



HAL
open science

Comparing semantics for temporal STIT logic

Roberto Ciuni, Emiliano Lorini

► **To cite this version:**

Roberto Ciuni, Emiliano Lorini. Comparing semantics for temporal STIT logic. *Logique et Analyse*, 2018, 61 (243), pp.299-339. hal-02378366

HAL Id: hal-02378366

<https://hal.science/hal-02378366>

Submitted on 25 Nov 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Open Archive Toulouse Archive Ouverte

OATAO is an open access repository that collects the work of Toulouse researchers and makes it freely available over the web where possible

This is an author's version published in:

<http://oatao.univ-toulouse.fr/24701>

Official URL

<http://virthost.vub.ac.be/lnaweb/ojs/index.php/LogiqueEtAnalyse/article/view/2146>

To cite this version: Ciuni, Roberto and Lorini, Emiliano *Comparing semantics for temporal STIT logic*. (2018) *Logique et Analyse*, 61 (243). 299-339. ISSN 2295-5836

Any correspondence concerning this service should be sent to the repository administrator: tech-oatao@listes-diff.inp-toulouse.fr

Comparing semantics for temporal STIT logic

Roberto Ciuni¹ and Emiliano Lorini²

¹Department FISPPA, Section of Philosophy, University of Padova, Italy

²IRIT-CNRS, University of Toulouse, France*

Abstract

In this paper we establish equivalence results for the different semantics for the temporal STIT logic T-STIT, that includes temporal operators and the group agency operator for the grand coalition, and we study a semantics for temporal STIT that is based on the concept of interpreted system à la Fagin et al. We discuss the descriptive adequacy of the above semantics in capturing a given game-theoretical scenario where information about the players is included, and we compare them with traditional BT+AC semantics. Also, we discuss the extension of T-STIT with full groups and the corresponding operators, and we discuss the distinction between frames that impose additivity and superadditivity on the choices of arbitrary groups.

1 Introduction

STIT logics (the logics of *seeing to it that*) are the logics of sentences of the form “agent i sees to it that φ is true”, where ‘seeing to it that’ is interpreted as an agent-relative *modal operator*. They have been first introduced by Belnap et al. in a number of papers culminating in [2]. The logic presented there includes the ‘historical necessity’ operator \Box , the temporal operators G and H of Linear Temporal Logic (LTL)—together with their duals F and P—and ‘STIT operators’ for individual agency of the form $[i]$. In [18] Horty extends this language by introducing group STIT operators of the form $[J]$ expressing sentences like “group J sees to it that φ is true”, where a group J is a set of individual agents.¹ Some recent research [20] has focused on a fragment of Horty’s language and semantics where only a STIT operator for the collective action of all agents appears alongside $[i]$, \Box , G and H. One reason for this is that, under the definition of the choice of group from [18], full group STIT logic is undecidable and not finitely axiomatizable [17], while fragments that express only the collective action of all agents are decidable and finitely axiomatizable [26]. Another reason concerns some

*The authors wish to thank two anonymous reviewers for their helpful comments. This paper was written while Roberto Ciuni was a Marie Curie IE Fellow with the WADOXA project at the Institute of Logic, Language and Computation, University of Amsterdam (2015–2016).

¹Also [2, 281–291] approaches collective (‘joint’) agency. Horty’s variant of group STIT is the most established in formal investigation on STIT—see for instance [17].

complications brought in by full groups. We briefly discuss the two points later on in this introduction.

In the last decade, STIT logics have attracted more and more attention in the area of multi-agent systems (MAS). The reason for this is that they provide a rigorous formal tool to reason about agents' interaction and collective agency. This interest has also prompted a number of formal proposals that modify the original semantics from [2] or define a language that is altogether different from the one introduced in [18]. As a by-product of this proliferation, however, a systematic overview is now missing. Besides, these semantics—as well as the traditional one—lack a property that proves desirable in computation-oriented analysis: they are not *computationally grounded*. That is, none of them displays a clear correspondence between states in the computing system and configurations in the semantical description of the logic—see [11, p. 151]. We believe that lack of a global picture and focused computation-oriented analysis may prove a hindrance for a full encounter between STIT logics and MAS.

In the present paper, we compare three alternative semantics w.r.t. the ‘temporal STIT logic’ T-STIT proposed in [20]. Besides, we relate the alternative semantics to the traditional semantics from [2]. Finally, we propose a computationally grounded semantics for temporal STIT that is based on the concept of interpreted system by Fagin et al. [12]. The paper has two aims. First, it aims at clarifying the relations among the different semantics for STIT logics w.r.t. a logic that aims at keeping a balance between the expressive power from [18] and convenient formal properties. Second, it aims at providing a semantics that can show in a clear way the connection between STIT theory and MAS.

In the remainder of the introduction we give a more detailed description of our topic and present the structure of the paper. Beside, we motivate our modeling choices, and summarize the rationale of our results.

Structure of the paper The original semantics for STIT logics is proposed by Belnap et al. in [2] and is based on the so-called BT+AC structures: branching-time (BT) structures—or *trees*—endowed with a set of agents and a function determining the *possible choices* of each agent at any moment (AC is indeed acronym of ‘Agents and Choices’). Recently, alternative semantics have been proposed: [9] and [20] all propose semantics that are not based on BT+ACs. In particular, [9] proposes to impose *bundles* on BT+ACs, thus performing a Henkin move on those structures, and the temporal Kripke STIT frames introduced in [20] present a one-sorted Kripke frame where two different relations are defined on points: one defines the temporal dimension of the frame, the other allows for reasoning about (historical) possibility and necessity. To these temporal structures, these semantics add a *function of choice* [9] or a set of *choice-equivalence relations* [20] that allows for reasoning about ‘seeing to it that’.² Beside, a further representation of agency and time can be offered by considering a natural extension of the Kamp frames from [27], where *times* and *worlds* are defined as primitives, and each world is endowed with a temporal order.

Here, we discuss and generalize the structures from [20], [9], and structures for

²The definition of the single-agent choice b-trees from [9] is equivalent to a definition in terms of a choice-equivalence relation (see below). Here, we use a relational semantics for choice b-trees rather than a function of choice, since this allows us for a more uniform presentation.

agency and time based on Kamp frames. In doing this, we relax some conditions on the structures from [9]—for instance, we drop discreteness—and extend them to a multi-agent setting. We discuss this choice later in this introduction, but it is important to clarify its rationale now. It will make our result more general, which is in turn a virtue in view of the global picture of STIT theory we would like to contribute to.

In Section 2, we consider the logic T-STIT introduced in [20], which includes operators $[i]$ and $[Agt]$ for individual agency and the collective action of all agents, as well as the operators G and H for ‘always in the future’ and ‘always in the past’, respectively, and the operator \Box of historical necessity. Also, we prove that the three semantics are equivalent relative to T-STIT. In Section 3, we propose the interpreted system semantics for T-STIT. We prove that, under adequate restrictions, our semantics in terms of interpreted systems is equivalent to the alternative semantics defined in Section 2. In Section 4, we sum up the relations among the different semantics presented in the previous two sections, and we show that the equivalence results proved there do not extend to the semantics based on BT+AC from [2]. In Section 5, we discuss the extension of the structures from [20] with choice-equivalence relations for *groups* of agents beyond the grand coalition and corresponding group STIT operators, and we highlight how this affects the equivalence results from Section 2. In Section 6 we briefly discuss the descriptive adequacy of the alternative semantics, an issue that has not yet been tackled, at least to the authors’ knowledge. Section 7 draws conclusions and opens directions for future research. We have placed all the proofs in the technical appendix at the end of the paper, for the sake of readability.

Rationale Here we present some motivations for our modeling choices and the structures we discuss in the paper. Along with different structures, different languages for ‘seeing to it that’ and time have been proposed. In particular, [4] and [9] propose *fused stit-temporal operators* expressing, among other notions, “agent i sees to it that next (sooner or later, always in the future)”. Such languages can just express ‘non-instantaneous agency’. By contrast, the language of T-STIT from [20] can express ‘instantaneous agency’: it includes separate operators for ‘seeing to it that’ and for the future (or past), that can combine to form stit-temporal formulae.

The rationale of T-STIT is that the language from [18] is still the most widespread in the community at the crossing between philosophy of agency and MAS. At the same time, when it comes to a balance between expressive power and relevant formal properties, T-STIT proves better than the apparatus from [18]. Indeed, the language from [18] including operators for all definable groups has been proved undecidable and non-finitely axiomatizable.³ By contrast, T-STIT from [20] has an elegant axiomatization (see Section 2) and its decidability is still an open issue.⁴ The languages from [4] and [9] are of course interesting per se, but their comparison with the structures presented here requires another paper, on which we wish to devote some future research.

Our focus on T-STIT also implies some modifications on the structures from [9]. First, we drop *discreteness* for reasons mentioned above. Second, we extend the choice

³As we have mentioned, the result by [17] presupposes the definition of the choice of a group given in [18]. We discuss another possible definition in Section 5.

⁴Decidability is also an open issues for the logic from [4], which also includes group STIT operators, but this logic, as well as the one from [9], is less expressive than a corresponding logic for instantaneous agency.

b-trees from [9] with choice-equivalence relations for many agents and the grand coalition. Third, we drop the condition from [9] that choice b-trees have a first element. These moves allow us to gain a higher generality. Indeed, temporal frames with discrete time and a first element are just specific examples of a more general family of temporal frames. Since we aim at a systematic overview on the semantics for temporal STIT logic, our pursue of a higher generality is justified, we believe.

Also, the distinction between *additivity* and *superadditivity* for collective choices (see Section 2 for the two notions) brings complications that are not relevant for the comparison we are going to carry out in this paper. Here, we impose *additivity* to the choices of the grand coalition, while other semantics—see [4]—impose *superadditivity*, but this choice has no bearing on the the results of the paper: T-STIT does not distinguish additive frames from superadditive ones (see Sections 2.3 and 2.5). In Section 5, we show that the distinction between additivity and superadditivity becomes relevant when considering all groups of agents.

Significance of the results We conclude with some remarks on the relevance of our results. The significance of the *equivalence result* in Section 2 is connected to the variety of current semantics for temporal STIT logics, which raises a natural question about their relations. In particular, the result guarantees that the three semantics yield exactly the same logic—namely T-STIT. The *descriptive adequacy* of the alternative semantics deserves discussion, in our view, because STIT logics aims at being a formal theory of agents’ interaction, and not only a computational device. This is in turn the rationale of Section 6.

Finally, *computationally grounded semantics* are interesting as far as they bridge the computational insights on multi-agent systems from Artificial Intelligence (AI) and the foundational approach in terms of Kripke semantics (and extensions) from the logical and philosophical investigations on agency and interaction. In this context, the significance of our results in Section 3 lies in the possibility of tightening the connection between STIT theory and the area of MAS. Indeed, *interpreted systems* have witnessed a number of applications in areas of AI and MAS such as security protocols [14] and the blockchain protocol [15]. The interpreted system semantics from Section 3 could then open a new potential area of application for STIT logics.

The reason why the interpreted system semantics for STIT of Section 3 is studied apart from the three alternative semantics presented in Section 2 is that the former makes some assumptions that are not made in the three semantics of Section 2, namely the assumption about discreteness of time and the assumption that time has a beginning. Notice that in the original semantics for STIT by Belnap et al. [2] based on BT+AC structures, time was not assumed to be discrete. In this sense, the concept of time modeled in the interpreted system semantics for STIT is more restrictive than the concept of time modeled in the original semantics for STIT by Belnap et al.

2 Three semantics for temporal STIT

The alternative semantics that we study here are the semantics from [20], a semantics based on extensions of Kamp frames, and (with some little variations) the semantics

from [9], respectively. We will see them in this section and hint at their formal rationale in Section 6—where we also discuss their descriptive adequacy.⁵

In many recent semantics for STIT—such as [1], [4] and [20]—relations of *choice equivalence* (one per agent) play the role of the *function of choice* originally proposed in [2, 18] and applied in [9]. We go along this line here, without loss of generality: that the *relational* treatment and the *functional* one are equivalent is already observed in [2].

2.1 Syntax

We consider the language of the logic T-STIT introduced in [20] that extends the language of atemporal individual STIT from [1, 2] with: (i) the future tense and past tense operators, and (ii) Horty’s operator of group agency for the *grand coalition* (the coalition of all agents).

Let Atm be a countably infinite set $\{p, q, \dots\}$ of propositional atoms, and let Agt be a finite set $\{1, \dots, n\}$ of agents. The language $\mathcal{L}_{T-STIT}(Atm, Agt)$ of the logic T-STIT is the set of formulae defined by the following BNF:

$$\varphi ::= p \mid \neg\varphi \mid \varphi \wedge \varphi \mid [i]\varphi \mid [Agt]\varphi \mid \Box\varphi \mid G\varphi \mid H\varphi$$

where p ranges over Atm and i ranges over Agt . The other Boolean constructions \top , \perp , \vee , \rightarrow and \leftrightarrow are defined from \neg and \wedge in the standard way. Operators of the form $[i]$ are so-called Chellas’ STIT operators, named after their proponent [8].⁶ The formula $[i]\varphi$ has to be read ‘agent i ensures φ regardless of what the other agents do’, or—more briefly—‘agent i sees to it that φ ’. $[Agt]$ is a group STIT operator and $[Agt]\varphi$ has to be read ‘all agents see to it that φ by acting together’. The duals of $[i]$ and $[Agt]$ are $\langle i \rangle$ and $\langle Agt \rangle$, respectively.

\Box is the usual necessity (or inevitability) operator, with $\Box\varphi$ reading ‘ φ is necessarily true’. Its dual is \Diamond , with $\Diamond\varphi$ reading ‘it is possibly true that φ ’. Finally, G and H are tense operators from linear temporal logic. $G\varphi$ means ‘ φ will always be true in the future’ and $H\varphi$ means ‘ φ has always been true in the past’. Their duals are F and P , respectively. $F\varphi$ reads ‘ φ will be true at least once in the future’ and $P\varphi$ reads ‘ φ was true at least once in the past’.

Combinations of the above operators allows for interesting expressive power. For instance, $[i]F\varphi$ expresses the situations where ‘agent i sees to it that φ will be true at least once’. One advantage of temporal STIT is that it can express a gap in time between choice and outcome (see example above). It is reasonable to assume such a gap, but atemporal STIT has no way to express it. Also, presence of \Box , G and H makes \mathcal{L}_{T-STIT} an extension of the language from [24, 25] for branching-time logic.

⁵It is also worth mentioning Wölf [30], that studies an alternative semantics based on so-called $T \times W$ -based agent frames, which are extensions of the $T \times W$ -frames studied in [27]. Such frames can be described as Kamp frames that have the same temporal order for all the different worlds (or ‘histories’). We do not consider the proposal in [30] here, since $T \times W$ -frames contain a mechanism of synchronization of worlds that makes them extremely different from the non-synchronized branching-time structure that we are studying here. An investigation of this option would take a full paper.

⁶More precisely, this is the name Belnap and Horty gave to the operator in [19], due to the similarities with the operator introduced in [8].

$\Diamond[i]\varphi$ reads ‘agent i is able to see to it that φ . Notice that Chellas’ stit operator and possibility operator together allows for the definition of the *deliberative* stit operator (see [1]): $[i \text{ dstit}]\varphi = [i]\varphi \wedge \Diamond\neg\varphi$.

Our language $\mathcal{L}_{\text{T-STIT}}$ is basically the language presented in [18], with the exception of *group operators*, that we leave out in order to keep a balance between expressive power and other formal properties. Indeed, [17] has proved that a STIT logic including the group operators from [18] is *not finitely axiomatizable*. Inclusion of the grand coalition operator $[Agt]$ is justified in light of this balance: T-STIT is finitely axiomatizable, as proved in [20].

2.2 Temporal Kripke STIT frames

We first consider the semantics of T-STIT based on the notion of temporal Kripke STIT frame, which was first introduced in [20]. The concept of a temporal Kripke STIT frame extends Zanardo’s concept of an Ockhamist frame [31] with a choice component. Ockhamist frames are one-sorted Kripke frames with an equivalence relation and a linear temporal order imposed on points. The addition of relations of choice-equivalence makes the frames suitable for modeling agency. The resulting semantics has the advantage of being closer to the standard semantics of modal logic (first-order quantification over point-like primitives) than the other semantics for temporal STIT.

For every set X , binary relation \mathcal{R} on X and $x \in X$, let $\mathcal{R}(x) = \{y \in X \mid x\mathcal{R}y\}$. Moreover, for all binary relations \mathcal{R}_1 and \mathcal{R}_2 on X , let $\mathcal{R}_1 \circ \mathcal{R}_2$ be the composition of \mathcal{R}_1 and \mathcal{R}_2 .⁷

Definition 1 (Temporal Kripke STIT frame) *The class \mathbb{K} of temporal Kripke STIT frames (TKSFs, for short) includes all tuples*

$\mathcal{K} = (X, \mathcal{R}_\square, \{\mathcal{R}_i\}_{i \in \text{Agt}}, \mathcal{R}_{\text{Agt}}, \mathcal{R}_G, \mathcal{R}_H)$ *where:*

- X is a nonempty set of points, which we will call time-points;
- \mathcal{R}_\square , every \mathcal{R}_i and \mathcal{R}_{Agt} are equivalence relations between time-points in X such that:
 - (C1) for all $i \in \text{Agt}$: $\mathcal{R}_i \subseteq \mathcal{R}_\square$;
 - (C2) for all $x_1, \dots, x_n \in X$: if $x_i \mathcal{R}_\square x_j$ for all $i, j \in \{1, \dots, n\}$ then $\bigcap_{1 \leq i \leq n} \mathcal{R}_i(x_i) \neq \emptyset$; (remember that $\{1, \dots, n\} = \text{Agt}$)
 - (C3) for all $x \in X$: $\mathcal{R}_{\text{Agt}}(x) = \bigcap_{i \in \text{Agt}} \mathcal{R}_i(x)$;
- \mathcal{R}_G and \mathcal{R}_H are binary relations between points in X such that \mathcal{R}_G is serial and transitive, \mathcal{R}_H is the inverse relation of \mathcal{R}_G (i.e., $\mathcal{R}_H = \mathcal{R}_G^{-1} = \{(x, y) \mid y\mathcal{R}_G x\}$), and:
 - (C4) for all $x, y, z \in X$: if $y, z \in \mathcal{R}_G(x)$ then $z \in \mathcal{R}_G(y)$ or $y \in \mathcal{R}_G(z)$ or $z = y$;
 - (C5) for all $x, y, z \in X$: if $y, z \in \mathcal{R}_H(x)$ then $z \in \mathcal{R}_H(y)$ or $y \in \mathcal{R}_H(z)$ or $z = y$;
 - (C6) $\mathcal{R}_G \circ \mathcal{R}_\square \subseteq \mathcal{R}_{\text{Agt}} \circ \mathcal{R}_G$;

⁷The composition $\mathcal{R}_1 \circ \mathcal{R}_2$ of \mathcal{R}_1 and \mathcal{R}_2 is formally defined as $\{(x, y) \in X^2 \mid \exists z \in X : \mathcal{R}_1(x, z) \text{ and } \mathcal{R}_2(z, y)\}$.

(C7) for all $x, y \in X$: if $y \in \mathcal{R}_\square(x)$ then $y \notin \mathcal{R}_G(x)$.

$\mathcal{R}_\square(x)$ can be conceived as the set of time-points that are *alternative* to the time-point x . An equivalence class $\mathcal{R}_\square(x)$ is also called a *moment*. $\mathcal{R}_G(x)$ defines the set of time-points that are in the *future* of time-point x , and $\mathcal{R}_H(x)$ is the set of time-points that are in the *past* of time-point x .

\mathcal{R}_i is a relation of choice-equivalence: $x\mathcal{R}_iy$ means that x and y are compatible with the very same choice of agent i . The set $\mathcal{R}_i(x) = \{y \in X \mid x\mathcal{R}_iy\}$ is the *choice* of agent i at x . For every time-point x , the set $\mathcal{R}_{Agt}(x)$ identifies the choice of group Agt at x .

Conditions C1–C3 and C6 define the concept of a choice. C1 states that an agent can only choose among possible alternatives: if $x\mathcal{R}_iy$, then x and y are mutually alternative. As a consequence, for every time-point x , \mathcal{R}_i induces a *partition* of the set $\mathcal{R}_\square(x)$. C2 expresses the so-called *independence of agents*: every choice of an agent i overlaps with every choice of an agent j such that $j \neq i$. C3 defines the choice of all agents (at x) as equal to the pointwise intersection of the individual choices of the agents in Agt (at x).⁸ The principle is usually called *additivity* in the literature on STIT logic. Finally, C6 guarantees that no choice can set apart two time-lines before they stop having time-points x and y that are mutually alternative. This corresponds to the property of *no choice between undivided histories* given in STIT logic [2, Chap. 7].

Conditions C4–C5 and C7 define the temporal order. C7 states that if two time-points x and y are *alternative* one to another, then none of them can be in the future of the other one. Due to the reflexivity of \mathcal{R}_\square , this implies that \mathcal{R}_G is irreflexive. Let $\mathcal{T}(x) = \mathcal{R}_H(x) \cup \{x\} \cup \mathcal{R}_G(x)$ be the set of time-points that are temporally related with time-point x . C4 and C5, together with the irreflexivity and transitivity of \mathcal{R}_G , guarantee that \mathcal{R}_G is a *strict linear order* on the set $\mathcal{T}(x)$. For every time-point x in X , we use ℓ_x as short for the linearly ordered set $(\mathcal{T}(x), \mathcal{R}_G)$ and set $TL = \{\ell_x \mid x \in X\}$. Notice that, if $y \in \mathcal{R}_H(x) \cup \{x\} \cup \mathcal{R}_G(x)$, then $\ell_x = \ell_y$. Every time-line ℓ_x is infinite, due to the seriality of \mathcal{R}_G . Due to the linearity of \mathcal{R}_G and \mathcal{R}_H , there is a functional mapping from time-points to time-lines, as for every $x \in X$ there exists exactly one time-line $\ell_x \in TL$ that includes x .

A temporal Kripke STIT model (TKSM, for short) $\mathcal{M}^\mathcal{K}$ is a pair $(\mathcal{K}, \mathcal{V})$, where \mathcal{K} is a TKSF and $\mathcal{V} : Atm \rightarrow 2^X$ is a valuation function associating every atom with a set of time-points. We assume that a TKSM satisfies the following condition:

(C8) for all $x, y \in X$ and for all $p \in Atm$: if $x\mathcal{R}_\square y$ then $x \in \mathcal{V}(p)$ iff $y \in \mathcal{V}(p)$.

C8 states that if two points belong to the same moment then they agree on the atoms. It is worth noting that this condition is not assumed by [20] and, more generally, in the context of STIT logic. However, it is generally assumed in the context of Prior’s Ockhamist branching-time logic (also called Prior’s Ockhamist logic of historical necessity) [24, 25] and of Full Computation Tree Logic (CTL*) [23]. In particular, under condition C8, Prior’s Ockhamist branching time logic (OBTL) as defined by Reynolds in [24, 25] is nothing but the fragment of T-STIT without the individual agency operators $[i]$ and the group agency operator $[Agt]$.

⁸This corresponds to the notion of joint action proposed by Horty in [18], where the joint action of a group is described in terms of the result that the agents in the group bring about by acting together.

Relative to TKSMs, the truth conditions for formulae in $\mathcal{L}_{\text{T-STIT}}$ are defined as follows.

$$\begin{aligned}
\mathcal{M}^{\mathcal{K}}, x \models p &\Leftrightarrow x \in \mathcal{V}(p) \\
\mathcal{M}^{\mathcal{K}}, x \models \neg\varphi &\Leftrightarrow \mathcal{M}^{\mathcal{K}}, x \not\models \varphi \\
\mathcal{M}^{\mathcal{K}}, x \models \varphi \wedge \psi &\Leftrightarrow \mathcal{M}^{\mathcal{K}}, x \models \varphi \text{ and } \mathcal{M}^{\mathcal{K}}, x \models \psi \\
\mathcal{M}^{\mathcal{K}}, x \models \Box\varphi &\Leftrightarrow \forall y \in \mathcal{R}_{\Box}(x) : \mathcal{M}^{\mathcal{K}}, y \models \varphi \\
\mathcal{M}^{\mathcal{K}}, x \models [i]\varphi &\Leftrightarrow \forall y \in \mathcal{R}_i(x) : \mathcal{M}^{\mathcal{K}}, y \models \varphi \\
\mathcal{M}^{\mathcal{K}}, x \models [\text{Agt}]\varphi &\Leftrightarrow \forall y \in \mathcal{R}_{\text{Agt}}(x) : \mathcal{M}^{\mathcal{K}}, y \models \varphi \\
\mathcal{M}^{\mathcal{K}}, x \models \mathbf{G}\varphi &\Leftrightarrow \forall y \in \mathcal{R}_{\mathbf{G}}(x) : \mathcal{M}^{\mathcal{K}}, y \models \varphi \\
\mathcal{M}^{\mathcal{K}}, x \models \mathbf{H}\varphi &\Leftrightarrow \forall y \in \mathcal{R}_{\mathbf{H}}(x) : \mathcal{M}^{\mathcal{K}}, y \models \varphi
\end{aligned}$$

The notions of validity and satisfiability are defined as usual. We use $\models_{\mathcal{K}}$ to denote TKSF-validity. The semantics of T-STIT is exemplified in Figure 1. In the TKSM $\mathcal{M}^{\mathcal{K}}$ represented there, x is a given point, and the (big) thick rectangles are classes of *alternatives* of given points—thus, for instance, the thick rectangle at the center of the figure is the class $\mathcal{R}_{\Box}(x)$ of the alternatives of x . The choices of agents 1 and 2 are represented by columns and rows, respectively—thus, for instance, the column including x is the choice $\mathcal{R}_1(x)$ of agent 1 at point x , and the row including x is the choice $\mathcal{R}_2(x)$ of agent 2 at point x . The (small) dotted rectangles represent the choices of all agents (1 and 2) together—thus for instance, the dotted rectangle around x is the choice $\mathcal{R}_{\text{Agt}}(x) = \mathcal{R}_{\{1,2\}}(x)$ of agents 1 and 2 at point x . Arrows represent the temporal relation $\mathcal{R}_{\mathbf{G}}$. As it is easy to check, $\mathcal{M}^{\mathcal{K}}, x \models [2]p$, since p holds at every time-point in agent 2’s choice at x . Also $\mathcal{M}^{\mathcal{K}}, x \models [1](p \vee q)$, because either p or q hold at every point in agent 1’s choice at x .

In our setting, all the axioms of Linear Time Logic LTL hold for \mathbf{G} , \mathbf{H} and their duals, and that the axioms of S5 hold for \Box , $[i]$ and $[\text{Agt}]$. The axiomatization of T-STIT in Subsection 2.5 includes such axioms, together with important principles for agency and its connection with time and necessity. We just highlight some of them here.

Formula $\mathbf{F}\Diamond\varphi \rightarrow \langle \text{Agt} \rangle \mathbf{F}\varphi$ corresponds to C6 (*no choice between undivided histories*). The operator $[\text{Agt}]$ in turn proves indispensable to express the condition, which provides the rationale for introducing it here.

Formula $(\Diamond[1]\varphi_1 \wedge \dots \wedge \Diamond[n]\varphi_n) \rightarrow \Diamond([1]\varphi_1 \wedge \dots \wedge [n]\varphi_n)$ corresponds to C2 and expresses the independence of agency. The compatibility of all the choices of an agent with all choices of the other agent is in turn illustrated in Figure 1, where each choice of 1 at x overlaps with each choice of 2 at some time-point that is alternative to x .

Formula $([1]\varphi_1 \wedge \dots \wedge [n]\varphi_n) \rightarrow [\text{Agt}](\varphi_1 \wedge \dots \wedge \varphi_n)$ corresponds to the condition that, for all $x \in X : \mathcal{R}_{\text{Agt}}(x) \subseteq \bigcap_{i \in \text{Agt}} \mathcal{R}_i(x)$. The condition is known as *superadditivity*. Our definition of TKSFs includes the stricter condition C3, (*additivity*) imposing $\mathcal{R}_{\text{Agt}}(x) = \bigcap_{i \in \text{Agt}} \mathcal{R}_i(x)$ for all $x \in X$. This does not make a difference with respect to our axiom, however. Indeed, it is proved in [20, Lemma 9, p. 381] that, given our language $\mathcal{L}_{\text{T-STIT}}$, we obtain the same notion of validity if we consider TKSFs from Definition 1 and if we replace C3 with superadditivity in Definition 1. Finally, $\Box\varphi \rightarrow [i]\varphi$ corresponds to C1.




Figure 1: Example of TKSM

2.3 Kamp agent frames

We now consider a semantics for T-STIT based on the Kamp frames that are discussed in [27] and attributed to Hans Kamp. These frames include *times* and *worlds* as primitives, and each world is provided with its own (linear) temporal structure.

Contrary to the other structures we investigate here, Kamp frames have been paid little attention in the STIT literature—an exception is [2], where the descriptive adequacy of Kamp frames as models for time is discussed. That notwithstanding, Kamp frames enjoy interesting logical features. In particular, Kamp frames are *two-sorted* structures that allow for *first-order* quantification over their main entities. This is a very convenient feature, since first-order theories are more technically tractable than second-order theories. Due to this, we believe that Kamp agent frames are worth investigating as mathematical models for agency, and deserve consideration.

Definition 2 (Kamp agent frame) *The class \mathbb{F} of Kamp agent frames includes all tuples $\mathcal{F} = (W, \mathcal{O}, \{\sim_t\}_{t \in T}, \{\sim_{(t,i)}\}_{t \in T, i \in \text{Agt}}, \{\sim_{(t, \text{Agt})}\}_{t \in T})$ where:*

- W is a non-empty set of worlds;
- \mathcal{O} is a function mapping every world w in W to an infinite⁹ linearly ordered set $\langle T_w, <_w \rangle$, where $<_w$ is a linear ordering on T_w . As usual, we let $>_w$ be the inverse relation of $<_w$. Also, we define $T = \bigcup_{w \in W} T_w$ and for all $t \in T$, $W_t = \{w \in W \mid t \in T_w\}$;

⁹The fact that every set T_w is infinite corresponds to the seriality of the relation \mathcal{R}_G in the definition of TKSF.

- for all $t \in T$ and for all $i \in \text{Agt}$, \sim_t , $\sim_{\langle t, i \rangle}$ and $\sim_{\langle t, \text{Agt} \rangle}$ are equivalence relations on W_t that satisfy the following four constraints:

(TW1) for all $t \in T$ and for all $i \in \text{Agt}$: $\sim_{\langle t, i \rangle} \subseteq \sim_t$;

(TW2) for all $t \in T$ and for all $w_1, \dots, w_n \in W$: if $w_i \sim_t w_j$ for all $i, j \in \{1, \dots, n\}$ then $\bigcap_{1 \leq i \leq n} \sim_{\langle t, i \rangle}(w_i) \neq \emptyset$;

(TW3) for all $t \in T$ and for all $w \in W$: $\sim_{\langle t, \text{Agt} \rangle}(w) = \bigcap_{i \in \text{Agt}} \sim_{\langle t, i \rangle}(w)$;

(TW4) for all $t, t' \in T$ and for all $w, w' \in W$: if $t <_w t'$ and $w \sim_{t'} w'$, then $w \sim_{\langle t, \text{Agt} \rangle} w'$.

For every time t , $w \sim_t w'$ states that w and w' match at t . What we have here is a number of world-relative temporal orders determined by $<_w$; the different worlds in W are related together on the ground of (shared) temporal orders. $w \sim_{\langle t, i \rangle} w'$ reads ‘ w is choice-equivalent to w' at t for i ,’ and $w \sim_{\langle t, \text{Agt} \rangle} w'$ reads ‘ w is choice-equivalent to w' at t for Agt .’ Thus, $\sim_{\langle t, i \rangle}(w)$ and $\sim_{\langle t, \text{Agt} \rangle}(w)$ are the choice of i at t which includes w and the collective choice of all agents at t which includes w , respectively. TW1–TW3 are the Kamp-counterparts of Constraints C1–C3 in Definition 1, and TW4 is the counterpart of no choice between undivided histories (Constraint C7): if two worlds match at a later time t' , they are *now* in the same choice of Agt . TW1, TW3 and TW4 together imply that, if $t <_w t'$ and $w \sim_{t'} w'$, then $w \sim_t w'$.¹⁰ Thus, our construction of Agt ’s choices and the choices of the agents is enough to guarantee that, if matching at a time t , different worlds also match at any previous time. Also, the three conditions imply that, if $w \sim_t w'$, then $\{t' | t' \in T_w \text{ and } t' <_w t\} = \{t' | t' \in T_{w'} \text{ and } t' <_{w'} t\}$. This guarantees that if two worlds match at t , the times earlier than t at w are the same as the times earlier than t at w' . This in turn guarantees the weaker condition that the orderings of w and w' are isomorphic *at least up to* t .

Notice that it is possible to replace *additivity* in Kamp agent frames (the equality condition in TW3 below) with *superadditivity* by imposing the weaker ‘subset condition’ $\sim_{\langle t, \text{Agt} \rangle}(w) \subseteq \sim_{\langle t, i \rangle}(w)$ (the choice of the grand coalition is *equal or smaller than* the overlap of the choices of all the agents). This condition is in turn the counterpart in Kamp agents frames of $\mathcal{R}_{\text{Agt}}(x) \subseteq \bigcap_{i \in \text{Agt}} \mathcal{R}_i(x)$ from Section 2.2. Theorem 1 and [20, Lemma 9, p. 381] make it clear that this difference has no impact on the interpretation of $\mathcal{L}_{\text{T-STIT}}$: the Kamp agent frames including additivity and those including superadditivity are equivalent with respect to our language.

A Kamp agent model $\mathcal{M}^{\mathcal{F}}$ is a pair (\mathcal{F}, π) , where \mathcal{F} is a Kamp agent frame and $\pi : \text{Atm} \rightarrow 2^T$ is a valuation function associating every atom with a set of times (those at which the atom is true). Formulae of $\mathcal{L}_{\text{T-STIT}}$ are evaluated w.r.t. a Kamp agent

¹⁰This is compatible with w and w' not matching at a later time: as time goes on, two worlds can stop matching. Thus, two worlds that are mutually alternative *now* can have different futures, but they share their past.

model $\mathcal{M}^{\mathcal{F}}$ and a time-world pair $\langle t, w \rangle$ in $\mathcal{M}^{\mathcal{F}}$ such that $t \in T_w$:

$$\begin{aligned}
\mathcal{M}^{\mathcal{F}}, \langle t, w \rangle \models p &\Leftrightarrow t \in \pi(p) \\
\mathcal{M}^{\mathcal{F}}, \langle t, w \rangle \models \neg\varphi &\Leftrightarrow \mathcal{M}^{\mathcal{F}}, \langle t, w \rangle \not\models \varphi \\
\mathcal{M}^{\mathcal{F}}, \langle t, w \rangle \models \varphi \wedge \psi &\Leftrightarrow \mathcal{M}^{\mathcal{F}}, \langle t, w \rangle \models \varphi \text{ and } \mathcal{M}^{\mathcal{F}}, \langle t, w \rangle \models \psi \\
\mathcal{M}^{\mathcal{F}}, \langle t, w \rangle \models \Box\varphi &\Leftrightarrow \forall w' \in W : \text{if } w \sim_t w' \text{ then } \mathcal{M}^{\mathcal{F}}, \langle t, w' \rangle \models \varphi \\
\mathcal{M}^{\mathcal{F}}, \langle t, w \rangle \models [i]\varphi &\Leftrightarrow \forall w' \in W : \text{if } w \sim_{\langle t, i \rangle} w' \text{ then } \mathcal{M}^{\mathcal{F}}, \langle t, w' \rangle \models \varphi \\
\mathcal{M}^{\mathcal{F}}, \langle t, w \rangle \models [Agt]\varphi &\Leftrightarrow \forall w' \in W : \text{if } w \sim_{\langle t, Agt \rangle} w' \text{ then } \mathcal{M}^{\mathcal{F}}, \langle t, w' \rangle \models \varphi \\
\mathcal{M}^{\mathcal{F}}, \langle t, w \rangle \models G\varphi &\Leftrightarrow \forall t' \in T : \text{if } t <_w t' \text{ then } \mathcal{M}^{\mathcal{F}}, \langle t', w \rangle \models \varphi \\
\mathcal{M}^{\mathcal{F}}, \langle t, w \rangle \models H\varphi &\Leftrightarrow \forall t' \in T : \text{if } t >_w t' \text{ then } \mathcal{M}^{\mathcal{F}}, \langle t', w \rangle \models \varphi
\end{aligned}$$

Validity and satisfiability are defined as usual, with $\models_{\mathcal{F}} \varphi$ denoting the fact that a formula φ is valid relative to the class of Kamp agent models.

2.4 Choice b-trees

We now consider choice b-trees, which were first introduced in [9]. Such structures are based on bundled trees of temporal logics, which [9] augments with agent-relative relations. The structures we are presenting here generalize those from [9] in three respects. First, they include many agents and the grand coalition Agt , while [9] considers only structures with one agent. Second, we assume that choice b-trees have no root (*i.e.*, no \prec -first element for the relation \prec below). Third, we do not impose any constraint on the cardinality of time, while the structures in [9] impose discreteness.¹¹ The upgrade to a multi-agent setting is just a natural extension of the setting presented in [9], and the generalization to structures having no root is justified by the convenient logical properties that this guarantees when past operators are accounted for. Theorem 1 below easily applies to the discrete-time versions of temporal STIT frames and our generalized choice b-trees, and thus we will not pursue a focused formal comparison between such structures here.

Bundled trees impose a selection over the histories of (full) trees $\mathfrak{T} = (M, \prec)$, where M is a non-empty set of moments and \prec is a serial, irreflexive, transitive and past-linear relation on M . We first introduce *trees*, *bundles* and *bundled trees*, and we then define choice b-trees.

Definition 3 (Trees) *A (full) tree is a tuple $\mathfrak{T} = (M, \prec)$, where:*

- M is a non-empty set of moments;
- \prec is a relation on M that is serial, irreflexive, transitive and past-linear (*i.e.*, for all $m, m', m'' \in M$ if $m' \prec m$ and $m'' \prec m$ then $m' = m''$ or $m' \prec m''$ or $m'' \prec m'$). We let \succ be the inverse relation of \prec .

¹¹Also, we replace the function of choice with choice-equivalence classes, with a move that is usual in STIT logics and will have no bearings in what follows.

The notion of a *history* is also crucial in such structures:

Definition 4 (Histories) Histories are sets h, h', \dots of moments that are linearly ordered by \prec and are maximal for inclusion. $H_{\mathfrak{T}}$ is the set of all histories in the tree \mathfrak{T} , and H_m is the set of histories h such that $m \in h$ (the histories “passing through m ”)—we omit reference to the given tree, in this case.

Definition 5 (Bundles) B is a bundle on \mathfrak{T} iff $B \subseteq H_{\mathfrak{T}}$, and

(B0) for every $m \in M$, there is a history $h \in B$ such that $m \in h$.¹²

We write B_m to denote the set of histories in B which pass through m , that is $B \cap H_m$. Condition B0 ensures that $B \cap H_m \neq \emptyset$. A *bundled tree* is a pair (\mathfrak{T}, B) where \mathfrak{T} is a tree and B is a bundle on \mathfrak{T} .¹³ We are now ready to define *choice b-trees*.

Definition 6 (Choice b-trees) The class \mathbb{B} of choice b-trees include all tuples $\mathcal{B} = (\mathfrak{T}, B, \{\sim_{\langle m, i \rangle}\}_{m \in M, i \in \text{Agt}}, \{\sim_{\langle m, \text{Agt} \rangle}\}_{m \in M})$ where:

- (\mathfrak{T}, B) is a bundled tree;
- for all $m \in M$ and for all $i \in \text{Agt}$, $\sim_{\langle m, i \rangle}$ and $\sim_{\langle m, \text{Agt} \rangle}$ are equivalence relations on B_m that satisfy the following conditions:
 - (B1) for all $m \in M$ and for all $h_1, \dots, h_n \in B_m$: $\bigcap_{1 \leq i \leq n} \sim_{\langle m, i \rangle}(h_i) \neq \emptyset$;
 - (B2) for all $m \in M$ and for all $h \in B_m$: $\sim_{\langle m, \text{Agt} \rangle}(h) = \bigcap_{i \in \text{Agt}} \sim_{\langle m, i \rangle}(h)$;
 - (B3) for all $m, m' \in M$ and for all $h, h' \in B$: if $m \prec m'$ and $h, h' \in B_{m'}$, then $h, h' \in B_m$ and $h \sim_{\langle m, \text{Agt} \rangle} h'$.

$\sim_{\langle m, i \rangle}$ and $\sim_{\langle m, \text{Agt} \rangle}$ read ‘ h is choice-equivalent to h' for agent i at m ’ and ‘ h is choice-equivalent to h' for Agt at m ,’ respectively. B1–B2 are the counterparts of Constraints C2–C3 in Definition 1, while B3 is the counterpart of C6.

Again, we can replace *additivity* (condition B2) with *superadditivity* by imposing the weaker condition $\sim_{\langle m, \text{Agt} \rangle}(h) \subseteq \bigcap_{i \in \text{Agt}} \sim_{\langle m, i \rangle}(h)$. Theorem 2 and [20, Lemma 9, p. 381] secure that this does not change the logic: the choice b-trees including additivity and those including superadditivity are equivalent with respect to our language.

A choice b-model $\mathcal{M}^{\mathcal{B}}$ is a pair (\mathcal{B}, v) , where \mathcal{B} is a choice b-tree and $v : \text{Atm} \rightarrow 2^M$ is a valuation function associating atoms with sets of moments. Formulae of $\mathcal{L}_{\text{T-STIT}}$ are evaluated w.r.t. a choice b-model $\mathcal{M}^{\mathcal{B}}$ and a moment-history pair $\langle m, h \rangle$ in $\mathcal{M}^{\mathcal{B}}$

¹²The usual definition of a bundle allows for *branches*, that is possibly non-maximal \prec -chains of moments to be considered, and it then imposes that, if a branch is in a bundle, then all its sub-branches and all its linear extensions are in the bundle. We do not need to impose the latter in our definition, since we are defining bundles in terms of maximal \prec -chains of moments.

¹³It is clear by Definitions 3 and 5 that *full trees* \mathfrak{T} are a limit-case of bundled trees, namely those bundled trees where $B_m = H_m$ for every $m \in M$.

such that $h \in B$ and $m \in h$:

$$\begin{aligned}
\mathcal{M}^{\mathcal{B}}, \langle m, h \rangle \models p &\Leftrightarrow m \in v(p) \\
\mathcal{M}^{\mathcal{B}}, \langle m, h \rangle \models \neg\varphi &\Leftrightarrow \mathcal{M}^{\mathcal{B}}, \langle m, h \rangle \not\models \varphi \\
\mathcal{M}^{\mathcal{B}}, \langle m, h \rangle \models \varphi \wedge \psi &\Leftrightarrow \mathcal{M}^{\mathcal{B}}, \langle m, h \rangle \models \varphi \text{ and } \mathcal{M}^{\mathcal{B}}, \langle m, h \rangle \models \psi \\
\mathcal{M}^{\mathcal{B}}, \langle m, h \rangle \models \Box\varphi &\Leftrightarrow \forall h' \in B_m : \mathcal{M}^{\mathcal{B}}, \langle m, h' \rangle \models \varphi \\
\mathcal{M}^{\mathcal{B}}, \langle m, h \rangle \models [i]\varphi &\Leftrightarrow \forall h' \in B_m : \text{if } h \sim_{\langle m, i \rangle} h' \text{ then } \mathcal{M}^{\mathcal{B}}, \langle m, h' \rangle \models \varphi \\
\mathcal{M}^{\mathcal{B}}, \langle m, h \rangle \models [Agt]\varphi &\Leftrightarrow \forall h' \in B_m : \text{if } h \sim_{\langle m, Agt \rangle} h' \text{ then } \mathcal{M}^{\mathcal{B}}, \langle m, h' \rangle \models \varphi \\
\mathcal{M}^{\mathcal{B}}, \langle m, h \rangle \models G\varphi &\Leftrightarrow \forall m' \in h : \text{if } m \prec m' \text{ then } \mathcal{M}^{\mathcal{B}}, \langle m', h \rangle \models \varphi \\
\mathcal{M}^{\mathcal{B}}, \langle m, h \rangle \models H\varphi &\Leftrightarrow \forall m' \in h : \text{if } m \succ m' \text{ then } \mathcal{M}^{\mathcal{B}}, \langle m', h \rangle \models \varphi
\end{aligned}$$

Validity and satisfiability are defined as usual, with $\models_{\mathcal{B}} \varphi$ denoting the fact that a formula φ is valid relative to the class of choice b-models.

The traditional BT+AC structures for STIT logics given by Belnap et al. [2] are based on *full* trees. As we have already mentioned, a full tree is a *bundled* tree such that its bundle B coincides with the set $H_{\mathfrak{T}}$ of histories in the tree. Thus, full trees are presupposed by bundled trees and, conceptually, more basic than the latter. Here we consider a version of BT+ACs that includes the individual choices of many agents and the collective choices of the set Agt of all the agents. Also, we will consider a relational version of the structures. This will bring no real deflection from BT+ACs: it is easy to see that the equivalence relations for choice naturally induce a function of choice as defined in [2, Def. 11], and vice versa (an example of this is [2, Def. 12, p. 214]).

The presentation of BT+AC structures—BT+ACs, for short—is extremely simple, since the only difference with choice b-trees is in the fact that $B_m = H_m$ for all $m \in M$ —that is, we simply need to drop B and extend our consideration to *all* maximal \prec -chains defined in \mathfrak{T} .

Definition 7 (BT + AC's) *The class \mathbb{T} of BT + AC's includes all those choice b-trees \mathcal{B} such that $B = H_{\mathfrak{T}}$.*

A BT+AC model is a choice-b model (\mathcal{B}, v) such that \mathcal{B} is a BT+AC. The truth conditions for the formulae in \mathcal{L}_{T-STIT} w.r.t. BT+AC models obtain from the truth conditions w.r.t. choice b-models by substituting any occurrence of B_m with H_m .

2.5 Equivalence between the three semantics

The following theorems highlight that the semantics based on Kamp agent frames, that based on TKSFs and that based on bundled trees are equivalent w.r.t. \mathcal{L}_{T-STIT} .

Theorem 1 *Let φ be a T-STIT formula. Then, $\models_{\mathcal{K}} \varphi$ iff $\models_{\mathcal{F}} \varphi$.*

Theorem 2 *Let φ be a T-STIT formula. Then, $\models_{\mathcal{K}} \varphi$ iff $\models_{\mathcal{B}} \varphi$.*

PC	All tautologies of classical propositional calculus
S5(<i>i</i>)	All S5-principles for the operators [<i>i</i>]
S5(\Box)	All S5-principles for the operator \Box
S5(<i>Agt</i>)	All S5-principles for the operator [<i>Agt</i>]
KD4(G)	All KD4-principles for the operator G
K(H)	All K-principles for the operator H
($\Box \rightarrow i$)	$\Box\varphi \rightarrow [i]\varphi$
(<i>i</i> \rightarrow <i>Agt</i>)	$([1]\varphi_1 \wedge \dots \wedge [n]\varphi_n) \rightarrow [Agt](\varphi_1 \wedge \dots \wedge \varphi_n)$
(AIA)	$(\Diamond[1]\varphi_1 \wedge \dots \wedge \Diamond[n]\varphi_n) \rightarrow \Diamond([1]\varphi_1 \wedge \dots \wedge [n]\varphi_n)$
(Conv_{G,H})	$\varphi \rightarrow \mathbf{GP}\varphi$
(Conv_{H,G})	$\varphi \rightarrow \mathbf{HF}\varphi$
(Connected_G)	$\mathbf{PF}\varphi \rightarrow (\mathbf{P}\varphi \vee \varphi \vee \mathbf{F}\varphi)$
(Connected_H)	$\mathbf{FP}\varphi \rightarrow (\mathbf{P}\varphi \vee \varphi \vee \mathbf{F}\varphi)$
(NCUH)	$\mathbf{F}\Diamond\varphi \rightarrow \langle \mathbf{Agt} \rangle \mathbf{F}\varphi$
(AtmMom)	$p \rightarrow \Box p$ if $p \in \mathit{Atm}$
(MP)	$\frac{\varphi, \varphi \rightarrow \psi}{\psi}$
(IRR)	$\frac{(\Box\neg p \wedge \Box(\mathbf{G}p \wedge \mathbf{H}p)) \rightarrow \varphi}{\varphi}, \text{ provided } p \text{ does not occur in } \varphi$

Figure 2: Axiomatization of T-STIT

As with any other theorem of the paper, the proofs are presented in Appendix A, but we give a hint of the proof strategy here. Our proofs establish, first, that every Kamp agent model and every choice b-tree model determine a TKSM, and vice-versa and, second, that the correspondences preserve satisfiability. In a nutshell, we obtain the equivalence result among the different notions of validity by establishing a stronger result of correspondence among structures. We will use the same strategy of proof for Theorem 3 in Section 3.

As proved in [20, Lemma 9, p. 381], the notion of validity of the language $\mathcal{L}_{\text{T-STIT}}$ relative to TKSFs of Definition 1 and the notion of validity of $\mathcal{L}_{\text{T-STIT}}$ relative to *superadditive* TKSFs—in which the condition C3 is replaced by the weaker condition $\mathcal{R}_{Agt}(x) \subseteq \mathcal{R}_i(x)$ —are equivalent. Moreover, it is trivial to adapt the proof of Theorem 1 to show that the notion of validity of $\mathcal{L}_{\text{T-STIT}}$ relative to *superadditive* TKSFs and the notion of validity of $\mathcal{L}_{\text{T-STIT}}$ relative to *superadditive* Kamp agent frames—in which the condition TW3 is replaced by the weaker condition $\sim_{\langle t, Agt \rangle}(w) \subseteq \sim_{\langle t, i \rangle}(w)$ —are equivalent. Consequently, as we have emphasized in Section 2.3, the notion of validity relative to TKSFs and the notion of validity relative to *superadditive* TKSFs are equivalent when considering the language $\mathcal{L}_{\text{T-STIT}}$.

Figure 2 presents an axiomatization of T-STIT that is sound and complete w.r.t. to TKSFs. A proof of the completeness theorem can be found in [20].¹⁴ An interesting

¹⁴[20] proves completeness for a version of T-STIT that does not satisfy condition C8. The proof provided there can be easily adjusted in such a way that the resulting canonical model satisfies the

aspect of this axiomatization is that it exploits a variant of the Gabbay’s irreflexivity rule that has been widely used in the past for proving completeness results for different kinds of temporal logic in which time is supposed to be irreflexive (see, e.g., [13, 31, 25, 28]). The idea of the irreflexivity rule (**IRR**) is that the special kind of irreflexivity for the relation \mathcal{R}_G expressed by the Constraint C7 in Definition 1, although not definable in terms of an axiom, can be characterized in an alternative sense by means of the rule (**IRR**). This rule is perhaps more comprehensible if we consider its contrapositive: if $\neg\varphi$ is satisfiable and p does not occur in φ , then $\Box\neg p \wedge \Box(Gp \wedge Hp) \wedge \neg\varphi$ is satisfiable as well, because the temporal relation is irreflexive. Thus, (**IRR**) preserves validity.

An immediate consequence of Theorems 1 and 2 is that the axiomatization in Figure 2 is sound and complete also w.r.t. to Kamp agent frames and choice b-trees.

3 An interpreted system semantics for T-STIT

In this section we define a semantics for temporal STIT based on Fagin et al.’s notion of interpreted system [12]. The notion of interpreted system has been widely employed to provide semantics for epistemic temporal logics. We show that this semantics is equivalent to the temporal STIT semantics in terms of *discrete* temporal Kripke STIT frames *with initial point*.

The following definition introduces the basic components of interpreted systems.

Definition 8 (Global states, runs and systems) L_e is the set of possible (local) states for the environment and L_1, \dots, L_n are the sets of local states for the agents. $\mathcal{G} \subseteq L_e \times L_1 \times \dots \times L_n$ is a set of global states. A run over \mathcal{G} is a function from the natural numbers to \mathcal{G} . Runs are denoted by symbols r, r', \dots . We write $r_i(k)$ to denote the local state of agent i at the global state $r(k)$ and $r_e(k)$ to denote the local state of the environment at the global state $r(k)$. A system \mathcal{S} over \mathcal{G} is a set of runs over \mathcal{G} . Given a system \mathcal{S} , we call $\mathcal{P} = \mathcal{S} \times \mathbb{N}$ the set of points.

In the context of STIT, the set L_i of local states for agent i should be conceived as the set of actions that agent i can choose (i.e., agent i ’s action repertoire).

An interpreted system is nothing but a system supplemented with a valuation function that specifies the set of points in which a given propositional atom p is true.

Definition 9 (Interpreted system) The class of interpreted systems includes all tuples $\mathcal{I} = (\mathcal{S}, \omega)$ where \mathcal{S} is a system and ω is a valuation function $\omega : \text{Atm} \rightarrow 2^{\mathcal{P}}$.

The following definition introduces equivalence relations for the agents and for the environment.

Definition 10 (Equivalence relations for agents and the environment) Let $\mathcal{I} = (\mathcal{S}, \omega)$ be an interpreted system. Then, the equivalence relations \equiv_e and \equiv_i on \mathcal{P} for the environment and for agent i are defined as follows. For all $(r, k), (r', k') \in \mathcal{P}$:

- $(r, k) \equiv_e (r', k')$ iff $r_e(k) = r'_e(k')$ and $k = k'$,

condition that propositions are moment-determinate. This suffices to adapt the proof from [20] to the version of T-STIT that we present in this paper.

- $(r, k) \equiv_i (r', k')$ iff $(r, k) \equiv_e (r', k')$ and $r_i(k) = r'_i(k')$,
- $(r, k) \equiv_{Agt} (r', k')$ iff $(r, k) \equiv_i (r', k')$ for all $i \in Agt$.

where $(r, k) \equiv_e (r', k')$ means that the environment is the same at (r, k) and (r', k') . Moreover, for all $i \in Agt$, $(r, k) \equiv_i (r', k')$ means that (r, k) and (r', k') are choice-equivalent for agent i . Note that the equivalence relation \equiv_e for the environment is completely determined by the local state of the environment and the time,¹⁵ whereas the equivalence relation \equiv_i for an agent i is determined by the local state of the environment, the time and the local state of the agent. In other words, (r, k) and (r', k') are choice-equivalent for agent i if and only if, the environment, the time and i 's choice are the same at (r, k) and (r', k') .

In order to provide an interpreted system semantics for temporal STIT, we need the following three properties.

Definition 11 (Environment determinacy) *The interpreted system $\mathcal{I} = (\mathcal{S}, \omega)$ satisfies environment determinacy iff:*
for all $(r, k), (r', k') \in \mathcal{P}$ and for all $p \in Atm$, if $(r, k) \equiv_e (r', k')$ then $(r, k) \in \omega(p)$ iff $(r', k') \in \omega(p)$.

Definition 12 (Choice independence) *The interpreted system $\mathcal{I} = (\mathcal{S}, \omega)$ satisfies choice independence iff:*
for all $(r_1, k_1), \dots, (r_n, k_n) \in \mathcal{P}$, if $(r_i, k_i) \equiv_e (r_j, k_j)$ for all $i, j \in \{1, \dots, n\}$ then $\bigcap_{1 \leq i \leq n} \equiv_i (r_i, k_i) \neq \emptyset$.

Definition 13 (No choice between undivided runs) *The interpreted system $\mathcal{I} = (\mathcal{S}, \omega)$ satisfies no choice between undivided runs iff:*
for all $(r, k), (r', k') \in \mathcal{P}$, if $(r, k) \equiv_e (r', k')$ and $h < k$ then there exists $h' < k'$ such that $(r, h) \equiv_{Agt} (r', h')$.

Note that the previous property of no choice between undivided runs implies that if $(r, k) \equiv_e (r', k')$ and $k, k' > 0$, then $(r, k - 1) \equiv_{Agt} (r', k' - 1)$. More generally, we have that if $(r, k) \equiv_e (r', k')$ and $k, k' \geq h > 0$ then $(r, k - h) \equiv_{Agt} (r', k' - h)$. Since by Definition 10 $(r, k) \equiv_e (r', k')$ implies $k = k'$, we also have that if $(r, k) \equiv_e (r', k')$ and $k, k' \geq h > 0$ then $(r, k - h) \equiv_{Agt} (r', k' - h)$ and $k - h = k' - h$.

Environment determinacy is the counterpart of Constraint C8 in the definition of TKSM. It states that if the environment is the same at two points, then the two points agree on the atoms. In this sense, the environment determines the truth values of the atoms. *Choice independence* and *no choice between undivided runs* are respectively the counterparts of Constraints C2 and C6 in the definition of TKSM.

The following definition introduces the concept of STIT interpreted system.

Definition 14 (STIT interpreted system) *A STIT interpreted system is an interpreted system $\mathcal{I} = (\mathcal{S}, \omega)$ that satisfies environment determinacy, choice independence and no choice between undivided runs.*

¹⁵The equivalence relation \equiv_e depends on time since we assume that a complete description of the environment includes a specification of its temporal properties.

The truth conditions for formulae in $\mathcal{L}_{T\text{-STIT}}$ w.r.t. STIT interpreted systems are defined as follows:

$$\begin{aligned}
(\mathcal{I}, r, k) \models p &\Leftrightarrow (r, k) \in \omega(p) \\
(\mathcal{I}, r, k) \models \neg\varphi &\Leftrightarrow (\mathcal{I}, r, k) \not\models \varphi \\
(\mathcal{I}, r, k) \models \varphi \wedge \psi &\Leftrightarrow (\mathcal{I}, r, k) \models \varphi \text{ and } (\mathcal{I}, r, k) \models \psi \\
(\mathcal{I}, r, k) \models \Box\varphi &\Leftrightarrow \forall (r', k') \in \mathcal{P} : \text{if } (r, k) \equiv_e (r', k') \text{ then } (\mathcal{I}, r', k') \models \varphi \\
(\mathcal{I}, r, k) \models [i]\varphi &\Leftrightarrow \forall (r', k') \in \mathcal{P} : \text{if } (r, k) \equiv_i (r', k') \text{ then } (\mathcal{I}, r', k') \models \varphi \\
(\mathcal{I}, r, k) \models [Agt]\varphi &\Leftrightarrow \forall (r', k') \in \mathcal{P} : \text{if } (r, k) \equiv_{Agt} (r', k') \text{ then } (\mathcal{I}, r', k') \models \varphi \\
(\mathcal{I}, r, k) \models \mathbf{G}\varphi &\Leftrightarrow \forall k' \in \mathbb{N} : \text{if } k' > k \text{ then } (\mathcal{I}, r, k') \models \varphi \\
(\mathcal{I}, r, k) \models \mathbf{H}\varphi &\Leftrightarrow \forall k' \in \mathbb{N} : \text{if } k' < k \text{ then } (\mathcal{I}, r, k') \models \varphi
\end{aligned}$$

We compare STIT interpreted systems with a specific subclass of temporal Kripke STIT models, namely the class of temporal Kripke STIT models with discrete time and initial point.

Definition 15 (Temporal Kripke STIT model with discrete time and initial point)

A temporal Kripke STIT frame with discrete time and initial point (*DTKSM*) is a tuple $\mathcal{M}^{\mathcal{DK}} = (X, \mathcal{R}_\Box, \{\mathcal{R}_i\}_{i \in Agt}, \mathcal{R}_{Agt}, \mathcal{R}_X, \mathcal{V})$ where:

- X is a nonempty set of time-points;
- \mathcal{R}_\Box , every \mathcal{R}_i and \mathcal{R}_{Agt} are equivalence relations on X , \mathcal{R}_X is a serial and deterministic relation on X ;

and which satisfies the Constraints C1, C2, C3, C6, C7 and C8 in the definition of TKS \mathcal{M} , namely:

- (C1) for all $i \in Agt$: $\mathcal{R}_i \subseteq \mathcal{R}_\Box$;
- (C2) for all $x_1, \dots, x_n \in X$: if $x_i \mathcal{R}_\Box x_j$ for all $i, j \in \{1, \dots, n\}$ then $\bigcap_{1 \leq i \leq n} \mathcal{R}_i(x_i) \neq \emptyset$;
- (C3) for all $x \in X$: $\mathcal{R}_{Agt}(x) = \bigcap_{i \in Agt} \mathcal{R}_i(x)$;
- (C6) $\mathcal{R}_G \circ \mathcal{R}_\Box \subseteq \mathcal{R}_{Agt} \circ \mathcal{R}_G$;
- (C7) for all $x, y \in X$: if $y \in \mathcal{R}_\Box(x)$ then $y \notin \mathcal{R}_G(x)$;
- (C8) for all $x, y \in X$ and for all $p \in Atm$: if $x \mathcal{R}_\Box y$ then $x \in \mathcal{V}(p)$ iff $y \in \mathcal{V}(p)$;

plus the fact that \mathcal{R}_Y is deterministic and the following constraint:

- (C9) for all $x \in X$: there exists $y \in \mathcal{R}_H(x)$ such that $\mathcal{R}_H(y) = \emptyset$;

where $\mathcal{R}_G = \mathcal{R}_X^+$ is the transitive closure of \mathcal{R}_X , \mathcal{R}_H is the inverse relation of \mathcal{R}_G and \mathcal{R}_Y is the inverse relation of \mathcal{R}_X .

The main difference between TKSMs and DTKSMs is that in the latter the accessibility relation for the future \mathcal{R}_G is built from the accessibility relation \mathcal{R}_X for the ‘next’ time point. In particular, the accessibility relation for the future is defined to be the transitive closure of the accessibility relation for the ‘next’ operator. Note that the relation \mathcal{R}_X is serial and deterministic because every point is assumed to have *exactly* one temporal successor. Furthermore, note that the previous Constraint C6 implies the following property:

(C10) for all $x, y \in X$: if $x\mathcal{R}_\square y$ and $\mathcal{R}_H(x) = \emptyset$ then $\mathcal{R}_H(y) = \emptyset$.

The previous Constraints C6 and C7 together with the fact that \mathcal{R}_Y is deterministic imply the following property:

(C11) $\mathcal{R}_X \circ \mathcal{R}_\square \subseteq \mathcal{R}_{Agt} \circ \mathcal{R}_X$.

Finally, by the Constraint C1 and C3 together with the previous properties C10 and C11, we can prove by induction that the following property holds:

(C12) for all $x, y \in X$: if $x\mathcal{R}_\square y$ then there exists $x', y' \in X$ and $k \in \mathbb{N}$ such that $x\mathcal{R}_Y^k x', y\mathcal{R}_Y^k y'$ and $\mathcal{R}_H(x') = \mathcal{R}_H(y') = \emptyset$,

where \mathcal{R}_Y^k is the k -composition of the relation \mathcal{R}_Y defined inductively by (i) $\mathcal{R}_Y^0 = \{(x, x) | x \in X\}$, and (ii) $\mathcal{R}_Y^{k+1} = \mathcal{R}_Y^k \circ \mathcal{R}_Y = \mathcal{R}_Y \circ \mathcal{R}_Y^k$.

Condition **C12** tells that two points in the same moment are associated with the same time. Notice that the condition corresponds to property $k = k'$ in the definition of the equivalence relations \equiv_e from Definition 10.

Truth conditions of formulae relative to DTKSMs are the same as truth conditions of formulae relative to TKSMs given in Section 2.2.

Let $\models_{\mathcal{I}} \varphi$ and $\models_{\mathcal{DK}} \varphi$ denote that formula φ is valid relative to STIT interpreted systems and valid relative to DTKSMs, respectively. As the following theorem highlights, these two notions of validity are equivalent.

Theorem 3 *Let φ be a T-STIT formula. Then, $\models_{\mathcal{I}} \varphi$ iff $\models_{\mathcal{DK}} \varphi$.*

The previous result can be easily generalized to the class of *discrete* Kamp agent frames and to the class of *discrete* choice b-trees. The class of *discrete* Kamp agent frames includes all Kamp agent frames of Definition 2 satisfying the extra-condition that every ordered set $\langle T_w, <_w \rangle$ is isomorphic to the set of natural numbers. Similarly, the class of *discrete* choice b-trees includes all choice b-trees of Definition 6 satisfying the extra-condition that every history h is isomorphic to the set of natural numbers. From the class of *discrete* choice b-trees, the class of *discrete* BT+AC structures is easily defined by taking every *discrete* choice b-tree whose bundle coincides with the set of all histories. Let $\models_{\mathcal{DF}} \varphi$ and $\models_{\mathcal{DB}} \varphi$ denote that formula φ is valid relative to discrete Kamp agent frames and valid relative to discrete choice b-trees, respectively.

As the following two theorems highlight, DTKSMs, *discrete* Kamp agent frames and *discrete* choice b-trees provide equivalent semantics for the language \mathcal{L}_{T-STIT} . The proofs of these two theorems follow the general lines of the proofs of Theorems 1 and 2 given in the appendix. Thus, we do not give them here.

Theorem 4 *Let φ be a T-STIT formula. Then, $\models_{\mathcal{DK}} \varphi$ iff $\models_{\mathcal{DF}} \varphi$.*




Figure 3: Summary of relationships between semantics

Theorem 5 *Let φ be a T-STIT formula. Then, $\models_{DK} \varphi$ iff $\models_{DB} \varphi$.*

The previous Theorems 3, 4 and 5 together highlight the equivalence between the four semantics w.r.t. \mathcal{L}_{T-STIT} , namely the semantics in terms of DTKSMs, discrete Kamp agent frames, discrete choice b-trees and STIT interpreted systems.

4 Summary of the results

Figure 3 summarizes the relationships between the different semantics for T-STIT studied in the paper. The left side of the figure highlights the equivalence between the three semantics in terms of TKSMs, Kamp agent models and choice b-models. The right side highlights the equivalence between the different semantics for T-STIT in which time is assumed to be discrete. The inclusion between the two sides shows that the set of validities obtaining in the semantics from the left side is included in the set of validities obtaining in the semantics from the right side.

The figure also highlights that the set of validities for the semantics with not necessarily discrete time is strictly included in the set of validities for the semantics with discrete time. The following is an example of a formula of the language \mathcal{L}_{T-STIT} which is valid with respect to the discrete-time semantics presented in Section 3, in which the temporal order is isomorphic to the natural numbers, but which is not valid with respect to the semantics with not necessarily discrete time:

$$(\varphi \wedge H\varphi) \rightarrow FH\varphi \quad (1)$$

It is worth noting that while DTKSM-validity coincides with validity in discrete Kamp agent models and discrete choice b-models, DTKSM-validity and validity in discrete BT+AC models do not coincide. This immediately follows from the fact that bundled trees and full trees do not validate exactly the same formulae involving G and \Box . A well-known example discussed in [7] is the following formula:

$$\Box G(\varphi \rightarrow \Diamond F\varphi) \rightarrow \Diamond G(\varphi \rightarrow F\varphi) \quad (2)$$

This formula is valid on full trees with discrete time, but not on bundled trees with discrete time. Since the formula is also in $\mathcal{L}_{\text{T-STIT}}(\text{Atm}, \text{Agt})$, the result immediately extends to discrete choice b-validity (DTKSM-validity, discrete Kamp agent-validity) and discrete BT+AC validity. Figure 4 illustrates a DTKSM where $\Box G(\varphi \rightarrow \Diamond F\varphi) \rightarrow \Diamond G(\varphi \rightarrow F\varphi)$ fails.




Figure 4: Example of DTKSM that falsifies $\Box G(p \rightarrow \Diamond Fp) \rightarrow \Diamond G(p \rightarrow Fp)$

In order to show that the preceding formula (2) is valid on discrete BT+ACs, assume that $\mathcal{M}^{\mathcal{B}}, \langle m_0, h \rangle \models \Box G(\varphi \rightarrow \Diamond F\varphi)$ for an arbitrary discrete BT+AC model $\mathcal{M}^{\mathcal{B}}$, moment m_0 and history h passing through m_0 . This semantically means that:

- (i) for all $h' \in B_{m_0}$ and for all $m \in h'$ if $m_0 \prec m$ and $\mathcal{M}^{\mathcal{B}}, \langle m, h' \rangle \models \varphi$ then there exists $h'' \in B_m$ and $m' \in h''$ such that $m \prec m'$ and $\mathcal{M}^{\mathcal{B}}, \langle m', h'' \rangle \models \varphi$.

We distinguish two cases. As for the first case, let us assume that $\mathcal{M}^{\mathcal{B}}, \langle m_0, h \rangle \models \Box G\neg\varphi$. The latter implies $\mathcal{M}^{\mathcal{B}}, \langle m_0, h \rangle \models \Diamond G\neg\varphi$ which in turn implies $\mathcal{M}^{\mathcal{B}}, \langle m_0, h \rangle \models \Diamond G(\varphi \rightarrow F\varphi)$. As for the second case, let us assume that $\mathcal{M}^{\mathcal{B}}, \langle m_0, h \rangle \models \Diamond F\varphi$, which semantically means:

- (ii) there exists $h' \in B_{m_0}$ and a lower set X of $\mathcal{F}(h', m_0)$ such that X has order type 1 and $\mathcal{M}^{\mathcal{B}}, \langle m, h' \rangle \models F\varphi$ for all $m \in X$,

where $\mathcal{F}(h', m_0) = (\{m_0\} \cup \{m \in h' \mid m_0 \prec m\}, \prec)$ and where X is a lower set of $\mathcal{F}(h', m_0)$ iff X is a subset of $\mathcal{F}(h', m_0)$ such that for all $m \in X$ and for all $m' \in \mathcal{F}(h', m_0)$, if $m' \prec m$ then $m' \in X$. For every $k \in \mathbb{N}$, let

$$Z_{m_0, k, \varphi}$$

be the set of all pairs (h, X) such that $h \in B_{m_0}$, X is a lower set of $\mathcal{F}(h, m_0)$ of order type k , and $\mathcal{M}^{\mathcal{B}}, \langle m, h \rangle \models F\varphi$ for all $m \in X$.

Item (i) together with the fact that $\mathcal{M}^{\mathcal{B}}$ is a discrete BT+AC model based on a full tree imply that the following holds for every $k \in \mathbb{N}$:

(iii) for every $(h, X) \in Z_{m_0, k, \varphi}$, there exists $(h', X') \in Z_{m_0, k+1, \varphi}$ such that $X \subset X'$.

Items (ii) and (iii), together with the axiom of choice, guarantee that we can find $h_1, h_2, \dots \in B_{m_0}$ and $X_1, X_2, \dots \subseteq M$ such that, for every $k \in \mathbb{N}$:¹⁶

- $X_k \subset X_{k+1}$,
- X_k has order type k ,
- X_k is a lower set of $\mathcal{F}(h_k, m_0)$, and
- $\mathcal{M}^{\mathcal{B}}, \langle m, h_k \rangle \models \text{F}\varphi$ for all $m \in X_k$.

Let

$$h_X = \bigcup_{k \in \mathbb{N}} X_k.$$

Clearly, h_X is a history in B_{m_0} such that $\mathcal{M}^{\mathcal{B}}, \langle m, h_X \rangle \models \text{F}\varphi$ for all $m \in \mathcal{F}(h_X, m_0)$. The latter means that $\mathcal{M}^{\mathcal{B}}, \langle m_0, h_X \rangle \models \text{GF}\varphi$. From this, we get $\mathcal{M}^{\mathcal{B}}, \langle m_0, h \rangle \models \diamond \text{GF}\varphi$, which in turn implies $\mathcal{M}^{\mathcal{B}}, \langle m_0, h \rangle \models \diamond \text{G}(\varphi \rightarrow \text{F}\varphi)$.

Other examples of formulae of the language $\mathcal{L}_{\text{T-STIT}}(\text{Atm}, \text{Agt})$ that are valid on discrete BT+AC models but not on DTKSMs, discrete Kamp agent models and discrete choice b-models are the following variants of the previous formula (2) in which the operator \diamond in the antecedent is replaced either by the operator $\langle i \rangle$ or by the operator $\langle \text{Agt} \rangle$:

$$\Box \text{G}(\varphi \rightarrow \langle i \rangle \text{F}\varphi) \rightarrow \diamond \text{G}(\varphi \rightarrow \text{F}\varphi) \quad (3)$$

$$\Box \text{G}(\varphi \rightarrow \langle \text{Agt} \rangle \text{F}\varphi) \rightarrow \diamond \text{G}(\varphi \rightarrow \text{F}\varphi) \quad (4)$$

5 Extension with full groups

As emphasized in Section 2.5, Lemma 9 from [20, p. 381] shows that we get the same logic, namely T-STIT, if we impose *superadditivity* instead of *additivity* on the construction of \mathcal{R}_{Agt} in TKSFs. Just to recall these notions: in TKSFs, the former equates with the condition that, for all $x \in X$, $\mathcal{R}_{\text{Agt}}(x) \subseteq \bigcap_{i \in \text{Agt}} \mathcal{R}_i(x)$, while the latter equates with the condition C3 from Definition 1 in Section 2.2: $\mathcal{R}_{\text{Agt}}(x) = \bigcap_{i \in \text{Agt}} \mathcal{R}_i(x)$ for all $x \in X$.

Things change if we extend TKSFs with choice-equivalence relations for all the definable groups and corresponding group STIT operators, that is, all the members of $2^{\text{Agt}^*} = 2^{\text{Agt}} \setminus \emptyset$ in each TKSF. In this section, we briefly explore how this extension affects the equivalence results from Subsection 2.5.

¹⁶The axiom of choice is used here to ensure that we can choose the element (h_1, X_1) from $Z_{m_0, 1, \varphi}$, which is guaranteed to be non-empty because of item (ii). Moreover, for every $k \in \mathbb{N}$, we use the axiom of choice to choose the element (h_k, X_k) from the set $\{(h, X) \in Z_{m_0, k, \varphi} \mid X_{k-1} \subset X\}$, which is guaranteed to be non-empty because of item (iii).

Let us extend $\mathcal{L}_{\text{T-STIT}}$ with many group STIT operators $[J], [I], \dots$, for the members $J, I \dots$ of 2^{Agt^*} resulting into the following language $\mathcal{L}_{\text{T-GSTIT}}(\text{Atm}, \text{Agt})$ of the logic T-GSTIT (temporal group STIT):

$$\varphi ::= p \mid \neg\varphi \mid \varphi \wedge \varphi \mid [J]\varphi \mid \Box\varphi \mid \mathbf{G}\varphi \mid \mathbf{H}\varphi$$

where p ranges over Atm and J ranges over 2^{Agt^*} .

Moreover, let us define two classes of group TKSFs: *additive* group TKSFs and *super-additive* group TKSFs. *Additive* group TKSFs extend TKSFs of Definition 1 with one relation \mathcal{R}_J per group J in 2^{Agt^*} and generalize the condition C3 of Definition 1 with the following condition:

$$(\mathbf{C3}') \text{ for all } x \in X \text{ and } J \in 2^{\text{Agt}^*}: \mathcal{R}_J(x) = \bigcap_{i \in J} \mathcal{R}_i(x).$$

We denote the class of additive group TKSFs by \mathcal{AGK} .

Super-additive TKSFs consists in generalizing the class of frames introduced in [5] with accessibility relations for the temporal operators \mathbf{G} and \mathbf{H} . Specifically, *super-additive* group TKSFs extend TKSFs of Definition 1 with one relation \mathcal{R}_J per group J after replacing the condition C1 with the following condition:

$$(\mathbf{C1}') \text{ for all } J \in 2^{\text{Agt}^*}: \mathcal{R}_J \subseteq \mathcal{R}_{\Box},$$

and the condition C3 with the following two conditions:

$$(\mathbf{C3}'') \text{ for all } I, J \in 2^{\text{Agt}^*}: \text{if } I \subseteq J \text{ then } \mathcal{R}_J \subseteq \mathcal{R}_I,$$

$$(\mathbf{C3}''') \text{ for all } J \in 2^{\text{Agt}^*} \text{ such that } J \neq \text{Agt}: \mathcal{R}_{\Box} \subseteq \mathcal{R}_J \circ \mathcal{R}_{\text{Agt} \setminus J}.$$

We denote the class of super-additive group TKSFs by \mathcal{SGK} .

Condition C1' generalizes Condition C1 in Definition 1 to groups. Condition C3'' is the condition of *coalitional monotonicity* meaning that bigger groups have finer choice partitions. Condition C3''' just means that every alternative can be attained by combining the current choice of a group and the appropriate choice of its complement. Conditions C1', C3'' and C3''' are also properties of additive group TKSFs. Indeed, Condition C1' is entailed by Condition C1 in Definition 1 and Condition C3' above. Condition C3'' is entailed by Condition C3'. Condition C3''' is entailed by Condition C2 in Definition 1, Condition C3' above and the fact that every \mathcal{R}_J is an equivalence relation. To prove this, suppose $x \mathcal{R}_{\Box} y$ and without loss of generality let $J = \{1, \dots, k\}$ and $\text{Agt} \setminus J = \{k, \dots, n\}$ with $k < n$. By Condition C2 in Definition 1, we have $\mathcal{R}_1(x) \cap \dots \cap \mathcal{R}_k(x) \cap \dots \cap \mathcal{R}_k(y) \cap \dots \cap \mathcal{R}_n(y) \neq \emptyset$. By Condition C3' the latter is equivalent to $\mathcal{R}_J(x) \cap \mathcal{R}_{\text{Agt} \setminus J}(y) \neq \emptyset$. Since \mathcal{R}_J and $\mathcal{R}_{\text{Agt} \setminus J}$ are equivalence relations, it follows that $x \mathcal{R}_J \circ \mathcal{R}_{\text{Agt} \setminus J} y$. This shows that the class of *additive* group TKSFs is indeed a subclass of the class of *super-additive* group TKSFs.

Truth conditions of formulae of the language $\mathcal{L}_{\text{T-GSTIT}}(\text{Atm}, \text{Agt})$ are as the ones given in Section 2.2 for the boolean operators, the temporal operators and the operator \Box plus the following one for the group STIT operator $[J]$:

$$\mathcal{M}^\tau, x \models [J]\varphi \iff \forall y \in \mathcal{R}_J(x) : \mathcal{M}^\tau, y \models \varphi$$

where τ is either \mathcal{AGK} or \mathcal{SGK} depending on whether we evaluate formulae relative to additive group TKSFs or to super-additive group TKSFs.

There exists an atemporal STIT formula which distinguishes *additive* group TKSFs from *super-additive* ones. Specifically, the following formula is valid w.r.t. *additive* group TKSFs and invalid w.r.t. *super-additive* group TKSFs:

$$[J][I]\varphi \rightarrow [J \cap I]\varphi \quad (5)$$

The proof of the validity requires just a straightforward adaptation of the right-left direction from [17, Theorem 13], to which we refer the reader. The following *super-additive* group TKSM shows that the formula is invalid w.r.t. *super-additive* group TKSFs. Let $\text{Agt} = \{1, 2, 3\}$ and let $\text{Atm} = \{p\}$. Moreover, let $M^{\text{SGK}} = (X, \mathcal{R}_\square, \{\mathcal{R}_J\}_{J \in 2^{\text{Agt}^*}}, \mathcal{R}_G, \mathcal{R}_H)$ be such that:

- $X = \bigcup_{\alpha \in \mathbb{N}} X_\alpha$ with $X_\alpha = \{x_\alpha, y_\alpha\}$ for all $\alpha \in \mathbb{N}$,
- $\mathcal{R}_\square = \mathcal{R}_{\{1\}} = \mathcal{R}_{\{2\}} = \mathcal{R}_{\{3\}} = X_1 \times X_1 \cup \bigcup_{\alpha > 1} \{(x_\alpha, x_\alpha), (y_\alpha, y_\alpha)\}$,
- $\mathcal{R}_{\{1,2\}} = \mathcal{R}_{\{1,3\}} = \mathcal{R}_{\{2,3\}} = \mathcal{R}_{\{1,2,3\}} = \bigcup_{\alpha \in \mathbb{N}} \{(x_\alpha, x_\alpha), (y_\alpha, y_\alpha)\}$,
- $\mathcal{R}_G = \{(x_\alpha, x_\beta) | \alpha, \beta \in \mathbb{N} \text{ and } \alpha \leq \beta\} \cup \mathcal{R}_H = \{(y_\alpha, y_\beta) | \alpha, \beta \in \mathbb{N} \text{ and } \alpha \leq \beta\}$,
- $\mathcal{V}(p) = \{x_\alpha | \alpha \in \mathbb{N}\}$, and
- \mathcal{R}_H is the inverse relation of \mathcal{R}_G .

It is easy to check that $M^{\text{SGK}}, x_1 \models \{[1, 2]\}[1, 3]p$ but $M^{\text{SGK}}, x_1 \not\models \{[1]\}p$.

Of course, we can adapt the definitions above to obtain additive and super-additive group Kamp agent frames and choice b-trees, respectively. It is straightforward to check that the previous observation also applies to the distinction between additive Kamp agent frames (choice b-trees) and super-additive Kamp agent frames (choice b-trees), namely the fact that formula $[J][I]\varphi \rightarrow [J \cap I]\varphi$ is valid w.r.t. the former and invalid w.r.t. the latter.

Besides, it is easy to adapt the proofs of Theorem 1 and Theorem 2 to the notions of validity in additive group TKSFs, additive group Kamp agent frames, additive group choice b-trees, and to the super-additive versions of such structures.

This implies that the equivalence results are preserved if the structures that we are taking into account are upgraded to choices of groups by adopting additivity (or super-additivity) for all of them. By contrast, equivalence is lost if we compare any frame for group agency with additivity with frames with super-additivity.

This does not undermine the generality of our results, however. The comparison we carried in this paper aims at proving that we obtain the same logic if we vary over three different *temporal* structures while keeping the same principle of agency *fixed*. Thus, the fact that super-additivity and additivity determine two different logics does not affect the generality of the results we have established.

The full *superadditive* group STIT logic from [4] and the full *additive* group STIT logics based on $\mathcal{L}_{\text{T-GSTIT}}$ enjoy different formal properties. Indeed, [4] presents a finite axiomatization of the former, whose decidability is still an open issue, while the latter is undecidable and not finitely axiomatizable [17]. Also, decidability of the full (superadditive) group logic from [4] is still an open issue.

It is easy to see why the negative results from [17] do not apply, as they stand, to the superadditive logic from [4]. The results from [17] establish that, for $k \geq 3$, the *product logic* $S5^k$ can be embedded in an *additive* T-GSTIT with $k \geq 3$ agents. Since $S5^k$ is undecidable and not finitely axiomatizable for $k \geq 3$, the embedding transfers these properties to any T-GSTIT with $k \geq 3$. By contrast, the group operators from [4] are KD-type operators¹⁷ and do not verify the so-called Church-Rosser axiom, which prevents them from embedding any product logic. This, the results from [17] have no implications for the logic from [4].

Another open issue concerns the role that superadditivity and additivity are supposed to play in the formal properties of the two logics. In particular, [26] conjectures that a non-temporal fragment of full group T-GSTIT is decidable if and only if the fragment does not capture the difference between additive and superadditive models. If this is correct, then superadditive T-GSTIT should be undecidable. By contrast, a conjecture from [21] is that the non-temporal fragment of *superadditive* T-GSTIT should be decidable. Even if this would be true, however, the conjecture from [21] presupposes a language that is weaker than $\mathcal{L}_{\text{T-GSTIT}}$, and then the conjecture has no direct implications of the decidability of superadditive T-GSTIT. In light of this, we feel that our choice to focus on the language $\mathcal{L}_{\text{T-STIT}}$ rather than $\mathcal{L}_{\text{T-GSTIT}}$ is justified.

6 Discussion

The alternative semantics from Section 2 are appealing because they are convenient mathematical objects. Indeed, the semantics based on TKSFs quantifies over time-points only, and that based on Kamp agent frames quantifies over times and worlds. These entities are introduced as primitives, not as defined set-theoretical objects. As a consequence, the two semantics keep quantification at first-order level and this in turn allows for the application of convenient techniques, such as Sahlqvist techniques for proving completeness, that do not apply to more common structures for indeterminist time and agency, such as BT+ACs.

As for the semantics based on choice b-trees, they quantify over sets of primitives (the histories from Definition 4). Thus they involve a second-order quantification, which prevents application of Sahlqvist techniques and usually brings some complexity in the Henkin construction for completeness results. However, bundling is basically the kind of move that defines a Henkin semantics for second-order logic: by performing it, we restrict the domain over which the variables range to a possibly *proper* subset of the entire domain in question (in our case, the maximal \prec -chains defined on a tree). As it is well known, this move downgrades expressivity of second-order quantification to that of many-sorted first-order quantification (we refer the reader to [10] and [29] for this). This is, in turn, the kind of quantification we have in Kamp agent frames.

Summing up, we get clear technical benefits from the alternative semantics from Section 2. However, little attention has been paid to their *descriptive adequacy*—with a significant exception that we discuss later in this section. One natural question in this respect is ‘Do the alternative semantics presented here validate desirable principles

¹⁷This is due to the fact that the operators from [4] combine quantification over choices with quantification over next moments

of agency and time?’ The axiomatization in Figure 2 gives a positive answer to this question: the alternative semantics validate established principles for time, modality, and seeing to it that, while providing natural mixing principles.

Here we approach two other questions concerning descriptive adequacy:

1. Can the semantics from Section 2 prove better than BT+ACs in describing some scenarios of multi-agent interaction?
2. Do their ontologies provide a sound picture of agency and time?

We answer these two questions in the following two subsections. In order to do this, we need to talk about choices, and thus we will extend our language with special choice-formulas. These allow us to express what action the relevant agent is choosing and prove to be special STIT formulae that hold exactly at the worlds that are compatible with a given choice. A similar device is implicit in [18], where choices are often expressed through formulas—see for instance [18, 55-58].

6.1 Adequacy in capturing game scenarios

We discuss a game-theoretical scenario that can be expressed if DTKSMs of Section 3 are the structures on which our language is interpreted, while it cannot be expressed if discrete BT+ACs are used. Our scenario will be a particular Infinitely Repeated Prisoner’s Dilemma (IRPD) with two agents 1 and 2 where: (i) agent 1 exploits a *grim strategy*, in the sense that she starts cooperating and stops cooperating at some *undetermined* future time-point after the other agent decides to defect; (ii) agent 2 is a *free rider*, in the sense that she will necessarily defect at some future time-point; (iii) there is no specific threshold after which agent 1 stops cooperating forever in case agent 2 decides to defect. Feature (iii) corresponds to a kind of indeterminism that can be modeled via DTKSMs and, consequently, via discrete choice b-trees and discrete Kamp agent frames but not via discrete BT+AC structures. This is what we are going to show in the rest of the section.

IRPD is an Infinitely Repeated Game (IRG), that is a repetition of a *constituent game*, in our case the Prisoner’s Dilemma (PD), with no last game-round. Moves are simultaneous. In PD, each agent i from the set $\{1, 2\}$ has two actions available: either she cooperates with the other agent by not confessing the crime (c) or she defects by confessing (d).¹⁸ The formal definition of a PD implies that (d, d) is a Nash Equilibrium, that it is not Pareto optimal, and that d is a strictly dominating action for each agent.

The basic dynamic at stake with IRPD is well known: in the one-shot PD, each agent $i \in \{1, 2\}$ has an incentive to play c —which would determine a Pareto optimal state¹⁹ and she is somehow ‘punished’ by the other if she plays d . In particular, this implies that each agent will defect if the other breaks cooperation first. There are many possible strategies to enact this kind of punishment. In our example, we will encounter

¹⁸We are not interested in assigning specific utilities to the action profiles: any specification that respects the definition of PD will do. The crucial point of the definition is that for each $i \in \{1, 2\}$, we have $u_i(d, c) > u_i(c, c) > u_i(d, d) > u_i(c, d)$, where u is a utility function from action-profiles to real numbers.

¹⁹In order to get this, IRPD imposes the additional condition $2u(c, c) > u(d, c) + u(c, d)$, which makes cooperation better than alternating cooperation and defection. The condition is often found in standard examples of the single-round PD, but it is not implied by its definition.

a variant of the so-called *grim strategy*: one starts cooperating and, after one or more defections by the other agent, she stops cooperating and never gets back to cooperation.

In order to formally represent the example, we need to talk about choices over alternative actions in $\mathcal{L}_{\text{T-STIT}}(Atm, Agt)$. To this aim, we extend our language with syntactical elements that describe the action chosen by a certain agent. This game-theoretical extension of STIT is interesting in its own right and displays connections with other recent work in STIT, such as [16], but here we will limit ourselves to introducing the semantical basics and their adequacy with respect to minimal requirements. We wish to provide a more structured insight on the extension of $\mathcal{L}_{\text{T-STIT}}(Atm, Agt)$ in future research.

For every agent $i \in Agt$, we introduce a non-empty set Act_i . Elements of Act_i are the actions available to (or choosable by) agent i . Moreover, we extend the definition of DTKSM of Section 3 with a function

$$play_i : X \longrightarrow Act_i$$

for each agent i that associates every time-point x to the action that agent i chooses at x . We impose the following condition on the function $play_i$: for all $x, y \in X$, $x \mathcal{R}_i y$ if and only if $play_i(x) = play_i(y)$. The condition allows us to label an agent's choice with the name of the action chosen by the agent: two time-points x and y belong to the same choice of agent i , if and only if agent i chooses the same action at x and y .

Once actions are defined in our (extended) DTKSMs, we extend the language $\mathcal{L}_{\text{T-STIT}}(Atm, Agt)$ with the resources to express them. We extend the language by choice-formulae of the form $\text{pl}(i, a)$ where i ranges over Agt and a ranges over Act_i and which has to read "agent i chooses the action a ". Truth conditions of choice-formulae are given by the following clause:

$$\mathcal{M}^{\mathcal{DK}}, x \models \text{pl}(i, a) \iff play_i(x) = a$$

It is clear from the semantics that the formulae $\text{pl}(i, a) \rightarrow [i]\text{pl}(i, a)$ and $(\text{pl}(1, a_1) \wedge \dots \wedge \text{pl}(n, a_n)) \rightarrow [Agt](\text{pl}(1, a_1) \wedge \dots \wedge \text{pl}(n, a_n))$ are valid. These two formulae are expected by a suitable encoding of properties of games into a logic of agency. The first states that an agent's action takes place because it is chosen by the agent. The second states that a joint action of a coalition takes place because it is chosen by the agents in the coalition. The first validity, together with Axiom T for $[i]$, highlights that the choice-formula $\text{pl}(i, a)$ expresses agency, since it holds if and only if agent i *sees to it that* it holds.

Notice that choice-formulae $\text{pl}(i, a)$ are not *atoms*. Indeed, $\text{pl}(i, a) \rightarrow \Box\text{pl}(i, a)$ does not hold in general, contrary to what Condition C8 from Definition 15 imposes on atoms. This difference between atoms and the new choice-formulae is justified by the fact that choices have a 'trace of futurity' somehow: a STIT choice is a way to constrain the future course of events into a subclass of such courses.

Let us now come back to the particular case of IRPD that we want to describe by means of $\mathcal{L}_{\text{T-STIT}}$ *plus* the new choice-formulae. As emphasized above, it is characterized by three features: (i) necessarily, agent 1 starts cooperating but in case of defection of agent 2, sooner or later she will stop cooperating forever; (ii) agent 2 will necessarily defect at some future time-point; (iii) there is no specific threshold after which agent

1 stops cooperating forever in case agent 2 decides to defect. Features (i) and (ii) together imply that, necessarily, at some future time-point agent 1 will cooperate and will stop cooperating forever afterwards. In our extended DTKSMs, we can capture this in formal terms as follows:

- (1) for every $y \in \mathcal{R}_\square(x)$, there exists $z \in \mathcal{R}_G(y)$ such that $play_1(z) = c$ and, for every $u \in \mathcal{R}_G(z)$, $play_1(u) = d$.

The kind of indeterminism expressed by feature (iii) guarantees that, for every possible time-line ℓ_x in which agent 1 will stop cooperating at a certain point, there exists an alternative time-line ℓ_y in which agent 1 will stop cooperating later than in ℓ_x . A time-line clearly corresponds to an infinite sequence of action profiles of the constituent PD, an action profile being a tuple of actions one per each agent in the game. In our extended DTKSMs, we can capture this in formal terms as follows:

- (2) for every $y \in \mathcal{R}_\square(x)$ and for every $z \in \mathcal{R}_G(y)$, if $play_1(z) = c$ then, there exists $u \in \mathcal{R}_\square(z)$ and $v \in \mathcal{R}_G(u)$, such that $play_1(v) = c$.

At the syntactic level the previous conditions (1) and (2) are expressed by the following formula:

$$\square F(\text{pl}(1, c) \wedge G\neg\text{pl}(1, c)) \wedge \square G(\text{pl}(1, c) \rightarrow \diamond F \text{pl}(1, c))$$

As observed in Section 4, such formula is satisfiable in DTKSMs, in discrete choice b-trees and in discrete Kamp agent frames but not in discrete BT+ACs. This highlights that to model the kind of indeterminism expressed by the previous feature (iii), one would need to exclude from the tree of the repeated game the branch that would otherwise arise as the limit of all those *finite* repetitions of cooperation by agent 1, that agent 1 does not wish to exclude, because the threshold after which agent 1 stops cooperating forever has not been determined yet.

Discussing the example Here, we discuss why, in our view, the ability to provide a model fulfilling conditions (i)–(iii) is a virtue of DTKSMs over discrete BT+ACs. In particular, we show that DTKSMs can accommodate two equally legitimate views on the use of formal modeling, while discrete BT+ACs rule out one of these views.

A crucial point of our example is that conditions (i)–(ii) exclude the option that, for every subsequent round of the IRG, agent 1 will cooperate again if she is cooperating at all. Accordingly, the relevant DTKSM does not satisfy the formula $\diamond G(\text{pl}(1, c) \rightarrow F \text{pl}(1, c))$. Of course, it is *compatible* with the *rules* of the game that agent 1 keeps (or comes back to) cooperating, and so one could wonder why we should exclude the option.

The answer to this point depends on the *modeling purposes* one has. For instance, in proposing a formal model for a given IRG, we may want to *predict* how a given *instance* of the repeated game in question can evolve. It is reasonable to say that our predictions will depend on *the information we* (external modelers) *have about the given instance of the IRG in question*. This information will include the rules of the game (actions available and resulting action-profiles, rationality of the players, and so on), but it can also include hard information on the kind of players we are modeling. This is what

happens in our example, where condition (ii) tells us that agent 2 is a *free rider*, and condition (i) tells us that agent 1 reacts to defection by stopping cooperation, sooner or later. If we (modelers) receive this information—as we do in the example above—then we know that the scenario where agent 1 keeps (or comes back to) cooperating cannot hold, since it is incompatible with conditions (i)–(ii). As a consequence, if our aim is to provide a model that predicts how our particular example—which includes (i)–(iii)—can evolve, then excluding the scenario where agent 1 keeps (or comes back to) cooperating (if she is cooperating at all) is a legitimate move.²⁰

A different answer may come if we assume some other perspective on the use of formal models. For instance, one can hold the view that our models should represent how the choices of different agents operate on a given set of objective possibilities. These are, in turn, possibilities that are due to how the world is at a given moment. The intuition here is that, at each moment, the choices of the agents act so as to exclude some of the objective possibilities that are defined in the model. This view requires that all such possibilities are represented in the model, as a background against which the choices are made.²¹ In the case of an IRG, what is objectively possible is determined by the rules of the game, and the view above turns to imply that we should capture everything that is *compatible* with such rules. From this point of view, excluding the above scenario—where agent 1 keeps (or comes back to) cooperating at some point—is not a legitimate move, since the scenario is compatible with the rules of IRPDs.

Thus, we have (at least) two different views on what we do when we model a game (or some other kind of multi-agent interaction, in principle), and they suggest different answers to our initial question. It is reasonable to hold that both views are legitimate in their own right. There is nothing wrong in building a model in order to predict (or describe) a specific instance of a phenomenon, and there is nothing wrong with a more general focus that wishes to cast agency against *all* the objective possibilities.

Consistently with this, we believe that we should opt for the modeling tool that can accommodate both views. From this point of view, DTKSMs score better than discrete BT+ACs. Indeed, we have DTKSMs where—just to stick to our example— $\Box G(\text{pl}(1, c) \rightarrow \Diamond F \text{pl}(1, c))$ is satisfied and $\Diamond G(\text{pl}(1, c) \rightarrow F \text{pl}(1, c))$ is not. This captures the situations where conditions (i)–(iii) are taken into account, and it fares well with the first methodological view on modeling that we have introduced. However, we also have DTKSMs where both $\Box G(\text{pl}(1, c) \rightarrow \Diamond F \text{pl}(1, c))$ and $\Diamond G(\text{pl}(1, c) \rightarrow F \text{pl}(1, c))$ are satisfied—which fares well with the second methodological view that we have introduced. This follows from the fact that every discrete choice b-tree corresponds to a DTKSM, and that this correspondence preserves satisfiability—see proof of Theorem 2 in the Appendix. Since discrete BT+ACs are special cases of discrete choice b-trees, there are DTKSMs satisfying $\Box G(\text{pl}(1, c) \rightarrow \Diamond F \text{pl}(1, c)) \rightarrow \Diamond G(\text{pl}(1, c) \rightarrow F \text{pl}(1, c))$.

By contrast, discrete BT+ACs rule out the possibility to proceed along the lines of the first methodological view. Since they make $\Box G(\text{pl}(1, c) \rightarrow \Diamond F \text{pl}(1, c)) \rightarrow \Diamond G(\text{pl}(1, c) \rightarrow F \text{pl}(1, c))$ valid, they do not let us model a scenario where conditions

²⁰Notice that the exclusion of the scenario above recalls the application of *integrity constraints* [22] in Artificial Intelligence—in particular, in *knowledge representation*. These constraints select away some states in a structure as non-admitted in the description of a given phenomenon.

²¹We thank one anonymous referee for pointing this perspective to us.

(i)–(iii) are satisfied—or, more generally, a scenario where these conditions are admitted as a part of the game description.

In sum, DTKSMs accommodate both methodological views, and let us switch from the one to the other in accordance with our modeling purposes of the moment. By contrast, discrete BT+ACs can accommodate just one of the two views above. Reasonably, a toolkit that is flexible enough to fit different (legitimate) modeling purposes is to be preferred over a toolkit that fits just one (legitimate) modeling purpose. As a consequence, we believe that the expressive ability exemplified in the example of the IRPD above is a virtue of DTKSMs over discrete BT+ACs.

An analogous discussion in temporal logic A comparison between *trees* and *Kamp frames*—or equivalent structures—has been presented in temporal logic by Belnap and colleagues—see [2, 7A.6]. The example from [2] concerns a particular radium atom a that, at each moment, may decay or not.²² The example assumes that, for every moment, if a has not yet decayed, then it is possible that will decay in the next moment, and it is possible that a will not decay in the next moment.²³ A consequence of this is that (\star) every sequence of ‘no-decay moments’ can be extended by at least one moment. This situation is compatible with *full trees*, on the one hand, and *Ockhamist frames*, *Kamp frames*, and *bundled trees*, on the other. Consider now the reasonable sentence $(\star\star)$ ‘at the starting moment, it is inevitable that a will decay after a finite sequence of moments’.²⁴ As noticed in [2], (\star) and $(\star\star)$ can be satisfied together in an Ockhamist or Kamp frame, or in a bundled tree, but not in a full tree.

[2, 7A.6] reacts to this by deploying the second methodological view that we have introduced above: they seem to imply that $(\star\star)$ is an objective possibility and that, as such, it should be accounted for.²⁵ We do not disagree with [2]. Simply, we believe that their insight should not be generalized to every situation one could wish to model. After all, the discussion from the previous section has shown that we might want to go with a different methodological view, which may on specific occasions require structures that are not based on full trees.

6.2 Sound picture of agency and time

The above discussion seems to shed some light on the picture of agency and time in the different semantics presented here, and their respective soundness. BT+AC include *trees*, which are usually assumed as a sound picture of time [2, 27]. Also, we can associate *moments* of a BT+AC with (single instances of the) constituent game from an IRG and *histories* being associated with maximal and linear *sequences of repeated constituent games*. However, as we have shown in the previous section, BT+ACs rule

²²The example from [2] involves a continuous time, with the relevant moments being marked by ticks—which can be in turn be put in correspondence with the naturals. Here, we simplify the scenario by assuming that, for every history, the moments in the history can be put in correspondence with the naturals.

²³In symbols, the assumption implies $\Box G(\neg p \rightarrow (\Diamond F p \wedge \Diamond F \neg p))$, where $\neg p$ is ‘atom a does not decay’.

²⁴In symbols, this assumption implies $\Box F(\neg p \wedge G p)$.

²⁵Another reaction, due to both [2] and [27], is that (\star) and $(\star\star)$ are mutually incompatible. We agree with [32] that this view actually *presupposes* the choice of full trees over Kamp frames, Ockhamist frames, and bundled trees, and thus it does not help settle the issue.

out some scenarios that are consistent with the definition of an IRPD, and this somehow limits the soundness of the picture of agency they offer.

One natural question is whether the alternative semantics from Section 2 can capture entities from IRGs and temporal entities in a sound way. In this respect, choice b-trees seem to score some good points: they provide the same intuitive picture of time as BT+ACs and model some game-theoretical scenario better than BT+ACs do—see previous section. The interesting point here is *bundling*. Indeed, bundling amounts to leaving out some histories in a tree, and this allows us to leaving out those sequences of action profiles that are not relevant for the description of a given IRG (see our example in the previous section, where endless repetition of cooperation by agent 1 is excluded). In this sense, choice b-trees may provide the exact *temporal structure* of a given IRG, while BT+AC seems to provide a fixed temporal structure imposing some constraint on the development of an IRG.

As for Kamp agent frames, *times* can be thought of as instances of the constituent game (‘game-rounds’) and *worlds* can be thought as the *admitted* maximal linear sequences of action profiles from the given game. Two such sequences w and w' would then be *matching* at game-round t if the two sequences evolve in exactly the same way up to (and including) round t , and possibly get distinct later.

In TKSFs, time-points correspond to moment-history pairs in choice b-trees, and in turn, in TKSFs, *moments* and *histories* can be constructed out of *time-points* and *time-lines*: a moment m is a class $\mathcal{R}_\square(x)$ of mutually alternative time-points (see Section 2), a TKSF-history h_x originating from x is a class of moments $\mathcal{R}_\square(y)$ such that y is in the time-line of x (see proof of Theorem 2 in the technical appendix at the end of the paper). The relation between a *time-point* x in a TKSF and the corresponding *moment* $\mathcal{R}_\square(x)$ seems to parallel the relation between the occurrence of a specific action profile of the constituent game in a given IRG and a complete description of the constituent game in which all possible action profiles are represented. As emphasized in the previous section, a natural interpretation of time-lines in a TKSF is to conceive them as infinite sequences of action profiles of the constituent game in a given IRG.

7 Conclusion

We have provided an equivalence result for three existing semantics of temporal STIT (T-STIT): the Kripke-style semantics for T-STIT recently introduced in [20], a semantics based on the concept of Kamp frame and a bundled tree semantics for T-STIT. In the second part of the paper, we have introduced a semantics for T-STIT based on Fagin et al.’s concept of interpreted system and proven its equivalence with a restricted class of temporal Kripke STIT models with discrete time and initial point.

We plan to devote future research to the investigation of the interpreted system semantics for T-STIT introduced in Section 3. In particular, we plan to find a complete axiomatization for T-STIT relative to this semantics and to compare it with the existing axiomatization for T-STIT relative to the class of TKSMs given in [20]. Finally, we plan to investigate additive and superadditive full group STIT logics based on fused temporal-stit operators in the style of [4] and [9]. Their decidability is still an open issue, (see Section 5) and a positive result would make these logics a convenient tool to reason about coalitional agency.

References

- [1] BALBIANI, P., HERZIG, A., and TROQUARD, N., ‘Alternative axiomatics and complexity of deliberative STIT theories’, *Journal of Philosophical Logic*, 37 (2008), 4, 387–406.
- [2] BELNAP, N., PERLOFF, M., and XU, M., *Facing the future: agents and choices in our indeterminist world*, Oxford Univ. Press, 2001.
- [3] BENTHEM, J. VAN, ‘Rational dynamics and epistemic logic in games’, *International Game Theory Review*, 9 (2007), 1, 13–45.
- [4] BROERSEN, J., ‘Deontic epistemic stit logic distinguishing modes of mens rea’, *Journal of Applied Logic*, 9 (2011), 2, 127–152.
- [5] BROERSEN, J., HERZIG, A., and TROQUARD, N., ‘What groups do, can do, and know they can do: an analysis in normal modal logics’, *Journal of Applied Non-Classical Logics*, 19 (2009), 3, 261–290.
- [6] BROERSEN, J., HERZIG, A., and TROQUARD, N., ‘Embedding Alternating-time Temporal Logic in Strategic STIT logic of agency’, *Journal of Logic and Computation*, 16 (2006), 5, 559–578.
- [7] BURGESS, J., ‘Logic and time’, *Journal of Symbolic Logic*, 44 (1979), 556–582.
- [8] CHELLAS, B. J., ‘Time and modality in the logic of agency’, *Studia Logica*, 51 (1992), 485–517.
- [9] CIUNI, R., and ZANARDO, A., ‘Completeness of a branching-time logic with possible choices’, *Studia Logica*, 96 (2010), 3, 393–420.
- [10] HENKIN, L., ‘Completeness in the Theory of Types’, *The Journal of Symbolic Logic*, 15 (1950), 2, 81–91.
- [11] HOEK, VAN DER W., and WOOLRIDGE, M., ‘Towards a logic of rational agency’, *Logic Journal of the IGPL*, 11 (2003), 2, 135–159.
- [12] FAGIN, R., HALPERN, J., MOSES, Y., and VARDI, M., *Reasoning about Knowledge*, MIT Press, Cambridge, 1995.
- [13] GABBAY, D. M., HODKINSON, I. M., and REYNOLDS, M., *Temporal Logic: Mathematical Foundations and Computational Aspects*, volume 1. Clarendon Press, Oxford, 1994.
- [14] HALPERN, J. Y., and PUCELLA, R., ‘Modeling adversaries in a logic for security protocol analysis’, *Logical Methods in Computer Science*, 8 (2012), 1, 1–26.
- [15] HALPERN, J. Y., and PASS, R., ‘A Knowledge-Based Analysis of the Blockchain Protocol’, in *Proceedings of the Sixteenth Conference on Theoretical Aspects of Rationality and Knowledge (TARK 2017)*, EPTCS 251 (2017), 324–335.

- [16] HERZIG, A., and LORINI, E., ‘A dynamic logic of agency I: STIT, Capabilities and Powers’, *Journal of Logic, Language and Information*, 19 (2010), 1, 89–121.
- [17] HERZIG, A., and SCHWARZENTRUBER, F., ‘Properties of logics of individual and group agency’, *Advances in modal logic*, 7 (2008), 133–149.
- [18] HORTY, J. F., *Agency and Deontic Logic*, Oxford University Press, Oxford, 2001.
- [19] HORTY, J. F., and BELNAP, N., ‘The deliberative STIT: A study of action, omission, and obligation’, *Journal of Philosophical Logic*, 24(6) (1995), 583–644.
- [20] LORINI, E., ‘Temporal stit logic and its application to normative reasoning’, *Journal of Applied Non-Classical Logics*, 23 (2013), 4, 372–399.
- [21] PAYETTE, G., ‘Decidability of an Xstit Logic’, *Studia Logica*, 102 (2014), 3, 577–677.
- [22] REITER, R., ‘On integrity constraints’, *Proceedings of the 2nd TARK conference*, Morgan Kaufmann Publishers Inc, p. 97–111.
- [23] REYNOLDS, M., ‘An axiomatization of full computation tree logic’, *Journal of Symbolic Logic*, 66 (2001), 3, 1011–1057.
- [24] REYNOLDS, M., ‘Axioms for branching time’, *Journal of Logic and Computation*, 12 (2002), 4, 679–697.
- [25] REYNOLDS, M., ‘An axiomatization of Prior’s ockhamist logic of historical necessity’, in *Advances in Modal Logic*, vol. 4, King’s College Publications, 2003, pp. 355–370.
- [26] SCHWARZENTRUBER, F., ‘Complexity Results of STIT Fragments’, *Studia Logica*, 100 (2012), 5, 1001–1045.
- [27] THOMASON, R., ‘Combinations of tense and modality’, in Gabbay, D. and Guenther, F. (eds.), *Handbook of Philosophical Logic*, vol. 2, Reidel, Dordrecht, 1984. 2nd edition.
- [28] VON KUTSCHERA, F., ‘T x W completeness’, *Journal of Philosophical Logic*, 26 (1997), 241–250.
- [29] VÄÄNÄNEN, J., ‘Second-order logic and the foundations of mathematics’, *Bulletin of Symbolic Logic*, 7 (2001), 4, 504–520.
- [30] WÖLF, S., ‘Propositional Q-logic’, *Journal of Philosophical Logic*, 31 (2002), 387–414.
- [31] ZANARDO, A., ‘Branching-time logic with quantification over branches: The point of view of modal logic’, *Journal of Symbolic Logic*, 61 (1996), 1, 143–166.
- [32] ZANARDO, A., ‘Quantification over Sets of Possible Worlds in Branching-Time Semantics’, *Studia Logica*, 82 (2006), 3, 379–400.

A Technical Appendix

A.1 Proof of Theorem 1

(\Rightarrow) We first prove the left-to-right direction of the theorem, that is, we show that if φ is satisfiable in a Kamp agent model, then it is satisfiable in a TKSM.

Let $\mathcal{M}^{\mathcal{F}} = (W, \mathcal{O}, \{\sim_t\}_{t \in T}, \{\sim_{\langle t, i \rangle}\}_{t \in T, i \in \text{Agt}}, \{\sim_{\langle t, \text{Agt} \rangle}\}_{t \in T}, \pi)$ be a Kamp agent model and let $\langle t_0, w_0 \rangle$ be a time-world pair such that $\mathcal{M}^{\mathcal{F}}, \langle t_0, w_0 \rangle \models \varphi$. From the Kamp agent model $\mathcal{M}^{\mathcal{F}}$ we define the corresponding structure $\mathcal{M}^{\mathcal{K}} = (X, \mathcal{R}_{\square}, \{\mathcal{R}_i\}_{i \in \text{Agt}}, \mathcal{R}_{\text{Agt}}, \mathcal{R}_{\mathcal{G}}, \mathcal{R}_{\text{H}}, \mathcal{V})$ with:

- $X = \{\langle t, w \rangle \mid w \in W \text{ and } t \in T_w\}$,
- $\mathcal{R}_{\square} = \{\langle \langle t, w \rangle, \langle t', w' \rangle \rangle \mid t = t' \text{ and } w \sim_t w'\}$,
- $\mathcal{R}_i = \{\langle \langle t, w \rangle, \langle t', w' \rangle \rangle \mid t = t' \text{ and } w \sim_{t, i} w'\}$,
- $\mathcal{R}_{\text{Agt}} = \{\langle \langle t, w \rangle, \langle t', w' \rangle \rangle \mid t = t' \text{ and } w \sim_{\langle t, \text{Agt} \rangle} w'\}$,
- $\mathcal{R}_{\mathcal{G}} = \{\langle \langle t, w \rangle, \langle t', w' \rangle \rangle \mid w = w' \text{ and } t <_w t'\}$ and \mathcal{R}_{H} is the inverse relation of $\mathcal{R}_{\mathcal{G}}$,
- for all $p \in \text{Atm}$ and $\langle t, w \rangle \in X$, $\langle t, w \rangle \in \mathcal{V}(p)$ iff $t \in \pi(p)$.

It is a routine task to check that $\mathcal{M}^{\mathcal{K}}$ is a TKSM. Moreover, by induction on the structure of φ , it is easy to check that for every T-STIT formula φ and for every time-world pair $\langle t, w \rangle$ in X , $\mathcal{M}^{\mathcal{F}}, \langle t, w \rangle \models \varphi$ iff $\mathcal{M}^{\mathcal{K}}, \langle t, w \rangle \models \varphi$. Hence, from the initial assumption $\mathcal{M}^{\mathcal{F}}, \langle t_0, w_0 \rangle \models \varphi$, we have $\mathcal{M}^{\mathcal{K}}, \langle t_0, w_0 \rangle \models \varphi$.

(\Leftarrow) Now, we prove the right-to-left direction of the theorem, that is, we show that if φ is satisfiable in a TKSM, then it is satisfiable in a Kamp agent model.

Let $\mathcal{M}^{\mathcal{K}} = (X, \mathcal{R}_{\square}, \{\mathcal{R}_i\}_{i \in \text{Agt}}, \mathcal{R}_{\text{Agt}}, \mathcal{R}_{\mathcal{G}}, \mathcal{R}_{\text{H}}, \mathcal{V})$ be a TKSM and x_0 a point in X such that $\mathcal{M}^{\mathcal{K}}, x_0 \models \varphi$. Let us remind the reader of the definition of the set $\mathcal{T}(x) = \mathcal{R}_{\text{H}}(x) \cup \{x\} \cup \mathcal{R}_{\mathcal{G}}(x)$ of the time-points that are temporally related with time-point x , and of the time-line ℓ_x as the linearly ordered set $(\mathcal{T}(x), \mathcal{R}_{\mathcal{G}})$. Also, recall the definition of TL as the set $\{\ell_x \mid x \in X\}$. When the subscript x is not needed, we will feel free to drop it. Also, for every ℓ , let us define $T_{\ell} = \{\mathcal{R}_{\square}(y) \mid y \in \ell\}$. From the TKSM $\mathcal{M}^{\mathcal{K}}$ we define the corresponding structure $\mathcal{F} = (W, \mathcal{O}, \{\sim_t\}_{t \in T}, \{\sim_{\langle t, i \rangle}\}_{t \in T, i \in \text{Agt}}, \{\sim_{\langle t, \text{Agt} \rangle}\}_{t \in T})$ where:

- $W = TL$,
- for all $w \in W$, $T_w = \{\mathcal{R}_{\square}(x) \mid x \in w\}$,
- for all $w \in W$ and for all $t, t' \in T$, $t <_w t'$ iff $\exists x \in t, \exists y \in t'$ such that $x \mathcal{R}_{\mathcal{G}} y$,
- for all $t \in T$ and for all $w, w' \in T$, $w \sim_t w'$ iff $t \in T_w \cap T_{w'}$,
- for all $t \in T$ and for all $w, w' \in T$, $w \sim_{\langle t, i \rangle} w'$ iff:
 - $t \in T_w \cap T_{w'}$, and
 - $\exists x, y \in t$ such that $x \mathcal{R}_i y, x \in w$ and $y \in w'$,

- for all $t \in T$ and for all $w, w' \in T$, $w \sim_{\langle t, \text{Agt} \rangle} w'$ iff:
 - $t \in T_w \cap T_{w'}$, and
 - $\exists x, y \in t$ such that $x \mathcal{R}_{\text{Agt}} y$, $x \in w$ and $y \in w'$.

It is a routine task to check that \mathcal{F} is a Kamp agent frame. Because of the Constraint C7 in the definition of TKSF, we have that, for all $t \in T$ and for all $w \in W$, if $t \cap w \neq \emptyset$ then $t \cap w$ is a singleton. We write t_w to denote the unique element in $t \cap w$, when $t \cap w \neq \emptyset$. Moreover, for all $x \in X$, $\{\langle t, w \rangle \mid t \cap w = \{x\}\}$ is also a singleton. We write $\langle t, w \rangle_x$ to denote the unique element in $\{\langle t, w \rangle \mid t \cap w = \{x\}\}$.

For all $p \in \text{Atm}$, let:

- $\pi(p) = \{t \in T \mid \exists x \in t \text{ such that } \mathcal{M}^{\mathcal{K}}, x \models p\}$.

Clearly, $\mathcal{M}^{\mathcal{F}} = (\mathcal{F}, \pi)$ is a Kamp agent model. Moreover, by induction on the structure of φ , it is easy to check that for every T-STIT formula φ and for every point $x \in X$, $\mathcal{M}^{\mathcal{K}}, x \models \varphi$ iff $\mathcal{M}^{\mathcal{F}}, \langle t, w \rangle_x \models \varphi$. Hence, from the initial assumption $\mathcal{M}^{\mathcal{K}}, x_0 \models \varphi$, we have $\mathcal{M}^{\mathcal{F}}, \langle t, w \rangle_{x_0} \models \varphi$.

A.2 Proof of Theorem 2

We proceed as in the proof of Theorem 1.

(\Rightarrow) We first prove the left-to-right direction of the theorem, that is, we show that if φ is satisfiable in a choice b-tree model, then it is satisfiable in a TKSM.

Let $\mathcal{M}^{\mathcal{B}} = (\mathfrak{T}, B, \{\sim_{\langle m, i \rangle}\}_{m \in M, i \in \text{Agt}}, \{\sim_{\langle m, \text{Agt} \rangle}\}_{m \in M}, v)$ be a choice b-tree model and let $\langle m_0, h_0 \rangle$ be a moment-history pair such that $\mathcal{M}^{\mathcal{B}}, \langle m_0, h_0 \rangle \models \varphi$. From the choice b-tree model $\mathcal{M}^{\mathcal{B}}$ we define the corresponding structure $\mathcal{M}^{\mathcal{K}} = (X, \mathcal{R}_{\square}, \{\mathcal{R}_i\}_{i \in \text{Agt}}, \mathcal{R}_{\text{Agt}}, \mathcal{R}_{\mathcal{G}}, \mathcal{R}_{\mathcal{H}}, \mathcal{V})$ with:

- $X = \{\langle m, h \rangle \mid m \in M \text{ and } h \in B_m\}$,
- $\mathcal{R}_{\square} = \{(\langle m, h \rangle, \langle m', h' \rangle) \mid m = m' \text{ and } h, h' \in B_m\}$,
- $\mathcal{R}_i = \{(\langle m, h \rangle, \langle m', h' \rangle) \mid m = m' \text{ and } h \sim_{\langle m, i \rangle} h'\}$,
- $\mathcal{R}_{\text{Agt}} = \{(\langle m, h \rangle, \langle m', h' \rangle) \mid m = m' \text{ and } h \sim_{\langle m, \text{Agt} \rangle} h'\}$,
- $\mathcal{R}_{\mathcal{G}} = \{(\langle m, h \rangle, \langle m', h' \rangle) \mid h = h' \text{ and } m \prec m'\}$ and $\mathcal{R}_{\mathcal{H}}$ is the inverse relation of $\mathcal{R}_{\mathcal{G}}$,
- for all $p \in \text{Atm}$ and $\langle m, h \rangle \in X$, $\langle m, h \rangle \in \mathcal{V}(p)$ iff $m \in v(p)$.

It is a routine task to check that $\mathcal{M}^{\mathcal{K}}$ is a TKSM. Moreover, by induction on the structure of φ , it is easy to check that for every T-STIT formula φ and for every moment-history pair $\langle m, h \rangle$, $\mathcal{M}^{\mathcal{B}}, \langle m, h \rangle \models \varphi$ iff $\mathcal{M}^{\mathcal{K}}, \langle m, h \rangle \models \varphi$. Hence, from the initial assumption $\mathcal{M}^{\mathcal{B}}, \langle m_0, h_0 \rangle \models \varphi$, we have $\mathcal{M}^{\mathcal{K}}, \langle m_0, h_0 \rangle \models \varphi$.

(\Leftarrow) Now, we prove the right-to-left direction of the theorem, that is, we show that if φ is satisfiable in a TKSM, then it is satisfiable in a choice b-tree model.

Let $\mathcal{M}^{\mathcal{K}} = (\mathcal{K}, \mathcal{V})$ be a TKSM based on the TKSF $\mathcal{K} = (X, \mathcal{R}_{\square}, \{\mathcal{R}_i\}_{i \in \text{Agt}}, \mathcal{R}_{\text{Agt}}, \mathcal{R}_{\mathcal{G}}, \mathcal{R}_{\mathcal{H}})$ and x_0 a point in X such that $\mathcal{M}^{\mathcal{K}}, x_0 \models \varphi$.

For every $x \in X$, we define the x -relative TKSF-history h_x in \mathcal{K} : $h_x = \{\mathcal{R}_{\square}(y) \mid y \in \mathcal{T}(x)\}$. Moreover, we define the set $B_{\mathcal{K}} = \{h_x \mid x \in X\}$ of TKSF-histories in \mathcal{K} . As we shall see, $B_{\mathcal{K}}$ is a bundle in $\mathcal{B}^{\mathcal{K}}$. Notice that the definition of h_x differs from the definition of ℓ_x since the former has set-theoretical constructions over time-points as its elements, while the latter has time-points themselves as elements. As usual, we will omit the subscript x when it is not needed.

From the TKSM $\mathcal{M}^{\mathcal{K}}$ we define the corresponding structure $\mathcal{B} = (\mathfrak{T}, B, \{\sim_{\langle m, i \rangle}\}_{m \in M, i \in \text{Agt}}, \{\sim_{\langle m, \text{Agt} \rangle}\}_{m \in M})$ where:

- $M = \{\mathcal{R}_{\square}(x) \mid x \in X\}$,
- for all $m, m' \in M$, $m \prec m'$ iff $\exists x \in m$ and $\exists y \in m'$ such that $x \mathcal{R}_{\mathcal{G}} y$,
- $B = B_{\mathcal{K}}$,
- for all $m \in M$, $h \sim_{\langle m, i \rangle} h'$ iff:
 - $h, h' \in B_m$, and
 - $\exists x, y \in m$ such that $x \mathcal{R}_i y$, $h_x = h$ and $h_y = h'$,
- for all $m \in M$, $h \sim_{\langle m, \text{Agt} \rangle} h'$ iff:
 - $h, h' \in B_m$, and
 - $\exists x, y \in m$ such that $x \mathcal{R}_{\text{Agt}} y$, $h_x = h$ and $h_y = h'$.

It is a routine task to check that \mathcal{B} is a choice b-tree.

For all, $p \in \text{Atm}$ let:

- $v(p) = \{m \in M \mid \exists x \in m \text{ such that } \mathcal{M}^{\mathcal{K}}, x \models p\}$.

Clearly, $\mathcal{M}^{\mathcal{B}} = (\mathcal{B}, v)$ is a choice-b model. Moreover, by induction on the structure of φ , it is easy to check that for every T-STIT formula φ and for every point $x \in X$, $\mathcal{M}^{\mathcal{K}}, x \models \varphi$ iff $\mathcal{M}^{\mathcal{B}}, \langle \mathcal{R}_{\square}(x), h_x \rangle \models \varphi$. Hence, from the initial assumption $\mathcal{M}^{\mathcal{K}}, x_0 \models \varphi$, we have $\mathcal{M}^{\mathcal{B}}, \langle \mathcal{R}_{\square}(x_0), h_{x_0} \rangle \models \varphi$.

A.3 Proof of Theorem 3

(\Rightarrow) Let $\mathcal{M}^{\mathcal{D}\mathcal{K}} = (X, \mathcal{R}_{\square}, \{\mathcal{R}_i\}_{i \in \text{Agt}}, \mathcal{R}_{\text{Agt}}, \mathcal{R}_{\mathcal{X}}, \mathcal{V})$ be a DTKSM and let $x_0 \in X$ such that $\mathcal{M}^{\mathcal{D}\mathcal{K}}, x_0 \models \varphi$.

Let L_e be an arbitrary set of local states of the environment and g_e an injection from X/\mathcal{R}_{\square} (i.e., the quotient set of X by \mathcal{R}_{\square}) to L_e . The fact that g_e is an injection implies that the size of L_e is equal to or higher than the size of X/\mathcal{R}_{\square} .

Let L_i be an arbitrary set of local states of agent i . Moreover, let $Y \in X/\mathcal{R}_{\square}$. Then, g_i^Y is an injection from Y/\mathcal{R}_i to L_i . The fact that g_i^Y is an injection implies that the size of L_i is equal to or higher than the size of Y/\mathcal{R}_i .

Let $\mathcal{G} = L_e \times L_1 \times \dots \times L_n$ be the set of global systems.

Because of the Constraints C7 and C9 in Definition 15, the fact that $\mathcal{R}_{\mathcal{X}}$ is serial and deterministic, the fact that $\mathcal{R}_{\mathcal{V}}$ is deterministic and that $\mathcal{R}_{\mathcal{G}}$ is defined to be the

transitive closure of \mathcal{R}_X , for every time-line $\ell_x \in TL$ we can define a corresponding *bijective* function $\gamma_x : \mathbb{N} \rightarrow \mathcal{T}(x)$ such that for all $y \in X$ and for all $k \in \mathbb{N}$:

$$\gamma_x(k) = y \text{ iff } |\mathcal{R}_H(y)| = k.$$

Thus, for every time-line $\ell_x \in TL$, we can define a corresponding run $r_x : \mathbb{N} \rightarrow \mathcal{G}$ such that:

- for all $k \in \mathbb{N}$ and for all $x \in X$,
 $r_x(k) = (g_e(\mathcal{R}_\square(\gamma_x(k))), g_1^{\mathcal{R}_\square(\gamma_x(k))}(\mathcal{R}_i(\gamma_x(k))), \dots, g_n^{\mathcal{R}_\square(\gamma_x(k))}(\mathcal{R}_n(\gamma_x(k))))$.

Let $\mathcal{S} = \{r_x : x \in X\}$ be the corresponding system.

Let ω be the valuation function defined as follows:

- for all $p \in \text{Atm}$ and for all $(r_x, k) \in \mathcal{P}$, $(r_x, k) \in \omega(p)$ iff $\gamma_x(k) \in \mathcal{V}(p)$.

It is straightforward to prove that the interpreted system $\mathcal{I} = (\mathcal{S}, \omega)$ is a STIT interpreted system. Moreover, by induction on the structure of φ , it is a routine task to verify that $(\mathcal{I}, r_{x_0}, \gamma_{x_0}^{-1}(x_0)) \models \varphi$.

(\Leftarrow) Let $\mathcal{I} = (\mathcal{S}, \omega)$ be a STIT interpreted system and (r_0, k_0) a point in \mathcal{I} such that $(\mathcal{I}, r_0, k_0) \models \varphi$. Let us define the tuple $\mathcal{M}^{\mathcal{DK}} = (X, \mathcal{R}_\square, \{\mathcal{R}_i\}_{i \in \text{Agt}}, \mathcal{R}_{\text{Agt}}, \mathcal{R}_X, \mathcal{V})$ as follows:

- $X = \mathcal{P}$,
- for all $(r, k), (r', k') \in X$, $(r, k) \mathcal{R}_\square (r', k')$ iff $(r, k) \equiv_e (r', k')$,
- for all $(r, k), (r', k') \in X$, $(r, k) \mathcal{R}_i (r', k')$ iff $(r, k) \equiv_i (r', k')$,
- for all $(r, k), (r', k') \in X$, $(r, k) \mathcal{R}_{\text{Agt}} (r', k')$ iff $(r, k) \equiv_{\text{Agt}} (r', k')$,
- for all $(r, k), (r', k') \in X$, $(r, k) \mathcal{R}_X (r', k')$ iff $r = r'$ and $k' = k + 1$,
- for all $p \in \text{Atm}$, $\mathcal{V}(p) = \omega(p)$.

It is straightforward to prove that $\mathcal{M}^{\mathcal{DK}}$ is a DTKSM. Moreover, by induction on the structure of φ , it is a routine task to verify that $\mathcal{M}^{\mathcal{DK}}, (r_0, k_0) \models \varphi$.