
A substructural epistemic resource logic: theory and modelling applications

DIDIER GALMICHE, *Université de Lorraine, CNRS, LORIA, 54000, Nancy, France.*

PIERRE KIMMEL, *Université de Lorraine, CNRS, LORIA, 54000, Nancy, France.*

DAVID PYM, *University College London, London, WC1E 6BT, UK.*

Abstract

We present a substructural epistemic logic, based on Boolean BI, in which the epistemic modalities are parametrized on agents' local resources. The new modalities can be seen as generalizations of the usual epistemic modalities. The logic combines Boolean BI's resource semantics—we introduce BI and its resource semantics at some length—with epistemic agency. We illustrate the use of the logic in systems modelling by discussing some examples about access control, including semaphores, using resource tokens. We also give a labelled tableaux calculus and establish soundness and completeness with respect to the resource semantics.

1 Introduction

The concept of resource is important in many fields including, among others, computer science, economics and security. For example, in operating systems, processes access system resources such as memory, files, processor time and bandwidth, with correct resource usage being essential for the robust function of the system. The internet can be regarded as a giant, dynamic net of resources, in which Uniform Resource Locators refer to located data and code.

In recent years, the concept of resource has been studied and analysed in computer science through the bunched logic, BI [21, 30, 36] and its variants, such as Boolean BI (BBI) [25] and bunched modal logics [13, 15], and applications, such as Separation Logic [25, 37].

The truth-functional, Kripke semantics of these logics, based on preordered partial monoids is sketched below. However, before proceeding to describe this semantics, it is perhaps worth observing that this choice of structure for BI's models can be motivated directly in terms of natural requirements for the properties of a notion of resource. Assuming a set of resource elements, we expect to be able to

- combine two resource elements, to give a new resource element, and
- compare two resource elements, to determine which is the greater.

It is also natural to expect that the combination of elements be partial and this is indeed amply justified by leading examples. These simple assumptions, that around are cleanly captured by preordered partial monoids, have led to a remarkably useful 'resource semantics'. The need for partiality arises in two ways. Conceptually, we observe that in our semantics of resources it is quite natural to expect that not all combinations of resource elements will exist (Separation Logic [25, 37] provides an immediate and compelling example). Second, partiality is technically convenient for BI's metatheory [21].

These considerations lead to a semantics for BI based on partially ordered partial monoids of worlds,

$$\mathbf{R} = (R, \sqsubseteq, \bullet, e).$$

Here composition of resources is captured by the partial monoidal operation, \bullet , with unit e , and comparison of resources is captured by the partial order \sqsubseteq . Where defined, this structure is required to satisfy the bifunctionality condition that if $r_1 \sqsubseteq s_1$ and $r_2 \sqsubseteq s_2$, then $r_1 \bullet r_2 \sqsubseteq s_1 \bullet s_2$. Let us note that \downarrow denotes definedness of the composition.

Given such structures, the logic BI of bunched implications—see, e.g., [21, 30, 34, 36]—which freely combines intuitionistic propositional additives with intuitionistic propositional multiplicatives—has its Kripke semantics given by the following satisfaction relation, where V is an interpretation of propositional letters in $\wp(R)$, in the usual way:

$$\begin{array}{llll} r \models p & \text{iff} & r \in V(p) & r \models \phi \vee \psi & \text{iff} & r \models \phi \text{ or } r \models \psi \\ r \models \perp & \text{never} & & r \models \phi \wedge \psi & \text{iff} & r \models \phi \text{ and } r \models \psi \\ r \models \top & \text{always} & & r \models \phi \rightarrow \psi & \text{iff} & \text{for all } r \sqsubseteq s, \\ r \models \neg\phi & \text{iff} & r \not\models \phi & & & \text{if } s \models \phi, \text{ then } s \models \psi \end{array}$$

$$\begin{array}{ll} r \models I & \text{iff} \quad e \sqsubseteq r \\ r \models \phi * \psi & \text{iff} \quad \text{there exist } r_1, r_2 \in R \text{ s.t. } r_1 \bullet r_2 \downarrow, r \sqsubseteq r_1 \bullet r_2 \text{ and} \\ & \quad r_1 \models \phi \text{ and } r_2 \models \psi \\ r \models \phi -* \psi & \text{iff} \quad \text{for all } r' \in R, \text{ if } r \bullet r' \downarrow \text{ and } r' \models \phi, \\ & \quad \text{then } r \bullet r' \models \psi. \end{array}$$

This resource semantics for BI—i.e., the interpretation of BI’s semantics in terms of resources—underpins its applications to Separation Logic—and its family of derivatives; see [18, 19] for an extensive discussion—and is mainly concerned with sharing and separation.

Specifically, Separation Logic is usually given as a presentation (often using Hoare triples) of a *specific theory of BBI* for a language of memory cells and pointers with a model based on the stack and the heap [25]. Versions of Separation Logic that are based on (intuitionistic) BI, as given above, are also possible [25].

In BBI, [25, 37], the additives are classical, so that the order is collapsed to equality in the partial monoid. Thus, we have

$$\begin{array}{ll} r \models \phi \rightarrow \psi & \text{iff} \quad \text{if } r \models \phi, \text{ then } r \models \psi \\ r \models I & \text{iff} \quad e = r \\ r \models \phi * \psi & \text{iff} \quad \text{there exist } r_1, r_2 \in R \text{ s.t. } r_1 \bullet r_2 \downarrow, r = r_1 \bullet r_2 \text{ and} \\ & \quad r_1 \models \phi \text{ and } r_2 \models \psi. \end{array}$$

The semantics described above is otherwise unchanged.

Thus, sharing of resources is captured by additive connectives, such as \wedge , while separation of resources is captured by multiplicative connectives, such as $*$. These connectives are the logical kernels of the family of separation logics, with resources being interpreted in various ways, such as memory regions, [25, 37], or elements of other particular monoids of resources [9]. This semantic view of resource stands in stark contrast to the ‘number-of-uses’ reading of Linear Logic (LL)’s proof theory [23]. We shall return to this point in the sequel, where we consider the evolution of a model of system of resources.

This framework of resource semantics has also been extended into modal logic. Specifically, we can set up a conservative extension (a ‘Logic of Separating Modalities’ or LSM [15]) of the

modal logic S4 which adds *multiplicative modalities*—modalities that are parametrized on (local) resources. These modalities are defined relative to two-dimensional worlds, one of which captures the S4 accessibility relation and one of which supports the resource parametrization.

Roughly speaking, an LSM model is a 4-tuple $(W, R, \mathbf{R}, \mathcal{V})$, where W is a set of worlds, R is a partial monoid of ‘resources’ (Res, \bullet, e) , $\mathbf{R} \subseteq (W \times Res) \times (W \times Res)$ is a reflexive and transitive relation and \mathcal{V} is an interpretation of propositional letters in $\wp(W \times Res)$. Then, using the both dimensions of ‘worlds’ to handle, respectively, both classical modality and resource parametrization, we have

$$w, r \models \Diamond_s \phi \text{ iff there exist } w' \in W \text{ and } r' \in R \text{ such that } r \bullet s \downarrow, \\ (w, r \bullet s) \mathbf{R}(w', r') \text{ and } w', r' \models \phi$$

$$w, r \models \Box_s \phi \text{ iff for all } w' \in W \text{ and all } r' \in R, \text{ if } r \bullet s \downarrow \text{ and} \\ (w, r \bullet s) \mathbf{R}(w', r'), \text{ then } w', r' \models \phi.$$

Here, s is the local resource, associated with the modality, and r , in the model, is the ambient resource. The modalities are read as asserting that ϕ is possibly (respectively, necessarily) true at the world (w, r) subject to the availability of additional resource s .

Note that two other pairs of modalities are derivable from these:

- The basic additive modalities:

$$w, r \models \Diamond \phi \text{ iff there exist } w' \in W \text{ and } r' \in R \text{ such that } (w, r) \mathbf{R}(w', r') \\ \text{and } w', r' \models \phi$$

$$w, r \models \Box \phi \text{ iff for all } w' \in W \text{ and all } r' \in R, \text{ if } (w, r) \mathbf{R}(w', r') \text{ then} \\ w', r' \models \phi.$$

- Multiplicative modalities with undetermined additional resource parameters:

$$w, r \models \Diamond_\bullet \phi \text{ iff there exist } w' \in W \text{ and } s, r' \in R \text{ such that } r \bullet s \downarrow, \\ (w, r \bullet s) \mathbf{R}(w', r'), \text{ and } w', r' \models_{\mathcal{M}} \phi$$

$$w, r \models \Box_\bullet \phi \text{ iff for all } w' \in W \text{ and all } s, r' \in R, \text{ if } (r \bullet s \downarrow \text{ and} \\ (w, r \bullet s) \mathbf{R}(w', r')) \text{ then } w', r' \models \phi.$$

Full details of the derivations of these modalities may be found in [15] (Lemma 6), where the conservativity of LSM over S4 is also established (in Section 5). The key feature of BI as a modelling tool (and hence of its specific model Separation Logic) is its control of the representation and handling of resources provided by the resource semantics and the associated proof systems. Notice that, in the semantics given above, the components of the additive conjunction, \wedge , share resources whereas the truth condition for the multiplicative conjunction, $*$, requires separate resources for each component. Notice also that this interpretation extends to the multiplicative implication as follows: \multimap can be seen as (the type of) a function that combines the resource required to support itself with the resource required to support its argument to give the resource required to support the application of the function to its argument (see [30, 31]). Finally, notice also that we do *not* assume (in the manner of hybrid logic) the existence of an atomic proposition for each element ‘ s ’ of the set Res with $r \models s$ iff $r = s$: from the perspective of resource semantics, such an assumption—the motivations for

which would be somewhat technical and essentially syntactic—is not well supported. In particular, we would argue that such an assumption obscures the natural structure of the modalities that we wish to explore and, moreover, imposes a constraint on the relationship between worlds and their properties that we do not wish to take in general. We will return to this point briefly in Section 2.

BI's sequent proof systems employ *bunches*, with two context-building operations: one for the additives—characterized by \wedge , which admits weakening and contraction—and one for the multiplicatives—characterized by $*$, which admits neither weakening nor contraction. Bunches are not finite sequences of formulae but rather are finite trees, with formulae at the leaves and the context-building operations at the internal vertices. For the details of the set-up, see [30, 31, 36].

In this set-up, we have the following right rules for the conjunctions and their corresponding implications, \rightarrow and \multimap :

$$\frac{\Gamma \vdash \phi \quad \Delta \vdash \psi}{\Gamma ; \Delta \vdash \phi \wedge \psi} \wedge R \quad \text{and} \quad \frac{\Gamma ; \phi \vdash \psi}{\Gamma \vdash \phi \rightarrow \psi} \rightarrow R$$

and

$$\frac{\Gamma \vdash \phi \quad \Delta \vdash \psi}{\Gamma , \Delta \vdash \phi * \psi} *R \quad \text{and} \quad \frac{\Gamma , \phi \vdash \psi}{\Gamma \vdash \phi \multimap \psi} \multimap R.$$

Again, details may be found in the references given above.

In this setting, the structural rules of Weakening and Contraction arise as follows:

$$\frac{\Gamma(\phi) \vdash \chi}{\Gamma(\phi ; \psi) \vdash \chi} W \quad \text{and} \quad \frac{\Gamma(\phi ; \phi) \vdash \psi}{\Gamma(\phi) \vdash \psi} C.$$

In the former rule, the leaf ϕ is replaced by the bunch $\phi ; \psi$ and, in the latter rule, the sub-bunch (in the evident sense) $\phi ; \phi$ is replaced by the formula ϕ . In both cases, $;$ (rather than $,$) is used. Again, details may be found in the references given above.

The soundness and completeness of BI's proof systems for the semantics given above is established in [30, 36] and elsewhere and via labelled tableaux in [21], and the completeness of BBI for the partial monoid semantics described above is discussed comprehensively in [27].

The idea of resource semantics as it derives from BI and its models and its use as modelling tool is discussed extensively in [35], in an article that is intended to be widely accessible to logicians and computer scientists.

Girard's LL [23] also decomposes the logical connectives into additive and multiplicative forms (for classical and intuitionistic conjunction and disjunction, but not for intuitionistic implication). However, it does so in a very different way from BI. Instead of employing bunches to allow control of the structural rules, LL introduces the so-called exponentials $!$ and $?$ —modalities, similar to S4's \Box and \Diamond)—which have the following left and right rules:

$$\frac{\Gamma, \phi \vdash \Delta}{\Gamma, !\phi \vdash \Delta} !L \quad \frac{! \Gamma \vdash \phi, ?\Delta}{! \Gamma \vdash !\phi, ?\Delta} !R$$

$$\frac{! \Gamma, \phi \vdash ?\Delta}{! \Gamma, ?\phi \vdash ?\Delta} ?L \quad \frac{\Gamma \vdash \phi, \Delta}{\Gamma \vdash ?\phi, \Delta} ?R.$$

Then the structural rules of Weakening and Contraction arise as

$$\frac{\Gamma \vdash \Delta}{\Gamma, !\phi \vdash \Delta} WL \quad \frac{\Gamma \vdash \Delta}{\Gamma \vdash ?\phi, \Delta} WR$$

and

$$\frac{\Gamma, !\phi, !\phi \vdash \Delta}{\Gamma, !\phi \vdash \Delta} \quad CL \qquad \frac{\Gamma \vdash ?\phi, ?\phi, \Delta}{\Gamma \vdash ?\phi, \Delta} \quad CR.$$

Restricting to a single-conclusioned calculus for intuitionistic LL, we have just the !.

At this point, we may ask what is the relationship between BI and LL. The short answer is that they are essentially incomparable. This is explained in detail in the references given above (e.g., [30, 34, 35]), but the essential point can be seen in terms of their differing treatments of intuitionistic implication. In BI, which can be considered to freely combine intuitionistic propositional logic and multiplicative propositional LL, intuitionistic implication is present directly. In LL, intuitionistic implication, $\phi \supset \psi$, is represented using Girard's translation

$$\phi \supset \psi = !\phi \multimap \psi. \quad (1)$$

Such a representation does not exist in BI. This can be seen, as described in [30, 34, 35], using an argument based on category-theoretic models of BI's proofs. Specifically, BI's proofs are modelled by bi-cartesian doubly closed categories, and there is no endofunctor ! on such a category that satisfies (the interpretation of) (1).

Returning briefly to truth-functional semantics and its resource interpretation, we remark that LL's recently developed Kripke semantics [12] does not, as it stands, admit a direct resource interpretation of the kind outlined above. The possibility of such interpretations is an interesting issue.

Modal extensions of BI, such as MBI [1, 9], DBI and DMBI [13], have been proposed to introduce dynamics into resource semantics. In recent work, the idea of introducing agents, together with their knowledge, into the resource semantics has led to an Epistemic Separation Logic, called ESL, in which epistemic possible worlds are considered as resources [14]. This logic corresponds to an extension of BBI with a knowledge modality, \mathbf{K}_a , such that $\mathbf{K}_a\phi$ means that the agent a knows that ϕ holds.

Various previous works on epistemic logics consider the concept of resource, using a variety of approaches. They include [3, 24, 29]. Here we aim to explore more deeply the idea of epistemic reasoning [16] in the context of resource semantics and its associated logic, by taking the basic epistemic modality \mathbf{K}_a and parametrizing it with a resource s , with the associated introduction of relations not only between resources, according to an agent, but also between composition of resources in different ways. The parametrizing resource may be thought of as being associated with, or local to, the agent. This approach leads to the definition of two new modalities \mathbf{L}_a^s and \mathbf{M}_a^s and, consequently, to a new logic in which, as a leading example, we can obtain an account of access to resources and its control, whether they be pieces of knowledge, locations or other entities. We call this logic *Epistemic Resource Logic* or ERL.

In Section 2, we set up the logic ERL by a semantic definition and, in Section 3, we give the key conservative extension properties of the logic and also introduce a useful sublogic, ERL*. In Section 4, we explain how to use the logic to model and reason about the relationship between a security policy—in the context of access control—and the system to which it is applied (cf. Schneier's gate problem [38]). Our application to systems security policy stands in contrast to other work (e.g., [33]) in which epistemic logic has been applied to the analysis of cryptographic protocols. We complete this section with other examples, including joint access and semaphores, which illustrate the applicability of ERL in these perspectives. In Section 5, we set up a labelled tableaux calculus for ERL and establish soundness with respect to ERL's semantic definition and also completeness from a countermodel extraction method. Let us note that we apply the approach and techniques already used for designing such labelled tableaux for other modal extensions of BBI

[13–15]. Details of the arguments are provided in the appendices. Our arguments encompass also the sublogic ERL*.

Further work will be devoted to further study of the logic and its variants, including intuitionistic and dynamic systems, to local reasoning for resource-carrying agents [25, 37], to connections with other approaches to modelling the relationship between policy and implementation in system management [39] and to approaches involving logics for layered graphs [1, 10]. The work presented here builds upon and strongly develops early ideas presented in [20].

2 An ERL

Epistemic logic is the logic of knowledge and belief. It is concerned with what agents know and believe. The knowledge and beliefs of agents are represented using modalities which assert the truth of propositions relative to agents' judgements of the relationship between worlds [16]. In the setting of resource semantics, worlds are interpreted as representing available resources and agents make judgements about the equivalence of resources.

The language \mathcal{L} of the ERL is obtained by adding two new modal operators \mathbf{L} and \mathbf{M} to the BI language. In order to define the language of ERL, we introduce the following structures: a finite set of agents A ; a finite set of resources Res , with a particular element, e ; an internal composition operator \cdot on Res ($\cdot : Res \times Res \rightarrow Res$); and a countable set of propositional symbols $Prop$. The language \mathcal{L} of ERL is defined as follows:

$$\phi ::= p \mid \perp \mid \top \mid \neg\phi \mid \mathbf{I} \mid \phi \vee \psi \mid \phi \wedge \psi \mid \phi \rightarrow \psi \mid \phi * \psi \mid \phi \multimap \psi \mid \mathbf{L}_a^s \phi \mid \mathbf{M}_a^s \phi,$$

where $p \in Prop$, $a \in A$ and $s \in Res$.

In this context we call s the agent's *local resource*. We also define the following operators: $\tilde{\mathbf{M}}_a^s \phi \equiv \neg \mathbf{M}_a^s \neg \phi$ and $\tilde{\mathbf{L}}_a^s \phi \equiv \neg \mathbf{L}_a^s \neg \phi$. The meanings of these connectives are defined in the sequence of definitions that follow below. For simplicity, we write rs instead of $r \cdot s$ and so write $\mathbf{L}_a^{rs} \phi$ instead of $\mathbf{L}_a^{r \cdot s} \phi$.

Note that we introduce modalities that depend on agents and resources and compare them with previous work on an epistemic extension of BBI [14]. With a slight abuse of notation, we have explicit resources in the language syntax: just as in [15], we must assume that the resource elements present in the syntax of the modalities have counterparts in the partial resource monoid (PRM) semantics. This design choice has consequences both for the expressivity of the logic and for the formulation of the tableaux calculus. In the sequel, \downarrow denotes definedness and \uparrow undefinedness.

DEFINITION 1 (Partial resource monoid).

A PRM is a structure $\mathcal{R} = (R, \bullet)$ such that

- R is a set of *resources* such that $Res \subseteq R$ (which notably means that $e \in R$), and
- $\bullet : R \times R \rightarrow R$ is an operator on R such that, for all $r_1, r_2, r_3 \in R$,
 - \bullet is an extension of \cdot : if $r_1, r_2, r_3 \in Res$, then $r_1 = r_2 \cdot r_3$ iff $r_1 = r_2 \bullet r_3$,
 - e is a neutral element: $r_1 \bullet e \downarrow$ and $r_1 \bullet e = r_1$,
 - \bullet is commutative: if $r_1 \bullet r_2 \downarrow$, then $r_2 \bullet r_1 \downarrow$ and $r_2 \bullet r_1 = r_1 \bullet r_2$ and
 - \bullet is associative: if $r_1 \bullet (r_2 \bullet r_3) \downarrow$, then $(r_1 \bullet r_2) \bullet r_3 \downarrow$ and $(r_1 \bullet r_2) \bullet r_3 = r_1 \bullet (r_2 \bullet r_3)$.

We call e the *unit resource* and \bullet the *resource composition*. Henceforth, $\wp(R)$ denotes the powerset of R .

Note that we implicitly consider that the resource composition \bullet is compatible with equality between resources. That means that if $r_1 = r_2$ and $r_1 \bullet r_3 \downarrow$, then $r_2 \bullet r_3 \downarrow$ and $r_2 \bullet r_3 = r_1 \bullet r_3$ (right-composition property of \bullet). We also have the left-composition since \bullet is commutative.

DEFINITION 2 (Model).

A model is a triple $\mathcal{M} = (\mathcal{R}, \{\sim_a\}_{a \in A}, V)$ such that

- $\mathcal{R} = (R, \bullet)$ is a PRM,
- for all $a \in A$, $\sim_a \subseteq R \times R$ is an equivalence relation and
- $V : \text{Prop} \rightarrow \wp(R)$ is a valuation function.

We can place this logic in the context of our previous work on modal [9, 10] and epistemic extensions of (Boolean) BI [13, 14]. In [14], an epistemic extension of BBI, called ESL, is introduced. In this logic, there is just one epistemic modality, K_a , which allows the knowledge of an agent a to be expressed. The modalities employed in this system and those employed in the system presented herein stand in contrast to the modalities of the system LSM described in Section 1 in that they make essential use of the notion of agent in their definition.

More formally, the semantics of this modality is defined by $r \models_{\mathcal{M}} K_a \phi$ if and only if, for all r' such that $r \sim_a r'$, $r' \models_{\mathcal{M}} \phi$, where r and r' are semantic worlds (or resources) and \sim_a is a relation between worlds that expresses that they are equivalent from the point of view of the agent a . The parametrization of modalities on resources derives from ideas that are conveniently expressed in e.g., [9, 10].

In this paper, we aim to develop the idea in order to consider a modality like K_a and to parametrize it on a resource s , requiring the world relation to be of the form $r \bullet s \sim_a r'$ or $r \sim_a r' \bullet s$ or even $r \bullet s \sim_a r' \bullet s$. Then, in the spirit of ESL, we define a new logic from BBI that allows us to model not only relations between resources according to an agent but also how those relations are restricted by resources. We can also consider the resources upon which the agent's relation are parametrized to be local to the agent.

In this spirit, we define two new modalities $\mathbf{L}_a^s \phi$ and $\mathbf{M}_a^s \phi$, with the notation building on the usual one in epistemic logic, for which we have the following semantics expressing two forms of the agent's contingency for truth in the presence of composable resources:

1. $\mathbf{L}_a^s \phi$ expresses that the agent, a , can establish the truth of ϕ using a given resource whenever the ambient resource, r , can be combined with the agent's local resource, s , to yield a resource that a judges to be equivalent to that given resource.

In other words $\mathbf{L}_a^s \phi$ is true relative to the ambient resource, r , iff for a 's views of the combination of the ambient resource, r , and its local resource, s , ϕ is true. More formally we have

$$r \models_{\mathcal{M}} \mathbf{L}_a^s \phi \text{ iff if } r \bullet s \downarrow \text{ then for all } r' \in R, \text{ if } r \bullet s \sim_a r', \text{ then } r' \models_{\mathcal{M}} \phi.$$

2. $\mathbf{M}_a^s \phi$ expresses that the agent, a , can establish the truth of ϕ if there exists a resource that can be combined with its local resource, s , such that a judges the combined resource to be equivalent to the ambient resource, r .

In other words, $\mathbf{M}_a^s \phi$ is true relative to the ambient resource, r , iff for a 's views, the ambient resource is the combination of the local resource, s , with another resource that makes ϕ true. More formally we have

$$r \models_{\mathcal{M}} \mathbf{M}_a^s \phi \text{ iff there exists } r' \in R \text{ such that } r' \bullet s \downarrow \text{ and } r \sim_a r' \bullet s \text{ and } r' \bullet s \models_{\mathcal{M}} \phi.$$

ERL can thus be seen as a particular epistemic logic that provides new modalities which model access to resources, whether they are interpreted as pieces of knowledge, locations or otherwise.

Note that we could obtain operators with similar semantics by taking ESL [14] and adding it the hybrid operators of the hybrid logic HyBBI [4]. Such a new logical framework would allow us to use symbols, called nominals, that force a formula to be valid for a specific resource. Namely, if we consider a nominal n_s forcing the resource s , we then could define the modality $\mathbf{L}_a^s\phi$ by $\mathbf{L}_a^s\phi \equiv n_s \multimap \mathbf{K}_a\phi$ and we recover the semantics given in this section for this modality. Moreover, we could also define the modality $\mathbf{M}_a^s\phi$ by $\mathbf{M}_a^s\phi \equiv \tilde{\mathbf{K}}_a((\top * n_s) \wedge \phi)$. Observations like this are quite common for logics of the kinds considered here, but our view is that conceptual clarity, rather than syntactic ingenuity, should drive the design choices.

This hybrid approach based on nominals represents a significant technical addition to our semantic assumptions that is not justified by the motivations of resource semantics, adding a confusion between resources and propositions that we consider to be inconvenient for our intended modelling applications. Moreover, we would argue that the identities between the modalities that are induced obscure rather than elucidate their meaning—although we would concede that the identities may be of use in mechanical implementations—and lead to a less elegant analysis. Furthermore, working with the hybrid semantics requires additional work in setting the tableaux-based metatheory for the logic, as discussed in Section 5.

It therefore seems appropriate to add the epistemic operators systematically in a clean semantic setting.

DEFINITION 3 (Satisfaction and validity).

Let $\mathcal{M} = (\mathcal{R}, \{\sim_a\}_{a \in A}, V)$ be a model. The satisfaction relation $\models_{\mathcal{M}} \subseteq R \times \mathcal{L}$ is defined, for all $r \in R$, as follows:

$$\begin{array}{llll}
 r \models_{\mathcal{M}} p & \text{iff} & r \in V(p) & r \models_{\mathcal{M}} \phi \vee \psi & \text{iff} & r \models_{\mathcal{M}} \phi \text{ or } r \models_{\mathcal{M}} \psi \\
 r \models_{\mathcal{M}} \perp & \text{never} & & r \models_{\mathcal{M}} \phi \wedge \psi & \text{iff} & r \models_{\mathcal{M}} \phi \text{ and } r \models_{\mathcal{M}} \psi \\
 r \models_{\mathcal{M}} \top & \text{always} & & r \models_{\mathcal{M}} \phi \rightarrow \psi & \text{iff} & \text{if } r \models_{\mathcal{M}} \phi, \text{ then } r \models_{\mathcal{M}} \psi \\
 r \models_{\mathcal{M}} \neg\phi & \text{iff} & r \not\models_{\mathcal{M}} \phi & & & \\
 \\
 r \models_{\mathcal{M}} \mathbf{I} & \text{iff} & r = e & & & \\
 r \models_{\mathcal{M}} \phi * \psi & \text{iff} & \text{there exist } r_1, r_2 \in R \text{ s.t. } r_1 \bullet r_2 \downarrow, r_1 \bullet r_2 = r \text{ and } r_1 \models_{\mathcal{M}} \phi \text{ and } r_2 \models_{\mathcal{M}} \psi & & & \\
 r \models_{\mathcal{M}} \phi \multimap \psi & \text{iff} & \text{for all } r' \in R, \text{ if } r \bullet r' \downarrow \text{ and } r' \models_{\mathcal{M}} \phi, \text{ then } r \bullet r' \models_{\mathcal{M}} \psi & & & \\
 \\
 r \models_{\mathcal{M}} \mathbf{L}_a^s\phi & \text{iff} & \text{if } r \bullet s \downarrow \text{ then for all } r' \in R, \text{ if } r \bullet s \sim_a r', \text{ then } r' \models_{\mathcal{M}} \phi & & & \\
 r \models_{\mathcal{M}} \mathbf{M}_a^s\phi & \text{iff} & \text{there exists } r' \in R \text{ such that } r' \bullet s \downarrow \text{ and } r \sim_a r' \bullet s \text{ and } r' \bullet s \models_{\mathcal{M}} \phi. & & &
 \end{array}$$

A formula ϕ is *valid*, denoted $\models \phi$, if and only if, for any model \mathcal{W} and any resource r , we have $r \models_{\mathcal{M}} \phi$.

PROPOSITION 1 (Satisfaction for the secondary modalities).

Let $\mathcal{M} = (\mathcal{R}, \{\sim_a\}_{a \in A}, V)$ be a model, and let $r \in R$. The following statements hold:

1. $r \models_{\mathcal{M}} \tilde{\mathbf{L}}_a^s\phi$ iff if $r \bullet s \downarrow$ then there exists $r' \in R$ such that $r \bullet s \sim_a r'$ and $r' \models_{\mathcal{M}} \phi$;
2. $r \models_{\mathcal{M}} \tilde{\mathbf{M}}_a^s\phi$ iff for all $r' \in R$, if $r' \bullet s \downarrow$ and $r \sim_a r' \bullet s$, then $r' \bullet s \models_{\mathcal{M}} \phi$.

PROOF. Consider the first part, 1. $\tilde{\mathbf{L}}_a^s\phi \equiv \neg\mathbf{L}_a^s\neg\phi$, so $r \models_{\mathcal{M}} \tilde{\mathbf{L}}_a^s\phi$ iff $r \models_{\mathcal{M}} \neg\mathbf{L}_a^s\neg\phi$ iff $r \not\models_{\mathcal{M}} \mathbf{L}_a^s\neg\phi$ iff there exists $r' \in R$ s.t. $r \bullet s \sim_a r'$ and $r' \not\models_{\mathcal{M}} \neg\phi$ iff there exists $r' \in R$ s.t. $r \bullet s \sim_a r'$ and $r' \models_{\mathcal{M}} \phi$. Proof of 2 is similar. \square

More intuitively, we can see that $\tilde{\mathbf{L}}_a^s\phi$ expresses that the agent, a , can establish the truth of ϕ if there exists a resource such that the combination of the ambient resource, r , and the local resource, s , is judged by a to be equivalent to that resource. Similarly, $\tilde{\mathbf{M}}_a^s\phi$ expresses that the agent, a , can establish the truth of ϕ using a resource that is the combination of its local resource, s , with any resource such that a judges the combined resource to be equivalent to the ambient resource, r . We shall see later that these dual modalities can be also useful for modelling systems.

Returning to the possible representation of the modalities in a hybrid version of ESL, we could then define these modalities as follows: $\tilde{\mathbf{L}}_a^s\phi \equiv (\top * n_s) \wedge \tilde{\mathbf{K}}_a\phi$ and $\tilde{\mathbf{M}}_a^s\phi \equiv \mathbf{K}_a((\top * n_s) \rightarrow \phi)$, with n_s being a nominal forcing the resource s . As we have previously explained, here we aim at avoiding confusion between resources (which are part of the model) and propositions (which are part of the language) that we consider to be inconvenient for our intended modelling applications.

Note that the first point of the definition of \bullet , in Definition 1, implies that the three other definitions (neutral element, commutativity and associativity) extend to \cdot , so that the following are semantically equivalent (i.e., every valid formula in the one is valid in the other) for any agent a and any resources r, s and t : $\mathbf{L}_a^{re}\phi \equiv \mathbf{L}_a^r\phi$, $\mathbf{L}_a^{rs} \equiv \mathbf{L}_a^{sr}$ and $\mathbf{L}_a^{r(st)} \equiv \mathbf{L}_a^{(rs)t}$. Of course, such equivalences also hold for $\mathbf{M}\phi$, $\tilde{\mathbf{L}}\phi$, and $\tilde{\mathbf{M}}\phi$.

3 Some properties of ERL

We show that ERL is a conservative extension of BBI and Epistemic Logic (EL) and that, in the presence of additional properties of the partial resource monoid (Definition 1), there are some noteworthy relationships between modalities.

We consider two fragments of ERL. First, ERL_{BBI} —corresponding to BBI [25]—with $A = \emptyset$ on the language \mathcal{L}_{BBI} defined as \mathcal{L} excluding the \mathbf{L}_a^s and \mathbf{M}_a^s operators. Second, ERL_{EL} —corresponding to the epistemic logic EL consisting of classical propositional additives and the basic epistemic operator \mathbf{K}_a [16]—with $\text{Res} = \{e\}$, on the language \mathcal{L}_{EL} defined as \mathcal{L} excluding \mathbf{I} , $*$ and $\neg*$ and with \mathbf{L}_a^s and \mathbf{M}_a^s , replaced by the operator \mathbf{K}_a , which is defined, for all agents a , by $\mathbf{K}_a\phi = \mathbf{L}_a^e\phi = \mathbf{M}_a^e\phi$.

PROPOSITION 2 (ERL is a conservative extension of BBI and EL).

If, in every model of BBI, the neutral element of the composition is the element e of Res , then ERL_{BBI} is semantically equivalent to BBI. If the agent sets are the same for the two languages, ERL_{EL} is semantically equivalent to the epistemic logic EL.

We now consider some properties of ERL; specifically, the way in which the different operators behave when they are used together in formulae. One interesting property we might require in our semantics, which is based on monoidal structure, is the compatibility of \sim_a and \bullet . More precisely, we might require that if two resources are equivalent for an agent a , then the composition with a third resource be transferred through this equivalence.

Although such a property can be very useful, it introduces, from the modelling perspective, some quite strong properties: the transmission of properties of resources through agent-dependent equivalence is a strong assertion regarding agents' private accesses and should be avoided when modelling some security properties.

Considering these concerns, we take this extra property to be optional and identify it in a sublogic of ERL which we call ERL^* .

DEFINITION 4

The logic ERL^* is defined as ERL with the addition of the following property to the partial resource monoid (Definition 1): For any agent a and any resources $r, r' \in R$, if $r \bullet s \downarrow$ and $r \sim_a r'$, then $r' \bullet s \downarrow$ and $r \bullet s \sim_a r' \bullet s$. It is called the compatibility of \sim_a with \bullet .

Note that we use the logic ERL^* in the security modelling examples that we develop in the next section.

LEMMA 1

Let $a \in A$ be an agent, $s, t \in \text{Res}$ be resources and ϕ be a formula of ERL^* . We have the following properties:

1. $\mathbf{L}_a^s(\mathbf{L}_a^t\phi) \equiv \mathbf{L}_a^{st}\phi$
2. $\mathbf{M}_a^s(\mathbf{M}_a^t\phi) \rightarrow \mathbf{M}_a^t\phi$
3. $\mathbf{L}_a^s\phi \rightarrow \tilde{\mathbf{M}}_a^t(\mathbf{L}_a^s\phi)$
4. $\mathbf{M}_a^t(\tilde{\mathbf{L}}_a^s\phi) \rightarrow \tilde{\mathbf{L}}_a^s\phi$
5. $\tilde{\mathbf{L}}_a^t(\tilde{\mathbf{L}}_a^s\phi) \equiv \tilde{\mathbf{L}}_a^{ts}\phi$
6. $\tilde{\mathbf{M}}_a^s\phi \rightarrow \tilde{\mathbf{M}}_a^t(\tilde{\mathbf{M}}_a^s\phi)$
7. $\mathbf{L}_a^e\phi \equiv \tilde{\mathbf{M}}_a^e\phi$

PROOF. First consider 1. Let \mathcal{W} be a model and r be a resource. Suppose that $r \models_{\mathcal{M}} \mathbf{L}_a^s(\mathbf{L}_a^t\phi)$. Then we have $r \bullet s \downarrow$ and, for any $r' \in R$ such that $r \bullet s \sim_a r'$, we have $r' \models_{\mathcal{M}} \mathbf{L}_a^t\phi$. Thus, $r \bullet s \downarrow$ and, for any $r' \in R$ such that $r \bullet s \sim_a r'$, $r' \bullet t \downarrow$, and for any $r'' \in R$ such that $r' \bullet t \sim_a r''$, we have $r'' \models_{\mathcal{M}} \phi$. Consider $r''' \in R$ such that $r \bullet s \bullet t \sim_a r'''$. By reflexivity, we obtain $r \bullet s \sim_a r \bullet s$. Then with $r' = r \bullet s$ and $r'' = r'''$, we have $r \bullet s \bullet t \downarrow$ and $r''' \models_{\mathcal{M}} \phi$. Thus, $r \models_{\mathcal{M}} \mathbf{L}_a^{st}\phi$, and we can deduce that $\mathbf{L}_a^s(\mathbf{L}_a^t\phi) \rightarrow \mathbf{L}_a^{st}\phi$.

Now suppose that $r \models_{\mathcal{M}} \mathbf{L}_a^{st}\phi$. Then $r \bullet s \bullet t \downarrow$ and, for any r''' such that $r \bullet s \bullet t \sim_a r'''$, we have $r''' \models_{\mathcal{M}} \phi$. As $r \bullet s \bullet t \downarrow$, we have $r \bullet s \downarrow$. Let $r' \in R$ be such that $r \bullet s \sim_a r'$. Then, by compatibility, $r' \bullet t \downarrow$ and $r \bullet s \bullet t \sim_a r' \bullet t$. Let r'' be such that $r' \bullet t \sim_a r''$. Then, by transitivity, we have $r \bullet s \bullet t \sim_a r''$. Then, with $r''' = r''$, we have $r'' \models_{\mathcal{M}} \phi$. We obtain $r \bullet s \downarrow$ and, for any $r' \in R$ such that $r \bullet s \sim_a r'$, $r' \bullet t \downarrow$ and for any $r'' \in R$ such that $r' \bullet t \sim_a r''$, we have $r'' \models_{\mathcal{M}} \phi$. Then we have $r \models_{\mathcal{M}} \mathbf{L}_a^s(\mathbf{L}_a^t\phi)$, and then we can deduce $\mathbf{L}_a^{st}\phi \rightarrow \mathbf{L}_a^s(\mathbf{L}_a^t\phi)$. Finally, we have $\mathbf{L}_a^s(\mathbf{L}_a^t\phi) \equiv \mathbf{L}_a^{st}\phi$.

Now consider 6. Let \mathcal{W} be a model and r be a resource. Suppose that $r \models_{\mathcal{M}} \tilde{\mathbf{M}}_a^s\phi$. Then, for any r' such that $r' \bullet s \downarrow$ and $r \sim_a r' \bullet s$, we have $r' \bullet s \models_{\mathcal{M}} \phi$. Let r'' such that $r'' \bullet t \downarrow$ and $r \sim_a r'' \bullet t$ and r''' such that $r'' \bullet s \downarrow$ and $r'' \bullet t \sim_a r''' \bullet s$. By transitivity we deduce that $r \sim_a r''' \bullet s$ and if we fix $r' = r'''$ we have $r''' \bullet s \models_{\mathcal{M}} \phi$. As it is true for any r''' such that $r'' \bullet s \downarrow$ and $r'' \bullet t \sim_a r''' \bullet s$, we have $r'' \models_{\mathcal{M}} \tilde{\mathbf{M}}_a^s\phi$. As it is true that, for any r'' such that $r'' \bullet t \downarrow$ and $r \sim_a r'' \bullet t$, we have $r \models_{\mathcal{M}} \tilde{\mathbf{M}}_a^t(\tilde{\mathbf{M}}_a^s\phi)$, then for any resource r in any model \mathcal{W} , $\tilde{\mathbf{M}}_a^s\phi \rightarrow \tilde{\mathbf{M}}_a^t(\tilde{\mathbf{M}}_a^s\phi)$ is valid.

Note that the reverse implication, $\tilde{\mathbf{M}}_a^t(\tilde{\mathbf{M}}_a^s\phi) \rightarrow \tilde{\mathbf{M}}_a^s\phi$, is not valid. In fact, if $r \models_{\mathcal{M}} \tilde{\mathbf{M}}_a^t(\tilde{\mathbf{M}}_a^s\phi)$, ϕ is validated by all $r'' \bullet s$ such that $r \sim_a r' \bullet t$ and $r' \bullet t \sim_a r'' \bullet s$. But to have $r \models_{\mathcal{M}} \tilde{\mathbf{M}}_a^s\phi$, we must have $r'' \bullet s \models_{\mathcal{M}} \phi$ for all r'' such that $r \sim_a r'' \bullet s$, and not only for those for which the equivalence by \sim_a is built from t . Then there is no equivalence between $\tilde{\mathbf{M}}_a^s\phi$ and $\tilde{\mathbf{M}}_a^t(\tilde{\mathbf{M}}_a^s\phi)$.

All of the other cases are proved in similar ways. \square

We can complete our language with another modality $\mathbf{N}_a^s\phi$ that could also be helpful for our modelling perspective. From this modality, that is a variant of $\mathbf{L}_a^s\phi$, we can also derive $\tilde{\mathbf{N}}_a^s\phi$ such that $\tilde{\mathbf{N}}_a^s\phi \equiv \neg \mathbf{N}_a^s \neg \phi$.

$\mathbf{N}_a^s\phi$ expresses that the agent, a , can establish the truth of ϕ using any resource combined with its local resource, s , provided a judges that combination to be equivalent to the combination of the local resource, s , with the ambient resource, r . In other words, $\mathbf{N}_a^s\phi$ is true relative to the ambient resource r iff for a 's views of the combination of the ambient resource r and its local resource s , ϕ is true.

More formally we have the following:

$$r \models_{\mathcal{M}} \mathbf{N}_a^s \phi \quad \text{iff} \quad \text{if } r \bullet s \downarrow \text{ then for all } r' \in R \text{ s.t. } r' \bullet s \downarrow \text{ if } r \bullet s \sim_a r' \bullet s, \text{ then } r' \bullet s \models_{\mathcal{M}} \phi.$$

We can build $\mathbf{N}_a^s \phi$ from the previous main modalities as follows.

PROPOSITION 3

We have $\mathbf{N}_a^s \phi \equiv \mathbf{L}_a^s(\tilde{\mathbf{M}}_a^s \phi)$.

PROOF. Consider that $r \models_{\mathcal{M}} \mathbf{L}_a^s(\tilde{\mathbf{M}}_a^s \phi)$ iff, for all $r' \in R$, if $r \bullet s \sim_a r'$, then $r' \models_{\mathcal{M}} \tilde{\mathbf{M}}_a^s \phi$ iff, for all $r' \in R$, if $r \bullet s \sim_a r'$, then, for all $r'' \in R$, if $r' \sim_a r'' \bullet s$, then $r'' \bullet s \models_{\mathcal{M}} \phi$ iff, for all $r', r'' \in R$, if $r \bullet s \sim_a r'$ and $r' \sim_a r'' \bullet s$, then $r'' \bullet s \models_{\mathcal{M}} \phi$ iff (by the transitivity of \sim_a), for all $r'' \in R$, if $r \bullet s \sim_a r'' \bullet s$, then $r'' \bullet s \models_{\mathcal{M}} \phi$ iff $r \models_{\mathcal{M}} \mathbf{N}_a^s \phi$. \square

4 Modelling access control with the logic ERL*

In this section, we illustrate how to use ERL and its special sublogic ERL*, in modelling access control situations.

Security policies, such as those for access control, are often formulated separately from the architectural context in which they are intended to be applied. This can lead to the existence of vulnerabilities. Specifically, when a particular security policy is applied to a particular system, the security properties of the resulting system may not be as intended.

We aim to illustrate that the new operators \mathbf{L}_a^s and \mathbf{M}_a^s are appropriate for modelling situations where the access to resources (whether they are locations or pieces of data) is central. Indeed, both operators can be used to specify (in a slight different flavour) whether a resource verifies a property in agent's a perspective, granted that the local resource s is present.

Before developing our examples, we recall that there exists a body of work based on LL and multiset rewriting for modelling some access control problems in specific situations. For example, multiset rewriting has been used to characterize security protocols [7]. Our aim here, however, is to provide a more general framework that can be a modelling tool in many situations rather than be an ad hoc creation specific to a context. Even if such a framework based on LL and modalities for authorization and knowledge exists [22], we consider the differences between LL and BBI that make the later a more convenient tool for modelling. Both are able to model aspects of the properties of resources, but in LL propositions represent resources while in BBI (and, indeed, in BI) propositions represent properties of resources that can be expressed within the Kripke structures supporting resource semantics. LL focuses on the production and consumption—essentially counting—of resources while BBI focuses on separation and sharing of properties on resources. Modal extensions of BBI extend this view to incorporate the production and consumption of resources via the effects of actions in action modalities [13, 15].

Because—as explained in the introduction and in a substantial body of literature [34]—the semantics of BBI can be interpreted as being a theory of resources and their properties, we can directly use resources as tokens in our modelling of systems [8]. Of particular note in this paper is the use of *local* resources. For example, s in $r \models \mathbf{L}_a^s$ is of the same nature but does not have the same role, as the ambient resource r . This allows a simple integration of new actors of a system into a modelling using ERL and avoids the creation of new formal elements of a more ad hoc nature.

4.1 Modelling distributed systems

The construction of mathematical models always involves design choices. Our approach is guided by the approach to modelling distributed systems articulated in [1, 9]. This approach builds upon the

observation that, from a slightly abstract yet convenient point of view, the key structural components of a distributed systems are the following:

- *Locations*. The basic architecture of the system is considered to be described by a collection of connected places. Mathematically, we need some topological structure, with directed graphs being perhaps the most commonly useful set-up.
- *Resources*. Resources are situated at the locations identified in the system's architecture. They are the components of the system that are manipulated—i.e., consumed, created, moved and so on—as the system evolves in order to deliver the services that it is intended to provide. Mathematically, we take the ‘resource monoids’ adopted in, e.g., the semantics of BI, in Separation Logic and, indeed, in ERL. In the intuitionistic versions of these logics, we take a partially ordered (or sometimes preordered) partial monoid of resources. As we have seen in Section 1, the monoidal composition then captures the combination of resource elements and the ordering captures the comparison of resource elements. In the classical versions, we drop the ordering and work just with combination.
- *Processes*. The services that a system provides are delivered by the execution of processes, during which resources are manipulated. Mathematically, in formal generality, we can describe processes using an algebraic calculus of processes. In [8], we have employed a variation of Milner's basic system, SCCS [28], adapted to capture the interaction with resources and locations.

In addition, we require the following concept:

- *Environment*. When a system is modelled, it is necessary to decide what is its boundary. Things that are outside of the boundary are not represented in detail within the model. Nevertheless, the model must interact with its environment. Mathematically, this can be represented stochastically, using specified probability distributions to capture events at the boundary.

The structural components collectively represent the state of a system and can be used to define a process algebra with an operational semantics that defines their co-evolution as actions occur [1, 8, 9]:

$$L, R, E \xrightarrow{a} L', R', E'.$$

When building models in this style, it is necessary to set up a notion of *signature* for a model. For basic actions a and locations L , we define an evolution

$$\mu(a, L, R) = (L', R')$$

that specifies the effect of a on the resource R at this L . We call μ a *modification function*.

In this setting, there is an associated modal logic with a satisfaction relation of the form

$$L, R, E \models \phi,$$

which includes both additive and multiplicative action modalities [1, 8, 9]. Additive action modalities yield formulae of the form $[a]\phi$, with a truth condition along the following lines:

$$L, R, E \models [a]\phi \text{ iff for all } E \xrightarrow{a} E', L', R', E' \models \phi,$$

where we need the condition, part of the signature of the model, to the effect that the occurrence of the action a causes the evolution of L to L' and R to R' [1, 8, 9]. The multiplicative modalities allow actions to carry around local resources that can be combined with the ambient resource—so

we consider $L, R, E \models [a]_S \phi$ and form $R' \circ S'$ in the definiens of the satisfaction clause—to enable the evolution [1, 8, 9].

The logic is used both to constrain the model, through situation-specific logical properties, and to express desired or undesired properties of the system that are to be checked.

In the setting of modelling access control using ERL, locations, resources and processes can all be represented, although we can make some simplifications.

- *Locations.* The examples we consider implicitly employ location architectures, but they are sufficiently simple that they can also be handled implicitly in the formalization, often through the treatment of resources.
- *Resources.* The resource elements considered carry the structure of resource monoids, and we make essential use of this in the models.
- *Processes.* Our examples only deal with the actions that are required to the instantiation of epistemic modalities. Nevertheless, we provide discussions of how our examples can be understood in the location–resource–process context.

In this setting, we elide the modelling of environment: since we are not seeking to build executable models, this simplification is of little or no consequence for our present purposes. In these senses, we are making use of a fairly pure version of resource semantics.

We employ a range of examples of security modelling using this approach. We begin, in Section 4.2, with ‘Schneier’s gate’, which illustrates the policy–architecture gap, and then consider a core systems–security situations of joint access control, in Section 4.3, and semaphores, in Section 4.4.

4.2 The ‘Schneier’s gate’ problem

Consider the example of ‘Schneier’s gate’ [38], wherein a security system is ineffective because of the existence of a side-channel that allows a control to be circumvented. Here a facility that is intended to be secured is protected by a barrier that prevents cars from entering into the facility. The barrier may be controlled by a token—such as a card, a remote or a code—the holding of which distinguishes authorized personnel from intruders. If, however, the barrier itself is surrounded by ground that can be traversed by a vehicle, without any kind of fence or wall, then any car can drive around it (whether it is with malicious intent or just by laziness of getting through the security procedure) and the access control policy, as implemented by the barrier and the tokens, is undermined. So, the access control policy—that only authorized personnel, in possession of a token, may take vehicles into the facility—is undermined by the architecture of the system to which it is applied.

We show how ERL* can be used to model, and so reason about, the situation described above (following [38]), illustrating how such situations can be identified by logical analysis. Related analyses, employing logical models of layered graphs, can be found in [11].

We follow the approach to distributed systems modelling sketched in Section 4.1 and elaborated in [1, 8, 9]. We start with a simple model, depicted in Figure 2, and gradually refine it. We model just a facility protected by an access barrier. We will need the following key components:

- *Locations.* We assume, for what is an architecturally simple model, just three locations: *outside* and *inside* of the area guarded by the barrier and the barrier itself. In this simple setting, there is no need to incorporate an explicit representation of locations into our model’s worlds.
- *Resources.* There are just three types of resources: vehicles (cars), access tokens, which are required to operate the barrier, and a marker for the presence of the barrier.

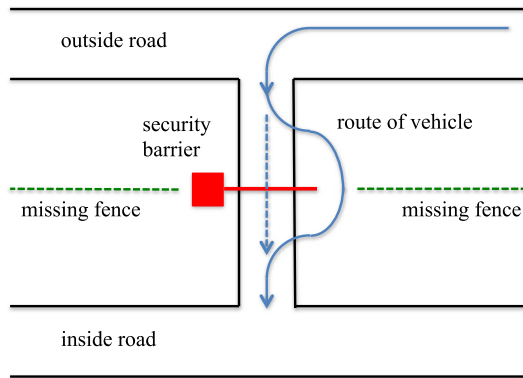


FIGURE 1 A depiction of the ‘Schneier’s gate’ problem.




FIGURE 2 Barrier problem, base case

- *Processes.* In this simple setting, we do not need to employ the full, quite complex, structure of a process algebra; rather, the actions of a logic with action modalities—in particular, the action modalities of ERL*, with their epistemic semantics—will suffice.

In fact, our treatment of resource in this epistemic-logic setting is a little more subtle. From the modelling perspective, the resources we have exposed here are diverse in nature: there is a material token (key or card for instance), there are cars, and a just a marker for the presence and well-functioning of the barrier. This diversity raises the question of the meaning and value of the unit resource, e . We finesse this problem by accepting that resources encompass a variety of different objects, but we can also employ the epistemic nature of our logic and consider that resources represent not objects as such but rather the knowledge that a given object is in our system.

A vehicle having the appropriate access token should be able to get inside. We consider the following sets of resources, agents, and logical properties of resources/system states:

$$Res = \{e, b, t, c\}, A = \{\alpha\}, Prop = \{O, J\}.$$

Here we have the following:

- the atomic propositions O and J , respectively, express the state of being *outside* and *inside* the facility—we use J instead of I to avoid confusion with I , the unit operator;
- a resource element b is taken as a marker for the presence and well-functioning of the barrier;
- a token, required to operate the barrier, is denoted by a resource element t and vehicles (cars) are denoted by resource elements c, c' , etc.;

- for simplicity we are assuming that all resource elements are of the same sort; i.e., are elements of the same resource monoid; this will cause no formal difficulty in this simple setting, though richer examples might require more care in this respect;
- $u \models_{\mathcal{M}} O$ means that u is outside the facility, and $v \models_{\mathcal{M}} J$ means that v is inside;
- the agent α is a generic one that represents a user of the system; i.e., say, the vehicle/driver that approaches the access control point. The resources b and t represent tokens that stand respectively for the barrier and the access token of the users.

So, c can be viewed as an abstract token marking the presence of a car and t the presence of the required access device in this car. Thus, resources act as an abstraction layer of our system. In this view, it follows that it is easy to see e as the absence of information (nothing is known of the system).

We have the following property: $O \rightarrow \mathbf{L}_{\alpha}^{bt}J$. According to the semantics, based on a resource monoid R , $c \models_{\mathcal{M}} O \rightarrow \mathbf{L}_{\alpha}^{bt}J$ just in case if $c \models_{\mathcal{M}} O$, then, for every $c' \in R$ such that $c \bullet b \bullet t \sim_{\alpha} c'$, $c' \models_{\mathcal{M}} J$. Thus, the combination of the two tokens grants access to the inside. The use of the token b for the presence of the barrier helps in modelling a situation in which the barrier is completely shut or is broken (in which case entering would not be possible). Note that the formulae $O \rightarrow \mathbf{L}_{\alpha}^tJ$, $O \rightarrow \mathbf{L}_{\alpha}^bJ$ and $O \rightarrow \mathbf{L}_{\alpha}^eJ$ are not valid because we cannot enter if the barrier is shut, if we have no access token or both.

The use of the operator \mathbf{L}_{α}^s in this situation is illustrative. First, consider what differences the use of other operators would make. If we were to state $O \rightarrow \widetilde{\mathbf{M}}_{\alpha}^{bt}J$, then it would mean that anyone outside can get (without condition) inside and acquire the two access tokens. This is of course not what we expect. On the other hand, using \mathbf{N}_{α}^s has an interesting effect. $O \rightarrow \mathbf{N}_{\alpha}^{bt}J$ requires not only that an entering agent have the expected tokens but also that those tokens remain active once they are inside. This is slightly different from our first approach: we do not know if the tokens are still active once the agent is inside.

We can also consider which of the additive implication, \rightarrow , and the multiplicative, $\rightarrow*$, would be the better modelling choice in this example. For the first approach, \rightarrow seems quite sufficient. Indeed, if we assert $O \rightarrow \mathbf{L}_{\alpha}^{bt}J$ as valid, then any resource satisfies it. So, if we have a car c such that $c \models_{\mathcal{M}} O$, we also have $c \models_{\mathcal{M}} O \rightarrow \mathbf{L}_{\alpha}^{bt}J$, and then we get the expected $c \models_{\mathcal{M}} \mathbf{L}_{\alpha}^{bt}J$.

However, if we consider more complex properties, the situation is different. Imagine, e.g., an environment that is composed not only of the car c but also another entity or piece of information, o . Our epistemic context is thus $o \bullet c$. If we have $c \models_{\mathcal{M}} O$ and if $O \rightarrow \mathbf{L}_{\alpha}^{bt}J$ is valid, then we get $c \models_{\mathcal{M}} \mathbf{L}_{\alpha}^{bt}J$. As we do not have $o \bullet c \models_{\mathcal{M}} O$, we cannot deduce that $o \bullet c \models_{\mathcal{M}} \mathbf{L}_{\alpha}^{bt}J$.

If instead we assume that the property $O \rightarrow* \mathbf{L}_{\alpha}^{bt}J$ is valid, then we have, in particular, $o \models_{\mathcal{M}} O \rightarrow* \mathbf{L}_{\alpha}^{bt}J$ and, together with $c \models_{\mathcal{M}} O$, we can deduce $o \bullet c \models_{\mathcal{M}} \mathbf{L}_{\alpha}^{bt}J$, as desired. So, the use of $\rightarrow*$ instead of \rightarrow is much more useful in more complex systems; as it allows us to set aside, as with Separation Logic's Frame Rule, some of the entities of our system and still apply the property.

Now we introduce agents to the model (see Figure 3). The first model may seem crude, because a single resource is used to model the access of any agent. So, we seek to benefit from the logic that allows us to take agents into account.

We change the model by defining a detailed set of agents, $A = \{\alpha, \beta, \gamma\}$ and now take three agents or users, α , β and γ . Each user should have its own access token, and the resource set is modified accordingly: $Res = \{e, b, t_{\alpha}, t_{\beta}, t_{\gamma}, c\}$. Now the slightly different formula $O \rightarrow \mathbf{L}_{\alpha}^{bt_{\alpha}}J$ is valid for any agent $a \in A$. So, e.g., $O \rightarrow \mathbf{L}_{\alpha}^{bt_{\alpha}}J$ is valid, which means that α can get inside with his own token, but $O \rightarrow \mathbf{L}_{\alpha}^{bt_{\beta}}J$ is not, which means α cannot use β 's token.

Now consider the case in which the access is controlled and the agents are supposed to cross the barrier only if they have the appropriate access device. We want to capture the fact that the




FIGURE 3 Barrier problem with agents.




FIGURE 4 Barrier problem with a shortcut.




FIGURE 5 Barrier problem with a fence.

system can actually be flawed (as mentioned in the problem presentation). It is actually quite easy to do, because being able to circumvent the barrier just means being able to access the inside of the complex without any token. We could be a little more specific by imagining that some agents know the shortcut (or dare to use it) and others do not (see Figure 4). In the previous setting, suppose that the agent β is aware of the shortcut and is disposed to use it. Our new set of properties should now be the following:

$$\left\{ O \rightarrow \mathbf{L}_a^{bt_a} J \text{ (for every } a \in A), O \rightarrow \mathbf{L}_\beta^e J \right\}.$$

The unit resource e expresses a direct access (with no resource needed). Note how the use of agents can help us to express different security policies in the same model.

We can reasonably suppose that such a flawed system would be quickly dealt with, e.g., by installing a fence that would prevent going around the barrier (see Figure 5). We could, of course, just model that by removing our last addition and get back to the intended policy, but it is more interesting to encode it by a formula. For example, we might then also describe a fault in the fence (or its removal). To do so, we can simply add a propositional formula F that is valid for any resource provided there is a fence preventing the passage of ‘rogue’ agents. Our system then becomes

$$\left\{ O \rightarrow \mathbf{L}_a^{bt_a} J \text{ (for every } a \in A), O \wedge \neg F \rightarrow \mathbf{L}_\beta^e J \right\}.$$

Having established a system of formulae that describes our modelling situation quite clearly, we can seek some properties of the model. The idea is to establish a property of the system that goes

beyond its basic definition. For example, we may want to check that every agent inside the facility has passed the barrier and has in its possession its access token. This means that we must prove that, for every agent $a \in A$, $J \rightarrow \mathbf{M}_a^{bt_a} J$.

Indeed, if $c \models_{\mathcal{M}} J \rightarrow \mathbf{M}_a^{bt_a} J$, this means that if $c \models_{\mathcal{M}} J$, then there exists $c' \in R$ such that $c \sim_a c' \bullet b \bullet t_a$ and $c' \bullet b \bullet t_a \models_{\mathcal{M}} J$, which expresses that every resource representing a car that is inside must in fact be equivalent, for an agent $a \in A$, to a resource that is inside *and* is composed with both the appropriate token t_a and the barrier token b . This is exactly what we wanted to capture.

Notice that this particular property is not verified by the system we described in our set-up. Indeed, noted previously, specifying entrance with $r \models_{\mathcal{M}} O \rightarrow \mathbf{L}_a^{bt_a} J$ makes J be satisfied by any resource r' such that $r \bullet b \bullet t_a \sim_a r'$. We can see that r' does not contain b and t_a . The use of $\mathbf{N}_a^{bt_a}$ instead solves this problem: we then have $r \bullet b \bullet t_a \sim_a r' \bullet b \bullet t_a$ and $r' \bullet b \bullet t_a \models_{\mathcal{M}} J$, as required.

So far, we have considered only simple situations, mainly one car crossing the barrier in various situations. Of course, we may wish to consider more complex models and establish similar properties. For example, we may want to see what happens if several cars are modelled together in the system.

We have the sets of properties in the form of implications stated before. To state there is a car in the system, we just assert that the formula O is valid. Then, by looking at the semantics of our formulae, we create a resource c which satisfies that formula. In order to have several cars, we might at first be tempted to assert something like $O \wedge O \wedge O$ (for three cars). However, given our semantics, we have trivially that $O \wedge O \wedge O \equiv O$, which is inconvenient for our modelling purpose. It is better to state $O * O * O$, using the multiplicative conjunction, instead. Then, to satisfy this formula, we need indeed three resources c_1, c_2, c_3 and we have $c_1 \bullet c_2 \bullet c_3 \models_{\mathcal{M}} O * O * O$ —i.e., for each car to gain access, a token is required for that car. Then, using $*$ as described above, we can see the system evolve as cars are allowed inside. Thus, the use of $*$ is particularly relevant to model several instances of a same object.

Of course, we could easily enrich this model to make more distinctions between different cars and their different properties, but the essentials of the model would remain the same.

4.3 Joint access

One of the most common problems of access control is joint access, and we propose to model a very simple example with our logic. The background for this example can be found in many films about the cold war era: the situation is that a critical system—such as one that controls the release of nuclear weapons, as in ‘Crimson Tide’ [5]—is secured by two different keys, each one held by a different operator. For the system to unlock, it is necessary that both operators activate their keys simultaneously. We provide a logical analysis of this situation.

From our systems modelling perspective, we can set this up quite simply, as depicted in Figure 6.

Some of the modelling choices made here are quite obvious: we need two agents and two associated resources representing their keys. So, we take $A = \{\alpha, \beta\}$ and $Res = \{k_1, k_2, e\}$. Implicitly, the formulae will express that α is associated to k_1 and β to k_2 . Also implicitly, we are employing four locations, $l_1 - l_4$, so that we can sketch a system model as

$$\begin{array}{lcl}
 l_1, k_1, \alpha : \text{Unlock}_1 : 0 & \xrightarrow{\alpha} & l_3, k_1, \text{Unlock}_1 : 0 \\
 l_2, k_2, \beta : \text{Unlock}_2 : 0 & \xrightarrow{\beta} & l_3, k_2, \text{Unlock}_2 : 0 \\
 l_3, k_1 \bullet k_2, \underbrace{\text{Unlock}_1 : 0 \times \text{Unlock}_2 : 0}_{\stackrel{\text{def}}{=} \text{Unlock}} & \xrightarrow{\alpha\beta} & l_4, k_1 \bullet k_2, 0,
 \end{array}$$




FIGURE 6 Joint access.

where $l_3 \bullet l_3 \stackrel{\text{def}}{=} l_3$, and where the modification function of the model, which describes how the keys move from location to location, is given by

- $\mu(\alpha, l_3, k_1 \bullet k_2) = (l_4, k_1 \bullet k_2)$,
- $\mu(\alpha, l_1, k_1) = (l_3, k_1)$ and
- $\mu(\alpha, l_2, k_2) = (l_3, k_2)$.

Focussing on our logical modelling, and suppressing for now the location architecture, we must express the fact that each agent—representing here a simplified notion of process—must use its key. Of course, as the whole point of the example is to illustrate how two separate accesses unlock the system, thus each use of key must be modelled with a different formula. We propose the following formulae for this purpose:

$$\mathbf{M}_\alpha^{k_1} \top \text{ and } \mathbf{M}_\beta^{k_2} \top.$$

We use the atomic formula \top since we do not need to access any property—rather we need only to update α and β 's accessible worlds to express that k_1 and k_2 are now activated. If we consider $\mathbf{M}_\alpha^{k_1} \top$ for instance, then if $r \models_{\mathcal{M}} \mathbf{M}_\alpha^{k_1} \top$, then there exists a resource r' such that $r \sim_\alpha r' \bullet k_1$ and $r' \bullet k_1 \models_{\mathcal{M}} \top$. Given this last statement, we have that there exists r' such that $r \sim_\alpha r' \bullet k_1$. Thus, with this formula we have stated that α can reach a state in which k_1 is activated. The second formula states the same for β and k_2 .

We must express that whenever both keys are present, the system can be unlocked. We could consider using a formula such as $\tilde{\mathbf{M}}_\alpha^{k_1 k_2} U$, where U is an atomic formula expressing that the system is unlocked. However, we can see at once that this choice is problematic. Indeed, this formula is dependent on α , but the point of joint access is that none of the agents involved is responsible on its own for the activation of the device. Moreover, should we decide to proceed with such a formula, it would fail to do the required job— k_2 is brought in the system by β and only α is present in the formula. Obviously, using β instead of α raises the same problems (symmetrically).

It seems, therefore, that our model lacks (at least) an agent. We introduce an omnipotent agent o (and thus $A = \{\alpha, \beta, o\}$). The idea is to have an agent that can see and use whatever α and β can, without the two sharing knowledge or potential action. This agent can be interpreted either as a global authority or just as a modelling of the device itself (the computer that accepts the keys and executes the order). Now, with this extra agent, $\tilde{\mathbf{M}}_o^{k_1 k_2} U$ seems to be an acceptable candidate for modelling the unlocking of the system. This states that whichever state reachable for o that contains k_1 and k_2 triggers the unlocking. However, we still need to express o 's capability. To do that, we introduce the following set of formulae:

$$\{\mathbf{M}_a^s \phi \rightarrow \mathbf{M}_o^s \phi \mid a \in A, s \in Res, \phi \in \mathcal{L}\}.$$

This expresses that any access to a resource by an agent through the modality \mathbf{M} can be transferred to o . Of course, in a more general setting, we could state similar things for the other operators, but, in this very particular example, only \mathbf{M} will be useful.

Finally, in order to the system to work, we need to activate both keys simultaneously. The first approach could be to append the two key-activation with an \wedge : $\mathbf{M}_\alpha^{k_1} \top \wedge \mathbf{M}_\beta^{k_2} \top$. This does not produce the desired result. Indeed, if $r \models_{\mathcal{M}} \mathbf{M}_\alpha^{k_1} \top \wedge \mathbf{M}_\beta^{k_2} \top$, then we get $r \sim_\alpha r' \bullet k_1$ and $r \sim_\beta r'' \bullet k_2$ and we intended to have the combination of k_1 and k_2 , which is not obvious here. Thus, the best way is in fact to use $\mathbf{M}_\alpha^{k_1} \top * \mathbf{M}_\beta^{k_2} \top$. More than the simple correctness of our modelling, this use of $*$ is quite convincing, as we aimed to model the separated use of two keys.

Thus, we have modelled our situation as follows:

1. $\forall ag \in A, \forall s \in Res, \forall \phi \in \mathcal{L}, \mathbf{M}_{ag}^s \phi \rightarrow \mathbf{M}_o^s \phi$;
2. $\mathbf{M}_\alpha^{k_1} \top * \mathbf{M}_\beta^{k_2} \top$;
3. $\tilde{\mathbf{M}}_o^{k_1 k_2} U$.

We can check that this has the desired effect, i.e., that whenever both keys are present, the system can be unlocked. Consider a resource r that forces (2) and (3). The forcing of (3), unpacked, means

$$\text{for all } r' \text{ such that } r \sim_o r' \bullet k_1 \bullet k_2, r' \bullet k_1 \bullet k_2 \models_{\mathcal{M}} U.$$

On the other side, unpacking of (2) gives

$$\text{there exist } r_1, r_2 \text{ such that } r = r_1 \bullet r_2 \text{ and } r_1 \models_{\mathcal{M}} \mathbf{M}_\alpha^{k_1} \top \text{ and } r_2 \models_{\mathcal{M}} \mathbf{M}_\beta^{k_2} \top.$$

We can then instantiate (1) twice, with $ag = \alpha, s = k_1$ and $\phi = \top$, then with $ag = \beta, s = k_2$ and $\phi = \top$ to get

$$\text{there exist } r_1, r_2 \text{ such that } r = r_1 \bullet r_2 \text{ and } r_1 \models_{\mathcal{M}} \mathbf{M}_\alpha^{k_1} \top \text{ and } r_2 \models_{\mathcal{M}} \mathbf{M}_\beta^{k_2} \top.$$

Unpacking this, we get

$$\text{there exist } r_1, r_2, r'_1, r'_2 \text{ such that } r = r_1 \bullet r_2 \text{ and } r_1 \sim_o r'_1 \bullet k_1 \text{ and } r_2 \sim_o r'_2 \bullet k_2.$$

By the compatibility of \bullet and \sim , we obtain that $r \sim_o r'_1 \bullet k_1 \bullet r_2$ and then that $r \sim_o r'_1 \bullet k_1 \bullet r'_2 \bullet k_2$, which by commutativity is $r \sim_o r'_1 \bullet r'_2 \bullet k_1 \bullet k_2$. Then we have $r'_1 \bullet r'_2 \bullet k_1 \bullet k_2 \models_{\mathcal{M}} U$, as required.

4.4 Semaphores

Another important example of modelling in access control is concerned with concurrency in parallel programming. We have described in the introduction how Separation Logic, built on BI, is a powerful and efficient tool to model memory management. We propose, in this section, an example of a similar work with ERL* in which we use it to model programs accessing memory and the particular example of simple concurrency with semaphores.

First, we establish the general basis of our modelling approach. We consider a multi-processor (or a set of different systems) which is seeking to run multiple programs or tasks with a limited amount of memory space.

- The set R of resources will represent the memory of the system, Res being a subset of the memory specified for each problem. e always denotes an empty set of information in the memory. Thus, in this example, we again suppress location, conflating it with resource.

- The set of agents A represents all the different threads or processes which are running the tasks.
- Two parts, m and m' , of the memory are linked by the relationship \sim_α if the access to m is equivalent to the access to m' for the process α .
- Finally, we use propositions of ERL* to model programs run by the thread. Thus, when we write $m \models_{\mathcal{M}} P$, we mean that the memory stored in m is used to run the program P .

Just as in the example of joint access, we can set up our modelling of semaphores in the context of our general approach to systems modelling. We suppress the details here, preferring to use the simplified approach afforded by the logical tools introduced in this paper, but see [15] for examples of similar models that more closely follow the system modelling approach.

So, consider how to model semaphores in this context. Recall that semaphores are simple bits of program which use flags or tokens to ensure that a specific portion of program, called *critical section*, is always accessed by at most one process. We use an arbitrary set of agents A , and the set of resources $Res = \{e, t\}$, where t is a token marking the entrance into the critical section. We also have two propositions C and NC , the former being the critical section of code, the latter being all the non-critical part of the code. Note that, here, the agents correspond to processes.

We consider the following formulae, which constrain the model, for any arbitrary process $\alpha \in A$:

1. *Guard*: for any $\alpha', \alpha'' \in A$ s.t. $\alpha' \neq \alpha''$, $\tilde{\mathbf{L}}_{\alpha'}^t \top \rightarrow \neg \tilde{\mathbf{L}}_{\alpha''}^t \top$;
2. *In*: $NC \rightarrow \mathbf{L}_{\alpha}^t C$;
3. *Out*: $C \rightarrow ((\neg \mathbf{M}_{\alpha}^t \top) \wedge \mathbf{M}_{\alpha}^e NC)$.

The *Guard* formulae, true for any two different processes α' and α'' , ensure that two processes cannot enter a critical section together. Indeed, if, for any *Guard* formula, we have that $m \models_{\mathcal{M}} \text{Guard}$, then, if there is m' such that $m \bullet t \sim_{\alpha'} m'$, there is no m'' such that $m \bullet t \sim_{\alpha''} m''$. That is, for any process p' which has the token t in memory, no other process p'' can get the token.

The *In* formula specifies that the process α enters the critical section. If we have that $m \models_{\mathcal{M}} \text{In}$, then, if $m \models_{\mathcal{M}} NC$, then, for any m' such that $m \bullet t \sim_{\alpha} m'$, we have that $m' \models_{\mathcal{M}} C$. That is, if a process is running the non-critical section, the addition of the token t gives it access to a memory state sufficient to run the critical section.

Symmetrically, the *Out* formula expresses the exit of p from a critical section. If $m \models_{\mathcal{M}} \text{Out}$, then, if $m \models_{\mathcal{M}} C$, then, for all m' such that $m \sim_{\alpha} m' \bullet t$, $m' \bullet t \not\models_{\mathcal{M}} \top$. That is, there is no m' such that $m \sim_{\alpha} m' \bullet t$. This allows us to delete t from the memory accessible by α . The second part of the formula, $\mathbf{M}_{\alpha}^e NC$, states that there is a state m'' such that $m \sim_{\alpha} m''$ and $m'' \models_{\mathcal{M}} NC$; i.e., α gets back into non-critical section.

No memory state that satisfies NC after C has been executed can have t in it. So, once this formula is taken into account, either p can continue to execute C or go into NC and release the token t . We can now see whether the guard we proposed is sufficient to ensure us that no two processes can get the critical section together. We do that in a simple way, by introducing the (new) formula $NC * NC$. If we have $m \models_{\mathcal{M}} NC * NC$, then we have $m = m_1 \bullet m_2$, with $m_1 \models_{\mathcal{M}} NC$ and $m_2 \models_{\mathcal{M}} NC$. This is a fair representation of two processes running the non-critical section in parallel, each one using a different part of the memory (cf. the treatment of concurrent composition in [1, 9] and in Concurrent Separation Logic [32]).

Now consider a process α_1 and suppose it has access to the token; i.e., there exists m'_1 such that $m_1 \bullet t \sim_{\alpha_1} m'_1$. If *In* is valid, then we have in particular that $m_1 \models_{\mathcal{M}} \text{In}$ and thus we have $m'_1 \models_{\mathcal{M}} C$. Now, α_1 is executing the critical section with m'_1 . Could another process α_2 access the critical section with m_2 ? The guard should avoid it. Indeed, if *Guard* is valid, then we have $m \models_{\mathcal{M}} \text{Guard}$. Yet, we

have established that $m_1 \bullet t \sim_{\alpha_1} m'_1$. We also have that $m = m_1 \bullet m_2$ and, by right composition, we have $m_1 \bullet m_2 \bullet t \sim_{\alpha_1} m'_1 \bullet m_2$; thus, $m \bullet t \sim_{\alpha_1} m'_1 \bullet m_2$. By applying $m \models_{\mathcal{M}} \text{Guard}$ with $\alpha' = \alpha_1$ and $\alpha'' = \alpha_2$, we have that there is no m' such that $m \bullet t \sim_{\alpha_2} m'$. Now, if α_2 were to access the critical section with m_2 , then we should have m'_2 such that $m_2 \bullet t \sim_{\alpha_2} m'_2$. Then we should have that $m \bullet t \sim_{\alpha_2} m'_2 \bullet m_1$ which would contradict what we stated before. Thus, α_2 cannot enter the critical section.

However, once in this situation, as we have $m'_1 \models_{\mathcal{M}} C$, we can use *Out* to let α_1 out of the critical section. As $m'_1 \models_{\mathcal{M}} \text{Out}$, we generate $m'_1 \models_{\mathcal{M}} \neg \mathbf{M}'_{\alpha_1} \top$ and $m'_1 \models_{\mathcal{M}} \mathbf{M}'_{\alpha_1} NC$. The first tells us that there is no m' such that $m'_1 \sim_{\alpha_1} m' \bullet t$. But, in our premiss, we have that $m'_1 \sim_{\alpha_1} m_1 \bullet t$. Those two facts are contradictory. Thus, if we want to use this formula, we have to delete the relation $m'_1 \sim_{\alpha_1} m_1 \bullet t$. This guarantees that t is no longer in α_1 's grasp. The second part, $m'_1 \models_{\mathcal{M}} \mathbf{M}'_{\alpha_1} NC$, gives us a new memory state m''_1 such that $m'_1 \sim_{\alpha_1} m''_1$ and $m''_1 \models_{\mathcal{M}} NC$. Thus, α_1 is back in non-critical state. Note that once $m'_1 \sim_{\alpha_1} m_1 \bullet t$ is deleted, the guard ceases to be applicable, and nothing prevents α_2 from entering the critical section this time.

4.5 Evolution in LL, BI and ERL

It is perhaps worthwhile pausing at this point to compare the representation of system evolution that is available here with that which is available in LL. First, we should note that the nature of the system model employed here is quite different from that which would derive from a representation based on LL. Second, in our setting, as we have explained, we employ a truth-functional instantiation of the general distributed systems modelling approach based on concepts of location, resource and process. In the examples of this paper, the account of process is very limited, being restricted to the actions of epistemic agents (with no rich process-theoretic structure). Third, as a result of these design choices, the readily available account of evolution requires unpacking the truth-functional semantics, which can be seen in terms of tableaux proofs (as presented in Section 5). Experience from, e.g., Separation Logic [37] suggests that the presence (as in BBI and ERL and ERL*) of a negation with the standard classical semantics is a very useful modelling tool.

In contrast, representations using LL's sequent calculus, such as the logic programming approach described in [2, 26], employ a less rich modelling perspective—restricted to proofs of sequences of resource manipulations—but then give a very direct operational reading of evolution in this restricted setting. A proof-theoretic treatment of some underlying ideas in LL may be found in [6]. Note, however, that BI includes MILL as a fragment (as we have seen) and that the basic propositional systems for BI can be presented as sequent calculi with well-understood relationships with LL. Within the multiplicative fragment of BI, the same readings of resource evolution can, of course, be obtained—we do not consider it worthwhile to rehearse these readings in the context of our examples, which are intended to illustrate resource semantics. We conjecture, therefore, that it is possible to give (perhaps labelled) sequent calculi for ERL and ERL* that would provide a similar operational reading of evolution (see the remarks at the beginning of Section 5) to that which is available in LL or the multiplicative fragment of BI.

To set up a precise correspondence between these evolutions and the semantic representation of resource is an interesting issue.

A brief comparison with 'epistemic linear logic' [22]—which is about modelling access control in LL—is perhaps also worthwhile. Again, this work benefits from the syntactic structures of LL as basis for representing evolution in the setting of the restricted model of systems that is naturally treated syntactically by LL. Again, in contrast, we begin from a more comprehensive systems semantics—which accommodates a very general notion of resource, including ambient

system resources and resources that are local to agents—and treat similar examples in this restricted instance. Again, we might expect sequent calculi for ERL and ERL* to capture a similar treatment of evolution to that provided by LL.

5 A tableaux calculus for ERL

In this section, we provide a labelled calculus for ERL in the spirit of the calculi previously developed for BI [21] and BBI [27] that are based on labels and label constraints allowing the capture of the semantics of these logics inside the corresponding calculus. In the case of BBI, a specific completeness proof, based on an oracle, has been developed in [27].

Similar labelled calculi have been proposed also for some modal and epistemic extensions of BI and BBI [13–15]. In these cases, the calculus design, used for BBI, is applied with specific labels and constraints issued from a semantic analysis of the considered logic. In the case of the labelled calculus for ESL [14], which is an epistemic extension of BBI, we deal with constraints that are parametrized by agents but do not handle the presence of resources in the scope of the modal operators (the local resources).

While herein we provide a tableaux calculus in the continuation of previous works on modal bunched logics, we note also that we could design a labelled sequent calculus for ERL and ERL* that would also be used to provide an operational reading of evolution through proof construction as in some LL fragments. However, our aim in this section is only to provide, by applying an approach and some proof methods already developed for other modal bunched logics, a labelled tableaux calculus for our logic—both in order to establish its metatheory and as a general reasoning tool.

For the present work, we must introduce labels that correspond to the local resources embedded in operators. As we shall see, we do that through a subset Λ_r of labels that is in bijection with the set of local resources Res . Similar techniques have been used with the logic LSM [15], which extends BBI with resource-parametrized S4 modalities. Likewise, the proofs of soundness and completeness of the calculus with respect to the semantics introduced in Section 2 are similar to the ones for ESL, mainly addressing the need to take the set Λ_r into account. Revisiting the remarks in Section 2 about the possibility of working with a hybrid semantics and then relating ERL to a hybrid version of ESL, we remark that the design of a hybrid tableau calculus would require some specific work about using nominals and formulas to replace labels and constraints—and this replacement introduces more complexity and undermines the strong links with the resource semantics that is central in our approach.

First, we introduce labels and constraints that correspond, respectively, to resources and to the equality and equivalence relations on resources and agents. Next, we develop labelled tableaux for ERL. Then, we establish soundness with respect to the resource semantics, giving the details of the proof in the appendix. Finally, we consider countermodel extraction and completeness, again giving the details of the proof in the appendix.

5.1 Labels and constraints

We consider a finite set of constants Λ_r such that $|\Lambda_r| = |Res| - 1$. On it we build an infinite countable set of (resource) constants γ_r such that $\Lambda_r \subset \gamma_r$, and then $\gamma_r = \Lambda_r \cup \{c_1, c_2, \dots\}$. Concatenation of lists is denoted by \oplus ; \square denotes the empty list. A *resource label* is a word built on γ_r , where the order of letters is not taken into account, i.e., a finite multiset γ_r and by ϵ the empty word. For example, xy is the composition of the resource labels x and y . We say that x is a *resource*

Rules for resource constraints:

$$\frac{}{\varepsilon \simeq \varepsilon} \langle \varepsilon \rangle \quad \frac{x \simeq y}{y \simeq x} \langle s_r \rangle \quad \frac{xy \simeq xy}{x \simeq x} \langle d_r \rangle \quad \frac{x \simeq y \quad y \simeq z}{x \simeq z} \langle t_r \rangle$$

$$\frac{x \simeq y \quad yk \simeq yk}{xk \simeq yk} \langle c_r \rangle \quad \frac{x \simeq_u y}{x \simeq x} \langle k_r \rangle$$

Rules for agent constraints:

$$\frac{x \simeq x}{x \simeq_v x} \langle r_a \rangle \quad \frac{x \simeq_u y}{y \simeq_u x} \langle s_a \rangle \quad \frac{x \simeq_u y \quad y \simeq_u z}{x \simeq_u z} \langle t_a \rangle \quad \frac{x \simeq_u y \quad x \simeq k}{k \simeq_u y} \langle k_a \rangle$$

FIGURE 7 Rules for constraint closure (for any $u \in A$).

sublabel of y if and only if there exists z such that $xz = y$. The set of resource sublabels of x is denoted $\mathcal{E}(x)$.

We define a function $\lambda : Res \rightarrow A_r$ such that

1. $\lambda(\varepsilon) = \varepsilon$;
2. for all $r \in Res \setminus \{\varepsilon\}$, $\lambda(r) \in A_r$; and
3. λ is injective. $r = r'$.

Note that λ is trivially a bijection between Res and $A_r \cup \{\varepsilon\}$.

DEFINITION 5 (Constraints).

A *resource constraint* is an expression of the form $x \simeq y$, where x and y are resource labels. An *agent constraint* is an expression of the form $x \simeq_u y$, where x and y are resource labels and u belongs to the set of agents A .

A *set of constraints* is any set \mathcal{C} that contains resource constraints and agent constraints. Let \mathcal{C} be a set of constraints. The (resource) *domain* of \mathcal{C} is the set of all resource sublabels that appear in \mathcal{C} ; i.e.,

$$\mathcal{D}_r(\mathcal{C}) = \bigcup_{x \simeq y \in \mathcal{C}} (\mathcal{E}(x) \cup \mathcal{E}(y)) \cup \bigcup_{x \simeq_u y \in \mathcal{C}} (\mathcal{E}(x) \cup \mathcal{E}(y)).$$

Let \mathcal{C} be a set of constraints. The (resource) *alphabet* $\mathcal{A}_r(\mathcal{C})$ of \mathcal{C} is the set of resource constants that appear in \mathcal{C} . In particular, $\mathcal{A}_r(\mathcal{C}) = \gamma_r \cap \mathcal{D}_r(\mathcal{C})$. Now we introduce, in Figure 7, the rules for constraint closure that allow us to capture the properties of the models into the calculus.

DEFINITION 6 (Closure of constraints).

Let \mathcal{C} be a set of constraints. The closure of \mathcal{C} , denoted $\bar{\mathcal{C}}$, is the least relation closed under the rules of Figure 7 such that $\mathcal{C} \subseteq \bar{\mathcal{C}}$.

There are six rules ($\langle \varepsilon \rangle$, $\langle s_r \rangle$, $\langle d_r \rangle$, $\langle t_r \rangle$, $\langle c_r \rangle$ and $\langle k_r \rangle$) that produce resource constraints and four rules ($\langle r_a \rangle$, $\langle s_a \rangle$, $\langle t_a \rangle$ and $\langle k_a \rangle$) that produce agent constraints. We note that v , introduced in the rule $\langle r_a \rangle$, must belong to the set of agents A .

PROPOSITION 4

The following rules can be derived from the rules of constraint closure:

$$\frac{xk \simeq y}{x \simeq x} \langle pl \rangle \quad \frac{x \simeq yk}{y \simeq y} \langle pr \rangle \quad \frac{xk \equiv_u y}{x \simeq x} \langle ql \rangle \quad \frac{x \equiv_u yk}{y \simeq y} \langle qr \rangle$$

$$\frac{x \equiv_u y \quad x \simeq x' \quad y \simeq y'}{x' \equiv_u y'} \langle wa \rangle$$

COROLLARY 1

Let \mathcal{C} be a set of constraints and $u \in A$ be an agent.

1. $x \in \mathcal{D}_r(\overline{\mathcal{C}})$ iff $x \simeq x \in \overline{\mathcal{C}}$ iff $x \equiv_u x \in \overline{\mathcal{C}}$.
2. If $xy \in \mathcal{D}_r(\overline{\mathcal{C}})$, $x' \simeq x \in \overline{\mathcal{C}}$ and $y' \simeq y \in \overline{\mathcal{C}}$, then $xy \simeq x'y' \in \overline{\mathcal{C}}$.

PROPOSITION 5

Let \mathcal{C} be a set of constraints. We have $\mathcal{A}_r(\mathcal{C}) = \mathcal{A}_r(\overline{\mathcal{C}})$.

LEMMA 2 (Compactness).

Let \mathcal{C} be a (possibly infinite) set of constraints.

1. If $x \simeq y \in \overline{\mathcal{C}}$, then there is a finite set \mathcal{C}_f such that $\mathcal{C}_f \subseteq \mathcal{C}$ and $x \simeq y \in \overline{\mathcal{C}_f}$.
2. If $x \equiv_u y \in \overline{\mathcal{C}}$, then there is a finite set \mathcal{C}_f such that $\mathcal{C}_f \subseteq \mathcal{C}$ and $x \equiv_u y \in \overline{\mathcal{C}_f}$.

5.2 Labelled tableaux for ERL

We now define a labelled tableaux calculus for ERL in the spirit of previous works [14, 17, 21, 27] by using similar definitions and results but based on the specific label and constraints definitions.

DEFINITION 7

A *labelled formula* is a 3-tuple of the form $(S\phi : x)$ such that $S \in \{\mathbb{T}, \mathbb{F}\}$, $\phi \in \mathcal{L}$ is a formula and $x \in \Lambda_r$ is a resource label. A *constrained set of statements* (CSS) is a pair $\langle \mathcal{F}, \mathcal{C} \rangle$, where \mathcal{F} is a set of labelled formulae and \mathcal{C} is a set of constraints, satisfying the following property, denoted P_{css} ,

$$\text{if } (S\phi : x) \in \mathcal{F}, \text{ then } x \simeq x \in \overline{\mathcal{C}}(P_{css}).$$

A CSS $\langle \mathcal{F}, \mathcal{C} \rangle$ is *finite* if \mathcal{F} and \mathcal{C} are finite. The relation \preceq is defined by $\langle \mathcal{F}, \mathcal{C} \rangle \preceq \langle \mathcal{F}', \mathcal{C}' \rangle$ iff $\mathcal{F} \subseteq \mathcal{F}'$ and $\mathcal{C} \subseteq \mathcal{C}'$. We write $\langle \mathcal{F}_f, \mathcal{C}_f \rangle \preceq_f \langle \mathcal{F}, \mathcal{C} \rangle$ when $\langle \mathcal{F}_f, \mathcal{C}_f \rangle \preceq \langle \mathcal{F}, \mathcal{C} \rangle$ holds and $\langle \mathcal{F}_f, \mathcal{C}_f \rangle$ is finite, meaning that \mathcal{F}_f and \mathcal{C}_f are both finite.

PROPOSITION 6

For any CSS $\langle \mathcal{F}_f, \mathcal{C} \rangle$, where \mathcal{F}_f is finite, there exists $\mathcal{C}_f \subseteq \mathcal{C}$ such that \mathcal{C}_f is finite and $\langle \mathcal{F}_f, \mathcal{C}_f \rangle$ is a CSS.

PROOF. By induction on the number of labelled formulae of \mathcal{F}_f and by Lemma 2. \square

Figure 8 presents the rules of tableaux calculus for ERL. Note that ‘ c_i and c_j are new label constants’ means $c_i \neq c_j \in \gamma_r \setminus (\mathcal{A}_r(\mathcal{C}) \cup \Lambda_r)$.

DEFINITION 8 (Tableau for ERL).

Let $\langle \mathcal{F}_0, \mathcal{C}_0 \rangle$ be a finite CSS. A *tableau* for $\langle \mathcal{F}_0, \mathcal{C}_0 \rangle$ is a list of CSSs, called *branches*, inductively built according to the following rules:

$$\begin{array}{c}
\frac{(\mathbb{T}\mathbb{I} : x) \in \mathcal{F}}{\langle \emptyset, \{x \simeq \varepsilon\} \rangle} \langle \mathbb{T}\mathbb{I} \rangle \\
\\
\frac{(\mathbb{T}\neg\phi : x) \in \mathcal{F}}{\langle \{(\mathbb{F}\phi : x)\}, \emptyset \rangle} \langle \mathbb{T}\neg \rangle \quad \frac{(\mathbb{F}\neg\phi : x) \in \mathcal{F}}{\langle \{(\mathbb{T}\phi : x)\}, \emptyset \rangle} \langle \mathbb{F}\neg \rangle \\
\\
\frac{(\mathbb{T}\phi \wedge \psi : x) \in \mathcal{F}}{\langle \{(\mathbb{T}\phi : x), (\mathbb{T}\psi : x)\}, \emptyset \rangle} \langle \mathbb{T}\wedge \rangle \quad \frac{(\mathbb{F}\phi \wedge \psi : x) \in \mathcal{F}}{\langle \{(\mathbb{F}\phi : x), \emptyset\} \mid \{(\mathbb{F}\psi : x)\}, \emptyset \rangle} \langle \mathbb{F}\wedge \rangle \\
\\
\frac{(\mathbb{T}\phi \vee \psi : x) \in \mathcal{F}}{\langle \{(\mathbb{T}\phi : x)\}, \emptyset \mid \{(\mathbb{T}\psi : x)\}, \emptyset \rangle} \langle \mathbb{T}\vee \rangle \quad \frac{(\mathbb{F}\phi \vee \psi : x) \in \mathcal{F}}{\langle \{(\mathbb{F}\phi : x), (\mathbb{F}\psi : x)\}, \emptyset \rangle} \langle \mathbb{F}\vee \rangle \\
\\
\frac{(\mathbb{T}\phi \rightarrow \psi : x) \in \mathcal{F}}{\langle \{(\mathbb{F}\phi : x)\}, \emptyset \mid \{(\mathbb{T}\psi : x)\}, \emptyset \rangle} \langle \mathbb{T}\rightarrow \rangle \quad \frac{(\mathbb{F}\phi \rightarrow \psi : x) \in \mathcal{F}}{\langle \{(\mathbb{T}\phi : x), (\mathbb{F}\psi : x)\}, \emptyset \rangle} \langle \mathbb{F}\rightarrow \rangle \\
\\
\frac{(\mathbb{T}\phi * \psi : x) \in \mathcal{F}}{\langle \{(\mathbb{T}\phi : c_i), (\mathbb{T}\psi : c_j)\}, \{x \simeq c_i c_j\} \rangle} \langle \mathbb{T}* \rangle \quad \frac{(\mathbb{F}\phi * \psi : x) \in \mathcal{F} \text{ and } x \simeq yz \in \bar{C}}{\langle \{(\mathbb{F}\phi : y)\}, \emptyset \mid \{(\mathbb{F}\psi : z)\}, \emptyset \rangle} \langle \mathbb{F}* \rangle \\
\\
\frac{(\mathbb{T}\phi -* \psi : x) \in \mathcal{F} \text{ and } xy \simeq xy \in \bar{C}}{\langle \{(\mathbb{F}\phi : y)\}, \emptyset \mid \{(\mathbb{T}\psi : xy)\}, \emptyset \rangle} \langle \mathbb{T}-* \rangle \quad \frac{(\mathbb{F}\phi -* \psi : x) \in \mathcal{F}}{\langle \{(\mathbb{T}\phi : c_i), (\mathbb{F}\psi : xc_i)\}, \{xc_i \simeq xc_i\} \rangle} \langle \mathbb{F}-* \rangle \\
\\
\frac{(\mathbb{T}\mathbb{L}'_u \phi : x) \in \mathcal{F} \text{ and } x\lambda(r) =_u y \in \bar{C}}{\langle \{(\mathbb{T}\phi : y)\}, \emptyset \rangle} \langle \mathbb{T}\mathbb{L}' \rangle \quad \frac{(\mathbb{F}\mathbb{L}'_u \phi : x) \in \mathcal{F}}{\langle \{(\mathbb{F}\phi : c_i)\}, \{x\lambda(r) =_u c_i\} \rangle} \langle \mathbb{F}\mathbb{L}' \rangle \\
\\
\frac{(\mathbb{T}\mathbb{M}'_u \phi : x) \in \mathcal{F}}{\langle \{(\mathbb{T}\phi : c_i \lambda(r))\}, \{x =_u c_i \lambda(r)\} \rangle} \langle \mathbb{T}\mathbb{M}' \rangle \quad \frac{(\mathbb{F}\mathbb{M}'_u \phi : x) \in \mathcal{F} \text{ and } x =_u y \lambda(r) \in \bar{C}}{\langle \{(\mathbb{F}\phi : y \lambda(r))\}, \emptyset \rangle} \langle \mathbb{F}\mathbb{M}' \rangle \\
\\
\frac{(\mathbb{T}\mathbb{N}'_u \phi : x) \in \mathcal{F} \text{ and } x\lambda(r) =_u y \lambda(r) \in \bar{C}}{\langle \{(\mathbb{T}\phi : y \lambda(r))\}, \emptyset \rangle} \langle \mathbb{T}\mathbb{N}' \rangle \quad \frac{(\mathbb{F}\mathbb{N}'_u \phi : x) \in \mathcal{F}}{\langle \{(\mathbb{F}\phi : c_i \lambda(r))\}, \{x \lambda(r) =_u c_i \lambda(r)\} \rangle} \langle \mathbb{F}\mathbb{N}' \rangle \\
\\
\frac{(\mathbb{T}\tilde{\mathbb{L}}'_u \phi : x) \in \mathcal{F}}{\langle \{(\mathbb{T}\phi : c_i)\}, \{x \lambda(r) =_u c_i\} \rangle} \langle \mathbb{T}\tilde{\mathbb{L}}' \rangle \quad \frac{(\mathbb{F}\tilde{\mathbb{L}}'_u \phi : x) \in \mathcal{F} \text{ and } x\lambda(r) =_u y \in \bar{C}}{\langle \{(\mathbb{F}\phi : y)\}, \emptyset \rangle} \langle \mathbb{F}\tilde{\mathbb{L}}' \rangle \\
\\
\frac{(\mathbb{T}\tilde{\mathbb{M}}'_u \phi : x) \in \mathcal{F} \text{ and } x =_u y \lambda(r) \in \bar{C}}{\langle \{(\mathbb{T}\phi : y \lambda(r))\}, \emptyset \rangle} \langle \mathbb{T}\tilde{\mathbb{M}}' \rangle \quad \frac{(\mathbb{F}\tilde{\mathbb{M}}'_u \phi : x) \in \mathcal{F}}{\langle \{(\mathbb{F}\phi : c_i \lambda(r))\}, \{x =_u c_i \lambda(r)\} \rangle} \langle \mathbb{F}\tilde{\mathbb{M}}' \rangle \\
\\
\frac{(\mathbb{T}\tilde{\mathbb{N}}'_u \phi : x) \in \mathcal{F}}{\langle \{(\mathbb{T}\phi : c_i \lambda(r))\}, \{x \lambda(r) =_u c_i \lambda(r)\} \rangle} \langle \mathbb{T}\tilde{\mathbb{N}}' \rangle \quad \frac{(\mathbb{F}\tilde{\mathbb{N}}'_u \phi : x) \in \mathcal{F} \text{ and } x\lambda(r) =_u y \lambda(r) \in \bar{C}}{\langle \{(\mathbb{F}\phi : y \lambda(r))\}, \emptyset \rangle} \langle \mathbb{F}\tilde{\mathbb{N}}' \rangle
\end{array}$$

Note: c_i and c_j are new label constants, with $c_i, c_j \notin \Lambda_r$.

FIGURE 8 Rules of the tableaux calculus for ERL.

1. The one branch list $[\langle \mathcal{F}_0, \mathcal{C}_0 \rangle]$ is a tableau for $\langle \mathcal{F}_0, \mathcal{C}_0 \rangle$;
2. If the list $\mathcal{T}_m \oplus [\langle \mathcal{F}, \mathcal{C} \rangle] \oplus \mathcal{T}_n$ is a tableau for $\langle \mathcal{F}_0, \mathcal{C}_0 \rangle$ and

$$\frac{\text{cond} \langle \mathcal{F}, \mathcal{C} \rangle}{\langle \mathcal{F}_1, \mathcal{C}_1 \rangle \mid \dots \mid \langle \mathcal{F}_k, \mathcal{C}_k \rangle}$$

is an instance of a rule of Figure 8 for which $\text{cond}\langle\mathcal{F},\mathcal{C}\rangle$ is fulfilled, then the list $\mathcal{T}_m \oplus [(\mathcal{F} \cup \mathcal{F}_1, \mathcal{C} \cup \mathcal{C}_1); \dots; (\mathcal{F} \cup \mathcal{F}_k, \mathcal{C} \cup \mathcal{C}_k)] \oplus \mathcal{T}_n$ is a tableau for $\langle\mathcal{F}_0, \mathcal{C}_0\rangle$.

A *tableau* for the formula ϕ is a *tableau* for $\{(\mathbb{F}\phi : c_1)\}, \{c_1 \simeq c_1\}$.

We remark that a tableau for a formula ϕ verifies the property (P_{CSS}) of Definition 7 (by the rule $\langle r_a \rangle$) and any application of a rule of Figure 8 provides also a tableau that verifies the property (P_{CSS}) (in particular, by Corollary 1).

In this calculus, we have two particular set of rules. The first set is composed of the rules $\langle \mathbb{T}\mathbb{I} \rangle$, $\langle \mathbb{T}\ast \rangle$, $\langle \mathbb{F}\ast \rangle$, $\langle \mathbb{F}\mathbb{L} \rangle$, $\langle \mathbb{F}\tilde{\mathbb{M}} \rangle$, $\langle \mathbb{F}\mathbb{N} \rangle$, $\langle \mathbb{T}\tilde{\mathbb{L}} \rangle$, $\langle \mathbb{T}\mathbb{M} \rangle$ and $\langle \mathbb{T}\tilde{\mathbb{N}} \rangle$, that introduce new label constants (c_i and c_j) and new constraints, except for $\langle \mathbb{T}\mathbb{I} \rangle$ that only introduces a new constraint. The second set is composed of the rules $\langle \mathbb{F}\ast \rangle$, $\langle \mathbb{T}\ast \rangle$, $\langle \mathbb{T}\mathbb{L} \rangle$, $\langle \mathbb{T}\tilde{\mathbb{M}} \rangle$, $\langle \mathbb{T}\mathbb{N} \rangle$, $\langle \mathbb{F}\tilde{\mathbb{L}} \rangle$, $\langle \mathbb{F}\mathbb{M} \rangle$ and $\langle \mathbb{F}\tilde{\mathbb{N}} \rangle$, that have a condition on the closure of constraints. To apply one of these rules we choose a label which satisfies the condition and then apply the corresponding rule. Otherwise, we cannot apply the rule.

DEFINITION 9 (Closure conditions).

A CSS $\langle\mathcal{F}, \mathcal{C}\rangle$ is *closed* if one of the following conditions holds, where $\phi \in \mathcal{L}$:

1. $\langle \mathbb{T}\phi : x \rangle \in \mathcal{F}$, $\langle \mathbb{F}\phi : y \rangle \in \mathcal{F}$ and $x \simeq y \in \bar{\mathcal{C}}$;
2. $\langle \mathbb{F}\mathbb{I} : x \rangle \in \mathcal{F}$ and $x \simeq \epsilon \in \bar{\mathcal{C}}$;
3. $\langle \mathbb{F}\mathbb{T} : x \rangle \in \mathcal{F}$;
4. $\langle \mathbb{T}\perp : x \rangle \in \mathcal{F}$.

A CSS is *open* if it is not closed. A tableau for ϕ is *closed* if all its branches (i.e., all of its CSSs) are closed and a *tableaux proof* for ϕ is a closed tableau for ϕ .

Closed branches are marked with \times and open branches are marked with \circ .

Example. Let us consider the formula $\tilde{\mathbb{M}}_a^s \phi \rightarrow \tilde{\mathbb{M}}_a^r (\tilde{\mathbb{M}}_a^s \phi)$. To build the corresponding tableau, we start with the CCS $\{(\mathbb{F}\tilde{\mathbb{M}}_a^s \phi \rightarrow \tilde{\mathbb{M}}_a^r (\tilde{\mathbb{M}}_a^s \phi) : c_1)\}, \{c_1 \simeq c_1\}$ and with the following representation of the formula set \mathcal{F} and the constraints set \mathcal{C} :

$$\begin{array}{c} [\mathcal{F}] \\ \sqrt{1} (\mathbb{F}\tilde{\mathbb{M}}_a^s \phi \rightarrow \tilde{\mathbb{M}}_a^r (\tilde{\mathbb{M}}_a^s \phi) : c_1) \end{array} \quad \begin{array}{c} [\mathcal{C}] \\ c_1 \simeq c_1 \end{array}$$

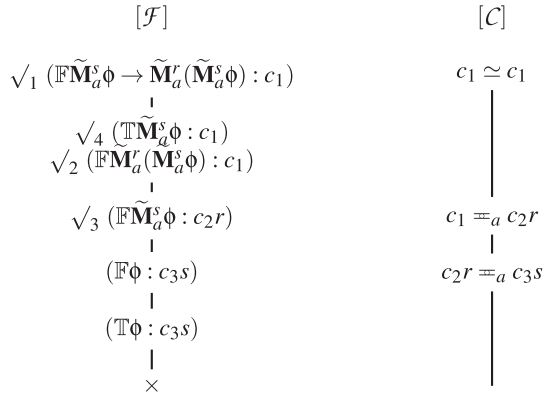
We then apply the rules of our tableaux method, respecting the priority order, and we obtain the tableau of Figure 9. We omit the λ and write r for $\lambda(r)$, for any resource.

Note that we mark with $\sqrt{\quad}$ the steps of the tableau construction. The main steps are the following: first apply the rule $\langle \mathbb{F}\rightarrow \rangle$ ($\sqrt{1}$) and then obtain two formulae both with $\tilde{\mathbb{M}}$ as operator. According to the priority rules, first apply the $\langle \mathbb{F}\tilde{\mathbb{M}} \rangle$ rule ($\sqrt{2}$), which generates a new formula, a new resource label c_2 and the constraint $c_1 \simeq_a c_2 r$. Then apply the $\langle \mathbb{F}\tilde{\mathbb{M}} \rangle$ rule again ($\sqrt{3}$), which generates a new formula, a new resource label c_3 and the constraint $c_2 r \simeq_a c_3 s$. We must now apply the $\langle \mathbb{T}\tilde{\mathbb{M}} \rangle$ rule ($\sqrt{4}$) and then we need a resource label z such that $c_1 \simeq_a z s \in \bar{\mathcal{C}}$.

Now, having closure by rule $\langle t_a \rangle$ with agent a , we generate the constraint $c_1 \simeq_a c_3 s$, and thus apply the rule with $z = c_1$ and generate $\langle \mathbb{T}\phi : c_3 s \rangle$. As we also have $\langle \mathbb{F}\phi : c_3 s \rangle$, we have a closed branch and thus a closed tableau.

5.3 Soundness of the calculus

We start by proving the soundness property of the tableaux calculus. The proof is similar to the soundness proof developed for BI tableaux and some recent extensions [13, 14, 17, 21]. We remind here the key notions and more detailed proofs are given in Appendix A.


 FIGURE 9 Tableau for $\tilde{\mathbf{M}}_a^s\phi \rightarrow \tilde{\mathbf{M}}_a^r(\tilde{\mathbf{M}}_a^s\phi)$.

The main point is the notion of *realizability* of a CSS $\langle \mathcal{F}, \mathcal{C} \rangle$, meaning that there exists a model \mathcal{M} and an embedding $(|\cdot|)$ from the resource labels to the resource set of \mathcal{M} such that if $(\mathbb{T}\phi : x) \in \mathcal{F}$, then $|x| \models_{\mathcal{M}} \phi$, and if $(\mathbb{F}\phi : x) \in \mathcal{F}$, then $|x| \not\models_{\mathcal{M}} \phi$.

DEFINITION 10 (Realization).

Let $\langle \mathcal{F}, \mathcal{C} \rangle$ be a CSS. A *realization* of it is a pair $(\mathcal{M}, |\cdot|)$ where $\mathcal{M} = (\mathcal{R}, \{\sim_a\}_{a \in A}, V)$ is a model and $|\cdot| : \mathcal{D}_r(\mathcal{C}) \rightarrow \mathcal{R}$ such that

- for any $r \in \text{Res}$, we have $|\lambda(r)| = r$,
- $|\epsilon| = e$,
- $|\cdot|$ is a total function (for all $x \in \mathcal{D}_r(\mathcal{C})$, $|x|$ is defined),
- if $xy \in \mathcal{D}_r(\mathcal{C})$, then $|x| \bullet |y| \downarrow$ and $|x| \bullet |y| = |xy|$,
- if $(\mathbb{T}\phi : x) \in \mathcal{F}$, then $|x| \models_{\mathcal{M}} \phi$,
- if $(\mathbb{F}\phi : x) \in \mathcal{F}$, then $|x| \not\models_{\mathcal{M}} \phi$,
- if $x \simeq y \in \mathcal{C}$, then $|x| = |y|$, and
- if $x \equiv_u y \in \mathcal{C}$, then $|x| \sim_u |y|$.

We say that a CSS is *realizable* if there exists a realization of this CSS. We say that a tableau is *realizable* if at least one of its branches is realizable.

PROPOSITION 7

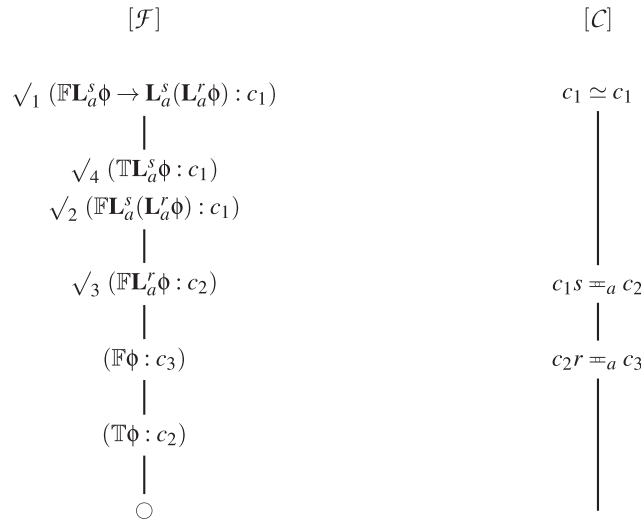
Let $\langle \mathcal{F}, \mathcal{C} \rangle$ be a CSS and $\mathcal{R} = (\mathcal{M}, |\cdot|)$ be a realization of it. \mathcal{R} is also a realization of $\langle \mathcal{F}, \bar{\mathcal{C}} \rangle$, and then

1. for all $x \in \mathcal{D}_r(\bar{\mathcal{C}})$, $|x|$ is defined,
2. if $x \simeq y \in \bar{\mathcal{C}}$, then $|x| = |y|$ and
3. if $x \equiv_u y \in \bar{\mathcal{C}}$, then $|x| \sim_u |y|$.

LEMMA 3

The rules of the tableaux method for ERL preserve realizability.

PROOF. By induction on the structure of realizable tableaux. See [15] for a similar argument and Appendix A for more details. \square

FIGURE 10 Tableau for $\mathbf{L}_a^s \phi \rightarrow \mathbf{L}_a^s(\mathbf{L}_a^r \phi)$.

LEMMA 4

Closed branches are not realizable.

PROOF. By a case analysis of closed branches that are realizable. See [15] for a similar argument and Appendix A for more details. \square

THEOREM 1 (Soundness).

Let ϕ be a formula of ERL. If there exists a tableaux proof for ϕ , then ϕ is valid.

PROOF. We suppose that there exists a proof for ϕ . Then there is a closed tableau \mathcal{T}_ϕ for the CSS $\mathcal{C} = \langle \{(\mathbb{F}\phi : c_1)\}, \{c_1 \simeq c_1\} \rangle$. Now suppose that ϕ is not valid. Then there is a countermodel $\mathcal{M} = (\mathcal{R}, \{\sim_a\}_{a \in A}, V)$ and a resource $r \in R$ such that $r \not\models_{\mathcal{M}} \phi$. Let $\mathfrak{R} = (\mathcal{M}, | \cdot |)$ such that $|c_1| = r$. As \mathfrak{R} is a realization of \mathcal{C} , by Lemma 3, \mathcal{T}_ϕ is realizable. Moreover, by Lemma 4, \mathcal{T}_ϕ cannot be closed, which is absurd because \mathcal{T}_ϕ is a proof and then is closed by definition. Therefore, ϕ is valid. \square

5.4 Countermodel generation and completeness of the calculus

Before proceeding to establish completeness, we consider a countermodel extraction method for our calculus that is adapted from a method proposed in [27].

Countermodel generation. The method transforms the sets of resource and agent constraints of a branch $\langle \mathcal{F}, \mathcal{C} \rangle$ into a model \mathcal{M} such that, if $(\mathbb{T}\phi : x) \in \mathcal{F}$, then $\rho_x \models_{\mathcal{M}} \phi$ and, if $(\mathbb{F}\phi : x) \in \mathcal{F}$, then $\rho_x \not\models_{\mathcal{M}} \phi$, where ρ_x is the representative of the equivalence class of x .

The method is based mainly on the definition on a particular CSS $\langle \mathcal{F}, \mathcal{C} \rangle$, called a *Hintikka* CSS. For more details, see Appendix B. This approach for countermodel extraction is proposed and illustrated for other bunched logics in [13–15, 17, 21] and adapted to our ERL logic.

Example. We give an example of countermodel extraction by considering $A = \{a\}$ and $Res = \{e, r\}$ and the formula $\mathbf{L}_a^s \phi \rightarrow \mathbf{L}_a^r \mathbf{L}_a^s \phi$, which is not valid. By applications of the tableaux rules, we obtain the tableau of Figure 10.


We see that, in step 4, we can only find c_2 as suitable label for $c_1 s \approx_a x$ and thus the tableau is not closed. The only branch of this tableau is a Hintikka CSS and we extract this countermodel using Definition 13.

We have $\mathcal{M} = (\mathcal{R}, \{\sim_a\}_{a \in A}, V)$, where

- $R = \text{Rep}(\mathcal{D}_r(\bar{\mathcal{C}})) \cup \text{Res} = \{e, r, s, \rho_{c_1}, \rho_{c_2}, \rho_{c_3}, \rho_{c_1\lambda(s)}, \rho_{c_2\lambda(r)}\}$
- The resource composition:

•	e	r	s	ρ_{c_1}	ρ_{c_2}	ρ_{c_3}	$\rho_{c_1\lambda(s)}$	$\rho_{c_2\lambda(r)}$
e	e	r	s	ρ_{c_1}	ρ_{c_2}	ρ_{c_3}	$\rho_{c_1\lambda(s)}$	$\rho_{c_2\lambda(r)}$
r	r	\uparrow	\uparrow	\uparrow	$\rho_{c_2\lambda(r)}$	\uparrow	\uparrow	\uparrow
s	s	\uparrow	\uparrow	$\rho_{c_1\lambda(s)}$	\uparrow	\uparrow	\uparrow	\uparrow
ρ_{c_1}	ρ_{c_1}	\uparrow	$\rho_{c_1\lambda(s)}$	\uparrow	\uparrow	\uparrow	\uparrow	\uparrow
ρ_{c_2}	ρ_{c_2}	$\rho_{c_2\lambda(r)}$	\uparrow	\uparrow	\uparrow	\uparrow	\uparrow	\uparrow
ρ_{c_3}	ρ_{c_3}	\uparrow	\uparrow	\uparrow	\uparrow	\uparrow	\uparrow	\uparrow
$\rho_{c_1\lambda(s)}$	$\rho_{c_1\lambda(s)}$	\uparrow	\uparrow	\uparrow	\uparrow	\uparrow	\uparrow	\uparrow
$\rho_{c_2\lambda(r)}$	$\rho_{c_2\lambda(r)}$	\uparrow	\uparrow	\uparrow	\uparrow	\uparrow	\uparrow	\uparrow

- The equivalence relation, reflexivity is not represented:



- $V(\phi) = \{\rho_{c_2}\}$.

We can easily verify that we have a countermodel of $\mathbf{L}_a^s \phi \rightarrow \mathbf{L}_a^r(\mathbf{L}_a^r \phi)$.

1. As $\rho_{c_2} \in V(\phi)$, we have $\rho_{c_2} \models \phi$.
2. As $\{x \in R \mid \rho_{c_1} \bullet s \sim_a x\} = \{\rho_{c_2}\}$, we have by (1), $\rho_{c_1} \models_{\mathcal{M}} \mathbf{L}_a^s \phi$.
3. As $\rho_{c_3} \notin V(\phi)$, we have $\rho_{c_3} \not\models \phi$.
4. As $\rho_{c_2} \bullet r = \rho_{c_2\lambda(r)} \sim_a \rho_{c_3}$, by (3), we have $\rho_{c_2} \not\models_{\mathcal{M}} \mathbf{L}_a^r \phi$.
5. As $\rho_{c_1} \bullet s = \rho_{c_1\lambda(s)} \sim_a \rho_{c_2}$, by (4), we have $\rho_{c_1} \not\models_{\mathcal{M}} \mathbf{L}_a^s(\mathbf{L}_a^r \phi)$.
6. By (2) and (5), we conclude that $\rho_{c_1} \not\models_{\mathcal{M}} \mathbf{L}_a^s \phi \rightarrow \mathbf{L}_a^r(\mathbf{L}_a^r \phi)$.

Completeness. The proof of completeness is an extension of the corresponding proof proposed for BBI [27] to the epistemic connectives of our logic. It consists in building, using a fair strategy, a Hintikka CSS from a formula for which there is no tableaux proof that is a sequence of labelled formulae in which all labelled formulae occur infinitely many times, and also an oracle that is a set of non-closed CSS with some specific properties. Then, assuming there is no tableaux proof for ϕ , we build a Hintikka CSS and deduce from it that ϕ is not valid.

THEOREM 2 (Completeness).

Let ϕ be an ERL formula. If ϕ is valid, then there exists a tableaux proof for ϕ .

PROOF. The proof is an extension of the corresponding proof proposed for BBI [27] to the epistemic connectives of our logic. More details are given in Appendix C. \square

To complete this section, we show how we can define a tableaux calculus for the sublogic ERL^* .

DEFINITION 11 (Tableaux for ERL^*).

The tableaux calculus for ERL^* is defined exactly as the tableaux calculus for ERL , with the addition of the following rule to Definition 6:

$$\frac{x \vDash_u y \quad yk \simeq yk}{xk \vDash_u yk} \langle c_a \rangle$$

PROPOSITION 8

The tableaux calculus for ERL^* is sound and complete with respect to the semantics given in Sections 2 and 3.

PROOF. The proof is the same as the one for ERL except that the new rule $\langle c_a \rangle$ must be considered each time the closure of constraints is concerned. This addition does not cause any difficulties with proofs since this rule is a direct translation of the specific property of ERL^* as described in Definition 4. \square

6 Conclusions

We have presented a substructural epistemic logic, based on BBI , in which the epistemic modalities, which extend the usual epistemic modalities, are parametrized on the agent's local resource. The logic represents the first step in developing an epistemic resource semantics. This step is illustrated through examples that explore the gap between policy and implementation in access control. We have also provided a system of labelled tableaux for the logic and established soundness and completeness.

Much further work is suggested. First, we might consider the theory, pragmatics, and interpretation of the epistemic modalities with resource semantics, including aspects of local reasoning for resource-carrying agents [25, 37], concurrency [32]. Second, we might consider logical theory, including proof systems, model-theoretic properties, and complexity. Connections with other approaches to modelling the relationship between policy and implementation in system management, such as those discussed in [39] and approaches involving logics for layered graphs [1, 10], should be explored.

A Soundness: proofs of lemmas

LEMMA 3

The rules of the tableaux method for ERL preserve realizability.

PROOF. By induction on the structure of realizable tableaux. See [15] for a similar argument. Let \mathcal{T} be a realizable tableau. By definition, \mathcal{T} has a realizable branch $\mathcal{B} = \langle \mathcal{F}, \mathcal{C} \rangle$. Let $\mathfrak{R} = (\mathcal{M}, |\cdot|)$ be a realization of the branch \mathcal{B} , where $\mathcal{M} = (\mathcal{R}, \{\sim_a\}_{a \in A}, V)$ and $|\cdot| : \mathcal{D}_r(\mathcal{C}) \rightarrow R$. If we apply a rule on a labelled formula of a branch that is not \mathcal{B} then \mathcal{B} is not modified, and then \mathcal{T} is realizable. Else, we consider each kind of formula on which the rule is applied.

- $(\text{TI} : x) \in \mathcal{F}$.

We have, by definition of realization, $|x| \Vdash_{\mathcal{M}} \text{I}$. Then $|x| = e$. As $|\epsilon| = e$ then $|x| = |\epsilon|$ and we remark that \mathfrak{R} is a realization of the new branch $\langle \mathcal{F}, \mathcal{C} \cup \{x \simeq \epsilon\} \rangle$.

- $(\text{T}\phi_1 * \phi_2 : x) \in \mathcal{F}$.

By realization, we have $|x| \Vdash_{\mathcal{M}} \phi_1 * \phi_2$. Then, by definition, there exist $r_1, r_2 \in R$ such that

$r_1 \bullet r_2 \downarrow$, $|x| = r_1 \bullet r_2$, $r_1 \models_{\mathcal{M}} \phi_1$ and $r_2 \models_{\mathcal{M}} \phi_2$. As c_i and c_j are new resource label constants, $|c_i|$ and $|c_j|$ are not defined. Moreover, as $c_i \neq c_j$, we can extend \mathfrak{R} by setting $|c_i| = r_1$ and $|c_j| = r_2$. As we have $|c_i| \bullet |c_j| \downarrow$ and, by implicit extension, $|x| = |c_i| \bullet |c_j| = |c_i c_j|$, we obtain a realization of $\langle \mathcal{F}, \mathcal{C} \cup \{x \simeq c_i c_j\} \rangle$, that is a realization of the branch $\langle \mathcal{F} \cup \{(\mathbb{T}c_i \cdot;)(\mathbb{T}\phi_2 : c_j)\}, \mathcal{C} \cup \{x \simeq c_i c_j\} \rangle$.

- $(\mathbb{F}\phi_1 * \phi_2 : x) \in \mathcal{F}$.

We have $|x| \not\models_{\mathcal{M}} \phi_1 * \phi_2$. By definition, for all $r_1, r_2 \in R$ such that $r_1 \bullet r_2 \downarrow$ and $|x| = r_1 \bullet r_2$, we have $r_1 \not\models_{\mathcal{M}} \phi$ or $r_2 \not\models_{\mathcal{M}} \psi$. The branch is expanded into two branches that are $\langle \mathcal{F} \cup \{(\mathbb{F}\phi : y)\}, \mathcal{C} \rangle$ and $\langle \mathcal{F} \cup \{(\mathbb{F}\psi : z)\}, \mathcal{C} \rangle$, where $x \simeq yz \in \bar{\mathcal{C}}$. By Proposition 7, $|x| = |yz|$. By definition of realization, $|\cdot|$ is total, then $|y| \bullet |z| \downarrow$ and $|yz| = |y| \bullet |z|$. Thus, $|y| \not\models_{\mathcal{M}} \phi$ or $|z| \not\models_{\mathcal{M}} \psi$. Therefore, \mathfrak{R} is a realization of at least one of the two new branches $\langle \mathcal{F} \cup \{(\mathbb{F}\phi : y)\}, \mathcal{C} \rangle$ or $\langle \mathcal{F} \cup \{(\mathbb{F}\psi : z)\}, \mathcal{C} \rangle$.

- $(\mathbb{T}\mathbf{L}'_u \phi : x) \in \mathcal{F}$ and $x\lambda(r) =_u y \in \bar{\mathcal{C}}$.

We have $|x| \models_{\mathcal{M}} \mathbf{L}'_u \phi$. By definition, for all $r' \in R$ such that $|x| \bullet r \sim_u r'$, we have $r' \models_{\mathcal{M}} \phi$. Moreover, as $x\lambda(r) =_u y \in \bar{\mathcal{C}}$, by Proposition 7, we have $|x\lambda(r)| \sim_u |y|$. By definition, $|x\lambda(r)| = |x| \bullet |\lambda(r)| = |x| \bullet r$. Thus, $|x| \bullet r \sim_u |y|$ and finally, we have $|y| \models_{\mathcal{M}} \phi$; thus, \mathcal{R} is a realization of the branch $\langle \mathcal{F} \cup \{(\mathbb{T}\phi : y)\}, \mathcal{C} \rangle$.

- $(\mathbb{F}\mathbf{L}'_u \phi : x) \in \mathcal{F}$.

We have $|x| \not\models_{\mathcal{M}} \mathbf{L}'_u \phi$. By definition, there exists $r' \in R$ such that $|x| \bullet r \sim_u r'$ and $r' \not\models_{\mathcal{M}} \phi$. As c_i is a new constraint, $|c_i|$ is not defined and we can choose $|c_i| = r'$ and we have $|c_i| \not\models_{\mathcal{M}} \phi$ and $|x| \bullet r \sim_u |c_i|$. By definition, $|x\lambda(r)| = |x| \bullet |\lambda(r)| = |x| \bullet r$. Thus, $|x\lambda(r)| \sim_u |c_i|$ and we have a realization of the branch $\langle \mathcal{F} \cup \{(\mathbb{F}\phi : c_i)\}, \mathcal{C} \cup \{x\lambda(r) =_u c_i\} \rangle$.

Other cases are proved similarly. □

LEMMA 4

Closed branches are not realizable.

PROOF. By a case analysis of closed branches that are realizable. See [15] for more details.

Let $\langle \mathcal{F}, \mathcal{C} \rangle$ a closed branch. We suppose that this branch is realizable. Let $\mathfrak{R} = (\mathcal{M}, |\cdot|)$ a realization of it. There are four cases:

- $(\mathbb{T}\phi : x) \in \mathcal{F}$, $(\mathbb{F}\phi : y) \in \mathcal{F}$ and $x \simeq y \in \bar{\mathcal{C}}$. By Proposition 7, as the branch is realizable, we must have $|x| \models_{\mathcal{M}} \phi$, $|y| \not\models_{\mathcal{M}} \phi$ and $|x| = |y|$, which is absurd.
- $(\mathbb{F}\mathbf{I} : x) \in \mathcal{F}$ and $x \simeq \epsilon \in \bar{\mathcal{C}}$. By Proposition 7, as the branch is realizable, we must have $|x| \not\models_{\mathcal{M}} \mathbf{I}$ and $|x| = |\epsilon|$. By Definition 3, we have $e \neq |x|$ and by Definition 10 we have $|x| = e$, which is absurd.
- $(\mathbb{F}\mathbf{T} : x) \in \mathcal{F}$. By Proposition 7, as the branch is realizable, we must have $|x| \not\models_{\mathcal{M}} \mathbf{T}$, which is absurd by Definition 3.
- $(\mathbb{T}\perp : x) \in \mathcal{F}$. By Proposition 7, as the branch is realizable, we must have $|x| \models_{\mathcal{M}} \perp$, which is absurd by Definition 3. □

As all cases are absurd, we conclude that $\langle \mathcal{F}, \mathcal{C} \rangle$ is not realizable.

B Countermodel extraction method

We propose a countermodel extraction method, first designed in [27] for BBI, that consists in transforming the sets of resource and agent constraints of a branch $\langle \mathcal{F}, \mathcal{C} \rangle$ into a model \mathcal{M} such that if $(\mathbb{T}\phi : x) \in \mathcal{F}$ then $\rho_x \models_{\mathcal{M}} \phi$ and if $(\mathbb{F}\phi : x) \in \mathcal{F}$ then $\rho_x \not\models_{\mathcal{M}} \phi$, where ρ_x is the representative of the equivalence class of x . First, we define when a CSS $\langle \mathcal{F}, \mathcal{C} \rangle$ is a *Hintikka CSS*.

DEFINITION 12 (Hintikka CSS).

A CSS $\langle \mathcal{F}, \mathcal{C} \rangle$ is a *Hintikka CSS* iff, for any formula $\phi, \psi \in \mathcal{L}$, any resource $r \in Res$, any resource label $x, y, z \in \Lambda_r$ and any agent $u \in A$:

1. $(\mathbb{T}\phi : x) \notin \mathcal{F}$ or $(\mathbb{F}\phi : y) \notin \mathcal{F}$ or $x \simeq y \notin \bar{\mathcal{C}}$
2. $(\mathbb{F}\mathbb{I} : x) \notin \mathcal{F}$ or $x \simeq \epsilon \notin \bar{\mathcal{C}}$
3. $(\mathbb{F}\mathbb{T} : x) \notin \mathcal{F}$
4. $(\mathbb{T}\perp : x) \notin \mathcal{F}$
5. If $(\mathbb{T}\mathbb{I} : x) \in \mathcal{F}$, then $x \simeq \epsilon \in \bar{\mathcal{C}}$
6. If $(\mathbb{T}\neg\phi : x) \in \mathcal{F}$, then $(\mathbb{F}\phi : x) \in \mathcal{F}$
7. If $(\mathbb{F}\neg\phi : x) \in \mathcal{F}$, then $(\mathbb{T}\phi : x) \in \mathcal{F}$
8. If $(\mathbb{T}\phi \wedge \psi : x) \in \mathcal{F}$, then $(\mathbb{T}\phi : x) \in \mathcal{F}$ and $(\mathbb{T}\psi : x) \in \mathcal{F}$
9. If $(\mathbb{F}\phi \wedge \psi : x) \in \mathcal{F}$, then $(\mathbb{F}\phi : x) \in \mathcal{F}$ or $(\mathbb{F}\psi : x) \in \mathcal{F}$
10. If $(\mathbb{T}\phi \vee \psi : x) \in \mathcal{F}$, then $(\mathbb{T}\phi : x) \in \mathcal{F}$ or $(\mathbb{T}\psi : x) \in \mathcal{F}$
11. If $(\mathbb{F}\phi \vee \psi : x) \in \mathcal{F}$, then $(\mathbb{F}\phi : x) \in \mathcal{F}$ and $(\mathbb{F}\psi : x) \in \mathcal{F}$
12. If $(\mathbb{T}\phi \rightarrow \psi : x) \in \mathcal{F}$, then $(\mathbb{F}\phi : x) \in \mathcal{F}$ or $(\mathbb{T}\psi : x) \in \mathcal{F}$
13. If $(\mathbb{F}\phi \rightarrow \psi : x) \in \mathcal{F}$, then $(\mathbb{T}\phi : x) \in \mathcal{F}$ and $(\mathbb{F}\psi : x) \in \mathcal{F}$
14. If $(\mathbb{T}\phi * \psi : x) \in \mathcal{F}$, then $\exists y, z \in \Lambda_r, x \simeq yz \in \bar{\mathcal{C}}$ and $(\mathbb{T}\phi : y) \in \mathcal{F}$ and $(\mathbb{T}\psi : z) \in \mathcal{F}$
15. If $(\mathbb{F}\phi * \psi : x) \in \mathcal{F}$, then $\forall y, z \in \Lambda_r, x \simeq yz \in \bar{\mathcal{C}}$ implies $(\mathbb{F}\phi : y) \in \mathcal{F}$ or $(\mathbb{F}\psi : z) \in \mathcal{F}$
16. If $(\mathbb{T}\phi * \psi : x) \in \mathcal{F}$, then $\forall y \in \Lambda_r, xy \in \mathcal{D}_r$ implies $(\mathbb{F}\phi : y) \in \mathcal{F}$ or $(\mathbb{T}\psi : xy) \in \mathcal{F}$
17. If $(\mathbb{F}\phi * \psi : x) \in \mathcal{F}$, then $\exists y \in \Lambda_r, xy \in \mathcal{D}_r$ and $(\mathbb{T}\phi : y) \in \mathcal{F}$ and $(\mathbb{F}\psi : xy) \in \mathcal{F}$
18. If $(\mathbb{T}\mathbb{L}_u^r \phi : x) \in \mathcal{F}$, then $\forall y \in \Lambda_r, x\lambda(r) =_u y \in \bar{\mathcal{C}}$ implies $(\mathbb{T}\phi : y) \in \mathcal{F}$
19. If $(\mathbb{F}\mathbb{L}_u^r \phi : x) \in \mathcal{F}$, then $\exists y \in \Lambda_r, x\lambda(r) =_u y \in \bar{\mathcal{C}}$ and $(\mathbb{F}\phi : y) \in \mathcal{F}$
20. If $(\mathbb{T}\mathbb{M}_u^r \phi : x) \in \mathcal{F}$, then there exists $y \in \Lambda_r, x =_u y\lambda(r) \in \bar{\mathcal{C}}$ and $(\mathbb{T}\phi : y\lambda(r)) \in \mathcal{F}$
21. If $(\mathbb{F}\mathbb{M}_u^r \phi : x) \in \mathcal{F}$, then for all $y \in \Lambda_r, x =_u y\lambda(r) \in \bar{\mathcal{C}}$ implies $(\mathbb{F}\phi : y\lambda(r)) \in \mathcal{F}$
22. If $(\mathbb{T}\mathbb{N}_u^r \phi : x) \in \mathcal{F}$, then for all $y \in \Lambda_r, x\lambda(r) =_u y\lambda(r) \in \bar{\mathcal{C}}$ implies $(\mathbb{T}\phi : y\lambda(r)) \in \mathcal{F}$
23. If $(\mathbb{F}\mathbb{N}_u^r \phi : x) \in \mathcal{F}$, then there exists $y \in \Lambda_r, x\lambda(r) =_u y\lambda(r) \in \bar{\mathcal{C}}$ and $(\mathbb{F}\phi : y\lambda(r)) \in \mathcal{F}$
24. If $(\mathbb{T}\tilde{\mathbb{L}}_u^r \phi : x) \in \mathcal{F}$, then there exists $y \in \Lambda_r, x\lambda(r) =_u y \in \bar{\mathcal{C}}$ and $(\mathbb{T}\phi : y) \in \mathcal{F}$
25. If $(\mathbb{F}\tilde{\mathbb{L}}_u^r \phi : x) \in \mathcal{F}$, then for all $y \in \Lambda_r, x\lambda(r) =_u y \in \bar{\mathcal{C}}$ implies $(\mathbb{F}\phi : y) \in \mathcal{F}$
26. If $(\mathbb{T}\tilde{\mathbb{M}}_u^r \phi : x) \in \mathcal{F}$, then $\forall y \in \Lambda_r, x =_u y\lambda(r) \in \bar{\mathcal{C}}$ implies $(\mathbb{T}\phi : y\lambda(r)) \in \mathcal{F}$
27. If $(\mathbb{F}\tilde{\mathbb{M}}_u^r \phi : x) \in \mathcal{F}$, then $\exists y \in \Lambda_r, x =_u y\lambda(r) \in \bar{\mathcal{C}}$ and $(\mathbb{F}\phi : y\lambda(r)) \in \mathcal{F}$
28. If $(\mathbb{T}\tilde{\mathbb{N}}_u^r \phi : x) \in \mathcal{F}$, then there exists $y \in \Lambda_r, x\lambda(r) =_u y\lambda(r) \in \bar{\mathcal{C}}$ and $(\mathbb{T}\phi : y\lambda(r)) \in \mathcal{F}$
29. If $(\mathbb{F}\tilde{\mathbb{N}}_u^r \phi : x) \in \mathcal{F}$, then for all $y \in \Lambda_r, x\lambda(r) =_u y\lambda(r) \in \bar{\mathcal{C}}$ implies $(\mathbb{F}\phi : y\lambda(r)) \in \mathcal{F}$.

Conditions 1–4 ensure that a Hintikka CSS is not closed and conditions 5–29 ensure that it is saturated (no new tableaux rule can be applied).

To extract countermodels, we must manipulate equivalence classes. The equivalence class of $x \in \mathcal{D}_r(\bar{\mathcal{C}})$, denoted $[x]$, is the set $[x] = \{y \in \Lambda_r \mid x \simeq y \in \bar{\mathcal{C}}\}$. Moreover, the function ρ that extracts a representative from a class is defined for any class $[x]$ by $\rho([x]) = r$ if $\exists r \in Res/\lambda(r) \in [x]$ and by $\rho([x]) = y$ with y an arbitrary element of $[x]$ otherwise. We note that $\rho_x = \rho([x])$ and that $Rep(\mathcal{D}_r(\bar{\mathcal{C}}))$, the set of all representatives of $\mathcal{D}_r(\bar{\mathcal{C}})$, is given by $Rep(\mathcal{D}_r(\bar{\mathcal{C}})) = \{\rho_x \mid x \in \mathcal{D}_r(\bar{\mathcal{C}})\}$.

LEMMA 5

For any set of constraints \mathcal{C} , we have $e \in Rep(\mathcal{D}_r(\bar{\mathcal{C}}))$ and $\rho_e = e$.

DEFINITION 13 (Function Ω).

Let $\langle \mathcal{F}, \mathcal{C} \rangle$ be a Hintikka CSS. The function Ω associates to $\langle \mathcal{F}, \mathcal{C} \rangle$ a 3-tuple $\Omega(\langle \mathcal{F}, \mathcal{C} \rangle) = (\mathcal{R}, \{\sim_a\}_{a \in A}, V)$, where $\mathcal{R} = (R, \bullet)$, such that

- $R = Rep(\mathcal{D}_r(\bar{\mathcal{C}})) \cup Res$,

- if $\alpha \notin \text{Rep}(\mathcal{D}_r(\bar{\mathcal{C}}))$ or $\beta \notin \text{Rep}(\mathcal{D}_r(\bar{\mathcal{C}}))$, then $\alpha \bullet \beta = \uparrow$, else, $\alpha = \rho_x$ and $\beta = \rho_y$, and we have

$$\rho_x \bullet \rho_y = \begin{cases} \uparrow & \text{if } xy \notin \mathcal{D}_r(\bar{\mathcal{C}}) \\ \rho_{xy} & \text{otherwise,} \end{cases}$$
- for all $a \in A$, $\alpha \sim_a \beta$ iff $\alpha = \rho_x$ and $\beta = \rho_y$ and $x \equiv_a y \in \bar{\mathcal{C}}$ and
- $\alpha \in V(p)$ iff $\alpha = \rho_x$ and there exists $y \in \Lambda_r$ such that $y \simeq x \in \bar{\mathcal{C}}$ and $(\mathbb{T}p : y) \in \mathcal{F}$.

LEMMA 6

Let $\langle \mathcal{F}, \mathcal{C} \rangle$ be a Hintikka CSS. $\Omega(\langle \mathcal{F}, \mathcal{C} \rangle)$ is a model.

LEMMA 7

Let $\langle \mathcal{F}, \mathcal{C} \rangle$ be a Hintikka CSS and $\mathcal{M} = \Omega(\langle \mathcal{F}, \mathcal{C} \rangle) = (\mathcal{R}, \{\sim_a\}_{a \in A}, V)$, where $\mathcal{R} = (R, \bullet)$. For any formula $\phi \in \mathcal{L}$, any agent $a \in A$ and any $x, y \in \mathcal{D}_r(\bar{\mathcal{C}})$, we have the following: (1) If $(\mathbb{F}\phi : x) \in \mathcal{F}$, then $\rho_x \not\models_{\mathcal{M}} \phi$; (2) If $(\mathbb{T}\phi : x) \in \mathcal{F}$, then $\rho_x \models_{\mathcal{M}} \phi$.

LEMMA 8

Let $\langle \mathcal{F}, \mathcal{C} \rangle$ be a Hintikka CSS such that $(\mathbb{F}\phi : x) \in \mathcal{F}$. The formula ϕ is not valid and $\Omega(\langle \mathcal{F}, \mathcal{C} \rangle)$ is a countermodel of ϕ .

PROOF. Let $\langle \mathcal{F}, \mathcal{C} \rangle$ be a Hintikka CSS such that $(\mathbb{F}\phi : x) \in \mathcal{F}$. Let $\mathcal{K} = \Omega(\langle \mathcal{F}, \mathcal{C} \rangle)$. By Lemma 6, \mathcal{K} is a model. As $\langle \mathcal{F}, \mathcal{C} \rangle$ is a CSS, then by (P_{css}) and Corollary 2, $x \in \mathcal{D}_r(\bar{\mathcal{C}})$. Thus, by Lemma 7, we have $\rho_x \not\models_{\mathcal{M}} \phi$. Therefore, \mathcal{K} is a countermodel of the formula ϕ and we can conclude that ϕ is not valid. \square

C Proof of completeness

This proof is an extension of the proof for BBI [27] to the epistemic connectives of our logic. It consists in identifying two things: first, a Hintikka CSS, using a fair strategy, from a formula for which there is no tableaux proof, i.e., a sequence of labelled formulae in which all labelled formulae occur infinitely many times; second, an oracle, i.e., a set of non-closed CSSs with some specific properties.

DEFINITION 14 (Fair strategy).

A *fair strategy* is a sequence of labelled formulae and agent constraints $(S_i)_{i \in \mathbb{N}}$ in $(\{\mathbb{T}, \mathbb{F}\} \times \mathcal{L} \times \Lambda_r) \cup (\Lambda_r \times A \times \Lambda_r)$ such that all labelled formulae and all agent constraints occur infinitely many times in this sequence, i.e., $\{i \in \mathbb{N} \mid S_i \equiv (\mathbb{S}F : x)\}$ and $\{i \in \mathbb{N} \mid S_i \equiv x\lambda(r) \equiv_u y\}$ are infinite, for any $(\mathbb{S}F : x) \in \{\mathbb{T}, \mathbb{F}\} \times \mathcal{L} \times \Lambda_r$ and any $x\lambda(r) \equiv_u y \in \Lambda_r \times A \times \Lambda_r$.

PROPOSITION 9

There exists a fair strategy.

PROOF. Let $X = (\{\mathbb{T}, \mathbb{F}\} \times \mathcal{L} \times \Lambda_r) \cup (\Lambda_r \times A \times \Lambda_r)$. As *Prop* is countable then \mathcal{L} is countable. Moreover, Λ_r is countable (remember that γ_r is countable). Therefore, X is countable. So $\mathbb{N} \times X$ is countable and there exists a surjective function $\varphi : \mathbb{N} \rightarrow \mathbb{N} \times X$. Let $p : \mathbb{N} \times X \rightarrow X$ defined by $p(i, x) = x$ and $u = p \circ \varphi$. We show that u is a fair strategy by showing that for any $x \in X$, $u^{-1}(\{x\})$ is infinite. Let $x \in X$. $u^{-1}(\{x\}) = \varphi^{-1}(p^{-1}(\{x\}))$. But $p^{-1}(\{x\}) = \{(i, x) \mid i \in \mathbb{N}\}$ so $p^{-1}(x)$ is infinite. As φ is surjective $\varphi^{-1}(p^{-1}(\{x\}))$ is also infinite. \square

DEFINITION 15

Let \wp be a set of CSS.

1. \wp is \preceq -closed if $\langle \mathcal{F}, \mathcal{C} \rangle \in \wp$ holds whenever $\langle \mathcal{F}, \mathcal{C} \rangle \preceq \langle \mathcal{F}', \mathcal{C}' \rangle$ and $\langle \mathcal{F}', \mathcal{C}' \rangle \in \wp$ holds.
2. \wp is of finite character if $\langle \mathcal{F}, \mathcal{C} \rangle \in \wp$ holds whenever $\langle \mathcal{F}_f, \mathcal{C}_f \rangle \in \wp$ holds for every $\langle \mathcal{F}_f, \mathcal{C}_f \rangle \preceq_f \langle \mathcal{F}, \mathcal{C} \rangle$.
3. \wp is *saturated* if, for any $\langle \mathcal{F}, \mathcal{C} \rangle \in \wp$ and any instance

$$\frac{\text{cond}(\mathcal{F}, \mathcal{C})}{\langle \mathcal{F}_1, \mathcal{C}_1 \rangle \mid \dots \mid \langle \mathcal{F}_k, \mathcal{C}_k \rangle}$$

of a rule of Figure 8, if $\text{cond}(\mathcal{F}, \mathcal{C})$ is fulfilled, then $\langle \mathcal{F} \cup \mathcal{F}_i, \mathcal{C} \cup \mathcal{C}_i \rangle \in \wp$ for at least one $i \in \{1, \dots, k\}$.

DEFINITION 16 (Oracle).

An *oracle* is a set of non-closed CSSs that is \preceq -closed, of finite character and saturated.

LEMMA 9

There exists an oracle which contains every finite CSS for which there exists no closed tableau.

PROOF. The proof is an adaptation for our epistemic modalities of the corresponding proof schema in [13, 27]. The proof given in [13] provides the necessary notions to develop this proof in detail. \square

To prove completeness, we consider a formula φ for which there exists no proof and we show that there exists a countermodel for this formula.

The proof depends on finding a way to obtain a Hintikka CSS. By Lemma 9, there exists an oracle which contains every finite CSS for which there exists no closed tableau. We denote by \wp this oracle. By Proposition 9, there exists a fair strategy. We denote by \mathcal{S} this strategy and \mathcal{S}_i the i th formula or agent constraint of \mathcal{S} . As \mathcal{T}_0 cannot be closed then its unique branch belongs to the oracle, i.e., $\langle \{(\mathbb{F}\varphi : c_1)\}, \{c_1 \simeq c_1\} \rangle \in \wp$.

We build a sequence $\langle \mathcal{F}_i, \mathcal{C}_i \rangle_{i \geq 0}$ whose limit is a Hintikka CSS, as follows:

- $\langle \mathcal{F}_0, \mathcal{C}_0 \rangle = \langle \{(\mathbb{F}\varphi : c_1)\}, \{c_1 \simeq c_1\} \rangle$;
- \mathcal{S}_i is a labelled formula of the form $(\mathbb{S}F : x)$:
 - If $\langle \mathcal{F}_i \cup \{(\mathbb{S}F : x)\}, \mathcal{C}_i \rangle \notin \wp$, then $\langle \mathcal{F}_{i+1}, \mathcal{C}_{i+1} \rangle = \langle \mathcal{F}_i, \mathcal{C}_i \rangle$;
 - If $\langle \mathcal{F}_i \cup \{(\mathbb{S}F : x)\}, \mathcal{C}_i \rangle \in \wp$, then $\langle \mathcal{F}_{i+1}, \mathcal{C}_{i+1} \rangle = \langle \mathcal{F}_i \cup \{(\mathbb{S}F : x)\} \cup F_e, \mathcal{C}_i \cup \mathcal{C}_e \rangle$ such that F_e and \mathcal{C}_e are given by

\mathbb{S}_i	F_i	F_e	\mathcal{C}_e
\mathbb{T}	\mathbb{I}	\emptyset	$\{x \simeq \epsilon\}$
\mathbb{T}	$\phi * \psi$	$\{(\mathbb{T}\phi : \mathbf{a}), (\mathbb{T}\psi : \mathbf{b})\}$	$\{x \simeq \mathbf{ab}\}$
\mathbb{F}	$\phi \multimap \psi$	$\{(\mathbb{T}\phi : \mathbf{a}), (\mathbb{F}\psi : x\mathbf{a})\}$	$\{x\mathbf{a} \simeq x\mathbf{a}\}$
\mathbb{F}	$\mathbf{L}_u^r \phi$	$\{(\mathbb{F}\phi : \mathbf{a})\}$	$\{x\lambda(r) \simeq_u \mathbf{a}\}$
\mathbb{T}	$\mathbf{M}_u^r \phi$	$\{(\mathbb{T}\phi : \mathbf{a}\lambda(r))\}$	$\{x \simeq_u \mathbf{a}\lambda(r)\}$
\mathbb{F}	$\mathbf{N}_u^r \phi$	$\{(\mathbb{F}\phi : \mathbf{a}\lambda(r))\}$	$\{x\lambda(r) \simeq_u \mathbf{a}\lambda(r)\}$
\mathbb{T}	$\mathbf{L}_u^r \phi$	$\{(\mathbb{T}\phi : \mathbf{a})\}$	$\{x\lambda(r) \simeq_u \mathbf{a}\}$
\mathbb{F}	$\mathbf{M}_u^r \phi$	$\{(\mathbb{F}\phi : \mathbf{a}\lambda(r))\}$	$\{x \simeq_u \mathbf{a}\lambda(r)\}$
\mathbb{T}	$\mathbf{N}_u^r \phi$	$\{(\mathbb{T}\phi : \mathbf{a}\lambda(r))\}$	$\{x\lambda(r) \simeq_u \mathbf{a}\lambda(r)\}$
<i>Otherwise</i>		\emptyset	\emptyset

with $\mathbf{a} = c_{2i+2}$ and $\mathbf{b} = c_{2i+3}$.

- S_i is an agent constraint of the form $x\lambda(r) =_u y$:
 - If $\gamma_r \cap (\mathcal{E}(x) \cup \mathcal{E}(y)) \not\subseteq \{c_1, \dots, c_{2i+1}\}$, then $\langle \mathcal{F}_{i+1}, \mathcal{C}_{i+1} \rangle = \langle \mathcal{F}_i, \mathcal{C}_i \rangle$;
 - If $\langle \mathcal{F}_i, \mathcal{C}_i \cup \{x\lambda(r) =_u y\} \rangle \notin \wp$ then $\langle \mathcal{F}_{i+1}, \mathcal{C}_{i+1} \rangle = \langle \mathcal{F}_i, \mathcal{C}_i \rangle$;
 - If $\langle \mathcal{F}_i, \mathcal{C}_i \cup \{x\lambda(r) =_u y\} \rangle \in \wp$, then $\langle \mathcal{F}_{i+1}, \mathcal{C}_{i+1} \rangle = \langle \mathcal{F}_i, \mathcal{C}_i \cup \{x\lambda(r) =_u y\} \rangle$.

PROPOSITION 10

For any $i \in \mathbb{N}$, the following properties hold:

1. $(\mathbb{F}\varphi : c_1) \in \mathcal{F}_i$ and $c_1 \simeq c_1 \in \mathcal{C}_i$;
2. $\mathcal{F}_i \subseteq \mathcal{F}_{i+1}$ and $\mathcal{C}_i \subseteq \mathcal{C}_{i+1}$;
3. $\langle \mathcal{F}_i, \mathcal{C}_i \rangle_{i \geq 0} \in \wp$;
4. $\mathcal{A}_r(\mathcal{C}_i) \subseteq \{c_1, c_2, \dots, c_{2i+1}\}$.

The limit CSS $\langle \mathcal{F}_\infty, \mathcal{C}_\infty \rangle$ of $\langle \mathcal{F}_i, \mathcal{C}_i \rangle_{i \geq 0}$ is defined by $\mathcal{F}_\infty = \bigcup_{i \geq 0} \mathcal{F}_i$, $\mathcal{C}_\infty = \bigcup_{i \geq 0} \mathcal{C}_i$.

PROPOSITION 11

The following properties hold:

1. $\langle \mathcal{F}_\infty, \mathcal{C}_\infty \rangle \in \wp$;
2. For any labelled formula $(\mathbb{S}\phi : x)$, if $\langle \mathcal{F}_\infty \cup \{(\mathbb{S}\phi : x)\}, \mathcal{C}_\infty \rangle \in \wp$, then $(\mathbb{S}\phi : x) \in \mathcal{F}_\infty$;
3. For any agent constraint $x\lambda(r) =_u y$, if $\langle \mathcal{F}_\infty, \mathcal{C}_\infty \cup \{x\lambda(r) =_u y\} \rangle \in \wp$, then $x\lambda(r) =_u y \in \mathcal{C}_\infty$.

LEMMA 10

The limit CSS is an Hintikka CSS.

PROOF. By Proposition 11, $\langle \mathcal{F}_\infty, \mathcal{C}_\infty \rangle \in \wp$. We must verify that all conditions of Definition 12 hold. \square

THEOREM 3 (Completeness).

Let φ be a formula. If φ is valid, then there exists a proof for φ .

PROOF. Similar to the proof of the corresponding result in [15]. We suppose that there is no proof for the formula φ and show that φ is not valid. The method which we present here allows us to build a limit CSS $\langle \mathcal{F}_\infty, \mathcal{C}_\infty \rangle$ that, by Lemma 10, is a Hintikka CSS. By property 1 of Proposition 10, $(\mathbb{F}\varphi : c_1) \in \mathcal{F}_i$, for any $i \geq 0$. By the definition of a limit CSS, $(\mathbb{F}\varphi : c_1) \in \mathcal{F}_\infty$. By Lemma 8, φ is not valid. \square

Acknowledgements

We are grateful to Simon Docherty and to the anonymous referees for their comments on earlier drafts of this paper. We also thank many colleagues, including particularly James Brotherston, Johan van Benthem and Peter O’Hearn, among many, who have commented on documents related to this document.

References

- [1] G. Anderson and D. Pym. A calculus and logic of bunched resources and processes. *Theoretical Computer Science*, **614**, 63–96, 2016.
- [2] J.-M. Andreoli. Logic programming with focusing proofs in linear logic. *Journal of Logic and Computation*, **2**, 297–347, 1992.

- [3] A. Baltag, B. Coecke and M. Sadrzadeh. Epistemic actions as resources. *Journal of Logic and Computation*, **17**, 555–585, 2006.
- [4] J. Brotherston and J. Villard. Parametric completeness for separation theories. In *ACM Symposium on Principles of Programming Languages, POPL 41*, pp. 453–464, San Diego, CA, 2014.
- [5] J. Bruckheimer (Producer) and T. Scott (Director). *Crimson Tide*. Hollywood Pictures, 1995.
- [6] S. Castellán and N. Yoshida. Causality in linear logic. In *Proc. FoSSaCS 2019*, pp. 150–168. Vol. 11425 of LNCS, Springer, 2019.
- [7] I. Cervesato. Typed multiset rewriting specifications of security protocols. *Electronic Notes in Theoretical Computer Science*, **40**, 8–51, 2001.
- [8] M. Collinson and D. Pym. Algebra and logic for resource-based systems modelling. *Mathematical Structures in Computer Science*, **19**, 959–1027, 2009.
- [9] M. Collinson, B. Monahan and D. Pym. *A Discipline of Mathematical Systems Modelling*. College Publications, 2012.
- [10] M. Collinson, K. McDonald and D. Pym. Layered graph logic as an assertion language for access control policy models. *Journal of Logic and Computation*, **27**, 41–80, 2017.
- [11] M. Collinson, K. McDonald and D. Pym. A substructural logic for layered graphs. *Journal of Logic and Computation*, **24**, 953–988, 2014. Erratum at <https://doi.org/10.1093/logcom/exv019>.
- [12] D. Coumans, M. Gehrke and L. van Rooijen. Relational semantics for full linear logic. *Journal of Applied Logic*, **12**, 50–66, 2014. doi: doi.org/10.1016/j.jal.2013.07.005.
- [13] J.-R. Courtault and D. Galmiche. A modal separation logic for resource dynamics. *Journal of Logic and Computation*, **28**, 733–778, 2018. doi: [10.1093/logcom/exv031](https://doi.org/10.1093/logcom/exv031).
- [14] J.-R. Courtault, H. van Ditmarsch and D. Galmiche. An epistemic separation logic. In *22nd International Workshop on Logic, Language, Information, and Computation, WoLLIC 2015*, pp. 156–173. Vol. 9160 of LNCS, Springer, Bloomington, IN, 2015.
- [15] J.-R. Courtault, D. Galmiche and D. Pym. A logic of separating modalities. *Theoretical Computer Science*, **637**, 30–58, 2016. doi: [10.1016/j.tcs.2016.04.040](https://doi.org/10.1016/j.tcs.2016.04.040).
- [16] H. van Ditmarsch, J. Y. Halpern, W. van der Hoek and B. Kooi, eds. *Handbook of Epistemic Logic*. College Publications, 2015.
- [17] S. Docherty and D. Pym. Intuitionistic layered graph logic. In *Proc. IJCAR 2016*, pp. 469–486. Vol. 9706 of LNCS, Springer, Coimbra, Portugal, 2016.
- [18] S. Docherty and D. Pym. A stone-type duality theorem for separation logic via its underlying bunched logics. *Electronic Notes in Theoretical Computer Science*, **336**, 101–118, 2018.
- [19] S. Docherty and D. Pym. A stone-type duality theorem for separation logic via its underlying bunched logics. *Logical Methods in Computer Science*, **15**, 27:1–27:51, 2019.
- [20] D. Galmiche, P. Kimmel and D. Pym. A substructural epistemic resource logic. In *Proc. ICLA 2017*, pp. 106–122. Vol. 10119 of LNCS, Springer, 2017.
- [21] D. Galmiche, D. Méry and D. Pym. The semantics of BI and resource tableaux. *Mathematical Structures in Computer Science*, **15**, 1033–1088, 2005.
- [22] D. Garg, L. Bauer, K. D. Bowers, F. Pfenning and M. K. Reiter. A linear logic of authorization and knowledge. In *11th European Symposium on Research in Computer Security, ESORICS 2006*, pp. 297–312. Vol. 4189 of LNCS, Springer, 2006.
- [23] J. Y. Girard. Linear logic. *Theoretical Computer Science*, **50**, 1–102, 1986.
- [24] J. Halpern and R. Pucella. Modeling adversaries in a logic for security protocol analysis. In *Formal Aspects of Security, FASec 2002*, pp. 115–132. Vol. 2629 of LNCS, Springer, 2003.

- [25] S. Ishtiaq and P. O’Hearn. BI as an assertion language for mutable data structures. In *28th ACM Symposium on Principles of Programming Languages (POPL)*, pp. 14–26, London, 2001.
- [26] J. Hodas and D. Miller. Logic programming in a fragment of intuitionistic linear logic. *Information and Computation*, **110**, 327–365, 1994.
- [27] D. Larchey-Wendling. The formal strong completeness of partial monoidal Boolean BI. *Journal of Logic and Computation*, **26**, 605–640, 2014.
- [28] R. Milner. Calculi for synchrony and asynchrony. *Theoretical Computer Science*, **25**, 269–310, 1983.
- [29] P. Naumov and J. Tao. Budget-constrained knowledge in multiagent systems. In *Proc. AAMAS 2015*, pp. 219–226, International Foundation for Autonomous Agents and Multiagent Systems, 2015.
- [30] P. O’Hearn and D. Pym. The logic of bunched implications. *The Bulletin of Symbolic Logic*, **5**, 215–244, 1999.
- [31] P. O’Hearn. On bunched typing. *Journal of Functional Programming*, **13**, 747–796, 2003.
- [32] P. W. O’Hearn. Resources, concurrency and local reasoning. *Theoretical Computer Science*, **375**, 271–307, 2007.
- [33] R. Pucella. Knowledge and security. Chapter 12 of [16], pp. 591–655.
- [34] D. Pym. *The Semantics and Proof Theory of the Logic of Bunched Implications*. Applied Logic Series, vol. **26**. Kluwer Academic Publishers, 2002.
- [35] D. Pym. Resource semantics: logic as a modelling technology. *ACM SIGLOG News*, **6**, 5–41, April 2019.
- [36] D. Pym, P. O’Hearn and H. Yang. Possible worlds and resources: the semantics of BI. *Theoretical Computer Science*, **315**, 257–305 Erratum: p. 22, l. 22 (preprint), p. 285, l.-12 (TCS): ‘, for some P , $Q \equiv P; P'$ ’ should be ‘ $P \vdash Q$ ’.
- [37] J. Reynolds. Separation logic: a logic for shared mutable data structures. In *IEEE Symposium on Logic in Computer Science, LICS 2002*, pp. 55–74, Copenhagen, Denmark, July 2002.
- [38] B. Schneier. *The Weakest Link*. https://www.schneier.com/blog/archives/2005/02/the_weakest_lin.html. Schneier on Security (<https://www.schneier.com>, 2005).
- [39] B. Toninho and L. Caires. A spatial-epistemic logic for reasoning about security protocols. In *8th Int. Workshop on Security Issues in Concurrency, SecCo 2010*, Electronic Proceedings in Theoretical Computer Science (EPTCS) (arXiv.org), 2010.

Received 00 Month 20xx