Theory of Mind and Epistemic Planning for Human-Robot Collaboration

Thomas Bolander, DTU Compute, Technical University of Denmark

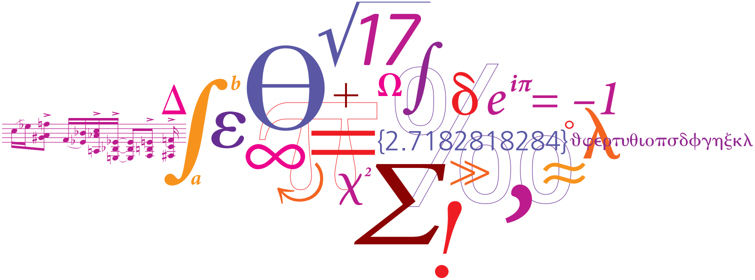

DTU Compute
Department of Applied Mathematics and Computer Science

Testing Theory of Mind: the Sally-Anne test

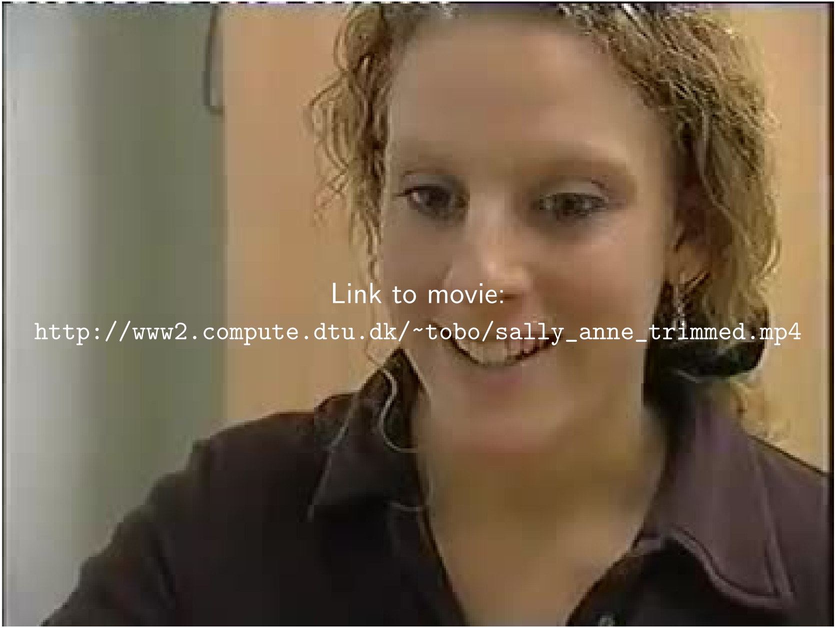

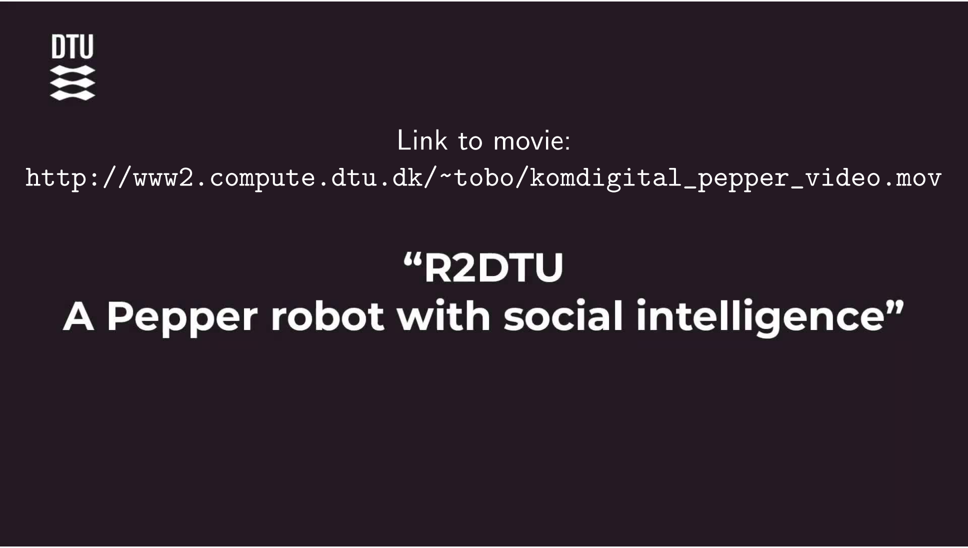

(KomDigital: R2DTU – A Pepper robot, 25 November 2020)

Epistemic planning =
automated *planning* + Theory of Mind reasoning

Aim: To compute plans that can take the mental states of other agents into account.

Essentially: (Decentralised) **multi-agent planning** in environments with (potentially higher-order) **information asymmetry**.

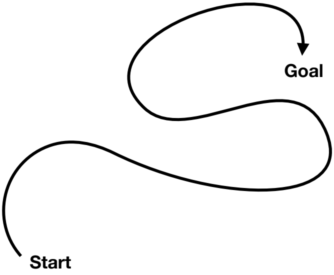

Automated planning

+

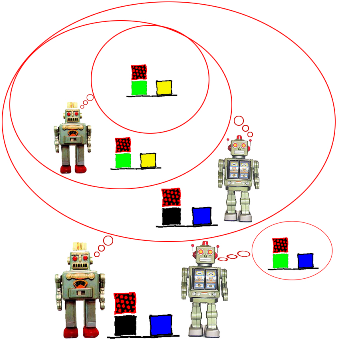

Logical reasoning about the mental states of other agents

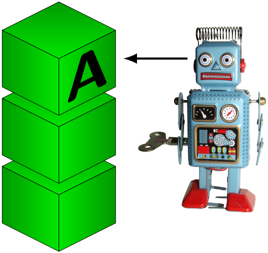

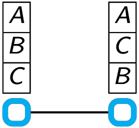

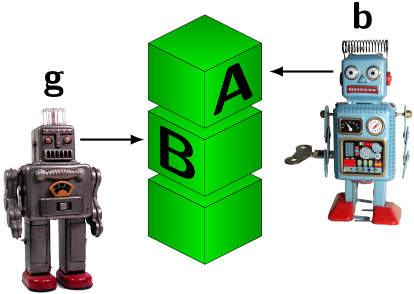

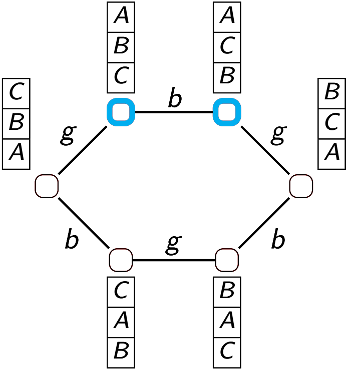

Epistemic states: Multi-pointed epistemic models of multi-agent S5. Nodes are **worlds**. **Designated worlds**: ▢ (those considered possible by planning agent).

The coordinated attack problem in dynamic epistemic logic (DEL)

Two generals (agents), a and b. They want to coordinate an attack, and only win if they attack simultaneously.

d: "general a will attack at dawn".

m_i: the messenger is at general i (for $i = a, b$).

Initial **epistemic state**:

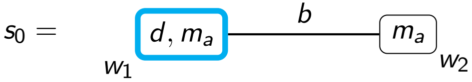

Nodes are **worlds**, edges are **indistinguishability edges** (reflexive loops not shown).

The coordinated attack problem in dynamic epistemic logic (DEL)

Recall: d means "a attacks at dawn"; m_i means messenger is at general i.

Available **epistemic actions** (aka **action models** aka **event models**):

$$a{:}send = \underset{e_1}{\boxed{\begin{array}{ll} \text{pre}: & d \wedge m_a \\ \hline \text{post}: & m_b \wedge \neg m_a \end{array}}} \underset{a}{\text{———}} \underset{e_2}{\boxed{\begin{array}{ll} \text{pre}: & \top \\ \hline \text{post}: & \neg m_a \wedge \neg m_b \end{array}}}$$

And symmetrically an epistemic action $b{:}send$. We read $i{:}\alpha$ as "agent i does α".

Nodes are **events**, and each event has a **precondition** and a **postcondition** (effect). The precondition is an epistemic formula and the postcondition is a conjunction of literals.

[Baltag et al., 1998, van Ditmarsch and Kooi, 2008]

The product update in dynamic epistemic logic

$s_0 =$

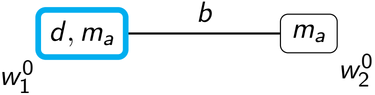

$s_0 \models K_a d \wedge \neg K_b d$

$a{:}send =$

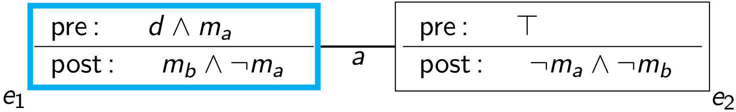

$s_0 \otimes a{:}send =$

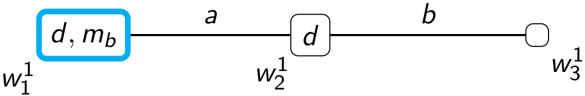

$s_0 \otimes a{:}send \models K_a d \wedge K_b d \wedge \neg K_a K_b d$

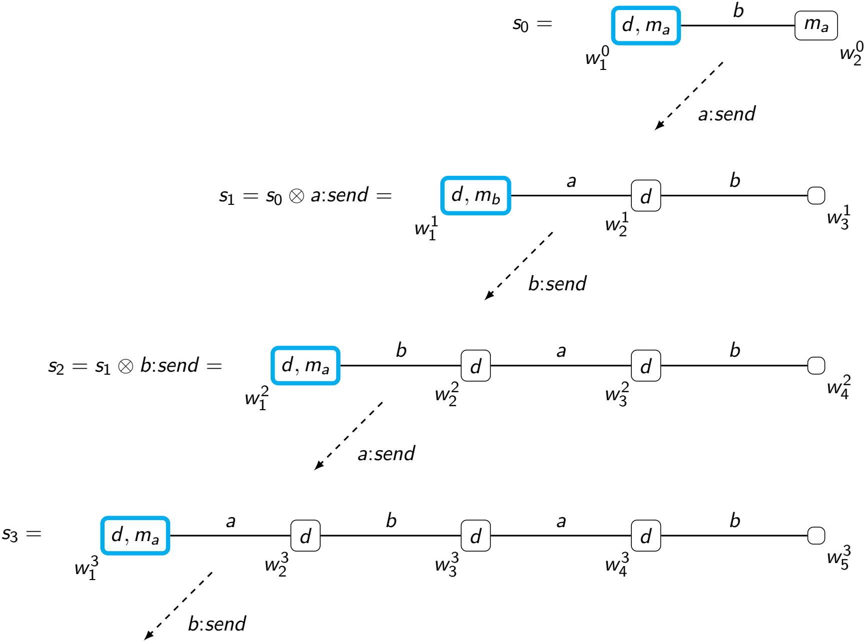

$s_0 =$
d, m_a
m_a
b
w_1^0
w_2^0
a:send
$s_1 = s_0 \otimes a$:send $=$
d, m_b
d
a
b
w_1^1
w_2^1
w_3^1
b:send
$s_2 = s_1 \otimes b$:send $=$
d, m_a
d
d
b
a
b
w_1^2
w_2^2
w_3^2
w_4^2
a:send
$s_3 =$
d, m_a
d
d
d
a
b
a
b
w_1^3
w_2^3
w_3^3
w_4^3
w_5^3
b:send

Epistemic planning tasks

Definition. An **epistemic planning task** (or simply a **planning task**) $T = (s_0, A, \gamma)$ consists of an epistemic state s_0 called the **initial state**; a finite set of epistemic actions A; and a **goal formula** γ of the epistemic language.

Definition. A (sequential) **solution** to a planning task $T = (s_0, A, \gamma)$ is a sequence of actions $\alpha_1, \alpha_2, \ldots, \alpha_n$ from A such that for all $1 \leq i \leq n$, α_i is applicable in $s_0 \otimes \alpha_1 \otimes \cdots \otimes \alpha_{i-1}$ and

$$s_0 \otimes \alpha_1 \otimes \alpha_2 \otimes \cdots \otimes \alpha_n \models \gamma.$$

Example. Let s_0 be the initial state of the coordinated attack problem. Let $A = \{a{:}send, b{:}send\}$. Then the following are planning tasks:

1. $T = (s_0, A, Cd)$, where C denotes common knowledge. It has no solution.
2. $T = (s_0, A, E^n d)$, where E denotes "everybody knows" and $n \geq 1$. It has a solution of length n.

[Bolander et al., 2020]

Epistemic planning example: Get the cube

- **Objects**: $\mathcal{O} = \{b_1, b_2, c\}$, two boxes b_1 and b_2, and a cube c.
- **Agents**: $\mathcal{A} = \{h, a\}$, a human h and a robot r. The robot is the planning agent.
- **Atomic propositions**: $In(x, y)$ means x is in y, where $x, y \in \mathcal{O} \cup \mathcal{A}$ (when $y \in \mathcal{A}$, it means y is holding x).

Initial epistemic state:

$$s_0 = \quad \boxed{In(c, b_1)} \overset{h}{\text{———}} \boxed{In(c, b_2)}$$

The goal is for the human to hold the cube, $In(c, h)$.

Actions specialised for the case of $\mathcal{O} = \{b_1, b_2, c\}$.

Agent i (semi-privately) **peeks** into box x:

$$i{:}peek(x) = \quad \begin{tikzcd} \boxed{\text{pre}: \quad In(c,x)} \arrow[r, dash, "\mathcal{A} - \{i\}"'] & \boxed{\text{pre}: \quad \neg In(c,x)} \end{tikzcd}$$

Agent i (publicly) **picks up** object x from y:

$$i{:}pickup(x,y) = \quad \boxed{\begin{array}{ll} \text{pre}: & In(x,y) \\ \hline \text{post}: & In(x,i) \wedge \neg In(x,y) \end{array}}$$

Agent i (publicly) **puts** object x in y:

$$i{:}putdown(x,y) = \quad \boxed{\begin{array}{ll} \text{pre}: & In(x,i) \\ \hline \text{post}: & In(x,y) \wedge \neg In(x,i) \end{array}}$$

Agent i (publicly) **announces** that formula φ is true:

$$i{:}ann(\varphi) = \quad \boxed{\text{pre}: \quad \varphi}$$

Get the cube: Planning task and solutions

The planning task T has the actions of the previous slide and initial state s_0 and goal γ given by:

$$s_0 = \begin{tikzcd} \boxed{In(c, b_1)} \arrow[r, dash, "h"] & \boxed{In(c, b_2)} \end{tikzcd} \qquad \gamma = In(c, h)$$

Solution to T, by robot R:

$$\begin{tikzcd}
s_0 = & \boxed{In(c, b_1)} \arrow[r, dash, "h"] \arrow[d, dashed, "r{:}pickup(c, b_1)"] & \boxed{In(c, b_2)} \\
s_1 = s_0 \otimes r{:}pickup(c, b_1) = & \boxed{In(c, r)} \arrow[d, dashed, "r{:}putdown(c, h)"] & \\
s_2 = s_1 \otimes r{:}putdown(c, h) = & \boxed{In(c, h)} &
\end{tikzcd}$$

Applicability, perspective shifts, implicit coordination

Seemingly simpler solution: $h{:}\mathit{pickup}(c, b_1)$. But intuitively, this shouldn't work, since the human doesn't know the cube is in box 1...

Applicability: An action α is **applicable** in a state s if for each designated world w of s there is a designated event e of α with $w \models \mathit{pre}(e)$.

Perspective shift: The **perspective shift** of state s to agent i, denoted s^i, is achieved by closing under the indistinguishability relation of i. We call s^i the **perspective** of agent i on state s.

Example. $h{:}\mathit{pickup}(c, b_1)$ is not applicable in s_0 from h's perspective.

Implicitly coordinated solution to planning task: Each action has to be applicable from the perspective of the acting agent; and the product update $s \otimes i{:}\alpha$ is replaced by $s^i \otimes i{:}\alpha$.

Get the cube: Implicit coordination

Joint solution to T, by robot R, implicitly coordinated:

$$
\begin{tikzcd}
s_0 = & \boxed{In(c, b_1)} \arrow[r, dash, "h"] \arrow[d, dashed, "r{:}ann(In(c, b_1))"] & \boxed{In(c, b_2)} \\
s_1 = s_0 \otimes r{:}ann(In(c, b_1)) = & \boxed{In(c, b_1)} \arrow[d, dashed, "h{:}pickup(c, b_1)"] & \\
s_2 = s_1 \otimes h{:}pickup(c, b_1) = & \boxed{In(c, h)} &
\end{tikzcd}
$$

If purely epistemic actions (announcements) have a lower cost than ontic actions (moving things around), the solution above is the only optimal one.

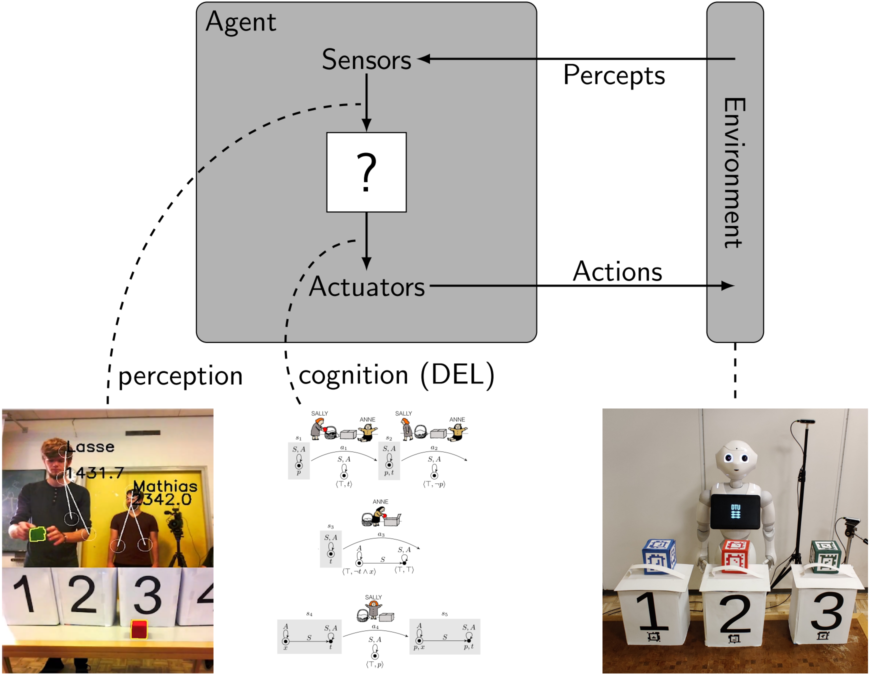
Agent
Sensors
Percepts
?
Environment
Actuators
Actions
perception
cognition (DEL)
Lasse
1431.7
Mathias
2342.0
1
2
3
SALLY
ANNE
SALLY
ANNE
ANNE
SALLY
DTU
1
2
3

Perception layer: detectors, world model and events

Detectors: Detect a specific kind of feature such as faces (dlib CNN face recognition), markers (AprilTag fiducial markers), and body poses (OpenPose).

Spatial world model: Keeps track of the spatial position of physical entities using the detectors. Physical entities are split into *objects* $\mathcal{O}$ and *agents* $\mathcal{A}$.

Events: Changes in the spatial world model triggers *events*:

- $\texttt{Appear}(c)/\texttt{Disappear}(c)$: World model starts/stops tracking entity c.
- $\texttt{pickup}(i, c)$: Agent i picks up object c. Triggered by hand of i entering bounding box of c.
- $\texttt{put}(i, c, b)$: Agent i puts object c in container b.

From perception layer to cognition layer: Every event is translated into its corresponding epistemic action and applied to the current epistemic state via the product update.
E.g. $\texttt{put}(i, c, b) \curvearrowright i{:}\mathit{putdown}(c, b)$.

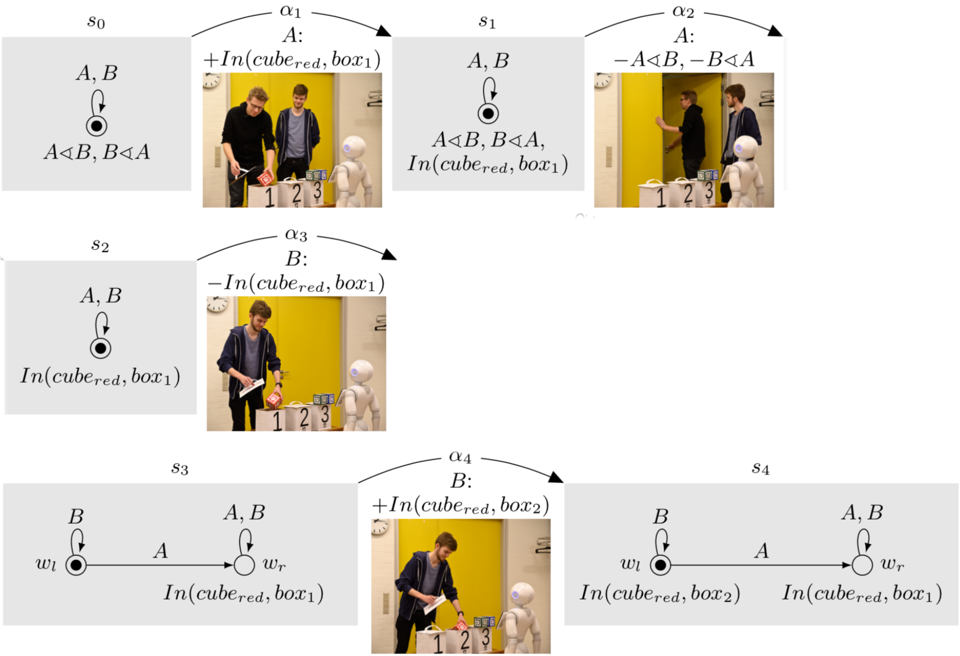
s_0
A, B
$A \triangleleft B, B \triangleleft A$
α_1
A:
$+In(cube_{red}, box_1)$
s_1
A, B
$A \triangleleft B, B \triangleleft A,$
$In(cube_{red}, box_1)$
α_2
A:
$-A \triangleleft B, -B \triangleleft A$
s_2
A, B
$In(cube_{red}, box_1)$
α_3
B:
$-In(cube_{red}, box_1)$
s_3
B
w_l
A
A, B
w_r
$In(cube_{red}, box_1)$
α_4
B:
$+In(cube_{red}, box_2)$
s_4
B
w_l
$In(cube_{red}, box_2)$
A
A, B
w_r
$In(cube_{red}, box_1)$

Helpful announcements

- We add announcements, so the robot can be helpful by announcing facts.
- The robot does epistemic planning with implicit coordination: multi-agent planning with perspective shifts [Nebel et al., 2019, Bolander et al., 2018, Engesser et al., 2017].

Example. Consider the following action sequence:

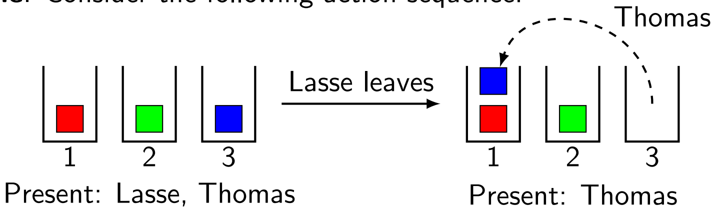

If I say "I want two cubes in the same box", nothing happens. Lasse arrives and says the same. Now the robot replies: "It is already true".

Afterwards Lasse says: "I want three cubes in the same box". The robot replies: "Box 3 is empty".

References I

Baltag, A., Moss, L. S. and Solecki, S. (1998).
The Logic of Public Announcements and Common Knowledge and Private Suspicions.
In Proceedings of the 7th Conference on Theoretical Aspects of Rationality and Knowledge (TARK-98), (Gilboa, I., ed.), pp. 43–56, Morgan Kaufmann.

Bolander, T., Charrier, T., Pinchinat, S. and Schwarzentruber, F. (2020).
DEL-based Epistemic Planning: Decidability and Complexity.
Artificial Intelligence *287*, 1–34.

Bolander, T., Engesser, T., Mattmüller, R. and Nebel, B. (2018).
Better Eager Than Lazy? How Agent Types Impact the Successfulness of Implicit Coordination.
In Proceedings of the 16th International Conference on Principles of Knowledge Representation and Reasoning (KR 2018) AAAI Press.

Engesser, T., Bolander, T., Mattmüller, R. and Nebel, B. (2017).
Cooperative Epistemic Multi-Agent Planning for Implicit Coordination.
In Proceedings of Methods for Modalities number 243 in Electronic Proceedings in Theoretical Computer Science pp. 75–90,.

Nebel, B., Bolander, T., Engesser, T. and Mattmüller, R. (2019).
Implicitly Coordinated Multi-Agent Path Finding under Destination Uncertainty: Success Guarantees and Computational Complexity.
Journal of Artificial Intelligence Research *64*, 497–527.

van Ditmarsch, H. and Kooi, B. (2008).
Semantic Results for Ontic and Epistemic Change.
In Logic and the Foundation of Game and Decision Theory (LOFT 7), (Bonanno, G., van der Hoek, W. and Wooldridge, M., eds), Texts in Logic and Games 3 pp. 87–117, Amsterdam University Press.