Logical Omniscience As Infeasibility

Sergei Artemov[a,*], Roman Kuznets[b,*]

[a] *CUNY Graduate Center, 365 Fifth Avenue, New York, NY 10016, United States*
[b] *Institut für Informatik und angewandte Mathematik, Universität Bern*
Neubrückstrasse 10, 3012 Bern, Switzerland

Abstract

Logical theories for representing knowledge are often plagued by the so-called Logical Omniscience Problem. The problem stems from the clash between the desire to model rational agents, which should be capable of simple logical inferences, and the fact that any logical inference, however complex, almost inevitably consists of inference steps that are simple enough. This contradiction points to the fruitlessness of trying to solve the Logical Omniscience Problem qualitatively if the rationality of agents is to be maintained. We provide a quantitative solution to the problem compatible with the two important facets of the reasoning agent: rationality and resource boundedness. More precisely, we provide a test for the logical omniscience problem in a given formal theory of knowledge. The quantitative measures we use are inspired by the complexity theory. We illustrate our framework with a number of examples ranging from the traditional implicit representation of knowledge in modal logic to the language of justification logic, which is capable of spelling out the internal inference process. We use these examples to divide representations of knowledge into logically omniscient and not logically omniscient, thus trying to determine how much information about the reasoning process needs to be present in a theory to avoid logical omniscience.

Keywords: epistemic logic, logical omniscience, justification logic, complexity theory

---

1. Introduction

Since Hintikka's seminal work [21] in 1962, reasoning about knowledge has entered the realm of formal logic. Following Hintikka, knowledge or belief is represented by a modal operator $\Box$, which is sometimes denoted K for knowledge or B for belief to emphasize the intended meaning. To make formulations simpler (and following the tradition of the book "Reasoning About Knowledge" [15], which also reasoned about belief), we use the term *knowledge* when describing any situation where the fact is subjectively considered true by the agent, even when the fact is objectively false, which makes it belief rather than knowledge. This enables us to sidestep the philosophical discussion of how exactly knowledge is different from belief, a discussion re-energized in 1963 by a short but influential paper of Gettier [17]. In sum, to keep things neutral and general, we use the modality $\Box$ when speaking about Hintikka-type knowledge and/or belief.

Hintikka uses A for objective truth of a statement A and $\Box A$ for the subjective perception of A to be true by the agent. In a multi-agent environment, each agent a has its own modality $\Box_a$. This framework naturally provides a means for describing the subjective perception of one agent about the informational state of other agents by simply nesting modalities: $\Box_a \Box_b A$ says that agent a thinks that agent b considers A to be true. Such considerations are necessary, for instance, in game theory, where agents implement their strategies based on their current information about the world, including their understanding of which information is available to their adversaries.

---

*Corresponding author

Email addresses: sartemov@gc.cuny.edu (Sergei Artemov), kuznets@iam.unibe.ch (Roman Kuznets)

URL: http://web.cs.gc.cuny.edu/~sartemov/ (Sergei Artemov), https://sites.google.com/site/kuznets/ (Roman Kuznets)

This representation of knowledge, however, suffers from a problem noted by Hintikka himself. Despite the ongoing discussion of the exact laws governing knowledge, virtually every modal logic that has been viewed as a logic of knowledge is a *normal modal logic.*[1] Normality means that this logic validates the distribution of knowledge over implication

$$\Box(A \to B) \to (\Box A \to \Box B) \tag{1}$$

and the rule of knowledge necessitation

$$\vdash A \qquad \Longrightarrow \qquad \vdash \Box A \ . \tag{2}$$

These two postulates imply that knowledge is closed under at least the propositional reasoning: (1) means that the agent is capable of performing modus ponens, and (2) says that it can recognize valid facts of the theory. Both postulates can be justified by the assumption that agents are rational reasoners. Working with normal modal logics yields the possibility of using the semantics of Kripke models, which have proved to be a convenient and intuitively clear tool for reasoning about knowledge, based on the Leibnizian supposition of multiple possible worlds.

These postulates, however, have an unrealistic consequence that agents become logically omniscient: not only do the agents know all valid facts, but they must also possess knowledge of all logical consequences of contingent facts they know. This *Logical Omniscience Problem* was identified in [13, 14, 21, 22, 28, 29]. While any scientific representation is always a simplification and/or an idealization of the represented phenomenon, this particular idealization goes too much contrary to our experience. The paradigmatic examples of agents failing to meet logical omniscience benchmarks are finding a non-losing strategy in chess and factoring the product of two very large primes. Both tasks should be achievable if the agent is logically omniscient but are outside of reach for humans and computers alike. For this reason, modal logics of knowledge are sometimes called logics of knowability or of implicit knowledge. They describe the information that is available to the agent in principle: for example, all the information necessary to compute the factorization of a large integer is already present in the integer, which does not necessarily make the computation feasible.

To understand why finding a non-losing strategy is hard, it is helpful to compare agents who have found it with those who couldn't. For this purpose, chess is not an optimal example because of the lack of the former type of agents. Game theoretically, tic-tac-toe is not much different from chess. But the search space for tic-tac-toe is much smaller and the non-losing strategies are well-known and can be easily remembered. Unlike in the case of chess, some players successfully implement these strategies and others don't. While a win against a child can indeed be explained by his/her imperfect reasoning apparatus, a win over an adult is harder to explain within Hintikka's original framework. The only explanation for winning in tic-tac-toe is that the losing player is not logically omniscient.

Naturally, the discrepancy between Hintikka's idealized reasoning agents and real-world imperfect reasoners has received much attention in logic, epistemology, distributed systems, artificial intelligence, game theory and economics, etc. An incomplete list of references includes but is not limited to [1, 7, 11–14, 16, 18, 19, 23, 25–27, 30–32, 34, 35, 37, 38, 40]. These approaches differ in the identified source of cognitive deficiency that prevents agents from being logically omniscient. Following [15, Ch. 9], we now identify the major trends. Most of the approaches add features to standard Kripke models, thereby adding conditions on the attainment of knowledge by the agent. The resulting types of models include syntactic structures, semantic structures (neighborhood semantics), impossible-world structures, awareness structures, and algorithmic knowledge structures:

- Knowledge described in syntactic structures is similar to the situation of the Chinese room thought experiment [36]. The agent possesses a list of facts. Knowledge is fully determined by this list: a fact is known if it is contained in the list. The Chinese room argument has been used, for instance, to describe how a person can win a game of chess, not knowing the rules, by following pre-written instructions from the list. When this paradigm is applied to tic-tac-toe, the agent may lose because it has not been given instructions how to win.

[1]Unless the modal logic has been specifically designed to avoid the Logical Omniscience Problem.

- Neighborhood semantics, or Montague–Scott semantics, is strictly more general than Kripke models. In Kripke models, knowledge comes from the observation that a fact is true in all situations the agent considers possible. In neighborhood semantics, which is still based on possible worlds, the agent knows propositions, i.e., sets of possible worlds. Accordingly, a formula is known if it represents one of the propositions known to the agent. It is, perhaps, harder to relate actual reasoning processes to this view of knowledge, which seems to require that agents possess a high level of introspection. Indeed, the agent must have a rather complete knowledge of the whole possible-world structure, not only of those worlds the agent considers possible. Similarly, the studied fact is to be inspected in all situations, not only in those considered possible. Unlike most other approaches, neighborhood semantics has some minimal level of rationality embedded in it: with every known fact, an agent knows all equivalent facts because this semantics is extensional. While seeming perfectly rational, this does smuggle at least some level of logical omniscience: it is not possible for the agent to know simple tautologies but to be ignorant of more complex ones. In particular, the difference between tic-tac-toe and chess ceases to exits.

- Impossible-world structures are based on the agent's ability to imagine inconsistent or incomplete worlds, i.e., worlds where a fact may be both true and false or neither true nor false respectively, which may distort the agent's knowledge of the world. When modeled in this way, an agent playing tic-tac-toe may lose because of considering a game position that is completely impossible and of making moves to avoid this imaginary position.

- Awareness structures provide a more intuitive view of the shortcomings of the reasoner. An agent may not know a fact that is fully within the agent's grasp simply because the agent has not taken this fact into account. We are all familiar with this "why didn't I think of that?" moment. For instance, an adult may be too lazy to think about a winning strategy and make a bad move not being aware it is a losing move. Arguably, this is what usually happens when a chess player fails to avoid a preventable checkmate by overlooking the adversary's move that leads to this checkmate.

- Algorithmic knowledge structures provide another agent-centered rather than model-centered approach that supplies the agent with (typically unspecified) algorithms for computing knowledge. Thus, the lack of knowledge can be explained by the mismatch between the available algorithms and the facts the knowledge of which is sought. For instance, a student not familiar with predicate calculus may not be able to recognize the validity of the famous "Socrates is mortal" argument. A defeat in tic-tac-toe may result from the agent's inability to count to three.

The feature unifying these approaches is their qualitativeness. Hintikka's definition of knowledge as truth in all possible worlds is supplied (or replaced) with a filter that allows or denies knowledge of a fact based on its presence

- in the list of known facts,
- in the list of known propositions,
- in the list of facts true in each of the impossible worlds the agent imagines possible,
- in the list of facts the agent is aware of, or
- in the list of facts computable by an algorithm.

Predictably, the installation of such a filter destroys the rationality of the agent [20]. The only exception is neighborhood semantics, which makes equivalent statements indistinguishable [15], as mentioned above. It is sometimes possible to tinker with the installed filter to bring the rationality back, e.g., by requiring that the agent which is aware of the fact A be also aware of all subformulas of A. But the overall result is typically all or nothing: either too many logically omniscient features are preserved because of the restrictions on the filter or the agents remain too irrational. As a result, certain situations with insufficient knowledge, especially situations involving computers pre-programmed for a task they are unable to solve, cannot be reasonably explained within these approaches. The difference between tic-tac-toe and chess is one of such phenomena.

The algorithmic knowledge approach can, perhaps, be injected with the quantitative measures necessary to distinguish between the two games by restricting the set of algorithms available to the agent according to the computational complexity of the task. Such restrictions, however, will have to remain on the semantic side. It is not clear how to compute the logic of a particular set of complexity-bound algorithms given that the quantitative measures remain outside of the language of the theory.

Very similar concerns led Duc [11] to introduce effort for gaining knowledge and measure thereof into the object language. In the dynamic-epistemic logic[2] presented in Chapter 4 of his PhD thesis, the knowledge of an agent need not be deductively closed. Instead, an additional modal operator $\langle F \rangle$ is employed, which is similar to the existential future modality of tense logic and essentially represents the effort modality. For instance, in Duc's framework, the normality axiom of modal logic becomes

$$\Box(A \to B) \to (\Box A \to \langle F \rangle \Box B)$$

and says that the agent who knows both $A \to B$ and A may learn B in the future/after applying some effort but need not know B at the current moment. Moreover, the knowledge of B in the future is not guaranteed because the agent is under no obligation to apply the necessary effort and not every reasoning path leads to acquiring knowledge of B. However, according to Duc himself, this framework is rather limited. For instance, it does not allow the effort modality to occur within the scope of the knowledge modality. As a result of this restriction, Duc's dynamic-epistemic logics are shown to correspond to ordinary modal logics of knowledge only with respect to the formulas of modal depth 1 and "the capacities of the dynamic-epistemic systems to model meta-reasoning are rather restricted." More importantly, while the idea that effort, or time, is needed to gain knowledge is relevant for the chess example, the quantitative measures are still missing and tic-tac-toe remains indistinguishable from chess.

To overcome these shortcomings, in Chapter 5 of [11], Duc introduces logics of algorithmic knowledge[3]. The idea is to consider a series of knowledge operators

$$\Box^0, \Box^1, \Box^2, \ldots, \Box^n, \ldots, \Box^\exists \quad ,$$

where numerical superscripts represent the number of reasoning steps sufficient for gaining knowledge.[4] Accordingly, $\Box^0$ represents the actual knowledge the agent possesses and $\Box^\exists$ corresponds to the modality $\Box$ in the standard approach, i.e., to the implicit knowledge obtainable by the agent in principle. This is definitely a step towards explaining the difference between tic-tac-toe and chess. Unfortunately, to the best of our knowledge, this project remains unfinished, with no convenient semantics and with an open decidability problem.[5]

It should be noted that Duc has expanded the definition of what it means for an agent to be rational: the agent should be able not only to perform simple inference actions but also to "compute the complexities of certain reasoning problems in order to infer when something can be known." While the formulation is somewhat vague, it still represents a stronger form of introspection than positive and negative introspection: the agent should know how hard it would be to obtain knowledge about "certain reasoning problems." And it is clear that formalizing this postulate within the object language would make things overly complicated: Peano arithmetic plus complexity theory is not something one needs as a subsystem for describing knowledge.

So while Duc introduced a measuring mechanism into the object language, the counting is still done externally, on the metalevel rather than by the agent itself. Most properties of knowledge achievable in n steps are formulated using unspecified partial functions that describe worst-case complexities of unspecified classes of formulas. The agents are supposed to be able to compute these complexities metalogically [11,

[2] No relation to dynamic epistemic logics in the sense of [10], such as public announcement, action model logics, etc.

[3] No relation to algorithmic knowledge structures from [15].

[4] The idea of introducing the number of numerical steps into the language was to a certain extent already present in [12].

[5] A more recent development towards resource-sensitive logical framework can be found in [2]. However, the complexity of inference is one of the resources not taken into account. In particular, the equivalence rule from [2] postulates that any resource bound b sufficient for enforcing an outcome φ must also be sufficient for any provably equivalent outcome ψ, i.e., $\vdash \varphi \leftrightarrow \psi$ implies $\vdash [C^b]\varphi \leftrightarrow [C^b]\psi$. For instance, any resource bound sufficient for non-losing in tic-tac-toe would also be sufficient for non-losing in chess.

p. 48]. Unfortunately, no details are provided: it is not even clear what data the agent can use for "computing a formula." It is described as "all what [the agent] knows." So most probably, the agents are assumed to have a knowledge set of a finite number of formulas, and complexities state how many steps are sufficient to derive another formula from this set.

Duc should be credited with the idea to employ complexity theory, but this idea remained largely undeveloped. Indeed, implementing this idea meets with certain obstacles. First of all, knowledge in n steps is still to a certain extent a knowability statement: this is the knowledge the agent can achieve after n steps if a correct reasoning strategy is chosen. The agent can still choose to sit on its hands or try a bad reasoning strategy, ending with no information after n time units. One might think that such a non-determinism should be well-suited for at least logics like S5, known to be NP-complete. But it is satisfiability problem that is NP-complete. If one speaks about derivability from the set of formulas already known, a task dual to satisfiability, then the complexity class is co-NP-complete, which (modulo standard complexity-theory assumptions) cannot be solved even in non-deterministic polynomial time.

Thus, it seems that determining knowledge even of propositional facts would require cost functions that are not polynomial. And requiring exponential computations may not be the best way of avoiding logical omniscience. It should also be mentioned that, in the absence of a decidability proof for these algorithmic knowledge theories, it is not even clear that such cost functions exist when reasoning is performed within the theory itself, rather than within its propositional fragment or standard modal logic simplification.

2. Tests for Non-Logical-Omniscience

Like Duc in his algorithmic knowledge approach, we are proposing to look at logical omniscience quantitatively. But we want to use complexity theory rather than abstract away from it. This paper continues the line of research started in [5, 6] and repeats some of the results from there. The main innovation of our approach is a way of testing theories for logical omniscience. We illustrate the need for such a test using Duc's dynamic-epistemic logics as an example. These logics have been specifically designed to avoid logical omniscience. Logical omniscience can manifest itself in several forms, and Duc shows these manifestations not to be present in his dynamic-epistemic logics. For instance,

$$\vdash A \quad \not\Longrightarrow \quad \vdash \Box A \ ,$$
$$\vdash A \to B \quad \not\Longrightarrow \quad \vdash \Box A \to \Box B \ .$$

However, being rational reasoners, agents can learn things by applying effort, which means that

$$\vdash A \quad \Longrightarrow \quad \vdash \langle F \rangle \Box A \ ,$$
$$\vdash A \to B \quad \Longrightarrow \quad \vdash \langle F \rangle \Box A \to \langle F \rangle \Box B \ .$$

So these logics, while not suffering from the usual forms of logical omniscience with respect to $\Box$, have logically omniscient features as long as the rationality of agents comes into play, i.e., with respect to $\langle F \rangle \Box$. More precisely, Duc's dynamic-epistemic logics describe two types of knowledge:

- one avoids the logical omniscience sickness by forfeiting rationality postulates;
- the other views agents as rational reasoners while making them susceptible to logical omniscience.

This should not be too surprising. Despite the overall vagueness of what it means for an agent to be rational and what it means to be logically omniscient, the two concepts seem too close to be easily disentangled from each other. For instance, [15] lists various forms of logical omniscience[6]:

1. knowledge of valid formulas: $\Vdash A \quad \Longrightarrow \quad \vdash \Box A$,

[6] [15] also mentions the so-called "full logical omniscience," which is formulated in terms of classes of models and, thus, is not directly related to agents' internal reasoning processes.

2. closure under logical implication: $A \Vdash B \quad \Longrightarrow \quad \vdash \Box A \to \Box B$,

3. closure under logical equivalence: $A \dashv\Vdash B \quad \Longrightarrow \quad \vdash \Box A \to \Box B$,

4. closure under valid implication: $\Vdash A \to B \quad \Longrightarrow \quad \vdash \Box A \to \Box B$,

5. closure under material implication: $\vdash \Box(A \to B) \wedge \Box A \to \Box B$,

6. closure under conjunction: $\vdash \Box A \wedge \Box B \to \Box(A \wedge B)$.

The first four forms of logical omniscience require knowledge to be derived from the validity of certain facts or implications, which need not be the ability possessed by every rational agent. The last two forms, the closures under material implication and under conjunction, on the other hand, speak directly to the rationality of the agent by postulating the capability to perform simple inference steps: namely, the modus ponens and the conjunction introduction rules respectively. In fact, given an efficient algorithm for deciding the validity of a formula, the difference between the first four and the last two forms of logical omniscience becomes quantitative: one inference step for the latter as opposed to a well-defined but potentially lengthy reasoning process for the former.

Thus, given these considerations and multiple attempts to avoid logical omniscience, it seems impossible to model an agent that is simultaneously rational and free of the above-mentioned forms of logical omniscience. The best hope is to model both the rational knowledge and the non-logically-omniscient knowledge within the same logic, as is done in [11].

These considerations in no way diminish the usefulness of already existing frameworks. Rather, they point out that not all forms of logical omniscience have been sufficiently studied. We, therefore, suggest viewing logical omniscience as a complexity problem and, in particular, including quantitative elements to replace the knowledge/ignorance switch of the above-mentioned approaches with a numerical scale—in order to bridge the gap between rationality and the absence of logical omniscience.

The agent we want to model is first and foremost rational. It possesses sufficient and consistent knowledge of the subject matter to be able to reason about it. The agent is also motivated to work towards obtaining knowledge necessary for performing some tasks. A good approximation of such an agent is a correctly written computer program tasked with computing new information, such as the factorization of a given integer or the best next move in a chess game. Such a program is aware of all relevant information, possesses all relevant subroutines to extract new knowledge, and is not capable of self-deception. The only thing stopping this program from gaining knowledge is a lack of resources. We need to define what these resources are and how we can quantify them.

Since our emphasis is on the rationality of an agent, we identify the agent with the inference system it uses. Since the agent is to reason about knowledge, we call such systems *epistemic systems*. In particular, we assume that some formulas of such an epistemic system, which we call *knowledge assertions*, stipulate knowledge of other formulas. It is also possible to have several types of knowledge assertions in one epistemic system, e.g., $\Box A$ and $\langle F \rangle \Box A$ in Duc's dynamic-epistemic logics. Our goal is to grade formulas based on how feasible it is for the agent to learn them. Let us first give the general idea of the feasibility we mean, with technical details to be supplied later.

Weak Knowledge-Feasibility Test (Artemov, 2005)**.** *An epistemic system E* weakly avoids infeasible logical omniscience *(with respect to knowledge assertions of a certain kind) if the following condition holds: whenever an E-valid knowledge assertion $\mathcal{A}$ (of the specified kind) stipulates that* formula F is known*, this F can be feasibly proved in E.*

Attaining knowledge of F, in our view, is equated with inferring F in the epistemic system. And there must be a feasible way of achieving this. But feasibility is typically considered relative to some input rather than in absolute terms. For us, such a unit for measuring feasibility is the size of the knowledge assertion, i.e., the cost the agent has to pay for formulating the fact of knowing F. Thus, the agent's knowledge becomes naturally graded by the ability to (internally) stipulate the existence of this knowledge. Why should it be the size of the knowledge assertion and not the size of the fact the knowledge of which is being asserted? Because merely stating Fermat's Last Theorem helped no one to obtain knowledge of its validity. The proof has remained elusive despite intensive attempts at finding it—a classical infeasibility scenario. By contrast, the λ-term produced by Coq as a proof of the Four Color Theorem provides sufficient information for the proof to be verified independently and, thus, for attaining knowledge of the theorem. More precisely, a computer capable of processing the term should be able to restore and verify the proof, while a human most probably can do neither.

Under this view, the standard epistemic representation of knowledge should (and will be shown to) fail the Knowledge-Feasibility Test outlined above because the difficulty of obtaining knowledge is not reflected in the language but is replaced with a simple modality $\Box$ or $\langle F \rangle \Box$ in Duc's dynamic-epistemic logics.

However, the Knowledge-Feasibility Test still possesses a flavor of knowability, similar to all the systems of [11]. A proof can be feasibly found, provided the agent chooses the right reasoning path. In the Four Color Theorem example, the λ-term provides enough information to remove the non-determinacy. Thus, a stronger requirement can be imposed:

Strong Knowledge-Feasibility Test (Artemov, 2005)**.** *An epistemic system E* strongly avoids infeasible logical omniscience *(with respect to knowledge assertions of a certain kind) if there exists a feasible deterministic algorithm that constructs a proof of F given an E-valid knowledge assertion $\mathcal{A}$ (of the specified kind) stipulating that* formula F is known.

Thus, feasible knowledge of a non-losing strategy in tic-tac-toe should encode enough information to retrieve the proof that the strategy is non-losing. An average adult is capable of storing this information, provided the proof makes use of the multiple symmetries of the game board. A child may not be able to store so much structured information, perhaps because his/her reasoning theory does not include the laws of symmetry. A child can be taught one of the non-losing strategies, Chinese-room style, but this would have to do more with trustful announcements than with internal knowledge acquisition. In other words, there first must be someone to achieve this knowledge to be able to transfer it to the child. In chess, no such entity exists till this day. We claim that the reason is that the complexity of the proof that a given strategy is non-losing still remains infeasible with the current state of technology.

3. Knowledge-Feasibility Tests, Formally

Using the notion of proof systems due to Cook and Reckhow [8, 33], we give the following definition that we try to make as general as possible to be able to capture the multitude of existing epistemic languages:

Definition 3.1. An *epistemic reasoning system* L consists of

- a language $\mathcal{L}$;
- a measure of size for formulas $|\cdot|\colon \mathcal{L} \to \omega$;
- a polynomial-time computable *proof function* $p\colon \Sigma^* \to \mathcal{L}$ from strings in an alphabet Σ, called *proofs*, to formulas in the language $\mathcal{L}$ of the epistemic reasoning system. An $\mathcal{L}$-formula A is called L-valid, written $\mathsf{L} \vdash A$, iff $A \in p(\Sigma^*)$, i.e., if it has a proof;
- a measure of size for proofs, which is a function $\ell\colon \Sigma^* \to \omega$;
- pairwise disjoint sets $K\mathcal{L}_1, \dots, K\mathcal{L}_N \subseteq \mathcal{L}$, designated as N types of *knowledge assertions*. We denote the set of all knowledge assertions by $K\mathcal{L} := K\mathcal{L}_1 \cup \dots \cup K\mathcal{L}_N$. Each knowledge assertion $\mathcal{A} \in K\mathcal{L}$ has an intended meaning *formula F is known* for a unique formula F;
- a function $OK\colon K\mathcal{L} \to \mathcal{L}$ that extracts these unique formulas F, the *objects of knowledge* of given knowledge assertions $\mathcal{A}$, and that must be
  - computable in time polynomial in $|\mathcal{A}|$ and

– L-validity preserving: for any $\mathcal{A} \in K\mathcal{L}$,

$$\mathsf{L} \vdash \mathcal{A} \qquad \Longrightarrow \qquad \mathsf{L} \vdash OK(\mathcal{A}) \ .$$

The object-of-knowledge extracting function is extended to sets of formulas in the standard way: for any $X \subseteq K\mathcal{L}$,

$$OK(X) := \{OK(F) \mid F \in X\} \ .$$

We require that knowledge assertions of each type be rich enough to express knowledge of any formula of the language: $OK(K\mathcal{L}_i) = \mathcal{L}$.

We use the notation $Th(\mathsf{L}) := \{F \in \mathcal{L} \mid \mathsf{L} \vdash F\}$ to denote the set of valid statements of L.

Most of the epistemic reasoning systems we consider in this paper have either one or two types of knowledge assertions. In the case of a single type of knowledge assertions, $K\mathcal{L} = K\mathcal{L}_1$, and, thus, the subscript can be omitted. However, a multi-agent environment naturally calls for at least one type of knowledge assertions per agent. For each type of knowledge assertions in a given epistemic reasoning system, we define its K-fragment:

Definition 3.2. Let L be an epistemic reasoning system with N types of knowledge assertions $K\mathcal{L}_1, \dots, K\mathcal{L}_N$. The corresponding *K-fragments* $\mathsf{KL}_1, \dots, \mathsf{KL}_N$ are defined as valid knowledge assertions of the respective types: $\mathsf{KL}_i := K\mathcal{L}_i \cap Th(\mathsf{L})$.

By definition of an epistemic reasoning system, $OK(\mathsf{KL}_i) \subseteq Th(\mathsf{L})$. If, in addition, $OK(\mathsf{KL}_i) = Th(\mathsf{L})$, the K-fragment KL_i is called *complete*. In other words, a K-fragment is complete if it provides valid knowledge assertions for all valid facts of the epistemic reasoning system L.

The completeness of K-fragments is related to the knowledge of valid formulas. Since the knowledge of a valid formula should, in principle, be achievable by a rational agent's reasoning, there should be a formula validating this process.

Definition 3.3. Weak Knowledge-Feasibility Test (Weak Test): An epistemic reasoning system L *weakly avoids infeasible logical omniscience* if there exists a polynomial P such that for any L-valid knowledge assertion $\mathcal{A} \in \mathsf{KL}$, there is a proof of $OK(\mathcal{A})$ in L of size bounded by $P(|\mathcal{A}|)$.

Strong Knowledge-Feasibility Test (Strong Test): An epistemic reasoning system L *strongly avoids infeasible logical omniscience* if there is a polynomial-time deterministic algorithm that, for any L-valid knowledge assertion $\mathcal{A} \in \mathsf{KL}$, is capable of restoring a proof of $OK(\mathcal{A})$ in L.

Remark 3.4. Since a polynomial-time algorithm can only restore a polynomial-size proof, the strong avoidance of infeasible logical omniscience is indeed stronger than the weak avoidance thereof. In particular, any epistemic reasoning system incapable of avoiding infeasible logical omniscience weakly is also incapable of avoiding it strongly.

Both tests are parameterized by the epistemic reasoning system, in particular, by the way the size of formulas is measured. In addition, the Weak Test depends on the size measure used for proofs. Thus, one should be careful when applying these tests to existing epistemic reasoning systems. Feasibility can be easily trivialized if, for instance, the size of every proof is taken to be 1.

The most natural complexity measure for proofs represented by words in the alphabet Σ is word length. Apart from this universal measure, applicable to any abstract proof system, we may want to use proof-size measures tailored to a particular type of inference systems. For instance, Hilbert-style derivations can be measured by the number of steps, i.e., the number of formulas in the derivation written linearly, or by the number of logical symbols in the derivation. The latter measure differs from the word length in that all atomic expressions, e.g., propositional atoms, are taken to have size 1 even though their representation in any finite alphabet necessitates the use of indices, which contribute to the word length.[7]

[7] Parentheses are not counted as logical symbols either, but this does not affect the knowledge-feasibility tests because the parenthesis overhead is at most linear and can be disposed of completely if Polish or a similar notation is used.

It is natural to use the same complexity measures for formulas, except that any one formula represents only one step of a Hilbert derivation, which is not a very useful measure. Thus, we only consider the number of logical symbols or the word length as reasonable measures of formula size.

4. Infeasibility of Standard Epistemic Logics

In this section, we outline the relationship between the complexity of a K-fragment of an epistemic reasoning system and the feasibility of the type of knowledge represented by this fragment. We then use this relationship to demonstrate that no epistemic reasoning system that avoids infeasible logical omniscience can be constructed for standard epistemic logics (under reasonable size measures).

Since it is not guaranteed that a given reasoning system is computationally optimal for a given type of knowledge, we introduce a notion of equivalence among epistemic reasoning systems.

Definition 4.1. Epistemic reasoning systems L_1 and L_2 are *equivalent* if they share the object language $\mathcal{L}$, the types of knowledge assertions $K\mathcal{L}_1, \dots, K\mathcal{L}_N$, and the OK function and validate the same formulas, although they may differ in the proof function used and the measures chosen for formulas and proofs.

The following theorem states that an inference system of feasible knowledge can be based on any efficient algorithm for obtaining knowledge:

Theorem 4.2. Let KL_i be a complete K-fragment of an epistemic reasoning system L with respect to knowledge assertions of type $K\mathcal{L}_i$.

1. If KL_i is in NP, there exists an epistemic reasoning system L' equivalent to L that weakly avoids infeasible logical omniscience with respect to $K\mathcal{L}_i$.
2. If KL_i is in P, there exists an epistemic reasoning system L' equivalent to L that strongly avoids infeasible logical omniscience with respect to $K\mathcal{L}_i$.

Proof. We construct the equivalent epistemic reasoning system by providing a new proof function and using the word-length measure for both formulas and proofs. In both cases, there must exist a (non-)deterministic polynomial-time Turing machine M for deciding KL_i.

1. In the case of the non-deterministic Turing machine M, we construct the proof system as follows: each proof consists of an element $\mathcal{A} \in K\mathcal{L}_i$, followed by a sequence of choices made by M, when given $\mathcal{A}$ as its input. Clearly, given the set of choices along one of M's branches, a deterministic polynomial-time Turing machine E can emulate the non-deterministic M. This deterministic machine E outputs $OK(\mathcal{A})$ if the corresponding branch of M's computation is successful. Otherwise, E outputs a fixed L-valid formula. Thus, E outputs only L-valid formulas. Since $OK(\cdot)$ is polynomially computable, so is the function computed by E. Moreover, the range of E contains all L-valid formulas because KL_i is complete. It remains to note that for each valid knowledge assertion $\mathcal{A} \in \mathsf{KL}_i$, there must exist at least one successful run of M that is polynomial in $|\mathcal{A}|$ and, as such, involves only polynomially many choices. Therefore, in the constructed epistemic inference system, there exists a proof of $OK(\mathcal{A})$ polynomial in $|\mathcal{A}|$.

2. In the case of the deterministic Turing machine M, the proof system can be made even simpler. The proofs are just knowledge assertions from $K\mathcal{L}_i$. Given such a knowledge assertion $\mathcal{A} \in K\mathcal{L}_i$, E first runs the given M to determine whether $\mathcal{A}$ is valid. If M succeeds, E outputs $OK(\mathcal{A})$. Otherwise, E outputs a fixed L-valid formula. Again, E is a polynomial-time computable function with its range being the set of all L-valid formulas. Since a proof of $OK(\mathcal{A})$ is nothing but $\mathcal{A}$, finding the proof of $OK(\mathcal{A})$ given $\mathcal{A}$ is trivial, while verifying that it is a proof can be done in polynomial time by the given deterministic Turing machine M. □

The converse to Theorem 4.2 does not in general hold. The fact that an epistemic reasoning system avoids infeasible logical omniscience enables us to guess/find proofs of $OK(\mathcal{A})$ feasible in $|\mathcal{A}|$. But guessing/computing $\mathcal{A} \in OK^{-1}(F)$ for a given formula F may not be feasible or even possible. Thus, given F and a knowledge-feasible epistemic reasoning system, it is unclear how to obtain a proof of F. Additional assumptions are necessary to formulate a partial converse:

Theorem 4.3 (Partial converse to Theorem 4.2). Let L be an epistemic reasoning system in the language $\mathcal{L}$ with word length used as a measure of size for both formulas and proofs. Let KL_i be a complete K-fragment of L with respect to knowledge assertions of type $K\mathcal{L}_i$.

1. If L weakly avoids infeasible logical omniscience with respect to KL_i and there exists a polynomial P such that, for all $\mathcal{A} \in K\mathcal{L}$,
$$|\mathcal{A}| \leq P(|OK(\mathcal{A})|) \ , \tag{3}$$
then $Th(\mathsf{L})$ is in NP.
2. If L strongly avoids infeasible logical omniscience with respect to KL_i and there exists a function $\mathsf{K}\colon \mathcal{L} \to K\mathcal{L}_i$ computable by a deterministic polynomial algorithm and such that, for any formula F, it outputs
   - a knowledge assertion about F, i.e., $\mathsf{K}(F) \in OK^{-1}(F)$,
   - moreover, a valid knowledge assertion in case a formula is valid itself, i.e., $\mathsf{K}(F) \in \mathsf{KL}_i$ whenever $\mathsf{L} \vdash F$,

   then $Th(\mathsf{L})$ is in P.

Proof. 1. By completeness of KL, for every valid $F \in Th(\mathsf{L})$, there must be a valid $\mathcal{A} \in \mathsf{KL}$ such that $OK(\mathcal{A}) = F$. Since L weakly avoids infeasible logical omniscience, there must be a proof of F polynomial in the size of this $\mathcal{A}$. But according to (3), $|\mathcal{A}|$ itself is polynomial in $|F|$. Hence, there is a proof of F polynomial in $|F|$ that can be non-deterministically guessed and then deterministically verified in polynomial time by the given epistemic reasoning system.

2. Given $F \in \mathcal{L}$, $\mathcal{A} = \mathsf{K}(F)$ is computed first. Then, the polynomial algorithm guaranteed by the strong avoidance of infeasible logical omniscience is used to construct a proof of F from $\mathcal{A}$. If F is valid, so is $\mathcal{A}$; therefore, the algorithm outputs an actual proof of F that can be verified by the given epistemic reasoning system. Thus, if the verification is not successful, it means that F is not valid. $\square$

This partial converse is not very elegant but is useful in proving strong negative results about certain epistemic logics. The nature of the argument is that a logic with a sufficiently high lower complexity bound and insufficiently informative knowledge assertions cannot avoid infeasible logical omniscience no matter which inference method the agent uses.

Corollary 4.4. Let a monomodal epistemic logic ML employ $\Box F$ as the statement asserting the knowledge of formula F, let modal necessitation rule be admissible and invertible, i.e.,
$$\vdash F \qquad \Longleftrightarrow \qquad \vdash \Box F \ ,$$
and let ML be conservative over classical propositional calculus (CPC). Let L be an epistemic reasoning system with $Th(\mathsf{L}) = \mathsf{ML}$ and with word length used as the size measure for both formulas and proofs.

1. L cannot weakly avoid the infeasible logical omniscience, unless NP=co-NP.
2. L cannot strongly avoid the infeasible logical omniscience, unless P=NP.

In particular, common epistemic modal logics suffer from the infeasible form of logical omniscience (modulo the stated complexity theory assumptions), in addition to all the other forms of logical omniscience.

Proof. For this very simple type of knowledge assertions, we have
$$K\mathcal{L} = \{\Box F \mid F \in \mathcal{L}\} \ ,$$
$$OK(\Box F) = F \ ,$$
$$OK^{-1}(F) = \{\Box F\} \ .$$

If an epistemic reasoning system L with $Th(\mathsf{L}) = \mathsf{ML}$ weakly avoids infeasible logical omniscience, then the derivability of F is equivalent to the existence of an L-proof of F, polynomial in $|OK^{-1}(F)|$, which is the same as being polynomial in $|F|$ since $|\, OK^{-1}(F)| = |\Box F| = |F| + 1$. If L strongly avoids infeasible logical omniscience, this proof can be found by a deterministic algorithm polynomial in $|F|$.

Guessing a polynomial-size proof and verifying it can be viewed as an NP decision procedure, whereas finding a proof and verifying it is a P algorithm for deciding ML. It remains to note that any modal logic conservative over CPC is co-NP-hard. □

The word length is the real computational measure relevant for automated computation. The situation with more human measures is not much better.

Theorem 4.5. Let a monomodal epistemic logic ML employ $\Box F$ as the statement asserting the knowledge of formula F and be conservative over CPC. Let L be a modal Hilbert-style epistemic reasoning system with $Th(\mathsf{L}) = \mathsf{ML}$, i.e., let a derivation be a list of formulas, let modus ponens and modal necessitation be the only inference rules, and let axioms be given as a finite number of axiom schemes. Let word length be used as the size measure for formulas and the number of formulas in the list as the size measure for proofs. L cannot avoid the infeasible logical omniscience (neither weakly nor strongly), unless NP=co-NP.

Proof. Since a logic conservative over CPC is co-NP-hard, it is sufficient to construct an non-deterministic polynomial-time algorithm for deciding validity in ML based on the assumption that infeasible logical omniscience is weakly avoidable. So let us assume that there exists a polynomial P such that any valid formula F has a Hilbert derivation of $P(|\Box F|)$ formulas, i.e., a derivation with the number of formulas polynomial in $|F|$. We cannot directly guess these formulas in an NP manner because their sizes can *a priori* be arbitrarily large.

To avoid guessing an arbitrary formula, we will instead be guessing derivable schemas and use unification and modified Robinson's algorithm (see [9]) to do the proof schematically. For an arbitrary formula F, non-deterministically guess the structure of a Hilbert proof of F, i.e., for each of the polynomially many formulas, guess whether it is an axiom, or the conclusion of a modus ponens rule, or the conclusion of a necessitation rule. For each rule, also guess which of the other formulas was(were) used as its premise(s); for each axiom, guess to which of the finitely many axiom schemes it belongs. This gives us the structure of the derivation tree or, rather, of the derivation dag because we use the common linear notation for Hilbert proofs that enables to use one formula as a premise for several rules.

We write each axiom in the guessed dag skeleton in the form of the corresponding axiom scheme using variables over formulas (variables in different axioms must be distinct). Then, starting from the axioms, we can restore the proof in a schematic way. Where a necessitation rule needs to be used, we prefix the formula with $\Box$. A case of modus ponens is less trivial. Suppose modus ponens is to be used on schemes $X \to Y$ and Z. Then, unify X with Z using modified Robinson's algorithm from [9] and apply the resulting most general unifier (m.g.u.) to Y. If modus ponens is to be used on schemes U and Z, with U not being an implication, abort the procedure and return failure since modus ponens is not applicable when the first argument is not an implication.[8] Eventually, at the root of the tree, we will obtain the most general form of formulas that can be proved using derivations with this particular dag structure. Unify this form with the formula F.

All unifications can be done in quadratic time of the size of all the formula dags in the derivation dag according to the complexity of modified Robinson's algorithm [9]. Each axiom scheme at the beginning has a constant size, and the number of axioms and rules is polynomial in $|\Box F|$; hence, the whole unification procedure is polynomial in $|F|$. Thus, we have presented an NP algorithm deciding ML. □

5. On the Way to Feasible Knowledge

The assumptions used to prove the unavoidability of infeasible logical omniscience are rather general. We require conservativity over CPC, which seems uncontroversial. Requiring modus ponens to be admissible is

[8] In principle, there exists a scheme unifiable with implications which is not in the implication format itself. It is a single variable over formulas. However, such a scheme cannot possibly occur in the described algorithm because such an occurrence would mean that all formulas of the monomodal language are derivable and ML is inconsistent, which would contradict its conservativity over CPC.

then a natural extension of accepting classical reasoning. Modal necessitation is more controversial because allowing the modal necessitation rule introduces logical omniscience directly, at least in the form of knowledge of valid formulas, and hence is a bad decision. But, by using a particular inference system, a rational agent should be able to derive the valid facts of the system. So we argue that there must be some form of knowledge assertion postulating knowledge of valid facts. Thus, it is the form of knowledge assertions used in modal logic, rather than their existence, that should be blamed for logical omniscience. We discuss the issue of necessitation in more detail in Sect. 5.1.

Modal necessitation guarantees the completeness of the K-fragment for the monomodal language. Duc in [11] chose another path and made the K-fragment with respect to knowledge assertions $\Box A$ incomplete: $\Box A$ is typically not derivable in his dynamic-epistemic logic. However, concerned with the rationality of the agent, Duc enriched the language to create a new form of knowledge assertions, $\langle F \rangle \Box A$. The corresponding K-fragment is still incomplete. It is only complete in the presence of the additional postulate $\langle F \rangle \neg \langle F \rangle \neg A \to \neg \langle F \rangle \neg \langle F \rangle A$[9] and only when restricted to propositional tautologies:

$$\vdash_{CPC} A \qquad \Longrightarrow \qquad \vdash \langle F \rangle \Box A \ .$$

However, even this propositional form of completeness shows that Duc's epistemic reasoning system suffers from infeasible logical omniscience with respect to $\langle F \rangle \Box$ if formulas are measured by the word length and Hilbert derivations by the number of formulas, unless co-NP=NP. To show this, it is sufficient to adapt the proof of Theorem 4.5 to the Hilbert systems from [11]. The only difficulty would be the treatment of the axioms that are schematic but only with respect to what Duc calls *knowledge-persistent formulas*. But it should be straightforward to modify the unification algorithm to account for such a restriction without changing its efficiency.

To summarize, dynamic-epistemic systems from [11] describe two types of knowledge assertions:

1. $\Box A$, which presupposes no rationality of the agent, similarly to the approaches from [15], and
2. $\langle F \rangle \Box A$, which suffers from infeasible logical omniscience, similarly to ordinary epistemic logics.

The reason the proof of Theorem 4.5 is easily applicable is because $\langle F \rangle$ only signifies the presence of effort but does not quantify the amount of effort necessary to gain knowledge. In our terms, the size of a formula asserting knowledge of F is only $|F| + O(1)$, which is not informative enough.

Duc's algorithmic knowledge systems are intended to amend this deficiency. Unfortunately, the completeness of their K-fragments is certainly restricted but in ways dependent on unspecified assumptions about unspecified classes of formulas. We are, thus, unable to provide an analysis of these algorithmic knowledge systems in our framework.

Instead, we will now introduce and explore several new epistemic languages, gradually increasing the information provided by knowledge assertions. The goal of this exploration is to outline a boundary between feasible and infeasible knowledge representations. To keep things simple and manageable, we concentrate on agents whose reasoning strength corresponds to S4, a well-studied and, more importantly, well-behaved epistemic modal logic. We also fix a particular Hilbert-style axiomatization of S4 with finitely many axiom schemes and use it also for more refined theories of knowledge we will consider, sometimes in a properly modified form. The choice of propositional axioms is irrelevant: there is always a polynomial-time reduction, as we know from proof complexity theory. The modal axioms we consider are

$\Box(A \to B) \to (\Box A \to \Box B)$,

$\Box A \to A$,

$\Box A \to \Box\Box A$.

The rules are modus ponens and modal necessitation.

5.1. Removing Necessitation

We start by clarifying the issue of modal necessitation. As mentioned above, it is a bad rule to have if one aims at avoiding logical omniscience. It seems to be the only irresponsible postulate in S4. However,

[9] This postulate guarantees, among other things, that if both premises of an instance of modus ponens can initially be learned by the agent, then the effort to learn one premise does not preclude applying effort to learn the other premise.

removing necessitation does not improve the situation as long as some minimal rationality requirements are maintained. While postulating that the agent knows all validities of the theory may seem too much, we argue that a rational agent should at least know the axioms of its own inference system. Thus, we restrict the modal necessitation rule to axioms of the theory:

$$\frac{A \text{ is an axiom}}{\Box A} \ .$$

It is easy to show that this restriction does not affect the set of theorems: the modal necessitation rule is admissible in this proof system. Thus, Cor. 4.4 remains applicable and the agent using the restricted form of necessitation, which we will call *axiom necessitation* from now on, remains equally logically omniscient, including with respect to infeasible logical omniscience. Cor. 4.4 concerns measuring size as word length. But measuring the size of proofs as the number of steps in a Hilbert derivation is obviously no better. Every instance of axiom necessitation is also an instance of full modal necessitation, making proofs in the restricted system at least as long as those in the standard Hilbert system for S4, which was shown in Theorem 4.5 to be incapable of avoiding infeasible logical omniscience as far as the number of formulas in a proof is concerned.

In short, although an inference system with axiom necessitation postulates only incremental increase in the agent's knowledge, the only change is the increased proof length. This cleansing the postulates of logical omniscience only results in reasoning becoming even costlier, which makes the agent more logically omniscient.

Remark 5.1. The observation that modal necessitation can be restricted to the axioms is, of course, not new. It belongs to the modal logic folklore and has been used all around the modal logic world for various purposes. In particular, it is exploited in [11].

Remark 5.2. The given simple form of the restricted necessitation rule is due to the positive introspection axiom. The modal logic K, for instance, would require a slightly more convoluted formulation.

5.2. Polarizing Knowledge

Knowledge in knowledge assertions is used in two ways: as an assumption and a conclusion. For instance, $\Box(A \wedge B) \to \Box A$ says that the knowledge of a conjunction implies the knowledge of its conjuncts. Thus, the knowledge of A is to be obtained via reasoning while the knowledge of $A \wedge B$ is assumed as given. More generally, if $\Box A$ is read existentially as *there exists an inference for* A and the resulting first-order statement about these inferences is skolemized, any negative occurrence of $\Box$ becomes an assumption about preexisting information/inference while the positive occurrences will represent inferences made by the agent based on this preexisting knowledge. It should then be clear that the logical omniscience problem only concerns the positive occurrences: it does not matter how hard it is to obtain knowledge if we assume it as given. Let us introduce this distinction into the object language.

Consider a bimodal language $\mathcal{L}_-^+$ with modalities $\Box^-$ and $\Box^+$ the formulas of which are obtained by explicitly marking every modality occurrence in a monomodal formula with its polarity. Thus, there is a 1-1 correspondence between monomodal formulas and polarized formulas of $\mathcal{L}_-^+$. One can easily define polarizing and depolarizing functions that translate formulas between the languages. Consider the inference system for $\mathcal{L}_-^+$ that is a polarized version of a Hilbert-style proof system for S4. The axioms are polarized axioms of S4, including

- $\Box^-(\overline{A} \to \overline{B}) \to (\Box^- A \to \Box^+ B)$,
- $\Box^- \overline{A} \to A$,
- $\Box^- \overline{A} \to \Box^+ \Box^+ A$,

where $\overline{C}$ is the result of inverting the polarities of all modalities in C. The rules are polarized modus ponens and axiom necessitation for $\Box^+$:

$$\frac{A \quad \overline{A} \to B}{B} \qquad \text{and} \qquad \frac{A \text{ is an axiom}}{\Box^+ A} \ .$$

The knowledge assertions for this language have the form $\Box^+ A$. It should be clear that the arguments given earlier for S4 are still applicable to this system, given that S4 easily reduces to it.

Remark 5.3. Duc's dynamic-epistemic logics exhibit a resemblance to this system if his $\Box$ is understood as $\Box^-$ and his $\langle F\rangle\Box$ as $\Box^+$.

5.3. Counting the Steps

Since restricting one-step knowledge acquisition to axioms does not improve chances of avoiding logical omniscience, it seems reasonable to count the number of steps. Let the language include countably many modalities $\Box_1, \ldots, \Box_n, \ldots$. The size of formulas is defined in the standard way, except that $|\Box_n A| := |A| + n$. A knowledge assertion of this language is any formula $\Box_n A$. It asserts the knowledge of A and claims that one needs n "reasoning steps" to learn A. Axioms can be learned in one step. If A can be learned in n steps, then one additional step, counting the steps, gives the agent the knowledge of $\Box_n A$. As for modus ponens, if the agent needs m steps to learn one premise and n steps to learn the other premise, then the conclusion would require $m + n + 1$ steps, counting the application of modus ponens as one final step.

The resulting axioms, which extend our fixed classical propositional axiomatics, are

- $\Box_n(A \to B) \to (\Box_m A \to \Box_{m+n+1} B)$,
- $\Box_n A \to \Box_{n+1}\Box_n A$,
- $\Box_n A \to A$,
- $\Box_n A \to \Box_m A$ if $m > n$,

plus the axiom necessitation rule

$$\frac{A \text{ is an axiom}}{\Box_1 A}$$

as the only rule aside from modus ponens. Let us call this system $\mathsf{S4}_{\text{step}}$. We conjecture that this system is a complete match to the standard $\mathsf{S4}$. However, proving this would be beyond the scope of this paper. The example of a similar system, developed by Wang in [39], requires rather intricate work to prove what Wang calls a temporalization theorem. This theorem states that every modality in an $\mathsf{S4}$ theorem can be indexed by an appropriate positive integer to make the statement valid in Wang's logic, which is similar in ideology to $\mathsf{S4}_{\text{step}}$ although counts steps differently.[10] It is, however, easy to prove by induction on the derivation length that $\mathsf{S4}_{\text{step}}$ satisfies the following rationality postulate:

$$\vdash_{\mathsf{S4}_{\text{step}}} A \qquad \Longrightarrow \qquad (\exists n) \quad \vdash_{\mathsf{S4}_{\text{step}}} \Box_n A \ .$$

Given enough time, an agent can learn every valid fact of the theory. Thus, the K-fragment we defined is complete.

Using methods developed by N. Krupski in [24], we can prove that $\vdash_{\mathsf{S4}_{\text{step}}} \Box_n A$ iff $\Box_n A$ can be derived in the following calculus $\mathsf{S4}^*_{\text{step}}$, which is similar to the $*$-calculus from [24]:

$$\frac{\Box_n(A \to B) \quad \Box_m A}{\Box_{n+m+1} B} \ , \qquad \frac{\Box_n A}{\Box_{n+1}\Box_n A} \ ,$$

$$\frac{}{\Box_1 A} \quad \text{if } A \text{ is an axiom}, \qquad \frac{\Box_n A}{\Box_m A} \quad \text{if } m > n.$$

This calculus has a property that any derivation of $\Box_n A$ contains no more than n formulas. This observation paves way to an epistemic reasoning system for $\mathsf{S4}_{\text{step}}$ that avoids logical omniscience if proofs are measured by the number of formulas in them. In this epistemic reasoning system, a proof of A is any $\mathsf{S4}^*_{\text{step}}$-derivation of $\Box_n A$ for some positive n. The weak knowledge-feasibility test is passed because the size of a valid knowledge assertion $\Box_n A$, which states the knowledge of A, is greater than n, which is an upper bound on the number of formulas in a derivation of $\Box_n A$, considered to be a proof of A.

Thus, already this very crude epistemic logic with no guidance for finding knowledge other than how hard it should be, already enables us to avoid infeasible logical omniscience, albeit in a rather artificial reasoning

[10] Another ideologically similar system is Duc's algorithmic-knowledge logics, but they lack the explicit mechanism for establishing $\Box_n A$.

system and with a proof size measure that is useless for actual computations. The artificial reasoning system can, in fact, be improved. Any derivation of $\Box_n A$ in $\mathsf{S4}^*_{\mathrm{step}}$ can be transformed into a derivation of A in $\mathsf{S4}_{\mathrm{step}}$ with only linear increase in size. Indeed, each rule of $\mathsf{S4}^*_{\mathrm{step}}$, except for the axiom necessitation, is derivable in $\mathsf{S4}_{\mathrm{step}}$ by the corresponding axiom and one or two applications of modus ponens, which increases the number of formulas at most threefold. Using the axiom $\Box_n A \to A$ and modus ponens, we now obtain a derivation of A in the original Hilbert-style calculus. Thus, the epistemic reasoning system based on this calculus and with the number of formulas as a proof-size measure also weakly avoids infeasible logical omniscience, a notable improvement compared to standard epistemic modal logics, where infeasible logical omniscience is unavoidable in any reasonable shape or form.

Remark 5.4. It should be noted that the weight of $\Box_n$ in a formula could be made smaller than n: writing a natural number only requires logarithmic number of digits. This was, however, not our intention. The goal was to use the number itself as the weight. For purists who would want to use the real word-length size measure instead, one should use unary notation for the indices of $\Box$.

Remark 5.5. In the $\mathsf{S4}^*_{\mathrm{step}}$ calculus, we have allowed formulas to be reused, as is common in Hilbert-style calculi. This was not necessary. The upper bound on the number of formulas in a derivation of $\Box_n A$ naturally holds even if reusing is prohibited.

5.4. Structuring the Steps

While it may be possible to construct an epistemic reasoning system for $\mathsf{S4}_{\mathrm{step}}$ that weakly avoids logical omniscience with respect to the word-length measure on proofs, the methods used in such a construction can be better explained by means of another example that we introduce now. After all, our goal is not to advocate one definitive epistemic logic with an optimal feasibility-to-conciseness ratio. Rather, we are trying to illustrate the notion of infeasible logical omniscience using a range of examples.

We've already added a measure of how many inference steps are necessary for acquiring knowledge of a given fact. We now add information about the configuration of these steps. Let knowledge modalities be constructed with the help of the following grammar:

$$k ::= [\bullet] \mid [(k \cdot k)] \mid [!\,k] \enspace .$$

Here $\bullet$ plays the role of $\Box_1$, i.e., knowledge of axioms, which is attainable in one step. The two operations on knowledge modalities correspond to applying modus ponens and positive introspection. We define

$$|\bullet| := 1 = |\Box_1| \enspace , \qquad |k_1 \cdot k_2| := |k_1| + |k_2| + 1 \enspace , \qquad |!\,k| := |k| + 1 \enspace ,$$

and $|[k]A| := |k| + |A|$. A knowledge assertion is any formula $[k]A$ asserting the knowledge of A.

The non-propositional axioms are

- $[k_1](A \to B) \to ([k_2]A \to [k_1 \cdot k_2]B)$,
- $[k]A \to [!\,k][k]A$,
- $[k]A \to A$,

plus the axiom necessitation rule

$$\frac{A \text{ is an axiom}}{[\bullet]A} \enspace .$$

Let us call this system $\mathsf{S4}_\bullet$. This system is likely to be weaker than $\mathsf{S4}$, which is easy to amend, but we choose not to do it here to keep things as simple as possible in order to better illustrate the properties of our knowledge-feasibility tests. It is, however, easy to show that all statements provable in $\mathsf{S4}_\bullet$ are $\mathsf{S4}$-compliant.

Note that there is a straightforward size-preserving translation from $\mathsf{S4}_\bullet$ to $\mathsf{S4}_{\mathrm{step}}$ that preserves the propositional part and replaces $[k]A$ with $\Box_{|k|}A^\circ$, where A° is a translation of A. This translation maps derivations in $\mathsf{S4}_\bullet$ to derivations in $\mathsf{S4}_{\mathrm{step}}$ of the same size. An argument similar to the case of $\mathsf{S4}_{\mathrm{step}}$ shows that

$$\vdash_{\mathsf{S4}_\bullet} A \qquad \Longrightarrow \qquad (\exists k) \quad \vdash_{\mathsf{S4}_\bullet} [k]A \enspace ,$$

meaning that the K-fragment of $\mathsf{S4}_\bullet$ is also complete.

Similar to the case of $\mathsf{S4}_{\text{step}}$, the calculus

$$\frac{[k_1](A \to B) \quad [k_2]A}{[k_1 \cdot k_2]B}\,, \qquad \frac{[k]A}{[!\,k][k]A}\,,$$
$$\frac{}{[\bullet]A} \quad \text{if } A \text{ is an axiom,}$$

which we call $\mathsf{S4}^*_\bullet$, is sound and complete for the K-fragment of $\mathsf{S4}_\bullet$.

The additional information that knowledge modalities of $\mathsf{S4}_\bullet$ provide is the structure of the $\mathsf{S4}^*_\bullet$-derivation. In the case of $\mathsf{S4}_{\text{step}}$, we knew how many formulas could be used. Here, in addition to the formula-size of the derivation, we know its tree structure. We do not know which axioms are to be used in the leaves of the tree, but, given a knowledge assertion $[k]A$ derivable in $\mathsf{S4}_\bullet$, we can non-deterministically guess the correct axiom scheme for each leaf and then proceed down the tree using unification in the same way as described in the proof of Theorem 4.5. The only difference is that here variables over knowledge modalities are to be used in addition to variables over formulas. Unlike in Theorem 4.5, here this process enables us to non-deterministically guess an $\mathsf{S4}^*_\bullet$-derivation of size polynomial in $|k|$ because all axiom schemes have size $O(1)$ and unification can be done in polynomial time. Thus, there exists an $\mathsf{S4}^*_\bullet$-derivation of $[k]A$ polynomial in the size of $|k|$ with respect to the word-length measure. The transformation from an $\mathsf{S4}^*_\bullet$-derivation to an $\mathsf{S4}_\bullet$-derivation remains linear.

To summarize, for the logic $\mathsf{S4}_\bullet$, it is possible to design an epistemic reasoning system that weakly avoids infeasible logical omniscience with respect to the word-length size measure for proofs. This size measure is computationally more relevant than the number of formulas. And this epistemic reasoning system can even be Hilbert-style.

5.5. Justification logic

The source of non-determinism in the search for proof in $\mathsf{S4}_\bullet$ was the uncertainty about the axioms used in the leaves of the $\mathsf{S4}^*_\bullet$ tree. Thus, the next natural step is to add this information. One of the simplest ways of doing this is to replace one $\bullet$ with several knowledge modalities, one per axiom scheme. This makes the derivation construction process from the preceding subsection fully deterministic without changing its polynomial complexity. Thus, the resulting logic strongly avoids infeasible logical omniscience. Such a logic is isomorphic to the $+$-free fragment of the logic of proofs with injective constant specification [3–6].

6. Future Work

Our approach covers only absolute knowledge assertions of the form "F is known," which corresponds to the derivability of $\Box F$ in the underlying epistemic logic. However, some cases require extending this analysis to hypothetical reasoning of the form "$\Gamma \Rightarrow F$ is known," which corresponds to the derivability of $\Box F$ from the set of assumptions Γ. Such an analysis would require extending the Cook–Reckhow theory to hypothetical reasoning.

Acknowledgments

Roman Kuznets is supported by Swiss National Science Foundation grant PZ00P2–131706. The authors would like to thank Galina Savukova for her help with the linguistic aspects of the paper.

References

[1] N. Alechina, B. Logan, Ascribing beliefs to resource bounded agents, in: Proceedings of the first international joint conference on Autonomous Agents and Multiagent Systems, AAMAS 2002, Bologna, Italy, July 15–19, 2002, Part 2, ACM Press, 2002, pp. 881–888.

[2] N. Alechina, B. Logan, H. Nga Nguyen, A. Rakib, Logic for coalitions with bounded resources, Journal of Logic and Computation 21 (2011) 907–937.

[3] S.N. Artemov, Operational modal logic, Technical Report MSI 95–29, Cornell University, 1995.

[4] S.N. Artemov, Explicit provability and constructive semantics, Bulletin of Symbolic Logic 7 (2001) 1–36.
[5] S.N. Artemov, R. Kuznets, Logical omniscience via proof complexity, in: Z. Ésik (Ed.), Computer Science Logic, 20th International Workshop, CSL 2006, 15th Annual Conference of the EACSL, Szeged, Hungary, September 25–29, 2006, Proceedings, volume 4207 of *Lecture Notes in Computer Science*, Springer, 2006, pp. 135–149.
[6] S.N. Artemov, R. Kuznets, Logical omniscience as a computational complexity problem, in: A. Heifetz (Ed.), Theoretical Aspects of Rationality and Knowledge, Proceedings of the Twelfth Conference (TARK 2009), ACM, Stanford University, California, 2009, pp. 14–23.
[7] R.J. Aumann, Reasoning about knowledge in economics, in: J.Y. Halpern (Ed.), Theoretical Aspects of Reasoning about Knowledge, Proceedings of the 1986 Conference, March 19–22, 1986, Monterey, California, Morgan Kaufmann, 1986, p. 251. Abstract of invited talk.
[8] S.A. Cook, R. Reckhow, On the lengths of proofs in the propositional calculus (preliminary version), in: Conference Record of Sixth Annual ACM Symposium on Theory of Computing, Papers Presented at the Symposium, Seattle, Washington, April 30–May 2, 1974, ACM Press, 1974, pp. 135–148.
[9] J. Corbin, M. Bidoit, A rehabilitation of Robinson's unification algorithm, in: R.E.A. Mason (Ed.), Information Processing 83, Proceedings of the 9th World Computer Congress, Paris, France, September 19–23, 1983, North-Holland/IFIP, 1983, pp. 909–914.
[10] H. van Ditmarsch, W. van der Hoek, B. Kooi, Dynamic Epistemic Logic, volume 337 of *Synthese Library*, Springer, 2007.
[11] H.N. Duc, Resource-Bounded Reasoning about Knowledge, Ph.D. thesis, Fakultät für Mathematik und Informatik, Universität Leipzig, 2001.
[12] J.J. Elgot-Drapkin, M. Miller, D. Perlis, Memory, reason, and time: The step-logic approach, in: R. Cummins, J. Pollock (Eds.), Philosophy and AI, Essays at the Interface, MIT Press, 1991, pp. 79–104.
[13] R. Fagin, J.Y. Halpern, Belief, awareness, and limited reasoning: Preliminary report, in: A.K. Joshi (Ed.), Proceedings of the 9th International Joint Conference on Artificial Intelligence, IJCAI 1985, volume 1, Morgan Kaufmann, 1985, pp. 491–501.
[14] R. Fagin, J.Y. Halpern, Belief, awareness, and limited reasoning, Artificial Intelligence 34 (1987) 39–76.
[15] R. Fagin, J.Y. Halpern, Y. Moses, M.Y. Vardi, Reasoning about Knowledge, MIT Press, 1995.
[16] R. Fagin, J.Y. Halpern, M.Y. Vardi, A nonstandard approach to the logical omniscience problem, Artificial Intelligence 79 (1995) 203–240.
[17] E.L. Gettier, Is justified true belief knowledge?, Analysis 23 (1963) 121–123.
[18] J.Y. Halpern, Y. Moses, A guide to the modal logics of knowledge and belief: Preliminary draft, in: A.K. Joshi (Ed.), Proceedings of the 9th International Joint Conference on Artificial Intelligence, IJCAI 1985, Los Angeles, CA, August 1985, volume 1, Morgan Kaufmann, 1985, pp. 480–490.
[19] J.Y. Halpern, Y. Moses, A guide to completeness and complexity for modal logics of knowledge and belief, Artificial Intelligence 54 (1992) 319–379.
[20] J.Y. Halpern, R. Pucella, Dealing with logical omniscience: Expressiveness and pragmatics, Artificial Intelligence 175 (2011) 220–235.
[21] J. Hintikka, Knowledge and Belief: An Introduction to the Logic of the Two Notions, Cornell University Press, 1962.
[22] J. Hintikka, Impossible possible worlds vindicated, Journal of Philosophical Logic 4 (1975) 475–484.
[23] K. Konolige, A Deduction Model of Belief, Research Notes in Artificial Intelligence, Morgan Kaufmann, 1986.
[24] N.V. Krupski, On the complexity of the reflected logic of proofs, Theoretical Computer Science 357 (2006) 136–142.
[25] H.J. Levesque, A logic of implicit and explicit belief, in: R.J. Brachman (Ed.), Proceedings of the Fourth National Conference on Artificial Intelligence, AAAI 1984, Austin, Texas, August 6–10, 1984, AAAI Press, 1984, pp. 198–202.
[26] R. Montague, Universal grammar, Theoria 36 (1970) 373–398.
[27] R.C. Moore, Reasoning about knowledge in artificial intelligence, in: J.Y. Halpern (Ed.), Theoretical Aspects of Reasoning about Knowledge, Proceedings of the 1986 Conference, March 19–22, 1986, Monterey, California, Morgan Kaufmann, 1986, p. 81. Abstract of invited talk.
[28] Y. Moses, Resource-bounded knowledge, in: M.Y. Vardi (Ed.), Proceedings of the Second Conference on Theoretical Aspects of Reasoning about Knowledge, March 7–9, 1988, Pacific Grove, California, Morgan Kaufmann, 1988, pp. 261–275.
[29] R. Parikh, Knowledge and the problem of logical omniscience, in: Z.W. Ras, M. Zemankova (Eds.), Methodologies for Intelligent Systems, Proceedings of the Second International Symposium, ISMIS 1987, Charlotte, North Carolina, USA, October 14–17, 1987, North-Holland/Elsevier, 1987, pp. 432–439.
[30] R. Parikh, Logical omniscience, in: D. Leivant (Ed.), Logic and Computational Complexity, International Workshop LCC '94, Indianapolis, IN, USA, October 13–16, 1994, Selected Papers, volume 960 of *Lecture Notes in Computer Science*, Springer, 1995, pp. 22–29.
[31] R. Parikh, Logical omniscience and common knowledge; WHAT do we know and what do WE know?, in: R. van der Meyden (Ed.), Theoretical Aspects of Rationality and Knowledge, Proceedings of the Tenth Conference (TARK 2005), National University of Singapore, Singapore, 2005, pp. 62–77.
[32] R. Parikh, Sentences, belief and logical omniscience, or what does deduction tell us?, The Review of Symbolic Logic 1 (2008) 459–476.
[33] P. Pudlák, The lengths of proofs, in: S.R. Buss (Ed.), Handbook of Proof Theory, volume 137 of *Studies in Logic and the Foundations of Mathematics*, Elsevier, 1998, pp. 547–637.
[34] V. Rantala, Impossible worlds semantics and logical omniscience, Acta Philosophica Fennica 35 (1982) 106–115.
[35] D. Scott, Advice in modal logic, in: K. Lambert (Ed.), Philosophical Problems in Logic, Reidel, 1970, pp. 143–173.
[36] J.R. Searle, Minds, brains, and programs, Behavioral and Brain Sciences 3 (1980) 417–424.

[37] H.S. Shin, T. Williamson, Representing the knowledge of Turing machines, Theory and Decision 37 (1994) 125–146.
[38] M.Y. Vardi, On epistemic logic and logical omniscience, in: J.Y. Halpern (Ed.), Theoretical Aspects of Reasoning about Knowledge, Proceedings of the 1986 Conference, March 19–22, 1986, Monterey, California, Morgan Kaufmann, 1986, pp. 293–305.
[39] R.J. Wang, Temporalizing modal epistemic logic, in: S.N. Artemov, A. Nerode (Eds.), Logical Foundations of Computer Science, International Symposium, LFCS 2013, San Diego, CA, USA, January 6–8, 2013, Proceedings, volume 7734 of *Lecture Notes in Computer Science*, Springer, 2013, pp. 359–371.
[40] H. Wansing, A general possible worlds framework for reasoning about knowledge and belief, Studia Logica 49 (1990) 523–539.