D. Klein
E. Pacuit

Changing Types: Information Dynamics for Qualitative Type Spaces

Abstract. Many different approaches to describing the players' knowledge and beliefs can be found in the literature on the epistemic foundations of game theory. We focus here on non-probabilistic approaches. The two most prominent are the so-called Kripke- or Aumann- structures and knowledge structures (non-probabilistic variants of Harsanyi type spaces). Much of the recent work on Kripke structures has focused on dynamic extensions and simple ways of incorporating these. We argue that many of these ideas can be applied to knowledge structures as well. Our main result characterizes precisely when one *type* can be transformed into another *type* by a specific type of information update. Our work in this paper suggest that it would be interesting to pursue a theory of "information dynamics" for knowledge structures (and eventually Harsanyi type spaces).

Keywords: Dynamic Epistemic Logic, Epistemic Models of Games, Harsanyi Type Spaces

1. Introduction

The[1] central thesis of the epistemic program in game theory is that the basic mathematical model of a game situation should include an explicit parameter describing the players' *informational attitudes*.[2] See [7] for the relevant references and a discussion of the key results, and [23] for an introduction to this literature. Games are played in specific *informational contexts*, in which players have specific knowledge and beliefs about each other.[3] Many differ-

Presented by **Name of Editor**; *Received* December 1, 2005

[1]This research was funded by the NWO Vidi project 016.094.345. The authors would like to thank the anonymous referees for their extremely thorough comments which greatly improved this paper.

[2]This is, of course, something of a truism regarding games of *incomplete* or *imperfect* information. But, the thesis is intended to apply to *all* game situations. See [8, Section 5] for a precise description about the crucial differences between an epistemic model of a game and a *Bayesian game*.

[3]This is nicely explained by Adam Brandenburger and Amanda Friedenberg ([9, pg. 801]): "In any particular structure, certain beliefs, beliefs about beliefs, ..., will be present and others won't be. So, there is an important implicit assumption behind the choice of a structure. This is that it is "transparent" to the players that the beliefs in the [type] structure — and only those beliefs — are possible....The idea is that there is a "context" to the strategic situation (eg., history, conventions, etc.) and this "context" causes the

ent formal models have been used to represent such informational contexts of a game (see [6, 20, 5], and references therein, for a discussion). In this paper, we are not only interested in structures that describe the informational context of a game, but how these structures can *change* in response to the players' observations, communicatory acts or other dynamic operations of information change (cf. [3]).

We focus our attention on the players' *hard information* about the game (which we refer to as *knowledge* following standard terminology in the game theory and epistemic logic literature) and its dynamics. Broadly speaking, there are two different types of models that have been used to describe the players' knowledge (and beliefs) in a game situation. Both types of models include a nonempty set S of *states of nature* (elements of S are intended to represent possible outcomes of a game situation).[4] The first type of models are the so-called *Aumann-* or *Kripke-structures* [2, 18]. These structures describe the players' knowledge in terms of an *epistemic indistinguishability* relation over a (finite) set of states W. The second type of models are the knowledge structures of [16, 15], which are non-probabilistic variants of *Harsanyi type spaces* [19]. [5] The key concept here is a **type** which describes the players' infinite hierarchy of knowledge (i.e., what the players know about the ground facts, what the players know about each others knowledge of the ground facts, what players know about what each other know about each others knowledge of the ground facts, and so on.) The precise relationship between these two types of models was clarified in [16, 15].

Our goal in this paper is to show how to adapt recent work modeling information change on Kripke structures as a product update with an *event model* [12] to the more general setting where the players' knowledge is represented using knowledge structures. To the best of our knowledge, this is the first attempt to develop a theory of information change for knowledge structures in the style of recent work on dynamic epistemic logic. Our main result (Theorem 3.17) characterizes precisely when a type in a fixed knowledge structure can be transformed into another type in that structure using the product update operation.

There are two main motivations for this technical study. The first is to explore generalizations of the product update operation. This is done

players to rule out certain beliefs."

[4]Often, it is assumed that the elements of S can be described by some logical language (for example, propositional logic), but this is not crucial for us in this paper.

[5]See [25] for a modern introduction to type spaces as models of beliefs and [22] for a discussion of Harsanyi's classic paper.

in Section 3.1 where we also generalize a result of [10] characterizing when a Kripke structure can be transformed into another Kripke structure by a product update. The second motivation for this work is to initiate a study of information dynamics for epistemic models of games. The players' information in a game can change in two ways. First, the players' knowledge and beliefs change during the play of a sequential game (for example, they learn about the choices of the other players as the game is played). The second way that the players' information can change is in response to some exogenous event. For example, during a poker game, a player may accidentally drop his cards or a gust of wind may allow a subset of the players to see certain cards. Of course, one may argue that game-theoretic models should abstract away from these payoff irrelevant events. We agree that the type of events we have in mind here are irrelevant to a game-theoretic analysis. But, these events do change the *context*[6] of a game by revealing or hiding important information to all or some of the players. This paper is a first step towards a more general project that uses the dynamic epistemic logic framework to represent changes in the informational context of a game.

Our paper is organized as follows. Section 2 provides the necessary background on (dynamic) epistemic logic and knowledge structures. Note that this Section was written for a reader already familiar with the key concepts and definitions. Consult [3] and [15] for motivations and a broader discussion of the literature. Our main result is in Section 3.2 with the technical preliminaries found in Section 3.1. We conclude in Section 4 with a discussion of topics for future research.

2. Background

2.1. A Primer on Dynamic Epistemic Logic

We assume the reader is familiar with the basics of (dynamic) epistemic logic, and so, we only give the key definitions here (see the textbooks [18, 3] for an introduction to the subsequent definitions). Let I be the finite set of players and At a (finite or infinite) set of atomic propositions.[7]

Definition 2.1 (Epistemic Language). The **epistemic language**, denoted $\mathcal{L}_{EL}$, is the smallest set of formulas generated by the following grammar:

$$\varphi ::= p \mid \neg\varphi \mid \varphi \wedge \varphi \mid K_i\varphi$$

[6] Here, we take the "context" of a game to be *all* events that influence the players' beliefs in the game situation.

[7] Atomic propositions are intended to represent properties of states of nature.

where $p \in \mathsf{At}$ and $i \in I$. Define $L_i\varphi$ as the dual of K_i (i.e., $L_i\varphi := \neg K_i \neg\varphi$) and the other boolean connectives (e.g., $\vee, \rightarrow$) as usual. ◁

The intended interpretation of $K_i\varphi$ is "agent i knows that φ (is true)". The standard semantics for $\mathcal{L}_{EL}$ are Kripke structures.

Definition 2.2 (Kripke Structure). A **Kripke structure** (for a set of atomic propositions At) is a tuple $\langle W, \{R_i\}_{i\in I}, V\rangle$ where W is a set of states, $R_i \subseteq W \times W$ is an equivalence relation[8], and $V : \mathsf{At} \rightarrow \wp(W)$ is a valuation function. To simplify notation, we may write $w \in \mathcal{M}$ when $w \in W$. ◁

Formulas of $\mathcal{L}_{EL}$ are interpreted at states in a Kripke model in the standard way, we only remind the reader of the definition for the knowledge modality:

$$\mathcal{M}, w \models K_i\varphi \text{ iff } \text{ for all } v \in W \text{ if } wR_iv \text{ then } \mathcal{M}, v \models \varphi$$

The central idea of *dynamic* epistemic logic is to describe events that change a situation and the (uncertain) perceptions of these events by the agents' as a so-called event model.

Definition 2.3 (Event Model). An **event model** is a tuple $\langle E, \{Q_i\}_{i\in I}, \mathsf{pre}\rangle$ where E is a set of basic events, $Q_i \subseteq E \times E$ is an equivalence relation[9] and $\mathsf{pre} : E \rightarrow \mathcal{L}_{EL}$ assigns to each primitive event a formula that serves as a **precondition** for that event. We write $e \in \mathcal{E}$ if e is an event in $\mathcal{E}$. ◁

The primitive events represent the basic observations available to the agents in a dynamic situation. Similar to Kripke structures, uncertainty about which events are taking place is represented by relations Q_i. Given our assumptions that each Q_i is an equivalence relation, the intended interpretation of eQ_if is that agent i cannot distinguish between events e and f. The key operation of *product update* describes how to incorporate into a Kripke structure $\mathcal{M}$ (describing an epistemic situation) the epistemic event described by an event model $\mathcal{E}$.

Definition 2.4 (Product Update). The **product update** of a Kripke model $\mathcal{M} = \langle W, \{R_i\}_{i\in I}, V\rangle$ and an event model $\mathcal{E} = \langle E, \{Q_i\}_{i\in I}, \mathsf{pre}\rangle$ is a Kripke model $\mathcal{M} \oplus \mathcal{E} = \langle W', \{R'_i\}_{i\in I}, V'\rangle$ defined as follows:

- $W' = \{(w, e) \in W \times E \mid \mathcal{M}, w \models \mathsf{pre}(e)\}$

[8] In this paper, we restrict attention structures where the epistemic relations are equivalence relations. These are known in the literature as **S5**-structures or Aumann structures.

[9] To keep things manageable for this initial study, we restrict attention to event models with equivalence relations. For much of what follows, this assumption is not crucial.

- $(w,e)R_i'(w',e')$ iff wR_iw' and eQ_ie'
- $(w,e) \in V'(p)$ iff $w \in V(p)$ $\lhd$

This operation (together with variants appropriate for modeling belief and preference change) has been extensively studied in the literature. We do not provide an overview of this literature here: See [12, 3] for an extensive analysis. Rather, the focus is on how to understand this theory of information dynamics in the context of models of knowledge (and beliefs) typically found in the game theory literature. We need one additional notion from the general theory of modal logic.

Definition 2.5 (Bisimulation). Suppose that $\mathcal{M}_1 = \langle W_1, R_1, V_1\rangle$ and $\mathcal{M}_2 = \langle W_2, R_2, V_2\rangle$ are Kripke structures. A nonempty relation $Z \subseteq W_1 \times W_2$ is a **bisimulation** provided for all $w_1 \in W_1$ and $w_2 \in W_2$, if w_1Zw_2 then:

(atomic harmony) For all $p \in \mathsf{At}$, $w_1 \in V_1(p)$ iff $w_2 \in V_2(p)$.

(zig) If $w_1R_1v_1$ then there is a $v_2 \in W_2$ such that $w_2R_2v_2$ and v_1Zv_2.

(zag) If $w_2R_2v_2$ then there is a $v_1 \in W_1$ such that $w_1R_1v_1$ and v_1Zv_2.

We write $\mathcal{M}_1, w_1 \leftrightarroweq \mathcal{M}_2, w_2$ if there is a bisimulation relating w_1 with w_2. We write $\mathcal{M}_1 \leftrightarroweq \mathcal{M}_2$ if there is a bitotal bisimulation between $\mathcal{M}_1$ and $\mathcal{M}_2$, that is a bisimulation Z such that for every $v \in \mathcal{M}_1$ there is some $W \in \mathcal{M}_2$ with vZw and vice versa. The relation Z is called a **simulation from $\mathcal{M}_1$ to $\mathcal{M}_2$**, denoted $\mathcal{M}_1, w_1 \rightrightarrows \mathcal{M}_2, w_2$, if Z satisfies the **atomic harmony** and **zig** properties. Z is called **total** provided for each $w_1 \in W_1$ there is a $w_2 \in W_2$ such that w_1Zw_2. Finally, Z is called **functional** if it is total and a function from W_1 to W_2 (i.e. for every $w_1 \in W_1$ and $w_2, \tilde{w}_2 \in W_2$ it is the case that w_1Zw_2 and $w_1Z\tilde{w}_2$ implies $w_2 = \tilde{w}_2$). $\lhd$

2.2. Knowledge Structures

Knowledge structures were introduced in [15] as an alternative semantics for the basic epistemic language $\mathcal{L}_{EL}$.[10] They are non-probabilistic versions of Harsanyi type spaces which are the predominant model of knowledge and beliefs in the literature on the epistemic foundations of game theory ([8] offers some explanation about why this is the case).

The key concept is a **κ-world** (also called a **type** in the game theory literature) describing the players' infinite hierarchy of knowledge (belief) of a given state of affairs.

---

[10] See [15] for an extended discussion of knowledge structures aimed at game theorists. Fagin [14] and Fagin and Vardi [17] show how variants of knowledge structures can provide an elegant semantics for many modal logics.

DEFINITION 2.6 (κ-world). Let S be a (finite or infinite) nonempty set (whose elements are called states). A **κ-world** is a vector of functions $\mathbf{f} = \langle f_0, f_1, f_2 \ldots \rangle$ of length κ (a possibly infinite ordinal) defined inductively as follows.

- A **1-world** is a vector $\langle f_0 \rangle$ where f_0 is a state of nature (i.e., $f_0 \in S$).[11]
- For $\kappa > 1$ of the form $\kappa = \lambda + 1$ (i.e. κ is a successor ordinal) a *κ-world* is a vector $\langle f_0 \ldots f_\lambda \rangle$ such that $\langle f_i \mid i < \lambda \rangle$ is a λ-world and f_λ is a function from the set of agents I to the power set of the set of λ-worlds over S (i.e., $f_\lambda : I \to \wp(\mathcal{F}_\lambda(S))$, where $\mathcal{F}_\lambda(S)$ denotes the set of all λ-worlds over S) that satisfies the following conditions. Let $\mathbf{f}_{<\beta}$ denote the initial segment of $\mathbf{f}$ of length β.

  Extendability If $0 < \alpha < \lambda$, then $\mathbf{g} \in f_\alpha(i)$ iff there is some $\mathbf{h} \in f_\lambda(i)$ such that $\mathbf{g} = \mathbf{h}_{<\alpha}$ (i.e., higher-order worlds are extensions of lower-order worlds and every lower-order world has at least one higher-order extension).

  In addition, since we intend κ-worlds to represent the **knowledge** of the players, we impose two additional conditions:

  Correctness For each agent $i \in I$, $\mathbf{f}_{<\lambda} \in f_\lambda(i)$ (i.e., every agent must consider the actual state of the world possible).

  Introspection For all $i \in I$, if $\langle g_0, g_1, \ldots \rangle \in f_\kappa(i)$, then $g_\lambda(i) = f_\lambda(i)$, for all λ with $0 < \lambda < \kappa$ (i.e., players cannot consider states possible that differ in their description of their own lower-order beliefs). $\triangleleft$

- Finally, for κ a limit ordinal a *κ-world* is a vector of functions $\langle f_i \mid i < \kappa \rangle$ such that for every $\lambda < \kappa$ the vector $\langle f_i \mid i < \lambda \rangle$ is a λ-world.

We denote the set of all κ-worlds over S by $\mathcal{F}_\kappa(S)$.

The intended interpretation is that $f_\kappa(i) \subseteq \mathcal{F}_\kappa(S)$ is the set of all κ-worlds player i considers possible. Then, κ-worlds $\mathbf{f}$ are descriptions of the state of affairs and the players higher-order knowledge (up to level κ). Thus, we can interpret the basic epistemic language at κ-worlds. For simplicity, we assume there is an atomic proposition E for every subset of the set of

[11] For the comparison with epistemic logic, it is useful to think of the set of states S as the set of propositional valuations on a set At of atomic propositions. In this case f_0 would be a propositional valuation function.

states S (i.e., $\mathsf{At} = \wp(S)$). This language is interpreted as follows:

$$\begin{aligned}
\mathbf{f} \models E &\Leftrightarrow f_0 \in E \\
\mathbf{f} \models \neg\varphi &\Leftrightarrow \mathbf{f} \not\models \varphi \\
\mathbf{f} \models \varphi \wedge \psi &\Leftrightarrow \mathbf{f} \models \varphi \text{ and } \mathbf{f} \models \psi \\
\mathbf{f} \models K_i\varphi &\Leftrightarrow \text{ for each } \mathbf{g} \in f_l(i) : \mathbf{g} \models \varphi \\
&\qquad \text{where } l \text{ is the quantifier depth}^{12} \text{ of } \varphi.
\end{aligned}$$

There is an alternative way of defining truth of the knowledge modality by defining an accessibility relation on $\mathcal{F}_\kappa(S)$, which transforms $\mathcal{F}_\kappa(S)$ into a Kripke model. We can then use the standard definition of a modal operator. For a κ-world $\mathbf{f} = \langle f_0, f_1, \ldots\rangle$, let $\mathbf{f}^i = \langle f_1(i), f_2(i), \ldots\rangle$ (note that the state of nature is not part of $\mathbf{f}^i$) and define a relation $\sim_i$ on the $\mathcal{F}_\kappa(S)$ as follows: $\mathbf{f} \sim_i \mathbf{g}$ iff $\mathbf{f}^i = \mathbf{g}^i$ (equality is defined component-wise). If $\mathbf{f} \sim_i \mathbf{g}$ then we say $\mathbf{f}$ and $\mathbf{g}$ are equivalent according to agent i. It is easy to see that these relations are equivalence relations. They turn $\mathcal{F}_\kappa(S)$ into a Kripke structure (with $\mathsf{At} = \wp(S)$ and the valuation function V defined by $w \in V(E)$ iff $w \in E$). Fagin *et al.* show ([16, Theorem 2.4]) that the interpretation of the epistemic language given above coincides with the interpretation of the epistemic language obtained by interpreting $\langle \mathcal{F}_\kappa(S), \{\sim_i\}_{i\in I}, V\rangle$ as a Kripke structure. So, there are two equivalent ways to interpret the basic epistemic language on the set $\mathcal{F}_\kappa(S)$ of κ-worlds. In the remainder of the paper, we will use whichever definition is most convenient.

We are interested in general maps between Kripke structures and knowledge structures. To this end, we fix a set of atomic propositions At and assume that the state space S is the set of propositional valuations of At, i.e., $S = \wp(\mathsf{At})$. To simplify our exposition, we identify $p \in \mathsf{At}$ with $\{e \in S \mid p \in e\} \subseteq S$, i.e. the set of valuations containing p.

The key observation is that every Kripke structure can be naturally associated with a substructure of $\langle \mathcal{F}_\omega(S), \{\sim_i\}_{i\in I}, V\rangle$. The mapping is defined as follows.[13]

Definition 2.7 (Embedding from Kripke structures to knowledge structures). Let $\mathcal{M} = \langle W, \{R_i\}_{i\in N}, V\rangle$ be a Kripke structure. We associate with

[12] Quantifier depth is defined as usual by induction on the structure of $\varphi \in \mathcal{L}_{EL}$: Formally, $qd(p) = 0$, $qd(\neg\varphi) = qd(\varphi)$, $qd(\varphi \wedge \psi) = \max(qd(\varphi), qd(\psi))$, and $qd(K_i\varphi) = 1 + qd(\varphi)$.

[13] The mapping is a functional simulation but in general not a bisimulation onto its image. Nonetheless, it is a natural mapping in the sense that when applied to connected components $\mathcal{K}$ of $\langle \mathcal{F}_\omega(S), \{\sim_i\}_{i\in I}, V\rangle$ it is simply the embedding of $\mathcal{K}$ into $\langle \mathcal{F}_\omega(S), \{\sim_i\}_{i\in I}, V\rangle$

each state $w \in W$ in $\mathcal{M}$ an ω-world $\mathbf{f}_{\mathcal{M},w} = \langle f_0^w, f_1^w, f_2^w, \ldots \rangle$ where the f_α^w are defined by synchronous induction on all worlds $w \in W$:

- $f_0^w = \{p \mid w \in V(p)\}$.
- To define the function f_{k+1}^w assume inductively that f_0^x, f_1^x, f_2^x, …, f_k^x have been defined for all worlds $x \in W$ (k a natural number). Then, $f_{k+1}^w(i) = \{\langle f_0^x, f_1^x, \ldots f_k^x \rangle \mid wR_ix\}$.

Define the map $r : W \to \mathcal{F}_\omega(\wp(\mathsf{At}))$ as $r(w) = \mathbf{f}_{\mathcal{M},w}$. $\triangleleft$

For every ordinal λ we can continue the construction to get a vector $\langle f_i^x \mid i < \lambda \rangle$. Thus this map naturally generalizes to maps $r_\lambda : W \to \mathcal{F}_\lambda(\wp(\mathsf{At}))$ for every ordinal λ. To simplify notation, assume for the rest of the paper that $S = \wp(\mathsf{At})$ and that S is finite. The map r_κ gives a precise way to connect the class of all Kripke structures to a single structure $\mathcal{M}^\kappa = \langle \mathcal{F}_\kappa(S), \{\sim_i\}_{i \in I}, V \rangle$ for any κ. The following observation is immediate from the relevant definitions.

Observation 2.8. *Let $\mathcal{M} = \langle W, \{R_i\}_{i \in I}, V \rangle$ be a Kripke structure and $\mathcal{M}^\kappa$ be the structure $\langle \mathcal{F}_\kappa(S), \{\sim_i\}_{i \in I}, V \rangle$.*

1. *The relation $wZ\mathbf{f}$ iff $r_\kappa(w) = \mathbf{f}$ is a functional simulation from $\mathcal{M}$ into $\mathcal{M}^\kappa$, but, in general, is not a bisimulation.*
2. *There is an ordinal λ, depending on $\mathcal{M}$ such that Z is a bisimulation if $\kappa \geq \lambda$.*[14]
3. *In particular, if $\mathcal{M}$ is finite, then there is a bisimulation between $\mathcal{M}$ and $r(\mathcal{M}) = \langle r[W], \{\sim_i\}, V \rangle$. Moreover, $r(\mathcal{M})$ is the minimal bisimulation contraction of $\mathcal{M}$, i.e. the Kripke model of minimal cardinality that allows for a total bisimulation to $\mathcal{M}$.*

Proof. *i*) The functionality of Z is obvious, since r_κ is a function. Atomic harmony holds by definition of f_0^w. To see that zig holds let $v_0, v_1 \in M$ with $v_0 R_i v_1$ and $w \in \mathcal{M}^\kappa$ with $v_0 Z w$. Since Z is functional we have $w = \mathbf{f}_{\mathcal{M},v_0}$ An induction shows that $f_i^{v_0} = f_i^{v_1}$ for every $i \leq \kappa$, thus $\mathbf{f}_{\mathcal{M},v_0}(i) = \mathbf{f}_{\mathcal{M},v_1}(i)$. Thus by definition of $\sim_i$ we have $\mathbf{f}_{\mathcal{M},v_0} \sim_i \mathbf{f}_{\mathcal{M},v_1}$. By definition of Z we also have $v_1 Z \mathbf{f}_{\mathcal{M},v_1}$, thus zig holds. Example 3.10 of [15] shows that that Z is in general not a bisimulation.

[14] In fact, for $\mathcal{M} = \mathcal{F}_\kappa(S)$ we have $\lambda = \kappa$. Moreover the model $\mathcal{M}^\kappa$ is a terminal object in the category of Kripke models over At with total simulations as morphisms. Though $\mathcal{F}_\kappa(S)$ and $\mathcal{F}_\lambda(S)$ are not bisimilar for $\kappa \neq \lambda$.

ii) Choose λ' such that for all $v, w \in \mathcal{M}$ holds: If there is some μ such that $r_\mu(v) \neq r_\mu(w)$, then $r_{\lambda'}(v) \neq r_{\lambda'}(w)$ and let $\lambda := \lambda' + \omega$. We have to show that zag holds: Let vZw with Z defined as above and let $w \sim_i w'$. We have to show that there is some $v' \in \mathcal{M}$ with $r_\lambda(v') = w'$. Indeed, since $w \sim_i w'$ we have for all $\mu < \lambda$ that $w' \restriction \mu \in w_\mu(i)$. By the construction of r_λ this implies that for every $\mu < \lambda$ there is some $v' \in \mathcal{M}$ such that $w' \restriction \mu = r_\mu(v')$. By the choice of λ' and the extendability condtion, we have that $\exists \mu \in [\lambda'; \lambda] : r_\mu(v') \in w_\mu(i)$ implies $\forall \mu \in [\lambda'; \lambda] : r_\mu(v') \in w_\mu(i)$. In particular we have by the limit condition that $r_\lambda(v') = w'$ as desired. See chapter 3 of [15] for more details.
iii) Obvious from *ii*) and the definition of r_ω.

3. Information Dynamics on Knowledge Structures

Our aim is to examine natural transitions between types in a knowledge structure. These transitions are intended to represent some type of reasoning process or information update. For this initial study, we focus on the operation of product update (restricted to equivalence relations as in Definition 2.4).

3.1. Technical Preliminaries: Generalized Product Update

Our first contribution is to define a sequence of products $\times_{N_n}$ between *Kripke structures*. The idea to apply product update between Kripke structures (rather than Kripke structures and event models) was initially proposed by Jan van Eijck and colleagues [13]. We follow the same basic idea, although our approach differs in a technical, but crucial, way.

In order to generalize the product update operation so that it applies between two Kripke structures, we must replace the precondition function with something appropriate for merging two Kripke structures. Our approach is to explicitly mark which of the formulas we are interested in, and treat these formulas as atomic propositions. [15] Fix a set I of players and At of atomic propositions (for simplicity assume both are finite).

Definition 3.1 (Language extension). 1. Let $\mathcal{T} \subseteq \mathcal{L}_{EL}$ with $\mathsf{At} \subseteq \mathcal{T}$. For every $\varphi \in \mathcal{T}$ we introduce a new constant $\check{\varphi}$ called the name of φ. Let

[15] In general, this type of language extension can be used to model agents with limited memory. For instance, this is needed for an analysis of situations such as the sum and product riddle involving the dialogues: *A: I don't know* φ. *B: I* knew *you didn't know before you said that* (cf. [24] for an analysis of this puzzle in Public Announcement Logic).

$\check{\mathcal{T}} := \{\check{\varphi} | \varphi \in \mathcal{T}\}$. The **language extension** with $\mathcal{T}$, denoted by $\mathcal{L}_{EL}^{\mathcal{T}}$, is the epistemic language with $\check{\mathcal{T}}$ as atomic propositions. By a slight abuse of notation we write p instead of $\check{p}$ for $p \in \mathsf{At} \subseteq \mathcal{T}$. We denote the valuation function over the language $\mathcal{L}_{EL}^{\mathcal{T}}$ by $V_{\mathcal{T}}$. As usual, we omit the subscript when it is clear from the context.

2. Let $\mathcal{M} = \langle W, \{R_i\}_{i\in I}, V\rangle$ be a Kripke model with atomic propositions At and let $\mathcal{T} \subseteq \mathcal{L}_{EL}$ with $\mathsf{At} \subseteq \mathcal{T}$. Then $\mathcal{M}$ can naturally be interpreted as a Kripke model over $\mathcal{L}_{EL}^{\mathcal{T}}$ by defining $V_{\mathcal{T}}$ as: $w \in V_{\mathcal{T}}(\check{\varphi})$ iff $\mathcal{M}, w \vDash \varphi$. We denote $\mathcal{M}$ viewed over $\mathcal{L}_{EL}^{\mathcal{T}}$ by $\mathcal{M}^{\mathcal{T}}$. $\triangleleft$

In $\oplus$-updates every state v in the event model comes with a (generally complex) formula φ that is the precondition for v to occur. That is (w, v) is only defined if $\mathcal{M}, w \vDash \mathsf{pre}(v)$. This is exactly the idea of the $\times_{\mathcal{T}}$ update defined below: pairs of states are in the new model only if they agree on the formulas in $\mathcal{T}$.

Definition 3.2 (Product update). i) Let $\mathcal{T} \subseteq \mathcal{L}_{EL}$ with $\mathsf{At} \subseteq \mathcal{T}$. Let $\mathcal{M} = \langle W, \{R_i\}_{i\in I}, V\rangle$ and $\mathcal{M}' = \langle W', \{R_i'\}_{i\in I}, V'\rangle$ be two Kripke models over $\mathcal{L}_{EL}^{\mathcal{T}}$. The *product model* $\mathcal{M} \times \mathcal{M}' = \langle W'', \{R_i''\}_{i\in I}, V''\rangle$ over $\mathcal{L}_{EL}^{\mathcal{T}}$ is defined as follows:

- $W'' = \{(w, w') \mid w \in W, w' \in W'$ and for all $\check{\varphi} \in \check{T} : w \in V_{\check{\mathcal{T}}}^{\mathcal{M}}(\check{\varphi})$ iff $w' \in V_{\check{\mathcal{T}}}^{\mathcal{M}'}(\check{\varphi})$;
- $(w, w')R_i''(v, v')$ iff wR_iv and $w'R_i'v'$; and
- $(w, w') \in V_{\check{\mathcal{T}}}''(\check{\varphi})$ iff $w \in V_{\check{\mathcal{T}}}^{\mathcal{M}}(\check{\varphi})$ (and thus also $w' \in V_{\check{\mathcal{T}}}'(\check{\varphi})$).

ii) The **generalized product update** of $\mathcal{M}$ and $\mathcal{M}'$ over $\mathcal{T}$, denoted by $\mathcal{M} \times_{\mathcal{T}} \mathcal{M}'$ is the model $\mathcal{M} \times \mathcal{M}'$ as defined above interpreted as a model over $\mathcal{L}_{EL}$. (That is: removing all atoms $\check{\varphi}$ with $\varphi \in \mathcal{T} \setminus \mathsf{At}$ and identifying $\check{p}$ with p for all $p \in \mathsf{At}$.) $\triangleleft$

We write $\mathcal{M} \times_{\mathcal{T}} \mathcal{M}'$ where $\mathcal{M}$ and $\mathcal{M}'$ are Kripke models over $\mathcal{L}_{EL}$, meaning that we interpret $\mathcal{M}$ and $\mathcal{M}'$ as being models over $\check{\mathcal{T}}$ and do the $\times_T$-update as defined above. The procedure that we follow to compute this product runs as follows:

1. Pick a set $\mathcal{T}$ of statements to keep track of,
2. Build the Product in $\mathcal{L}_{EL}^{\mathcal{T}}$, and
3. Remove the additional information, i.e., restrict the valuation function from $\check{\mathcal{T}}$ to At.

The following example demonstrates this procedure.

EXAMPLE 3.3. Let $\mathcal{T} = \{p, K_1 p, K_2 p, K_1 \neg p, K_2 \neg p\}$. Then the product of the two models is calculated as follows.

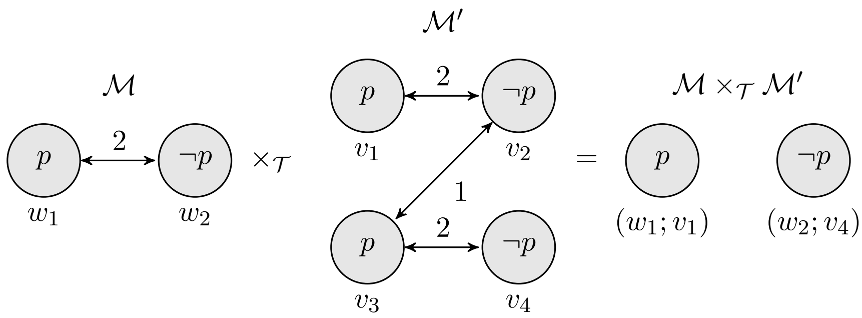

Note that the reflexive and transitive arrows are not drawn in the above picture for simplicity. The set $\mathcal{T}$ is rich enough to uniquely describe all knowledge assignments of level at most one. Thus, the product reflects a merging of models taking into account the agents' first-order information. The fragments of $\mathcal{T}$ true at the individual worlds are:

$$\mathcal{M}, w_1 \vDash \{p, K_1 p\} \qquad \mathcal{M}, w_2 \vDash \{K_1 \neg p\} \qquad \mathcal{M}', v_1 \vDash \{p, K_1 p\}$$
$$\mathcal{M}', v_2 \vDash \emptyset \qquad \mathcal{M}', v_3 \vDash p \qquad \mathcal{M}', v_4 \vDash \{K_1 \neg p\}$$

The only pairs satisfying the same fragment of $\mathcal{T}$ are (w_1, v_1) and (w_2, v_4). Observe that in the model $\mathcal{M} \times_{\mathcal{T}} \mathcal{M}'$ we have:

$$\mathcal{M} \times_{\mathcal{T}} \mathcal{M}', (w_1; v_1) \vDash \{p, K_1 p, K_2 p\}$$

which is different from the fragment of $\mathcal{T}$ satisfied by $\mathcal{M}, w_1$.

In general, taking a generalized product update consists of two steps: The first is picking a set of statements $\mathcal{T} \supseteq \mathsf{At}$ that one wants to keep track of and extending the language to $\mathcal{L}_{EL}^{\mathcal{T}}$. The second is to do generalized product update $\times_{\mathcal{T}}$, that is the normal product $\times$ over $\mathcal{L}_{EL}^{\mathcal{T}}$ followed by omitting all the information about the valuation of $\check{\mathcal{T}} \setminus \mathsf{At}$, i.e., making the newly created model an $\mathcal{L}_{EL}$ model again. The above example shows that the $\times_{\mathcal{T}}$ product does not preserve higher order information.

REMARK 3.4. *There are epistemic models $\mathcal{K}, w$ and $\mathcal{L}, v$ over $\mathcal{L}_{EL}$ a fragment $\mathcal{T}$ of $\mathcal{L}_{EL}$ and some $\varphi \in \mathcal{T} \setminus \mathsf{At}$ such that $(v, w) \in \mathcal{K} \times \mathcal{L}$ (the product over $\mathcal{L}_{EL}^{\mathcal{T}}$) and $\mathcal{K} \times \mathcal{L}, (v, w) \vDash \check{\varphi}$, but $\mathcal{K} \times \mathcal{L}, (v, w) \nvDash \varphi$. (Where in the last formula φ is evaluated as a formula of $\mathcal{L}_{EL}$.)*

There is a close connection between generalized product update and the $\oplus$-update. In both cases, the result is not the complete cartesian product between the two state spaces, but a subset that is characterized by a certain set of formulae. The precise connection between the two concepts is clarified by the following lemma.

LEMMA 3.5. *For every event model $\mathcal{E}$ there is some fragment $\mathcal{T} \subseteq \mathcal{L}_{EL}$ and a Kripke model $\mathcal{M}'$ (for the language $\mathcal{L}_{EL}^{\mathcal{T}}$) such that $\oplus\mathcal{E}$ is the same as $\times_{\mathcal{T}}\mathcal{M}'$ (i.e., for all Kripke models $\mathcal{M}$, $\mathcal{M} \oplus \mathcal{E}$ is isomorphic to $\mathcal{M} \times_{\mathcal{T}} \mathcal{M}'$).*

PROOF. Let $\mathcal{E} = \langle E, \{Q_i\}_{i\in I}, \mathsf{pre}\rangle$ be an event model. Let $\mathcal{T}$ be the set $\{\mathsf{pre}(e) \mid e \in E\} \cup \mathsf{At}$. Construct the model $\mathcal{M}' = \langle W', \{R'_i\}, V'\rangle$ as follows: Let W' be the set of pairs (e, L_e) where $e \in E$ and $L_e \subseteq \mathcal{T}$ is a maximally consistent subset of $\mathcal{T}$ containing $\mathsf{pre}(e)$. The relations R'_i are defined as $(e, L_e)R_i(e', L'_e)$ iff eQ_ie', and the valuation V' is defined by L_e (i.e., $(e, L_e) \in V(\check{\varphi})$ provided $\check{\varphi} \in L_e$). It is easy to check that this $\mathcal{M}'$ has the desired properties.

COROLLARY 3.6. *If there is an upper bound for the quantifier depths of the preconditions in the event model $\mathcal{E}$ (i.e., the set $\{qd(\mathsf{pre}(e)) \mid e \in \mathcal{E}\}$ has an upper bound) then the set $\mathcal{T}$ in the above lemma can be chosen finite. This holds in particular if $\mathcal{E}$ is finite.*

PROOF. Let n be an upper bound for the quantifier depths of $\{\mathsf{pre}(e) \mid e \in E\}$. Recall that $\mathcal{F}_n(\wp(\mathsf{At}))$ is finite, and so, there are characteristic formulae φ_t for every $t \in \mathcal{F}_n(\wp(\mathsf{At}))$ (that is, $\mathcal{F}_n(\wp(\mathsf{At})), s \vDash \varphi_t \Leftrightarrow s = t$). Let $\mathcal{T} := \{\varphi_e \mid e \in \mathcal{F}_n(\wp(\mathsf{At}))\} \cup \mathsf{At}$ and construct the model $\mathcal{M}'$ as follows:

$$W' := \{(e,t)|e \in E, t \in \mathcal{F}_n(\wp(\mathsf{At})) \text{ and } \mathcal{F}_n(\wp(\mathsf{At})), t \models \mathsf{pre}(e)\},$$

let $(e, t)R'_i(e', s)$ if eQ_ie', and define V' as:

$$(e, t) \in V'(\check{\varphi}) \text{ iff } \mathcal{F}_n(\wp(\mathsf{At}), t \vDash \varphi$$

The sets $S = \{\varphi_t \mid t \in \mathcal{F}_n(\wp(\mathsf{At}))\}$ chosen above are special in that these sets reflect all possible knowledge assignments up to depth n. We denote the resulting set of formulas by N_n (i.e., $N_n = \{\varphi_t \mid t \in \mathcal{F}_n(\wp(\mathsf{At}))\} \cup \mathsf{At}$).

REMARK 3.7. *i) In the above proof, we can turn $\mathcal{M}'$ into an event model $\mathcal{E}'$ by letting $\mathsf{pre}(e, t) = \varphi_t$. In this case we have $\mathcal{M} \times_{N_n} \mathcal{M}' = \mathcal{M} \oplus \mathcal{E}'$ for all $\mathcal{M}$. In particular the model $\mathcal{E}'$ is a special event model that* *only has preconditions from N_n. This follows a general pattern: The initial strength of arbitrary event models is that they allow for a very intuitive description of events in a multi-agent setting. However, from a technical point of view arbitrary event models can be difficult to handle. Therefore it sometimes proves useful to translate arbitrary event models into a certain subclass of event models which are easier to work with. For instance, [11] defined a class of canonical event models that are useful for studying when two event models are* equivalent.

ii) The translation of an event model into a Kripke model blurs the distinction between static descriptions of situations and descriptions of events.

There is an interesting peculiarity of the $\times_{\mathcal{T}}$-products. Obviously, $\times_{\mathcal{T}}$ is commutative, but the following example shows that it is not associative.[16]

EXAMPLE 3.8. This example is similar to Example 3.3. Suppose that $\mathcal{T} = \{p, K_1 p, K_2 p, K_1\neg p, K_2\neg p\}$. Consider the following $\mathcal{L}_{EL}$-models which we interpret as $\mathcal{L}_{EL}^{\mathcal{T}}$-models.

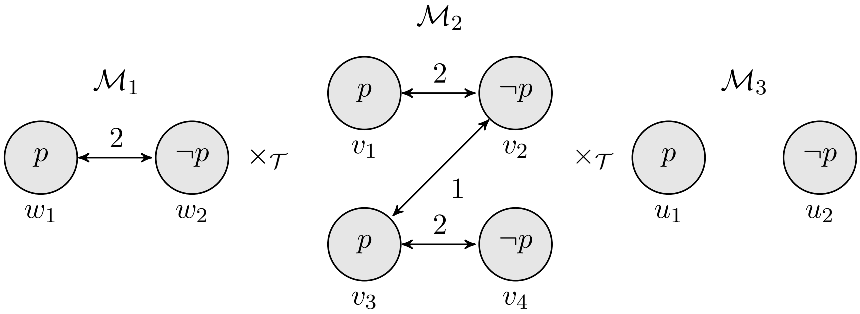

We now show that $(\mathcal{M}_1 \times_{\mathcal{T}} \mathcal{M}_2) \times_{\mathcal{T}} \mathcal{M}_3 \neq \mathcal{M}_1 \times_{\mathcal{T}} (\mathcal{M}_2 \times_{\mathcal{T}} \mathcal{M}_3)$. As we already noted in the previous example (Example 3.3), $\mathcal{M}_1 \times_{\mathcal{T}} \mathcal{M}_2 = \mathcal{M}_3$. In particular, $(\mathcal{M}_1 \times_{\mathcal{T}} \mathcal{M}_2) \times_{\mathcal{T}} \mathcal{M}_3 = \mathcal{M}_3 \times_{\mathcal{T}} \mathcal{M}_3 = \mathcal{M}_3$ where the last equivalence holds since u_1 and u_2 satisfy different formulas from $\mathcal{T}$.

On the other hand, note that the following formulas from $\mathcal{T}$ are true at states in $\mathcal{M}_3$:

$$\mathcal{M}_3, u_1 \vDash \{p, K_1 p, K_2 p\} \qquad \mathcal{M}_3, u_2 \vDash \{K_1\neg p, K_2\neg p\}$$

[16] In general, it is clear that the process of consecutive learning is not commutative. One's actions in some event B can depend on having learned A before. In our formalization, the non-associativity captures this intuition: $(A \times_S B) \times_S C$ is to be read as being in situation A and learning B, then C, whereas $A \times_S (B \times_S C) = A \times_S (C \times_S B)$ corresponds to learning B and C at a time. A similar phenomena has been noticed in the belief merging literature (cf. [21, Section 5.1]).

However, there are no states in $\mathcal{M}_2$ satisfying precisely these formulas, so $\mathcal{M}_2 \times_\mathcal{T} \mathcal{M}_3 = \emptyset$ and consequently $\mathcal{M}_1 \times_\mathcal{T} (\mathcal{M}_2 \times_\mathcal{T} \mathcal{M}_3) = \emptyset$. Thus, we have $(\mathcal{M}_1 \times_\mathcal{T} \mathcal{M}_2) \times_\mathcal{T} \mathcal{M}_3 \neq \mathcal{M}_1 \times_\mathcal{T} (\mathcal{M}_2 \times_\mathcal{T} \mathcal{M}_3)$. [17]

The interpretation of this statement is that first learning $\mathcal{E}$ and then learning $\mathcal{E}'$ is different to learning $\mathcal{E}$ and $\mathcal{E}'$ *at the same time*. To be more precise, we have $(\mathcal{E} \times_\mathcal{T} \mathcal{F}) \times_\mathcal{T} \mathcal{G} \neq \mathcal{E} \times_\mathcal{T} (\mathcal{F} \times_\mathcal{T} \mathcal{G}) \neq \mathcal{E} \times_\mathcal{T} \mathcal{F} \times_\mathcal{T} \mathcal{G}$[18] This non-associativity shows that our framework is rich enough to distinguish between consecutive learning and receiving all information at once.

These observations should be contrasted with the theory developed in [13]. The authors of [13] are concerned with updates where all preconditions are boolean combinations of the ground variables (describing non-epistemic facts about the state of the world). Learning facts about the world is associative (cf. [13, Theorem 1]), whereas learning facts about the players' previous knowledge is not!

Van Eijck *et al.* [13] study the monoid generated by $\times_{\mathsf{At}}$ products. Our primary goal in this paper is to understand how the $\oplus$-update works in type spaces. To that end, we first generalize a result from [10].

THEOREM 3.9. *Let $\mathcal{M}_1$ be a Kripke structure such that for any $v, w \in \mathcal{M}$ there is an epistemic formula φ distinguishing v and w (i.e. $\mathcal{M}, w \vDash \varphi$ and $\mathcal{M}, v \vDash \neg\varphi$). Let $\mathcal{M}_2$ be an arbitrary Kripke structures. Then there is a set of formulas $\mathcal{T}$ and $\mathcal{L}_{EL}^{\mathcal{T}}$-Kripke structure $\mathcal{M}'$ such that $\mathcal{M}_1 \times_\mathcal{T} \mathcal{M}' \leftrightarroweq \mathcal{M}_2$ if and only if there is a total simulation from $\mathcal{M}_2$ to $\mathcal{M}_1$. Furthermore, if the model $\mathcal{M}_1$ is finite the set $\mathcal{T}$ can be chosen finite.*

PROOF. The direction from left to right is easy: Let $\mathcal{M}'$ and $\mathcal{T}$ be such that $\mathcal{M}_1 \times_\mathcal{T} \mathcal{M}' = \mathcal{M}_2$. It is easy to see that the map $\mathcal{M}_1 \times_\mathcal{T} \mathcal{M}' \to \mathcal{M}_1$ sending every pair (w, w') to w is a functional, hence total, simulation.

For the direction from right to left: Let Z be a total simulation from $\mathcal{M}_2$ to $\mathcal{M}_1$. First we define a Kripke model $\mathcal{M}^\circ = \langle W^{\mathcal{M}^\circ}, \{R_i^{\mathcal{M}^\circ}\}_{i\in I}, V^{\mathcal{M}^\circ}\rangle$:

- $W^{\mathcal{M}^\circ} = \{(t_1, t_2) | t_i \in \mathcal{M}_i, i = 1, 2 \text{ and } t_1 Z t_2\}$
- $(t_1, t_2) R_i^{\mathcal{M}^\circ} (s_1, s_2)$ iff $t_1 R_i^{\mathcal{M}_1} s_1$ and $t_2 R_i^{\mathcal{M}_2} s_2$
- $(t_1, t_2) \in V^{\mathcal{M}^\circ}(p)$ iff $t_2 \in V^{\mathcal{M}_2}(p)$ (and thus also $t_1 \in V^{\mathcal{M}_1}(p)$)

---

[17]There are examples where both $(\mathcal{M}_1 \times_\mathcal{T} \mathcal{M}_2) \times_\mathcal{T} \mathcal{M}_3$ and $\mathcal{M}_1 \times_\mathcal{T} (\mathcal{M}_2 \times_\mathcal{T} \mathcal{M}_3)$ are non-empty; however, they are more complicated while making the same point.

[18]Here $\mathcal{E} \times_\mathcal{T} \mathcal{F} \times_\mathcal{T} \mathcal{G}$ is the obvious generalization of $\times_T$ where all tuples (e, f) in the definition are replaced by triples (e, f, g).

First we show that the model $\mathcal{M}^\circ$ is bisimilar to $\mathcal{M}_2$. We show that the projection map π_2 mapping every $(t_1, t_2) \in \mathcal{M}^\circ$ to $t_2 \in \mathcal{M}_2$ is a bitotal bisimulation (recall Definition 2.5). The atom condition is clear. For forth assume that $(t_1, t_2)\pi_2 t_2$ and that $(t_1, t_2)R_i^{\mathcal{M}^\circ}(s_1, s_2)$. By the definition of $R_i^{\mathcal{M}^\circ}$ we have $t_2 R_i^{\mathcal{M}_2} s_2$ and by definition of π_2 we have $(s_1, s_2)\pi_2 s_2$, thus forth is fulfilled.

Similarly, for back assume that $(t_1, t_2)\pi_2 t_2$ and that $t_2 R_i^{\mathcal{M}_2} s_2$. Since Z is a total simulation and $t_1 Z t_2$ holds by the construction of $\mathcal{M}^\circ$, there is some $s_1 \in \mathcal{M}_1$ with $s_1 Z s_2$ and $t_1 R_i^{\mathcal{M}_1} s_1$. But this means that $(s_1, s_2) \in \mathcal{M}^\circ$ and that $(t_1, t_2)R_i^{\mathcal{M}^\circ}(s_1, s_2)$, thus proving the back condition.

Since $\mathcal{M}_2 \leftrightarroweq \mathcal{M}^\circ$, it suffices to show that there is some $\mathcal{M}'$ with $\mathcal{M}_1 \times_{\mathcal{T}} \mathcal{M}' = \mathcal{M}^\circ$.

Note, that the projection $\pi_1 : \mathcal{M}^\circ \to \mathcal{M}_1$ sending each pair (t_1, t_2) to t_1 is a functional left simulation. The atom condition is clear, and the rest can be shown with arguments similar to ones given above.

Now, pick a set $\mathcal{T}^* \subseteq \mathcal{L}_{EL}$ that contains a distinguishing formula for any $v, w \in \mathcal{M}_1$ and let $\mathcal{T} := \mathcal{T}^* \cup \mathsf{At}$. Turn $\mathcal{M}^\circ$ into an $\mathcal{L}_{EL}^{\mathcal{T}}$-model $\mathcal{M}'$ by defining: $(t_1, t_2) \in V^{\mathcal{T}}(\check{\varphi})$ iff $\mathcal{M}_1, t_1 \vDash \varphi$. Since $\mathcal{T}^*$ is separating, $s_1 \in \mathcal{M}_1$ and $(t_1, t_2) \in \mathcal{M}^\circ$ satisfy the same $\check{\mathcal{T}}$-formulae iff $s_1 = t_1$. Therefore $\mathcal{M}_1 \times_{\mathcal{T}} \mathcal{M}' = \mathcal{M}^\circ$ as desired. Furthermore, if $\mathcal{M}_1$ is finite, then the set $\mathcal{T}^*$ can be chosen finite, thus proving the last statement.

Remark 3.10. *[10] contains a proof for a similar statement about $\oplus$-updates in the finite case. However, the generalization to infinite Kripke models does not hold for the $\oplus$-update.*

Remark 3.11. *Note that the model $\mathcal{M}'$ constructed in the right-to-left direction of the prove of Lemma* 3.5 *is in general not a $\mathcal{L}_{EL}$ model that is simply interpreted as an $\mathcal{L}_{EL}^{\mathcal{T}}$ model. That is: There is in general some $\varphi \in \mathcal{T}$ and some $w \in \mathcal{M}'$ such that $\mathcal{M}', w \vDash \check{\varphi}$ but $\mathcal{M}', w \nvDash \varphi$ (where $\check{\varphi}$ is an atom and φ is a formula evaluated in $\mathcal{M}'$ interpreted as a Kripke model over* At *(i.e. only containing atoms from $\{\check{p} \mid p \in \mathsf{At}\}$). That is: to gain the expressive power of updating with an arbitrary event model, one needs the class of all $\mathcal{L}_{EL}^{\mathcal{T}}$-models. Interestingly enough, this is no longer true when we restrict ourselves to the class of finite Kripke structures. There, the full expressive power of the class of all $\oplus$-updates is already given by the class of all finite Kripke models over $\mathcal{L}_{EL}$ together with the set of all $\times_{N_n}$ products for $n \in \omega$. More formally, we have the following fact (whose straightforward, but tedious, proof we leave out).*

Fact 3.12. *Let $\mathcal{K} = \langle W, (R_i)_i, V \rangle$ be a finite Kripke model without bisimilar* *points and let $\mathcal{L} = \langle W', R'_i, V' \rangle$ be a finite Kripke model such that $\mathcal{L}$ is obtainable from $\mathcal{K}$ by a product update. Then, there is some $\mathcal{T} \supseteq \mathsf{At}$ and some Kripke model $\mathcal{M}$ over $\mathcal{L}_{EL}$ such that $\mathcal{K} \times_{\mathcal{T}} \mathcal{M} = \mathcal{L}$.*

3.2. Characterization Result

As discussed in the previous section, every $\oplus$-update can be written as a $\times_{\mathcal{T}}$-update over a language in which the formulas in $\mathcal{T}$ are treated as atomic propositions. This will help us represent the product update in knowledge structures.

First, we need an equivalent to the extension of atomic propositions on types: For $n \in \mathbb{N}$ let S_n denote the set of all possible n-worlds, thus $S_n = \mathcal{F}_n(S)$ and $S_0 = S$). Technically, this is redundant, though it helps conceptually to distinguish $\mathcal{F}_n(S)$ as a type space generated by S and S_n which is the same type space reinterpreted as new set of atoms. By switching between those interpretations, every $n + k$ world over S can be seen as a k-world over S_n and thus there is a canonical embedding $\mathcal{F}_\omega(S) \to \mathcal{F}_\omega(S_n)$.[19]

For any two Kripke models $\mathcal{K}, v$ and $\mathcal{L}, w$ we have defined the product update $(\mathcal{K} \times \mathcal{L}, (v, w))$ over the unextended language $\mathcal{L}_{EL}$ above. Furthermore, we have seen that there is some κ such that r_κ is a bisimulation of $\mathcal{K}$ onto its image. Since $r_\kappa(v)$ is obviously in the image of r_κ this implies that parts of $\mathcal{K}$ are somehow coded in $r_\kappa(v)$. The idea of the following definition is that we can unravel enough information about $\mathcal{K}$ and $\mathcal{L}$ from $r_\kappa(v)$ and $r_\kappa(w)$ to determine $r_\kappa((v, w))$. We define a product $\times_0$ below and we will show later (lemma 3.15) that $r_\kappa((v, w)) = r_\kappa(v) \times_0 r_\kappa(w)$. As with the original definition of a κ-world (see 2.6), the definition is by induction.

DEFINITION 3.13. Suppose that $n \in \mathbb{N}$ and $\mathbf{f}, \mathbf{g} \in \mathcal{F}_\omega(S)$. Then the $\times_0$-product $(\mathbf{f} \times_0 \mathbf{g}) \in \mathcal{F}_\omega(S) \cup \{\emptyset\}$ is defined as follows:

- $(\mathbf{f} \times_0 \mathbf{g})_0 = \langle f_0 \rangle$ iff $f_0 = g_0$ and $\emptyset$ otherwise.
- $(\mathbf{f} \times_0 \mathbf{g})_m(i) = \{(\mathbf{f}' \times_0 \mathbf{g}')_{m-1} \mid \mathbf{f}' \in f_m(i), \mathbf{g}' \in g_m(i)\}$

This definition can be lifted to an analogue of the generalized product update: The operator $\times_n$ will correspond to a product update with $\mathcal{T} = N_n$. First observe that the above definition of $\times_0$ works equally well if all S are replaced by S_n. As in the case of the generalized product update, the $\times_n$ update implicitly consists of two steps: First a product update between two

[19] Note that this map is not surjective for $n \geq 1$: For instance the introspection conditions of $\mathcal{F}_{k+1}(S)$ gives some limitations on which elements of $\mathcal{F}_2(S_k)$ can come from $\mathcal{F}_{k+1}(S)$.

elements of $\mathcal{F}_\kappa(S_n)$ followed by a removal of information, i.e. a projection from S_n to S. As with general product updates, the definition contracts these two steps into one:

DEFINITION 3.14. Let $\bar{\pi} : S_n \to S$ be the projection map sending the tuple $\langle f_0 \dots f_{n-1} \rangle$ to f_0. Define $\times_n : \mathcal{F}_\omega(S_n) \times \mathcal{F}_\omega(S_n) \to \mathcal{F}_\omega(S)$ as follows:

- $(\mathbf{f} \times_n \mathbf{g})_0 = \langle s_0 \rangle$ iff $\bar{\pi}(f_0)) = \bar{\pi}(g_0) = s_0$, and $\emptyset$ otherwise.
- $(\mathbf{f} \times_n \mathbf{g})_m(i) = \{(\mathbf{f}' \times_n \mathbf{g}')_{m-1} \mid \mathbf{f}' \in f_m(i), \mathbf{g}' \in g_m(i)\}$. ◁

The following lemma describes the relationship between the $\times_{N_n}$-product and the $\times_n$-product. Basically, the $\times_{N_n}$ product of two Kripke models $(\mathcal{K}, w)$ and $(\mathcal{L}, v)$ carries the same information as the $\times_n$-product on the types $r(v)$ and $r(w)$.

For technical convenience we need a definition before we state the lemma: Recall that $N_n \backslash \mathsf{At}$ was chosen to be a set of characteristic formulae for $\mathcal{F}_n(S)$. In particular, every state w in a Kripke structure $\mathcal{K}$ over $\mathcal{L}_{EL}$ satisfies exactly one formula of $N_n \setminus \mathsf{At}$. In particular for any Kripke model $\mathcal{L}$ over $\mathcal{L}_{EL}^{N_n}$ we have that $(v, w) \in \mathcal{K} \times_{N_n} \mathcal{L}$ implies that there is exactly one $\check{\varphi} \in N_n \setminus \mathsf{At}$ with $w \in V(\check{\varphi})$. We call Kripke models over $\mathcal{L}_{EL}^{N_n}$ satisfying this property **admissible**. Since every $e \in \mathcal{F}_n(S)$ satisfies exactly one formula from $N_n \backslash \mathsf{At}$ we have that every state of nature in an admissible Kripke model corresponds to exactly one $e \in \mathcal{F}_n(S)$ and we can define a map r' from admissible Kripke models to $\mathcal{F}_\omega(S_n)$ in the same way as we defined r.

LEMMA 3.15. *Let $n \in \omega$ and let $\mathcal{K}, \mathcal{L}$, be Kripke models over $\mathcal{L}_{EL}^{N_n}$. Let $v \in \mathcal{K}$, $w \in \mathcal{L}$ satisfaying the same N_n-formulae. Let $(v, w) \in \mathcal{K} \times_{N_n} \mathcal{L}$ denote the product of v and w in $\mathcal{K} \times_{N_n} \mathcal{L}$. Then we have $r((v, w)) = r'(v) \times_n r'(w)$, i.e., the following diagram commutes:*

$$\begin{array}{ccc} \mathcal{K}, \mathcal{L} & \xrightarrow{\times_{N_n}} & \mathcal{K} \times_{N_n} \mathcal{L} \\ \downarrow{\scriptstyle r', r'} & & \downarrow{\scriptstyle r} \\ \mathcal{F}_\omega(S_n), \mathcal{F}_\omega(S_n), & \xrightarrow{\times_n} & \mathcal{F}_\omega(S) \end{array}$$

PROOF. Let $n \in \mathbb{N}$ and $v \in \mathcal{K}$, $w \in \mathcal{L}$ satisfying the same N_n-formulae. We inductively show that $(r'(v) \times_n r'(w))_k = r(v, w)_k$. for $k = 0$ this is trival: If v and w satisfy the same atomic propostions over $\check{N}_n$ we have $(r'(v) \times_n r'(w))_0 = r((v, w))_0 = \{p \in \mathsf{At} : v \in V^{\mathcal{K}}(p)\}$. If they satisfy different atomic propositions we have $(v, w) \notin \mathcal{K} \times_{N_n} \mathcal{L}$ and $r'(v) \times_n r'(w) = \emptyset$. Now assume the statement holds for $k - 1$ and let $i \in I$ (the set of agents).

First, we show $r(v,w)_k(i) \subseteq (r'(v) \times_n r'(w))_k(i)$. Let $x \in r((v,w))_k(i)$, thus x is a $k-1$-world. By construction of the map r there is some $\tilde{x}$ in $\mathcal{K} \times_{N_n} \mathcal{L}$ such that $\tilde{x}R_i(v,w)$ and $r(\tilde{x})_{k-1} = x$. Thus there are $x_1 \in \mathcal{K}$ and $x_2 \in \mathcal{L}$ such that the product of x_1 and x_2 in $\mathcal{K} \times_{N_n} \mathcal{L}$ is $\tilde{x}$ - in particular $x_1 R_i v$ and $x_2 R_i w$ and x_1 and x_2 satisfy the same N_n-formulae. In particular, $r'(x_1) \times_n r'(x_2) \neq \emptyset$ and by induction we have that $(r(x_1,x_2))_{k-1} = (r'(x_1) \times_n r'(x_2))_{k-1}$. On the other hand, we have $r'(x_1)_{k-1} \in r'(v)_k(i)$ and similarly for x_2 and w by the construction of r'. In particular, we have $x = (r'(x_1) \times_n r'(x_2))_{k-1} \in (r'(v) \times_n r'(w))_k(i)$ as desired, thus proving the first direction.

The argument for the reverse inclusion $r(v,w)_k(i) \supseteq (r'(v) \times_n r'(w))_k)(i)$ is similar: Let $x \in (r('v) \times_n r'(w))_k(i)$. Then there are $\tilde{x}_1 \in r'(v)$ and $\tilde{x}_2 \in r'(w)$ such that $(r'(\tilde{x}_1) \times_n r'(\tilde{x}_2))_{k-1} = x$ and such that there are $x_1 \in \mathcal{K}, x_2 \in \mathcal{L}$ such that $r'(x_i) = \tilde{x}_i$ and $x_1 R_i v$ and $x_2 R_i w$ hold. Since $\tilde{x}_1 \times_n \tilde{x}_2$ exists, x_1 and x_2 satisfy the same N_n-formulae. In particular there is some (x_1,x_2) in $\mathcal{K} \times_{N_n} \mathcal{L}$ with $(x_1,x_2)R_i(v,w)$. By construction of r we have $r((x_1,x_2))_{k-1} \in r((v,w))_k$ and by induction we have $r((x_1,x_2))_{k-1} = x$, thus proving the reverse direction.

Note that the calculation of $\mathbf{f} \times_n \mathbf{g}$ from types $\mathbf{f}$ and $\mathbf{g}$ is computationally efficient: In order to calculate the k-th level of $\mathbf{f} \times_n \mathbf{g}$ only the first $n+k$ levels of $\mathbf{f}$ and $\mathbf{g}$ are required.

The above definition of $\times_n$ updates gives a way of modeling dynamics on a type space—thus, opening up the field of epistemic game theory to belief dynamics. Event models were designed as a very intuitive and natural tool for representing epistemic events in a multiagent setting. The translation of event models into the corresponding pair of Kripke models and a product relation $\times_{N_n}$, and further into a type and a relation $\times_n$ allows us to calculate the change of epistemic status brought about by an event model $\mathcal{E}$.

On the other hand, every product update with a finite event model can be written as a $\times_n$-update, thus it suffices to understand the structure of $\times_n$ to study product updates. Thus, $\mathcal{F}_\omega(\wp(S))$ is not only a universal Kripke model in the static sense, together with the products $\times_n$ is also universal in that it incorporates all potential updates.

On Kripke structures, translating event models into types allows us to study updating events as separate entities without any reference to a ground type. Furthermore, the translation blurs the distinction between types as static descriptions of epistemic states and knowledge changing events.

One natural and important question is: Given two types $\mathbf{f}$ and $\mathbf{g}$, is there a possible piece of incoming information that transforms $\mathbf{f}$ into $\mathbf{g}$?

The intuition behind the answer given by the following theorem is: In the entire model, the agents are assumed to be omniscient and non-forgetting. Thus, an event cannot add any uncertainty about the state of nature, it can only remove some states from the sets of possible states. In contrast, for the higher order information, essentially anything is possible as long as it is compatible with individuals gaining new information about the state of nature. In particular, an epistemic event may increase the uncertainty about other agents' types. This idea is captured by the following definition.

DEFINITION 3.16 (Admissibility of Types). For a type $\mathbf{f} \in \mathcal{F}_\alpha(S)$ we say that a type $\mathbf{g}$ is **admissible** for $\mathbf{f}$ iff

- $f_0 = g_0$;
- for all agents i: $g_1(i) \subseteq f_1(i)$; and
- for $\alpha > 1$: If $\mathbf{h} \in g_\alpha(i)$ then there is some $\mathbf{h}' \in f_\alpha(i)$ such that h is admissible for $\mathbf{h}'$. $\triangleleft$

Our characterization theorem is similar to Theorem 3.9.

THEOREM 3.17. *Let $\mathbf{f}, \mathbf{g} \in \mathcal{F}_\alpha(S)$ be types such that $\mathbf{g}$ is obtainable by an update from $\mathbf{f}$, i.e. there is some n and some $\mathbf{h} \in \mathcal{F}_\alpha(S_n)$ such that $\mathbf{f} \times_n \mathbf{h} = \mathbf{g}$. Then $\mathbf{g}$ is admissible for $\mathbf{f}$. If the submodel of $\mathcal{F}_\omega(S)$ generated by $\mathbf{f}$ is finite also the converse holds true.*

Before we can prove this theorem, we recall the following result from infinite combinatorics.

THEOREM 3.18. *(König's Lemma) Let T be an infinite, finitely branching tree. Then, T has an infinite branch.*

PROOF. Construct an infinite branch $\langle x_0, x_1, \ldots \rangle$ as follows: x_0 is the root. For $i > 0$: If $x_0, \ldots x_i$ are already in the branch, pick a successor x_{i+1} of x_i that on an infinite path (since the tree is finitely branching such a successor always exists). Then $\langle x_0, x_1, \ldots \rangle$ is an infinite branch.

PROOF. of Theorem 3.17 The first statement is straightforward: Let $\mathcal{F}$ and $\mathcal{G}$ be the epistemic submodels of $\mathcal{F}_\omega(S)$ induced by $\mathbf{f}$ and $\mathbf{g}$, respectively. Assume that there is some $\mathbf{h} \in \mathcal{F}_\omega(S_n)$ such that $\mathbf{f} \times_n \mathbf{h} = \mathbf{g}$. By Lemma 3.15, this is equivalent to saying that $\mathcal{F} \times_{N_n} \mathcal{H} = \mathcal{G}$, where $\mathcal{F}, \mathcal{G}, \mathcal{H}$ are the generated Kripke models (over $\mathcal{L}_{EL}^{N_n}$) from $\mathbf{f}, \mathbf{g}$, and $\mathbf{h}$. By Theorem 3.9 there is a total simulation S from G to F. We inductively show that every $\mathbf{g}' \in G$ is admissible for every $\mathbf{f}' \in F$ with $\mathbf{f}' S \mathbf{g}'$. The 0th-level is clear by definition of a simulation. Now it suffices to show that definition of admissibility is fulfilled at the 1st level: Since we do this for all $\mathbf{g}' \in G$ the rest follows from the inductive definition of admissibility and the map r. To see that admissibility is fulfilled at the $1st$ level, let $\mathbf{h} \in G$ with $\mathbf{g}' \sim_i \mathbf{h}$. By definition, there is a $\mathbf{h}' \in \mathcal{F}$ with $\mathbf{f}' \sim_i \mathbf{h}'$. Thus, every state of nature that is conceivable for agent i in $\mathcal{G}$ via $\mathbf{h}$ is also conceivable in $\mathcal{F}$ via $\mathbf{h}'$ - this is exactly the definition of being admissible in the first level.

For the second statement let $\mathbf{g}$ be admissible for $\mathbf{f}$ and let the submodel of $\mathcal{F}_\omega(\wp(S))$ generated by $\mathbf{g}$ be finite. Again, let $\mathcal{F}$ and $\mathcal{G}$ be the Kripke submodels of $\mathcal{F}_\omega(S)$ induced by $\mathbf{f}$ and $\mathbf{g}$. Define the Relation Z between $\mathcal{F}$ and $\mathcal{G}$ as $\mathbf{f}' Z \mathbf{g}'$ iff $\mathbf{g}' \in \mathcal{G}$ is admissible for $\mathbf{f}' \in \mathcal{F}$. We will show that Z is a total simulation from $\mathcal{G}$ to $\mathcal{F}$, thus showing that $\mathcal{G}$ is obtainable by $\mathcal{F}$ via update (again using Theorem 3.9 and Lemma 3.15).

By assumption, $\mathbf{g}$ is admissible for $\mathbf{f}$. We show that whenever $\mathbf{g}' \in \mathcal{G}$ is admissible for $\mathbf{f}' \in \mathcal{F}$ and $\tilde{\mathbf{g}} \sim_i \mathbf{g}'$, then there is some $\tilde{\mathbf{f}} \sim_i \mathbf{f}'$ such that $\tilde{\mathbf{g}}$ is admissible for $\tilde{\mathbf{f}}$. This proves that Z is a left simulation. To see that Z is total, note that for every $\mathbf{g}'$ in $\mathcal{G}$ there is a chain $\mathbf{g} \sim_{i_1} \mathbf{g}_1 \sim_{i_2} \dots \sim_{i_n} \mathbf{g}'$ connecting $\mathbf{g}$ with $\mathbf{g}'$. Let $\mathbf{g}' \in \mathcal{G}$ be admissible for $\mathbf{f}' \in \mathcal{F}$ and $\tilde{\mathbf{g}} \sim_i \mathbf{g}'$. We construct an ω-tree $(T, \prec)$ as follows: The k-th level consists of all those types in $f'_{k+1}(i)$ that enlarge $\tilde{g}_k$. The $\prec$-relation is defined as $\mathbf{r} \prec \mathbf{s}$ iff $\mathbf{r}$ is an initial segment of $\mathbf{s}$. By definition of the admissibility relation, every finite level of T is non-empty. Since the state of nature is considered finite, every nonempty level is also finite. Thus, by König's lemma T has an infinite path P. By construction, $\tilde{\mathbf{f}} = \bigcup_{\mathbf{r} \in P} \mathbf{r}$ is a type and $\tilde{\mathbf{g}}$ is admissible for $\tilde{\mathbf{f}}$. Since $\mathcal{F}$ is the substructure of $\mathcal{F}_\omega(S)$ induced by $\mathbf{f}$ (and thus by $\mathbf{f}'$) we have $\tilde{\mathbf{f}} \in \mathcal{F}$, thus the simulation Z relates $\tilde{\mathbf{g}}$ to $\tilde{\mathbf{f}}$.

Again, there is an obvious counterpart of Remark 3.11 allowing us to update with $\mathcal{F}(S)$ worlds rather than $\mathcal{F}(S_n)$ worlds, provided all the induced Kripke structures involved are finite. To be precise, we can show the following: Let $\mathbf{f}, \mathbf{g} \in \mathcal{F}_\omega(S)$ be such that the epistemic submodels of $\mathcal{F}_\omega(S)$ induced by $\mathbf{f}$ and $\mathbf{g}$ are finite. Then $\mathbf{g}$ is admissible for $\mathbf{f}$ if and only if there is some natural number n and some $\mathbf{h} \in \mathcal{F}_\omega(S)$ such that $\mathbf{f} \times_n \mathbf{h} = \mathbf{g}$.

4. Conclusion and Future Work

Many different formal models have been used to describe the players knowledge and beliefs in game-theoretic situations. The variety of models reflect different mathematical conventions used by the various sub-communities, as well as competing intuitions about how best to describe the players' beliefs and reasoning in a game situation. It is important to understand the precise relationship between the alternative modeling paradigms. In this paper, we focused on the two most prominent models found in the literature on the epistemic foundations of game theory: Kripke- or Aumann- structures and knowledge structures (non-probabilistic variants of Harsanyi type spaces).

There are two main contributions in this paper. The first is to initiate a study of "information dynamics" for knowledge structures in the style of recent work on *dynamic epistemic logic* (cf. [3]). Such a theory would further illustrate the subtle relationship between type spaces and Kripke structures (updating the discussion initiated in [15, 16]). In particular, it allows us to combine the strengths of both approaches and use event models as a tool to describe epistemic events. The main technical contribution is the definition of a product operation $\times_n$ on the type space $\mathcal{F}_\omega(S)$. We provide a procedure that allows us to translate arbitrary event models into types. Furthermore, we show that the $\times_n$ product is powerful enough to simulate all updates by event models. Furthermore, we prove a characterization theorem (Theorem 3.17) showing when a type can be transformed into another type by updates with an event model.

This is only an initial study. We see our work here opening up many different avenues of future research. In particular, we plan on investigating the following issues in the future.

- What happens if we allow only updating types from a certain subclass of $\mathcal{F}_\alpha(S_n)$ (for example, finite epistemic models $\langle \mathcal{F}_\alpha(S_n), \{\sim_i\}_{i\in I}, V\rangle$)?
- What are the "behavioral" implications of our main characterization theorem (Theorem 3.17)? For example, if a strategy is rational for a type $\mathbf{f}$ in a game G, does that strategy remain rational for all types that are admissible for $\mathbf{f}$?
- How do we extend the ideas developed in this paper to Harsanyi type spaces, where the beliefs are represented by probability measures? The first step is to generalize the dynamic epistemic logic framework to settings where beliefs are represented by probabilities. Fortunately, this has largely been done (see [4, 1] for details). A very interesting direction for future research is to explore how to use the probabilistic event models and product update operation of [4] to prove a result analogous to our main characterization theorem (Theorem 3.17) for Harsanyi type spaces.
- The relation "obtainable by an update" together with our extended theorem (see Remark 3.11) turns the set of finite induced submodels of $\mathcal{F}_S(w)$ into an algebra. Can we characterize this algebra?

References

[1] Aceto, Luca, Wiebe van der Hoek, Anna Ingólfsdóttir, and Joshua Sack, 'Sigma algebras in probabilistic epistemic dynamics', in *Proceedings of the 13th conference on Theoretical Aspects of Rationality and Knowledge, ACM, 2011, pp. 191–199.*

[2] Aumann, Robert, 'Interactive epistemology I: Knowledge', *International Journal of Game Theory*, 28 (1999), 263–300.

[3] van Benthem, Johan, *Logical Dynamics of Information and Interaction*, Cambridge University Press, 2011.

[4] van Benthem, Johan, Jelle Gerbrandy, and Barteld Kooi, 'Dynamic update with probabilities', *Studia Logica: An International Journal for Symbolic Logic*, 93 (2009), 1, 67 – 96.

[5] van Benthem, Johan, Eric Pacuit, and Olivier Roy, 'Toward a theory of play: A logical perspective on games and interaction', *Games*, 2 (2011), 1, 52 – 86.

[6] Bonanno, Giacomo, and Pierpaolo Battigalli, 'Recent results on belief, knowledge and the epistemic foundations of game theory', *Research in Economics*, 53 (1999), 2, 149–225.

[7] Brandenburger, Adam, 'Epistemic game theory: Complete information', in Steven N. Durlauf, and Lawrence E. Blume, (eds.), *The New Palgrave Dictionary of Economics*, Palgrave Macmillan, 2008.

[8] Brandenburger, Adam, 'Origins of epsitemic game theory', in Vincent F. Hendricks, and Olivier Roy, (eds.), *Epistemic Logic: Five Questions*, Automatic Press, 2010, pp. 59–69.

[9] Brandenburger, Adam, and Amanda Friedenberg, 'Self-admissible sets', *Journal of Economic Theory*, 145 (2010), 785 – 811.

[10] van Ditmarsch, Hans, and Tim French, 'Simulation and information: Quantifying over epistemic events', in John-Jules Meyer, and Jan Broersen, (eds.), *Knowledge Representation for Agents and Multi-Agent Systems*, vol. 5605 of *Lecture Notes in Computer Science*, Springer Berlin / Heidelberg, 2009, pp. 51–65.

[11] van Ditmarsch, Hans, Ji Ruan, and Rineke Verbrugge, 'Sum and product in dynamic epistemic logic', *Journal of Logic and Computation*, 18 (2008), 563–588.

[12] van Ditmarsch, Hans, Wiebe van der Hoek, and Barteld Kooi, *Dynamic Epistemic Logic*, Synthese Library, Springer, 2007.

[13] van Eijck, Jan, Yanjing Wang, and Floor Sietsma, 'Composing models', in *Proceedings of LOFT 2010*, 2010.

[14] Fagin, Ron, 'A quantitative analysis of modal logic', *Journal of Symbolic Logic*, 59 (1994), 1, 209 – 252.

[15] Fagin, Ron, John Geanakoplos, Joseph Halpern, and Moshe Vardi, 'The hierarchical approach to modeling knowledge and common knowledge', *International Journal of Game Theory*, 28 (1999), 3, 331 – 365.

[16] Fagin, Ron, Joseph Halpern, and Moshe Vardi, 'A model-theoretic analysis of knowledge', *Journal of the ACM*, 91 (1991), 2, 382 – 428.

[17] Fagin, Ron, and Moshe Vardi, 'An internal semantics for modal logic: Preliminary report', in *Proc. 17th ACM SIGACT Symposium on Theory of Computing*, 1985, pp. 305 – 315.

[18] Fagin, Ronald, Joseph Halpern, Yoram Moses, and Moshe Vardi, *Reasoning about Knowledge*, The MIT Press, 1995.

[19] Harsanyi, John, 'Games with incomplete informations played by 'bayesian' players', *Management Science*, 14 (1967), 159–182, 320–334, 486–502.

[20] van der Hoek, Wiebe, and Marc Pauly, 'Modal logic for games and information', in P. Blackburn, J. van Benthem, and F. Wolter, (eds.), *Handbook of Modal Logic*, vol. 3 of *Studies in Logic*, Elsevier, 2006, pp. 1077 – 1148.

[21] Maynard-Reid II, Pedrito, and Yoav Shoham, 'From belief revision to belief fusion', in *Proceedings of LOFT-98*, 1998.

[22] Myerson, Roger, 'Harsanyi's games with incomplete information', *Management Science*, 50 (2004), 12, 1818–1824.

[23] Perea, Andrés, *Epistemic Game Theory: Reasoning and Choice*, Cambridge University Press, 2012.

[24] Sietsma, Floor, and Jan van Eijck, 'Action emulation between canonical models', Proceedings of the Conference on Logic and the Foundations of Game and Decision Theory 2012.

[25] Siniscalchi, Marciano, 'Epistemic game theory: Beliefs and types', in S. Durlauf, and L. Blume, (eds.), *The New Palgrave Dictionary of Economics*, Palgrave Macmillan, Basingstoke, 2008.

Dominik Klein
Tilburg Center for Logic and Philosophy of Science
Tilburg University
P.O. Box 90153, 5000 LE
Tilburg, The Netherlands
`d.klein@uvt.nl`

Eric Pacuit
Tilburg Center for Logic and Philosophy of Science
Tilburg University
P.O. Box 90153, 5000 LE
Tilburg, The Netherlands
`e.j.pacuit@uvt.nl`