

Games, Actions, and Social Software

JAN VAN EIJCK AND RINEKE VERBRUGGE (EDS.)

SEPTEMBER 29, 2010

Contents

Chapter 1. On the Logic of Lying	5
1. What is a Lie?	5
2. The Logic of Communication and Change	7
3. The Logic of Lying in Public Discourse	10
4. Liar's Dice — Game-Theoretical Analysis	18
5. Liar's Dice — Doxastic Analysis	22
6. Conclusion	32
Bibliography	35

CHAPTER 1

On the Logic of Lying

Hans van Ditmarsch,¹ Jan van Eijck,² Floor Sietsma,³ Yanjing Wang⁴

ABSTRACT. We look at lying as an act of communication, where (i) the proposition that is communicated is not true, (ii) the utterer of the lie knows (or believes) that what she communicates is not true, and (iii) the utterer of the lie intends the lie to be taken as truth. Rather than dwell on the moral issues, we provide a sketch of what goes on logically when a lie is communicated. We present a complete logic of manipulative updating, to analyse the effects of lying in public discourse. Next, we turn to the study of lying in games. First, a game-theoretical analysis is used to explain how the possibility of lying makes such games interesting, and how lying is put to use in optimal strategies for playing the game. Finally, we give a matching logical analysis. Our running example of lying in games is liar's dice.

1. What is a Lie?

The church father St. Augustine, who wrote at length about lying in *De Mendacio* [3], holds a subtle view on what lying is and what it is not. We will take his view as our point of departure. Here is his famous quote on what lying is not.

For not every one who says a false thing lies, if he believes or opines that to be true which he says. Now between believing and opining there is this difference, that sometimes he who believes feels that he does not know that which he believes, (although he may know himself to be ignorant of a thing, and yet have no doubt at all concerning it, if he most firmly believes it:) whereas he who opines, thinks he knows that which he does not know. Now whoever utters that which he holds in his mind either as belief or as opinion, even though it be false, he lies not. For this he owes to the faith of his utterance, that he thereby produce that which he holds in his mind, and has in that way in which he produces it. Not that he is without fault, although he lie not, if either he believes what he ought not to believe, or thinks he knows what he knows not, even though it should be true: for he accounts an unknown thing for a known.

¹University of Sevilla. Email hvd@us.es

²CWI, Amsterdam. Email jve@cwi.nl

³CWI, Amsterdam. Email F.Sietsma@cwi.nl

⁴Peking University. Email wangyanjing@gmail.com

Augustine, *De Mendacio* (On Lying), ca. AD 395 [3]

And on what lying is:

Wherefore, that man lies, who has one thing in his mind and utters another in words, or by signs of whatever kind. Whence also the heart of him who lies is said to be double; that is, there is a double thought: the one, of that thing which he either knows or thinks to be true and does not produce; the other, of that thing which he produces instead thereof, knowing or thinking it to be false. Whence it comes to pass, that he may say a false thing and yet not lie, if he thinks it to be so as he says although it be not so; and, that he may say a true thing, and yet lie, if he thinks it to be false and utters it for true, although in reality it be so as he utters it. For from the sense of his own mind, not from the verity or falsity of the things themselves, is he to be judged to lie or not to lie. Therefore he who utters a false thing for a true, which however he opines to be true, may be called erring and rash: but he is not rightly said to lie; because he has not a double heart when he utters it, neither does he wish to deceive, but is deceived. But the fault of him who lies, is the desire of deceiving in the uttering of his mind; whether he do deceive, in that he is believed when uttering the false thing; or whether he do not deceive, either in that he is not believed, or in that he utters a true thing with will to deceive, which he does not think to be true: wherein being believed, he does not deceive though it was his will to deceive: except that he deceives in so far as he is thought to know or think as he utters.

Augustine, [3]

We cannot do better than to follow Augustine in assuming that intent to mislead is part of the definition of a liar. Thus, to us, lying that p is communicating p in the belief that $\neg p$ is the case, with the intent to be believed.

Note that we allow ourselves to deviate from Augustine a bit: we will use a logic that treats knowledge as true belief, and we will occasionally switch between knowledge and belief, so we will occasionally view lying about p as communication of p in the knowledge that $\neg p$ is the case. Intention to be believed can be modelled as a preference relation different from the accessibility relation for belief (plausibility). This is to account for the fact that people can believe things for which they have no preference, and vice versa.

Our analysis is not intended as a contribution to epistemology. We are aware of the philosophical difficulties with the treatment of knowledge as (justified) true belief [21]. But we endorse the following quote from the Stanford Encyclopedia of Philosophy:

Most likely, there isn't one single concept of knowledge that permits consensus on what the necessary and sufficient conditions of knowledge are. [27]

The moral sides to the issue of lying are clarified in the ninth of the ten commandments ('Thou shalt not bear false witness') and the fourth of the five Buddhist precepts ('I undertake the precept to refrain from false speech'). On the other hand, in the *Analects* of Confucius, Confucius is quoted as condoning a lie if its purpose is to preserve social structure:

The Governor of She said to Confucius, 'In our village we have an example of a straight person. When the father stole a sheep, the son gave evidence against him.' Confucius answered, 'In our village those who are straight are quite different. Fathers cover up for their sons, and sons cover up for their fathers. In such behaviour is straightness to be found as a matter of course.' *Analects*, 13.18.

Among philosophical treatises, the quoted text of Augustine is a classic.⁵ For more, see [12] and [2] and the references therein.

Rather than dwell on the moral side of the issue of lying, in this paper we will study its logic, focusing on simple cases of lying in game situations, and on a particular kind of public announcement that may or may not be deceptive (we call these 'manipulative updates'). Thus, we abstract from the moral issues; we feel that it is important to understand why lying is tempting (why and how it pays off) before addressing the choice between condemnation and absolution.

The rest of the paper is structured as follows. First, in Section 2, we link up to the generic logic of communication and change. Next, in Section 3, we develop our logic of lying in public discourse, treating a lie as an update with a communication originating from someone who believes it is not true. Next, we turn to lying in games, by analyzing the game of Liar's Dice, first in terms of game theory (Section 4), next in terms of (an implementation of) our logical system (Section 5). Section 6 concludes with a reflection on a crucial difference between lying in public discourse and lying in games: lying in games does not induce false beliefs, because the communications in a bluffing game are not deceptive or manipulative in our technical sense.

2. The Logic of Communication and Change

The logic of communication and change presented in [11] provides means to model communicative actions and actions that change the world, and their effects in given epistemic situations. In this section we introduce the syntax and semantics of the logic. In the next sections we show how this machinery can be put to use to analyse manipulation in public discourse and to describe what goes on in the game of Liar's Dice. Rather than use the enhanced version based on the proposal in [17] to use propositional dynamic logic as a logic of belief revision, we stick to simple doxastic models, with plausibility relations that satisfy the KD45 axioms.

DEFINITION 1.1 (Doxastic models). *Let a set of propositional variables P and a finite set of agents N be given. A doxastic model is a triple $M = (W, V, R)$ where W is a set of worlds, $V : W \rightarrow \mathcal{P}(P)$ assigns a valuation to each world $w \in W$, and $R : N \rightarrow \mathcal{P}(W^2)$*

⁵We found a site on the internet where the views of the writer *De Mendacio* on lying are expounded.

assigns an accessibility relation \xrightarrow{i} to each agent $i \in N$, satisfying transitivity, seriality and euclidicity. (A binary relation R is euclidian if it satisfies $\forall xyz((Rxy \wedge Rxz) \rightarrow Ryz)$.)

A pair $\mathbf{M} = (M, U)$ with $U \subseteq W$ is a multiple-pointed doxastic model, indicating that the actual world is among U .

Note that in a multiple-pointed doxastic model (M, U) , U is allowed to be empty, indicating that the model pictures a doxastic situation that is incompatible with reality.

If we also want to model the *intention* to deceive, we need to use doxastic *preference* models (W, V, R, S) , where S is a second relation for preference. Again, it is reasonable to let S satisfy the KD45 postulates. But rather than carry such preference relations along in the exposition, we will indicate at appropriate places how they can be dealt with.

[7] proposes to model doxastic actions as doxastic models, with valuations replaced by preconditions. (See also: [4, 5, 6, 8, 9, 13, 18, 19, 23].) [11] proposes to add substitutions for modelling change in the world (this proposal is based on [15]). See also [14]. The story of update logic with substitutions is retold for a fragment in [22].

DEFINITION 1.2 (Substitutions). \mathcal{L} substitutions are functions of type $\mathcal{L} \rightarrow \mathcal{L}$ that distribute over all language constructs, and that map all but a finite number of basic propositions to themselves. \mathcal{L} substitutions can be represented as sets of bindings

$$\{p_1 \mapsto \phi_1, \dots, p_n \mapsto \phi_n\}$$

where all the p_i are different. If σ is a \mathcal{L} substitution, then the set $\{p \in P \mid \sigma(p) \neq p\}$ is called its domain, notation $\text{dom}(\sigma)$. Use ϵ for the identity substitution. Let $\mathbf{Sub}_{\mathcal{L}}$ be the set of all \mathcal{L} substitutions.

This notion of ‘domain’ is from the logic programming tradition (see, e.g., [1]).

DEFINITION 1.3 (Doxastic Models under a Substitution). If

$$M = (W, V, R)$$

is a doxastic model and σ is a \mathcal{L} substitution (for an appropriate doxastic language \mathcal{L}), then V_M^σ is the valuation given by $\lambda p \in P \cdot \llbracket \sigma(p) \rrbracket^M$. In other words, V_M^σ assigns to p the set of worlds w in which $\sigma(p)$ is true. For $M = (W, V, R)$, call M^σ the model given by (W, V_M^σ, R) .

Note⁶ that the functor $\sigma \mapsto [\sigma]$ given by $[\sigma] : V \mapsto V^\sigma$ is contravariant, i.e., $[\tau \circ \sigma] = [\sigma] \circ [\tau]$. Cf. also [1].

DEFINITION 1.4 (Action models for a given language \mathcal{L}). Let a finite set of agents N and a doxastic language \mathcal{L} be given. An action model for \mathcal{L} is a quadruple $A = (W, \text{pre}, \text{sub}, R)$ where

- W is a set of action states or events,
- $\text{pre} : W \rightarrow \mathcal{L}$ assigns a precondition to each action state,
- $\text{sub} : W \rightarrow \mathbf{Sub}_{\mathcal{L}}$ assigns a substitution to each action state and

⁶With thanks to one of our reviewers.

- $R : N \rightarrow \mathcal{P}(W^2)$ assigns a transitive, serial and euclidic accessibility relation \xrightarrow{i} to each agent $i \in N$.

A pair $\mathbf{A} = (A, S)$ with $S \subseteq W$ is a multiple-pointed action model, indicating that the actual event that takes place is a member of S .

The doxastic language \mathcal{L} is defined as follows.

DEFINITION 1.5 (\mathcal{L}). Assume p ranges over the set of basic propositions P , i ranges over the set of agents N . The formulas of \mathcal{L} are given by:

$$\begin{aligned}\phi & ::= \top \mid p \mid \neg\phi \mid \phi_1 \wedge \phi_2 \mid [\alpha]\phi \mid [A, S]\phi, \\ \alpha & ::= i \mid ?\phi \mid \alpha_1 \cup \alpha_2 \mid \alpha_1; \alpha_2 \mid \alpha^*,\end{aligned}$$

where (A, S) is a multiple-pointed finite \mathcal{L} (action) model.

We employ the usual abbreviations. In particular, $\phi_1 \vee \phi_2$ is shorthand for $\neg(\neg\phi_1 \wedge \neg\phi_2)$, $\phi_1 \rightarrow \phi_2$ for $\neg(\phi_1 \wedge \neg\phi_2)$, $\langle \alpha \rangle \phi$ for $\neg[\alpha]\neg\phi$, $\langle A, S \rangle \phi$ for $\neg[A, S]\neg\phi$. Note that the standard doxastic language is a sublanguage of \mathcal{L} , with “[i] ϕ ” and “[$(\cup_{i \in N} i)^*$] ϕ ” interpreted as “ $B_i\phi$ ” and “ $C\phi$ ”, respectively.⁷

Let MOD be the class of multiple-pointed doxastic models and ACT the class of multiple-pointed finite \mathcal{L} models. Then \mathcal{L} -update is an operation of the following type:

$$\otimes : \text{MOD} \times \text{ACT} \rightarrow \text{MOD}.$$

The operation \otimes and the truth definition for \mathcal{L} are defined by mutual recursion, as follows.

DEFINITION 1.6 (Update, Truth). Given a multiple-pointed doxastic model (M, U) and an action model (A, S) , we define

$$(M, U) \otimes (A, S)$$

as

$$((W', V', R'), U'),$$

where

$$\begin{aligned}W' & ::= \{(w, s) \mid w \in W_M, s \in W_A, M \models_w \text{pre}_s\}, \\ V'(w, s) & ::= \{p \in P \mid M \models_w \text{sub}_s(p)\}, \\ (w, s) \xrightarrow{i} (w', s') \in R' & ::= w \xrightarrow{i} w' \in R_M \text{ and } s \xrightarrow{i} s' \in R_A, \\ U' & ::= \{(u, s) \mid u \in U, s \in S, (u, s) \in W'\},\end{aligned}$$

⁷The reason to employ multiple-pointed models for updating is that it allows us to handle choice. Suppose we want to model the action of testing whether ϕ , followed by a public announcement of the result. This involves choice: if the outcome of the test is affirmative, then do this, else do that. Choice is modelled in a straightforward way in multiple-pointed action models. Once we allow multiple-pointed action models, it is reasonable to also take our doxastic models to be multiple-pointed, with the multiple points constraining the whereabouts of the actual world.

and where the truth definition is given by:

$$\begin{aligned}
M \models_w \top & \quad \text{always} \\
M \models_w p & \quad := p \in V_M(w) \\
M \models_w \neg\phi & \quad := \text{not } M \models_w \phi \\
M \models_w \phi_1 \wedge \phi_2 & \quad := M \models_w \phi_1 \text{ and } M \models_w \phi_2 \\
M \models_w [\alpha]\phi & \quad := \text{for all } w' \text{ with } w \xrightarrow{\alpha} w' \text{ } M \models_{w'} \phi \\
M \models_w [A, S]\phi & \quad := (W', V', R') \models_{(w,s)} \phi \text{ for all } (w, s) \in U', \\
& \quad \text{where } ((W', V', R'), U') = (M, \{w\}) \otimes (A, S).
\end{aligned}$$

In this definition \xrightarrow{i} is the doxastic relation for i in the model M , $\xrightarrow{?}\phi$ is the relation $\{(w, w) \mid M \models_w \phi\}$, $\xrightarrow{\alpha_1 \cup \alpha_2}$ is the relation $\xrightarrow{\alpha_1} \cup \xrightarrow{\alpha_2}$, $\xrightarrow{\alpha_1; \alpha_2}$ is $\{(x, y) \mid \exists z (x \xrightarrow{\alpha_1} z) \& (z \xrightarrow{\alpha_2} y)\}$, and $\xrightarrow{\alpha^*}$ is the reflexive transitive closure of $\xrightarrow{\alpha}$.

In [11] it is shown that this logic is axiomatized by the axioms and inference rules of doxastic PDL, plus appropriate axioms for the doxastic relations (in this case, the KD45 axioms for the \xrightarrow{i} relations), plus a set of reduction axioms that are generated from the action models by means of a process of program transformation. In the next section we will use (appropriate fragments of) this logical system to model what goes on in manipulative communication, and in Section 5 we will employ it to analyze what goes on in Liar's Dice.

3. The Logic of Lying in Public Discourse

We get lied to in the public domain, all the time, by people who have an interest in obfuscating the truth. In 1993 the tobacco company Philip Morris tried to discredit a report on *Respiratory Health Effects of Passive Smoking* by founding, through a hired intermediary, a fake citizen's group called *The Advancement of Sound Science* or TASSC, to cast doubt on it. Exxon-Mobile used the same organisation to spread disinformation about global warming.⁸ Their main ploy: hang the label of 'junk science' on peer-reviewed scientific papers on smoking hazards or global warming, and promote propaganda disguised as research as 'sound science'. It worked beautifully for a while, until the *New York Times* exposed the fraud [24]. As a result, many educated people who should know better are still in doubt about the reality of global warming, or think the issues are just too hard for them to understand.

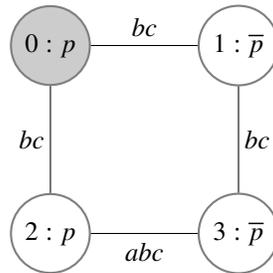
It has frequently been noted that the surest result of brainwashing in the long run is a peculiar kind of cynicism, the absolute refusal to believe in the truth of anything, no matter how well it may be established. In other words, the result of a consistent and total substitution of lies for factual truth is not that the lie will now be accepted as truth, and truth be defamed as lie, but that the sense by which we take our bearings in the real world—and the category of truth versus falsehood is among the mental means to this end—is being destroyed.

⁸See <http://www.exxonsecrets.org/html/orgfactsheet.php?id=6>.

Hannah Arendt, “Truth and Politics”, 1967 [2].

Below we will give a logic of possibly deceptive public speech acts, to model the effects of lying in politics. Logically, lying in public and lying in private are quite different, just as public announcement is quite different from private announcement. Incidentally, our modelling shows quite clearly that Arendt is right about the grave consequences of lying in politics (Proposition 1.10 below).

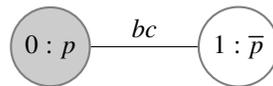
First, take the prototypical example of lying about p . Picture an initial situation where agent a knows that p , and agent a knows that agents b, c do not know that p . One way to picture this initial situation is like this:



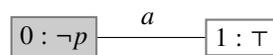
The grey shading indicates that 0 is the actual world. Two worlds are x -connected if there is an x -path from one to the other. E.g., 0 and 3 are b -connected. Note that agent a believes (or even: knows) that p , but agents b, c also consider it possible that agent a believes the opposite (which is the case in world 1), or that agent a does have no beliefs whatsoever about p (the situation in worlds 2 and 3).

In typical examples of bearing witness in court, the situation is often a bit different. In cases of providing an alibi, for example, the question ‘Was the accused at home with you during the evening of June 6th?’ is posed on the understanding that the witness is in a position to know the true answer, even if nobody can check that she is telling the truth.

Let us assume that everyone knows that a knows whether p . The picture now becomes:

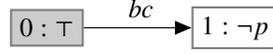


Assume agent a sends a group communication to b, c to the effect that $\neg p$. Would this be a correct communication model for the lie that $\neg p$?

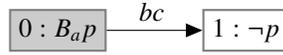


It is easy to see that this cannot be right. The result of this update is a model that has no actual worlds, i.e., an inconsistent model. The actual worlds of an update are pairs (w, e) where w is an actual world of the input doxastic model and e an actual event of the update model, and w satisfies the precondition of e . Since the actual world has p true, and the precondition of the actual action is $\neg p$, there are no such pairs.

Rather, the misleading communication should be modelled as a KD45 action model, as follows (reflexive arrows not drawn):

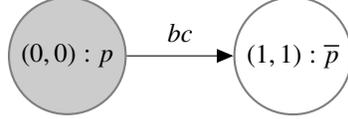


The misleading agent a knows that no truthful communication is being made, but the two agents b, c mistakenly believe that $\neg p$ is being truthfully asserted. The fact that the originator of the lie does believe that p is true can be taken on board as well, of course:



In the context of doxastic *preference* models, the precondition for the actual action could be extended even further, with the intent to mislead: in a 's most preferred worlds, his victims believe that $\neg p$. We will omit the formal details in the interest of readability.

Updating the initial model with this action model gives:

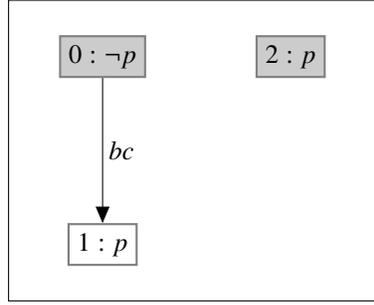


Note that the precondition $B_a p$ forces the actual event 0 to match with the actual world 0, so that the new model has an actual world $(0, 0)$. Similarly for world 1 and event 1.

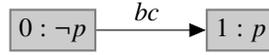
This is a model where a believes that p , where b, c mistakenly believe that $\neg p$, and where b, c also believe that a believes that $\neg p$. Note that the model is KD45: beliefs are still consistent ($[i]\phi \rightarrow \langle i \rangle \phi$ holds in the model), but the model is not truthful anymore (there are ϕ and i for which $[i]\phi \rightarrow \phi$ does not hold). The postulate of truthfulness has been replaced by the weaker postulate of consistency (the D postulate $[i]\phi \rightarrow \langle i \rangle \phi$).

This way to model lying suggests a natural generalization of the well-studied concept of a public announcement. In the logic of public announcements [25, 19], a public announcement $!\phi$ is always taken to be a *true* statement. Unfortunately, this is not completely realistic.

A more realistic version of public announcements leaves open the possibility of deceit, as follows. A possibly deceptive public announcement ϕ is a kind of ‘if then else’ action. In case ϕ is true, the announcement is a public update with ϕ , in case ϕ is false, the public is deceived into taking ϕ as true. The manipulative update with p , in a setting where the public consists of a, b, c , looks like this:



There are two actual events, one for the situation where p is true — in this case, the public is duly informed — and one for the situation where p is false — in this case the public is misled to believe that p . This action model can be simplified, as follows:



Call this the two-pointed manipulative update for ϕ . We will refer to this action model as U_ϕ . The variation on this action model where only event 0 is actual will be referred to as U_ϕ^0 . This action model denotes the lie with ϕ . The variant with only event 1 actual will be referred to as U_ϕ^1 . This action model denotes the public announcement with ϕ .

Let us introduce operations for these actions. The manipulative update with ϕ is denoted $\ddagger\phi$, and its two variants are denoted $;\phi$ (for the lie that ϕ) and $!\phi$ (for the public announcement that ϕ).

Now it turns out that the logic of individual belief and manipulative update, has a simple axiomatisation in terms of reduction axioms, just like the logic of individual knowledge and public announcement.

$$\phi ::= p \mid \neg\phi \mid \phi_1 \wedge \phi_2 \mid B_i\phi \mid [\ddagger\phi_1]\phi_2 \mid [;\phi_1]\phi_2 \mid [!\phi_1]\phi_2$$

Interpretation as sketched above:

- $[\ddagger\phi]\psi$ is true in a model M at a world w if ψ is true in both $(w, 0)$ and $(w, 1)$ of updated model $M \otimes U$.
- $[;\phi]\psi$ is true in a model M at a world w if ψ is true in $(w, 0)$ of updated model $M \otimes U^0$.
- $[!\phi]\psi$ is true in a model M at a world w if ψ is true in $(w, 1)$ of updated model $M \otimes U^1$.

Complete axiomatisation: the usual KD45 axioms for B_i , modus ponens, necessitation for B_i , $\ddagger\phi$, $;\phi$ and $!\phi$, and reduction axioms for the $[\ddagger\phi]$, $[;\phi]$, $[!\phi]$ modalities:

$$[\ddagger\phi]\psi \leftrightarrow [;\phi]\psi \wedge [!\phi]\psi$$

This defines the effect of $[\ddagger\phi]$ in terms of those of $[\!|\phi]$ and $[\!|\phi]$. Next, we have the usual reduction axioms for public announcement:

$$\begin{aligned} [\!|\phi]p &\leftrightarrow \phi \rightarrow p \\ [\!|\phi]\neg\psi &\leftrightarrow \phi \rightarrow \neg[\!|\phi]\psi \\ [\!|\phi](\psi_1 \wedge \psi_2) &\leftrightarrow [\!|\phi]\psi_1 \wedge [\!|\phi]\psi_2 \\ [\!|\phi]B_i\psi &\leftrightarrow \phi \rightarrow B_i[\!|\phi]\psi \end{aligned}$$

Finally, the reduction axioms for lying:

$$\begin{aligned} [\!|\phi]p &\leftrightarrow \neg\phi \rightarrow p \\ [\!|\phi]\neg\psi &\leftrightarrow \neg\phi \rightarrow \neg[\!|\phi]\psi \\ [\!|\phi](\psi_1 \wedge \psi_2) &\leftrightarrow [\!|\phi]\psi_1 \wedge [\!|\phi]\psi_2 \\ [\!|\phi]B_i\psi &\leftrightarrow \neg\phi \rightarrow B_i[\!|\phi]\psi \end{aligned}$$

The final axiom of this list is the most interesting: it expresses that believing ψ after a lie that ϕ amounts to the belief that a public announcement of ϕ implies ψ , conditioned by $\neg\phi$.

Since all these axioms have the form of equivalences, completeness of the calculus of manipulation and individual belief follows from a reduction argument, as in the case of public announcements with individual knowledge. We refer to [11] for a general perspective on proving communication logics complete by means of reduction axioms.

THEOREM 1.1. *The calculus of manipulation and individual belief is complete.*

Another way to see that the logic is complete is by means of the observation that this is the special case of the LCC logic where updates are restricted to manipulations, announcements and lies, and where doxastic programs are restricted to individual accessibilities. Then apply the equivalence between $[A, s]B_i\phi$ and $\text{pre}(A, s) \rightarrow \bigwedge_{\{t|s \rightarrow t\}} B_i[A, t]\phi$.

Interestingly, our logic of manipulation is closely related to the variation on public announcement that is used in [20] to analyse the ‘surprise exam puzzle’, where public announcement of ϕ is defined as an operation that restricts the doxastic alternatives of the agents to the worlds where ϕ is true, i.e., M^ϕ is the model where each R_i gets replaced by R_i^ϕ given by $R_i^\phi(w) = \llbracket \phi \rrbracket \cap R_i(w)$. Using $\dagger\phi$ for this alternative announcement, the corresponding reduction axiom is $[\dagger\phi]B_i\psi \leftrightarrow B_i(\phi \rightarrow [\dagger\phi]\psi)$. This same definition and reduction axiom can be found in [22], but neither author cites the other on this.

At first sight, this alternative semantics for announcement takes us outside of the framework sketched in Section 2 above. However, if $\dagger\phi$ is alternative announcement, then we have:

PROPOSITION 1.2. $M, w \models [\dagger\phi]\psi$ iff $M, w \models [\ddagger\phi]\psi$.

Alternative announcement turns out to be the same as manipulative updating, and our analysis can be viewed as a decomposition of alternative announcement into public lying and (regular) public announcement.

Regular public announcements can be expressed in terms of manipulative updating:

PROPOSITION 1.3. $\vdash [!\phi]\psi \leftrightarrow (\phi \rightarrow [\ddagger\phi]\psi)$.

The proof is by induction on ψ and is left to the reader.

It is the case that the logic of public announcement and the logic of manipulation have the same expressive power: this follows from the fact that they both reduce to multi-modal KD45. But note that the logic of manipulative updating has greater ‘action expressivity’ than the logic of public announcement: the logic of $[!\phi]$ has no means to express an operation mapping S5 models to KD45 models, and $[\ddagger\phi]$ is such an operation.

As an example of reasoning with the calculus, we use the axioms to show that a manipulative update followed by a belief is equivalent to a belief followed by the corresponding public announcement:

PROPOSITION 1.4. $\vdash [\ddagger\phi]B_i\psi \leftrightarrow B_i[!\phi]\psi$.

PROOF.

$$[\ddagger\phi]B_i\psi \leftrightarrow ([!_i\phi]B_i\psi \wedge [!\phi]B_i\psi) \leftrightarrow ((\neg\phi \rightarrow B_i[!\phi]\psi) \wedge (\phi \rightarrow B_i[!\phi]\psi)) \leftrightarrow B_i[!\phi]\psi.$$

□

An important difference between manipulative update and public announcement shows up when we work out the preconditions of inconsistency after an update. For announcements we get:

PROPOSITION 1.5. $\vdash [!\phi]\perp \leftrightarrow \neg\phi$.

PROOF.

$$\begin{aligned} [!\phi]\perp &\leftrightarrow [!\phi](p \wedge \neg p) \leftrightarrow ([!\phi]p \wedge [!\phi]\neg p) \leftrightarrow ([!\phi]p \wedge (\phi \rightarrow \neg[!\phi]p)) \\ &\leftrightarrow ((\phi \rightarrow p) \wedge (\phi \rightarrow \neg p)) \leftrightarrow \neg\phi \end{aligned}$$

□

This shows that a public announcement with ϕ leads to an inconsistent state iff the negation of ϕ is true. Similarly, it is easy to work out that a public lie that ϕ leads to an inconsistency iff ϕ is true, i.e., we can derive

PROPOSITION 1.6. $\vdash [!_i\phi]\perp \leftrightarrow \phi$.

Using this we can work out the preconditions for inconsistency after a manipulative update:

PROPOSITION 1.7. $\vdash [\ddagger\phi]\perp \leftrightarrow \perp$.

PROOF.

$$[\ddagger\phi]\perp \leftrightarrow ([!\phi]\perp \wedge [!_i\phi]\perp) \stackrel{\text{Prop 1.6}}{\leftrightarrow} (\neg\phi \wedge \phi) \leftrightarrow \perp$$

□

This means that a manipulative update in a consistent state will *never* lead to inconsistency.

The following proposition about public announcements can be proved by induction on ϕ . It shows that if we update with an inconsistency, the resulting model is inconsistent:

PROPOSITION 1.8. $\vdash [!\perp]\phi \leftrightarrow \top$.

In the case of manipulatively updating with an inconsistency, the result is not an inconsistent model, but a model where all accessibilities have vanished. In the particular case of an i -belief, we get:

PROPOSITION 1.9. $\vdash [\ddagger\perp]B_i\phi \leftrightarrow \top$.

PROOF.

$$[\ddagger\perp]B_i\phi \leftrightarrow ([!\perp]B_i\phi \wedge [;\perp]B_i\phi) \leftrightarrow (\top \wedge B_i[!\perp]\phi) \leftrightarrow B_i[!\perp]\phi \stackrel{\text{Prop 1.8}}{\leftrightarrow} B_i\top \leftrightarrow \top.$$

□

After a manipulative update with an inconsistency, the public will no longer be able to distinguish what is false from what is true.

Finally, the following proposition spells out under what conditions our ‘sense by which we take our bearings in the real world’ is destroyed. This happens exactly when we are manipulated into accepting as truth what flatly contradicts our firm belief:

PROPOSITION 1.10. $\vdash [\ddagger\phi]B_i\perp \leftrightarrow B_i\neg\phi$.

PROOF.

$$\begin{aligned} [\ddagger\phi]B_i\perp &\leftrightarrow ([!\phi]B_i\perp \wedge [;\phi]B_i\perp) \leftrightarrow ((\phi \rightarrow B_i[!\phi]\perp) \wedge (\neg\phi \rightarrow B_i[!\phi]\perp)) \\ &\leftrightarrow ((\phi \rightarrow B_i\neg\phi) \wedge (\neg\phi \rightarrow B_i\neg\phi)) \leftrightarrow B_i\neg\phi. \end{aligned}$$

□

The full logic of manipulative updating, to which we now turn, consists of doxastic PDL extended with manipulative updates, lies and announcements:

$$\begin{aligned} \alpha &::= i \mid ?\phi \mid \alpha_1; \alpha_2 \mid \alpha_1 \cup \alpha_2 \mid \alpha^* \\ \phi &::= p \mid \neg\phi \mid \phi_1 \wedge \phi_2 \mid [\alpha]\phi \mid [\ddagger\phi_1]\phi_2 \mid [;\phi_1]\phi_2 \mid [!\phi_1]\phi_2 \end{aligned}$$

Again, there is a complete axiomatisation: the axioms and rules of PDL, the axioms of KD45, necessitation for $[\ddagger\phi]$, $[;\phi]$, $[!\phi]$, and the following reduction axioms for the three update modalities.

The definition of \ddagger in terms of $;$ and $!$ is as before:

$$[\ddagger\phi]\psi \leftrightarrow [;\phi]\psi \wedge [!\phi]\psi$$

Reduction axioms for public announcement are as follows:

$$\begin{aligned}
[!\phi]p &\leftrightarrow \phi \rightarrow p \\
[!\phi]\neg\psi &\leftrightarrow \phi \rightarrow \neg[!\phi]\psi \\
[!\phi](\psi_1 \wedge \psi_2) &\leftrightarrow [!\phi]\psi_1 \wedge [!\phi]\psi_2 \\
[!\phi][i]\psi &\leftrightarrow [?\phi; i][!\phi]\psi \\
[!\phi][?\chi]\psi &\leftrightarrow [?\phi; ?\chi][!\phi]\psi \\
[!\phi][\alpha_1; \alpha_2]\psi &\leftrightarrow [!\phi][\alpha_1][\alpha_2]\psi \\
[!\phi][\alpha_1 \cup \alpha_2]\psi &\leftrightarrow [!\phi](\alpha_1\psi \wedge \alpha_2\psi) \\
[!\phi][\alpha^*]\psi &\leftrightarrow [\alpha'^*][!\phi]\psi \\
&\text{where } \alpha' \text{ such that } [!\phi][\alpha]\psi \leftrightarrow [\alpha'][!\phi]\psi
\end{aligned}$$

It can be shown by an inductive argument that for every doxastic program α , every announcement $!\phi$, and every postcondition ψ a doxastic program α' exists such that $[!\phi][\alpha]\psi \leftrightarrow [\alpha'][!\phi]\psi$. This α' , which does not have to be unique, can be found by applying the above reduction axioms.

Reduction axioms for public lies:

$$\begin{aligned}
[_i\phi]p &\leftrightarrow \neg\phi \rightarrow p \\
[_i\phi]\neg\psi &\leftrightarrow \neg\phi \rightarrow \neg[_i\phi]\psi \\
[_i\phi](\psi_1 \wedge \psi_2) &\leftrightarrow [_i\phi]\psi_1 \wedge [_i\phi]\psi_2 \\
[_i\phi][i]\psi &\leftrightarrow [?\neg\phi; i][!\phi]\psi \\
[_i\phi][?\chi]\psi &\leftrightarrow [?\neg\phi; ?\chi][!\phi]\psi \\
[_i\phi][\alpha_1; \alpha_2]\psi &\leftrightarrow [_i\phi][\alpha_1][\alpha_2]\psi \\
[_i\phi][\alpha_1 \cup \alpha_2]\psi &\leftrightarrow [_i\phi](\alpha_1\psi \wedge \alpha_2\psi) \\
[_i\phi][\alpha^*]\psi &\leftrightarrow [\alpha'; \alpha''^*][!\phi]\psi \\
&\text{where } \alpha' \text{ such that } [_i\phi][\alpha]\psi \leftrightarrow [\alpha'][!\phi]\psi \\
&\text{and } \alpha'' \text{ such that } [!\phi][\alpha]\psi \leftrightarrow [\alpha'']![!\phi]\psi
\end{aligned}$$

Again, it can be shown by an inductive argument that for every doxastic program α , every lie $_i\phi$, and every postcondition ψ , a doxastic program α' exists such that $[_i\phi][\alpha]\psi \leftrightarrow [\alpha'][!\phi]\psi$.

The α' and α'' in the axioms for α^* can be viewed as the transformed versions of the programs α , where the update operator acts as a doxastic program transformer. To give an example, suppose $\alpha = i \cup j$, and we want to calculate the way common belief of i and j is transformed by a public lie that ϕ . Then the transformed program for $i \cup j$ becomes $?\neg\phi; i \cup j$, i.e., we have:

$$[_i\phi][i \cup j]\psi \leftrightarrow [?\neg\phi; i \cup j][!\phi]\psi.$$

Similarly for the way common belief of i and j is transformed by a public announcement: the transformed program for $i \cup j$ becomes $?\phi; i \cup j$, and we have:

$$[!\phi][i \cup j]\psi \leftrightarrow [?\phi; i \cup j][!\phi]\psi.$$

Using these transformed programs, we see that the reduction axiom for $(i \cup j)^*$ takes the shape:

$$[! \phi][!(i \cup j)^*] \psi \leftrightarrow [? \neg \phi; i \cup j; (? \phi; i \cup j)^*][! \phi] \psi.$$

This expresses that after a lie with ϕ , i and j have a common belief that ψ iff in the model before the lie it holds that along all $i \cup j$ paths that start from a $\neg \phi$ world and that pass only through ϕ worlds, $[! \phi] \psi$ is true. Note that this is a ‘relativized common belief’ similar to the relativized common knowledge that is needed to get a reduction style analysis going of public announcement in the presence of common knowledge.

In fact, the style of axiomatisation that we have adopted is borrowed from the reduction axioms formulated in terms of program transformations, in [11]. In the same manner as in [11] we can derive:

THEOREM 1.11. *The calculus of manipulative updating is complete.*

In this section we have investigated the effect of lying in public discourse. In such a setting the agents usually assume the other agents tell the truth and in the event of a lie, the agents hearing the lie do not believe that the agent who announces the lie is actually lying. This causes them to believe a false thing. In Section 5 we will analyse lying in a different setting, where the agents are playing a game of Liar’s Dice and following a game strategy. But first, we will give a game-theoretical analysis of the game to see how lying affects a game’s outcome.

4. Liar’s Dice — Game-Theoretical Analysis

In his later years as a saint, Augustine held the opinion that lying, even in jest, is wrong, but as the young and playful sinner that he was before his turn to seriousness he may well have enjoyed an occasional game of dice. We will examine a simplified version of two-person Liar’s Dice, and show by means of a game-theoretical analysis that it is precisely the possibility of lying — using private information in order to mislead an opponent — that makes the game interesting.

In our simplified version of Liar’s Dice, the die is replaced by a coin. A typical move of the game is tossing a coin and inspecting the result while keeping it hidden from the other player. Here is a description of what goes on, and what the options of the two players are.

- Players 1 and 2 both stake one euro: Player 1 bets on heads, Player 2 bets on tails.
- Player 1 tosses a coin and observes the outcome (heads or tails), while keeping it concealed from player 2.
- Player 1 *announces* either \ddagger Head or \ddagger Tail.
- If 1 announces \ddagger Tail, she then simply loses her one euro to player 2 and game ends (for 1 bets on heads, so she announces defeat).
- If 1 announces \ddagger Head, she adds one euro to the stake and game continues.
- In response to \ddagger Head, 2 either *passes* (gives up) or *challenges* (“I don’t believe that, you liar”) and adds 1 euro to the stake.
- If 2 passes, 1 wins the stake, and the game ends.

- If 2 challenges, and the toss was heads, 1 wins the stake, otherwise 2 wins the stake. The game ends.

Player 1 has two information states: *Heads* and *Tails*, while player 2 has a single information state, for player 2 cannot distinguish the two possible outcomes of the toss. We will give a game theoretic analysis of how player 1 can exploit her ‘information advantage’ to the utmost, and of how player 2 can react to minimize her losses, on the assumption that the procedure is repeated a large number of times. The following picture gives the extensive game form. The first move is made by Chance; this move gives the outcome of the coin toss. Then player 1 reacts, letting her move depend on the toss outcome. Finally, player 2 decides whether to pass or challenge. This decision does not depend on the coin toss; player 2 cannot distinguish the state where 1 announced \ddagger *Head* on seeing heads from the state where she is bluffing. In the picture of the extensive game form (Figure 1) this is expressed by a dotted line.

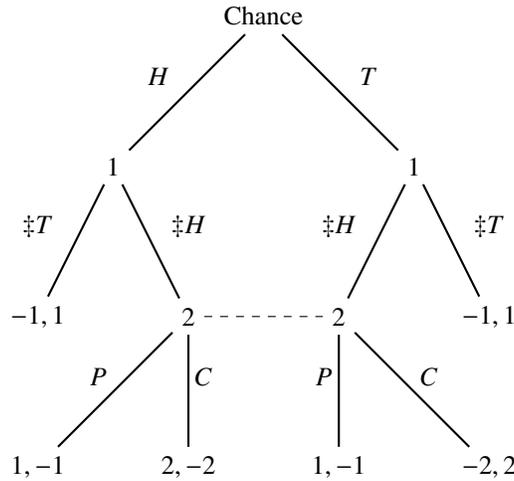


FIGURE 1. Extensive game form for Liar’s Toss game.

The leaves of the game tree indicate the payoffs. If the game sequence is *Heads*, \ddagger *Tail*, the payoffs are -1 euro for player 1 and 1 euro for player 2. The same for the sequence *Tails*, \ddagger *Tail*. Player 1 gets 1 euro and player 2 gets -1 euro for the sequences *Heads*, \ddagger *Head*, *Pass*, and *Tail*, \ddagger *Head*, *Pass* (these are the sequences where 2 gives up). The sequence *Heads*, \ddagger *Head*, *Challenge* is a win for player 1, with payoff 2 euros, and -2 euros for player 2. The sequence *Tails*, \ddagger *Head*, *Challenge*, finally, is a win for player 2, with payoff 2 euros, and -2 euros loss for player 1.

Player 1 has four strategies: $(\ddagger$ *Head*, \ddagger *Head*) (\ddagger *Head* in case of heads and in case of tails), $(\ddagger$ *Head*, \ddagger *Tail*) (\ddagger *Head* in case of heads, \ddagger *Tail* in case of tails), $(\ddagger$ *Tail*, \ddagger *Head*), and $(\ddagger$ *Tail*, \ddagger *Tail*). Player 2 has two strategies: *Pass* and *Challenge*. To find the strategic game form, one has to take the average of the expected payoffs for the two cases of *heads*

and *tails*. E.g., if player 1 plays (\dagger Head, \dagger Tail) and player 2 responds with *Challenge*, then in the long run in $\frac{1}{2}$ of the cases the outcome will be heads, and player 1 wins 2 euros, and in $\frac{1}{2}$ of the cases the outcome will be tails, and player 1 loses 1 euro (for her strategy is just to give up in such cases). Thus, the expected payoff is $\frac{1}{2} \times 2 - \frac{1}{2} \times 1 = \frac{1}{2}$ euro for player 1, and because the game is zero sum, $-\frac{1}{2}$ euro for player 2. The strategic game form is given by:

	Pass	Challenge
\dagger Head, \dagger Head	1,-1	0,0
\dagger Head, \dagger Tail	0,0	$\frac{1}{2}, -\frac{1}{2}$
\dagger Tail, \dagger Head	0,0	$-\frac{3}{2}, \frac{3}{2}$
\dagger Tail, \dagger Tail	-1,1	-1,1

It is easy to see that there is no pure strategy Nash equilibrium (a Nash equilibrium is a combination of strategies, one for each player, with the property that neither of the players can improve their payoff by unilaterally deviating from her strategy). Clearly, none of the eight strategy pairs has this property.

Now let's consider the strategy (\dagger Tail, \dagger Tail) for 1. This is the strategy of the doomed loser: even when the toss is heads the player still announces \dagger Tail. This is obviously *not* the best thing that 1 can do. *Always* announcing \dagger Head gives a much better payoff in the long run. In other words, the strategy (\dagger Tail, \dagger Tail) is strictly dominated by (\dagger Head, \dagger Head). Similar for the strategy of the unconditional liar: (\dagger Tail, \dagger Head). It is also strictly dominated by the strategy (\dagger Head, \dagger Head). Thus, we are left with:

	Pass	Challenge
\dagger Head, \dagger Head	1,-1	0,0
\dagger Head, \dagger Tail	0,0	$\frac{1}{2}, -\frac{1}{2}$

Suppose 1 plays (\dagger Head, \dagger Head) with probability p and (\dagger Head, \dagger Tail) with probability $1 - p$. Then her expected value is p for her first strategy, and $\frac{1}{2}(1 - p)$ for her second strategy. Any choice of p where the expected payoff for p is different from that for $1 - p$ can be exploited by the other player. Therefore, player 1 should play her first strategy with probability $p = \frac{1}{2}(1 - p)$, i.e., $p = \frac{1}{3}$, and her second strategy with probability $1 - p = \frac{2}{3}$. For player 2, we can reason similarly. Suppose 2 plays *Pass* with probability q and *Challenge* with probability $1 - q$. Again, the expected values for q and $1 - q$ should be the same, for otherwise this mixed strategy can be exploited by the other player. The expected value is $-q$ for her first strategy and $-\frac{1}{2}(1 - q)$ for her second strategy. Thus, she should play her first strategy with probability $q = \frac{1}{2}(1 - q)$, i.e., $q = \frac{1}{3}$. Neither player can improve on her payoff by unilateral deviation from these strategies, so the mixed strategy where 1 plays (\dagger Head, \dagger Head) in $\frac{1}{3}$ of the cases and 2 plays *Pass* in $\frac{1}{3}$ of the cases is a Nash equilibrium. In other words, the best that player 1 can do is always announcing the truth and raise the stakes when her toss is heads, and lying in one third of the cases when her toss is tails, and 2's best response to this is to Pass in one third of all cases and Challenge two thirds of the time.

The game-theoretic analysis yields that lying pays off for player 1, and that player 2, knowing this, may reasonably expect to catch player 1 on a lie in one sixth of all cases.

The value of the game is $\frac{1}{3}$ euro, and the solution is $\frac{1}{3} (\ddagger\text{Head}, \ddagger\text{Head}), \frac{2}{3} (\ddagger\text{Head}, \ddagger\text{Tail})$ as player 1's optimal strategy, and $\frac{1}{3}$ Pass, $\frac{2}{3}$ Challenge as player 2's optimal strategy. It is clear that the honest strategy $(\ddagger\text{Head}, \ddagger\text{Tail})$ is not the optimal one for player 1: given player 2 plays $\frac{1}{3}$ Pass and $\frac{2}{3}$ Challenge, the expected payoff for player 1 is only $\frac{1}{6}$ if she sticks to the honest strategy. Lying indeed pays off sometimes.

If we modify the game so that player 1 cannot lie anymore, by refusing her the privilege of having a peek at the toss outcome, the game immediately becomes a lot less interesting. In the extensive game form for this version, an extra dotted line indicates that player 1 cannot distinguish the outcome *Heads* from the outcome *Tails*. See Figure 2.

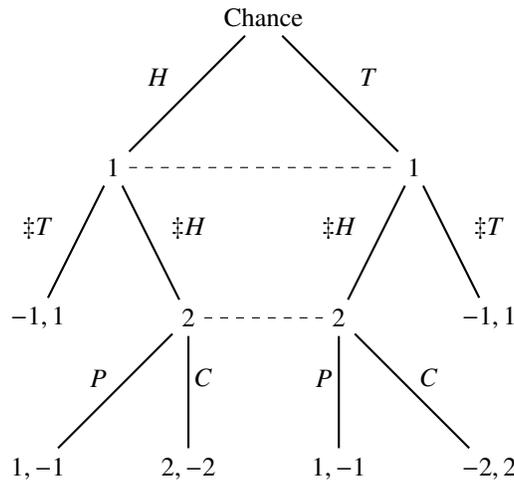


FIGURE 2. Modified game where player 1 has no information advantage.

Player 1 has just two strategies left, $\ddagger\text{Head}$ and $\ddagger\text{Tail}$, and the strategic form of the game becomes:

	Pass	Challenge
$\ddagger\text{Head}$	1,-1	0,0
$\ddagger\text{Tail}$	-1,1	-1,1

The strategy $\ddagger\text{Tail}$ for player 1 is weakly dominated by $\ddagger\text{Head}$, so it can be eliminated, and we are left with:

	Pass	Challenge
$\ddagger\text{Head}$	1,-1	0,0

The strategy pair $(\ddagger\text{Head}, \text{Challenge})$ is a Nash equilibrium. The game-theoretic analysis predicts that a rational player 1 will always play $\ddagger\text{Head}$, and a rational player 2 will always Challenge, and the game becomes a pure zero-sum game of chance. Surely, it is the possibility of lying that makes Liar's Dice an interesting game.

5. Liar's Dice — Doxastic Analysis

In the game of Liar's Dice, when player a announces *Heads* while she actually saw that the outcome of the toss was *Tails*, she is announcing something which she believes to be false with the intent to be believed. This certainly seems to be a lie. However, we usually do not condemn people who tell such a lie in a game as untruthful. In fact, in this game player a is supposed to lie sometimes, or she would never win. This is an important point: player a *intends* player b to believe her, but she probably does not expect it, because player b may very well expect player a to lie sometimes. In this section we will analyse the game of Liar's Dice from a doxastic viewpoint in order to answer the question: is lying really lying, when one is actually *supposed* to lie?

For our analysis we will use the doxastic model checker DEMO [16]. Using DEMO, we can automatically check the truth of formulas in a doxastic model. One of the authors, Floor Sietsma, has extended DEMO with factual changes to allow action models with substitutions and also with the possibility to store integer values (in the Bachelor's Thesis [26] dating from 2007). We will use this extended model checker. The code of this model checker is available from

<http://www.cwi.nl/~sietsma/DEMO/>.

We show how the game of Liar's Dice can be modelled using DEMO, and we demonstrate the doxastic models that we get if we trace a particular run of the game.

First we will closely examine the different actions that take place in the game and their representations as action models. Let p represent the value of a penny, with 1 signifying heads, and 0 signifying tails. Let agents a and b represent the two players, and let C_1 represent the contents of the purse of player a (C for cash), and C_2 that of player b , with natural number values representing the amounts in euros that each player has in her purse. These natural number registers are available in the new extension of DEMO. Let S_1, S_2 represent the money at stake for each player. Factual change can be thought of as assignment of new values to variables. This is an essential ingredient of the various actions in the game:

Initialisation: Both players put one euro at stake, and they both know this. $S_1 := 1, C_1 := C_1 - 1, S_2 := 1, C_2 := C_2 - 1$, together with public announcement of these factual changes.

Heads: Factual change of the propositional value of a penny p to 1, with private communication of the result to player a ($p = 1$ signifies heads).

Tails: Factual change of the propositional value of a penny p to 0, with private communication of the result to player a . ($p = 0$ signifies tails).

Announce: Player a announces either \ddagger *Head* or \ddagger *Tail*. There are several ways to model this and we will come back to this later.

Pass: Player b passes and loses, player a gets the stakes. $C_1 := C_1 + S_1 + S_2, S_1 := 0, S_2 := 0$.

Challenge: Public setting of $C_2 := C_2 - 1, S_2 := S_2 + 1$, followed by public announcement of the value of p . If the outcome is p then $C_1 := C_1 + S_1 + S_2$, otherwise $C_2 := C_2 + S_1 + S_2$ and in any case $S_1 := 0, S_2 := 0$.

We will show how these actions can be defined as doxastic action models in Haskell code using DEMO.

```
module Lies
where
import DEMOFCR
```

We first define the cash and stakes of each player as integer registers.

```
c1, c2, s1, s2 :: Reg
c1 = (Rg 1); c2 = (Rg 2)
s1 = (Rg 3); s2 = (Rg 4)
```

This declares four integer registers, and gives them appropriate names. The initial contents of the purses of the two players must also be defined. Let's assume both players have five euros in cash to start with.

```
initCash1, initCash2 :: Integer
initCash1 = 5
initCash2 = 5
```

Initialisation of the game: both players put one euro at stake. This is modelled by the following factual change: $S_1 := 1, C_1 := C_1 - 1, S_2 := 1, C_2 := C_2 - 1$. The representation of this in our modelling language is straightforward. We just represent the contents of the registers at startup.

```
initGame :: EM
initGame = (Mo
            [0]
            [(0, ([], [(s1, 1), (s2, 1),
                    (c1, (initCash1-1)), (c2, (initCash2-1))]))]
            [a, b]
            [(a, 0, 0), (b, 0, 0)]
            [0])
```

Tossing the penny is a factual change of p to 0 or 1. The penny is tossed secretly and before player a looks both players don't know the value of the penny. Because of this there are two worlds, one where p is set to 0 and one where p is set to 1, and neither of the two players can distinguish these worlds.

```

toss :: Integer -> FAM
toss c ags = (Mo
  [0,1]
  [(0, (Top, [(P 0, Neg (Top))], [])),
   (1, (Top, [(P 0, Top)], []))])
  ags
  [(ag, w, w') | w <- [0,1],
                 w' <- [0,1], ag <- ags]
  [c])

```

Note that the action model has a list that assigns to each world a triple consisting of a precondition, a change to the propositions, and a change to the registers. In world 0, the precondition is \top and the change is to set p to value $\neg\top$, i.e., \perp (and there is no change to the registers), and in world 1, the precondition is again \top and the change is to set p to value \top (and again, there is no change to the registers).

After the penny is tossed player a looks under the cup without showing the penny to player b . We define a generic function for computing the model of the action where a group of agents looks under the cup. These models consist of two worlds, one where p is true (heads) and one where p is false (tails), the agents in the group can distinguish these two worlds and the other agents cannot.

```

look :: [Agent] -> FAM
look group ags = (Mo
  [0,1]
  [(0, (p, [], [])), (1, (Neg(p), [], []))])
  ags
  [(ag, w, w') | w <- [0,1], w' <- [0,1],
                 ag <- (ags \\ group)] ++
  [(ag, w, w) | w <- [0,1], ag <- group])
  [0,1])

```

In this case, there are no changes to propositions or registers, but world 0 has precondition p , and world 1 has precondition $\neg p$.

Now we define the models of the situation after the penny has been tossed and player a has looked at the outcome, distinguishing the two outcomes of the toss:

```

headsg :: EM
headsg = upd (upd initGame (toss 1)) (look [a])

tailsg :: EM
tailsg = upd (upd initGame (toss 0)) (look [a])

```

Before looking at the way to model the announcement of an outcome of the toss by player a we will first define the action models for passing and challenging.

When player b passes, the stakes are added to player a 's cash: $C_2 := C_2 + S_1 + S_1, S_1 := 0, S_2 := 0$. Player b never gets to see the actual value of the penny so there are no changes in the knowledge of the agents about p . The model for this has only one world that indicates the changes in the stakes and cash.

```

pass :: FAM
pass ags = (Mo
  [0]
  [(0, (Top, [],
    [(s1, (I 0)),
     (s2, (I 0)),
     (c1, ASum [Reg c1, Reg s1, Reg s2])]))])
  ags
  [(ag, 0, 0) | ag <- ags]
  [0])

```

Note that here for the first time we see changes to the registers.

When player b decides to challenge player a , the cup is lifted and both players get to know the value of p . Then the stakes are added to the cash of player a in case of heads and player b in case of tails, together with one extra euro from the cash of player b that player b added to the stakes while challenging player a . So instead of $S_2 := S_2 + 1, C_2 := C_2 - 1$ and after that $C_1 := C_1 + S_1 + S_2$ in case of heads and $C_2 := C_2 + S_1 + S_2$ in case of tails, we use $C_1 := C_1 + S_1 + S_2 + 1, C_2 := C_2 - 1$ in case of heads and $C_2 := C_2 + S_1 + S_2$ in case of tails. The action model for this has one world for the case of heads and one world for the case of tails. Both players can distinguish these worlds because the cup was lifted, and the stakes are divided differently in the two worlds.

```

challenge :: FAM
challenge ags =
  Mo
  [0,1]
  [(0, (Neg(p), [],
        [(s1, (I 0)),
         (s2, (I 0)),
         (c2, ASum [Reg c2, Reg s1, Reg s2])])),
   (1, ( p , [],
        [(s1, (I 0)),
         (s2, (I 0)),
         (c2, ASum [Reg c2, I (-1)]),
         (c1, ASum [Reg c1, Reg s1, Reg s2, I 1])])))]
  ags
  [(ag,w,w) | w <- [0,1], ag <- ags]
  [0,1]

```

When player a announces $\ddagger\text{Head}$ or $\ddagger\text{Tail}$ the stakes change. In case of $\ddagger\text{Head}$ $C_1 := C_1 - 1, S_1 := S_1 + 1$ and in case of $\ddagger\text{Tail}$ $C_2 := C_2 + S_1 + S_2, S_1 := 0, S_2 := 0$.

```

announceStakes :: Integer -> FAM
announceStakes 0 ags =
  Mo
  [0]
  [(0, (Top, [], [(s1, (I 0)),
                  (s2, (I 0)),
                  (c2, ASum [Reg c2, Reg s1, Reg s2])])))]
  ags
  [(ag,0,0) | ag <- ags]
  [0]
announceStakes 1 ags =
  Mo
  [0]
  [(0, (Top, [], [(s1, ASum [Reg s1, I 1]),
                  (c1, ASum [Reg c1, I (-1)])])))]
  ags
  [(ag,0,0) | ag <- ags]
  [0]

```

Now the only thing we have to decide is how we will model the announcement of $\ddagger\text{Head}$ or $\ddagger\text{Tail}$. Suppose we would use the manipulative update $\ddagger p$ or $\ddagger \neg p$ for this. This would imply that the other player believes the claims that are made. However, in a real game of Liar's Dice player b knows that player a might very well be bluffing and she doesn't

really believe player a 's claim at all. So to correctly model the game we should not use the manipulative update. When player a makes an announcement this doesn't even change player b 's knowledge and beliefs because player b doesn't believe player a .

So instead of the manipulative update we should only use the model for changing the stakes to model the announcement:

```
announce :: Integer -> FAM
announce = announceStakes
```

Now player b doesn't know whether p is true but she knows she doesn't know:

```
bKnows :: Form
bKnows = Disj [(K b (Neg p)), (K b p)]
```

```
Lies> isTrue (upd tailsg (announce 0)) bKnows
False
Lies> isTrue (upd tailsg (announce 0)) (K b (Neg bKnows))
True
Lies> isTrue (upd headsg (announce 0)) bKnows
False
Lies> isTrue (upd headsg (announce 0)) (K b (Neg bKnows))
True
Lies> isTrue (upd tailsg (announce 1)) bKnows
False
Lies> isTrue (upd tailsg (announce 1)) (K b (Neg bKnows))
True
Lies> isTrue (upd headsg (announce 1)) bKnows
False
Lies> isTrue (upd headsg (announce 1)) (K b (Neg bKnows))
True
```

Note that since we did not use the manipulative update to model player a 's announcement (although it is easy to implement in DEMO, of course) the resulting models are still S5-models.

```
Lies> isS5Model (upd headsg (announce 1))
True
Lies> isS5Model (upd headsg (announce 0))
True
Lies> isS5Model (upd tailsg (announce 1))
True
Lies> isS5Model (upd tailsg (announce 0))
True
```

This means that no actual misleading is taking place at all! This is actually very plausible because player *b* knows that player *a*'s announcement might very well be false. This shows that lying only creates false belief if the person who lies is believed to be telling the truth.

Now we can use these action models to do a doxastic analysis of a game of Liar's Dice. The different possible games are:

- (1) Player *a* tosses tails and announces $\ddagger Tail$
- (2) Player *a* tosses heads and announces $\ddagger Tail$
- (3) Player *a* tosses tails and announces $\ddagger Head$ and player *b* passes
- (4) Player *a* tosses tails and announces $\ddagger Head$ and player *b* challenges
- (5) Player *a* tosses heads and announces $\ddagger Head$ and player *b* passes
- (6) Player *a* tosses heads and announces $\ddagger Head$ and player *b* challenges

The models for these games are:

```
game1, game2, game3, game4, game5, game6 :: EM
game1 = gsm (upd tailsg (announce 0))
game2 = gsm (upd headsg (announce 0))
game3 = gsm (upd (upd tailsg (announce 1)) pass)
game4 = gsm (upd (upd tailsg (announce 1)) challenge)
game5 = gsm (upd (upd headsg (announce 1)) pass)
game6 = gsm (upd (upd headsg (announce 1)) challenge)
```

We will now consider these six different cases in turn.

Game 1 is the game where player 1 tosses tails and admits this.

In this case both players stake one euro and player *b* wins the stakes, so in the end player *a* lost one euro and player *b* won one euro. This can be checked with DEMO:

```
Lies> isTrue game1 (eq (Reg c1) (ASum [I initCash1, I (-1)]))
True
Lies> isTrue game1 (eq (Reg c2) (ASum [I initCash2, I 1]))
True
```

Player *b* doesn't get to know what the value of the penny was:

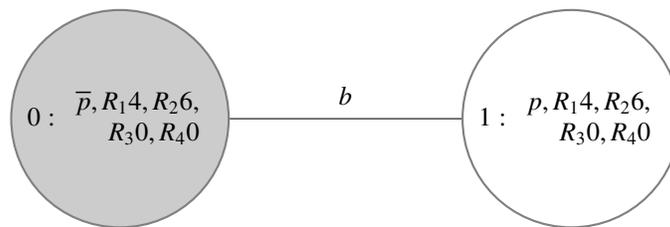
```
Lies> isTrue game1 bKnows
False
```

The model for game 1 is:

```
Lies> game1
Mo
[0,1]
[(0, ([], [(R1, 4), (R2, 6), (R3, 0), (R4, 0)])),
 (1, ([p], [(R1, 4), (R2, 6), (R3, 0), (R4, 0)]))]
```

```
[a,b]
[(a,0,0),(a,1,1),(b,0,0),(b,0,1),(b,1,0),(b,1,1)]
[0]
```

A picture of this model is below. There are two worlds, one where the toss was heads and one where it was tails. Player a can distinguish these worlds, player b cannot because player b never got to see the penny. In both worlds the cash of player a is 4 and that of player b is 6 euros, because the division of the stakes doesn't depend on the value of the penny.



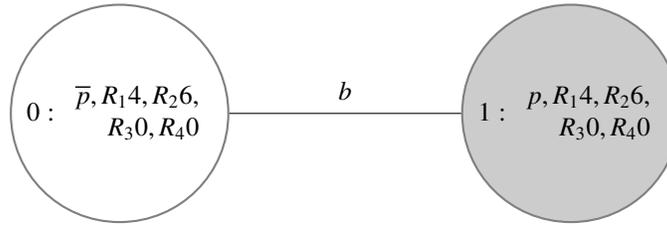
Game 2 is the game where player a falsely announces \ddagger Head. Just like in game 1, player a loses one euro and player b wins one euro, and player b doesn't get to know the value of the penny.

```
Lies> isTrue game2 (eq (Reg c1) (ASum [I initCash1,I (-1)]))
True
Lies> isTrue game2 (eq (Reg c2) (ASum [I initCash2,I 1]))
True
Lies> isTrue game2 bKnows
False
```

The model for this game is almost the same as for game 1: the difference is that now the world where p is true is actual instead of the world where p is false.

```
Lies> game2
Mo
[0,1]
[(0,([],[R1,4],[R2,6],[R3,0],[R4,0])),
 (1,([p],[R1,4],[R2,6],[R3,0],[R4,0]))]
[a,b]
[(a,0,0),(a,1,1),(b,0,0),(b,0,1),(b,1,0),(b,1,1)]
[1]
```

The picture of this model is:



The third game is the case where player a tosses tails but falsely announces \ddagger Head and player b passes. In this case player a stakes two euros and player b stakes one euro, and player a gets to keep the stakes, so the final payoff is that player a wins one euro and player b loses one euro:

```
Lies> isTrue game3 (eq (Reg c1) (ASum [I initCash1, I 1]))
True
Lies> isTrue game3 (eq (Reg c2) (ASum [I initCash1, I (-1)]))
True
```

Player b passes, so the cup is never lifted and player b doesn't know the value of the penny:

```
Lies> isTrue game3 bKnows
False
```

The model for this game is:

```
Lies> game3
Mo
[0, 1]
[(0, ([], [(R1, 6), (R2, 4), (R3, 0), (R4, 0)])),
 (1, ([p], [(R1, 6), (R2, 4), (R3, 0), (R4, 0)]))]
[a, b]
[(a, 0, 0), (a, 1, 1), (b, 0, 0), (b, 0, 1), (b, 1, 0), (b, 1, 1)]
[0]
```

This model has the same two worlds as the models for game 1 and 2 except for the changes in the player's cash.

In the fourth game, player a tosses tails but falsely announces \ddagger Head and player b challenges player a . This means that both players stake one extra euro and then the cup is lifted and player b gets the stakes.

In this case player b does know the value of the penny:

```
Lies> isTrue game4 bKnows
True
```

The payoffs are -2 euros for player a and 2 euros for player b :

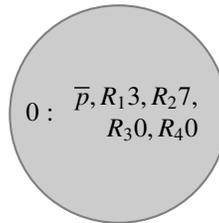
```
Lies> isTrue game4 (eq (Reg c1) (ASum [I initCash1, I (-2)]))
True
```

```
Lies> isTrue game4 (eq (Reg c2) (ASum [I initCash2,I 2]))
True
```

The model for this game is:

```
Lies> game4
Mo
[0]
[(0, ([], [(R1, 3), (R2, 7), (R3, 0), (R4, 0)]))]
[a, b]
[(a, 0, 0), (b, 0, 0)]
[0]
```

This model has only one world because none of the players consider any other world possible, because both players know the values of the penny. In this world p is false (because the toss was tails), player a 's cash is 3 euros and player b 's cash is 7 euros. A picture of this model is below.



The fifth game is the game where player a tosses heads and truthfully announces this and player b passes. In this case the cup isn't lifted so player b doesn't know the value of the penny again:

```
Lies> isTrue game5 bKnows
False
```

The payoffs are 1 for player a and -1 for player b :

```
Lies> isTrue game5 (eq (Reg c1) (ASum [I initCash1,I 1]))
True
Lies> isTrue game5 (eq (Reg c2) (ASum [I initCash2,I (-1)]))
True
```

The model for game 5 has two worlds again because player b doesn't know the value of the penny.

```
Lies> game5
Mo
[0, 1]
[(0, ([], [(R1, 6), (R2, 4), (R3, 0), (R4, 0)])),
 (1, ([p], [(R1, 6), (R2, 4), (R3, 0), (R4, 0)]))]
[a, b]
[(a, 0, 0), (a, 1, 1), (b, 0, 0), (b, 0, 1), (b, 1, 0), (b, 1, 1)]
```

[1]

In game 6 player *a* tosses heads and truthfully announces this and player *b* challenges player *a*. In this case both players add one extra euro to the stakes, the cup is lifted and player *a* gets to keep the stakes. The model for this has one world where *p* is true, player *a* has 7 euros and player *b* has 3 euros.

```
Lies> game6
Mo
[1]
[ (1, ([p], [(R1, 7), (R2, 3), (R3, 0), (R4, 0)])) ]
[a, b]
[ (a, 1, 1), (b, 1, 1) ]
[1]
```

In this case player *b* knows the value of the penny and the payoffs are 2 euros for player 1 and -2 euros for player 2:

```
Lies> isTrue game6 bKnows
True
Lies> isTrue game6 (eq (Reg c1) (ASum [I initCash1, I 2]))
True
Lies> isTrue game6 (eq (Reg c2) (ASum [I initCash2, I (-2)]))
True
```

Our analysis shows that even though in the game of Liar’s Dice, lying takes place according to the definition of Augustine, no misleading is taking place and the players are never duped into believing a falsehood. This is shown by the fact that all updates in the game are S5 updates.

6. Conclusion

There are still two discrepancies in the paper that we have to address. The first one is between our treatment of lying in public discourse and our treatment of lying in games. As we have seen, lying in public discourse can lead to KD45 models, which illustrates the fact that genuine misleading takes place. We argued that the players in a game like Liar’s Dice are never actually misled, so in a sense no real lying takes place here at all. Still, lying in public discourse and lying in games seem to be connected somehow.

The difference between the two settings could be seen as a difference in the *protocol* the agents are following. In public discourse, the agents usually assume that they are following the protocol “only speak the truth”. Therefore, when one of them deviates from the protocol by telling a lie, the others believe him and are misled. In the game of Liar’s Dice, the protocol is “say anything in order to improve your payoff”. Since all agents know that the others are following the protocol, they do not believe each other’s lies. The issue of protocol dynamics in epistemic modelling is explored further in [28].

The second discrepancy is between the game-theoretical analysis of lying in games in terms of mixed strategies that use probabilities, and the logical analysis in terms of truth

values. To see that these perspectives do still not quite match, consider the game situation where player 1 tosses the penny, observes the result, and announces ‘heads’. In our logical analysis this does *not* lead to the false belief of player 2 that the coin has landed heads; it does not lead to a belief change at all. But the game-theoretical analysis reveals that a rational agent would have formed a belief about the probability that the claim is true. So it seems that the logical analysis is still too crude.

This defect could be remedied by using probabilistic beliefs and probabilistic updates, in the style of [10], which would allow us to express the probability of actions in the game. With these, we can model the fact that the game-theoretical analysis in terms of mixed strategies is common knowledge. For if this is the case, it is common knowledge that if the toss is tail, then player 1 will announce ‘heads’ with probability $\frac{1}{3}$ and ‘tails’ with probability $\frac{2}{3}$.

Interestingly, this is also relevant for the first discrepancy. For why are the players not duped into believing falsehoods, in the game of Liar’s Dice? Because they look further than a single run of the game, and they know that as the game gets repeated they can adhere to mixed strategies. Therefore, an analysis in terms of manipulative probabilistic updates might work for both lying in public discourse and lying in games.

But there is need here for further work. Even if we switch to a probabilistic version of the logic of communication and change, we have to attach probabilities to the update actions that we start with. This leaves open the problem of how to use logic to *derive* the correct Nash equilibria in the first place. In future work we will explore the possibility of letting agents find such solutions by iterative playing of the game and updating their probabilities until a fixpoint representing an equilibrium is reached.

Other areas of future work are the connection of the logic of lying with belief revision and the modelling of agency. Believing a lie might have the consequence that an initial true belief is given up in favour of a false one. This will only happen, however, if the original true belief is held weakly enough to be replaced by the lie. In modelling a lie as a publicly announced falsehood that is believed by the audience we have left out the liar. To get the liar back into the picture, one has to analyse the preconditions for a lie, in terms of the doxastics of the input model. For agent i to be the originator of a lie, i has to believe ϕ and announce $\neg\phi$, so $B_i\phi$ is a precondition of the lying action. This issue will be taken up in future work.

Finally, we must mention the fact that in philosophy and logic there is a long standing interest in liar paradoxes. Now it seems that our language is not powerful enough to express such paradoxes. What happens if we add a mechanism for self reference to dynamic doxastic logic? Does this immediately lead to either incompleteness or inconsistency? What is the simplest possible way of expressing liar paradoxes in (an extension of) dynamic doxastic logic, and what happens as a result?

Acknowledgement Thanks to Rineke Verbrugge for comments and encouragement during the writing of this paper. We also benefitted from the helpful comments of two anonymous reviewers.

Bibliography

- [1] APT, K. Logic programming. In *Handbook of Theoretical Computer Science, Volume B*, J. v. Leeuwen, Ed. Elsevier, Amsterdam, 1990, pp. 493–574.
- [2] ARENDT, H. Truth and politics. In *Between Past and Future — Six Exercises in Political Thought*. Viking Press, 1967 (Penguin Classics Edition, 2006).
- [3] AUGUSTINE, S. De Mendacio. In *A Select Library of the Nicene and Post-Nicene Fathers of the Christian Church*, P. Schaff, Ed., vol. 3. Eerdmans, 1956, 1988. Translated by Rev. H. Browne. Available online at <http://www.newadvent.org/fathers/>.
- [4] BALTAG, A. A logic for suspicious players: epistemic action and belief-updates in games. *Bulletin of Economic Research* 54, 1 (2002), 1–45.
- [5] BALTAG, A., AND MOSS, L. Logics for epistemic programs. *Synthese* 139, 2 (2004), 165–224.
- [6] BALTAG, A., MOSS, L., AND SOLECKI, S. The logic of public announcements, common knowledge, and private suspicions. Tech. Rep. SEN-R9922, CWI, Amsterdam, 1999. With many updates.
- [7] BALTAG, A., MOSS, L., AND SOLECKI, S. The logic of public announcements, common knowledge, and private suspicions. Tech. rep., Dept of Cognitive Science, Indiana University and Dept of Computing, Oxford University, 2003.
- [8] BENTHEM, J. v. Language, logic, and communication. In *Logic in Action*, J. van Benthem, P. Dekker, J. van Eijck, M. de Rijke, and Y. Venema, Eds. ILLC, 2001, pp. 7–25.
- [9] BENTHEM, J. v. One is a lonely number: on the logic of communication. Tech. Rep. PP-2002-27, ILLC, Amsterdam, 2002.
- [10] BENTHEM, J. v., GERBRANDY, J., AND KOOI, B. Dynamic update with probabilities. Manuscript, ILLC, Amsterdam, 2006.
- [11] BENTHEM, J. v., VAN EIJCK, J., AND KOOI, B. Logics of communication and change. *Information and Computation* 204, 11 (2006), 1620–1662.
- [12] BOK, S. *Lying — Moral Choice in Public and Private Life*. The Harvester Press, Sussex, 1978.
- [13] DITMARSCH, H. v. *Knowledge Games*. PhD thesis, ILLC, Amsterdam, 2000.
- [14] DITMARSCH, H. v., VAN DER HOEK, W., AND KOOI, B. Dynamic epistemic logic with assignment. In *Proceedings of the Fourth International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 05)* (New York, 2005), ACM Inc., pp. 141–148.
- [15] EIJCK, J. v. Guarded actions. Tech. Rep. SEN-E0425, CWI, Amsterdam, December 2004. Available from <http://db.cwi.nl/rapporten/>.
- [16] EIJCK, J. v. DEMO — a demo of epistemic modelling. In *Interactive Logic — Proceedings of the 7th Augustus de Morgan Workshop*, J. v. Benthem, D. Gabbay, and B. Löwe, Eds., no. 1 in Texts in Logic and Games. Amsterdam University Press, 2007, pp. 305–363.
- [17] EIJCK, J. v., AND WANG, Y. Propositional Dynamic Logic as a logic of belief revision. In *Proceedings of Wollic’08* (2008), W. Hodges and R. de Queiros, Eds., no. 5110 in Lecture Notes in Artificial Intelligence, Springer, pp. 136–148. http://dx.doi.org/10.1007/978-3-540-69937-8_13.
- [18] FAGIN, R., HALPERN, J., MOSES, Y., AND VARDI, M. *Reasoning about Knowledge*. MIT Press, 1995.
- [19] GERBRANDY, J. *Bisimulations on planet Kripke*. PhD thesis, ILLC, 1999.
- [20] GERBRANDY, J. The surprise examination in dynamic epistemic logic. *Synthese* 155 (2007), 21–33.
- [21] GETTIER, E. Is justified true belief knowledge? *Analysis* 23 (1963), 121–123.
- [22] KOOI, B. Expressivity and completeness for public update logics via reduction axioms. *Journal of Applied Non-Classical Logics* 16, 2 (2007).
- [23] KOOI, B. P. *Knowledge, Chance, and Change*. PhD thesis, Groningen University, 2003.
- [24] MONTAGUE, P. A new disinformation campaign. *New York Times* (April 29 1998).

- [25] PLAZA, J. A. Logics of public communications. In *Proceedings of the 4th International Symposium on Methodologies for Intelligent Systems* (1989), M. L. Emrich, M. S. Pfeifer, M. Hadzikadic, and Z. W. Ras, Eds., pp. 201–216.
- [26] SIETSMA, F. Model checking for dynamic epistemic logic with factual change. Bachelor's Thesis, University of Amsterdam, 2007.
- [27] STEUB, M. The analysis of knowledge. In *Stanford Encyclopedia of Philosophy*, E. N. Zalta, Ed. Metaphysics Research Lab, CSLI, Stanford, 2006.
- [28] WANG, Y. Epistemic Modelling and Protocol Dynamics. PhD thesis, University of Amsterdam, 2010.