace as before; domain and function interpretations are untouched. Because the whole structure at the exceptional trace is N, equality remains a congruence there; and every atomic truth ensemble changes only on the common null ensemble E, so by Lemma A.13 below no twisted formula changes value at any world.

b) The strict layer is realized.: An atomic instance holds pointwise at every trace of the full space iff N satisfies it: if $N \models A$, every component satisfies A by the inclusion $\mathrm{Th}_{\mathrm{at}}(N) \subseteq \bigcap_i \mathrm{Th}_{\mathrm{at}}(M_i)$, so A holds on every band and at e_η by construction; if $N \not\models A$, then A fails at e_η. Compound flat formulas, mixing bare and boxed atoms, follow by Tarski induction over the closed-term domain, exactly as at the root of the fan: the bare atoms agree with N by the above, the boxed atoms with the family by the transfer lemma, and the outer classical structure is evaluated identically on both sides.

Lemma A.13 (Exception stability): If two computational models have the same domain and function-symbol interpretations, and for every atomic instance and valuation the symmetric difference of the two truth ensembles is negligible, then for every formula, valuation, and non-negligible world the two models agree on twisted satisfaction.

Proof: Induction on the matrix, mirroring the transfer lemma, carried for all valuations simultaneously; the quantifier clauses apply the induction hypothesis per domain element. Atoms: a pointwise witness set for one model, less the negligible difference, is one for the other, and remains non-negligible; the $\Box\Diamond$ clause is therefore insensitive in both directions. The pointwise clauses ($\wedge$, $\forall$) and the boxed clauses ($\neg$, $\rightarrow$, $\vee$, $\exists$) pass through the induction hypothesis world by world. □

By the lemma, the null-root modification leaves every twisted formula, at every world, exactly as the banded construction made it; in particular the transfer lemma holds verbatim. Together with the strict-layer verification this proves Theorem A.8 for enriched fans; the diagonal remark on failed alternatives is unaffected. Explicitly:

Theorem A.14 (Realization of enriched fans): Every countable enriched term fan $(N; \{M_i\}_i)$ with $\mathrm{Th}_{\mathrm{at}}(N) \subseteq \bigcap_i \mathrm{Th}_{\mathrm{at}}(M_i)$ is realized by a computational model in the relaxed class of Remark 7: at the full space, an atomic instance holds pointwise iff N satisfies it, $[\phi]$ holds iff every $M_i \models \phi$, and flat sentences evaluate exactly as at the root of the corresponding depth-one Kripke fan.

Applied to the fan built in the paper's exactness appendix from a maximal BC^+-consistent Henkin theory T—with $\mathrm{Th}_{\mathrm{at}}(N) := B$, the bare atoms of T, and the inclusion supplied by the adjoined (P)-instances—the theorem turns the abstract counter-model into a computational one: if $\Gamma \cup \mathsf{P} \nvdash_{\mathsf{BC}^+} \Phi$, with P the definitional (P)-instances supplying the inclusion, some computational Kripke model satisfies Γ—and, being pointwise, all of P—and refutes Φ at the full space. Taking the reduct from the Henkin expansion Σ^+ to the original signature, and translating along Lemma A.1, transfers the counter-model from the flat representatives to the original Γ andf Φ. This proves Theorem 6.2.

Remark 8 (The realized model respects the signature): Every predicate of the realized model is interpreted as a trace property: band membership and the null modification only shape trace sets. Trace properties are persistent, so the counter-model lies in the class of computational models interpreting the local predicates persistently—the class over which the definitional instances of (P) are sound. One construction therefore serves both readings of the theorem, the pointwise form being the stronger. In particular, no hypothesis can deny persistence: the (P)-instances are adjoined to Γ, so a theory containing $P(\vec{c}) \wedge \neg[P(\vec{c})]$ is inconsistent and no fan is built for it—matching the semantic side, where such a formula is unsatisfiable over the models of Definition 3.1.

E. The Sequent Calculus and Cut Elimination (Theorem 6.3)

1) The Extended Language, Matrices, and the Calculus: Throughout, the *extended fragment* $\mathcal{L}^+$ is generated from atoms $P(\vec{t})$ and $\bot$ by the classical connectives, the quantifiers, and the operator $[\,\cdot\,]$, applied to arbitrary $\mathcal{L}^+$-formulas. Formulas are read as strings of the ambient modal language through the extended twist: the twist clauses of Section III-B are applied with boxed formulas treated as atoms, so that on modality-free arguments $[\phi] = \phi^*$, while for instance $[[\phi]] = \Box\Diamond[\phi]$. Under this reading the transparent clauses remain *syntactic identities* for arbitrary arguments:

$$[\bot] = \bot, \qquad [\Phi \wedge \Psi] = [\Phi] \wedge [\Psi], \qquad [\forall x.\Phi] = \forall x.[\Phi],$$

so the calculus builds in the axioms (D) and (BF) and the derived conjunction principle: they cost nothing at the sequent level, and the Barcan content resurfaces only in the cut analysis below. Bare atoms are new: $P(\vec{t})$ is a formula on its own, its computational reading being truth at *every* trace of the world under evaluation, and its fan reading being truth in the strict layer N of an enriched fan $(N; \{M_i\}_i)$, where N is one further classical structure over the common domain constrained by atomic inclusion, $\mathrm{Th}_{\mathrm{at}}(N) \subseteq \bigcap_i \mathrm{Th}_{\mathrm{at}}(M_i)$.

Definition A.15 (Matrix): The *matrix* $m(\Phi)$ of an $\mathcal{L}^+$-formula is the modality-free classical formula obtained by erasing every occurrence of $[\,\cdot\,]$: $m(P(\vec{t})) = P(\vec{t})$, $m([\Phi]) = m(\Phi)$, and m commutes with $\bot$, the connectives and the quantifiers. On the flat fragment, $m([\phi]) = \phi$.

Sequents $\Gamma \vdash \Delta$ are pairs of finite multisets of $\mathcal{L}^+$-formulas. The calculus G, the sequent form of BC^+, has:

- initial sequents $A \vdash A$ for atomic A (bare or boxed), and $\bot, \Gamma \vdash \Delta$; general $\Phi \vdash \Phi$ is derivable;
- weakening and contraction on both sides;
- the rules of classical LK for the outer connectives and quantifiers, with the usual eigenvariable conditions;
- the box rules, for $k, r \geq 0$, bare atoms $P_1, \ldots, P_r$, and $\mathcal{L}^+$-formulas $\Phi_1, \ldots, \Phi_k, \Psi$:

$$\frac{P_1, \ldots, P_r,\ m(\Phi_1), \ldots, m(\Phi_k) \vdash_{\mathrm{LK}} m(\Psi)}{P_1, \ldots, P_r,\ [\Phi_1], \ldots, [\Phi_k] \vdash [\Psi]}\ (\Box)$$

$$\frac{P_1, \ldots, P_r,\ m(\Phi_1), \ldots, m(\Phi_k) \vdash_{\mathrm{LK}} \emptyset}{P_1, \ldots, P_r,\ [\Phi_1], \ldots, [\Phi_k] \vdash} \ (\Box\emptyset)$$

whose premise is derivability in classical LK for the modality-free language (kept cut-free inside by Gentzen's Hauptsatz whenever inspected);

- the cut rule, confined and eliminated below.

Bare atoms enter the inner sequent unchanged; boxed antecedents and the boxed succedent contribute their matrices. The box rules have no side formulas.

Lemma A.16 (Soundness of the box rules): The box rules are sound over enriched fans and over computational models.

Proof: Enriched fans: by the identities and the extended twist, $[\Phi]$ holds in a fan iff every component satisfies $m(\Phi)$ (induction on Φ; the inner boxes are absorbed exactly as in the Flattening schema, whose fan validity is immediate since components see only matrices). If the conclusion's antecedents hold, then $N \models P_j$ for each j, whence by atomic inclusion every component satisfies each P_j; every component also satisfies each $m(\Phi_i)$; classical soundness of the premise gives $m(\Psi)$ in every component, i.e. $[\Psi]$. For computational models, translate the cut-free LK premise into the corresponding first-order implication. Rule (N), repeated use of (K), and the conjunction principle yield the bracketed implication; pointwise truth of each bare passenger P_j implies $[P_j]$ because trace predicates are persistent, and Flattening replaces $[m(\Phi_i)]$ and $[m(\Psi)]$ by $[\Phi_i]$ and $[\Psi]$. Corollary 5.3 therefore validates $(\Box)$ in every computational model. The same argument with an empty inner succedent and $[\bot] = \bot$ validates $(\Box\emptyset)$. □

Remark 9 (Persistence is an instance): The persistence schema $P(\vec{x}) \to [P(\vec{x})]$—the definitional theory of local signatures, not an axiom of BC^+—is nevertheless derivable in G: it is the box inference with inner premise $P \vdash_{\mathrm{LK}} P$, followed by $\to$-R. Its converse is refuted below (Corollary A.23): the strict layer may be strictly smaller. No axiom is postulated, so no new cut pathology can enter through persistence—contrast the Barcan formula in full $\mathsf{FOS4}$, which must be postulated and then cut on.

Proposition A.17 (Flattening is invisible): For every classical context C over boxed atoms and all $\Phi_1, \ldots, \Phi_k$, both sequents

$$[C([\Phi_1], \ldots, [\Phi_k])] \vdash [C(\Phi_1, \ldots, \Phi_k)] \qquad \text{and its converse}$$

are single box inferences. Hence the Flattening schema of BC^+ requires no rule of G: the two sides have the same matrix $C(m(\Phi_1), \ldots, m(\Phi_k))$, and the inner premise is an initial sequent.

Proof: Immediate from Definition A.15: m erases the inner brackets first and the outer bracket last, so both sides have the matrix stated, and $(\Box)$ applies with the inner premise $C(m(\vec{\Phi})) \vdash_{\mathrm{LK}} C(m(\vec{\Phi}))$. □

This is the proof-theoretic form of a slogan the Hilbert side makes precise: the axioms T and 4, concentrated in the single Flattening schema of BC^+, are *free* at the sequent level. Nesting never costs anything, because the twist makes it matrix-preserving; in full S4, where nesting is not matrix-preserving, cut elimination must recurse through modal depth.

Proposition A.18 (Equivalence with BC^+): Let BC be classical first-order logic over $\mathcal{L}^+$ with (N), (K), (D) $[\bot] \to \bot$ and (BF) $\forall x.[\Phi] \to [\forall x.\Phi]$, and let BC^+ extend it by the Flattening schema (Section V); write (P) for the set of persistence instances $P(\vec{x}) \to [P(\vec{x})]$, the definitional theory of local signatures. Then for sentences, $\Gamma_0 \vdash \Phi$ is derivable in G with cut for some finite $\Gamma_0 \subseteq \Gamma$ iff Φ is derivable in BC^+ from Γ and (P): the bare passengers of the box rule build the persistence instances in.

Proof: From G to BC^+: read $\Gamma \vdash \Delta$ as $\bigwedge\Gamma \to \bigvee\Delta$; LK rules and cut are classically admissible; $(\Box)$ is derivable in BC^+: from the inner premise, (N) gives $[\bigwedge \vec{P} \wedge \bigwedge m(\vec{\Phi}) \to m(\Psi)]$, (K) and the $\wedge$-identity give $[\bigwedge \vec{P}] \wedge [\bigwedge m(\vec{\Phi})] \to [m(\Psi)]$, the adjoined (P)-instances supply $[\vec{P}]$ from $\vec{P}$, and Flattening converts $[m(\Phi_i)]$ to $[\Phi_i]$ and $[m(\Psi)]$ to $[\Psi]$; $(\Box\emptyset)$ likewise via $[\bot] = \bot$. Conversely: modus ponens is a cut, generalization is $\forall$-R, (N) is $(\Box)$ with $k = r = 0$, (K) is the $(\Box)$ instance $[\Phi \to \Psi], [\Phi] \vdash [\Psi]$ (inner premise: modus ponens on the matrices), persistence is Remark 9, and Flattening is Proposition A.17. □

2) Cut Confinement and the Barcan Obstruction: A *Barcan cut* is a cut whose cut formula is of the form $\forall x.[\Theta]$ $(= [\forall x.\Theta])$, whose left premise ends in $\forall$-R introducing it, and whose right premise is a box inference with $[\forall x.\Theta]$ among its antecedents. The *outer degree* of a formula counts the connectives and quantifiers of its outermost classical layer, boxed atoms (of any internal complexity) and bare atoms counting 0; the secondary measure on boxed atoms is the size of the matrix.

Theorem A.19 (Cut confinement): Every G-derivation can be transformed into one of the same end-sequent in which every cut is a Barcan cut.

Proof: Gentzen's procedure with the mix rule, by induction on $\langle$outer degree, matrix size$\rangle$ of the cut formula, with a subsidiary induction on heights. Parametric cases are the usual permutations; box inferences have no side formulas, so no cut formula is parametric in one. The principal–principal cases:

More explicitly, use the mix form

$$\frac{\Gamma \vdash \Delta, A^m \qquad A^n, \Pi \vdash \Lambda}{\Gamma, \Pi \vdash \Delta, \Lambda}\ (\mathrm{mix}),$$

which removes all displayed occurrences of A at once. Its primary rank is $\langle d_o(A), |m(A)| \rangle$, ordered lexicographically, and its secondary rank is the sum of the premise heights. This formulation absorbs the contraction cases: a terminal contraction on an occurrence of A is removed by regarding both contracted copies as principal occurrences of the same mix; weakenings introducing A are erased, and the other structural rules permute in the standard way. Before permuting across a quantified rule, rename its eigenvariable fresh for the opposite derivation and the end-sequent. If A is parametric in

the last outer LK inference of either premise, permute mix above that inference; the primary rank is unchanged and the height sum strictly decreases. Since a box inference has no side formulas, every formula occurrence in its conclusion is principal. Thus it remains only to verify the following principal cases.

Bare-atom cuts. A bare atom is principal on the right of no rule except an initial sequent (the box succedent is boxed), so the cut formula traces on the left to initial sequents and weakenings; contract or erase, as in the standard elimination of atomic cuts.

Both premises outer-logical. Gentzen's reductions verbatim, producing cuts of smaller outer degree.

Both premises box inferences (cut formula $[\Psi]$, any Ψ). The inner sequents cut classically on the common matrix $m(\Psi)$; one box inference concludes, the bare passengers of both premises riding along. The outer cut disappears entirely and no induction on modal structure is needed—nested cut formulas included, since $[[\Psi]]$ and $[\Psi]$ have the same matrix. Writing the two final inferences schematically as

$$\frac{\vec{P}, m(\vec{\Phi}) \vdash_{\mathrm{LK}} m(\Psi)}{\vec{P}, [\vec{\Phi}] \vdash [\Psi]} \qquad \frac{\vec{Q}, m(\Psi), m(\vec{\Theta}) \vdash_{\mathrm{LK}} m(\Xi)}{\vec{Q}, [\Psi], [\vec{\Theta}] \vdash [\Xi]},$$

classical cut followed by the Hauptsatz gives a cut-free inner derivation

$$\vec{P}, \vec{Q}, m(\vec{\Phi}), m(\vec{\Theta}) \vdash_{\mathrm{LK}} m(\Xi).$$

One application of $(\Box)$ gives the mixed conclusion. If the right box has empty succedent, the same construction ends with $(\Box\emptyset)$. Repeated passengers are contracted inside LK or after the box inference.

Left box, right outer. The boxed cut formula is consumed on the right by: $\forall$-L on $[\forall x.\Theta] = \forall x.[\Theta]$—instantiate the left premise's inner derivation at t, box, and cut on $[\Theta(t/x)]$, of smaller outer degree; $\wedge$-L on $[\Phi_1 \wedge \Phi_2] = [\Phi_1] \wedge [\Phi_2]$—split inside, box, cut on $[\Phi_i]$, of smaller matrix; or $\bot$-L on $[\bot] = \bot$—the inner sequent proves $\bot$, hence (no right rule for $\bot$ in cut-free LK) the empty succedent, and $(\Box\emptyset)$ with weakenings concludes without any cut. For the universal case, if the left inner derivation ends in a proof of $\forall x.m(\Theta)$, compose it with the classical derivation $\forall x.m(\Theta) \vdash m(\Theta)[t/x]$ and eliminate that inner cut; boxing yields the required left premise for a new outer cut on $[\Theta(t/x)]$. For conjunction, compose with the two projections from $m(\Phi_1) \wedge m(\Phi_2)$ and box the resulting inner derivations. These new cut formulas have strictly smaller outer degree (and, in the boxed atomic comparison, strictly smaller matrix size), so the primary induction applies. Eigenvariables are renamed before these transformations.

Left outer, right box. $\wedge$-R against a box consuming $[\Phi_1 \wedge \Phi_2]$: split the right premise's inner sequent, box, and cut on $[\Phi_1]$, $[\Phi_2]$, of smaller matrices. $\forall$-R against a box consuming $[\forall x.\Theta]$: the Barcan cut, *retained.* In the conjunction case, invert the inner $\wedge$-left rule in the right box premise, obtaining an inner derivation with both $m(\Phi_1)$ and $m(\Phi_2)$ as antecedents; one box inference then exposes $[\Phi_1]$ and $[\Phi_2]$, and the two premises of the outer $\wedge$-right inference are mixed with it successively. Both new cuts have smaller primary rank. In the universal case the left premise has the form

$$\frac{\Gamma \vdash \Delta, [\Theta(y/x)]}{\Gamma \vdash \Delta, \forall x.[\Theta]}(\forall R),$$

with y fresh, while the right box uses the inner assumption $\forall x.m(\Theta)$. Unlike $\wedge$-left, $\forall$-left is not invertible from a derivation using $\forall x.m(\Theta)$ to one using a single arbitrary instance $m(\Theta)[y/x]$; doing so would replace a strictly stronger assumption by a weaker one. Hence no rank-decreasing principal reduction exists. After eigenvariable renaming, the unreduced mix is exactly a Barcan cut in the sense defined above. All other reductions either remove the cut or replace it by cuts of smaller primary rank (or the same primary rank and smaller height sum), so the induction terminates and only Barcan cuts remain. □

Theorem A.20 (Barcan cuts are indispensable): For unary predicates A, B, the sequent

$$\mathrm{Seq} := \forall x.\big([A(x)] \vee [\neg B(x)]\big),\ [\exists x.B(x)] \ \vdash\ \exists x.[A(x)]$$

is valid, derivable with a single Barcan cut, and has no cut-free derivation in G. In particular G does not admit cut elimination, and Theorem A.19 is optimal.

Proof: Validity and the one-cut derivation are as before: the cut-free part derives $\forall x.([A(x)] \vee [\neg B(x)]) \vdash \exists x.[A(x)], \forall x.[\neg B(x)]$ by $\forall$-R, $\forall$-L at the eigenvariable, and $\vee$-L; the box inference $(\Box\emptyset)$ gives $[\forall x.\neg B], [\exists x.B] \vdash$; the cut on $\forall x.[\neg B(x)] = [\forall x.\neg B]$ is a Barcan cut. For the impossibility, suppose there were a cut-free derivation. Cut-free derivations have the outer subformula property, modulo term substitution by the quantifier rules and the transparent identities for $[\wedge]$ and $[\forall]$: this follows by induction on the derivation, since the outer LK rules have the usual property and a box rule has no side formulas. Hence no bare atom and no nested boxed atom occurs, and after inverting outer rules every box leaf has antecedents drawn from the finite families

$$[A(t)], \qquad [\neg B(t)], \qquad [\exists x.B(x)]$$

and has either one member of the same families as succedent or empty succedent.

We claim that the inner sequent of such a box leaf is derivable only when its succedent formula already occurs among its antecedents. If the succedent is empty, interpret the finitely many required $A(t)$ as true and the finitely many required $B(t)$ as false; when $\exists x.B(x)$ is an antecedent, add a fresh element d, distinct from the denotations of those finitely many terms, and set $B(d)$ true. Thus the antecedents are consistent. If the succedent is $A(s)$ and that formula is absent from the antecedent, use the same term-separated structure and set $A(s)$ false. If the succedent is $\neg B(s)$ and it is absent, set $B(s)$ true while keeping all antecedent $\neg B(t)$ true; syntactically distinct closed terms may be interpreted as distinct elements. If the succedent is $\exists x.B(x)$ and it is absent, interpret B as empty. In each non-identity case the antecedents are true and the succedent false, so classical soundness refutes

the inner sequent.

Every box leaf is therefore an identity with weakenings and can be replaced by an ordinary initial sequent. The alleged derivation then becomes a cut-free LK derivation in which $[A(t)]$, $[\neg B(t)]$, and $[\exists x.B]$ are read as uninterpreted atoms $a(t)$, $b(t)$, and c. It would derive

$$\forall x.(a(x) \vee b(x)), c \vdash \exists x.a(x),$$

which is classically invalid in the structure with a empty, b total, and c true. This contradiction proves that no cut-free derivation exists. □

Remark 10 (The Barcan obstruction, localized): The cut formula $[\forall x.\neg B]$ occurs among no outer subformulas of the end-sequent: the derivation must *invent* the Barcan-shaped intermediary "overwhelmingly, B holds nowhere". That the Barcan formula blocks cut elimination in unlabelled calculi for constant-domain quantified modal logics is well known, labelled calculi in the style of Negri being the standard repair; on the extended fragment the phenomenon is *confined* to exactly this shape, and—by Proposition A.17 and Remark 9—neither nesting nor persistence adds to it.

Call a sequent *outer-propositional* if its formulas are generated from bare atoms, boxed atoms $[\Phi]$ (with arbitrary, possibly nested and quantified Φ), and $\bot$ by $\neg, \wedge, \vee, \rightarrow$ only.

Theorem A.21 (Cut elimination on the outer-propositional part): Every derivable outer-propositional sequent has a cut-free derivation. Moreover, for atomic sequents,

$$\vec{P}, [\Phi_1], \ldots, [\Phi_k] \vdash \vec{Q}, [\Psi_1], \ldots, [\Psi_l] \text{ is derivable } \iff$$

$$Q_i \in \vec{P} \text{ for some } i, \quad \text{or} \quad \vec{P} \cup m(\vec{\Phi}) \vdash_{\mathrm{LK}} \emptyset,$$
$$\text{or} \quad \vec{P} \cup m(\vec{\Phi}) \vdash_{\mathrm{LK}} m(\Psi_j) \text{ for some } j,$$

a disjunction property extended to bare atoms.

Proof: Saturate backward with the invertible propositional rules, treating bare and boxed atoms as atoms; leaves are of the displayed atomic form. If a leaf satisfies none of the three disjuncts, refute it by an enriched fan: the theories $F_0 \cup \{\neg m(\Psi_j)\}$ with $F_0 := \vec{P} \cup m(\vec{\Phi})$ are each consistent; extend the signature by one common countable set of fresh constants and complete each to a maximal witnessed theory G_j; their term models M_j share the universe of closed terms and the syntactic functions. For the strict layer take the *minimal* structure N over the same universe, whose atomic diagram is exactly $\vec{P}$ (a single Henkin component suffices if $l = 0$). The inclusion constraint holds since $\vec{P} \subseteq F_0 \subseteq G_j$ for every j; $N \models \vec{P}$ by construction; each $[\Phi_i]$ holds since every component satisfies $m(\Phi_i)$; each $[\Psi_j]$ fails in the j-th component; and each Q_i fails in N simply because $Q_i \notin \vec{P}$. Note that $F_0 \vdash Q_i$ may well hold—the sequent $P, [P \rightarrow Q] \vdash Q$ is underivable and, by exactly this countermodel, invalid: truth in every component does not lift bare atoms into the strict layer. The propositional rules being invertible, the countermodel refutes the end-sequent, contradicting validity; so every leaf closes by an initial sequent, $(\Box\emptyset)$, or $(\Box)$ with weakenings, and the saturation tree is cut-free. The right-to-left directions of the equivalence are one initial sequent or one box inference. □

Proposition A.22 (Fitting collapse at the sequent level): $\vec{P}, [\Phi_1], \ldots, [\Phi_k] \vdash [\Psi]$ is derivable in G iff $\vec{P} \cup m(\vec{\Phi}) \vdash_{\mathrm{LK}} m(\Psi)$; when it is, a single box inference derives it. The internal logic of the box is classical first-order logic, computed on matrices.

Proof: The case $l = 1$, $\vec{Q} = \emptyset$ of Theorem A.21: a consistent F_0 with $F_0 \nvdash m(\Psi)$ is refuted by the one-component enriched fan $(\, N{=}M;\, \{M\}\,)$ on a term model M of a completion of $F_0 \cup \{\neg m(\Psi)\}$. □

Corollary A.23 (Persistence is one-way): $P \vdash [P]$ is a box inference, while $[P] \nvdash P$: the one-component enriched fan with $M \models P$ and the strict layer N interpreting P as false satisfies the atomic inclusion and refutes the converse—matching the computational side, where a single exceptional trace refutes $[P] \rightarrow P$.

Remark 11 (Comparison with full FOS4*):* The ledger against the ambient system now reads: the ordinary calculus for first-order S4 derives the converse Barcan formula cut-free but must postulate BF, upon which cuts are ineliminable and no sharper map exists; here BF is an initial sequent, persistence a rule instance, T and 4 a matrix-invisible schema—and *still* Barcan cuts are indispensable, which relocates the diagnosis from any axiom's status to the intrinsic non-analyticity of constant-domain reasoning. Against that common baseline the extended fragment keeps its three advantages: confinement with a matching lower bound; box–box cuts collapsing to one classical cut on matrices, with no recursion through modal depth even for nested formulas; and the outer-propositional elimination and Fitting collapse, of which full S4 has no analogue. On what is derivable, the fragment is exactly right for computational models while the full language provably is not (Proposition 6.1); on how proofs behave, the Barcan problem is confined, not repaired, and labelled calculi would cure it for both at the price of the machinery the fragment avoids.

Remark 12 (What the tools need): CryptoVampire's sequents are pure first-order logic under one implicit box: by Proposition A.22 their theory is classical logic, all cut elimination happening inside $\vdash_{\mathrm{LK}}$. Squirrel's local sequents live under one box with bare execution facts as passengers, and the propositional skeleton of its global reasoning falls under Theorem A.21. Barcan cuts arise exactly where outer quantifiers genuinely interact with the box—the quantifier-commutation reasoning that `const` mediates—and nowhere else: nesting is free, persistence is free, and the one proof-theoretic cost of the constant domain is confined and optimal.

F. The Propositional Case

We prove Proposition 6.4. Throughout, formulas are propositional; by Lemma A.1 we may take them flat, and we take the hypothesis set finite—for derivability this is finitarity; for the semantic direction, infinite Γ is the propositional instance of

the countable fan realization of Appendix D—so only finitely many atoms $p_1, \ldots, p_n$ occur, inside the brackets and outside.

a) Cut elimination and the disjunction property.: Saturate the given sequent backward with the invertible propositional rules; this terminates, with atomic leaves $\vec{P}, [\phi_1], \ldots, [\phi_k] \vdash \vec{Q}, [\psi_1], \ldots, [\psi_l]$ whose matrices are Boolean formulas over $p_1, \ldots, p_n$. A leaf satisfying one of the three disjuncts closes by an initial sequent or by one interface inference with weakenings. A leaf satisfying none is refutable: for each j pick a Boolean valuation v_j of the atoms satisfying $\vec{P} \cup \{m(\phi_1), \ldots, m(\phi_k)\} \cup \{\neg m(\psi_j)\}$, one valuation v_0 of $\vec{P} \cup \{m(\vec{\phi})\}$ if $l = 0$, and form the depth-one Kripke fan with endpoints the v_j and root valuation exactly $\vec{P}$. The root refutes the leaf, hence—the propositional rules being invertible—the end-sequent; and the fan is a finite persistent constant-domain S4 model, so by Theorem 5.2 this contradicts derivability. Thus if the sequent is derivable at all, every leaf closes, and the saturation tree itself is a cut-free derivation.

b) Computational realization.: For completeness, suppose Φ is not derivable in BC^+ from Γ and the (P)-instances; saturating $\vdash \bigwedge \Gamma \to \Phi$ yields an open leaf, hence a finite fan $(B; v_1, \ldots, v_m)$ satisfying Γ and refuting Φ at its root. We realize it computationally, over the standard tape space $2^{\mathbb{N}} \times 2^{\mathbb{N}}$ with the fair-coin measure as in Appendix D. Fix k with $2^k \geq m$ and partition $\Omega^{\mathcal{C}}$ by the first k bits of the honest tape into cells of fixed measure 2^{-k}; assign to each v_j one cell R_j, merging the spare cells into R_1. Interpret each trace atom p_i as true exactly at the traces of the cells R_j with $v_j(p_i) = 1$. Then set the complete valuation at the single trace $e_\eta := (0^\omega, 0^\omega)$ to be exactly B: an atom holds at e_η iff it lies in B. The ensemble $E = \{e_\eta\}_\eta$ is null, so no $[\phi]$ changes value.

At the full space, a bare atom holds at every trace iff it lies in B: atoms of B hold on every cell, since every v_j satisfies B, and at e_η by construction; other atoms fail on some cell or at e_η. And $[\phi]$ holds iff every $v_j \models \phi$: by Proposition 3.4, $[\phi]$ at $\Omega^{\mathcal{C}}$ says the traces falsifying ϕ form a negligible set; up to the null trace e_η that set is a union of cells of fixed positive measure, hence negligible iff empty of cells, iff no v_j refutes ϕ. This is also why no analogue of the first-order transfer machinery is needed: cell measures do not vary with η, so the diagonal phenomenon of the first-order proof cannot arise. Boolean combinations then evaluate at the full space exactly as at the root of the fan, so the realized model satisfies Γ and refutes Φ there, proving the contrapositive; the converse direction is Corollary 5.3.