



THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Reasoning about Imperfect Information Games in the Epistemic Situation Calculus

Citation for published version:

Belle, V & Lakemeyer, G 2010, Reasoning about Imperfect Information Games in the Epistemic Situation Calculus. in *Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2010, Atlanta, Georgia, USA, July 11-15, 2010*. AAAI Press, pp. 255-260.
<<http://www.aaai.org/ocs/index.php/AAAI/AAAI10/paper/view/1803>>

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Publisher's PDF, also known as Version of record

Published In:

Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2010, Atlanta, Georgia, USA, July 11-15, 2010

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.



Reasoning about Imperfect Information Games in the Epistemic Situation Calculus

Vaishak Belle* and Gerhard Lakemeyer

Dept. of Computer Science

RWTH Aachen

52056 Aachen

Germany

{belle,gerhard}@cs.rwth-aachen.de

Abstract

Approaches to reasoning about knowledge in imperfect information games typically involve an exhaustive description of the game, the dynamics characterized by a tree and the incompleteness in knowledge by information sets. Such specifications depend on a modeler's intuition, are tedious to draft and vague on where the knowledge comes from. Also, formalisms proposed so far are essentially propositional, which, at the very least, makes them cumbersome to use in realistic scenarios. In this paper, we propose to model imperfect information games in a new multi-agent epistemic variant of the situation calculus. By using the concept of only-knowing, the beliefs and non-beliefs of players after any sequence of actions, sensing or otherwise, can be characterized as entailments in this logic. We show how *de re* vs. *de dicto* belief distinctions come about in the framework. We also obtain a regression theorem for multi-agent beliefs, which reduces reasoning about beliefs after actions to reasoning about beliefs in the initial situation.

Introduction

Over many decades now, computer science and game theory have benefited from each other. Here, the AI viewpoint has been to focus on the *representation* and *reasoning* of games, motivated by the idea that the contributions will be applicable on general purpose autonomous agents, such as robots, participating in strategic situations. However, most papers follow economics in analyzing small games, investigating issues of rational behavior and cooperation. In his critical paper, Halpern (2003) accurately points out that when agents, rather than people, are involved in more *practical* settings, a range of new issues come into play; and what is "rational" may be obvious. The main problem perhaps is in *describing* the game, and beliefs about them. The goal of this paper is to explore and suggest a simple alternative to how beliefs should come about in *imperfect information games*.

Games are broadly classified into perfect information and imperfect information games. In the former, participating players are aware of the current game state defined over a history of actions (Chess). In the latter, the precise move is not known (Kriegspiel Chess, Poker) and the norm is

to partition the nodes in a game tree, called *information sets*, to mean epistemic alternatives after the corresponding history. Although nodes do capture certain aspects of knowledge, they do not allow players to infer what opponent strategies might be, given their knowledge. Beliefs, in game theory, are modeled more thoroughly in a *S5* variant called the *state space* approach (Osborne and Rubinstein 1994; Fagin et al. 1995) where, by state, a full description of the world is meant. But Halpern (2003) remarks on three severe limitations of states: (a) they do not tell us the actions possible, (b) they cannot model situations of the sort "after *a*, Bob does not know ϕ ", and (c) where is the knowledge coming from? Roughly, the tree representation fits one task, and states another; both fail on (c). The treatment is more worrisome in large games, since exhaustively defining information sets for the entire depth of the game tree is completely left to a modeler's intuition and error-prone (see (Halpern 2003) on problems thereof). In this regard, to precisely capture beliefs after actions, one needs a semantically motivated approach. Finally, the underlying model of belief used in games is not always a logic, and formalisms to reify states as a logical entity are essentially propositional. At the very least, modeling actions in realistic scenarios such as picking one of 52 cards in Poker, or even playing games with (possibly) infinite candidate moves at each state is cumbersome.

In a sense, first-order action formalisms like the situation calculus (McCarthy and Hayes 1969) or the fluent calculus (Thielscher 1999) have been known to be expressive enough to capture dynamics in realistic settings. Schulte and Delgrande (2004) propose an *isomorphic* mapping from game trees to situation calculus theories, but this approach has some problems (discussions deferred to later). Our perspective is that an axiomatic formulation, as is standard in KR communities, is preferable, as opposed to mappings, and a semantics must be proposed to appropriately bring about beliefs in these games. We argue that the axiom exercise is justified in avoiding an explicit specification of large game trees and their information sets. (We shortly present a card game that demonstrates just that.) To this end, our methodology is as follows. We propose to model imperfect information games in a multi-agent epistemic variant of the situation calculus. From the perspective of this paper, the relevant feature of the logic and its possible-world semantics is that a world is a *tree* of situations. Intuitively, we are able to in-

*Vaishak Belle is supported by a DFG scholarship from the graduate school GK 643.

Copyright © 2010, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

interpret the dynamics of the game as a sequence of actions. Thus, if a set of worlds is epistemically possible after some history, then the player has incomplete knowledge, accounting for information sets. Strictly speaking though, worlds are richer than reified game states, since they are not restricted only to facets about the game. This fits well with the paradigm of *Cognitive Robotics* (Lakemeyer and Levesque 2007), that of autonomous agents reasoning about all sorts of tasks, not just strategic ones. From a more representational viewpoint, we give a remedy to the issues raised by Halpern *viz.* the history is maintained in the semantics (clarifying (a)), and sensing functions account for (b) and (c). We also show in the paper that sensing functions can bring about *de re* and *de dicto* distinctions in belief. For example, if Alice reads her card, Bob does not immediately know what her card is, but knows that she knows.

The leverage of our machinery, to begin with, is that atomic properties of game states after any sequence of actions can be analyzed as *entailments*. We lift this idea for knowledge, by appealing to Levesque’s notion of only-knowing.¹ Briefly, Levesque introduced only-knowing to precisely capture the beliefs of a knowledge base (KB). By specifying the axioms that characterize the game and sentences opponents are taken to know as part of the agent’s KB, beliefs and non-beliefs after any sequence of actions are naturally analyzed as logical consequences of the KB. To summarize, the resulting model allows us to study games after arbitrary histories, all in the confines of the logic.

Entailments in our logic are made amenable to effective reasoning with a new technical result. Regression theorems for knowledge are currently available in two flavors: in the classical situation calculus setting (Scherl and Levesque 2003), and in a modal setting (Lakemeyer and Levesque 2004), both for one agent only. To this end, we obtain a regression property for a rich multi-agent epistemic characterization, which allows us to reduce reasoning about beliefs after actions to reasoning about what is known initially.

We finally remark that the semantical framework proposes for the first time a natural generalization of only-knowing to the many agent case for a quantified language with equality and actions. In a recent paper, we (Belle and Lakemeyer 2010) discuss how this new notion of only-knowing, in the *static* case i.e. no actions, relates to earlier attempts (Halpern and Lakemeyer 2001). Only-knowing has been shown to capture non-trivial nonmonotonic behavior (Levesque 1990). So the appeal of the semantics is to have game playing agents reason non-monotonically about the world and such non-monotonic mechanisms in other agents.

We are organized as follows. First, we introduce the new logic. Then, we show how knowledge about games (formulated axiomatically) is modeled and discuss a regression property. We conclude with related work. For space reasons, we only provide sketches for proofs.

The Logic \mathcal{GL}

The logic \mathcal{GL} is a multi-agent epistemic variant of the situation calculus. The situation calculus is a first-order for-

¹We use the terms “knowledge” and “belief” interchangeably, but allow agents to have false beliefs about the world.

mal and computational framework for the representation and reasoning of actions and knowledge in the dynamic environment (McCarthy and Hayes 1969). A second-order extension forms the theoretical foundations for the agent programming language *Golog* (Reiter 2001). Here we follow the relevant ideas of a reformulation of the situation calculus in a modal setting (Lakemeyer and Levesque 2004), and extend it to the many agent case. The language is a sorted first-order modal dialect with $=$, and is composed of formulas over symbols from the following vocabulary: first-order variables v, x, y, \dots , a countably infinite set of standard names n_1, n_2, \dots of sort action and object, rigid functions, rigid and fluent predicates of every arity, a distinguished fluent predicate *Poss*, $=$, and connectives \forall , \neg and \vee . We include two types of modal operators: epistemic operators B_i and O_i for $i = a, b$, and modal operators for actions $[v]$, \square .² Finally, distinguished fluent functions SF_a and SF_b (for sensing) are the only fluent functions in the language. For readability, the sort (action or object) of variables and arguments of functions and predicates is left implicit. Standard names should be understood as rigid designators for all objects and actions in the domain. As we will see, having standard names in the language simplifies the semantics and also enables *de dicto* vs. *de re* distinctions.

Terms are the least set of expressions formed from all standard names, first-order variables, and if t_1, \dots, t_k are terms then so is $g(t_1, \dots, t_k)$, where g is a k -ary function. We call a term $g(n_1, \dots, n_k)$ *primitive* if the n_i are standard names. Terms without variables are called *ground* terms. The *well-formed formulas* (wffs) are the least set of expressions such that $P(t_1, \dots, t_k)$ and $t_1 = t_2$ are (*atomic*) formulas, and $[t]\alpha$, $\alpha \vee \beta$, $\neg\alpha$, $\forall x. \alpha$, $\square\alpha$, $B_i\alpha$, $O_i\alpha$ are all formulas (where t is an action term, and α, β are formulas). We read $[t]\alpha$ as “ α holds after action t ”, $\square\alpha$ as “ α holds after every sequence of actions”, $B_i\alpha$ as “ i knows α ”, and $O_i\alpha$ as “all that i knows is α .” We call a formula without any free variables a *sentence*. A *primitive sentence* is an atomic formula, say $P(n_1, \dots, n_k)$, where the n_i , are standard names. A formula with no \square operators is called *bounded*, one without $[t]$ or \square is called *static*, and one without B_i , O_i , \square , $[t]$, *Poss*, or SF_i is called a *fluent* formula. A formula that does not mention any O_i operators for any i is called *basic*, and one that does not mention epistemic operators at all is called *objective*. Finally, we also have a notion of *depth* defined as:

Definition 1. The i -depth of a formula α , denoted $|\alpha|_i$, is defined inductively as (L_i stands for both B_i and O_i):

1. $|\alpha|_i = 1$ for atomic formulas,
2. $|\neg\alpha|_i = |\alpha|_i$, $|\forall x. \alpha|_i = |\alpha|_i$, $|[v]\alpha|_i = |\alpha|_i$
3. $|\square\alpha|_i = |\alpha|_i$, $|\alpha \vee \beta|_i = \max(|\alpha|_i, |\beta|_i)$,
4. $|L_i\alpha|_i = |\alpha|_i$, $|L_j\alpha|_i = |\alpha|_j + 1$ for $j \neq i$.

A formula has a depth k if $\max(a\text{-depth}, b\text{-depth}) = k$.

For example, the formula $B_a B_b B_b B_a P(x) \wedge B_b Q(x)$ has a depth of 4, a a -depth of 3 and a b -depth of 4.

²For ease of exposition, only two agents Alice (a) and Bob (b) are considered. Extensions to more agents is straightforward. Also, other connectives, such as \exists , \wedge , \equiv and \supset are taken for their usual syntactic abbreviations.

Semantics. Here the possible-world framework, on one hand, accounts for the dynamic changes in the world, and on the other, provides a natural setting for only-knowing in the many agent case. Informally, a world determines the truth value of fluent and rigid predicates after actions, and co-referring standard names of terms. We define a notion of epistemic states called k -structures, built from sets of such worlds. Formally, denoting the set of all possible finite action sequences with \mathcal{Z} , a world $w \in \mathcal{W}$ is any function from primitive sentences and \mathcal{Z} to $\{0, 1\}$ and from primitive terms and \mathcal{Z} to the standard names, and satisfying the rigidity constraint: if l is a rigid function then $w[l(t_1, \dots, t_k), z] = w[l(t_1, \dots, t_k), z']$ for all action sequences $z, z' \in \mathcal{Z}$. $|t|_w^z$ denotes the co-referring standard name for term t , world w and sequence z , defined as: (a) if t is a standard name then $|t|_w^z = t$, (b) $|g(t_1, \dots, t_k)|_w^z = w[g(n_1, \dots, n_k), z]$, where $n_i = |t_i|_w^z$.

Definition 2. A k -structure ($k \geq 1$), say e^k , for an agent is defined inductively as:

- $e^1 \subseteq \mathcal{W} \times \{\{\}\}$,
- $e^k \subseteq \mathcal{W} \times \mathbb{E}^{k-1}$, where \mathbb{E}^m is the set of all m -structures.

A e^1 for a , denoted e_a^1 , is intended to represent a set of worlds. A e^2 is of the form $\{\langle w, e_b^1 \rangle, \langle w', e_b^1 \rangle \dots\}$, and is to be read as "at world w , a believes b considers the worlds from e_b^1 possible but at w' she believes b to consider e_b^1 possible". This captures the notion that a only has partial knowledge about b , and so in different worlds, her beliefs about what b knows differ. Intuitively, e_a^k determines what (and all that) is believed about the world and b , to a depth k .

We shall henceforth call a e^k for Alice, a e^j for Bob and a world w as a (k, j) -model. Formulas are interpreted wrt. a (k, j) -model iff they have a maximal a -depth of k , and a maximal b -depth of j . The semantics for such formulas wrt. a (k, j) -model and a sequence of actions $z \in \mathcal{Z}$ is:

1. $e_a^k, e_b^j, w, z \models P(t_1, \dots, t_k)$ iff $w[P(n_1, \dots, n_k), z] = 1$ where $n_i = |t_i|_w^z$;
 2. $e_a^k, e_b^j, w, z \models (t_1 = t_2)$ iff n_1 and n_2 are identical where $n_i = |t_i|_w^z$;
 3. $e_a^k, e_b^j, w, z \models \neg \alpha$ iff $e_a^k, e_b^j, w, z \not\models \alpha$;
 4. $e_a^k, e_b^j, w, z \models \alpha \vee \beta$ iff $e_a^k, e_b^j, w, z \models \alpha$ or $e_a^k, e_b^j, w, z \models \beta$;
 5. $e_a^k, e_b^j, w, z \models \forall x. \alpha$ iff $e_a^k, e_b^j, w, z \models \alpha_n^x$ for all standard names n of the appropriate sort;
 6. $e_a^k, e_b^j, w, z \models [r]\alpha$ iff $e_a^k, e_b^j, w, z \cdot r \models \alpha$;
 7. $e_a^k, e_b^j, w, z \models \Box \alpha$ iff $e_a^k, e_b^j, w, z \cdot z' \models \alpha$ for all $z' \in \mathcal{Z}$.
- To interpret epistemic possibilities over actions, we define $w' \simeq_z^i w$ (read: w' and w agree on i 's sensing for z) inductively: (a) for $z = \langle \rangle$, $w \simeq_z^i w'$ for any two worlds, (b) for $z = z' \cdot r$, $w \simeq_{z' \cdot r}^i w'$ iff $w' \simeq_{z'}^i w$ and $w[SF_i(r), z'] = w'[SF_i(r), z']$. Now, the semantics for $B_a \alpha$ and $O_a \alpha$ is:
8. $e_a^k, e_b^j, w, z \models B_a \alpha$ iff for all worlds $w' \simeq_z^a w$, for all e_b^{k-1} , if $\langle w', e_b^{k-1} \rangle \in e_a^k$ then $e_a^k, e_b^{k-1}, w', z \models \alpha$;
 9. $e_a^k, e_b^j, w, z \models O_a \alpha$ iff for all worlds $w' \simeq_z^a w$, for all e_b^{k-1} , $\langle w', e_b^{k-1} \rangle \in e_a^k$ iff $e_a^k, e_b^{k-1}, w', z \models \alpha$.




Figure 1: The dealer (root node) gives a card each to Alice (black node) and Bob (white node). The left branch for Alice denotes claiming, and the right branch denotes giving in. Similarly, the left branch for Bob denotes challenging, and the right branch denotes conceding. Information sets are either rectangular, elliptical or dotted lines. Only Alice's payoffs are labeled.

(The semantics for B_b and O_b are given analogously.) For a sentence α and model e_a^k, e_b^j, w , we write $e_a^k, e_b^j, w \models \alpha$ to mean $e_a^k, e_b^j, w, \langle \rangle \models \alpha$. A sentence α of a, b -depth of k, j is *satisfiable* iff there is a (k, j) -model e_a^k, e_b^j, w such that $e_a^k, e_b^j, w \models \alpha$. The formula is *valid* iff α is true at all (k, j) -models. If Σ is any set of sentences (of maximal a, b -depth of k, j) and α is as above, we write $\Sigma \models \alpha$ (read: Σ entails α) if for every (k, j) -model e_a^k, e_b^j, w , if $e_a^k, e_b^j, w \models \alpha'$ for every $\alpha' \in \Sigma$ then $e_a^k, e_b^j, w \models \alpha$. It can also be shown that if α is valid, i.e. α is true at all (k, j) -models then it is also true at all (k', j') -models, for $k' \geq k, j' \geq j$ (Belle and Lakemeyer 2010). Finally, beliefs with k -structures are easily shown to satisfy $K45$ properties, and the Barcan formula, after any sequence of actions (Hughes and Crosswell 1972).

Lemma 3. If α is a formula then the following are valid (wrt models of appropriate depth):

1. $\Box(B_i \alpha \wedge B_i(\alpha \supset \beta) \supset B_i \beta)$, $\Box(\forall x. B_i \alpha \supset B_i(\forall x. \alpha))$
2. $\Box(B_i \alpha \supset B_i B_i \alpha)$, $\Box(\neg B_i \alpha \supset B_i \neg B_i \alpha)$

Reasoning about Games

Axioms that model a game, in general, are defined over sentences that hold initially, and sentences that define how actions behave. These axioms are Reiter-style basic action theories, but expressed as in (Lakemeyer and Levesque 2004) using modal operators. The components of a theory Σ are: (a) Σ_0 is any set of fluent sentences; (b) Σ_{pre} is a singleton sentence of the form $\Box Poss(v) \equiv \pi$, where π is a fluent formula; (c) Σ_{post} is a set of sentences of the form $\Box F(x) \equiv \gamma$, one for each fluent, where γ is a fluent formula; and (d) Σ_{sense} are sentences of the form $\Box SF_i(v) = x \equiv \varphi_i$, for sensing, where φ_i is a fluent formula. To account for incomplete knowledge, we will need to distinguish between what is true in the world and what the players know. This is handled in the paper by maintaining two theories, say Σ and Ω . To exemplify, and also motivate the need for a first-order formalism, we consider a simplified Poker variant from (Koller and Pfeffer 1997). We propose an axiomatization, and then consider how knowledge about games is modeled.

From a deck of three cards $\{J, Q, K\}$, a dealer gives a card each to a and b . Player a moves first by reading her card, and can concede (thereby ending the game and forfeiting a

dollar) or claiming a bet that she has a higher card. Now, b , whose turn it is to play, reads his card. If he concedes then he forfeits a dollar, and if he challenges a 's claim then he forfeits \$4 if a indeed has the higher card. If a 's claim was false, then he gains an equivalent sum. Intuitively, a (and b) initially consider all deals possible (labeled \ominus in the figure). After a reads the card, say J , she only considers states where Bob has K or Bob has Q (labeled \oplus). If she bets anyway, b reads his card (not represented here), say K , and then concludes Alice is claiming with a J or a Q (labeled \otimes), either way he has certain victory on challenging (labeled Δ). We axiomatize this game as follows:³

1. $\Omega_0 = \{card_a \neq card_b, card_a = J \vee card_a = Q \vee card_a = K, card_b = J \vee card_b = Q \vee card_b = K, \neg Win_a, \mathbf{R}\}$;
2. $\Sigma_0 = \Omega_0 \cup \{card_a = J, card_b = K\}$;
3. $\Sigma_{pre} = \{\Box Poss(v) \equiv true\}$ (for simplicity);
4. $\Sigma_{post} = \{\Box[v]Win_a \equiv (v = yield_b) \vee (v = chill_b \wedge card_a > card_b)\}$
5. $\Sigma_{sense} = \{SF_a(snC_a) = x \equiv card_a = x, SF_b(snC_a) = x \equiv x = \epsilon, \dots\}$

In English: we model actions bet_a , $chill_b$, $yield_a$, $yield_b$, snC_a and snC_b to mean a (and b) betting (or challenging), conceding and sensing on the cards dealt to them. Initially, we let constants (0-arity functions) $card_a$ and $card_b$ denote the cards that a and b are dealt, such that they cannot be the same. A fluent Win_a signifies a victory for a (for simplicity, we do not go over the quantitative aspects of the payoffs). The Poker variant is in fact a *zero-sum* game (Osborne and Rubinstein 1994), such that a victory for b is taken to be $\neg Win_a$. Also, we let \mathbf{R} mean any set of sentences that establishes a rank ordering on the cards to allow discussions of a "higher" card. Formally, using 2-ary predicate $>$ in an infix notation, we write $K > Q$, $K > J$ and $Q > J$ (they are taken to be standard names). Incomplete knowledge is modeled using two initial theories, such that Ω_0 decides the basic setup, and Σ_0 also specifies the precise card deal.

For simplicity, we do not place any restrictions on when an action can be executed.⁴ Σ_{post} says that if b concedes or he challenges a with an inferior card, Win_a should be made true. The crux of the theory are the sense axioms. Roughly, we desire a 's sensing on snC_a to inform her of the standard name that $card_a$ returns. But b 's sensing on the same action snC_a should give him no information about what card it is. We keep a special standard name for the latter, say ϵ . It is this feature that brings about not only *de re* and *de dicto* distinctions, but also updates what was not known before and is now known. Sense axioms for snC_b are defined analogously. The sensing results for other actions are vacuously set to ϵ .

We remark that in realistic scenarios, the entire 52-card deck is thrown in, and actions of a more complex type are involved. In this sense, it is not surprising that formalisms study simple variants. For illustration, the information set \ominus

³Free variables are implicitly universally quantified.

⁴For space reasons, we are terse on the particulars. In the formal description of an extensive game, a *player function* is defined as a mapping from histories to $\{a, b\}$. We could maintain fluents in Σ_{pre} to correspond to such a function, if needed.

will consist of 52×51 nodes for a full deck, and similarly, \oplus will consist of 51 nodes. Moreover, the behavior of beliefs is simple enough in this example, but if we would like to model an action such as player 1 revealing the details of her card only to odd numbered participants at the table, drafting information sets states is no longer immediate. On the other hand, treating these issues with axioms in a first-order language allows for fairly natural and, in particular, succinct representations.

We let the union of 2-5 be Σ , the complete basic action theory while Ω denotes the union of Ω_0 , Σ_{pre} , Σ_{post} and Σ_{sense} . They differ since $\Sigma_0 \supseteq \Omega_0$. Intuitively, Ω are the *rules of the game*, in the sense that they define the actions possible and consequences, but not the cards dealt. We are now ready to consider how epistemic reasoning of game dynamics is done.

Knowledge about Games. The main idea is we specify the theory believed at the objective level, and the theory that a believes b to believe, to a depth k (or j). With many applications, it seems intuitive to assume players to believe that opponents only know the rules Ω (in a fair game). Sensibly enough, they also assume opponents to hold similar beliefs of them. In symbols (to a depth k for a):

$$\mathbf{O}_a(\Omega \wedge \mathbf{O}_b(\Omega \wedge \mathbf{O}_a(\dots))) \quad (1)$$

To construct epistemic states about games, such as (1), we restrict to those possible worlds that obey the rules. Let $\mathcal{W}_\Omega = \{w \mid w \models \Omega\}$. Then we have:

Lemma 4. *Let $e^{*1} = \mathcal{W}_\Omega \times \{\{\}\}$, and $e^{*k} = \{\langle w, e^{*k-1} \rangle \mid w \in \mathcal{W}_\Omega\}$ be defined inductively. Then for any w' and e_b^{k+1} , $e^{*k}, e_b^{k+1}, w' \models (1)$.*

Returning to the example, we show that all features analyzed earlier come out naturally with the semantics. The following entailments are of maximal a, b -depth 3, so let \mathcal{M} be shorthand for e_a^k, e_b^j, w^* for $k, j \geq 3$, where the epistemic states are constructed as in Lem. 4 and $w^* \models \Sigma$ is the real world.

1. $\mathcal{M} \models \neg \exists x. \mathbf{B}_a(card_a = x) \wedge \neg \exists x. \mathbf{B}_a \neg(card_a = x)$
Initially, a does not know her card (label \ominus in the figure).
2. $\mathcal{M} \models [snC_a] \mathbf{B}_a(card_a = J)$
After sensing, a knows her card (label \oplus in the figure).
3. $\mathcal{M} \not\models [snC_a] \exists x. \mathbf{B}_b(card_a = x)$
After a senses, b does not have *de re* knowledge of a 's card.
4. $\mathcal{M} \models [snC_a] \mathbf{B}_b \exists x. \mathbf{B}_a(card_a = x)$
After a senses, b has *de dicto* knowledge of a 's card.
5. $\mathcal{M} \models [snC_a] \mathbf{B}_a \neg \exists x. \mathbf{B}_b(card_a = x)$
 a knows that her information is *private*, as it should be.
6. $\mathcal{M} \models [snC_a][bet_a][snC_b] \mathbf{B}_b([chill_b] \neg Win_a)$
After reading his card, b knows that he wins on challenging (label Δ in the figure).

Proof (Sketches): We abbreviate and write $\phi(x)$ to mean $(card_a = x)$. **1:** Clearly, there are worlds $w \models \phi(J)$ and $w' \models \phi(Q)$ such that $w, w' \in \mathcal{W}_\Omega$. Since $w' \simeq_a^1 w^*$, and $\langle w', e_b^{k-1} \rangle \in e_a^k$ by construction (for some e_b^{k-1}), we have $\mathcal{M} \not\models \mathbf{B}_a \phi(J)$. All other cases are analogous. **2:** Following snC_a , only worlds w such that $w[SF_a(snC_a), \langle \rangle] =$

$w^*[SF_a(snC_a), \langle \rangle]$ remain accessible for a ($w \simeq_{snC_a}^a w^*$). Since $w^* \models (SF_a(snC_a) = J) \equiv \phi(J)$, we have that $\mathcal{M}, snC_a \models B_a\phi(J)$. **3:** For all $w \in \mathcal{W}_\Omega$, $w \models SF_b(snC_a) = \epsilon$ including $w' \models \phi(Q)$, $w' \simeq_{snC_a}^b w^*$. Thus, $\mathcal{M}, snC_a \not\models \exists x. B_b\phi(x)$. **4:** By an analogous argument to **2**, for every $\langle w, e_a^{j-1} \rangle \in e_b^j$, we can show $e_a^{j-1}, e_b^j, w, snC_a \models \exists x. B_a\phi(x)$. Intuitively, b believes a to have learnt of her card at each world, but does not know which of these is the real world w^* . **5:** By an analogous argument to **3**, for every $\langle w, e_b^{k-1} \rangle \in e_a^k$, we can show $e_a^k, e_b^{k-1}, w, snC_a \not\models \exists x. B_b\phi(x)$. Therefore, $\mathcal{M}, snC_a \models B_a\neg\exists x. B_b\phi(x)$. ■

Recall the involved information set construction earlier, and the intuition needed for a trivial three card imperfect information game.⁵ In contrast, our analysis is simple, the semantic arguments are direct, and epistemic properties richer than those derivable from an analysis of the information sets are shown to be regular entailments. We now show that sentences in a large fragment of the language are regressible.

Regression. An equivalent way to the entailments (1-6) is to verify, for instance, if $O_a(\Omega \dots) \supset [snC_a]B_a(card_a = J)$. Reiter demonstrated that formulating fluents as successor state axioms (SSAs) has the benefit, besides solving the frame problem, of making logical theories amenable to a reasoning procedure called *regression*. The idea is to replace fluents and the preceding action term in the sentence with the *r.h.s.* of the SSAs until no more action terms are left. The resultant *regressed formula* must follow from the initial theory for the entailment to be true. Since our worlds are defined precisely as in (Lakemeyer and Levesque 2004), we are able to easily lift a result for regressing bounded (no \square) objective formulas. Denoting \mathcal{R} as the regression operator, and letting z be an action sequence, Lakemeyer and Levesque (2004) define $\mathcal{R}[z, \alpha]$ (the regression of α over z) *wrt.* Σ . (We skip this definition for space reasons.) In order to extend such a regression property for knowledge, we use the following theorem which formalizes knowledge updates in a SSA-like form.

Theorem 5. $\models \square[v]B_i(\alpha) \equiv \exists x. SF_i(v) = x \wedge B_i(SF_i(v) = x \supset [v](\alpha))$

Proof (“ \rightarrow ”): Suppose $e_a^k, e_b^j, w, z \models [r]B_a\alpha_r^v \wedge SF_a(r) = \varsigma$ for standard names r, ς . For all $w' \simeq_{z \cdot r}^a w$, if $\langle w', e_b^{k-1} \rangle \in e_a^k$, then $e_a^k, e_b^{k-1}, w', z \cdot r \models \alpha_r^v$. Since $w'[z, SF_a(r)] = w[z, SF_a(r)]$, we have $e_a^k, e_b^{k-1}, w', z \models SF_a(r) = \varsigma \supset [r]\alpha_r^v$. Thus, $e_a^k, e_b^j, w, z \models B_a(SF_a(r) = \varsigma \supset [r]\alpha_r^v)$. ■

It roughly says that what will be known after an action depends on what was known and what the future looks like contingent on sensing. Henceforth, regression instead of being relative to Σ , is defined for a pair of theories $\langle \Sigma, \Omega \rangle$.

⁵Game theory makes a distinction between games of imperfect information, and *incomplete* information. In the latter, there is uncertainty about the payoffs. By constructing epistemic states from worlds $w \models \Omega$, we are modeling games of imperfect information. If we also allowed worlds that satisfy different payoffs, it is possible to consider incomplete information games. See (Osborne and Rubinstein 1994; Schulte and Delgrande 2004).

We allow Σ and Ω to differ arbitrarily, and the idea is all sentences in the scope of beliefs are interpreted *wrt.* Ω , and those not in the scope of epistemic operators are interpreted *wrt.* Σ . Consequently, we redefine \mathcal{R} to mean $\mathcal{R}[\Omega, \Sigma, z, \alpha]$ and the definition for bounded objective formulas is as discussed above, using axioms and sentences from Σ . We add:

8. $\mathcal{R}[\Omega, \Sigma, z, SF_i(v) = x] = \mathcal{R}[\Omega, \Sigma, z, \varphi_i]$;

9. $\mathcal{R}[\Omega, \Sigma, z, B_a(\alpha)]$ is defined inductively on z as:

(a) $\mathcal{R}[\Omega, \Sigma, \langle \rangle, B_a\alpha] = B_a(\mathcal{R}[\Omega, \Omega, \langle \rangle, \alpha])$,

(b) $\mathcal{R}[\Omega, \Sigma, z \cdot r, B_a\alpha] = \mathcal{R}[\Omega, \Sigma, z, \beta_r^v]$.

where β is the *r.h.s.* of Thm 5. (Analogously for $B_b\alpha$.) We now discuss the main regression result.

Theorem 6. *If α is a basic bounded sentence of max. a, b-depth k, j then $\mathcal{R}[\Omega, \Sigma, \langle \rangle, \alpha]$ is a static sentence such that:*

$\Sigma \wedge O_a(\Omega \wedge O_b(\Omega \dots)) \wedge O_b(\Omega \dots) \models \alpha$ iff $\Sigma_0 \wedge O_a(\Omega_0 \wedge O_b(\dots)) \wedge O_b(\Omega_0 \dots) \models \mathcal{R}[\Omega, \Sigma, \langle \rangle, \alpha]$ where, $O_a(\dots)$ has a-depth k , and $O_b(\dots)$ has b-depth j .

Proof (Sketch): The idea is to show that for any k -structure such that $e_a^k, e_b^j, w' \models O_a(\Omega_0 \dots)$ there is a k -structure e^* , constructed inductively by replacing all worlds w in e_a^k , at all depths with w_Ω that is exactly like w except that $w_\Omega \models \Omega$. Then e^* , by Lem. 4, satisfies $O_a(\Omega \dots)$. Then using results from (Lakemeyer and Levesque 2004), we show that the *regression* of objective formulas θ (in the scope of belief operators) *wrt.* Ω is satisfied in e_a^k iff θ is satisfied in e^* . ■

Intuitively, contingent on the nesting of belief operators, an agent refers to the theory believed at the corresponding level to evaluate the interplay between the dynamics and fluents. This result is easily extended to a regression property that allows theories to differ arbitrarily in depths.⁶ One application of such a property is when players make assumptions on the limits of their opponent’s knowledge, and believe that the latter’s initial knowledge is only a subset (or superset) of what they initially believe. We remark that regression-based reasoning as in the above theorem can be applied to the evaluation of game strategies, among other things (Osborne and Rubinstein 1994). For example, verifying whether a strategy is *effective* in enabling some goal α amounts to determining whether α holds at the end of every candidate move by an opponent, all from the initial situation. And owing to Thm 6, such goals can include epistemic operators. In other words, a can reason if her actions lead to b learning of some critical information that would hamper her chances of victory; a relevant feature in card games that involve players exchanging or replacing cards.

Related Work and Conclusions

This paper advocates the use of the situation calculus as an alternate, and more expressive, representation to the extensive form, and a tool to reason about non-trivial dynamic and epistemic properties via regression. Our general idea relies on formulating actions in a logical framework, based on how they change the world.

⁶Then, \mathcal{R} is defined *wrt.* $k + j + 1$ theories, with the k corresponding to the k -levels in a ’s epistemic state, j for b ’s epistemic state and Σ . Note that the result is only applicable when only-knowing is used at all depths.

This idea of using logical approaches to reason about information change in imperfect information games is not new. In particular, Van Ditmarsch (2002) has proposed epistemic extensions to dynamic logic to reason about card games. More broadly, approaches that embed epistemic operators into Alternating-time Temporal Logic (ATL) (see (Jamroga and Ågotnes 2007) and references therein) are known to provide strategic operators to reason about action paths. In the latter, it is even possible to quantify strategies in the context of belief, using meta-language notions. But the main idea in these approaches is that the extensive form is reified as Kripke structures, not addressing our concerns about drafting complex game trees in the first place. It is sometimes not even clear when states are allowed to be grouped in an information set (Halpern 2003). In our view, beliefs should come out *naturally* given a *set of rules*. Also note that, knowledge updates in (Van Ditmarsch 2002) are handled using meta-theoretical Kripke frame operators, while ours can be established semantically using functional or truth value sensing. Finally, we believe that in many realistic settings, propositional approaches can be cumbersome. For instance, (Jamroga and Ågotnes 2007) consider a game where players guess codes. In general, such an action, even for a 6-digit number requires 10^6 propositions, and expressing a disjunction over them in a logical framework can be unrealistic. Also, consider the following *auction* scenario involving *a* and *b* bidding alternatively, and *a* has a simple strategy that she bids \$2 more than *b*'s bid. She starts the bid. In principle, at this stage, *b*'s bid can be completely arbitrary and the game tree must contain possibly infinitely many branches for the second move. In contrast, first-order theories can be defined over infinite domains, and indeed, this action is already representable in our logic.

The expressibility of the situation calculus in axiomatizing and reasoning about games is also investigated in (Schulte and Delgrande 2004; Finzi and Lukasiewicz 2007). While the line of work of (Finzi and Lukasiewicz 2007) is mostly in the integration of *Golog* and game-theoretic planning in stochastic games, (Schulte and Delgrande 2004) are closer to our approach. They propose two directions, an isomorphic translation from game trees to situation calculus theories, as well as an axiomatic solution to model games, as encouraged in this paper. However, from the epistemic standpoint, there are some problems. Mainly, they use the knowledge fluent from (Scherl and Levesque 2003) without modification for a many agent framework, and this is possible because they define a partition on the set of situations and assign it to players if it is their turn at move at the situations in the corresponding partition. The problem with a *single* partition is that if it is *a*'s turn to move at situation *s*, then *b*'s beliefs at *s* are undefined. Consequently, the only properties we can consider at a situation are objective beliefs, and this even, only if the player is scheduled to move at the situation. Thus, non-trivial multi-agent beliefs, and a regression property like the one in this paper are not possible.

We intend to consider a number of directions for the future. Reasoning about the quantitative aspects of payoffs, embedding ATL strategic operators, and modeling agents

with noisy actuators are some extensions. Note that, with the latter, game trees need to account for information sets contingent on the success/failure of actions. The difficulty of addressing such issues is another limitation we find in the extensive form. A version of the *representation theorem* (Lakemeyer and Levesque 2004) defined for *k*-structures would allow us to reduce multi-agent beliefs after actions to non-modal *first-order* reasoning about the initial situation. Our main interest lies in investigating results that compare the *compactness* of first-order theories to game trees, as formalisms to represent and reason about games.

References

- Belle, V., and Lakemeyer, G. 2010. Multi-agent only-knowing revisited. In *KR*. AAAI Press.
- Fagin, R.; Halpern, J. Y.; Moses, Y.; and Vardi, M. Y. 1995. *Reasoning About Knowledge*. The MIT Press.
- Finzi, A., and Lukasiewicz, T. 2007. Game-theoretic agent programming in Golog under partial observability. *Lecture Notes in Computer Science* 4314:113.
- Halpern, J., and Lakemeyer, G. 2001. Multi-agent only knowing. *Journal of Logic and Computation* 11(1):41.
- Halpern, J. Y. 2003. A computer scientist looks at game theory. *Games and Economic Behavior* 45(1):114–131.
- Hughes, G. E., and Cresswell, M. J. 1972. *An introduction to modal logic*. Methuen London,.
- Jamroga, W., and Ågotnes, T. 2007. Constructive knowledge: What agents can achieve under incomplete information. *J. of Applied Non-Classical Logics* 17(4):423–475.
- Koller, D., and Pfeffer, A. 1997. Representations and solutions for game-theoretic problems. *Artif. Int.* 94(1):167–215.
- Lakemeyer, G., and Levesque, H. J. 2004. Situations, si! situation terms, no! In *KR*, 516–526.
- Lakemeyer, G., and Levesque, H. 2007. Cognitive robotics. *Handbook of Knowledge Representation*. Elsevier.
- Levesque, H. J. 1990. All I know: a study in autoepistemic logic. *Artif. Intell.* 42(2-3):263–309.
- McCarthy, J., and Hayes, P. J. 1969. Some philosophical problems from the standpoint of artificial intelligence. In *Machine Intelligence*, volume 4, 463–502.
- Osborne, M. J., and Rubinstein, A. 1994. *A Course in Game Theory*. The MIT Press.
- Reiter, R. 2001. *Knowledge in action : logical foundations for specifying and implementing dynamical systems*. MIT Press.
- Scherl, R. B., and Levesque, H. J. 2003. Knowledge, action, and the frame problem. *Artif. Intell.* 144(1-2):1–39.
- Schulte, O., and Delgrande, J. 2004. Representing von neumann–morgenstern games in the situation calculus. *Annals of Mathematics and Artif. Intell.* 42(1-3):73–101.
- Thielscher, M. 1999. From situation calculus to fluent calculus: state update axioms as a solution to the inferential frame problem. *Artif. Intell.* 111(1-2):277–299.
- Van Ditmarsch, H. P. 2002. Descriptions of game actions. *J. of Logic, Lang. and Inf.* 11(3):349–365.