

DOING WITHOUT ACTION TYPES

HEIN DUIJF

Vrije Universiteit Amsterdam

JAN BROERSEN, ALEXANDRA KUNCOVÁ
and ALDO IVÁN RAMÍREZ ABARCA

Utrecht University

Abstract. This paper explores the analysis of ability, where ability is to be understood in the epistemic sense—in contrast to what might be called a causal sense. There are plenty of cases where an agent is able to perform an action that guarantees a given result even though she does not know which of her actions guarantees that result. Such an agent possesses the causal ability but lacks the epistemic ability. The standard analysis of such epistemic abilities relies on the notion of action types—as opposed to action tokens—and then posits that an agent has the epistemic ability to do something if and only if there is an action type available to her that she knows guarantees it. We show that these action types are not needed: we present a formalism without action types that can simulate analyses of epistemic ability that rely on action types. Our formalism is a standard epistemic extension of the theory of “seeing to it that”, which arose from a modal tradition in the logic of action.

§1. Introduction. The theory of “seeing to it that”, abbreviated to stit, grew out of a tradition that views agency as a modal notion. It originates from a series of papers by Belnap, Perloff, and Xu, culminating in their (Belnap *et al.*, 2001). In this tradition, the agency of an individual is characterized by a modal operator of the form $[i \text{ stit}] \varphi$, which is to be read as saying that agent i sees to it that φ obtains.¹ Stit theorists have proposed that the ability of an agent can reasonably be characterized by a combination of impersonal possibility (characterized by \diamond) and agency of the form $\diamond [i \text{ stit}] \varphi$, which is to be read as saying that it is possible that agent i sees to it that φ (Horty & Belnap, 1995).

Consider an agent who is instructed to pick one card from a face-down deck of cards. We may safely assume that the deck contains all 52 cards. The agent has the causal ability to pick the queen of diamonds even though she lacks the epistemic ability to do so. That is, even though she does not know which card is the queen of diamonds, the action of picking the queen of diamonds is available to her. We agree that this ability can be reasonably characterized by a formula of the form $\diamond [i \text{ stit}] \varphi$, where i depicts

Received: October 12, 2018.

2020 *Mathematics Subject Classification*: 03B42, 03B44, 03B45, 91A40.

Keywords and phrases: stit theory, knowledge, ability, action types, modal logic.

¹ In a recent survey, Lindström & Segerberg (2007, p. 1198) label this tradition the “logic of action without actions”, a label that is rather unfortunate, especially given the work on stit theory by Horty (2001), who defines a deontic operator in terms of a dominance ordering on the actions available to an agent.

the particular agent in this example. Let us call this the causal sense of ability. This paper focuses on the epistemic sense of ability, which requires knowledge.

When including knowledge in stit theory it is common to use the standard modal treatment akin to epistemic logic, where the knowledge of a given agent i is characterized by a modal operator of the form $K_i\varphi$.² Horty & Pacuit (2017) have recently argued that the epistemic stit formalism needs to incorporate *action types* to address certain puzzles about knowledge and action. More precisely, they argue that these puzzles demonstrate that the simple combinations of operators $\Diamond K_i[i \text{ stit}]\varphi$ and $K_i\Diamond[i \text{ stit}]\varphi$ both fail to characterize the epistemic sense of ability. On epistemic ability and action types, they write:

if the epistemic sense of ability requires that some single action must be known by i to guarantee the truth of φ , then this must be the *action type*, not one of its various tokens. (Horty & Pacuit, 2017, p. 626—notation adapted and emphasis added)

Horty and Pacuit therefore introduce action types into epistemic stit models and call the resulting models *labeled stit models*. They then characterize the epistemic agency of an individual by a new modal operator of the form $[i \text{ kstit}]\varphi$, which is to be read as saying that agent i sees to it that φ obtains, in an epistemic sense. In opposition to this approach, one of the main goals of our paper is to show that the stit formalism can address these puzzles about knowledge and action without making use of action types.

It is important to realize that Horty and Pacuit follow a by now standard approach to abilities under imperfect information (Horty & Pacuit, 2017, see the discussion in sec. 5.2). Indeed, action types have a long tradition in the literature on knowledge and action, especially in the literature on alternating-time temporal logic (Alur, Henzinger, & Kupferman, 2002), abbreviated to ATL, and its epistemic extension, standardly referred to as ATEL (see van der Hoek & Wooldridge, 2003; Jamroga & van der Hoek, 2004; Schobbens, 2004; Jamroga & Ågotnes, 2007). The central ability operator in such logics is of the form $\langle\langle i \rangle\rangle\varphi$, which is to be read as saying that agent i can ensure φ ; different such logics are intended to capture different senses of “can”.³ Simply stated, this tradition proposes that agent i is able, in an epistemic sense, to do φ if and only if there is an action type available to her that she knows guarantees φ .

In contrast to this tradition, in a series of papers Broersen (2008, 2011a, 2011b) analyzes the notion of *knowingly doing*, which does not rely on action types.⁴ He writes:

“knowingly doing” is an epistemic qualification concerning an action. (Broersen, 2011a, p. 144)

² Xu (2015) presents an excellent survey of stit theory and its various epistemic extensions.

³ In this tradition it is also common to include a coalitional ability operator of the form $\langle\langle A \rangle\rangle\varphi$, which is to be read as saying that the coalition A can ensure φ .

⁴ This undoubtedly resulted from an interaction between Broersen, Herzig, and Troquard (see Herzig, & Troquard, 2006; Broersen, Herzig, & Troquard, 2006b, 2007).

More precisely, an agent knowingly does φ if and only if she knows that she is seeing to it that φ . The concept of knowingly doing can hence be expressed by a simple combination of knowledge and agency: $K_i[i \text{ stit}]\varphi$.

Our central theorem provides a general correspondence result between the analysis of epistemic agency by Horty and Pacuit and the analysis of knowingly doing in Broersen's work. To achieve this, we systematically transform a labeled stit model to an epistemic stit model. We then show that the analysis of epistemic agency in a given labeled stit model corresponds to the analysis of knowingly doing in the transform epistemic stit model. This means that the analysis of epistemic ability of Horty & Pacuit (2017) can be simulated in epistemic stit theory without involving action types. The upshot of this is that epistemic ability, as studied by Horty & Pacuit (2017), can be characterized by a simple combination of impersonal possibility, knowledge, and agency of the form $\Diamond K_i[i \text{ stit}]\varphi$, which is to be read as saying that it is possible that the agent knowingly does φ .

Moreover, given that labeled stit models naturally relate to ATEL models, it is foreseeable that our central theorem can be extended to these latter models. Several results have already been established showing that various analyzes of ability in the ATL tradition can be simulated in stit theory (Broersen *et al.*, 2006a, 2007).⁵ Our central theorem indicates that the treatment of abilities under imperfect information in the ATEL tradition can be simulated in epistemic stit models without involving action types. The epistemic ability operator in the ATEL tradition again corresponds to the stit formula $\Diamond K_i[i \text{ stit}]\varphi$.

The paper is organized as follows. We start with providing a brief introduction to stit theory in §2. In §3, we introduce an epistemic extension of the basic stit framework and present the puzzles that motivated Horty and Pacuit to introduce action types into stit theory (§3.1). Readers familiar with standard stit theory can decide to skip §2 and those familiar with epistemic stit theory decide to skim over §3—with the exception of §3.1. In §4, we discuss Horty and Pacuit's analysis of epistemic ability and two prominent conditions on their models. In §5 we introduce the notion of knowingly doing and present the central result of the paper: any labeled stit model can be transformed into an epistemic stit model such that the analysis of epistemic agency in the former corresponds to that of knowingly doing in the latter. A discussion of two optional conditions on epistemic stit models (§5.1) and the condition of uniformity of available action types (§5.2) follows. In the final section the main findings are summarized, and we briefly reflect on the concept of action types and provide two small remarks on why we think it might be beneficial to do without action types. All claims are proved in the Appendix.

§2. Stit theory. The seminal contributions of Prior (1967) and Thomason (1970, 1984) gave rise to the theory of branching time that would later serve as the backbone

⁵ For example, Broersen *et al.* (2006a) show that (non-epistemic) ATL can be embedded in stit theory. The central coalitional strategic ability operator of ATL, expressed by $\langle\langle G \rangle\rangle\varphi$, corresponds to the strategic stit formula $\Diamond[G \text{ sstit}]\varphi$. (Where $[G \text{ sstit}]$ is a *strategic* group stit operator, rather than the stit operator used in this paper.)

for stit semantics (Belnap *et al.*, 2001; Horty, 2001).⁶ The branching-time models originate from a philosophical enquiry into the truth-values of temporal sentences, for example, so-called future contingencies. Belnap *et al.* (2001) present a detailed account of how our indeterministic world can be modelled.⁷ The fundamental idea is to represent the world as moments ordered in a tree of histories. (It is important to note a possible confusion: “histories” are taken to include future moments.) The upward branching of histories represents the openness of the future. Although histories may branch at a particular moment, it is conceivable that there are moments at which no history branches. The absence of backward branching represents the determinateness of the past, that is, the fact that every moment has only a single past sequence of events. Each history in such a tree-like structure represents a complete temporal evolution of the world.

A *branching-time model* involves a set of moments M , a set of histories $H \subseteq 2^M$, and a relation $<$ between moment/history pairs which represents the progression of events along a history. We use m, m' as variables for moments in M and h, h' as variables for histories in H . When a moment m and a history h satisfy $m \in h$, this can be taken to mean that m occurs on h , or that h passes through m . Because of indeterminacy, there may be multiple histories that pass through m , so we let $H_m = \{h \in H \mid m \in h\}$ denote the set of histories through m . We use m/h as variables for moment/history pairs that satisfy $m \in h$ or, equivalently, $h \in H_m$. It is common to call these moment/history pairs indices; we let Ind denote the set of indices. A given index includes the current moment and the complete temporal evolution of the world. Finally, a valuation function V assigns to each propositional variable $q \in \mathcal{P}$ the set of indices $V(q)$ where q obtains.

DEFINITION 2.1. (Branching-time Model) *A branching-time frame is a tuple $BTF = \langle M, H, < \rangle$, where M is the set of moments, $H \subseteq 2^M$ is the set of histories, and $< \subseteq Ind \times Ind$ is a relation on the set of indices. Moreover, BTF is required to satisfy the following:*

- for every history h , the ordering $<_h$ on h induced by $<$, viz. $m <_h m'$ iff $m/h < m'/h$, constitutes a linear ordering. (*Linear Histories*)
- for every moment m and all histories h, h' such that $m \in h \cap h'$, it holds that $\{m' \in h \mid m' <_h m\} = \{m' \in h' \mid m' <_{h'} m\}$. In other words, there is a single past sequence of events. (*Past Determinacy*)

⁶ Logics of bringing it about (Pörn, 1970; Kanger, 1971; Elgesem, 1993, 1997) are related to stit theory; the difference is that stit frameworks typically build on branching-time structures and impose the independence of agency requirement. The stit operator that takes center stage in this paper is known as the Chellas stit operator (named after Chellas, 1992), as opposed to the deliberative and the achievement stit operators (originally proposed by Belnap & Perloff, 1988 and later simplified by Horty & Belnap, 1995). Later, stit theory has been extended to include *strategic action* by Horty (2001, chap. 7) and Broersen (2009). For more recent surveys we would recommend (Broersen & Herzig, 2015) and (Xu, 2015).

⁷ Perloff & Belnap (2011, pp. 583–584) write: “Part of the idea of indeterminism as we conceive it is that at any given moment there are a variety of ways in which the world might proceed. Such possibilities are real, not merely epistemic; they are possibilities”.

A branching-time model $BTM = \langle M, H, <, V \rangle$ is a branching-time frame supplemented with a valuation $V : \mathcal{P} \rightarrow 2^{Ind}$.⁸

These branching-time models are typically used to provide semantics for a logical language that includes a past operator P , a future operator F , and a historical necessity operator \Box .⁹ Intuitively, $P\varphi$ is true at a moment/history pair m/h if and only if there is a moment m' before m on h where φ obtains; $F\varphi$ holds at a moment/history pair m/h if and only if there is a moment m' after m on h where φ obtains. The historical necessity operator $\Box\varphi$ expresses that φ holds at the current moment in the moment/history pair m/h regardless of how the future unfolds. The dual $\Diamond\varphi$ expresses that there is a possible way for the future to unfold such that φ holds now.

DEFINITION 2.2. (Evaluation Rules Temporal Formulas) *Let BTM be a branching-time model, and let φ be a formula constructed using propositional constants and temporal operators P, F , and \Box . Then the truth of φ , in a moment/history pair m/h in BTM , notation: $BTM, m/h \models \varphi$, is given by the following (suppressing the standard propositional clauses and the model BTM):*

$$\begin{aligned} m/h \models P\varphi & \text{ iff } \text{there is an } m' \in h \text{ satisfying } m' <_h m \text{ and } m'/h \models \varphi; \\ m/h \models F\varphi & \text{ iff } \text{there is an } m' \in h \text{ satisfying } m <_h m' \text{ and } m'/h \models \varphi; \\ m/h \models \Box\varphi & \text{ iff } \text{every } h' \in H_m \text{ satisfies } m/h' \models \varphi. \end{aligned}$$

Given these semantics, the idea that the future may still be open is represented by the invalidity of the formula $F\varphi \rightarrow \Box F\varphi$. In other words, there is a branching-time model BTM and a moment/history pair m/h such that $BTM, m/h \models F\varphi$ while $BTM, m/h \not\models \Box F\varphi$.¹⁰ Or, equivalently, $BTM, m/h \models F\varphi \wedge \Diamond \neg F\varphi$.

The fundamental idea of stit theory is to treat agency as a modality in these branching-time models. At a particular moment m we may view H_m as representing the possibilities that are still open. Conversely, the histories outside H_m are no longer possible, or accessible, at moment m . Given that the histories in H_m are still open, an agent's action, or choice, is viewed as restricting the possible histories to a subset K of H_m . Accordingly, "the agent sees to it that φ " means that the truth of φ is guaranteed by an action or choice K of the agent. When Ann empties her glass of milk, the nature of her action on this view is to constrain the possible histories to those where the glass of milk is emptied. Hence, an action is identified with a subset of the possible histories. This induces the reading that an agent sees to it that φ only if she performs an action that constrains the possible worlds to only φ -worlds.

A *branching-time agency model* supplements a branching-time model with a finite set of agents Ags and sets of available actions, one for each agent at each moment.

⁸ Branching-time models are often conceived as trees. However, for some purposes a branching-time model is better thought of as a forest consisting of several independent trees. This may, for instance, be important when adding an epistemic indistinguishability relation. For example, to allow for the possibility that an agent does not know the exact past, it may be advantageous to interpret a branching-time model as a forest rather than a tree.

⁹ The literature also discusses other temporal operators; the most common additional operators are the next operator and the until operator.

¹⁰ In line with the terminology of Horty (2001, p. 10), this means that the formula $F\varphi$ may not be *moment-determinate*: "A formula is moment determinate if it is always, at any moment, either settled true or settled false".

Given a moment m and an agent i , the set of available actions is given by a collection of subsets of the possible histories $Act_i^m \subset 2^{H_m}$.¹¹ Since a history can be viewed as the complete temporal evolution of the world, it includes the actions that the agents are performing. The particular action that agent i executes at a moment/history pair m/h is given by $Act_i^m(h)$, which is the action $K \in Act_i^m$ satisfying $h \in K$.¹²

DEFINITION 2.3. (Branching-time Agency Model) *A branching-time agency frame is a tuple BTAF = $\langle M, H, <, Ags, (Act_i^m) \rangle$, involving a branching-time frame $\langle M, H, < \rangle$, a finite set of agents Ags , and for each moment m and each agent $i \in Ags$ it holds that $Act_i^m \subset 2^{H_m}$ is a finite set of actions available to agent i at moment m , satisfying the following:*

- for every moment m and every agent i , Act_i^m constitutes a partition of H_m . (Partition)
- for every moment m , every agent i , and all histories $h, h' \in H_m$, if there is a moment m' such that $m' \in h \cap h'$ and $m <_h m'$, then $h \in Act_i^m(h')$. (No Choice between Undivided Histories)
- for every moment m , any set of histories $h_i \in H_m$, one for each agent $i \in Ags$, it holds that $\bigcap_{i \in Ags} Act_i^m(h_i) \neq \emptyset$.¹³ (Independence of Agency)

A branching-time agency model is a branching-time agency frame supplemented with a valuation $V : \mathcal{P} \rightarrow 2^{Ind}$.

Figure 1 depicts such a branching-time agency frame.¹⁴ For example, at m_1 , agents i and j both have two available actions, where each action is identified with a subset of the possible histories, in particular, $Act_i^{m_1} = \{\{h_1, h_6, h_7\}, \{h_2, h_3, h_4, h_5\}\}$. These branching-time agency models are used to interpret a logical language that includes agency operators $[i \text{ stit}]$, one for each agent i . Intuitively, $[i \text{ stit}]\varphi$ is true at a moment/history pair m/h if and only if the truth of φ is guaranteed by the action of the agent i . That is, agent i performs a certain action thereby constraining the possible histories to only those where φ holds. It may be useful to add that $[i \text{ stit}]\varphi$ may be interpreted, relative to a moment/history pair m/h , as “agent i guarantees that φ holds regardless of what the others do”.

DEFINITION 2.4. (Evaluation Rule Agency) *Let BTAM = $\langle M, H, <, Ags, (Act_i^m), V \rangle$ be a branching-time agency model. Then the evaluation rule for the stit operator at*

¹¹ The main stit theorists are ambiguous about the interpretation of Act as either choices or actions. According to Horty (1996, p. 274—emphasis added), “ Act is a device for representing the constraints that an agent is able to exercise upon the course of history at a given moment, the actions or choices open to him at that moment”. Belnap *et al.* (2001, pp. 33–34—notation adapted and emphasis added) writes: “The equivalence classes belonging to Act can be thought of as the possible choices or actions available”. For our current purposes, the elements of Act are best thought of as actions.

¹² The fact that there is a unique action $K \in Act_i^m$ satisfying $h \in K$ follows from the requirement that Act_i^m be a partition of H_m (see below).

¹³ See (Horty, 2001, sec. 2.4) for a useful discussion of the independence of agency requirement.

¹⁴ It may be helpful to point out a convention for interpreting the figures: whenever a formula φ is printed next to a history h emanating from a moment m , then this means that φ holds at that moment/history pair. For example, this means that q holds at m_2/h_4 .




Fig. 1. A branching-time agency frame.

a given moment/history pair m/h is given by the following (suppressing the model BTAM):

$$m/h \models [i \text{ stit}]\varphi \text{ iff for every } h' \in \text{Act}_i^m(h) \text{ it holds that } m/h' \models \varphi.$$

We will write $\langle i \text{ stit} \rangle \varphi$ for the dual of $[i \text{ stit}]\varphi$, i.e., for $\neg[i \text{ stit}]\neg\varphi$.

It might be helpful to briefly revisit the branching-time agency frame depicted in Figure 1. Let us focus on moment m_2 . At m_2 , agent j has two available actions, where each action is identified with a subset of the possible histories, in particular, $\text{Act}_j^{m_2} = \{\{h_2\}, \{h_3, h_4\}\}$. Note that at index m_2/h_4 it is the case that q holds even though it is not the case that agent j sees to it that q holds. That is, $m_2/h_4 \models q \wedge \neg[j \text{ stit}]q$. The second conjunct follows from the fact that q does not hold at m_2/h_3 . This illustrates that it is foreseeable that not everything that happens is brought about by an agent.¹⁵

§3. Knowledge in stit theory. The branching-time agency models serve as the backbone for epistemic stit models. An *epistemic stit model* supplements a branching-time agency model with a set of indistinguishability relations $\sim_i \subseteq \text{Ind} \times \text{Ind}$, one for each agent $i \in \text{Ags}$. The idea is that whenever $m/h \sim_i m'/h'$ holds, then agent i cannot distinguish these indices. Given such an indistinguishability relation, we can define the

¹⁵ Horty & Belnap (1995) propose that we characterize the ability of an individual by a combination of impersonal possibility and agency of the form $\diamond[i \text{ stit}]\varphi$, which is to be read as saying that it is possible that agent i sees to it that φ obtains (it should be noted that they use a deliberative stit operator). A model similar to the one depicted in Figure 1 reveals that this proposal escapes an objection raised by Kenny (1975, 1976) that is based on the example of a poor darts player who accidentally hits the bull’s-eye (see Horty & Belnap, 1995, pp. 610–611).




Fig. 2. The epistemic stit model that Horty & Pacuit (2017, Figure 3) use to depict the first coin example.

information set of agent i at a moment/history pair m/h as the set consisting of those moment/history pairs that the agent cannot distinguish from m/h . More specifically, this information set is given by $\{m'/h' \in Ind \mid m'/h' \sim_i m/h\}$.¹⁶ An epistemic stit model is thus best thought of as representing the possibilities, knowledge (or information), and available actions at several moments in time.

DEFINITION 3.1. (Epistemic Stit Models) *An epistemic stit frame is a tuple $\langle M, H, <, Ags, (Act_i^m), (\sim_i) \rangle$, involving a stit frame $\langle M, H, <, Ags, (Act_i^m) \rangle$ and for each agent $i \in Ags$ an indistinguishability relation $\sim_i \subseteq Ind \times Ind$, satisfying the following:*

- for every $i \in Ags$ the indistinguishability relation \sim_i is an equivalence relation on the set of indices Ind . (Equivalence)

An epistemic stit model is a stit frame supplemented with a valuation $V : \mathcal{P} \rightarrow 2^{Ind}$.

To illustrate epistemic stit models, let us consider the two examples that Horty & Pacuit (2017, pp. 623–625, Figures 3 and 4) discuss to motivate the introduction of action types into stit theory, here depicted in Figures 2 and 3. In each of these examples, an agent, called j , first places a coin on the table either heads up, by selecting K_1 , or tails up, by selecting K_2 . Afterwards the other agent, called i , is offered the option to bet on whether agent j put the coin heads up or tails up. This happens at either moment m_2 or moment m_3 , depending on agent j 's choice. If agent j chooses to place the coin heads up, then the resulting moment will be m_2 , and agent i bets heads by selecting K_3 and bets tails by selecting K_4 . If agent j chooses to place the coin tails up, then the resulting moment will be m_3 , and agent i bets heads by selecting K_5 and bets tails by selecting K_6 . If we let φ denote the state of affairs where agent i bets correctly, then we see that φ holds at m_2/h_1 and m_3/h_4 , and it is false at m_2/h_2 and m_3/h_3 .

¹⁶ In artificial intelligence research such indistinguishability relations are standardly used to represent an agent's knowledge (see Meyer & van der Hoek, 1995; Fagin *et al.*, 2003). They straightforwardly correspond to partition structures that are commonly employed in game theory and economics to model the information states of the players (see Aumann, 1999).




Fig. 3. The epistemic stit model that Horty & Pacuit (2017, Figure 4) use to depict the second coin example.

The two coin examples are identical with regard to their causal structure, but they differ in the information that is available to agent i . In the first coin example, agent i knows whether the coin has been placed heads up or tails up. This is represented by the fact that agent i can distinguish between moments m_2 and m_3 . This indistinguishability relation is represented in the figures by the dashed line between indices.¹⁷ More precisely, the indistinguishability relation for agent i is given by the equivalence classes $\{m_2/h_1, m_2/h_2\}$ and $\{m_3/h_3, m_3/h_4\}$. In the second coin example, agent i does not know whether the coin has been placed heads up or tails up. This is represented by the fact that agent i cannot distinguish between moments m_2 and m_3 . More precisely, the indistinguishability relation for agent i is given by the single equivalence class $\{m_2/h_1, m_2/h_2, m_3/h_3, m_3/h_4\}$. We will return to these models in §3.1.

It is common among logicians, computer scientists, and economists to use these epistemic models to provide semantics for a logical language that includes knowledge operators K_i , one for each agent i . Intuitively, $K_i\varphi$ is true at a moment/history pair m/h if and only if φ is true at every moment/history pair that agent i cannot distinguish from m/h . In other words, we can say that an agent knows something if her information set entails it.

DEFINITION 3.2. (Evaluation Rule Knowledge) *Let \mathcal{M} be an epistemic stit model. Then the evaluation rule for the knowledge operator at a given moment/history pair m/h is given by the following:*

$$\mathcal{M}, m/h \models K_i\varphi \quad \text{iff} \quad \text{for every } m'/h' \in \text{Ind satisfying } m/h \sim_i m'/h' \text{ it holds that } \mathcal{M}, m'/h' \models \varphi.$$

We will write $\hat{K}_i\varphi$ for the dual of $K_i\varphi$, i.e., for $\neg K_i\neg\varphi$.

¹⁷ It may be helpful to point out a convention for interpreting the figures: the indistinguishability relation is closed under reflexivity, transitivity, and symmetry.

3.1. Some puzzles of knowledge and action. To study abilities under uncertainty, the two coin examples play a central role in the work of Horty & Pacuit (2017) (see the models in Figures 2 and 3).¹⁸ They discard the idea that epistemic ability can be expressed by a simple combination of knowledge, impersonal possibility, and agency. Their refutation of such simple combinations builds on the idea that a theory of epistemic ability should discriminate between the two coin examples: the agent in the first coin example is able to see to it that φ in the epistemic sense, whereas she is not able to do so in the second coin example (where φ represents the state of affairs where agent i wins). They discuss and ultimately reject two plausible candidates for representing the idea that agent i is able, in the epistemic sense, to do φ : $K_i \diamond [i \text{ stit}] \varphi$ and $\diamond K_i [i \text{ stit}] \varphi$. Let us see why.

It should first be noted that their models demonstrate that agent i is able, in a purely causal sense, to see to it that φ , which is represented by the formula $\diamond [i \text{ stit}] \varphi$. After all, this formula is settled true at both m_2 and m_3 in their models.

Second, since $\diamond [i \text{ stit}] \varphi$ is settled true at both m_2 and m_3 in both models, it immediately follows that $K_i \diamond [i \text{ stit}] \varphi$ is also settled true at both m_2 and m_3 in both models. This means that the formula $K_i \diamond [i \text{ stit}] \varphi$ does not discriminate between the two models, whereas our judgment on epistemic ability does. This formula is hence not an accurate formalization of epistemic ability.

Third, note that $K_i [i \text{ stit}] \varphi$ is settled false at both m_2 and m_3 in both models. For example, in both their models, the agent cannot distinguish m_2/h_1 from m_2/h_2 which entails that $m_2/h_1 \models \neg K_i \varphi$, which in turn entails that $K_i [i \text{ stit}] \varphi$ does not hold at m_2/h_1 . This fact entails that $\diamond K_i [i \text{ stit}] \varphi$ is settled false at both m_2 and m_3 in both models. This means that this formula also does not discriminate between the two models, whereas our judgment on epistemic ability does. This formula hence also does not qualify as an accurate formalization of epistemic ability.

It is important to flag that this argument for the inaccuracy of both these formulas depends on the particular models that are used to represent the two coin examples. But are these *correct* models of these coin examples? We will argue in the remainder that these examples can be rectified if we transform the models of these examples in a systematic way. In light of our transformation and our central correspondence result, we will see that the proposal and arguments by Horty and Pacuit point out that $\diamond K_i [i \text{ stit}] \varphi$ is *actually an accurate formalization of epistemic ability*.

§4. Action types. Horty & Pacuit (2017, sec. 4) have recently imported ideas on action types and abilities under uncertainty into stit theory. Simply stated, they submit that agent i is able, in an epistemic sense, to do φ if and only if there is an action type available to her that she knows guarantees φ . We will first discuss their approach in detail and then establish a correspondence to standard epistemic stit models that eschew action types in §5.

Horty and Pacuit extend traditional epistemic stit models to *labeled* stit models by postulating a set of action types Tps and a label function Lbl that maps each action

¹⁸ Actually, they introduce a third game to emphasize the need to introduce action types into stit theory. In light of our general correspondence result (see next section), we can safely restrict our attention to the earlier two examples.

token K to a particular action type $Lbl(K) \in Tps$.¹⁹ We will often write $Lbl_i(m/h)$, instead of $Lbl(Act_i^m(h))$, to denote the action type that agent i performs at m/h . In addition, a partial execution function Exn is added, which maps each action type $\tau \in Tps$, moment m , and agent i to a particular action token $Exn_i^m(\tau) \in Act_i^m$.²⁰ It is partial because there may be moments at which a certain action type is not executable. A *labeled stit* model is a tuple of the form:

$$\langle M, H, <, Ags, (Act_i^m), (\sim_i), V, Tps, Lbl, Exn \rangle.$$

With the help of these action types, one can define a new epistemic stit operator, denoted by $[i \text{ kstit}]$, which is distinct from the standard stit operator, denoted by $[i \text{ stit}]$, and is meant to capture a subjective sense of agency. That is, an agent epistemically sees to it that φ if and only if she knows that the action type she performs guarantees φ . Or, equivalently, if and only if the action type she performs guarantees φ at every indistinguishable index. More formally:

DEFINITION 4.1. (Evaluation Rule Epistemic Agency) *Let \mathcal{M} be a labeled stit model. Then the evaluation rule for the kstit operator is given by:*

$$\mathcal{M}, m/h \models [i \text{ kstit}]\varphi \quad \text{iff} \quad \text{for every } m'/h' \text{ that satisfies } m'/h' \sim_i m/h \text{ and for any } m'/h'' \text{ that satisfies } Lbl_i(m/h) = Lbl_i(m'/h'') \text{ it holds that } m'/h'' \models \varphi.$$

It is important to emphasize that the analysis of epistemic ability of Horty & Pacuit (2017) relies on two constraints on the labeled stit models. First, they concede that their analysis of epistemic ability “makes most sense, and is most useful, when the indistinguishability relation ...can be thought of as a relation, not just between indices, but between moments themselves” (Horty & Pacuit, 2017, p. 631—condition C4). That is, they impose the restriction that agents cannot know more about the current moment/history pair than what is historically settled. This requirement can be characterized by the following semantic condition:

- for any agent i and for all $m/h, m'/h'_1$ if $m/h \sim_i m'/h'_1$, then for any m'/h'_2 it holds that $m/h \sim_i m'/h'_2$.
(Settled Knowledge)

This semantic condition can be characterized by an axiom scheme on epistemic frames:

PROPOSITION 4.2. *Let \mathcal{F} be an epistemic stit frame. Then, \mathcal{F} satisfies the settled knowledge condition if and only if*

$$\mathcal{F} \models K_i \varphi \rightarrow K_i \Box \varphi.$$

¹⁹ Horty & Pacuit (2017, p. 626) write: “Once action types have been introduced into stit logic, it is most natural to assume that it is the execution of these action types, rather than the performance of concrete action tokens, that falls most directly within the agent’s control”. Note that each action token is associated with exactly one action type. Horty & Pacuit (2017, p. 627) write: “it seems best to assume that the execution of different action types leads to the performance of different action tokens”.

²⁰ Horty & Pacuit (2017, p. 627) impose the following two execution/label constraints: (i) $Exn_i^m(Lbl_i(m/h)) = Act_i^m(h)$ and (ii) if $Exn_i^m(\tau)$ is defined, then $Lbl_i(Exn_i^m(\tau)) = \tau$.

It is important to flag that this condition aligns with existing frameworks in the computer science literature on knowledge and action. Models in the ATL tradition involve a set of states, and actions are thought of as transitions between such states. In light of the connections between stit theory and ATL (Broersen *et al.*, 2006a, 2007), a state in an ATL model can be thought of as a moment in a stit model. In the ATL tradition, imperfect information is usually modelled using an epistemic indistinguishability relation on states (see van der Hoek & Wooldridge, 2003), which hence corresponds to an epistemic indistinguishability relation on moments.

So, what would an epistemic indistinguishability relation on indices correspond to in ATL models? An index involves a moment and a history and therefore includes a representation of the entire temporal evolution of the world. Since a stit theorist can think of a history as a set of linearly ordered moments, ATL theorists can think of a history as a set of linearly ordered states.²¹ We postulate that an ATL theorist should think of an index in stit theory as a *dynamic state*, where a dynamic state involves the current state and the history of states that represent the temporal evolution of the world or system. (It is important to, once again, note a possible confusion: “histories” are taken to include future states.) In contrast, we propose that a moment be thought of as a *static state*, which represents the current state but does not include a representation of the temporal evolution of the world or system. Accordingly, we will call an epistemic stit model that satisfies (Settled Knowledge) a *static* epistemic stit model.

Second, to ensure that the definition of the *kstit* operator makes sense Horty & Pacuit (2017, p. 628—condition C1) stipulate that labeled stit frames must satisfy an extra condition that might be called *uniformity of available action types*: This condition says that whenever two indices are indistinguishable for a given agent, then the agent has the same action types available at these indices. Let us use Tps_i^m to denote the action types that are available to agent i at m in a given labeled stit model. That is, $Tps_i^m = \{\tau \in Tps \mid \text{there is a } h \in H_m \text{ such that } \tau = Lbl_i(m/h)\}$. The condition of uniformity of available action types can be expressed by the following semantic condition, which we’ll often abbreviate to (UAAT):

- for every $m/h, m'/h' \in Ind$ if $m/h \sim_i m'/h'$ then $Tps_i^m = Tps_i^{m'}$.²²
(Uniformity of Available Action Types)

With the exception of §5.2, we assume that labeled stit models meet this requirement in the remainder of the paper.

Horty and Pacuit’s new notion of epistemic agency is then used to characterize the notion of epistemic ability using a combination of epistemic agency and impersonal possibility:

$$\diamond [i \text{ kstit}] \varphi,$$

that is, this formula “can be taken to represent the idea that the agent i has the ability, in the epistemic sense, to see to it that φ ” (Horty & Pacuit, 2017, p. 630—notation adapted).

²¹ These are sometimes called paths or computations (i.e., see van der Hoek & Wooldridge, 2003).

²² This corresponds to the discussion of Herzig & Troquard (2006), more specifically Hypothesis 3 in p. 212 and Property 2 in p. 214.

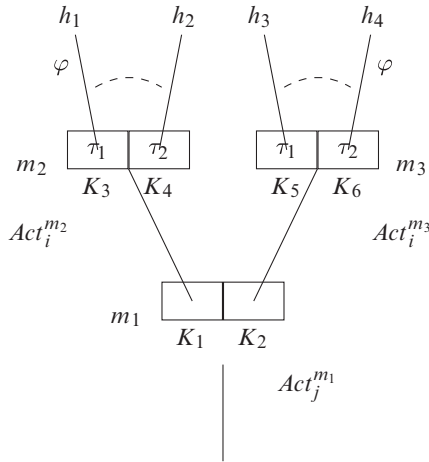


Fig. 4. The labeled stit model that Horty & Pacuit (2017) use to depict the first coin example.

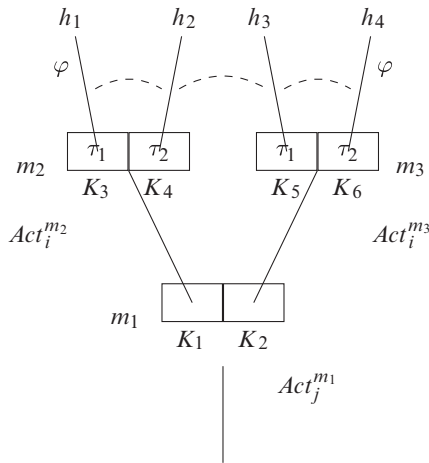


Fig. 5. The labeled stit model that Horty & Pacuit (2017) use to depict the second coin example.

It is helpful to note that the two coin examples of §3 can be represented using the labeled stit models of Figures 4 and 5, respectively. It is straightforward to verify that these labeled stit models give the desired answers regarding epistemic ability. In the first labeled stit model (Figure 4), it is easy to see that $m_2/h_1 \models [i \text{ kstit}] \varphi$, because φ holds at any index that is indistinguishable from m_2/h_1 and where agent i also executes action type τ_1 . Consequently, we can see that $\diamond [i \text{ kstit}] \varphi$ holds at any index that is based on moment m_2 or m_3 . This means that agent i is able, in the epistemic sense, to do φ , regardless of the choice of agent j . Moreover, in the second labeled stit model (Figure 5), it is easy to see that $m_2/h_1 \not\models [i \text{ kstit}] \varphi$, because agent i cannot distinguish m_3/h_3 from m_2/h_1 and executes the same action type at these indices, namely τ_1 . It then follows that $\diamond [i \text{ kstit}] \varphi$ never obtains and that agent i is not able, in the epistemic sense, to do φ , regardless of what agent j chose. The analysis of Horty and Pacuit thus solves the puzzles alluded to in §3.1.

§5. Knowingly doing and the general correspondence result. The source of our disagreement with the analysis of epistemic ability by Horty & Pacuit (2017) is that they end up imposing the semantic condition of settled knowledge on their models. There are two reasons against adopting this semantic condition. First, it is essential for the idea of agency akin to stit theory that alternative histories through a particular moment may witness different actions that the agents can perform. So, given moment/history pair m/h , if an agent cannot distinguish any of the moment/history pairs that emanate from that given moment m , then she does not know anything about the action that she performs at m/h . In particular, the semantic condition of settled knowledge rules out that any agent knows anything about the action she performs. This is one of the main reasons why we reject this assumption and prefer to think of epistemic indistinguishability as a relation between indices, which include histories, rather than between moments.²³

Second, abandoning the assumption of settled knowledge yields the benefit of being able to characterize the notion of *knowingly doing* by a simple combination of knowledge and agency. From the perspective of our formalism, the formula $K_i[i \text{ stit}] \varphi$ expresses that agent i knowingly sees to it that φ (Broersen, 2011a, sec. 3). Conversely, an agent *unknowingly* does φ if and only if (i) she is performing an action that guarantees φ and (ii) she considers it possible that she does not guarantee φ . Jan Broersen writes:

In general the things an agent does unknowingly vastly outnumber the things an agent *knows* it does. For instance, by sending an email, we may enforce many, many things we are not aware of, which are nevertheless the result of sending the email. All these things we do *unknowingly* by knowingly sending the email. (Broersen, 2011a, p. 145)

The concept of knowingly doing can be naturally related to Horty and Pacuit's epistemic notion of agency, which relies on action types. Given a labeled stit model, we can construct a transform epistemic stit model such that epistemic agency in the former corresponds to knowingly doing in the latter. This transformation only changes the indistinguishability relations. For a given agent, two indices are indistinguishable in the transform epistemic stit model if and only if these indices are indistinguishable in the original labeled stit model and the agent executes the same action type at each of these indices. More precisely:

DEFINITION 5.1. (Transform Epistemic Stit Model) *Suppose we are given a labeled stit model $\mathcal{M} = \langle M, H, <, \text{Ags}, (\text{Act}_i^m), (\sim_i), V, \text{Tps}, \text{Lbl}, \text{Exn} \rangle$. Then we can define a transform epistemic stit model $\mathcal{M}' = \langle M', H', <', \text{Ags}', (\bar{\text{Act}}_i^m), (\sim'_i), V' \rangle$ as follows:*

- $M' := M, H' := H, <' := <, \text{Ags}' := \text{Ags}, \text{ and } V' := V;$
- *for every $m \in M$ and every $i \in \text{Ags}$ we define $\bar{\text{Act}}_i^m := \text{Act}_i^m;$*
- *for every i we define $m/h \sim'_i m'/h'$ if and only if $m/h \sim_i m'/h'$ and $\text{Lbl}_i(m/h) = \text{Lbl}_i(m'/h').$ ²⁴*

²³ Herzig & Troquard (2006) were the first to model knowledge in stit theory by using indistinguishability relations between indices.

²⁴ Horty & Pacuit (2017) seem to discuss such a possibility under condition C5 but discard it because it is incompatible with their condition C4, which states that the indistinguishability relations hold between moments rather than indices. Our transform epistemic stit models

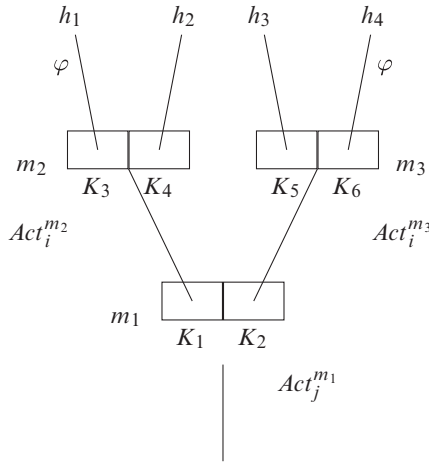


Fig. 6. The transform epistemic stit model of the labeled stit model of the first coin example (Figure 4).

To illustrate this transformation, let us present the transformations of the labeled stit models of the coin examples (see Figures 4 and 5) in Figures 6 and 7, respectively. It is straightforward to verify that these transform epistemic stit models give the desired answers regarding epistemic agency as modelled by the notion of knowingly doing. In the first transform epistemic stit model (Figure 6), note that the information state of agent i at index m_2/h_1 is given by the set $\{m_2/h_1\}$. It is then easy to see that $m_2/h_1 \models K_i[i \text{ stit}]\varphi$. Consequently, we can see that $\diamond K_i[i \text{ stit}]\varphi$ holds at any index that is based on moment m_2 or m_3 . This means that agent i is able to knowingly do φ , regardless of what agent j chose. Moreover, in the second transform epistemic stit model (Figure 7), it is easy to see that $m_2/h_1 \not\models K_i[i \text{ stit}]\varphi$, because agent i cannot distinguish m_3/h_3 from m_2/h_1 and $[i \text{ stit}]\varphi$ is falsified at m_3/h_3 . It then follows that $\diamond K_i[i \text{ stit}]\varphi$ never obtains and that agent i is not able to knowingly do φ , no matter what agent j chooses.

We propose that the simple formula $\diamond K_i[i \text{ stit}]\varphi$ expresses epistemic ability. The above discussion shows that our proposal adequately discriminates between the two coin examples—when they are modelled as in Figures 6 and 7 rather than Figures 2 and 3. In other words, the examples used by Horty and Pacuit point out that epistemic ability can be expressed by a simple combination of impersonal possibility, knowledge, and agency. Moreover, our formalism shows that an analysis of epistemic ability can do without action types.

These last observations regarding the relation between a labeled stit model and its transform epistemic stit model can be generalized. If we limit ourselves to standard stit formulas (without the knowledge operator),²⁵ it follows that the epistemic agency

seem to align with the work of Herzig & Troquard (2006, p. 212—altered notation) who discuss abilities under imperfect knowledge within stit theory and, given an agent i , have used relations R_i , one for each agent i , such that “if two moment/history pairs are linked by a R_i relation, the same type of choice is done by agent i at both indexes”.

²⁵ This restriction is mostly added for simplicity’s sake. We are confident that a more general result for the entire epistemic stit language can be given.

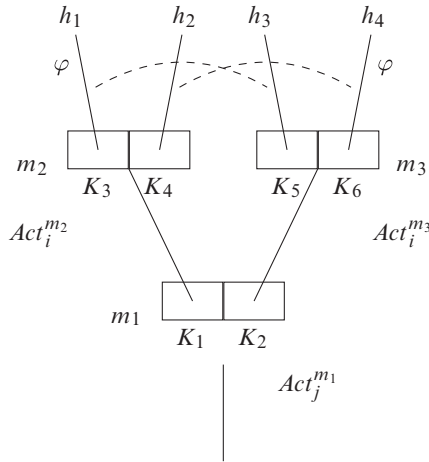


Fig. 7. The transform epistemic stit model of the labeled stit model of the second coin example (Figure 5).

operator in the labeled stit model is supported at an index *only if* the knowingly doing operator is supported at that index in the transform epistemic stit model.

PROPOSITION 5.2. *Let \mathcal{M} be a labeled stit model, and let \mathcal{M}' be the associated transform epistemic stit model. Let $i \in \text{Ags}$ be an agent, $\varphi \in \mathcal{L}_{\text{stit}}$ be a standard stit formula, and m/h be an index. Then*

$$\text{If } \mathcal{M}, m/h \models [i \text{ kstit}] \varphi, \text{ then } \mathcal{M}', m/h \models K_i [i \text{ stit}] \varphi.$$

This observation entails that whenever at a given index in a given labeled stit model it is the case that an agent sees to it that φ in the epistemic sense, then that agent knowingly does φ at that index in the transform epistemic stit model. Moreover, if an agent unknowingly does φ at a given index in the transform epistemic stit model, then that agent does not see to it that φ in the epistemic sense at that index in the labeled stit model. However, it may be that an agent knowingly does φ at a given index in the transform epistemic stit model while that agent does not see to it that φ , in the epistemic sense, at that index in the labeled stit model.²⁶

When we restrict our attention to *static* labeled stit models, then we can expand the previous correspondence result. If we limit ourselves to standard stit formulas (without the knowledge operator), then it follows that the kstit operator in a given static labeled stit model is supported at a given index *if and only if* the knowingly doing operator is supported at that index in the transform epistemic stit model. Moreover, under these assumptions, the knowledge operator K_i in a given static labeled stit model is supported at a given index *if and only if* the combination of the knowledge and historical necessity operators is supported in the transform epistemic stit model.

²⁶ The reader can easily verify that Figure 7 in the work of Horty & Pacuit (2017, p. 632) can be modified (by changing the truth value of A at m_2/h_3) to yield a labeled stit model where this holds for index m_1/h_1 . Another class of models in which this occurs are *label-inconsistent* labeled stit models—see the discussion at the end of this section.

THEOREM 5.3. *Let \mathcal{M} be a static labeled stit model, and let \mathcal{M}' be the associated transform epistemic stit model. Let $i \in \text{Ags}$ be an agent, $\varphi \in \mathcal{L}_{\text{stit}}$ be a standard stit formula, and m/h be an index. Then the following holds*

- (1) $\mathcal{M}, m/h \models [i \text{ kstit}] \varphi$ if and only if $\mathcal{M}', m/h \models K_i [i \text{ stit}] \varphi$;
- (2) $\mathcal{M}, m/h \models K_i \varphi$ if and only if $\mathcal{M}', m/h \models K_i \Box \varphi$.

This result yields several important observations. First, the epistemic notion of ability that Horty & Pacuit (2017) capture by the formula $\Diamond [i \text{ kstit}] \varphi$ corresponds to the formula $\Diamond K_i [i \text{ stit}] \varphi$ in our formalism. So whenever one wants to analyze the epistemic abilities of a given agent in a particular example, then if one adopts Horty and Pacuit's analysis in terms of static labeled stit models using the formula $\Diamond [i \text{ kstit}] \varphi$, then that analysis can be simulated in terms of epistemic stit models using the formula $\Diamond K_i [i \text{ stit}] \varphi$. Notice that the converse may not hold: the jury is out on whether there exist examples where epistemic ability can be analyzed in terms of epistemic stit models using the formula $\Diamond K_i [i \text{ stit}] \varphi$ yet this analysis cannot be simulated in terms of static labeled stit models using the formula $\Diamond [i \text{ kstit}] \varphi$.

Second, our framework captures epistemic ability without any appeal to action types. Horty & Pacuit (2017) argued that stit models need to incorporate action types to address certain puzzles about knowledge and action. We believe their argument is misguided due to their insistence that the indistinguishability relations hold between moments instead of indices. After all, as the above result proves, any static labeled stit model can be transformed into a (non-static) epistemic stit model in such a way that it addresses these puzzles about knowledge and action.

5.1. Two optional conditions on knowledge in stit theory. The transform epistemic stit models, that correspond to a static labeled stit model in light of our transformation, are restricted in important ways. To emphasize these restrictions we present two properties. They can be characterized by axiom schemas or by the corresponding semantic conditions on epistemic stit frames.²⁷ They highlight the flexibility and interpretation of the epistemic stit framework.

PROPOSITION 5.4. *Let \mathcal{F} be an epistemic stit model. Then*

(OAC) *We say that \mathcal{F} satisfies the own-action condition if and only if*

- the schema $K_i \varphi \rightarrow K_i [i \text{ stit}] \varphi$ is valid on \mathcal{F} .
- Or, equivalently, if and only if for all indices $m/h, m'/h'_1$, and m'/h'_2 , if $m/h \sim_i m'/h'_1$ and $h'_2 \in \text{Act}_i^{m'}(h'_1)$, then $m/h \sim_i m'/h'_2$.

(Unif-H) *We say that \mathcal{F} satisfies the uniformity of historical possibility property if and only if*

- the schema $\Diamond K_i \varphi \rightarrow K_i \Diamond \varphi$ is valid on \mathcal{F} .

²⁷ These correspondences can be algorithmically checked using the SQEMA algorithm (Conradie et al., 2006). Broersen (2011a, Propositions 3.1 and 3.2) uses a stit logic about affecting next states, so-called XSTIT, to model knowingly doing and discusses two additional properties: (i) knowledge about next states, which expresses that “agents cannot know more about next states than what is affected by the choices they have” (p. 145) and (ii) effect recollection, which expresses that “the effects of an action that is knowingly performed are known in the next state” (p. 145—it is a dynamic version of perfect recall).

- Or, equivalently, if and only if for all indices $m/h_1, m/h_2$, and m'/h'_1 , if $m/h_1 \sim_i m'/h'_1$, then there is a history h'_2 such that $m/h_2 \sim_i m'/h'_2$.

It may be helpful to note that the own-action condition is equivalent to $\langle i \text{ stit} \rangle \varphi \rightarrow \hat{K}_i \varphi$.²⁸ The own-action condition therefore corresponds to the simple frame condition that for any indices $m/h, m/h'$ if $h' \in \text{Act}_i^m(h)$ then $m/h \sim_i m/h'$. That is, any information state in the partition induced by \sim_i is a union of actions from Act_i .²⁹

The own-action condition expresses the idea that an agent cannot know *more* about the current moment/history pair than what she knows about what she herself brings about. In other words, agents can only know something about the current moment/history pair if it is the result of an action they themselves knowingly perform.

We might say that the own-action condition is a consequence of the assumption that agents cannot know what actions other agents perform concurrently.³⁰ The independence-of-agency condition (see Definition 2.3) guarantees that agents' choices always refine the choices of others. So if an agent knows about the choices of other agents, then she knows more about the future than what is guaranteed by her own choice.

The uniformity of historical possibility property expresses the idea that whenever it is possible for agent i to know φ , then she knows that φ is possible. It is easy to verify that the semantic condition also corresponds to $K_i \Box \varphi \rightarrow \Box K_i \varphi$.³¹ This means that whenever an agent knows that φ is historically settled, then it is historically settled that she knows φ . Or, equivalently, it rules out that agents know that a given proposition is historically settled in case it is possible that she does not know that proposition.

Why are these two conditions important for the analysis of epistemic ability? We can show that the transform epistemic stit models validate both of these conditions:

THEOREM 5.5. *Let \mathcal{M} be a static labeled stit model, and let \mathcal{M}' be the associated transform epistemic stit model. Let $i \in \text{Ags}$ be an agent, $\varphi \in \mathcal{L}_{\text{stit}}$ be a standard stit formula, and m/h be an index. Then the following holds*

$$\begin{aligned} \text{(OAC)} \quad & \mathcal{M}', m/h \models K_i \varphi \rightarrow K_i [i \text{ stit}] \varphi; \\ \text{(Unif-H)} \quad & \mathcal{M}', m/h \models \Diamond K_i \varphi \rightarrow K_i \Diamond \varphi. \end{aligned}$$

It may be helpful to remark that Theorems 5.3 and 5.5 jointly entail two further equivalences:

$$\begin{aligned} (1^*) \quad & \mathcal{M}, m/h \models [i \text{ kstit}] \varphi \quad \text{if and only if} \quad \mathcal{M}', m/h \models K_i \varphi; \\ (2^*) \quad & \mathcal{M}, m/h \models K_i \varphi \quad \text{if and only if} \quad \mathcal{M}', m/h \models \Box K_i \varphi. \end{aligned}$$

²⁸ This follows from the necessitation-rule and the fact that K_i is a transitive normal operator.

²⁹ The own-action condition also plays an important role in the philosophy of intentions. Duijf *et al.* (2019) gave an interpretation of this condition in stit semantics with intentions.

³⁰ It is questionable whether this property also holds under the assumption of common knowledge of rationality, which is standardly assumed in game theory. Under that assumption, it seems appropriate to say that an agent knows that her opponents are avoiding a ridiculous, yet possible, action. This would mean that she considers it impossible for her current action to be compatible with such ridiculous actions of her opponents, yet it is (historically) possible that her opponents perform these actions.

³¹ This follows from the fact that \Box is an S5 operator and K_i is a normal operator.

That is, epistemic agency in a given static labeled stit model is equivalent to knowledge in the transform epistemic stit model; and knowledge in a given static labeled stit model is equivalent to settled knowledge in the transform epistemic stit model.

These equivalences naturally relate to a distinction from economics between information at different stages of decision-making: “an agent’s *ex ante* knowledge is the information available to the agent without taking into account any actions she is currently executing, while the agent’s *ex interim* knowledge is information that does take into account whatever actions the agent is currently executing, along with the effects of these actions” (Horty & Pacuit, 2017, pp. 31–32—emphasis added).³² In light of these equivalences, we submit that the notion of knowledge in our formalism relates to *ex interim* knowledge while the simple combination of impersonal necessity and knowledge relates to *ex ante* knowledge.

Let us finish with briefly mentioning a few ramifications of this result for the analysis of epistemic ability. From the perspective of our formalism, the transform epistemic stit models validate the own-action condition and the uniformity of historical possibility condition. More precisely, this is the case if the initial labeled stit model satisfies the conditions (UAAT) and (Settled Knowledge) (see §4). This means that Horty and Pacuit’s analysis of epistemic ability only applies to epistemic stit models that validate the conditions (OAC) and (Unif-H). Moreover, Proposition 5.4 reveals that if one is interested in the subclass of epistemic stit models that validate one or the other property, then the resulting logic is obtained by adding the respective logical schemata.

It is unclear whether Horty and Pacuit’s analysis applies to *all* models that validate (OAC) and (Unif-H). In analogy with Theorem 5.3, an affirmative answer would require that one be able to transform any epistemic stit model that validates (OAC) and (Unif-H) into a static labeled stit model such that our analysis of epistemic ability, using the formula $\diamond K_i[i \text{ stit}]\varphi$, can be simulated by their analysis of epistemic ability, using the formula $\diamond[i \text{ kstit}]\varphi$.³³

Finally, it is unclear whether Horty and Pacuit’s analysis extends to models that violate any of these conditions. Our formalism is more flexible in that it does not rely on these limitations. In fact, we are unsure whether these conditions are desirable or natural to impose on the models. Although an elaborate discussion of our worries has to be postponed to another occasion, it may be helpful to, once again, note that the own-action condition can be thought of as the requirement that agents do not know *more* about the current moment/history pair than what she knows about what she herself brings about. It, however, seems plausible that there are cases where agents have knowledge about their own future actions or about the actions of other agents.³⁴ Within our formalism, one could hence study whether the characterization of epistemic

³² This distinction was coined by Aumann & Dreze (2008). In comparison, Lorini *et al.* (2014, p. 1314) write: “Each type of knowledge is defined with respect to the time of the agent’s choice: before one’s choice (*ex ante* knowledge), after one’s choice but before knowing the choices of others (*interim* knowledge), and after the choices of all agents have been made public (*ex post* knowledge).”

³³ We do conjecture that this affirmative answer can be given.

³⁴ For example, in Footnote 30 we point out that the own-action condition seems at odds with the standard game-theoretical assumption of common knowledge of rationality. In the next subsection, we will demonstrate that the uniformity of historical possibility property can be thought of as the condition of uniformity of available action types.

ability, by the formula $\diamond K_i[i \text{ stit}]\varphi$, is supported in cases where any of these conditions is dropped.³⁵

5.2. Uniformity of available action types. We would like to close this section by briefly revisiting the semantic condition of uniformity of available action types, abbreviated to: (UAAT) (see §4). The aim is two-fold: (i) we would like to investigate whether there exists an axiom schema that characterizes this semantic condition; and (ii) we would like to study its logical relation to the condition of uniformity of historical possibility (see Proposition 5.4). To address the first issue, we broaden our attention to the class of static labeled stit models that includes models that violate (UAAT). This semantic condition can then be characterized by an axiom schema:³⁶

PROPOSITION 5.6. *Let \mathcal{F} be a static labeled stit frame. Then \mathcal{F} satisfies the semantic condition of uniformity of available action types if and only if the following schema is valid on \mathcal{F} :*

$$\diamond[i \text{ kstit}]\varphi \rightarrow K_i \diamond[i \text{ stit}]\varphi.$$

Let \mathcal{F}' be the transform epistemic stit frame of \mathcal{F} . Then the semantic condition holds at \mathcal{F} only if the following schema is valid on \mathcal{F}' :

$$\diamond K_i[i \text{ stit}]\varphi \rightarrow K_i \diamond[i \text{ stit}]\varphi.$$

This proposition shows that a static labeled stit frame validates the semantic condition (UAAT) only if the formula $\diamond K_i[i \text{ stit}]\varphi \rightarrow K_i \diamond[i \text{ stit}]\varphi$ is valid on the transform epistemic stit frame.³⁷ However, it turns out that there are static labeled stit frames that violate the semantic condition (UAAT) even though that formula is valid in the corresponding transform stit frame. These labeled stit models are rather peculiar. More precisely, let us say that a labeled stit model is *label-consistent* if and only if whenever $m/h \sim_i m'/h'$ then $Tps_i^m \cap Tps_i^{m'} \neq \emptyset$.³⁸ Whenever a labeled stit model

³⁵ It may be helpful to note that Horty & Pacuit (2017, see p. 628 and p. 631) impose condition C1 to ensure that their definition of epistemic agency is well defined and propose that their notion of epistemic agency “makes most sense, and is most useful” under condition C4. It is thus not clear whether Horty and Pacuit’s analysis of epistemic ability can be extended to cases where any of these conditions are violated.

³⁶ In comparison, Horty & Pacuit (2017, p. 629) argue that the semantic condition (UAAT) can be reflected in the object language if we were to introduce special propositions of the form A_i^τ which represent the fact that agent i performs an action of type τ . In particular, every labeled stit model that satisfies (UAAT) would validate the formula $\diamond A_i^\tau \rightarrow K_i \diamond A_i^\tau$. However, Horty and Pacuit do not address the question of whether that schema would *characterize* the semantic condition (UAAT)—which does seem intuitive. In any case, our characterization of the semantic condition (UAAT) goes beyond Horty and Pacuit’s observation in that it does not rely on these special propositions.

³⁷ Broersen (2011a, p. 146) claims that this formula captures the intuition “that exertion of the same choice types should be possible in different states of an information set”.

³⁸ One could ask why we didn’t investigate condition C2 in the work of Horty & Pacuit (2017, pp. 628–629). Condition C2 requires that whenever $m/h \sim_i m'/h'$ then $Lbl_i(m/h) \in Tps_i^{m'}$. It is easy to see that condition C2 entails label-consistency. It is, however, easy to verify that conditions C2 and C4 jointly entail the semantic condition (UAAT). Hence we cannot find a (non-trivial) formula that characterizes the semantic condition (UAAT) for the class of static labeled stit models that validate C2. In contrast, the reader can easily verify that there are static labeled stit models that are label-consistent yet do not satisfy the semantic condition (UAAT).

violates label-consistency, then it can happen that a given agent cannot distinguish two indices even though there is no action type that is available at both of them. In other words, a labeled stit model is label-consistent if and only if whenever a given agent cannot distinguish two indices that emanate from two given moments, there is an action type that she can execute that is consistent with each of these moments.³⁹ As the reader can easily verify, there is a label-inconsistent static labeled stit frame that violates the semantic condition of uniformity of available action types even though the associated transform epistemic stit frame validates the formula $\Diamond K_i[i \text{ stit}]\varphi \rightarrow K_i\Diamond[i \text{ stit}]\varphi$. This is impossible for label-consistent static labeled stit frames:

PROPOSITION 5.7. *Let \mathcal{F} be a label-consistent static labeled stit frame. Let \mathcal{F}' be the transform epistemic stit frame of \mathcal{F} . Then \mathcal{F} satisfies the semantic condition of uniformity of available action types if and only if the following schema is valid on \mathcal{F}' :*

$$\Diamond K_i[i \text{ stit}]\varphi \rightarrow K_i\Diamond[i \text{ stit}]\varphi.$$

Let us finish with addressing the second issue, regarding the relation between (Unif-H) and (UAAT). In light of the above result, let us call the formula $\Diamond K_i[i \text{ stit}]\varphi \rightarrow K_i\Diamond[i \text{ stit}]\varphi$ the (UAAT) schema. It is then interesting to investigate the logical relation between the (UAAT) schema and the schema for the uniformity of historical possibility property (Unif-H). It is fairly easy to see that the validity of the (Unif-H) schema entails the validity of the (UAAT) schema. It may, however, be surprising that the converse also holds:⁴⁰

PROPOSITION 5.8. *Let \mathcal{M} be an epistemic stit model. Then*

1. *If \mathcal{M} validates the schema (Unif-H) $\Diamond K_i\varphi \rightarrow K_i\Diamond\varphi$, then \mathcal{M} validates the schema (UAAT) $\Diamond K_i[i \text{ stit}]\varphi \rightarrow K_i\Diamond[i \text{ stit}]\varphi$;*
2. *If \mathcal{M} validates the schema (UAAT) $\Diamond K_i[i \text{ stit}]\varphi \rightarrow K_i\Diamond[i \text{ stit}]\varphi$, then \mathcal{M} validates the schema (Unif-H) $\Diamond K_i\varphi \rightarrow K_i\Diamond\varphi$.*

The above discussion yields a consideration that speaks in favour of dropping (Unif-H). The above proposition demonstrates that whenever (Unif-H) holds, then this entails that the condition of uniformity of available action types holds. It is standard to distinguish between games of complete information and games of incomplete information. The latter drop the assumption of uniformity of available action types. That is, games of *incomplete* information include scenarios where I may not know which action types are available to myself and others. Since Horty and Pacuit's analysis of epistemic ability (and the treatment of abilities under imperfect information in ATEL) is based on the assumption of uniformity of available action types, it is unclear how this analysis can be applied or extended to abilities under incomplete

³⁹ The technical issue associated with label-inconsistent labeled stit models is the following. Whenever $Tps_i^m \cap Tps_i^{m'} = \emptyset$, then there are no indices m/h and m'/h' such that $m/h \sim'_i m'/h'$. That is, whenever two given moments are label-*inconsistent*, then the transform epistemic indistinguishability relation is such that any indices based on these respective moments are distinguishable. This means that the indistinguishability relation in a given label-*inconsistent* labeled stit model can differ substantially from the one in the transform epistemic stit model.

⁴⁰ Xu (2015, pp. 867–868) discusses similar schemas within the theory of XSTIT and leaves it as an open problem whether the validity of the (UAAT) schema entails the validity of the (Unif-H) schema (he calls these schemas (U-s) and (K-s), respectively).

information. From the perspective of our formalism, this yields the question of whether $\Diamond K_i[i \text{ stit}]\varphi$ adequately expresses that an agent is able to φ in the epistemic sense in epistemic stit models that violate (Unif-H). A detailed analysis of abilities under incomplete information within stit theory, and the particular question we raised, must be postponed to another occasion.

§6. Discussion. We have shown that analyzes of epistemic abilities can do without action types. We propose that the epistemic ability of an individual is characterized by a combination of impersonal possibility, knowledge, and agency of the form $\Diamond K_i[i \text{ stit}]\varphi$, which is to be read as saying that it is possible that agent i knowingly does φ . Our primary focus has been on some of the examples discussed by Horty & Pacuit (2017, Figures 3 and 4), and we justified our proposal by showing that it solves these puzzles not by adding action types but instead by amending the models in a systematic way. Moreover, we have presented a general correspondence result that shows that any labeled stit model—which includes action types—can be transformed into a standard epistemic stit model—which eschews action types—such that epistemic agency in the former corresponds to knowingly doing in the latter. This means that the analysis of epistemic agency and ability by Horty & Pacuit (2017) can be simulated within epistemic stit theory without adding action types.

Although our focus has been on the labeled stit models that were recently introduced by Horty & Pacuit (2017), we firmly believe that our result transfers naturally to the standard treatment of abilities under imperfect information, as studied within the ATL tradition. Hence, the analysis of ability under imperfect information in the ATL tradition can be simulated in epistemic stit theory. Moreover, the central ability operator in the ATEL tradition can then be characterized by the stit-formula $\Diamond K_i[i \text{ stit}]\varphi$. This means that our result is of importance to agent-based artificial intelligence techniques that rely on logical methods to study cooperation and coalitional abilities in a precise way (Wooldridge & Jennings, 1995).

It is important to remark that the analysis of epistemic agency in the labeled stit models presented by Horty & Pacuit (2017) is subject to several constraints. We have discussed the condition of uniformity of available action types and the condition that epistemic indistinguishability can be thought of as a relation between moments rather than indices. The epistemic stit models used in our formalism are more general and therefore provide a more general analysis of the epistemic sense of agency and ability in terms of the notion of knowingly doing.⁴¹

We would like to briefly reflect on the role played by action types in the analysis of epistemic ability. One of the motivations for the introduction of action types can be summarized as follows: in the standard stit formalism

⁴¹ One interesting research question is to investigate an epistemic sense of *deliberative* stit (see Horty & Belnap, 1995). The standard notion of deliberative stit includes a negative condition and can be expressed in our formalism by the formula $[i \text{ stit}]\varphi \wedge \Diamond \neg\varphi$, which is to be read as saying that agent i sees to it that φ while it is possible that φ does not obtain. With regard to an epistemic sense of deliberative stit, different negative conditions can be expressed in our framework using the knowledge and impersonal possibility operators (see also Broersen, 2011a, definition 4.2) (i) $\Diamond \neg\varphi$, which means that it is actually possible that φ does not obtain; (ii) $K_i \Diamond \neg\varphi$, which means that agent i knows that it is actually possible that φ does not obtain; (iii) $\hat{K}_i \Diamond \neg\varphi$, which means that agent i considers it possible that it is actually possible that φ does not obtain. (Note that (ii) entails (i), which entails (iii).)

there were no general, repeatable kinds of actions, or action types; there were only particular [action tokens], with nothing to group them together as actions of the same kind. (Horty & Pacuit, 2017, p. 617)

A given action type τ hence groups together several particular actions of that kind. Alternatively, one may say that any *set* of actions can be considered to be an action type. This would entail a rather uninformative picture in which what groups a set of particular actions together is that they are grouped together. But this uninformative picture is no less informative than the picture delivered by introducing action types. On that view, what groups several particular actions together is that they have the same label. The introduction of labels for action types hence merely presents the illusion of informativeness, unless we explicate why these actions deserve the same label.

It thus seems vital to ask what action types *are*. Phrased differently, we must ask what makes it the case that a given set of actions represents the same action type. From the perspective of our formalism, our central correspondence result shows that an action type is available to an agent *only if* it is possible that she knowingly does it, that is, only if she is able, in the epistemic sense, to perform an action of that kind.⁴² In other words, an action type is available to a given agent only if she is able, in an epistemic sense, to perform it. Although we have shown that this notion of an action type is dispensable for analyzes of epistemic ability, there might be other ideas about what action types are or why we need them.⁴³ The jury is out on whether action types are essential for other such usages.

We conclude this paper with two remarks on why we think that it might be beneficial to do without action types—one is more formal, the other conceptual. First, the introduction of action types introduces unnecessary complications into the truth conditions for epistemic agency. Our alternative proposal is to model epistemic agency by a combination of two standard S5 operators, of the form $K_i[i \text{ stit}]\varphi$. Since S5 modalities and standard possible-world semantics are well studied, achieving various meta-logical results seems within reach. As an illustration of this, we present a completeness result for multi-agent epistemic stit logic:⁴⁴

PROPOSITION 6.1. (Completeness Epistemic Stit) *We define the multi-agent epistemic stit language $\mathcal{L}_{\text{estit}}$ as follows:*

$$\varphi := q \mid \neg\varphi \mid \varphi \wedge \varphi \mid \Box\varphi \mid [i \text{ stit}]\varphi \mid K_i\varphi,$$

⁴² One of the reasons why this is not an equivalence is that Horty & Pacuit (2017, p. 626) assume that “all action types are primitive”. After all, there might be things of which it is possible that I knowingly do them even though doing so is not a primitive action type.

⁴³ For example, Horty & Pacuit (2017, p. 626) write that “action types are repeatable” and “general kinds of action”; as a typical example they mention the action type of “opening a window” (p. 617). This way of thinking about action types, however, sheds some doubt on some of their natural conditions. For example, is the condition of uniformity of available action types intuitive when we think of action types in this way? To justify this condition, one would need to argue that an agent cannot be unsure of whether this particular action type is available to her if it is, in fact, available to her. However, there are plenty of cases where an agent lacks the epistemic ability to open a given window (e.g., because she is unsure whether it is locked), even though she does have the action type of “opening a window” available to her (since the window is in fact not locked).

⁴⁴ We refer the interested reader to the appendix for more details.

where q ranges over a given countable set of propositions \mathcal{P} and i ranges over a given finite set of agents Ags .

Then there is a normal modal logic, axiomatized by a finite set of axioms and the inference rules *modus ponens* and *necessitation*, that is sound and complete with respect to the class of epistemic stit models.

Second, one of the benefits of standard (non-epistemic) stit theory is that it is easy to generalize individual agency to group agency. Suppose we were to extend the analysis of individual abilities under uncertainty that relies on action types to group abilities under uncertainty. To do so, we need to start with an analysis of what the available (primitive) group action types are.⁴⁵ In fact, this conceptual problem is overlooked by the literature on knowledge and group action, and we do not know of any work that addresses this issue.⁴⁶ Alternatively, suppose we were to start from the notion of individual abilities under uncertainty that does not rely on action types. We could then take advantage of the fact that stit theory naturally generalizes individual agency to group agency. Moreover, we can use various notions of group knowledge that are well established in the literature on multi-agent epistemic logic—including mutual, distributive, and common knowledge. The details of such an analysis of group abilities under uncertainty are left for future work.

Acknowledgments We thank the audiences at symposia in Amsterdam and Utrecht. Special thanks go to Jeff Horty and Allard Tamminga for inspiring discussions. We gratefully acknowledge funding from the European Research Council (the ERC-2013-CoG project REINS, no. 616512, to H. Duijf, J. Broersen and A.I. Ramírez Abarca; and the ERC-2017-CoG project SEA, no. 771074, to H. Duijf) and from the Dutch Research Council (NWO, the project DigTEp, no. 652.001.004, to J. Broersen and A. Kuncová).

§A. Appendix: Proofs.

PROPOSITION 4.2. *Let \mathcal{F} be an epistemic stit frame. Then, \mathcal{F} satisfies the settled knowledge condition if and only if*

$$\mathcal{F} \models K_i \varphi \rightarrow K_i \Box \varphi.$$

Proof. This correspondence can be checked using the algorithm SQEMA (Conradie *et al.*, 2006). □

⁴⁵ Note that it is already hard to say what the primitive action types of an individual agent are. Can you mention any *primitive* action type available to you at this very moment? This challenge is not meant to be decisive but only suggestive in undermining the idea that this is a simple task.

⁴⁶ The analysis of group abilities under uncertainty within the ATL tradition typically seems to take for granted that for a given group of agents, any tuple of individual action types, one for each group member, is a group action type, and vice versa (see Ågotnes *et al.*, 2015, for an excellent overview). By contrast, the outline of a theory of action by Broersen (2014, p. 50) perhaps comes close, although it relies on agency and *event types* rather than action types. His framework “allows for events of a given type for which it takes strictly more than one agent to perform them”.

PROPOSITION 5.2. *Let \mathcal{M} be a labeled stit model, and let \mathcal{M}' be the associated transform epistemic stit model. Let $i \in \text{Ags}$ be an agent, $\varphi \in \mathcal{L}_{stit}$ be a standard stit formula, and m/h be an index. Then*

If $\mathcal{M}, m/h \models [i \text{ kstit}] \varphi$, then $\mathcal{M}', m/h \models K_i [i \text{ stit}] \varphi$.

Proof. Let $\mathcal{M} = \langle M, H, <, \text{Ags}, (\text{Act}_i^m), (\sim_i), V, \text{Tps}, \text{Lbl}, \text{Exn} \rangle$ be a labeled stit model and let \mathcal{M}' denote the transform epistemic stit model of \mathcal{M} . Let $i \in \text{Ags}$ be an agent, let $\varphi \in \mathcal{L}_{stit}$ be a formula, and let $m/h \in \text{Ind}$ be an index.

First note that it is immediately clear that every index m/h supports exactly the same standard stit formulas in \mathcal{M} and \mathcal{M}' , because these models only differ with regard to the epistemic indistinguishability relations.

Suppose $\mathcal{M}, m/h \models [i \text{ kstit}] \varphi$. Take any $m''/h'' \in \text{Ind}$ such that there is an $m'/h' \in \text{Ind}$ satisfying $m'/h' \sim'_i m/h$ and $m''/h'' \in \text{Act}_i(m'/h')$. Note that this entails that $m'' = m'$. Then by definition of \mathcal{M}' we get $\text{Lbl}_i(m/h) = \text{Lbl}_i(m'/h')$ and $m'/h' \sim_i m/h$. Moreover, since $\text{Act}_i(m'/h') = \text{Act}_i(m'/h'')$, we have $\text{Lbl}_i(m'/h'') = \text{Lbl}_i(m/h)$. Hence, by presupposition we get $\mathcal{M}, m'/h'' \models \varphi$. It follows that $\mathcal{M}', m''/h'' \models \varphi$. Since m''/h'' was arbitrary, this proves that $\mathcal{M}', m/h \models K_i [i \text{ stit}] \varphi$. \square

THEOREM 5.3. *Let \mathcal{M} be a static labeled stit model, and let \mathcal{M}' be the associated transform epistemic stit model. Let $i \in \text{Ags}$ be an agent, $\varphi \in \mathcal{L}_{stit}$ be a standard stit formula, and m/h be an index. Then the following holds*

- (1) $\mathcal{M}, m/h \models [i \text{ kstit}] \varphi$ if and only if $\mathcal{M}', m/h \models K_i [i \text{ stit}] \varphi$;
- (2) $\mathcal{M}, m/h \models K_i \varphi$ if and only if $\mathcal{M}', m/h \models K_i \square \varphi$.

Proof. Let $\mathcal{M} = \langle M, H, <, \text{Ags}, (\text{Act}_i^m), (\sim_i), V, \text{Tps}, \text{Lbl}, \text{Exn} \rangle$ be a labeled stit model that respects the condition of settled knowledge and let \mathcal{M}' denote the transform epistemic stit model of \mathcal{M} . Let $\varphi \in \mathcal{L}_{stit}$ be a formula, and let $m/h \in \text{Ind}$ be an index.

First note that it is immediately clear that every index m/h supports exactly the same standard stit formulas in \mathcal{M} and \mathcal{M}' , because these models only differ with regard to the epistemic indistinguishability relations.

(1. \Rightarrow) See Proposition 5.2.

(1. \Leftarrow) Suppose $\mathcal{M}', m/h \models K_i [i \text{ stit}] \varphi$. Take any m'/h' such that $m'/h' \sim_i m/h$ and take any $m''/h'' \in \text{Ind}$ such that $\text{Lbl}_i(m/h) = \text{Lbl}_i(m'/h'')$. Since \mathcal{M} respects the condition of settled knowledge, we get that $m'/h'' \sim_i m/h$. By definition of \mathcal{M}' it follows that $m'/h'' \sim'_i m/h$, and therefore we get $\mathcal{M}', m'/h'' \models [i \text{ stit}] \varphi$. In particular, $\mathcal{M}', m'/h'' \models \varphi$. Therefore $\mathcal{M}, m'/h'' \models \varphi$. Since m'/h'' was arbitrary, this proves that $\mathcal{M}, m/h \models [i \text{ kstit}] \varphi$.

(2. \Rightarrow) Suppose $\mathcal{M}, m/h \models K_i \varphi$. Take any $m''/h'' \in \text{Ind}$ such that there is an $m'/h' \in \text{Ind}$ satisfying $m'/h' \sim'_i m/h$ and $m'' = m'$. Then by definition of \mathcal{M}' we get $\text{Lbl}_i(m/h) = \text{Lbl}_i(m'/h')$ and $m'/h' \sim_i m/h$. Moreover, since \mathcal{M} respects the condition of settled knowledge, we get $m'/h'' \sim_i m/h$. Hence, by presupposition we get $\mathcal{M}, m'/h'' \models \varphi$. It follows that $\mathcal{M}', m''/h'' \models \varphi$. Since m''/h'' was arbitrary, this proves that $\mathcal{M}', m/h \models K_i \square \varphi$.

(2. \Leftarrow) Suppose $\mathcal{M}', m/h \models K_i \square \varphi$. Take any m'/h' such that $m'/h' \sim_i m/h$. Since \mathcal{M} respects the condition of uniformity of available action types, there is an index $m''/h'' \in \text{Ind}$ such that $\text{Lbl}_i(m/h) = \text{Lbl}_i(m'/h'')$. Since \mathcal{M} respects the condition of settled knowledge, we get that $m'/h'' \sim_i m/h$. By definition of \mathcal{M}' we see that

$m/h \sim'_i m'/h''$. By presupposition we get $\mathcal{M}', m'/h'' \models \Box\varphi$, which entails $\mathcal{M}', m'/h' \models \varphi$. Hence $\mathcal{M}, m'/h' \models \varphi$. Since m'/h' was arbitrary, this entails that $\mathcal{M}, m/h \models K_i\varphi$. \square

PROPOSITION 5.4. *Let \mathcal{F} be an epistemic stit model. Then*

(OAC) *We say that \mathcal{F} satisfies the own-action condition if and only if*

- *the schema $K_i\varphi \rightarrow K_i[i \text{ stit}]\varphi$ is valid on \mathcal{F} .*
- *Or, equivalently, if and only if for all indices $m/h, m'/h'_1$, and m'/h'_2 , if $m/h \sim_i m'/h'_1$ and $h'_2 \in Act_i^{m'}(h'_1)$, then $m/h \sim_i m'/h'_2$.*

(Unif-H) *We say that \mathcal{F} satisfies the uniformity of historical possibility property if and only if*

- *the schema $\Diamond K_i\varphi \rightarrow K_i\Diamond\varphi$ is valid on \mathcal{F} .*
- *Or, equivalently, if and only if for all indices $m/h_1, m/h_2$, and m'/h'_1 , if $m/h_1 \sim_i m'/h'_1$, then there is a history h'_2 such that $m/h_2 \sim_i m'/h'_2$.*

Proof. These correspondences can be checked using the algorithm SQEMA (Conradie et al., 2006). \square

THEOREM 5.5. *Let \mathcal{M} be a static labeled stit model, and let \mathcal{M}' be the associated transform epistemic stit model. Let $i \in \text{Ags}$ be an agent, $\varphi \in \mathcal{L}_{\text{stit}}$ be a standard stit formula, and m/h be an index. Then the following holds*

$$\begin{aligned} \text{(OAC)} \quad & \mathcal{M}', m/h \models K_i\varphi \rightarrow K_i[i \text{ stit}]\varphi; \\ \text{(Unif-H)} \quad & \mathcal{M}', m/h \models \Diamond K_i\varphi \rightarrow K_i\Diamond\varphi. \end{aligned}$$

Proof. Let $\mathcal{M} = \langle M, H, <, \text{Ags}, (Act_i^m), (\sim_i), V, Tps, Lbl, Exn \rangle$ be a labeled stit model that respects the condition of settled knowledge and let \mathcal{M}' denote the transform epistemic stit model of \mathcal{M} . Let $\varphi \in \mathcal{L}_{\text{stit}}$ be a formula, and let $m/h \in \text{Ind}$ be an index.

(OAC) Suppose $\mathcal{M}', m/h \models K_i\varphi$. Take any m'/h'' such that there is a m'/h' satisfying $m/h \sim'_i m'/h'$ and $m'/h'' \in Act_i(m'/h')$. By definition of \mathcal{M}' it follows that $m/h \sim_i m'/h'$ and $Lbl_i(m/h) = Lbl_i(m'/h')$. Since \mathcal{M} respects the condition of settled knowledge, we get that $m/h \sim_i m'/h''$. Moreover, it follows that $Lbl_i(m'/h') = Lbl_i(m'/h'')$. Hence, by definition of \mathcal{M}' , we see that $m/h \sim'_i m'/h''$. By presupposition we get $\mathcal{M}', m'/h'' \models \varphi$. Since m'/h'' was arbitrary, this entails that $\mathcal{M}', m/h \models K_i[i \text{ stit}]\varphi$.

(Unif-H) Suppose $\mathcal{M}', m/h \models \Diamond K_i\varphi$. Take any m_1/h_1 such that $m_1/h_1 \sim'_i m/h$. We show that $\mathcal{M}', m_1/h_1 \models \Diamond\varphi$. By definition of \mathcal{M}' we get $m_1/h_1 \sim_i m/h$ and $Lbl_i(m_1/h_1) = Lbl_i(m/h)$. By presupposition there is an index m/h' such that $\mathcal{M}', m/h' \models K_i\varphi$. Since \mathcal{M} respects the condition of uniformity of available action types, there is an m_1/h'_1 such that $Lbl_i(m_1/h'_1) = Lbl_i(m/h')$. Since \mathcal{M} respects the condition of settled knowledge, we get that $m_1/h'_1 \sim_i m/h'$. By definition of \mathcal{M}' we get $m_1/h'_1 \sim'_i m/h'$. Since $\mathcal{M}', m/h' \models K_i\varphi$, this entails $\mathcal{M}', m_1/h'_1 \models \varphi$. Hence $\mathcal{M}', m_1/h_1 \models \Diamond\varphi$. Since m_1/h_1 was arbitrary, this entails that $\mathcal{M}', m/h \models K_i\Diamond\varphi$. \square

PROPOSITION 5.6. *Let \mathcal{F} be a static labeled stit frame. Then \mathcal{F} satisfies the semantic condition of uniformity of available action types if and only if the following schema is valid on \mathcal{F} :*

$$\Diamond[i \text{ kstit}]\varphi \rightarrow K_i\Diamond[i \text{ stit}]\varphi.$$

Let \mathcal{F}' be the transform epistemic stit frame of \mathcal{F} . Then the semantic condition holds at \mathcal{F} only if the following schema is valid on \mathcal{F}' :

$$\diamond K_i[i \text{ stit}]\varphi \rightarrow K_i \diamond [i \text{ stit}]\varphi.$$

Proof. Let $\mathcal{F} = \langle M, H, <, \text{Ags}, (\text{Act}_i^m), (\sim_i), \text{Tps}, \text{Lbl}, \text{Exn} \rangle$ be a static labeled stit frame, and let \mathcal{F}' be the associated transform epistemic stit frame.

(1. \Rightarrow) Suppose the semantic condition (UAAT) holds on \mathcal{F} . We need to show that for any valuation V the model $\mathcal{M} = \langle \mathcal{F}, V \rangle$ based on \mathcal{F} satisfies $\mathcal{M} \models \diamond [i \text{ kstit}]\varphi \rightarrow K_i \diamond [i \text{ stit}]\varphi$. Let $V : \mathcal{P} \rightarrow 2^{\text{Ind}}$ be a valuation and let \mathcal{M} be the static labeled stit model that consists of \mathcal{F} and V . Let $m/h \in \text{Ind}$ be any index. Assume $\mathcal{M}, m/h \models \diamond [i \text{ kstit}]\varphi$. This means that there exists an index m/h_1 such that $\mathcal{M}, m/h_1 \models [i \text{ kstit}]\varphi$. It follows that for any $m'/h'_1 \in \text{Ind}$ such that $m/h_1 \sim_i m'/h'_1$ and any m'/h''_1 such that $\text{Lbl}_i(m/h_1) = \text{Lbl}_i(m'/h''_1)$ we have $\mathcal{M}, m'/h''_1 \models \varphi$. Take any $m'/h' \in \text{Ind}$ such that $m/h \sim_i m'/h'$. We need to show that $\mathcal{M}, m'/h' \models \diamond [i \text{ stit}]\varphi$. Because \mathcal{F} is static and $m/h \sim_i m'/h'$, it follows that $m/h_1 \sim_i m'/h'$. Because the semantic condition (UAAT) holds on \mathcal{F} , there is an index m'/h'' such that $\text{Lbl}_i(m/h_1) = \text{Lbl}_i(m'/h'')$. We show that $\mathcal{M}, m'/h'' \models [i \text{ stit}]\varphi$. Take any $m'/h^* \in \text{Act}_i(m'/h'')$. Note that $m/h_1 \sim_i m'/h' \sim_i m'/h^*$. By definition of Lbl_i , it also follows that $\text{Lbl}_i(m', h^*) = \text{Lbl}_i(m'/h'') = \text{Lbl}_i(m/h_1)$. Since $\mathcal{M}, m/h_1 \models [i \text{ kstit}]\varphi$, we get $\mathcal{M}, m'/h^* \models \varphi$. Hence, it follows that $\mathcal{M}, m'/h'' \models [i \text{ stit}]\varphi$ and therefore $\mathcal{M}, m'/h' \models \diamond [i \text{ stit}]\varphi$. Since m'/h' was arbitrary, this entails that $\mathcal{M}, m/h \models K_i \diamond [i \text{ stit}]\varphi$, as desired.

(1. \Leftarrow) Suppose the schema $\diamond [i \text{ kstit}]\varphi \rightarrow K_i \diamond [i \text{ stit}]\varphi$ is valid on \mathcal{F} . Take any m/h and m'/h' in Ind such that $m/h \sim_i m'/h'$. Take any $m/h_1 \in \text{Ind}$. To prove that the semantic condition (UAAT) holds at \mathcal{F} , we need to show that $\text{Lbl}_i(m/h_1) \in \text{Tps}_i^{m'}$. Define $V(q) := \{m^*/h^* \mid \text{Lbl}_i(m/h_1) = \text{Lbl}_i(m^*/h^*)\}$, and let \mathcal{M} be the static labeled stit model that consists of \mathcal{F} and V . It then follows that $\mathcal{M}, m/h_1 \models [i \text{ kstit}]q$, by definition of V . Hence $\mathcal{M}, m/h \models \diamond [i \text{ kstit}]q$. Because the schema is assumed to be valid on \mathcal{F} , we get $\mathcal{M}, m/h \models K_i \diamond [i \text{ stit}]q$. In particular, it follows that $\mathcal{M}, m'/h' \models \diamond [i \text{ stit}]q$. Hence, there is an index m'/h'' such that $\mathcal{M}, m'/h'' \models [i \text{ stit}]q$. In particular, $\mathcal{M}, m'/h'' \models q$. By definition of V , it follows that $\text{Lbl}_i(m/h_1) = \text{Lbl}_i(m'/h'') \in \text{Tps}_i^{m'}$, as desired.

(2.) Note that Theorem 5.5. entails that if the static labeled stit frame \mathcal{F} validates the semantic condition (UAAT), then the schema $\diamond [i \text{ kstit}]\varphi \rightarrow K_i \diamond [i \text{ stit}]\varphi$ is valid on \mathcal{F} if and only if the schema $\diamond K_i [i \text{ stit}]\varphi \rightarrow K_i \diamond [i \text{ stit}]\varphi$ is valid on \mathcal{F}' . Hence, if the semantic condition (UAAT) holds at \mathcal{F} , then the schema $\diamond K_i [i \text{ stit}]\varphi \rightarrow K_i \diamond [i \text{ stit}]\varphi$ is valid on \mathcal{F}' . \square

PROPOSITION 5.7. Let \mathcal{F} be a label-consistent static labeled stit frame. Let \mathcal{F}' be the transform epistemic stit frame of \mathcal{F} . Then \mathcal{F} satisfies the semantic condition of uniformity of available action types if and only if the following schema is valid on \mathcal{F}' :

$$\diamond K_i [i \text{ stit}]\varphi \rightarrow K_i \diamond [i \text{ stit}]\varphi.$$

Proof. In light of Proposition 5.6, we only need to prove the right-to-left direction.

(\Leftarrow) Suppose the schema $\diamond K_i [i \text{ stit}]\varphi \rightarrow K_i \diamond [i \text{ stit}]\varphi$ is valid on \mathcal{F}' . Take any m/h and m'/h' in Ind such that $m/h \sim_i m'/h'$. Take any $m/h_2 \in \text{Ind}$. To prove that the semantic condition (UAAT) holds at \mathcal{F} , we need to show that $\text{Lbl}_i(m/h_2) \in \text{Tps}_i^{m'}$. Define $V'(q) := \{m^*/h^* \mid \text{Lbl}_i(m/h_2) = \text{Lbl}_i(m^*/h^*)\}$, and let \mathcal{M}' be the epistemic stit model that consists of \mathcal{F}' and V' . By definition of Lbl_i , this entails $\mathcal{M}' \models q \Leftrightarrow [i \text{ stit}]q$.

It then follows that $\mathcal{M}', m/h_2 \models K_i[i \text{ stit}]q$, by definition of V' and by definition of \sim'_i . Since \mathcal{F} is label-consistent, there are m/h_1 and m'/h'_1 such that $Lbl_i(m/h_1) = Lbl_i(m'/h'_1)$. Since \mathcal{F} is static and $m/h \sim_i m'/h'$, we get $m/h_1 \sim_i m'/h'_1$. By definition of \sim'_i , it follows that $m/h_1 \sim'_i m'/h'_1$. Since $\mathcal{M}', m/h_2 \models K_i[i \text{ stit}]q$, it follows that $\mathcal{M}', m/h_1 \models \diamond K_i[i \text{ stit}]q$. Because the schema $\diamond K_i[i \text{ stit}]\varphi \rightarrow K_i \diamond [i \text{ stit}]\varphi$ is assumed to be valid on \mathcal{F}' , we get $\mathcal{M}', m/h \models K_i \diamond [i \text{ stit}]q$. In particular, because $m/h_1 \sim'_i m'/h'_1$, it follows that $\mathcal{M}', m'/h'_1 \models \diamond [i \text{ stit}]q$. Hence, there is an index m'/h'' such that $\mathcal{M}, m'/h'' \models [i \text{ stit}]q$. In particular, $\mathcal{M}, m'/h'' \models q$. By definition of V' , it follows that $Lbl_i(m/h_2) = Lbl_i(m'/h'') \in Tps_i^{m'}$, as desired. \square

PROPOSITION 5.8. *Let \mathcal{M} be an epistemic stit model. Then*

1. *If \mathcal{M} validates the schema (Unif-H) $\diamond K_i \varphi \rightarrow K_i \diamond \varphi$, then \mathcal{M} validates the schema (UAAT) $\diamond K_i [i \text{ stit}]\varphi \rightarrow K_i \diamond [i \text{ stit}]\varphi$;*
2. *If \mathcal{M} validates the schema (UAAT) $\diamond K_i [i \text{ stit}]\varphi \rightarrow K_i \diamond [i \text{ stit}]\varphi$, then \mathcal{M} validates the schema (Unif-H) $\diamond K_i \varphi \rightarrow K_i \diamond \varphi$.*

Proof. 1. Follows from substituting $[i \text{ stit}]\varphi$ for φ .

2. Note that the validity of the (UAAT) schema entails the validity of the schema $\diamond K_i \langle i \text{ stit} \rangle \varphi \rightarrow K_i \diamond \langle i \text{ stit} \rangle \varphi$ (replace φ with $\langle i \text{ stit} \rangle \varphi$ and note that $\models [i \text{ stit}]\langle i \text{ stit} \rangle \varphi \leftrightarrow \langle i \text{ stit} \rangle \varphi$). Since $\models \varphi \rightarrow \langle i \text{ stit} \rangle \varphi$, it follows that $\mathcal{M} \models \diamond K_i \varphi \rightarrow K_i \diamond \langle i \text{ stit} \rangle \varphi$. Since $\models \langle i \text{ stit} \rangle \varphi \rightarrow \diamond \varphi$ and $\models \diamond \diamond \varphi \rightarrow \diamond \varphi$, we get $\mathcal{M} \models \diamond K_i \varphi \rightarrow K_i \diamond \varphi$, as desired. \square

PROPOSITION 6.1. (Completeness Epistemic Stit) *We define the multi-agent epistemic stit language $\mathcal{L}_{\text{stit}}$ as follows:*

$$\varphi ::= q \mid \neg \varphi \mid \varphi \wedge \varphi \mid \square \varphi \mid [i \text{ stit}]\varphi \mid K_i \varphi,$$

where q ranges over a given countable set of propositions \mathcal{P} and i ranges over a given finite set of agents Ags .

Then there is a normal modal logic, axiomatized by a finite set of axioms and the inference rules *modus ponens* and *necessitation*, that is sound and complete with respect to the class of epistemic stit models.

Proof. The following axiom schemas, in combination with the rules and axiom schemas for normal modal operators, provide a sound and complete Hilbert-style system for the validities on epistemic stit models:

(S5 Historical Necessity)	S5 for \square
(S5 Agency)	for each agent $i \in \text{Ags}$: S5 for $[i \text{ stit}]$
(Independence of Agency)	for any formulas φ_i , one for each $i \in \text{Ags}$ it holds that $\bigwedge_{i \in \text{Ags}} \diamond [i \text{ stit}]\varphi_i \rightarrow \diamond (\bigwedge_{i \in \text{Ags}} [i \text{ stit}]\varphi_i)$
(S5 Knowledge)	for each agent $i \in \text{Ags}$: S5 for K_i

The epistemic stit logic is a so-called fusion of epistemic logic and non-epistemic stit theory. The complete logical system for epistemic stit theory is therefore given by the simple combination of the logical systems for multi-agent epistemic logic and (non-epistemic) stit logic.

It is well established that the non-epistemic fragment is complete with respect to stit models (see, e.g., Belnap *et al.*, 2001), and the same goes for the epistemic fragment (see, e.g., Meyer & van der Hoek, 1995).⁴⁷ □

BIBLIOGRAPHY

- Ågotnes, T., Goranko, V., Jamroga, W., & Wooldridge, M. (2015). Knowledge and ability. In van Ditmarsch, H., Halpern, J. Y., van der Hoek, W., and Kooi, B., editors. *Handbook of Epistemic Logic*. London, UK: College Publications, pp. 543–589.
- Alur, R., Henzinger, T. A., & Kupferman, O. (2002). Alternating-time temporal logic. *Journal of the ACM*, **49**(5), 672–713.
- Aumann, R. J. (1999). Interactive epistemology. I: Knowledge. *International Journal of Game Theory*, **28**(3), 263–300.
- Aumann, R. J., & Dreze, J. H. (2008). Rational expectations in games. *The American Economic Review*, **98**(1), 72–86.
- Belnap, N., & Perloff, M. (1988). Seeing to it that: A canonical form for agentives. *Theoria*, **54**(3), 175–199.
- Belnap, N., Perloff, M., & Xu, M. (2001). *Facing the Future. Agents and Choices in Our Indeterminist World*. Oxford, UK: Oxford University Press.
- Broersen, J. (2008). A logical analysis of the interaction between ‘obligation-to-do’ and ‘knowingly doing’. In van der Meyden, R., and van der Torre, L., editors. *Deontic Logic in Computer Science*. Lecture Notes in Computer Science, Vol. 5076. Berlin, Heidelberg/Germany: Springer, pp. 140–154.
- . (2009). A stit-logic for extensive form group strategies. In Boldi, P., Vizzari, G., Pasi, G., and Baeza-Yates, R., editors. *Proceedings of the International Joint Conference on Web Intelligence and Intelligent Agent Technology*, Vol. 3. Washington, DC: IEEE Computer Society, pp. 484–487.
- . (2011a). Deontic epistemic stit logic distinguishing modes of mens rea. *Journal of Applied Logic*, **9**(2), 137–152.
- . (2011b). Making a start with the stit logic analysis of intentional action. *Journal of Philosophical Logic*, **40**(4), 499–530.
- . (2014). On the reconciliation of logics of agency and logics of event types. In Trypuz, R., editor. *Krister Segerberg on Logic of Actions*. Netherlands: Springer, pp. 41–59.
- Broersen, J., and Herzig, A. (2015). Using STIT theory to talk about strategies. In van Benthem, J., Ghosh, S., and Verbrugge, R., editors. *Models of Strategic Reasoning*. New York, NY: Springer, pp. 137–173.
- Broersen, J., Herzig, A., and Troquard, N. (2006a). Embedding alternating-time temporal logic in strategic logic of agency. *Journal of Logic and Computation*, **16**(5), 559–578.
- . (2006b). A stit-extension of ATL. In Fisher, M., van der Hoek, W., Konev, B., and Lisitsa, A., editors. *Logics in Artificial Intelligence*. Lecture Notes in Computer Science, Vol. 4160. Berlin, Heidelberg/Germany: Springer, pp. 69–81.

⁴⁷ Broersen (2008, Theorem 1) proves a similar result for epistemic XSTIT, the difference is that his logic concerns XSTIT.

- . (2007). A normal simulation of coalition logic and an epistemic extension. In Samet, D., editor. *Proceedings of the Eleventh Conference on Theoretical Aspects of Rationality and Knowledge*. New York, NY: ACM, pp. 92–101.
- Chellas, B. F. (1992). Time and modality in the logic of agency. *Studia Logica*, **51**(3/4), 485–517.
- Conradie, W., Goranko, V., & Vakarelov, D. (2006). Algorithmic correspondence and completeness in modal logic. I. The core algorithm SQEMA. *Logical Methods in Computer Science*, **2**(1), 1–26.
- Duijf, H., Broersen, J., & Meyer, J.-J. C. (2019). Conflicting intentions: Rectifying the consistency requirements. *Philosophical Studies*, **176**(4), 1097–1118.
- Elgesem, D. (1993). Action Theory and Modal Logic. Ph.D. Thesis, University of Oslo.
- . (1997). The modal logic of agency. *Nordic Journal of Philosophical Logic*, **2**(2), 1–46.
- Fagin, R., Moses, Y., Vardi, M. Y., & Halpern, J. Y. (2003). *Reasoning About Knowledge*. Cambridge, MA: MIT Press.
- Herzig, A., & Troquard, N. (2006). Knowing how to play: Uniform choices in logics of agency. In Stone, P., and Weiss, G., editors. *Proceedings of the Fifth International Conference on Autonomous Agents and Multi-agent Systems*. New York, NY: ACM, pp. 209–216.
- Horty, J. F. (1996). Agency and obligation. *Synthese*, **108**(2), 269–307.
- . (2001). *Agency and Deontic Logic*. New York, NY: Oxford University Press.
- Horty, J. F., & Belnap, N. (1995). The deliberative stit: A study of action, omission, ability, and obligation. *Journal of Philosophical Logic*, **24**(6), 583–644.
- Horty, J. F., & Pacuit, E. (2017). Action types in stit semantics. *The Review of Symbolic Logic*, **10**(4), 617–637.
- Jamroga, W., & Ågotnes, T. (2007). Constructive knowledge: What agents can achieve under imperfect information. *Journal of Applied Non-Classical Logics*, **17**(4), 423.
- Jamroga, W., & van der Hoek, W. (2004). Agents that know how to play. *Fundamenta Informaticae*, **63**(2–3), 185–220.
- Kanger, S. (1971). New foundations for ethical theory. In Hilpinen, R., editor. *Deontic Logic: Introductory and Systematic Readings*. Dordrecht, Netherlands: D. Reidel Publishing Company, pp. 36–58.
- Kenny, A. (1975). *Will, Freedom and Power*. Oxford, UK: Basil Blackwell.
- . (1976). Human abilities and dynamic modalities. In Manninen, J., and Tuomela, R., editors. *Essays on Explanation and Understanding*. Dordrecht, Netherlands: D. Reidel Publishing Company, pp. 209–232.
- Lindström, S., & Segerberg, K. (2007). Modal logic and philosophy. In Blackburn, P., Van Benthem, J., and Wolter, F., editors. *Studies in Logic and Practical Reasoning, volume 3 of Handbook of Modal Logic*. Netherlands: Elsevier, pp. 1149–1214.
- Lorini, E., Longin, D., & Mayor, E. (2014). A logical analysis of responsibility attribution: Emotions, individuals and collectives. *Journal of Logic and Computation*, **24**(6), 1313–1339.
- Meyer, J.-J. C., & van der Hoek, W. (1995). *Epistemic Logic for AI and Computer Science*. New York, NY: Cambridge University Press.
- Perloff, M., & Belnap, N. (2011). Future contingents and the battle tomorrow. *The Review of Metaphysics*, **64**(3), 581–602.

- Pörn, I. (1970). *The Logic of Power*. New York, NY: Barnes & Noble.
- Prior, A. N. (1967). *Past, Present and Future*. Oxford, UK: Clarendon Press.
- Schobbens, P.-Y. (2004). Alternating-time logic with imperfect recall. In van der Hoek, W., Lomuscio, A., de Vink, E., and Wooldridge, M., editors. *Logic and Communication in Multi-Agent Systems*, Vol. 85. Netherlands: Elsevier, pp. 82–93.
- Thomason, R. H. (1970). Indeterminist time and truth-value gaps. *Theoria*, **36**(3), 264–281.
- . (1984). Combinations of tense and modality. In Gabbay, D., and Guenther, F., editors. *Handbook of Philosophical Logic*, Vol. 165. Dordrecht, Netherlands: Springer, pp. 135–165.
- van der Hoek, W., & Wooldridge, M. (2003). Cooperation, knowledge, and time: Alternating-time temporal epistemic logic and its applications. *Studia Logica*, **75**(1), 125–157.
- Wooldridge, M., & Jennings, N. R. (1995). Intelligent agents: Theory and practice. *The Knowledge Engineering Review*, **10**(2), 115–152.
- Xu, M. (2015). Combinations of stit with ought and know. *Journal of Philosophical Logic*, **44**(6), 851–877.

DEPARTMENT OF PHILOSOPHY
VRIJE UNIVERSITEIT AMSTERDAM
AMSTERDAM, THE NETHERLANDS
E-mail: h.w.a.duijf@vu.nl

DEPARTMENT OF PHILOSOPHY AND RELIGIOUS STUDIES
UTRECHT UNIVERSITY
UTRECHT, THE NETHERLANDS
E-mail: j.m.broersen@uu.nl
E-mail: a.kuncova@uu.nl
E-mail: a.i.ramirezabarc@uu.nl