Synthese (2014) 191:371–408
DOI 10.1007/s11229-012-0215-7

Exploring the tractability border in epistemic tasks

Cédric Dégremont · Lena Kurzen · Jakub Szymanik

Received: 27 February 2012 / Accepted: 29 October 2012 / Published online: 19 December 2012

Abstract We analyse the computational complexity of comparing informational structures. Intuitively, we study the complexity of deciding queries such as the following: Is Alice's epistemic information strictly coarser than Bob's? Do Alice and Bob have the same knowledge about each other's knowledge? Is it possible to manipulate Alice in a way that she will have the same beliefs as Bob? The results show that these problems lie on both sides of the border between tractability (P) and intractability (NP-hard). In particular, we investigate the impact of assuming information structures to be partition-based (rather than arbitrary relational structures) on the complexity of various problems. We focus on the tractability of concrete epistemic tasks and not on epistemic logics describing them.

Keywords Epistemic logic · Computational complexity · Epistemic reasoning · Multi-agent systems

1 Introduction

Having a social life is engaging and complicated. Still some of its aspects seem harder than others. For instance, telling the truth is much easier than lying. Deception demands

C. Dégremont (✉) · J. Szymanik
Institute of Artificial Intelligence, University of Groningen, Groningen, The Netherlands
e-mail: cedric.uva@gmail.com

J. Szymanik
e-mail: jakub.szymanik@gmail.com

L. Kurzen
Institute for Logic, Language and Computation, University of Amsterdam, Amsterdam, The Netherlands
e-mail: lena.kurzen@gmail.com

the ability to simulate another's reaction in order to determine if a lie will be believable. It is even harder if the goal of the lie is to manipulate. To successfully inculcate an idea in someone's mind I need to predict the way she will revise her beliefs and knowledge. How hard is it to be a skilful liar and what makes it hard for others to manipulate our beliefs? We study the difficulty of these problems from a computer science perspective by formalising various notions of information similarity along the lines of epistemic logic.

Epistemic modal logics and their extensions are concerned with global and abstract problems in reasoning about information. They are designed to model a wide range of epistemic scenarios (cf. Fagin et al. 1995; Baltag and Moss 2004). As logics have to be quite complex in order to be able to express various problems in epistemic reasoning, it is not surprising that there are many intractability and even undecidability results in the literature (see e.g., Halpern and Vardi 1989; Van Benthem and Pacuit 2006 for a survey). Consequently, the issue of trade-off between expressivity and complexity plays a central role in the field of epistemic modal logics.

The existing complexity results of modal logics provide an overview of the difficulty of epistemic reasoning in modal logic frameworks from the abstract global perspective of the modeller. For instance, Baral and Zhang (2005) studies the computational complexity of model checking for knowledge update. The authors have proved that in general the problem is Σ_2^P-complete but there are subclasses of formulas for which the complexity is polynomial and such that the complexity of model-checking is NP-complete. Our paper in a sense follows this tradition but we focus on the semantic notions of similarity between information states. Moreover, while Baral and Zhang (2005) focuses entirely on single-agent S5 structures we investigate the impact of assuming different properties of the accessibility relation on the complexity of various problems, as well as situations involving more than one agent, with relations that cannot be simulated in a single-agent S5 framework.

An analysis of these problems calls for other similarity concepts between epistemic structures. Our main aim is also to initiate the study of the tractability of epistemic tasks rather than epistemic logics. As a result, we can identify a theoretical threshold in the difficulty of reasoning about information, similarly to how this has been done in the context of reasoning with quantifiers (cf. Pratt-Hartmann and Moss 2009; Szymanik 2010). In order to do this, we shift our perspective: instead of investigating the complexity of a given logic that can be used to describe certain tasks in epistemic reasoning, we study the complexity of the concrete tasks themselves, determining what computational resources are needed in order to perform the required reasoning.

At this point, we would like to make clear that we do not propose a new formal model for epistemic reasoning from an internal agent-oriented perspective. For two approaches to modelling epistemic scenarios, such as the muddy children puzzle, in a more concise way than standard epistemic logic models, we refer the reader to Gierasimczuk and Szymanik (2011a,b) and Wang (2010). For a version of epistemic logic in which the modeller is one of the agents, we refer to Aucher (2010). In this paper, we work with models from epistemic modal logic and investigate the complexity of various interesting specific problems that arise when reasoning about these semantic structures.

Focusing on specific problems, the complexity may be much lower since concrete problems involved in the study of multi-agent interaction are rarely as general as e.g., satisfiability. In most cases, checking whether a given property is satisfied in a given (minimal) epistemic scenario is sufficient. This may sound as if we study the model checking complexity of different properties in epistemic logic. Indeed, some of the simpler tasks and problems we consider boil down to data complexity of model checking epistemic formulas, i.e., the computational complexity with respect to the size of the model only, abstracting from the size of the formula. However, we want to point out that we study the problems in purely semantic terms and our complexity results are thus independent of how succinctly, if at all, the properties could be expressed in an (extended) epistemic modal logic. The problems we consider in this work take epistemic models and sometimes also additional parameters as input and ask whether the models satisfy certain properties or whether the models are in a certain relation (cf. Baral and Zhang 2005).

Many of the concrete problems we study turn out to be easily computable (in polynomial time). We call such problems tractable. Still, we will see that even in this perspective there are some intractable problems that demand exponential computations. We believe that the feasibility of epistemic tasks, and especially the divide between tractable and intractable problems, is an interesting new topic for a formal study. It can help in the empirical assessment of the cognitive plausibility of epistemic logic frameworks. First of all, modal logic should not postulate intractable models as a description of epistemic cognitive tasks that people can deal with in everyday life without any problems (cf. Van Rooij 2008). Moreover, the complexity results, that our study aims to identify, should correlate with the difficulties faced by human agents solving such tasks (cf. Verbrugge 2009; Szymanik and Zajenkowski 2010).

So in a sense, we aim to initiate a search for an appropriate perspective and complexity measures that describe in plausible ways the cognitive difficulties agents face while interacting. Certain experimental results in the literature explore similar directions for specific game settings (Feltovich 2000; Weber 2001; Meijering et al. 2012).

In this paper we investigate the computational complexity of various decision problems that are relevant for interactive reasoning in epistemic modal logic frameworks. In particular, we explore the complexity of comparing and manipulating information structures possessed by different agents.

With respect to the comparison of information structures we are interested in whether agents have similar information (about each other) or whether one of them has more information.

Information similarity and symmetry

- Is one agent's information strictly coarser than another agent's information?
- Do two agents have the same knowledge/belief about each other's knowledge/belief?

In a situation with diverse agents that have different information, the question arises as to whether it is possible that some of the agents can be provided with information so that afterward the agents have similar information.

Information manipulation

- Given two agents, is it possible to give some information to one of them such that as a result
  - both agents have similar information structures? (cf. Van Ditmarsch and French 2009)
  - one of them has more refined information than the other?

Determining the complexity of the above questions will then help to analyse how the complexity of various reasoning tasks is influenced by

- the choice of similarity notion taken for similarity of information structures,
- the choice of information structures,
- the number of agents involved.

The next section will introduce and motivate the concepts we will be using throughout the paper. As mentioned previously all our results will be stated in purely semantic terms. Nevertheless, most of the concepts are closely related to syntactic, modal logic, approaches to knowledge in multi-agent systems. We will therefore introduce such logics to illustrate the use of our semantic concepts and motivate them.

2 Preliminaries

2.1 Basic epistemic logic

We start by briefly giving some preliminaries of (epistemic) modal logic. We use relational structures from epistemic logic to model information (cf. Blackburn et al. 2001; Fagin et al. 1995). Kripke models can compactly represent the information agents have about the world and about the information possessed by the other agents. For a more exhaustive conceptual introduction to epistemic logic the reader can consult e.g., (Fagin et al. 1995, Chap. 2) (Table 1).

Definition 2.1 (*Kripke Models*) A *Kripke model* $\mathcal{M}$ based on a finite set of agents $\mathtt{N}$ is of the form $(W, (R_i)_{i \in \mathtt{N}}, \mathsf{V})$, where $W \neq \emptyset$ is a set of possible worlds, for each $i \in \mathtt{N}$, R_i is a binary accessibility relation on W. A modal language with a set of propositions PROP is interpreted over $\mathcal{M}$ by a valuation $\mathsf{V} : \text{PROP} \rightarrow \wp(W)$.

Table 1 EL (also called $S5_{\mathtt{N}}$) axiom system

PL	$\vdash \varphi$ if φ is a substitution instance of a tautology of propositional logic
For $i \in \mathtt{N}$,	
Nec	if $\vdash \varphi$, then $\vdash K_i\varphi$
K	$\vdash K_i (\varphi \rightarrow \psi) \rightarrow (K_i\varphi \rightarrow K_i\psi)$
T	$\vdash K_i\varphi \rightarrow \varphi$
4	$\vdash K_i\varphi \rightarrow K_i K_i\varphi$
5	$\vdash \neg K_i\varphi \rightarrow K_i\neg K_i\varphi$
MP	if $\vdash \varphi \rightarrow \psi$ and $\vdash \varphi$, then $\vdash \psi$

It is frequently assumed that information structures are partition-based (Aumann 1999; Fagin et al. 1995; Osborne and Rubinstein 1994):

Definition 2.2 An *epistemic model* is a (multi-agent) Kripke model such that for all $i \in \mathrm{N}$, R_i is an equivalence relation. (We usually write $\sim_i$ instead of R_i).

To explicitly talk about knowledge one may use the language of basic epistemic logic.

Definition 2.3 (*Syntax of* $\mathcal{L}_{\mathrm{EL}}$) The syntax of epistemic language $\mathcal{L}_{\mathrm{EL}}$ is recursively defined as follows:

$$\varphi := p|\neg\varphi|\varphi \vee \varphi|K_i\varphi$$

where $p \in \textsc{prop}$, $i \in \mathrm{N}$. We will write $\top$ for $p \vee \neg p$ and $\bot$ for $\neg\top$. Other connectives ($\wedge$, $\rightarrow$, $\leftrightarrow$) are defined in the usual way.

Definition 2.4 (*Semantics of* $\mathcal{L}_{\mathrm{EL}}$) We interpret $\mathcal{L}_{\mathrm{EL}}$ in the states of epistemic models as follows:

$\mathcal{M}, w \models p$	iff	$w \in V(p)$
$\mathcal{M}, w \models \neg\varphi$	iff	it is not the case that $\mathcal{M}, w \models \varphi$
$\mathcal{M}, w \models \varphi \vee \psi$	iff	$\mathcal{M}, w \models \varphi$ or $\mathcal{M}, w \models \psi$
$\mathcal{M}, w \models K_i\varphi$	iff	for all v such that $w \sim_i v$ we have $\mathcal{M}, v \models \varphi$

Given a model $\mathcal{M} = (W, (R_i)_{i\in\mathrm{N}}, \mathsf{V})$ and a formula $\varphi \in \mathcal{L}_{\mathrm{EL}}$ we write $||\varphi||^{\mathcal{M}} := \{w \in W|\mathcal{M}, w \models \varphi\}$. Whenever $\mathcal{M}$ is clear from context, we simply write $||\varphi||$ for $||\varphi||^{\mathcal{M}}$.

There is a close relation between the assumptions that are made about models and the validities of the epistemic language. In particular, it is well-known (for a formal discussion of the concept of *correspondence* used below, see e.g., Blackburn et al. 2001) that

- T-axiom: $K_i p \rightarrow p$ corresponds to reflexivity (veridicality: what is known is true),
- 4-Axiom: $K_i p \rightarrow K_i K_i p$, corresponds to transitivity (positive introspection),
- 5-Axiom: $\neg K_i p \rightarrow K_i\neg K_i p$, corresponds to Euclideanity (negative introspection) (Table 2).

Theorem 2.5 (see e.g., Fagin et al. 1995) *EL is strongly complete with respect to the class of epistemic models.*

Table 2 K axiom system

PL	$\vdash \varphi$ if φ is a substitution instance of a tautology of propositional logic
For $i \in \mathrm{N}$,	
Nec	if $\vdash \varphi$, then $\vdash \Box_i\varphi$
K	$\vdash \Box_i (\varphi \rightarrow \psi) \rightarrow (\Box_i\varphi \rightarrow \Box_i\psi)$
MP	if $\vdash \varphi \rightarrow \psi$ and $\vdash \varphi$, then $\vdash \psi$

Again, the reader can check Blackburn et al. (2001) for a formal definition of *strong completeness*. On arbitrary Kripke models, the notation $\Box_i\varphi$ is preferred to K_i, and validities are axiomatised by K, which is defined in Table 2.

Theorem 2.6 (see e.g., Blackburn et al. 2001) K *is strongly complete with respect to the class of Kripke models.*

As mentioned, in this work we will only work with the semantic structures of epistemic modal logic and not with the logics themselves. We refer the reader interested in epistemic modal logic to the literature (Van Ditmarsch et al. 2007; Van Benthem 2010, 2011).

Intuitively, with an epistemic interpretation, an accessibility relation R_i in a Kripke model encodes i's uncertainty: if wR_iv, then if the actual world was w then i would consider it possible that the actual world is v.

Notation 2.7

- We write $\mathcal{K}_i[w] := \{v \in W | wR_iv\}$ to denote i's knowledge set at w.
- We refer to $\{\mathcal{K}_i[w] | w \in W\}$ as the information partition of i.
- For epistemic models for one agent, we sometimes also write $[w]$ to denote the equivalence class of w under the relation $\sim$, i.e., $[w] = \{w' \in W \mid w \sim w'\}$.
- For any non-empty group of agents $G \subseteq \mathbb{N}$ we write $R_G = \cup_{i\in G}R_i$, R^*_G for the reflexive transitive closure of $\bigcup_{i\in G} R_i$ and $R^*_G[w] := \{v \in W | wR^*_Gv\}$.

Importantly, if $||\varphi|| \subseteq R^*_G[w]$ then for any $n \in \omega$ and sequence $i_0, \ldots, i_{n-1}$ with range G, $K_{i_0} \ldots K_{i_{n-1}}\varphi$ holds at w. Since, the conjunction of all finite sequences is not finitary, a common knowledge operator of the form $C_G\varphi$ with precisely the previously given semantics is considered. An immediate consequence of the definition, is that a logic with common knowledge will lack compactness. It is however possible to finitely axiomatise the logic, proving a weak completeness result. The reader can consult (Fagin et al. 1995, Chap. 3) for a statement and proof of that result.

An important question is to what extent modal languages can distinguish between different structures. Similarity concepts and conditions at which two structures look alike for a given language are discussed in Section 2.3. These concepts will play an important role throughout the paper. Before we turn to the issues, we first discuss how simple concepts of updates are usually incorporated in the preceding epistemic framework.

2.2 Epistemic update

Epistemic models are static—they represent the informational state of an agent in temporal isolation. We will now make the setting more dynamic by assuming that agents observe some incoming data and are allowed to revise their informational states. For a richer development about modal logic for epistemic update and examples, the reader can consult (Van Ditmarsch et al. 2007).

We start with a definition that formalises the notion of update with a formula φ.

Definition 2.8 The update of an epistemic model $\mathcal{M} = (W, (\sim_i)_{i\in\mathrm{N}}, V)$ with a formula φ, restricts $\mathcal{M}$ to those worlds that satisfy φ, formally $\mathcal{M}|\varphi = \mathcal{M}' := (W', (\sim'_i)_{i\in\mathrm{N}}, V')$, where

(1) $W' = \{w \in W \mid \mathcal{M}, w \models \varphi\}$;
(2) for each $i \in \mathrm{N}$, $\sim'_i = \sim_i \upharpoonright W'$;
(3) $V' = V \upharpoonright W'$.

This is a very restrictive type of update, intuitively corresponding to public announcements. For more general types of epistemic updates, the reader can consult (Baltag et al. 1998; Van Ditmarsch et al. 2007). Basic epistemic logic, as defined above, can be extended to account for this type of update with a specific 'action' expression of *public announcement*, written as $!\varphi$.

Definition 2.9 (*Syntax of* $\mathcal{L}_{\mathrm{PAL}}$) The syntax of epistemic language $\mathcal{L}_{\mathrm{PAL}}$ is defined as follows:

$$\varphi := p|\neg\varphi|\varphi \vee \varphi|K_i\varphi|[A]\varphi$$
$$A := !\varphi$$

where $p \in \mathrm{PROP}$, $i \in \mathrm{N}$.

Definition 2.10 (*Semantics of* $\mathcal{L}_{\mathrm{PAL}}$) For the epistemic fragment $\mathcal{L}_{\mathrm{EL}}$ the interpretation is given in Definition 2.4. The remaining clause of $\mathcal{L}_{\mathrm{PAL}}$ is as follows.

$$\mathcal{M}, w \models [!\varphi]\psi \text{ iff if } \mathcal{M}, w \models \varphi \text{ then } \mathcal{M}|\varphi, w \models \psi$$

An axiomatisation of public announcement logic (PAL) of $\mathcal{L}_{\mathrm{PAL}}$ can be composed of the previously given axioms of epistemic logic enriched with the following reduction axioms:

$$1 \vdash [!\varphi]p \leftrightarrow (\varphi \rightarrow p),\ \text{for } p \in \mathrm{PROP}$$
$$2 \vdash [!\varphi]\neg\psi \leftrightarrow (\varphi \rightarrow \neg[!\varphi]\psi)$$
$$3 \vdash [!\varphi](\psi \vee \xi) \leftrightarrow ([!\varphi]\psi \vee [!\varphi]\xi)$$
$$4 \vdash [!\varphi]K_i\psi \leftrightarrow (\varphi \rightarrow K_i[!\varphi]\psi)$$

Theorem 2.11 (Soundness of **PAL**, Plaza 1989; Gerbrandy 1999; Baltag and Moss 2004) *PAL is sound with respect to the class of epistemic models.*

The soundness of **PAL** guarantees that a complete compositional analysis can be carried out into the epistemic language. Every formula of the public announcement language is thus equivalent to a formula of the epistemic language, whose validities are already decided by the axiom system **EL**.

Corollary 2.12 **PAL** *is strongly complete with respect to the class of epistemic models.*

Obviously, the incoming information that triggers update need not be propositional, not even purely linguistic. It can be any *event* that itself has an epistemic structure. Public announcement can have the effect that some situations that were considered possible before can be eliminated. This motivates why we would like to have a way to talk about parts of a model, e.g., only those states in which some proposition is true. For this, we need the notion of a *submodel*.

Definition 2.13 (*Submodel*) We say that $\mathcal{M}' = (W', (R_i)'_{i\in\mathrm{N}}, \mathsf{V}')$ is a submodel of $\mathcal{M} = (W, (R_i)_{i\in\mathrm{N}}, \mathsf{V})$ iff $W' \subseteq W$, $\forall i \in \mathrm{N}$, $R'_i = R_i \cap (W' \times W')$, $\forall p \in \mathrm{PROP}$, $\mathsf{V}'(p) = \mathsf{V}(p) \cap W'$.

Definition 2.14 (*Generated submodel*) We say that $\mathcal{M}' = (W', (R_i)'_{i\in\mathrm{N}}, \mathsf{V}')$ is a generated submodel of $\mathcal{M} = (W, (R_i)_{i\in\mathrm{N}}, \mathsf{V})$ iff $W' \subseteq W$ and $\forall i \in \mathrm{N}$, $R'_i = R_i \cap (W' \times W')$, $\forall p \in \mathrm{PROP}$, $\mathsf{V}'(p) = \mathsf{V}(p) \cap W'$ and if $w \in W'$ and wR_iv then $v \in W'$. The submodel of $\mathcal{M}$ generated by $X \subseteq W$ is the smallest generated submodel $\mathcal{M}'$ of $\mathcal{M}$ with $X \subseteq Dom(\mathcal{M}')$.

The notion of horizon generalises that of an information set:

Definition 2.15 (*Horizon*) The *horizon* of i at $(\mathcal{M}, w)$ (notation: $(\mathcal{M}, w)^i$) is the submodel generated by $\mathcal{K}_i[w]$.

The domain of $(\mathcal{M}, w)^i$ thus contains all the states that can be reached from w by first doing one step along the relation of agent i and then doing any number of steps along the relations of any agents.

Fact 2.16 In any Kripke structures, for all $i, j \in \mathrm{N}$, if $\sim_i$ is reflexive, then $(\mathcal{M}, w)^i \subseteq (\mathcal{M}, w)^j$.

Proof Assume that there is a path $w \sim_j v_0 \sim_{k_1} v_1 \cdots \sim k_n v_n$ with $w, v_0, \ldots, v_n \in W$ and $\{j, k_1, \ldots k_n\}$, then by reflexivity of $\sim_i$, there is also a path $w \sim_i w \sim_j v_0 \sim_{k_1} v_1 \cdots \sim k_n v_n$ with $w, v_0, \ldots, v_n \in W$ and $\{j, k_1, \ldots k_n\}$. □

Horizons of agents with reflexive indistinguishability relations are thus identical.

As mentioned, this paper will not use directly syntactic notions. In terms of intuition, the important definition is that of knowledge K_i: at w, agent i knows that φ iff it is the case that φ is true in all states that i considers possible at w. In equivalent semantic terms, we could substitute a finite collection of subsets of W (call them *events*) to our propositional letters and their valuation, and declare that, at w, i knows that some event $E \subseteq W$ is the case iff $\mathcal{K}_i[w] \subseteq E$. Similarly, that E is common knowledge in a group G at w iff $R^*_G[w] \subseteq E$. In the next section, we discuss closely how semantic similarity concepts closely correspond to syntactic definability.

In the technical parts of this paper, we use complexity results from graph theory (see e.g., Garey and Johnson 1990). Here, we use the connection between Kripke models and graphs: graphs are essentially Kripke models without valuations, i.e., *frames* (Blackburn et al. 2001). For graphs, the notion of *induced subgraph* is just like that of submodel (Definition 2.13) without the condition for the valuations. The notion of *subgraph* is weaker than that of an induced subgraph as it allows that $R'_i \subset R_i \cap W' \times W'$.

2.3 Epistemic situations that look alike

In various parts of our investigation we will need a reasonable notion of two models, that is two epistemic situations, being similar. We make use of the notions of simulation, simulation equivalence and bisimulation. We motivate these notions in terms of well-known preservation and invariance results. For a detailed introduction to these concepts, the reader can consult (Blackburn et al. 2001, Chap. 2), in particular Section 2.2 of that book.

Definition 2.17 (*Simulation*) We say that a pointed Kripke model $(\mathcal{M}, w)$, with $\mathcal{M} = (W, (R_i)_{i\in\mathtt{N}}, \mathsf{V})$ and $w \in W$, is simulated by another pointed Kripke model $(\mathcal{M}', w')$ (denoted by $(\mathcal{M}, w) \sqsubseteq (\mathcal{M}', w')$) with $\mathcal{M}' = (W', (R'_i)_{i\in\mathtt{N}}, \mathsf{V}')$ and $w' \in W'$ if the following holds.

There exists a binary relation $Z \subseteq W \times W'$ such that wZw' and for any pair of states $(x, x') \in W \times W'$, whenever xZx' then for all $i \in \mathtt{N}$:

(1) x, x' verify the same proposition letters.
(2) if xR_iz in $\mathcal{M}$ then there exists $z' \in W'$ with $x'R'_iz'$ and zZz'.

We say that $\mathcal{M} = (W, (R_i)_{i\in\mathtt{N}}, \mathsf{V})$ is simulated by $\mathcal{M}' = (W', (R'_i)_{i\in\mathtt{N}}, \mathsf{V}')$ (denoted by $\mathcal{M} \sqsubseteq \mathcal{M}'$) if there are $w \in W$ and $w' \in W'$ such that $(\mathcal{M}, w) \sqsubseteq (\mathcal{M}', w')$. We say that a simulation $Z \subseteq W \times W'$ is *total* if for every $w \in W$, there is some $w' \in W'$ such that wZw', and for every $w' \in W'$, there is some $w \in W$ such that wZw'. If $\mathcal{M}$ is simulated by $\mathcal{M}'$ by means of a total simulation, we say $\mathcal{M} \sqsubseteq_{total} \mathcal{M}'$. Moreover, we say that $\mathcal{M} = (W, (R_i)_{i\in\mathtt{N}}, \mathsf{V})$ and $\mathcal{M}' = (W', (R'_i)_{i\in\mathtt{N}}, \mathsf{V}')$ are simulation equivalent if $\mathcal{M}$ simulates $\mathcal{M}'$ and $\mathcal{M}'$ simulates $\mathcal{M}$.

Theorem 2.18 (De Rijke 1993) *A formula of epistemic logic is equivalent to a positive existential formulas of epistemic logic iff its truth in pointed models is preserved under simulations.*

For details and a proof of this result, the reader should consult (Blackburn et al. 2001, Section 2.7).

Example 2.19 In order to get an intuitive idea of simulation, consider two pointed Kripke models $(\mathcal{M}, w)$, $(\mathcal{M}', w')$ both with one agent (Bob), and the accessibility relations representing the uncertainty of Bob. Then

$$(\mathcal{M}, w) \sqsubseteq (\mathcal{M}', w')$$

means that in $(\mathcal{M}, w)$ Bob has more refined knowledge than in $(\mathcal{M}', w)$, i.e., in $(\mathcal{M}', w')$ Bob has more uncertainty.

The following notion is stronger than simulation equivalence.

Definition 2.20 (*Bisimulation*) A local bisimulation between two pointed Kripke models with set of agents $\mathtt{N}$, $(\mathcal{M}, w)$ with $\mathcal{M} = (W, (R_i)_{i\in\mathtt{N}}, \mathsf{V})$ and $(\mathcal{M}', w')$ with $\mathcal{M}' = (W', (R'_i)_{i\in\mathtt{N}}, \mathsf{V}')$ is a binary relation $Z \subseteq W \times W'$ such that wZw' and also for any pair of worlds $(x, x') \in W \times W'$ whenever xZx' then for all $i \in \mathtt{N}$:

(1) x, x' verify the same proposition letters.
(2) if xR_iu in $\mathcal{M}$ then there exists $u' \in W'$ with $x'R'_iu'$ and uZu'.
(3) if $x'R'_iu'$ in $\mathcal{M}'$ then there exists $u \in W$ with xR_iu and uZu'.

We say that $\mathcal{M} = (W, (R_i)_{i\in \mathrm{N}}, \mathsf{V})$ and $\mathcal{M}' = (W', (R'_i)_{i\in \mathrm{N}}, \mathsf{V}')$ are bisimilar ($\mathcal{M}\underline{\leftrightarrow}\mathcal{M}'$) if there are $w \in W$ and $w' \in W'$ such that $(\mathcal{M}, w)\underline{\leftrightarrow}(\mathcal{M}', w')$. A bisimulation $Z \subseteq Dom(\mathcal{M}) \times Dom(\mathcal{M}')$ is *total* if for every $w \in Dom(\mathcal{M})$, there is some $w' \in Dom(\mathcal{M}')$ such that wZw', and for every $w' \in Dom(\mathcal{M}')$, there is some $w \in Dom(\mathcal{M})$ such that wZw'. Then we write $\mathcal{M}\underline{\leftrightarrow}_{total}\mathcal{M}'$.

Theorem 2.21 (Van Benthem 1983) *A formula of first-order logic is equivalent to the translation of a formula of epistemic logic iff it is invariant under bisimulations.*

It follows, for example, that adding a disjoint part to a model is not affecting the knowledge of the agents, even though it is impacting the information partition. Intuitively, these disjoint parts correspond to counterfactual knowledge: what an agent would know in some situation that is not obtaining. It follows from Van Benthem's theorem that we need a stronger language than epistemic logic to describe such counterfactual knowledge. For example we could strengthen our notion of bisimulation by requiring it to be *total*. In syntactic terms, this corresponds to adding a universal operator A with which we can describe counterfactual knowledge. Other concepts of similarities have nice syntactic characterisations. In the other direction, certain similarity concepts such as *isomorphism* are arguably too strong in syntactic terms. Still, isomorphism can be thought as a maximal requirement: if two epistemic structures are isomorphic, then they describe the same informational situation, no matter what epistemic language we take as standard to describe that situation.

2.4 Tractability

Some problems, although computable, nevertheless require too much time or memory to be feasibly solved by a realistic computational device. Computational complexity theory investigates the resources (time, memory, etc.) required for the execution of algorithms and the inherent difficulty of computational problems (Papadimitriou 1993). In particular, we want to identify efficiently solvable problems and draw a line between tractability and intractability. In general, the most important distinction is that between problems which can be computed in polynomial time with respect to their size, and those which are believed to have only exponential time algorithmic solutions (Edmonds 1965). This is exactly the tractability border we investigate in the paper. The class of problems of the first type is called PTIME (P for short); one can demonstrate that a problem belongs to this class if one can show that it can be computed by a deterministic Turing machine in polynomial time with respect to the size (length) of the input. Problems belonging to the second class are referred to as NP-hard. They are at least as difficult as problems belonging to the NPTIME (NP) class; this is the class of problems which can be computed by non-deterministic Turing machines in polynomial time. NP-complete problems are NP-hard problems belonging to NPTIME, hence they are intuitively the most difficult problems among the NPTIME problems.

Let us start proper complexity considerations with the notation used for comparing the growth rates of functions.

Definition 2.22 Let $f, g : \omega \longrightarrow \omega$ be any functions. We say that $f = O(g)$ if there exists a constant $c > 0$ such that $f(n) \leq cg(n)$ for almost all (i.e., all but finitely many) n.

Let $f : \omega \longrightarrow \omega$ be a natural number function. TIME(f) is the class of languages (problems) which can be recognised by a deterministic Turing machine in time bounded by f with respect to the length of the input. In other words, $L \in$ TIME(f) if there exists a deterministic Turing machine such that for every $x \in L$, the computation path of M on x is shorter than $f(n)$, where n is the length of x. TIME(f) is called a *deterministic computational complexity class*. A *non-deterministic complexity class*, NTIME(f), is the class of languages L for which there exists a non-deterministic Turing machine M such that for every $x \in L$ all branches in the computation tree of M on x are bounded by $f(n)$ and moreover M decides L. One way of thinking about a non-deterministic Turing machine bounded by f is that it first guesses the right answer and then deterministically checks if the guess is correct, in a time bounded by f.

SPACE(f) is the class of languages which can be recognised by a deterministic machine using at most $f(n)$ cells of the working-tape. NSPACE(f) is defined analogously. Below we define some well-known complexity classes, i.e., the sets of languages of related complexity. In other words, we can say that a complexity class is the set of problems that can be solved by a Turing machine using $O(f(n))$ of time or space resource, where n is the size of the input.

Definition 2.23
- LOGSPACE $= \bigcup_{k\in\omega}$ SPACE($k \log n$)
- PTIME $= \bigcup_{k\in\omega}$ TIME(n^k)
- NP $= \bigcup_{k\in\omega}$ NTIME(n^k)
- PSPACE $= \bigcup_{k\in\omega}$ SPACE(n^k)

If $L \in$ NP, then we say that L *is decidable* (*computable*, *solvable*) *in non-deterministic polynomial time* and likewise for other complexity classes.

PTIME $\subseteq$ NP but the question whether PTIME (P for short) is strictly contained in NPTIME (NP) is the famous Millennium Problem—one of the most fundamental problems in theoretical computer science and mathematics. Also: LOGSPACE $\subseteq$ PTIME and NP $\subseteq$ PSPACE but it is not known whether the hierarchy is strict.

The intuition that some problems are more difficult than others is formalised in complexity theory by the notion of a *reduction*. We will use polynomial time and logarithmic space many-one (Karp 1972) reductions.

Definition 2.24 We say that a function $f : A \longrightarrow A$ is a *polynomial* time (*logarithmic* space) *computable function* iff there exists a deterministic Turing machine computing $f(w)$ for every $w \in A$ in a polynomial time (logarithmic space).

Definition 2.25 A *problem* $L \subseteq \Gamma^*$ *is polynomial* time (*logarithmic* space) *reducible to a problem* $L' \subseteq \Gamma^*$ if there is a polynomial time (logarithmic space) computable function $f : \Gamma^* \longrightarrow \Gamma^*$ from strings to strings, such that

$$w \in L \iff f(w) \in L'.$$

We will call such function f a *polynomial* time respectively (*logarithmic* space) *reduction* of L to L'.

Definition 2.26 A language L is complete for a complexity class $\mathcal{C}$ if $L \in \mathcal{C}$ and every language in $\mathcal{C}$ is reducible to L.

Intuitively, if L is complete for a complexity class $\mathcal{C}$ then it is among the hardest problems in this class. For example, PTIME-complete problems (under a weaker notion of NC or log-space reductions) are the hardest problems among PTIME problems, and as a result are believed to be difficult to easily compute on parallel machines and solved in limited space. The notable property of NP-complete problems is that, even though any given solution for such problems can be quickly verified, there is no efficient way to find the solution in the first place. Interestingly, there are known problems that are neither NP-complete nor tractable, e.g., the graph isomorphism problem asking to decide whether two given finite graphs are isomorphic (Garey and Johnson 1990). To study further this interesting territory researchers defined a new class of problems, GI, that contains problems polynomially reducible to the graph isomorphism problem. If the graph isomorphism problem turns out computable in PTIME, then GI would equal PTIME. If the graph isomorphism problem is proved to be NP-complete then the polynomial hierarchy must collapse. There are many problems known to belong to GI and the class has recently found many applications in complexity analysis (Köbler et al. 1993).

3 Complexity of comparing and manipulating information

In this section, we give the results we obtained when studying the complexity of different epistemic reasoning tasks in the semantic structures of modal logics. The tasks we investigate deal with three different aspects.

Information similarity (Section 3.1).
Are the information structures of two agents similar?
Information symmetry (Section 3.2).
Do two agents have the same (similar) information about each other?
Information manipulation (Section 3.3).
Can we manipulate the information of one agent so that as a result she knows at least as much as another agent?

3.1 Information similarity

The first natural question we would like to address is whether an agent in a given situation has similar information to the one possessed by some other agent (in a possibly different situation). One very strict way to understand such similarity is through the use of *isomorphism*. As we wrote previously, we admit that from the epistemic logic perspective the notion is arguably too strong but as we are mainly interested in a systematic semantic study of similarity notions we look at it for the sake of completeness as a strongest notion of similarity.

For the general problem of checking whether two Kripke models are isomorphic, we can give the following complexity bounds, in the sense that the problem is equivalent to the graph isomorphism problem.

Problem 3.1 (*Kripke model isomorphism*)
Input Pointed Kripke models $(\mathcal{M}_1, w_1), (\mathcal{M}_2, w_2)$.

Question Are $(\mathcal{M}_1, w_1)$ and $(\mathcal{M}_2, w_2)$ isomorphic, i.e., is it the case that $(\mathcal{M}_1, w_1) \cong (\mathcal{M}_2, w_2)$?

Fact 3.2 Kripke model isomorphism is GI-complete.

Proof Kripke model isomorphism is equivalent to the variation of graph isomorphism with labelled vertices, which is polynomially equivalent to graph isomorphism (see e.g., Hoffmann 1982), and thus GI-complete. □

However, isomorphism is arguably a too restrictive notion of similarity. Bisimilarity is a weaker concept of similarity. As we take a modal logic perspective in this work and want to analyse the complexity of epistemic tasks on the semantic structures of epistemic modal logic, bisimilarity is a very natural choice of similarity.

Here the question arises as to whether working with S5 models—a common assumption in the epistemic logic and interactive epistemology literature—rather than arbitrary Kripke structures has an influence on the complexity of the task.

Problem 3.3 (*Epistemic model bisimilarity*)
Input Two pointed multi-agent epistemic S5 models $(\mathcal{M}_1, w_1), (\mathcal{M}_2, w_2)$.

Question Are the two models bisimilar, i.e., $(\mathcal{M}_1, w_1) \underline{\leftrightarrow} (\mathcal{M}_2, w_2)$?

To illustrate the difference between epistemic model bisimilarity and epistemic model isomorphism, consider a situation in which Alice considers several worlds possible with the same valuation. Bisimilarity does not distinguish between situations in which the number of those worlds differs, whereas isomorphism does.

Balcázar et al. (1992) have shown that deciding bisimilarity is P-complete for finite labelled transition systems. As epistemic models are just a special kind of labelled transition systems, we can use an algorithm that solves bisimilarity for labelled transition systems also for epistemic models. It follows that epistemic model bisimilarity is also in P.

Fact 3.4 Multi-agent epistemic S5 model bisimilarity can be computed in polynomial time with respect to the size of the input $(|\mathcal{M}_1| + |\mathcal{M}_2|)$.

Thus, multi-agent epistemic S5 model bisimilarity is in P. Now, of course the question arises if it is also P-hard.

Proposition 3.5 *Multi-agent epistemic S5 model bisimilarity is* P*-complete.*

Proof P membership is immediate from Fact 3.4. For P-hardness, we adapt the hardness proof of Balcázar et al. (1992). In the reduction from monotone alternating circuits, the labelled transition systems that are constructed are irreflexive. We transform them into corresponding S5 models for two agents using the method used in Halpern and Moses (1992) and replace every edge $w \rightarrow v$ by $w \sim_1 w' \sim_2 v$, keeping the valuation of w and v the same as before and making the valuation of w' the same as that of w. Additionally, reflexive loops have to be added. Bisimilarity of two finite irreflexive structures is invariant under this transformation. Moreover, note that for the replacement of the edges, we only need constant memory space. P-hardness follows. □

To summarise, while deciding Kripke model isomorphism lies on the tractability border, deciding whether two Kripke models are bisimilar is among the hardest problems that are known to be in P. For S5 epistemic models with at least two agents, we get the same results.

3.2 Information symmetry: knowing what others know

The preceding notions of similarity are very strong as they are about the similarity of whole information structures. In the context of analysing epistemic interactions between agents, weaker notions of similarity are of interest, as often already the similarity of some relevant parts of information are sufficient for drawing some conclusions. In general, the information that agents have about each other's information state plays a crucial role. We will now analyse the problem of deciding whether two agents' views about the interactive epistemic structure, and in particular about the knowledge of other agents, are equivalent. A first reading is simply to fix some fact $E \subseteq W$ and ask whether E is common knowledge in a group G. Clearly this problem is tractable.

Fact 3.6 Given a pointed multi-agent epistemic model $(\mathcal{M}, w)$, some $E \subseteq Dom(\mathcal{M})$ and a subset of agents $G \subseteq \mathtt{N}$, deciding whether E is common knowledge in the group G at w can be done in polynomial time.

Proof To decide if E is common knowledge among the agents in G, we can use a reachability algorithm to check if any of the states which are not in E (i.e., any state in $Dom(\mathcal{M}) \setminus E$) is reachable from w by a *path* along the relation $\cup_{j \in G} \sim_j$. If this is the case, the answer is *no*, otherwise the answer is *yes*, as then $\sim_G^* [w] \subseteq E$. □

If some fact is common knowledge between two agents, the information of the two agents about this fact can be seen as symmetric, in the sense that both agents have the same information about the fact and about the information they have about the fact. More generally, instead of fixing some specific fact of interest, an interesting question is whether an epistemic situation is symmetric with respect to two given agents, say Ann and Bob. In other words, is the interactive informational structure from Ann's perspective similar to how it is from Bob's perspective? We first introduce some notation that we will use for representing situations in which the information of two agents is exchanged, in the sense that each of the agents gets exactly the information that the other one had before.

Definition 3.7 For a Kripke model $\mathcal{M} = (W, (R_i)_{i\in\mathrm{N}}, \mathsf{V})$, with $j, k \in \mathrm{N}$, we write $\mathcal{M}[j/k]$ to be the model $(W, (R'_i)_{i\in\mathrm{N}}, \mathsf{V})$ for $R'_i = R_i$ for $i \notin \{j, k\}$, $R'_j = R_k$ and $R'_k = R_j$.

That is, in $\mathcal{M}[j/k]$ agent j gets the accessibility relation of k in $\mathcal{M}$ and vice versa.

The intuition is that in many multi-agent scenarios it can be interesting to determine if the situation is symmetric with respect to two given agents in the sense that those two agents have similar information about facts, other agents and also about each other. As a typical such situation consider a two-player card game. Here, it can be crucial for the strategic abilities of the players whether they both know equally little about each other's cards and whether they know the same about each other's information. From a modelling perspective, determining if the information of two agents is interchangeable can also be crucial if we want to find a succinct representation of the situation (cf. Chap. 7 of Wang 2010), as in some situations only explicitly representing one of the agents might be sufficient (cf. Gierasimczuk and Szymanik 2011a,b).

To formalise this property of information symmetry, we introduce the notion of *flipped* bisimulation for a pair of agents. The main difference with respect to a standard bisimulation is that for each step along the accessibility relation for one agent in one model, there has to be a corresponding step along the relation of the *other* agent in the other model.

Definition 3.8 We say that two pointed multi-agent epistemic models $(\mathcal{M}, w)$ and $(\mathcal{M}', w')$ (with set of agents N) are flipped bisimilar for agents $i, j \in \mathrm{N}$—in notation $(\mathcal{M}, w)\underline{\leftrightarrow}_f{}^{(ij)}(\mathcal{M}', w')$—iff $(\mathcal{M}, w)\underline{\leftrightarrow}(\mathcal{M}'[i/j], w')$.

So, two models are flipped bisimilar for two agents if after swapping the accessibility relations of the two agents in one of the models, the resulting model is bisimilar to the other model.

To help to get an intuition of this notion, we list two facts about flipped bisimilarity that follow directly from its definition.

Fact 3.9 For any pointed multi-agent Kripke models $(\mathcal{M}, w)$, $(\mathcal{M}', w')$ with set of agents N and agents $i, j \in \mathrm{N}$ the following hold.

- $(\mathcal{M}, w)\underline{\leftrightarrow}_f^{(ii)}(\mathcal{M}', w')$ iff $(\mathcal{M}, w)\underline{\leftrightarrow}(\mathcal{M}', w')$,
- $(\mathcal{M}, w)\underline{\leftrightarrow}_f^{(ij)}(\mathcal{M}', w')$ iff $(\mathcal{M}, w)\underline{\leftrightarrow}_f^{(ji)}(\mathcal{M}', w')$.

Moreover, in general we can have that $(\mathcal{M}, w)\underline{\not\leftrightarrow}_f^{(ij)}(\mathcal{M}, w)$, i.e., the relation of flipped bisimilarity is not reflexive, and thus not an equivalence relation.

Thus, flipped bisimilarity for the same agent is equivalent to regular bisimilarity. While the relation of flipped bisimilarity for a pair of agents is not reflexive, flipped bisimilarity is indeed symmetric with respect to the flipping of the agents.

Note that in general $(\mathcal{M}[i/j])[j/k]$ and $(\mathcal{M}[j/k])[i/j]$ are not bisimilar, neither are they in general flipped bisimilar for any pair of agents. It follows that we can have

$$(\mathcal{M}, w)\underline{\leftrightarrow}_f^{(ij)}(\mathcal{M}', w')\underline{\leftrightarrow}_f^{(jk)}(\mathcal{M}'', w'')\ (\mathcal{M}, w)\underline{\leftrightarrow}_f^{(jk)}(\mathcal{M}^*, w^*)\underline{\leftrightarrow}_f^{(ij)}(\mathcal{M}^{**}, w^{**})$$

while not having for any pair $n, n' \in \{i, j, k\}$, $(\mathcal{M}'', w'')\underline{\leftrightarrow}_f^{n,n'}(\mathcal{M}^{**}, w^{**})$. Similarly, it is not the case that for all models $(\mathcal{M}, w)$, $(\mathcal{M}', w')$, $(\mathcal{M}'', w'')$ with set of agents $\mathtt{N}$ and agents $i, j, k \in \mathtt{N}$, if $(\mathcal{M}, w)\underline{\leftrightarrow}_f^{(ij)}(\mathcal{M}', w')\underline{\leftrightarrow}_f^{(jk)}(\mathcal{M}'', w'')$, then $(\mathcal{M}, w)\underline{\leftrightarrow}_f^{(ik)}(\mathcal{M}'', w'')$. This is because performing two consecutive swaps of agents is in general not equivalent to performing one swap of agents.

In the context of epistemic multi-agent models, the following question arises: How does flipped bisimilarity relate to knowledge of the individual agents and common knowledge?

The following is immediate: if on a whole model it holds that everything that two individual agents know is common knowledge among them, then every state is flipped bisimilar (for these two agents) to itself. The intuition here is that if everything that the two individuals know is commonly known among them, then the two agents have exactly the same information and can thus be swapped.

Fact 3.10 If for a multi-agent S5 model $\mathcal{M} = (W, (\sim_i)_{i\in\mathtt{N}}, \mathsf{V})$, it holds that $\sim^*_{\{i,j\}} \subseteq \sim_i \cap \sim_j$ for some $i, j \in \mathtt{N}$, then for all $w \in W$, $(\mathcal{M}, w)\underline{\leftrightarrow}_f{}^{(ij)}(\mathcal{M}, w)$.

Does the other direction hold? Locally, even on S5 models, flipped self-bisimulation is much weaker than the property of individual knowledge being common knowledge: flipped self-bisimulation does not even imply that (shared) knowledge of facts is common knowledge:

Fact 3.11 There exists a multi-agent S5 model $\mathcal{M} = (W, (\sim_i)_{i\in\mathtt{N}}, \mathsf{V})$, such that for some $i, j \in \mathtt{N}$ we have that for some $w \in W$ it holds that $(\mathcal{M}, w)\underline{\leftrightarrow}_f^{(ij)}(\mathcal{M}, w)$, and for some $p \in \textsc{prop}$ we have that $\mathcal{M}, w \models K_i p$ and $\mathcal{M}, w \models K_j p$ but p is not common knowledge among i and j at w.

Proof Consider the model $\mathcal{M} = (W, (\sim_i)_{i\in\mathtt{N}}, \mathsf{V})$, where

- $W = \{w_{-2}, w_{-1}, w_0, w_1, w_2\}$,
- $\mathtt{N} = \{\text{Ann}, \text{Bob}\}$,
- $\sim_{\text{Ann}}$ is the smallest equivalence relation on W containing $\{(-2, -1), (0, 1)\}$, and $\sim_{\text{Bob}}$ is the smallest equivalence relation on W containing $\{(-1, 0), (1, 2)\}$,
- $\mathsf{V}(p) = \{w_{-1}, w_0, w_1\}$.

The following figure represents $\mathcal{M}$. The dashed rectangles are the equivalence classes for Ann and the dotted rectangles those of Bob.

It is easy to check that at state w_0 both Ann and Bob know that p: $\mathcal{K}_{\text{Ann}}[w_0] = \{w_0, w_1\} \subseteq \mathsf{V}(p)$ and $\mathcal{K}_{\text{Bob}}[w_0] = \{w_{-1}, w_0\} \subseteq \mathsf{V}(p)$. But p is not common knowledge between Ann and Bob at w_0: we have that $w_0 \sim_{\text{Ann}} w_1 \sim_{\text{Bob}} w_2$ and $w_2 \notin \mathsf{V}(p)$. Now it remains to show that $(\mathcal{M}, w_0)$ is Ann, Bob-flipped bisimilar to itself. We can define a flipped bisimulation as follows $Z = \{(w_n, w_{-n}) | w_n \in W\}$, i.e., $Z = \{(w_{-2}, w_2), (w_{-1}, w_1), (w_0, w_0), (w_1, w_{-1}), (w_2, w_{-2})\}$. It is easy to check that Z is indeed a flipped bisimulation for Ann and Bob.

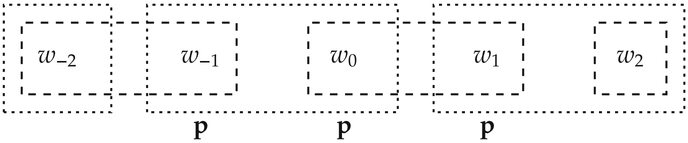

□

But required globally of every state, we do have the following converse: If for two agents we have flipped bisimilarity of every state to itself and the accessibility relations of the agents are transitive, then every fact that is known by at least one of the agents is immediately also common knowledge among the two agents.

Fact 3.12 For every Kripke model $\mathcal{M} = (W, (R_i)_{i \in \mathrm{N}}, \mathsf{V})$ with R_i and R_j being transitive for some $i,j \in \mathrm{N}$, it holds that for each $w \in W$ we have the following. Whenever the submodel $\mathcal{M}'$ of $\mathcal{M}$ generated by $\{w\}$ is such that for every state $w' \in Dom(\mathcal{M}')$ it holds that $(\mathcal{M}', w') \underline{\leftrightarrow}_f^{(ij)} (\mathcal{M}', w')$, then for any $p \in \mathrm{PROP}$, if at w at least one of the two agents i and j knows that p (i.e., $\mathsf{V}(p) \subseteq K_j[w]$ or $\mathsf{V}(p) \subseteq K_i[w]$), then p is common knowledge among i and j at w.

Proof Assume that for some model $\mathcal{M} = (W, (R_i)_{i \in \mathrm{N}}, \mathsf{V})$ it holds that R_i and R_j are transitive for some $i,j \in \mathrm{N}$. Now, assume that p is not common knowledge between i and j at w. It follows that we have a finite i,j-path leading to a state where p is false. Let $wRf(1)w_1Rf(2)\ldots Rf(n)w_n$ with $w_n \notin \mathsf{V}(p)$ and $f(k) \in \{i, j\}$ for all $k \leq n$ be a shortest such path. Then, by transitivity of R_i and R_j it has to be the case that for all k with $1 \leq k < n$, $f(k) \neq f(k+1)$. Without loss of generality assume that the path is of the form $wR_iw_1R_j \ldots R_iw_n$; the other cases are completely analogous. Now, as all the states in the path $wR_iw_1R_j \ldots R_iw_n$ are in $\mathcal{M}'$, by assumption for each w_k in the path we have $(\mathcal{M}', w_k) \underline{\leftrightarrow}_f^{(ij)} (\mathcal{M}', w_k)$. Then, in particular $(\mathcal{M}', w)$ is flipped i,j-bisimilar to itself. Then there has to be a path $wR_jw_1^1R_iw_2^1 \ldots R_jw_n^1$ with $w_n^1 \notin V(p)$. Then, we can continue this argument, as also $(\mathcal{M}', w_1^1)$ has to be flipped i, j-bisimilar to itself. Thus, there has to be some path $w_1^1R_jw_2^2R_j \ldots R_n^2w_n^2$. Then, by transitivity of R_j, $wR_jw_2^2$. Iterating this procedure, we will finally get that there is an R_j path from w to a state where p is false. Using the transitivity of R_j, we then conclude that $M, w \not\models K_jp$.

It remains to show that at w agent i does not know that p either. By assumption, wR_iw_1 and thus $(\mathcal{M}', w_1)$ has to be flipped i,j-bisimilar to itself. Thus, there has to be a path a $w_1R_iw_2'R_jw_3' \ldots R_jw_n'$ with $w_n' \notin V(p)$ Then, by transitivity, it follows from $wR_iw_1R_iw_2'$ that wR_iw_2'. Iterating this procedure, we get a state which is R_i-accessible from w where p is false. Hence, we conclude that at w neither i nor j knows that p. This concludes the proof. □

Let us recall the notion of an agent's horizon (Definition 2.15). It is the submodel generated by the information set of the agent: the *horizon* of i at (M, w) (notation: $(M, w)^i$) is the submodel generated by the set $\mathcal{K}_i[w]$.

We now analyse the complexity of deciding (flipped) bisimilarity of two agents' horizons at the same point in a model. We distinguish between S5 models and the class of all Kripke structures.

Proposition 3.13 *For horizon bisimilarity of multi-agent Kripke models we have the following complexity results*

(1) *For multi-agent S5 models* $(\mathcal{M}, w)$ *with set of agents* $\mathtt{N}$,
  (a) *deciding whether* $(\mathcal{M}, w)^i \underline{\leftrightarrow} (\mathcal{M}, w)^j$ *is trivial.*
  (b) *deciding whether* $(\mathcal{M}, w)^i \underline{\leftrightarrow}_f^{(ij)} (\mathcal{M}, w)^j$ *is in* P.
(2) *For multi-agent Kripke models* $(\mathcal{M}, w)$ *with set of agents* $\mathtt{N}$,
  (a) *deciding whether* $(\mathcal{M}, w)^i \underline{\leftrightarrow} (\mathcal{M}, w)^j$ *is* P*-complete.*
  (b) *deciding whether* $(\mathcal{M}, w)^i \underline{\leftrightarrow}_f^{(ij)} (\mathcal{M}, w)^j$ *is* P*-complete.*

Proof 1(a) follows from the fact that if the agents' accessibility relations are reflexive then the horizons of the agents are the same.

This is the case because $(\mathcal{M}, w)^i$ is the submodel generated by $\mathcal{K}_i[w]$, i.e., the submodel generated by the set of states that i considers possible at w. If at w, i considers w itself possible, the domain of this submodel will also contain the domain of the submodel generated by $\mathcal{K}_j[w]$. The argument for the other direction is analogous.

1(b) follows from the fact that deciding flipped bisimilarity of horizons in multi-agent S5 is polynomially equivalent to deciding (flipped) bisimilarity of multi-agent S5 models. Both decision problems of 2(a) and 2(b) are polynomially equivalent to deciding bisimilarity of multi-agent Kripke models because in general the horizons of two agents at a point in the model can be two completely disjoint submodels. □

Let us summarise the results we have on the complexity of deciding information symmetry. Both deciding whether a fact is commonly known and deciding horizon flipped bisimilarity in Kripke models are tractable, with the latter being among the hardest problems known to be tractable. Flipped bisimilarity of horizons remains P-hard even if we consider the horizons of two agents at the very same point in a model. For partition-based models, however, deciding bisimilarity of the horizons of two agents at the same point in a model is trivial, whereas for flipped bisimilarity, this is harder, but still tractable (in P).

The tasks we considered so far dealt with the comparison of agents' information states in given situations. Here, we were concerned with *static* aspects of agents' information. However, in many interactive situations *dynamic* aspects play a central role, as the information of agents can change while the agents interact. There are even interactive processes where information change can be the aim of the interaction itself, e.g., interactive deliberation processes. In such contexts the question arises as to whether it is possible to manipulate the information state of agents in a particular way.

3.3 Can we reshape an agent's mind into some desired informational state?

The problem that we investigate in this section is to decide whether new informational states (satisfying desired properties) can be achieved in certain ways. One immediate question is whether one can give some information to an agent (i.e., to restrict the agent's horizon) such that after the update the horizon is bisimilar to the horizon of some other agent. Concretely, we would like to know if there is any type of information that could reshape some agent's information in order to fit some desired new informational state or at least be similar to it. More precisely, we will consider information that *restricts* the horizon of an agent; we do not consider the process of changing an agent's information state by introducing more uncertainty. The processes we consider are related to those modelled by PAL, as introduced in Section 2.2. Deciding whether a formula $[!\varphi]\psi \in \mathcal{L}_{\text{PAL}}$ holds at a state in a model (i.e., model checking this formula) involves first checking if φ holds at the state and then relativising the original model to the set of states where φ holds and finally checking if ψ then holds at the current state. In order to put the complexity results of this section into perspective, note that for PAL it holds that given a pointed model and a formula, checking if the formula holds in the model can be done in time polynomial in the length of the formula and the size of the model (cf. Kooi and Van Benthem 2004 for polynomial model checking results for PAL with relativised common knowledge).

The model checking problem of PAL is about deciding whether getting a particular piece of information (i.e., the information that φ holds) has a certain effect (i.e., the effect of ψ being the case). In this section, we will investigate a more general problem which is about whether it is possible to restrict the model so that a certain effect is achieved. To be more precise, we consider the task of checking whether there is a *submodel* that has certain properties. This means that we determine if it is possible to purposely refine a model in a certain way. This question is in line with problems addressed by arbitrary PAL (APAL) and arbitrary event modal logic (Balbiani et al. 2008; Van Ditmarsch and French 2009). [1] Looking at the complexity results for such logics (see e.g., French and Van Ditmarsch 2008 for a proof of undecidability of SAT of APAL and Ågotnes et al. 2010 for PSPACE-completeness of the model-checking problem for APAL), we can already see that reasoning about the existence of information whose announcement has a certain effect seems to be hard. Our analysis will show whether this is also the case for concrete tasks about deciding whether a given model can be restricted so that it will have certain properties.

We start with the problem of checking whether there is a submodel of one model that is bisimilar to another one. On graphs, this is related to the problem of deciding if one graph contains a subgraph bisimilar to another graph. Note that in the problem referred to in the literature as "subgraph bisimulation" (Dovier and Piazza 2003), the subgraph can be any graph whose vertices are a subset of the vertices of the original graph, and the edges can be any subset of the edges of the original graph restricted to the subset of vertices. To be more specific, the problem investigated in Dovier and Piazza (2003) is the following:

> Given two graphs $G_1 = (V_1, E_1)$ and $G_2 = (V_2, E_2)$, is there a graph $G'_2 = (V'_2, E'_2)$ with $V'_2 \subseteq V_2$ and $E'_2 \subseteq E_2$ such that there is a total bisimulation between G'_2 and G_1?

[1] Note that in the current work, we focus on the semantic structures only and do not require that the submodel can be characterised by some formula in a certain epistemic modal language.

Since we want to investigate the complexity of reasoning about epistemic interaction using modal logic, we are interested in subgraphs that correspond to *relativisation* in modal logic: *induced* subgraphs. This leads us to an investigation of *induced subgraph bisimulation*.

Problem 3.14 (*Induced subgraph bisimulation*)
Input Two finite graphs $G_1 = (V_1, E_1)$, $G_2 = (V_2, E_2)$, $k \in \omega$.
Question Is there an induced subgraph of G_2 with $\geq k$ vertices that is totally bisimilar to G_1, i.e., is there a $V' \subseteq V_2$ with $|V'| \geq k$ *and* $(V', E_2 \cap (V' \times V')) \underline{\leftrightarrow}_{total} G_1$?

Even though the above problem looks very similar to the original subgraph bisimulation problem (whose NP-hardness is shown by reduction from Hamiltonian Path Garey and Johnson 1990), NP-hardness does not follow immediately.[2] Nevertheless, we show NP-hardness by reduction from independent set (Garey and Johnson 1990).

Proposition 3.15 *Induced subgraph bisimulation is* NP*-complete.*

Proof Showing that the problem is in NP is straightforward. Hardness is shown by reduction from independent set. First of all, let $I_k = (V_{I_k}, E_{I_k} = \emptyset)$ with $|V_{I_k}| = k$ denote a graph with k vertices and no edges. Given the input of independent set, i.e., a graph $G = (V, E)$ and some $k \in \omega$ we transform it into (I_k, G), k, as input for induced subgraph bisimulation.

Now, we claim that G has an independent set of size at least k iff there is some $V' \subseteq V$ with $|V'| \geq k$ and $(V', E \cap (V' \times V')) \underline{\leftrightarrow}_{total} I_k$.

From left to right, assume that there is some $S \subseteq V$ with $|S| = k$, and for all $v,v' \in S$, $(v,v') \notin E$. Now, any bijection between S and V_{I_k} is a total bisimulation between $G' = (S, E \cap (S \times S))$ and I_k, since $E \cap (S \times S) = \emptyset$ and $|S| = |V_{I_k}|$.

For the other direction, assume that there is some $V' \subseteq V$ with $|V'| = k$ such that for $G' = (V', E' = E \cap (V' \times V'))$ we have that $G' \underline{\leftrightarrow}_{total} I_k$. Thus, there is some total bisimulation Z between G' and I_k. Now, we claim that V' is an independent set of G of size k. Let $v,v' \in V'$. Suppose that $(v,v') \in E$. Since G' is an induced subgraph, we also have that $(v, v') \in E'$. Since Z is a total bisimulation, there is some $w \in I_k$ with $(v, w) \in Z$ and some w' with $(w, w') \in E_{I_k}$ and $(v', w') \in Z$. But this is a contradiction with $E_{I_k} = \emptyset$. Thus, V' is an independent set of size k of G. The reduction can clearly be computed in polynomial time. This concludes the proof. □

Now, an analogous result for Kripke models follows. Here, the problem is to decide whether it is possible to 'gently' restrict one model without letting its domain get smaller than k so that afterward it is bisimilar to another model. With an epistemic/doxastic interpretation of the accessibility relation, the intuitive interpretation is that we would like the new information to change the informational state of the agent as little as possible.

Problem 3.16 (*Submodel bisimulation for Kripke models*)
Input Kripke models $\mathcal{M}_1$, $\mathcal{M}_2$ with set of agents $\mathtt{N}$ and some $k \in \omega$.

[2] For induced subgraph bisimulation, a reduction from Hamiltonian Path seems to be more difficult, as does a direct reduction from the original subgraph bisimulation problem.

Question Is there a submodel $\mathcal{M}'_2$ of $\mathcal{M}_2$ with $|Dom(\mathcal{M}'_2)| \geq k$ such that $\mathcal{M}_1$ and $\mathcal{M}'_2$ are totally bisimilar i.e., $\mathcal{M}_1 \underline{\leftrightarrow}_{total} \mathcal{M}'_2$?

Corollary 3.17 *Submodel bisimulation for Kripke models is* NP-*complete.*

Proof Checking if a proposed model is indeed a submodel and has at least k states can be done in polynomial time. As also bisimilarity can be checked in polynomial time, membership of NP is immediate. NP-hardness follows from Proposition 3.15 as the problem of deciding induced subgraph bisimilarity can be reduced to submodel bisimilarity. □

We restrict ourselves to submodels and generated submodels as these are most natural for basic modal epistemic logics. For future work, it can also be interesting to consider more general model restrictions.

Since we are interested in the complexity of reasoning about the interaction of epistemic agents as it is modelled in (dynamic) epistemic logic, let us now see how the complexity of induced subgraph bisimulation changes when we make the assumption that models are partitional, i.e., that the relation is an equivalence relation, as it is frequently assumed in the AI or interactive epistemology literature. We will see that this assumption makes the problem significantly easier.

Proposition 3.18 *Let $G_1 = (V_1, E_1)$ and $G_2 = (V_2, E_2)$, be graphs with E_1 and E_2 reflexive. Induced subgraph bisimulation for G_1 and G_2 is in TIME(n).*

Proof In this proof, we will use the fact that $G_1 = (V_1, E_1) \underline{\leftrightarrow}_{total} G_2 = (V_2, E_2)$ if and only if it is the case that $V_1 = \emptyset$ iff $V_2 = \emptyset$. Let us prove this. From left to right, assume that $G_1 = (V_1, E_1) \underline{\leftrightarrow}_{total} G_2 = (V_2, E_2)$. Then since we have a total bisimulation, it must be the case that either $V_1 = V_2 = \emptyset$ or $V_1 \neq \emptyset \neq V_2$.

For the other direction, assume that $V_1 = \emptyset$ iff $V_2 = \emptyset$. Now, we show that in this case, $V_1 \times V_2$ is a total bisimulation between G_1 and G_2. If $V_1 = V_2 = \emptyset$, we are done. So, consider the case where $V_1 \neq \emptyset \neq V_2$. Let $(v_1, v_2) \in V_1 \times V_2$, and assume that $(v_1, v'_1) \in E_1$ for some $v'_1 \in V_1$. Since E_2 is reflexive, we know that there is some $v'_2 \in V_2$ such that $(v_2, v'_2) \in E_2$. Of course $(v'_1, v'_2) \in V_1 \times V_2$. The back condition is analogous. Since $V_1 \times V_2$ is total, we thus have $G_1 \underline{\leftrightarrow}_{total} G_2$. Hence, $G_1 = (V_1, E_1) \underline{\leftrightarrow}_{total} G_2 = (V_2, E_2)$ if and only if it is the case that $V_1 = \emptyset$ iff $V_2 = \emptyset$.

Therefore, for solving the induced subgraph bisimulation problem for input G_1 and G_2 with E_1 and E_2 being reflexive and $k \in \omega$, all we need to do is to go through the input once and check whether $V_1 = \emptyset$ iff $V_2 = \emptyset$, and whether $|V_2| \geq k$. If the answer to both is *yes* then we know that $G_1 \underline{\leftrightarrow}_{total} G_2$ and since $|V_2| \geq k$, we answer *yes*, otherwise *no*. □

Assuming the edge relation in a graph to be reflexive makes induced subgraph bisimulation a trivial problem because, unless its set of vertices is empty, every such graph is bisimilar to the graph $(\{v\}, \{(v, v)\})$. But for Kripke models, even for S5 models, this is of course not the case, as the bisimulation takes into account the valuation. Nevertheless, we will now show that also for single-agent S5 models, the problem of submodel bisimulation is significantly easier than in the case of arbitrary single-agent Kripke models. To be more precise, we will distinguish between two problems:

The first problem is *local* single-agent S5 submodel bisimulation. Here, we take as input two pointed S5 models. Then we ask whether there is a submodel of the second model that is bisimilar to the first one. Thus, the question is whether it is possible to restrict one of the models in such a way that there is a state in which the agent has exactly the same information as in the situation modelled in the other model. Note that in this problem we do not require the resulting model to be of a certain size.

Problem 3.19 (*Local single-agent S5 submodel bisimulation*)
Input A pointed S5 epistemic model $(\mathcal{M}_1, w)$ with $\mathcal{M}_1 = (W_1, \sim_1, \mathsf{V}_1)$ and $w \in W_1$, and an S5 epistemic model $\mathcal{M}_2 = (W_2, \sim_2, \mathsf{V}_2)$.
Question Is there a submodel $\mathcal{M}'_2 = (W'_2, \sim'_2, \mathsf{V}'_2)$ of $\mathcal{M}_2$ such that $(\mathcal{M}_1, w)\underline{\leftrightarrow}(\mathcal{M}'_2, w')$ for some $w' \in Dom(\mathcal{M}'_2)$?

We will show that this problem is tractable. First we introduce some notation.

Notation 3.20 Let $\mathcal{M} = (W, \sim, \mathsf{V})$ be a single-agent epistemic model. For the valuation function $\mathsf{V} : \text{PROP} \to W$, we define $\hat{\mathsf{V}} : W \to 2^{\text{PROP}}$, with $w \mapsto \{p \in \text{PROP} \mid w \in \mathsf{V}(p)\}$. Abusing notation, for $X \subseteq W$ we sometimes write $\hat{\mathsf{V}}(X)$ to denote $\{\hat{\mathsf{V}}(w) \mid w \in X\}$. We let $W/\sim$ denote the set of all equivalence classes of W for the relation $\sim$.

Proposition 3.21 *Local submodel bisimulation for single-agent pointed epistemic models is in* P.

Proof Given the input of the problem, i.e., a pointed epistemic model $\mathcal{M}_1, w$ with $\mathcal{M}_1 = (W_1, \sim_1, \mathsf{V}_1)$, and $w \in W_1$ and an epistemic model $\mathcal{M}_2 = (W_2, \sim_2, \mathsf{V}_2)$, we run the following procedure.

(1) For all $[w_2] \in W_2/\sim_2$ do the following:
   (a) Initialise the set $Z := \emptyset$.
   (b) for all $w' \in [w]$ do the following
      (i) For all $w'_2 \in [w_2]$ check if it is the case that $\hat{\mathsf{V}}_1(w') = \hat{\mathsf{V}}_2(w'_2)$. If this is the case, set $Z := Z \cup \{(w', w'_2)\}$.
      (ii) if there is no such w'_2, continue with 1 with the next element in $W_2/\sim_2$, otherwise we return Z and we stop.
(2) In case we didn't stop at 1(b)(ii), we can stop now, and return *no*.

This does not take more than $|\mathcal{M}_1| \cdot |\mathcal{M}_2|$ steps.

If the procedure has stopped at 2, there is no bisimulation with the required properties. To see this, note that if we stopped in 2, this means that there was no $[w_2] \in W_2/\sim_2$ such that for every state in $[w]$ there is one in $[w_2]$ in which exactly the same propositional letters are true. Thus, since we were looking for a bisimulation that is also defined for the state w, such a bisimulation cannot exist.

If the algorithm returned a relation Z, this is indeed a bisimulation between $\mathcal{M}_1$ and the submodel $\mathcal{M}'_2$ of $\mathcal{M}_2$ where $\mathcal{M}'_2 = (W'_2, \sim'_2, \mathsf{V}'_2)$, where

$$W'_2 = \{w_2 \in W_2 | \text{there is some } w_1 \in [w] \text{ such that } (w_1, w_2) \in Z\}$$

and $\sim'_2$ and V'_2 are the usual restrictions of $\sim_2$ and V_2 to W'_2. This follows from the two facts: First, for all pairs in Z it holds that both states satisfy exactly the same proposition letters. Second, since Z is total both on $[w]$ and on W'_2 and all the states in $[w]$ are connected to each other by $\sim_1$ and all states in W'_2 are connected to each other by $\sim'_2$, both the *forth* and *back* conditions are satisfied. This concludes the proof. □

The second problem we consider is *global* S5 submodel bisimulation, where the input are two models $\mathcal{M}_1$ and $\mathcal{M}_2$ and we ask whether there exists a submodel of $\mathcal{M}_2$ such that it is totally bisimilar to $\mathcal{M}_1$.

Problem 3.22 (*Global single-agent S5 submodel bisimulation*)
Input Two S5 epistemic models $\mathcal{M}_1 = (W_1, \sim_1, \mathsf{V}_1)$, $\mathcal{M}_2 = (W_2, \sim_2, \mathsf{V}_2)$.
Question Is there a submodel $\mathcal{M}'_2 = (W'_2, \sim'_2, \mathsf{V}'_2)$ of $\mathcal{M}_2$ such that $\mathcal{M}_1 \underline{\leftrightarrow}_{total} \mathcal{M}'_2$?

We can show that even though the above problem seems more complicated than Decision Problem 3.19, it can still be solved in polynomial time. The proof uses the fact that finding a maximum matching in a bipartite graph can be done in polynomial time (see e.g., Papadimitriou and Steiglitz 1982).

Theorem 3.23 *Global submodel bisimulation for single-agent epistemic models is in* P.

Before we give the proof, we do some pre-processing.

Definition 3.24 Given a single-agent epistemic model $\mathcal{M} = (W, \sim, \mathsf{V})$, $\mathcal{M}^{min\_cells}$ denotes a model obtained from $\mathcal{M}$ by the following procedure:

(1) Initialise X with $X := W/\sim$.
(2) Go through all the pairs in $X \times X$.
  (a) When you find $([w], [w'])$ with $[w] \neq [w']$ such that $\hat{\mathsf{V}}([w]) = \hat{\mathsf{V}}([w'])$, continue at 2 with $X := X - [w']$.
  (b) Otherwise, stop and return the model $\mathcal{M}^{min\_cells} := (\bigcup X, \sim', \mathsf{V}')$. By construction $\bigcup X \subseteq W$. Thus we can define $\sim'$ and V' as the respective restrictions of $\sim$ and V to $\bigcup X$, that is $\sim' = \sim \cap (\bigcup X \times \bigcup X)$ and $\mathsf{V}' = \mathsf{V}|_{\bigcup X}$.

Fact 3.25 With input $\mathcal{M} = (W, \sim, \mathsf{V})$, the procedure in Definition 3.24 runs in time polynomial in $|\mathcal{M}|$.

Proof Follows from the fact that the cardinality of $W/\sim$ is bounded by $|W|$; we only enter step 2 at most $|W|$ times, and each time do at most $|W|^2$ comparisons. □

Fact 3.26 The answer to total submodel bisimulation for single-agent epistemic models (Decision Problem 3.22) with input $\mathcal{M}_1 = (W_1, \sim_1, \mathsf{V}_1)$, $\mathcal{M}_2 = (W_2, \sim_2, \mathsf{V}_2)$ is *yes* iff it is with input ${\mathcal{M}_1}^{min\_cells} = (W_1, \sim_1, \mathsf{V}_1)$, $\mathcal{M}_2 = (W_2, \sim_2, \mathsf{V}_2)$.

Proof From left to right, we just need to restrict the bisimulation to the states of ${\mathcal{M}_1}^{min\_cells}$. For the other direction, we start with the given bisimulation and then extend it as follows. The states in a cell $[w']$ that was removed during the construction of ${\mathcal{M}_1}^{min\_cells}$ can be mapped to the ones of a cell $[w]$ in ${\mathcal{M}_1}^{min\_cells}$ with the same valuation. □

We can now prove Theorem 3.23.

Proof By Facts 3.25 and 3.26, transforming $\mathcal{M}_1$ into ${\mathcal{M}_1}^{min\_cells}$ can be done in polynomial time. Thus, without loss of generality, we can assume that $\mathcal{M}_1$ is already of the right shape; i.e., $\mathcal{M}_1 = {\mathcal{M}_1}^{min\_cells}$. Given the two models as input, we construct a bipartite graph $G = ((W_1/{\sim_1}, W_2/{\sim_2}), E)$ where E is defined as follows.

$$([w_1], [w_2]) \in E \textit{ iff } \hat{\mathsf{V}}_1([w_1]) \subseteq \hat{\mathsf{V}}_2([w_2]).$$

Claim 3.27 *The following are equivalent.*

(a) *There is a submodel $\mathcal{M}'_2$ of $\mathcal{M}_2$ such that $\mathcal{M}_1 \underline{\leftrightarrow}_{total} \mathcal{M}'_2$*
(b) *G has a matching of size $|W_1/{\sim_1}|$.*

Proof Assume that there is a submodel $\mathcal{M}'_2 = (W'_2, \sim'_2, \mathsf{V}'_2)$ of $\mathcal{M}_2$ such that $\mathcal{M}_1 \underline{\leftrightarrow}_{total} \mathcal{M}'_2$. Let Z be such a total bisimulation.

Note that since we assumed that $\mathcal{M}_1 = \mathcal{M}^{min\_cells}$ the following holds:

(1) For all $([w_1], [w_2]) \in W_1/{\sim_1} \times W_2/{\sim_2}$ it is the case that whenever $Z \cap ([w_1] \times [w_2]) \neq \emptyset$, then for all $[w'_1] \in W_1/{\sim_1}$ such that $[w'_1] \neq [w_1]$, $Z \cap ([w'_1] \times [w_2]) = \emptyset$.

Thus, the members of different equivalence classes in $W_1/{\sim_1}$ are mapped by Z to into different equivalence classes of $W_2/{\sim_2}$.

Now, we construct $\dot{E} \subseteq E$ as follows.

$$([w_1], [w_2]) \in \dot{E} \textit{ iff } ([w_1], [w_2]) \in E \textit{ and } ([w_1] \times [w_2]) \cap Z \neq \emptyset.$$

Then $|\dot{E}| \geq |W_1/{\sim_1}|$ because of the definitions E and $\dot{E}$ and the fact that Z is a bisimulation that is total on W_1. Now, if $|\dot{E}| = |W_1/{\sim_1}|$ then we are done since by definition of $\dot{E}$, for each $[w_1] \in W_1/{\sim_1}$ there is some $[w_2] \in W_2/{\sim_2}$ such that $([w_1], [w_2]) \in \dot{E}$. Then it follows from 1, that $\dot{E}$ is indeed a matching.

If $|\dot{E}| > |W_1/{\sim_1}|$ then we can transform $\dot{E}$ into a matching E' of size $|W_1/{\sim_1}|$: For each $[w_1] \in W_1/{\sim_1}$, we pick one $[w_2] \in W_2/{\sim_2}$ such that $([w_1], [w_2]) \in \dot{E}$ and put it into E' (note that such a $[w_2]$ always exists because by definition of $\dot{E}$, for each $[w_1] \in W_1/{\sim_1}$ there is some $[w_2] \in W_2/{\sim_2}$ such that $([w_1], [w_2]) \in \dot{E}$; moreover because of 1 all the $[w_2] \in W_2/{\sim_2}$ that we pick will be different). Then the resulting $E' \subseteq \dot{E} \subseteq E \subseteq (W_1/{\sim_1} \times W_2/{\sim_2})$ is a matching of G of size $|W_1/{\sim_1}|$. Thus, we have shown that if there is a submodel $\mathcal{M}'_2$ of $\mathcal{M}_2$ such that $\mathcal{M}_1 \underline{\leftrightarrow}_{total} \mathcal{M}'_2$ then G has a matching of size $|W_1/{\sim_1}|$.

For the other direction, assume that G has a matching $E' \subseteq E$ with $|E'| = |W_1/{\sim_1}|$. Then, recalling the definition of E, it follows that for all $[w] \in W_1/{\sim}$ there is some $[w'] \in W_2/{\sim_2}$ such that $([w], [w']) \in E'$ and thus $\hat{\mathsf{V}}_1([w]) \subseteq \hat{\mathsf{V}}_2([w'])$.

Let us define the following submodel $\mathcal{M}'_2$ of $\mathcal{M}_2$. $\mathcal{M}'_2 = (W'_2, \sim'_2, \mathsf{V}'_2)$, where

$$W'_2 = \left\{ w_2 \in W_2 | \text{ there is a } w \in W_1 \text{ with } \hat{\mathsf{V}}_1(w) = \hat{\mathsf{V}}_2(w_2) \text{ and } ([w], [w_2]) \in E' \right\}$$

and $\sim_2'$ and V_2' are the usual restrictions of $\sim_2$ and V_2 to W_2'.

Now, we define a relation $Z \subseteq W_1 \times W_2'$, which we then show to be a total bisimulation between $\mathcal{M}_1$ and $\mathcal{M}_2'$:

$$(w_1, w_2) \in Z \text{ iff } \hat{\mathsf{V}}(w_1) = \hat{\mathsf{V}}_2(w_2) \text{ and } ([w_1], [w_2]) \in E'.$$

Next, let us show that Z is indeed a bisimulation.

Let $(w_1, w_2) \in Z$. Then, by definition of Z, for every propositional letter p, we have $w_1 \in \mathsf{V}_1(p)$ iff $w_2 \in \mathsf{V}_2(p)$. Next, we check the forth condition. Let $w_1 \sim_1 w_1'$ for some $w_1' \in W_1$. Since $(w_1, w_2) \in Z$, and thus $([w_1], [w_2]) \in E'$, there is some $w_2' \in [w_2]$ such that $\hat{\mathsf{V}}_2(w_2') = \hat{\mathsf{V}}_1(w_1')$. Since $[w_1'] = [w_1]$ and $[w_2'] = [w_2]$, it follows that $([w_1'], [w_2']) \in E'$. Then $w_2' \in W_2'$, and $(w_1', w_2') \in Z$.

For the back condition, let $w_2 \sim_2 w_2'$, for some $w_2' \in W_2'$. Then by definition of W_2', there is some $w \in W_1$ such that $\hat{\mathsf{V}}_1(w) = \hat{\mathsf{V}}_2(w_2')$ and $([w], [w_2']) \in E'$. It follows that $(w, w_2') \in Z$. Now, we show that $w_1 \sim_1 w$. As the following holds: $([w], [w_2']) \in E'$, $[w_2] = [w_2']$, $([w], [w_2]) \in E'$ (because $(w_1, w_2) \in Z$) and E' is a matching, it follows that $[w] = [w_1]$. Thus, $w_1 \sim_1 w$.

Hence, Z is a bisimulation. It remains to show that Z is indeed total.

Let $w_1 \in W_1$. Since E' is a matching of size $|W_1/\!\sim_1|$, there is some $[w_2] \in W_2/\!\sim_2$ such that $([w_1], [w_2]) \in E'$. Thus, there is some $w_2' \in [w_2]$ such that $\hat{\mathsf{V}}_1(w_1) = \hat{\mathsf{V}}_2(w_2')$. Hence $w_2' \in W_2'$ and $(w_1, w_2') \in Z$. So Z is total on W_1.

Let $w_2 \in W_2'$. By definition of W_2', there is some $w \in W_1$ such that $\hat{\mathsf{V}}_1(w) = \hat{\mathsf{V}}_2(w_2)$ and $([w], [w_2]) \in E'$. Thus, by definition of Z, $(w, w_2) \in Z$. Therefore, Z is indeed a total bisimulation between $\mathcal{M}_1$ and $\mathcal{M}_2'$. This concludes the proof of Claim 3.27. □

Hence, given two models, we can transform the first one using the polynomial procedure of Definition 3.24 and then we construct the graph G, which can be done in polynomial time as well. Finally, we use a polynomial algorithm to check if G has a matching of size $M_1{}^{min\_cells}$. If the answer is yes, we return yes, otherwise no. This concludes the proof of Theorem 3.23. □

Now, the question arises whether the above results also hold for the multi-agent case.

Problem 3.28 (*Global multi-agent S5 submodel bisimulation*)
Input Two epistemic models $\mathcal{M}_1 = (W_1, (\sim_{1i})_{i \in \mathrm{N}}, \mathsf{V}_1)$, $\mathcal{M}_2 = (W_2, (\sim_{2i})_{i \in \mathrm{N}}, \mathsf{V}_2)$, for $\mathtt{N}$ being a finite set (*of agents*).
Question Is there a submodel $\mathcal{M}_2' = (W_2', (\sim_{2_i}')_{i \in \mathrm{N}}, \mathsf{V}_2')$ of $\mathcal{M}_2$ such that $\mathcal{M}_1 \underline{\leftrightarrow}_{total} \mathcal{M}_2'$?

Open Problem 3.29 Is *global multi-agent S5 submodel bisimulation* NP-hard?

We expect the answer to this question to be positive, as for S5, there seems to be a complexity jump between the single-agent case and the two-agent case: In case of the satisfiability problem of the logic, the one-agent logic is NP-complete, whereas as soon as we have at least two agents, we get PSPACE completeness. Similarly, in

Section 3.1, we showed in Proposition 3.5 that also for bisimilarity, there seems to be a complexity jump for S5 models when a second agent is added: the problem becomes P-hard and thus as hard as the problem for arbitrary Kripke models.

The idea behind these results is that arbitrary accessibility relations can be simulated by a concatenation of two equivalence relations. However, these techniques, as they have been used, e.g., by Halpern and Moses (1992), and do not seem to work for transforming models into S5-models for two agents such that the existence of submodels bisimilar to some model is preserved. The problem is caused by the fact that the resulting model has to be reflexive, in which case several states could be collapsed whereas they could not before the transformation. Thus, a coding of the existence of a successor and the existence of reflexive loops in the original model would be required to take care of this issue.[3]

Let us summarise our complexity results for problems related to deciding whether it is possible to restrict an agent's information structure so that after the restriction he will have similar information as another agent in some other situation. We showed that induced subgraph bisimulation is intractable (NP-complete). Using this, we could show that the same holds for submodel bisimulation of arbitrary Kripke models.

For partition-based graphs (with the edge relations being equivalence relations) however, we showed that the problem of induced subgraph bisimilarity is very easy: it is solvable in linear time if we are looking for a subgraph of a certain size with a total bisimulation. Deciding whether there is any subgraph with a total bisimulation is trivial for partition-based graphs, as then all such non-empty graphs are bisimilar. In fact, this already holds if the edge relation is reflexive.

Extending these results for S5 Kripke models, we could show that submodel bisimulation for single-agent models is not as trivial as for graphs, but still in P. For multi-agent S5 models, the problem remains open. We conjecture it to be polynomially equivalent to the problem for Kripke models in general. The technical challenge in showing this lies in the simulation of arbitrary accessibility relations by a combination of different equivalence relations. The idea would again be to use a method similar to that used, e.g., by Halpern and Moses (1992), and replace every edge by a composition of two equivalence relations. This technique can however not be applied in its standard way as some issues arise due to the fact that the resulting model has to be reflexive while in the original arbitrary Kripke structure this does not need to be the case. This means that the existence of reflexive loops in the original model would somehow have to be coded using a propositional letter so as not to lose the information during the transformation.

In dynamic systems with diverse agents, an interesting question is whether it is possible to give some information to one agent such that afterward she knows at least as much as some other agent. This is captured by an asymmetric notion, that of simulation (Definition 2.17). With this difference, the question can be raised of the effect on tractability and intractability of requiring simulation versus requiring bisimulation. With this motivation, we would like to explore the problem of induced subgraph simulation.

[3] Note that the situation in Proposition 3.5 is different, as there we started with models that were irreflexive and thus we did not need to take care of coding the information about loops and could apply the standard technique.

Problem 3.30 (*Induced subgraph simulation)*
Input Two finite graphs $G_1 = (V_1, E_1)$, $G_2 = (V_2, E_2)$, $k \in \omega$.

Question Is there an induced subgraph of G_2 with at least k vertices that is simulated by G_1, i.e., is there some $V' \subseteq V_2$ with $|V'| \geq k$ and $(V', E_2 \cap (V' \times V')) \sqsubseteq_{total} G_1$?

Proposition 3.31 *Induced subgraph simulation is* NP*-complete.*

Proof Showing that the problem is in NP is straightforward. Hardness is shown by reduction from independent set. First of all, let $I_k = (V_{I_k}, E_{I_k} = \emptyset)$ with $|V_{I_k}| = k$ denote a graph with k vertices and no edges. Given the input of independent set, i.e., a graph $G = (V, E)$ and some $k \in \omega$, we transform it into (I_k, G), k, as input for induced subgraph simulation.

Now, we claim that G has an independent set of size at least k iff there is some $V' \subseteq V$ with $|V'| \geq k$ and $(V', E \cap (V' \times V')) \sqsubseteq_{total} I_k$. From left to right, assume that there is some $S \subseteq V$ with $|S| = k$, and for all $v, v' \in S$, $(v, v') \notin E$. Now, any bijection between S and V_{I_k} is a total simulation (and in fact an isomorphism) between $G' = (S, E \cap (S \times S))$ and I_k, since $E \cap (S \times S) = \emptyset$ and $|S| = |V_{I_k}|$.

For the other direction, assume that there is some $V' \subseteq V$ with $|V'| = k$ such that for $G' = (V', E' = E \cap (V' \times V'))$ we have that $G' \sqsubseteq_{total} I_k$. Thus, there is some total simulation Z between G' and I_k. Now, we claim that V' is an independent set of G of size k. Let $v, v' \in V'$. Suppose that $(v, v') \in E$. Since G' is an induced subgraph, we have that $(v, v') \in E'$. Since Z is a total simulation, there is some $w \in I_k$ with $(v, w) \in Z$ and some w' with $(w, w') \in E_{I_k}$ and $(v', w') \in Z$. But this is a contradiction with $E_{I_k} = \emptyset$. Thus, V' is an independent set of size k of G. The reduction can clearly be computed in polynomial time. This concludes the proof. □

In De Nardo et al. (2009), it has been shown that given two graphs it is also NP-complete to decide if there is a subgraph (not necessarily an induced one) of one such that it is simulation equivalent to the other graph. Here, we show that this also holds if the subgraph is required to be an induced subgraph.

Problem 3.32 (*Induced subgraph simulation equivalence*)
Input Two finite graphs $G_1 = (V_1, E_1)$, $G_2 = (V_2, E_2)$, $k \in \omega$.

Question Is there an induced subgraph of G_2 with at least k vertices that is similar to G_1, i.e., is there some $V' \subseteq V_2$ with $|V'| \geq k$ and $(V', E_2 \cap (V' \times V')) \sqsubseteq_{total} G_1$ and $G_1 \sqsubseteq_{total} (V', E_2 \cap (V' \times V'))$?

Proposition 3.33 *Induced subgraph simulation equivalence is* NP*-complete.*

Proof For showing that the problem is in NP, note that we can use a simulation equivalence algorithm as provided in Henzinger et al. (1995). Hardness can again be shown by reduction from independent set. Given the input for independent set, i.e., a graph $G = (V, E)$ and some $k \in \omega$, we transform it into two graphs $I_k = (V_{I_k} = \{v_1, \ldots, v_k\}, E_{I_k} = \emptyset)$ and G, and we keep the $k \in \omega$. This can be done in polynomial time.

Now, we claim that G has an independent set of size k iff there is an induced subgraph of G with k vertices that is similar to I_k. From left to right, assume that

G has such an independent set S with $S \subseteq V$, $|S| = k$ and $E \cap S \times S = \emptyset$. Then $(S, \emptyset)$ is isomorphic to I_k since both have k vertices and no edges. Thus, they are also simulation equivalent.

For the other direction, assume that there is an induced subgraph $G' = (V', E')$ with $V' \subseteq V$, $|V'| = k$ and $E' = (V' \times V') \cap E$ such that G' is simulation equivalent to I_k. Suppose that there are $v, v' \in V'$ such that $(v, v') \in E$. Since G' is an induced subgraph, it must be the case that $(v, v') \in E'$, but since I_k simulates G', this leads to a contradiction since I_k does not have any edges. This concludes the proof. □

Along similar lines, we could consider further problems involving more interaction between agents, and ask, e.g., by Halpern and Moses (1992), whether one model represents the result of two agents sharing all their information.

As a corollary of the two previous propositions we get that for arbitrary Kripke models both submodel simulation and submodel simulation equivalence are NP-hard. An NP upper bound follows from the fact that given a relation between a model and a submodel of some other model, it can be checked in polynomial time if this relation is indeed a simulation.

Corollary 3.34 *Deciding submodel simulation and submodel equivalence for Kripke structures is* NP-*complete.*

For single-agent S5, we can use the methods as used in the proof of Theorem 3.23 in order to obtain a polynomial procedure for the single-agent case.

Proposition 3.35 *Deciding submodel simulation and submodel equivalence for single-agent S5 models is in* P.

Proof We use the procedure of the proof of Theorem 3.23. This also works for simulation and simulation equivalence because of the totality constraint and the fact that as we deal with S5 models, we only need to take care of the different valuations occurring in the equivalence classes. □

Let us now summarise our complexity analysis of tasks that involve checking whether in one situation an agent knows at least as much as another agent in a possibly different situation. We have shown that we can extend graph-theoretical complexity results about subgraph simulation equivalence to the case where the subgraph is required to be an induced subgraph. Via this technical result, we can then transfer the complexity bounds also for the problem of submodel simulation (equivalence) of Kripke models, which with an epistemic interpretation of the accessibility relation is the following problem: *decide whether it is possible to give information to one agent so that as a result he knows as least as much as some other agent*. In case of partition-based models (S5), for a single agent this problem can be solved in polynomial time analogously to how we have done it for submodel bisimulation. For the multi-agent case, the problem remains open, however. As for submodel bisimulation of multi-agent S5 model, the technical issue that would have to be solved for showing NP-hardness is caused by the reflexivity of the underlying relations.

4 Conclusions and further questions

We now summarise our main results, provide some conclusions and finish with some further questions.

4.1 Summary

We have identified concrete epistemic tasks related to the comparison and manipulation of information states of agents in possibly different situations. Interestingly, our complexity analysis shows that such tasks and decision problems live on both sides of the border between tractability and intractability. We now summarise our results for the different kinds of tasks we investigated.

4.1.1 Information similarity

Our results for the complexity of deciding whether information structures are similar can be found in Table 3.

If we take isomorphism as our similarity notion, then in general (without any particular assumptions on the Kripke structures representing agents' information) it is open whether checking if two information structures are similar is tractable. This follows from the fact that checking if two Kripke models are isomorphic is as hard as the graph isomorphism problem which is neither known to be in P nor known to be NP-hard. Thus, we can say that given the current knowledge, for isomorphism, deciding if two information structures are similar can be located on the border between tractability and intractability. We did not investigate the isomorphism problem for S5 but conjecture it to become as hard as Kripke model isomorphism (GI-complete) as soon as we have at least two agents.

Taking bisimilarity as the similarity notion, deciding if two structures are similar is among the hardest problems known to be tractable. If the models are based on partitions (S5), the problem is very easy in the single-agent case but also becomes P-hard in the multi-agent case.

4.1.2 Information symmetry

Table 4 summarises the results of our complexity analysis of tasks concerned with deciding whether the information of agents is symmetric, where symmetry can be understood in different ways.

Table 3 Complexity results for deciding information similarity

Problem	Tractable?	Comments
Kripke model isomorphism	Unknown	GI-complete
Epistemic model bisimilarity	Yes	P-complete in the multi-agent case

Table 4 Complexity results for deciding information symmetry

Problem	Tractable?	Comments
Common knowledge of a fact	Yes	Solvable using a reachability algorithm
Horizon bisimilarity (Kripke models)	Yes	P-complete for arbitrary models, even for horizons at the same point in the model
Flipped horizon bisimilarity (Kripke models)	Yes	P-complete, even for horizons at the same point in the model
Horizon bisimilarity (S5-models)	Yes	Trivial for horizons at the same point in a model
Flipped horizon bisimilarity (S5-models)	Yes	Problem does not get easier for horizons at the same point in a model

We started our investigation of information symmetry with the symmetry of two agents' knowledge about a given fact being true. This kind of symmetry arises if the fact is common knowledge among the two agents. Given an information structure, deciding if this is the case can be done using a reachability algorithm that checks for every state at which the fact is not true whether there is a path to it (via the union of the two relations of the agents) from the current state. This is the case if and only if the fact is not common knowledge between the two agents.

We then introduced the notion of (*epistemic*) *horizon*, which represents the submodel that is relevant for an agent at a given situation (i.e., at a given point in the model). The horizon of an agent in a situation is the submodel that is generated by the set of worlds the agent considers possible in that situation. When considering the epistemic reasoning of agents, our notion of horizon plays a crucial role as an agent's horizon contains exactly the possible worlds that the agent might take into consideration during his reasoning. We have shown that in general deciding if the horizons of two agents are bisimilar is exactly as hard as deciding bisimilarity of Kripke models. Without assuming reflexivity of the accessibility relation, deciding about the similarity of two agents' horizons does not get easier in the special case in which we compare horizons at the very same point in a model. As soon as the accessibility relations of the two agents under consideration are reflexive, however, the problem becomes completely trivial if we compare horizons at the same point in the model, as they are always identical. Thus, if we take information structures to be arbitrary Kripke models, then in general comparing horizons of agents in one given situation is as hard as comparing information structures in general. For S5 models, however, the situation is slightly different as horizon bisimilarity becomes trivial for horizons taken at the same point in a model.

For our investigation of information symmetry, we have introduced the notion of flipped bisimilarity, which captures the similarity of two models after swapping the information of two agents. For Kripke structures in general, the complexity of deciding flipped bisimilarity is just as for bisimilarity. For the special case in which two pointed structures are identical deciding bisimilarity is trivial but flipped bisimilarity can be as hard as it is for arbitrary pointed Kripke structures.

Our results for horizon comparison for arbitrary Kripke models show that both flipped bisimilarity and regular bisimilarity are P-complete, even if we take the horizon at the very same situation. Thus, comparing different agents' perspectives on the very same situation is as hard as comparing structures in general. Under the assumption of partition-based information structures (S5) however, we observed a significant difference between bisimilarity and flipped bisimilarity of horizons. While bisimilarity of horizons of different agents becomes trivial if they are taken at the very same situation (i.e., at the same point in the model), flipped bisimilarity stays as hard as it is for multi-agent S5 models in general.

Let us briefly summarise the technical facts that explain our complexity results as given in Table 4.

- Problems about information symmetry which can be solved by checking if certain states are reachable by (combinations of) agents' accessibility relations are relatively easy as they boil down to solving the reachability problem which is NL-complete.
- Problems involving bisimilarity of arbitrary models are among the hardest tractable problems and thus believed to be slightly easier than problems involving isomorphism of Kripke models as for isomorphism no polynomial algorithms are known.
- In the single-agent case, assuming S5 relations makes bisimilarity easier because checking for bisimilarity boils down to just comparing the propositional valuations of information cells.
- While flipped bisimilarity does not seem to be more complex than regular bisimilarity, the fact that flipped bisimilarity is in general not reflexive has the effect of making it harder than bisimilarity in the special case where we ask if a pointed model is flipped bisimilar to itself.

4.1.3 Information manipulation

Apart from the rather static problems about the comparison of information structures, we also investigated the complexity of tasks related to more dynamic aspects of information. In many interactive processes, the information of agents changes through time because agents can make observations or receive new information from other agents. Then an interesting question that arises is whether given an information state of an agent it is possible that through incoming information, the agent's information structure can change such that in the end the agent has similar information to some other agent.

Table 5 summarises the results of our complexity analysis of tasks concerned with the manipulation of information structures.

To determine the complexity of deciding whether it is possible to restrict an information structure in such a way such that it becomes similar to some other structure, we started by investigating the NP-complete graph-theoretical problem *subgraph bisimulation*. The problem is to decide whether one of two given graphs has a subgraph which is bisimilar to the other graph. We showed that it remains NP-complete if we require the subgraph to be an *induced* subgraph. This technical result then allowed us

Table 5 Complexity results for tasks about information manipulation

Problem	Tractable?	Comments
Kripke submodel bisimulation	No	NP-complete. Reduction from independent set
Single agent S5 submodel bisimulation	Yes	Local version easier; in general an algorithm for finding matchings in bipartite graphs can be used
Multi-agent S5 submodel bisimulation	Unknown	Conjectured to be NP-complete
Kripke submodel simulation (equivalence)	No	NP-complete. Reduction from independent set
Single agent S5 submodel simulation (equivalence)	Yes	Similar polynomial procedure as for single-agent S5 submodel bisimulation
Multi-agent S5 submodel simulation (equivalence)	Unknown	Same technical issues as for S5 submodel bisimulation

to show that for Kripke models, it is also NP-complete to decide if one given model has a submodel which is bisimilar to another given model. We then showed that this problem does indeed get easier if we have S5 structures with one agent only: we gave a polynomial procedure that uses the fact that computing whether there is a matching of a certain size in a bipartite graph can be done in polynomial time. This shows that deciding if an agent's information can be restricted in a certain way is easier under the assumption of S5 information structures. It remains open to show whether the problem also becomes intractable for S5 as soon as we have more than one agent. The technical issue which needs to be resolved here is to determine whether an arbitrary accessibility relation can be simulated by the composition of two equivalence relations in such a way that the existence of a submodel bisimilar to some other model is preserved. While it is relatively straightforward to make sure that in the model that results from the transformation the accessibility relations are symmetric and transitive, the requirement of reflexivity seems to cause some problems.

Instead of asking whether it is possible to give some information to an agent such that the resulting information structure is similar to some other structure, in many situations it might be sufficient to know if it is possible to manipulate the information of an agent such that he will know at least as much as some other agent. In more general terms, this leads us to the task of deciding whether it is possible to restrict some structure such that it becomes at least as refined as some other structure. Similar to the case of submodel bisimulation, we started by investigating the problem of induced subgraph simulation, which we showed to be NP-complete by reduction from *independent set*. Using this, we could then show that submodel simulation is NP-complete for Kripke models.

Under the assumption of S5 models, we can adapt the polynomial procedure that we had for single-agent S5 local submodel bisimulation for solving the analogous problem for simulation. This means that with S5 models, it is tractable to decide if we can restrict an information structure for one agent such that it becomes at least as refined as that of another agent in a given situation. We get an analogous result for simulation equivalence, a weaker notion of similarity than bisimulation. Whether on

S5 structures these problems become intractable as soon as we have models with at least two agents is open, and depends on the same technical issues as this problem for submodel bisimulation.

Let us briefly summarise the technical facts that explain our complexity results as listed in Table 5.

- For single-agent S5 models, submodel bisimilarity and simulation equivalence turned out to be solvable in polynomial time. Given a model A and a model B, we consider the bipartite graph that consists of the equivalence classes of both models and in which edges connect two information cells of each model if and only if all the valuations occurring in the first information cell, also occur in the second. We then used the fact that for single-agent S5 models, model A has a submodel bisimilar to model B if and only if the previous bipartite graph has a matching of size k, where k is the number of equivalence classes in model B.
- In general, submodel bisimilarity of Kripke models is NP-complete, as a graph having an induced subgraph bisimilar to the graph of k isolated points is equivalent to the graph having an independent set of size k.
- Whether submodel bisimilarity is NP-complete for S5 models with more than one agent depends on how arbitrary Kripke structures can be simulated using the combination of two equivalence classes in such a way that the existence of submodels bisimilar to another model is preserved.

4.2 Conclusions

From the above results, we conclude the following for the three classes of tasks that we have analysed.

Information similarity

- If information of agents is modelled by simple relational structures without any particular assumptions, then deciding if the information of agents in two different multi-agent situations is similar will in general be somewhere in between tractable but hard and the border to intractability.
- Under the assumption of S5 properties of the information structures, the complexity jump from easy to P-hard happens with the introduction of a second agent.

Information symmetry

- All problems we encountered are tractable, but nevertheless we were able to identify significant differences in their complexity.
  - Comparing the perspectives of agents in the very same situation becomes trivial as soon as we check for bisimilarity and the agents' accessibility relations are reflexive.
  - For checking if the agents have similar information *about each other* (captured by flipped bisimilarity of horizons) however, neither the assumption of reflexivity of the agents' relations nor considering horizons at the very same point in the model make the problem easier.
  - Thus, deciding if agents have similar information *about each other* can in certain cases be harder than deciding if agents have similar information.

Information manipulation

- For the problems we identified for deciding whether an information structure can be restricted in such a way that it will be in a certain relation to another model, we get the same pattern of complexity results for simulation, simulation equivalence and bisimulation.
  - Deciding whether a model can be restricted such that it is in one of those three relations to another model is tractable for single-agent S5 models and intractable in general.
  - Whether for S5 models the jump from being tractable to being intractable happens with the introduction of a second agent depends on whether we can simulate arbitrary relations by a combination of two equivalence relations while preserving the existence of submodels that are in a certain relationship to another model.

Comparing the three classes of tasks (about information similarity, symmetry and manipulation), information similarity is the easiest one in general if we stick to bisimulation as our notion of similarity. For information symmetry, all the problems we identified are tractable, with some special cases even being trivial, such as for reflexive models similarity of horizons at the same situation. Deciding if two agents have the same information about each other, however, does not become trivial unless the two agents are equal. For deciding whether it is possible to manipulate agents' information in a certain way, we considered problems of a wide variety of complexities ranging from very easy to NP-complete. Deciding if it is possible to restrict an information structure such that it becomes similar to or at least as refined as another is easiest if we only have one agent and assume S5 models. For arbitrary Kripke structures the problem is NP-complete. For multi-agent S5 models we conjecture it to be NP-complete as well. Locating the tractability border in epistemic tasks on modal logic frameworks, we conclude that for the static tasks concerning similarity and symmetry, most problems are tractable, whereas for the dynamic tasks involving the manipulation of information intractable tasks arise when we have multiple agents. In general, for S5 models, complexity jumps for various tasks seem to occur when a second agent is introduced.

Let us now come back to our research question.

Research Question 4.1

- Which parameters can make interaction difficult?
- How does the complexity of an interactive situation change when more participants enter the interaction or when we drop some simplifying assumptions on the participants themselves?

In the context of concrete tasks in reasoning about epistemic agents, we can give the following answers.

(1) The complexity of comparing the information of diverse agents crucially depends on the notion of similarity used.

(2) Under standard assumptions about knowledge (veridicality and full introspection), intractable tasks can become very easy.

(3) Moreover, under these assumptions, for various tasks a complexity jump occurs with the introduction of a second agent.
(4) Without any assumptions on information structures reasoning about a single agent seems to be already as hard as reasoning about multi-agent situations.

4.3 Further questions

Our work gives rise to some interesting questions for further investigation. Let us start with some technical open problems.

4.3.1 Does submodel bisimulation for S5 become intractable with two agents?

It remains open to show whether the problem also becomes intractable for S5 as soon as we have more than one agent. The technical issue which needs to be resolved here is to determine whether an arbitrary accessibility relation can be simulated by the composition of two equivalence relations in such a way that the existence of a submodel bisimilar to some other model is preserved.[4]

While it is relatively straightforward to make sure that in the model that results from the transformation the accessibility relations are symmetric and transitive, the requirement of reflexivity seems to cause some problems.

4.3.2 Does submodel simulation (equivalence) for S5 become intractable with two agents?

Whether on S5 structures, the problems of submodel simulation and submodel simulation equivalence become intractable as soon as we have models with at least two agents depends on the same technical issues as the problem for submodel bisimulation.

A more general problem that came up in our analysis is the following.

4.3.3 Is S5 models simulation (equivalence) at least as hard as bisimulation?

We did not investigate simulation and simulation equivalence of information structures. Here, an interesting general question arises as to whether also for epistemic (S5) models it holds that in general simulation (equivalence) is at least as hard as bisimulation as this holds for Kripke structures (Kučera and Mayr 2002).

4.3.4 Linking up to real epistemic reasoning

Besides the technical questions above, our results call for an empirical investigation of the tasks we identified in order to clarify the correspondence between our results and the cognitive difficulties involved in epistemic reasoning. For this, we note that the

[4] We stress that here we are concerned with *submodels*, i.e., in general these are *not* generated submodels.

formal concepts that we used in the decision problems (e.g., bisimilarity) were mostly motivated by the fact that they come up naturally in the context of modal logics.

However, for being able to draw conclusions about the complexity that real agents face in epistemic reasoning, it needs to be investigated which are cognitively adequate notions of similarity. One possibility would be to work out the connection between the similarity notions that we considered and those underlying analogical reasoning in interactive situations (cf. Besold et al. 2011).

Summing up our investigation so far, we have moved from a high-level perspective on the epistemic reasoning abilities of agents to an analysis of concrete tasks about information structures which represent the uncertainties that individuals have about the situation they are in. Since our work is originally motivated by the need of a formal theory of real interaction, this leads to the ultimate question about applicability of our analysis to real epistemic situations, and thus back to interaction in real life.

Acknowledgments Cédric Dégremont and Jakub Szymanik gratefully acknowledge the support of Vici Grant NWO-277-80-001.

References

Ågotnes, T., Balbiani, P., Van Ditmarsch, H., & Seban, P. (2010). Group announcement logic. *Journal of Applied Logic, 8*(1), 62–81.

Aucher, G. (2010). An internal version of epistemic logic. *Studia Logica, 94*(1), 1–22.

Aumann, R. J. (1999). Interactive epistemology I: Knowledge. *International Journal of Game Theory, 28*(3), 263–300.

Balbiani, P., Baltag, A., Van Ditmarsch, H., Herzig, A., Hoshi, T., & de Lima, T. (2008). 'Knowable' as 'known after an announcement'. *The Review of Symbolic Logic, 1*(03), 305–334.

Balcázar, J. L., Gabarró, J., & Santha, M. (1992). Deciding bisimilarity is P-complete. *Formal Aspects of Computing, 4*(6A), 638–648.

Baltag, A., & Moss, L. S. (2004). Logics for epistemic programs. *Synthese, 139*(2), 165–224.

Baltag, A., Moss, L. S., & Solecki, S. (1998). The logic of public announcements, common knowledge, and private suspicions. In *TARK '98*: *Proceedings of the 7th Conference on Theoretical Aspects of Rationality and Knowledge*, San Francisco, CA, USA, 1998 (pp. 43–56). Burlington, MA: Morgan Kaufmann Publishers Inc.

Baral, C., & Zhang, Y. (2005). Knowledge updates: Semantics and complexity issues. *Artificial Intelligence, 164*(1–2), 209–243.

Besold, T. R., Gust, H., Krumnack, U., Abdel-Fattah, A., Schmidt, M., & Kühnberger, K.-U.. (2011). An argument for an analogical perspective on rationality and decision-making. In R. Verbrugge, & J. van Eijck (Eds.), *Proceedings of the Workshop on Reasoning About Other Minds*: *Logical and Cognitive Perspectives* (*RAOM-2011*), *Groningen*, *The Netherland*, *11 July 2011*, volume 751 of *CEUR Workshop Proceedings*, (pp. 20–31). CEUR-WS.org.

Blackburn, P., De Rijke, M., & Venema, Y. (2001). *Modal logic*. Number 53 in Cambridge Tracts in Theoretical Computer Science. Cambridge: Cambridge University Press.

De Nardo, L., Ranzato, F., & Tapparo, F. (2009). The subgraph similarity problem. *IEEE Transactions on Knowledge and Data Engineering, 21*(5), 748–749.

De Rijke, M. (1993). *Extending Modal Logic*. PhD Thesis, ILLC, University of Amsterdam. ILLC Dissertation Series DS-93-04.

Dovier, A., & Piazza, C. (2003). The subgraph bisimulation problem. *IEEE Transactions on Knowledge and Data Engineering, 15*(4), 1055–1056.

Edmonds, J. (1965). Paths, trees, and flowers. *Canadian Journal of Mathematics, 17*, 449–467.

Fagin, R., Halpern, J. Y., Moses, Y., & Vardi, M. Y. (1995). *Reasoning about Knowledge*. Cambridge, MA: MIT Press.

Feltovich, N. (2000). Reinforcement-based vs. beliefs-based learning in experimental asymmetric-information games. *Econometrica, 68*, 605–641.
French, T., & Van Ditmarsch, H. (2008). Undecidability for arbitrary public announcement logic. In C. Areces & R. Goldblatt (Eds.), *Advances in Modal Logic* (pp. 23–42). College Publications.
Garey, M. R., & Johnson, D. S. (1990). *Computers and Intractability: A Guide to the Theory of NP-completeness*. New York: W. H. Freeman and Co..
Gerbrandy, J. (1999). *Bisimulations on Planet Kripke*. PhD Thesis, ILLC, University of Amsterdam. ILLC Dissertation Series DS-1999-01.
Gierasimczuk, N., & Szymanik, J. (2011a). Invariance properties of quantifiers and multiagent information exchange. In M. Kanazawa, A. Kornai, M. Kracht & H. Seki (Eds.), *Proceedings of 12th Meeting on Mathematics of Language*, volume 6878 of *Lecture Notes in Computer Science* (pp. 72–89). Berlin: Springer.
Gierasimczuk, N., & Szymanik, J. (2011b). A note on a generalization of the muddy children puzzle. In Krzysztof R. Apt, editor, *Proceedings of the 13th Conference on Theoretical Aspects of Rationality and Knowledge (TARK-2011), Groningen, The Netherlands, July 12–14, 2011*, pages 257–264. ACM
Halpern, J. Y., & Moses, Y. (1992). A guide to completeness and complexity for modal logics of knowledge and belief. *Artificial Intelligence, 54*, 319–379.
Halpern, J. Y., & Vardi, M. Y. (1989). The complexity of reasoning about knowledge and time. I. Lower bounds. *Journal of Computer and Systems Science, 38*(1), 195–237.
Henzinger, M. R., Henzinger, T. A., & Kopke, P. W. (1995). Computing simulations on finite and infinite graphs. In *FOCS '95*: *Proceedings of the 36th Annual Symposium on Foundations of Computer Science* (pp. 453–462). IEEE Computer Society Press.
Hoffmann, C. M. (1982). *Group-theoretic algorithms and graph isomorphism*, volume 136 of *Lecture Notes in Computer Science*. Berlin: Springer
Karp, R. M. (1972). Reducibility among combinatorial problems. In R. E. Miller & J. W. Thatcher (Eds.), *Complexity of Computer Computations* (pp. 85–103). New York: Plenum Press.
Köbler, J., Schöning, U., & Torán, J. (1993). *The Graph Isomorphism Problem: Its Structural Complexity*. Basel: Birkhauser Verlag.
Kooi, B., & Van Benthem, J. (2004). Reduction axioms for epistemic actions. In R. Schmidt, I. Pratt-Hartmann, M. Reynolds, & H. Wansing (Eds.), *Advances in Modal Logic 2004* (pp. 197–211). Manchester: Department of Computer Science, University of Manchester.
Kučera, A., & Mayr, R. (2002). Why is simulation harder than bisimulation?. In *CONCUR '02*: *Proceedings of the 13th International Conference on Concurrency Theory* (pp. 594–610), London, UK, 2002. Berlin: Springer.
Meijering, B., Van Rijn, H., Taatgen, N. A., & Verbrugge, R. (2012). What eye movements can tell about theory of mind in a strategic game. *PLoS ONE, 7*(9), e45961,09.
Osborne, M. J., & Rubinstein, A. (1994). *A Course in Game Theory*. Cambridge, MA: MIT Press.
Papadimitriou, C. H. (1993). *Computational Complexity*. Boston: Addison Wesley.
Papadimitriou, C. H., & Steiglitz, K. (1982). *Combinatorial Optimization: Algorithms and Complexity*. Upper Saddle River, NJ: Prentice-Hall Inc..
Plaza, J. A. (1989). Logics of public communications. In M. L. Emrich, M.S. Pfeifer, M. Hadzikadic, & Z. W. Ras (Eds.), *Proceedings of the Fourth International Symposium on Methodologies for Intelligent Systems: Poster session program* (pp. 201–216). Oak Ridge, TN: Oak Ridge National Laboratory.
Pratt-Hartmann, I., & Moss, L. S. (2009). Logics for the relational syllogistic. *The Review of Symbolic Logic, 2*(04), 647–683.
Van Benthem, J. (1983). *Modal Logic and Classical Logic*. Naples: Bibliopolis.
Van Benthem, J. (2010). *Modal Logic for Open Minds*. Number 199 in CSLI lecture notes. Stanford, CA: Center for the Study of Language and Information.
Van Benthem, J. (2011). *Logical Dynamics of Information Flow*. Cambridge, MA: Cambridge University Press.
Van Benthem, J., & Pacuit, E. (2006). The tree of knowledge in action: Towards a common perspective. In I. Hodkinson, G. Governatori, & Y. Venema, *Advances in Modal Logic* (Vol. 6). : College Publications.
Van Ditmarsch, H., & French, T. (2009). Simulation and information: Quantifying over epistemic events. In *Knowledge Representation for Agents and Multi-Agent Systems*: *First International Workshop*,

KRAMAS 2008, *Sydney*, *Australia*, *17 September 2008*, *Revised Selected Papers* (pp. 51–65), Berlin, Heidelberg, 2009. Berlin: Springer.

Van Ditmarsch, H., der Hoek, W., & Kooi, B. (2007). *Dynamic Epistemic Logic*. Dordrecht: Springer.

Van Rooij, I. (2008). The tractable cognition thesis. *Cognitive Science: A Multidisciplinary Journal, 32*(6), 939–984.

Szymanik, J. (2010). Computational complexity of polyadic lifts of generalized quantifiers in natural language. *Linguistics and Philosophy, 33*, 215–250.

Szymanik, J., & Zajenkowski, M. (2010). Comprehension of simple quantifiers. Empirical evaluation of a computational model. *Cognitive Science: A Multidisciplinary Journal, 34*(3), 521–532.

Verbrugge, R. (2009). Logic and social cognition. The facts matter, and so do computational models. *Journal of Philosophical Logic, 38*(6), 649–680.

Wang, Y. (2010). *Epistemic Modelling and Protocol Dynamics*. PhD Thesis, ILLC, University of Amsterdam. ILLC Dissertation Series DS-2010-06.

Weber, R. (2001). Behavior and learning in the "dirty faces" game. *Experimental Economics, 4*, 229–242.