

Knowledge and Control

Wiebe van der Hoek
Dept of Computer Science
University of Liverpool, UK
wiebe@csc.liv.ac.uk

Nicolas Troquard
School of Computer Science
and Electronic Engineering
University of Essex, UK
ntroqu@essex.ac.uk

Michael Wooldridge
Dept of Computer Science
University of Liverpool, UK
mjw@csc.liv.ac.uk

ABSTRACT

Logics of propositional control, such as van der Hoek and Wooldridge's CL-PC [14], were introduced in order to represent and reason about scenarios in which each agent within a system is able to exercise unique control over some set of system variables. Our aim in the present paper is to extend the study of logics of propositional control to settings in which these agents have incomplete information about the society they occupy. We consider two possible sources of incomplete information. First, we consider the possibility that an agent is only able to "read" a subset of the overall system variables, and so in any given system state, will have partial information about the state of the system. Second, we consider the possibility that an agent has incomplete information about which agent controls which variables. For both cases, we introduce a logic combining epistemic modalities with the operators of CL-PC, investigate its axiomatization, and discuss its properties.

Categories and Subject Descriptors

I.2.11 [Distributed Artificial Intelligence]: Multiagent Systems;
I.2.4 [Knowledge representation formalisms and methods]

General Terms

Theory

Keywords

epistemic logic, propositional control, partial observability

1. INTRODUCTION

The *Coalition Logic of Propositional Control* (CL-PC) was introduced by van der Hoek and Wooldridge as a formalism for reasoning about how agents and coalitions can exercise control in multi-agent environments [14]. The logic models situations in which each agent has control over some set of propositions; that is, each agent is associated with some set of propositions, and has the ability to assign a (truth) value to each of its propositions. In this way, valuations become possible worlds (see e.g., [15] for an early treatment of such modelings). The language of CL-PC provides modal constructs $\diamond_i \varphi$ to express the fact that, under the assumption that the rest of the system remains unchanged, agent i can assign values to the propositions under its control in such a way that

φ becomes true; these operators are closely related to the strategic ability operators in cooperation logics such as ATL [2] and Coalition Logic [11]. Since the logic was originally presented, a number of variants of CL-PC have been developed. For example: the logic DCL-PC is an extension to CL-PC in which agents are able to transfer the control of their variables by executing *transfer programs* [12]; and Gerbrandy studied generalisations of CL-PC, allowing for instance situations in which agents have "partial" control of propositions [7].

Our aim in this paper is to study one rather obvious aspect of propositional control logics that has hitherto been neglected: *the interaction between knowledge and control*. It is indeed surprising that this aspect of propositional control logics has not been previously studied in the literature. After all, the interaction between knowledge and ability has a venerable history in the artificial intelligence community, going back at least to the work of Moore in the late 1970s [10]. Moore was interested in *knowledge pre-conditions*: what an agent needs to know in order to be able to do something. To use a standard example, in order to be able to open a safe, you need to know the combination. He formalised a notion of ability that was able to capture such subtleties in a logic that combined elements of dynamic and epistemic logic. More recently, the interplay between ability and knowledge has been studied with respect to cooperation logics such as ATL and Coalition Logic. For example, van der Hoek and Wooldridge proposed ATEL, a variant of ATL extended with epistemic modalities [13]; and various authors developed variants of ATEL intended to rectify some counterintuitive properties of the original ATEL proposal [8, 1].

Epistemic logic is, ultimately, a logic modelling (un)certainly [6]. When we say an agent knows φ , we typically mean that the agent is certain about φ . This notion of uncertainty is elegantly captured in possible worlds semantics, where knowing φ means that φ is true in all worlds that the agent considers possible. If we turn to CL-PC, we can identify several different sources of uncertainty, as follows.

First, and most obviously, an agent may be uncertain about the value of the variables in the system. We call this type of uncertainty *partial observability*, and it is very naturally modelled by assigning to every agent a set of variables that the agent is able to "see". Partial observability interacts with control in several important ways. For example, if I control the variable q and my goal is to achieve the formula $p \leftrightarrow \neg q$, then if I can observe the value of p , I can readily choose a value for q that will result in my goal being achieved: I simply choose the opposite to the value of p . However, if I cannot see the value of p , then I am in trouble. Second, and perhaps more unusually, there may be uncertainty about *which agent controls which variables*. Here, we might conceivably have a situation in which an agent is able to bring about some state of affairs, but does not know that they are able to bring it about, because it is not

Cite as: Knowledge and Control, W. van der Hoek, N. Troquard and M. Wooldridge, *Proc. of 10th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2011)*, Tumer, Yolum, Sonenberg and Stone (eds.), May, 2–6, 2011, Taipei, Taiwan, pp. 719-726. Copyright © 2011, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

aware that it controls the appropriate variables.

The aim of the present paper is to develop extensions to CL-PC that are able to capture these two types of uncertainty. The remainder of the paper is structured as follows. After presenting some definitions that will be used throughout the remainder of the paper, in Section 2, we present the epistemic extension to CL-PC for the case that agents have complete knowledge about how the control of variables is actually distributed over the agents, but they may lack information about what is factually true. Subsequently, in Section 3, we then look at formalising the case where agents have full knowledge about factual truth, partial knowledge about who controls what, and are completely ignorant about other's information regarding control. We also sketch an even more general setting where both factual truth and control may be uncertain. We conclude in Section 4.

We begin with some definitions, which are used throughout the remainder of the paper. First, let $\mathbb{B} = \{true, false\}$ be the set of Boolean truth values. We assume that the domains we model contain a (finite, non-empty) set $N = \{1, \dots, n\}$ of agents ($|N| = n, n > 0$). The environment is also assumed to contain a (fixed, finite) set $\mathbb{A} = \{p, q, \dots\}$ of *Boolean variables*. Each agent $i \in N$ will be assumed to *control* some subset \mathbb{A}_i of atoms \mathbb{A} , with the intended interpretation that if $p \in \mathbb{A}_i$, then i has the unique ability to assign a value (*true* or *false*) to p . We require that the sets \mathbb{A}_i form a partition of \mathbb{A} , i.e., $\mathbb{A}_i \cap \mathbb{A}_j = \emptyset$ for $i \neq j$, and $\mathbb{A}_1 \cup \dots \cup \mathbb{A}_n = \mathbb{A}$. Thus every variable is controlled by some agent; and no variable is controlled by more than one agent. A *coalition* is simply a set of agents, i.e., a subset of N . We typically use C, C', \dots as variables standing for coalitions. Where $C \subseteq N$, we denote by \mathbb{A}_C the set of variables under the collective control of the agents in C : $\mathbb{A}_C = \bigcup_{i \in C} \mathbb{A}_i$. A *valuation* is a total function $\theta : \mathbb{A} \rightarrow \mathbb{B}$, which assigns a truth value to every Boolean variable. Let Θ denote the set of all valuations. Where C is a coalition, a C -valuation is a function $\theta_C : \mathbb{A}_C \rightarrow \mathbb{B}$; thus a C -valuation is a valuation to variables under the control of the agents in C . Given a set X of atoms and two valuations θ_1 and θ_2 , we write $\theta_1 \equiv_X \theta_2$ to mean that θ_1 and θ_2 agree on the value of all variables in X , i.e., $\theta_1(p) = \theta_2(p)$ for all $p \in X$.

2. PARTIAL OBSERVABILITY

In this section, we develop an *Epistemic Coalition Logic of Propositional Control with Partial Observability* – ECL-PC(PO) for short. This logic is essentially CL-PC extended with epistemic modalities K_i , one for each agent $i \in N$. These epistemic modalities have a conventional (S5) possible worlds semantics. The interpretation we give to epistemic accessibility relations is as follows. We assume each agent $i \in N$ is able to see a subset $V_i \subseteq \mathbb{A}$ of the overall set of Boolean variables; that is, it is able to correctly perceive the value of these variables. A valuation θ' is then i -accessible from valuation θ if θ and θ' agree on the valuation of variables visible to i , i.e., $\theta \equiv_{V_i} \theta'$. Formally, the language of ECL-PC(PO) is defined by the following BNF grammar:

$$\varphi ::= p \mid \neg\varphi \mid \varphi \vee \psi \mid \diamond_i \varphi \mid K_i \varphi$$

where $p \in \mathbb{A}$, and $i \in N$. As in CL-PC [14], a formula $\diamond_i \varphi$ means that i can assign values to the variables under its control in such a way that, assuming no other variables are changed, φ becomes true. As in epistemic logic [6], a formula $K_i \varphi$ means that the agent i knows φ .

The remaining operators of classical logic (“ \wedge ” – and, “ \rightarrow ” – implies, “ \leftrightarrow ” – iff) are assumed to be defined as abbreviations in terms of \neg, \vee as usual. We define the box dual operator of \diamond_i as: $\square_i \varphi \equiv \neg \diamond_i \neg \varphi$. We also assume the existential dual M_i (“maybe”)

of the K_i operator is defined as: $M_i \varphi \equiv \neg K_i \neg \varphi$. For coalitions, we define (this definition is justified in [14]):

$$\square_{\{1, \dots, k\}} \varphi \equiv \square_1 \dots \square_k \varphi.$$

Coming to the semantics, a *frame* for CL-PC is simply a structure $\langle N, \mathbb{A}_1, \dots, \mathbb{A}_n \rangle$, where N is the set of agents in the system, and each \mathbb{A}_i is the set of variables under the control of agent i ; a model for CL-PC combines such a frame with a valuation $\theta \in \Theta$, which gives an initial value for every Boolean variable [14]. Frames for ECL-PC(PO) extend CL-PC frames with a set of variables $V_i \subseteq \mathbb{A}$ for each agent $i \in N$. Formally, an ECL-PC(PO) frame, F , is a $(2n + 1)$ -tuple

$$F = \langle N, \mathbb{A}_1, \dots, \mathbb{A}_n, V_1, \dots, V_n \rangle, \text{ where}$$

- $N = \{1, 2, \dots, n\}$ is a (finite, nonempty) set of agents.
- The sets \mathbb{A}_i form a partition of \mathbb{A} .
- $V_i \subseteq \mathbb{A}$ is the set of atoms whose values are visible to i .

It will often make sense to assume $V_i \supseteq \mathbb{A}_i$, i.e., each agent can see the value of the variables it controls; however, we will not impose this as a requirement. We leave aside the question for now of what settings there are in which this assumption does not hold.

The truth value of an ECL-PC(PO) formula is inductively defined wrt. a frame F and a valuation θ by the following rules (\models^d stands for a ‘direct semantics’, [14]):

$$\begin{aligned} F, \theta \models^d p & \quad \text{iff} \quad \theta(p) = true & (p \in \mathbb{A}) \\ F, \theta \models^d \neg\varphi & \quad \text{iff} \quad F, \theta \not\models^d \varphi \\ F, \theta \models^d \varphi \vee \psi & \quad \text{iff} \quad F, \theta \models^d \varphi \text{ or } F, \theta \models^d \psi \\ F, \theta \models^d \diamond_i \varphi & \quad \text{iff} \quad \exists \theta' \in \Theta : \theta' \equiv_{\mathbb{A} \setminus \mathbb{A}_i} \theta \text{ s.t. } M, \theta' \models^d \varphi \\ F, \theta \models^d K_i \varphi & \quad \text{iff} \quad \forall \theta' \in \Theta : \theta' \equiv_{V_i} \theta \implies M, \theta' \models^d \varphi \end{aligned}$$

We denote the fact that φ is true in all models by $\models^d \varphi$. We let $\Lambda_1 = \{\varphi \mid \models^d \varphi\}$ be the logic of all the formulas valid in all ECL-PC(PO) models.

EXAMPLE 1. *Suppose we have a frame F with two agents, $N = \{1, 2\}$ and two Boolean variables, $\mathbb{A} = \{p, q\}$, with $\mathbb{A}_1 = V_1 = \{p\}$ and $\mathbb{A}_2 = \{q\}$ and $V_2 = \{p, q\}$. Thus agent 1 can only see the value of the variable it controls, while agent 2 can see the values of both variables. Let $\theta(p) = \theta(q) = true$. Now, we have:*

- $F, \theta \models^d \diamond_1(p \leftrightarrow \neg q)$
Agent 1 can set his variable p in such a way that p and q have different values.
- $F, \theta \models^d \neg K_1 q \wedge \neg K_1 \neg q \wedge K_1(K_2 q \vee K_2 \neg q)$
Agent 1 does not know the value of variable q , but he does know that 2 knows the value of q .
- $F, \theta \models^d K_1 \diamond_1(p \leftrightarrow \neg q) \wedge \neg \diamond_1 K_1(p \leftrightarrow \neg q)$
Agent 1 knows that he can make p and q take on different values (because he controls p , and hence can make it different to q in any given state). However, agent 1 cannot choose values for the variables he controls in such a way that he knows that p and q take on different values.
- $F, \theta \models^d K_2 \square_1((K_2 p \vee K_2 \neg p) \wedge (K_2 q \vee K_2 \neg q))$
Agent 2 knows that whatever truth values 1 chooses for her variables, 2 will know the value of p and of q .
- $F, \theta \models^d K_2((p \wedge q) \wedge \diamond_1(\neg p \wedge \diamond_2(\neg p \wedge \neg q)))$ *Agent 2 knows that $(p \wedge q)$ and that 1 can bring about that $\neg p$ which 2 can further narrow down to $(\neg p \wedge \neg q)$.*

CLPC	
(Prop)	φ , where φ is a propositional tautology
(K(\square))	$\square_i(\varphi \rightarrow \psi) \rightarrow (\square_i\varphi \rightarrow \square_i\psi)$
(T(\square))	$\square_i\varphi \rightarrow \varphi$
(B(\square))	$\varphi \rightarrow \square_i\diamond_i\varphi$
(empty)	$\square_\emptyset\varphi \leftrightarrow \varphi$
(comp \cup)	$\square_{C_1}\square_{C_2}\varphi \leftrightarrow \square_{C_1\cup C_2}\varphi$
(confl)	$\diamond_i\square_j\varphi \rightarrow \square_j\diamond_i\varphi$
(exclu)	$(\diamond_i p \wedge \diamond_i \neg p) \rightarrow (\square_j p \vee \square_j \neg p)$, where $j \neq i$
(actual)	$\bigvee_{i \in N} \diamond_i p \wedge \diamond_i \neg p$
(full \square)	$(\bigwedge_{p \in X} \diamond_i p \wedge \diamond_i \neg p) \rightarrow \diamond_i \varphi_X$
Knowledge	
(K(K))	$K_i(\varphi \rightarrow \psi) \rightarrow (K_i\varphi \rightarrow K_i\psi)$
(T(K))	$K_i\varphi \rightarrow \varphi$
(B(K))	$\varphi \rightarrow K_i M_i \varphi$
(A(K))	$K_i\varphi \rightarrow K_i K_i \varphi$
(incl)	$\square_N \varphi \rightarrow K_i \varphi$
(unif)	$M_i p \wedge M_i \neg p \rightarrow \square_N (M_i p \wedge M_i \neg p)$
(fullK)	$(\bigwedge_{p \in X} M_i p \wedge M_i \neg p) \rightarrow M_i \varphi_X$
Rules	
(MP)	from $\vdash \varphi \rightarrow \psi$ and $\vdash \varphi$ infer $\vdash \psi$
(Nec(\square))	from $\vdash \varphi$ infer $\vdash \square_i \varphi$

Figure 1: Axiomatics of Λ_1 . The meta-variable i ranges over N , C_1 and C_2 over 2^N , φ represents an arbitrary formula of ECL-PC(PO), p ranges over \mathbb{A} . φ_X is the conjunction of literals true in any valuation of $X \subseteq \mathbb{A}$.

When studying a new logic, there are two key computational problems we must consider: the *model checking* problem and the *satisfiability* problem. For ECL-PC(PO), the model checking problem is the problem of determining, for a given frame and valuation F, θ and formula φ , whether or not $F, \theta \models^d \varphi$. The satisfiability problem is the problem of determining whether, for a given formula φ there exists a frame F and valuation θ such that $F, \theta \models^d \varphi$. It was proved in [14] that both the model checking and satisfiability problems for the underlying logic CL-PC are PSPACE-complete. The fact that the model checking problem is PSPACE complete in fact yielded a decision problem for satisfiability: because the frames F are “small”, we can exhaustively search the set of possible frames and valuations for a formula, checking each pair in turn to see whether it satisfies the formula. Observe that the model checking problem for ECL-PC(PO) is trivially seen to be solvable in polynomial space. Then basically the same approach for satisfiability checking of CL-PC also works for ECL-PC(PO): the truth of a formula only depends on at most one more agent than is named in the formula, as with CL-PC [14], and so we can exhaustively examine each F, θ pair to see whether $F, \theta \models^d \varphi$. We may conclude:

THEOREM 1. *The model checking and satisfiability problems for ECL-PC(PO) are both PSPACE-complete.*

An axiomatization for ECL-PC(PO) is provided in Figure 1. Several points are in order with respect to this axiomatization. First, note that K_i is an S5 modality, and that the axiom 4 for the modality \square_i is an instance of axiom (comp \cup). With $K(\square)$, $T(\square)$ and $B(\square)$ this implies that \square_i is also an S5 modality.

LEMMA 1. *The axiomatization for Λ_1 in Figure 1 is sound.*

We now prove that this axiomatization is complete. This will be done using a normal form for ECL-PC(PO)-formulas.

DEFINITION 1. *We define $ctrls(i, p)$ as $(\diamond_i p \wedge \diamond_i \neg p)$ and $sees(i, p)$ as $(K_i p \vee K_i \neg p)$. Let $CTRL = \{ctrls(i, p) \mid i \in N \ \& \ p \in \mathbb{A}\}$ and $VIEW = \{sees(i, p) \mid i \in N \ \& \ p \in \mathbb{A}\}$.*

The elements of \mathbb{A} , $CTRL$ and $VIEW$ are called *basic propositions*. For any set Φ of basic propositions, call $L(\Phi) = \{x, \neg x \mid x \in \Phi\}$ the set of literals over Φ . For a basic proposition x , let $\ell(x) \in \{x, \neg x\}$. So e.g., $\ell(p) \rightarrow \diamond_i \ell(p)$ stands both for $p \rightarrow \diamond_i p$ and for $\neg p \rightarrow \diamond_i \neg p$. A propositional description π is a conjunction over $L(\mathbb{A})$ where each $p \in \mathbb{A}$ occurs exactly once. Let Π be the set of propositional descriptions. A control description γ is a conjunction over $CTRL$ such that for every $p \in \mathbb{A}$, there is exactly one $i \in N$ such that $ctrls(i, p)$ occurs in γ . Let Γ be the set of control descriptions. Finally, a visibility description ς is a conjunction over $L(VIEW)$, such that for every agent i and every atom $p \in \mathbb{A}$, either $sees(i, p)$ or $\neg sees(i, p)$ occurs in ς . Let Σ be the set of visibility descriptions. A full description is a conjunction $\pi \wedge \gamma \wedge \varsigma$, where π, γ and ς are as explained above.

Given a propositional description $\pi \in \Pi$, we shall note $\hat{\pi}^i$ the conjunction of literals in π that are under the control of agent i and $\tilde{\pi}^i$ the conjunction of literals in π that are not under its control. Of course $\pi \leftrightarrow \hat{\pi}^i \wedge \tilde{\pi}^i$. In the same vein, we shall note $\hat{\pi}^i$ the conjunction of literals in π that are seen by agent i and $\tilde{\pi}^i$ the conjunction of literals in π that are not seen by it. Again $\pi \leftrightarrow \hat{\pi}^i \wedge \tilde{\pi}^i$.

As its name suggests, a full description $(\pi \wedge \gamma \wedge \varsigma)$ fully characterises a situation: it specifies which atoms are true and which are false (this is π), it specifies which agents control which variables (through γ) and it specifies exactly which propositional variables each agent can see (through ς). So semantically, it is immediately clear that any formula will be a disjunction of such full descriptions (namely, descriptions of those situations where φ is true), but our task is now to show that this is derivable in the logic.

The next Lemma states a few theorems derivable within our axiomatic system, all of which are instrumental in the proofs of Theorem 2 and of Theorem 3.

LEMMA 2. *Let π, γ and ς be propositional, control and visibility descriptions, respectively (and so are their ‘primed’ version). For $P \subseteq \mathbb{A}$, let $\pi_1(L(P))$ be a conjunction over $L(P)$ and let $\pi_2(L(\mathbb{A} \setminus P))$ be a conjunction over $L(\mathbb{A} \setminus P)$.*

Then, the following are derivable in Λ_1 :

1. $\neg ctrls(i, p) \rightarrow (\ell(p) \rightarrow \square_i \ell(p))$
2. $sees(i, p) \rightarrow (\ell(p) \rightarrow K_i \ell(p))$
3. $\ell(ctrls(i, p)) \leftrightarrow \square_N \ell(ctrls(i, p))$
4. $\ell(sees(i, p)) \leftrightarrow \square_N \ell(sees(i, p))$
5. $\bigwedge_{p \in P} ctrls(i, p) \wedge \bigwedge_{p \notin P} \neg ctrls(i, p) \rightarrow \diamond_i \pi(L(P)) \wedge (\pi_2(L(\mathbb{A} \setminus P)) \rightarrow \square_i \pi_2(L(\mathbb{A} \setminus P)))$
6. $\diamond_i (\hat{\pi}^i \wedge \tilde{\pi}^i) \leftrightarrow \hat{\pi}^i$
7. $\bigwedge_{p \in P} sees(i, p) \wedge \bigwedge_{p \notin P} \neg sees(i, p) \rightarrow M_i \pi_2(L(\mathbb{A} \setminus P)) \wedge (\pi_1(L(P)) \rightarrow K_i \pi_1(L(P)))$
8. $M_i (\hat{\pi}^i \wedge \tilde{\pi}^i) \leftrightarrow \hat{\pi}^i$
9. $\square_N \varphi \leftrightarrow \square_i \square_N \varphi$
10. $\square_N \varphi \leftrightarrow K_i \square_N \varphi$
11. $(\pi \wedge \gamma \wedge \varsigma) \leftrightarrow (\pi \wedge \square_N \gamma \wedge \square_N \varsigma)$

THEOREM 2 (NORMAL FORM). *Every formula φ is provably equivalent to a disjunction of full descriptions, i.e., for every φ there exists a k and π_j, γ_j and ς_j ($1 \leq j \leq k$) such that*

$$\vdash \varphi \leftrightarrow \bigvee_{j \leq k} (\pi_j \wedge \square_N \gamma_j \wedge \square_N \varsigma_j) \quad (1)$$

PROOF. By Lemma 2.11, it follows from

$$\vdash \varphi \leftrightarrow \bigvee_{j \leq k} (\pi_j \wedge \gamma_j \wedge \varsigma_j)$$

which we prove now by induction on the structure of φ .

We will make use of the fact that the sets of propositional (Π), control (Γ) and visibility (Σ) descriptions are finite. Roughly speaking, a triple (π, γ, ς) represents a state. The idea behind the normal form is, that a formula can be represented by a subset $X \subseteq \Pi \times \Gamma \times \Sigma$, which translates in the language as a (typically large) disjunction of formulas of the form $\pi \wedge \gamma \wedge \varsigma$.

One base case is for φ being a basic proposition in Φ .

$$\vdash p \leftrightarrow \bigvee_{\substack{\pi_i \in \Pi \\ \pi_i \vdash p}} \bigvee_{\gamma_j \in \Gamma} \bigvee_{\varsigma_k \in \Sigma} (\pi_i \wedge \gamma_j \wedge \varsigma_k)$$

The statement $\pi_i \vdash p$ means that p appears as a positive literal in π_i . The two other base cases $\varphi = \text{ctrls}(i, p)$ and $\varphi = \text{sees}(i, p)$ are analogous.

Now we suppose for induction that φ can be transformed into an equivalent formula $\bigvee_{j \leq k} (\pi_j \wedge \gamma_j \wedge \varsigma_j)$.

Case $\psi = \neg\varphi$: “ ψ is represented by the complement of the states representing φ ”.

$$\vdash \psi \leftrightarrow \bigvee_{j \leq k} \bigvee_{\substack{(\pi, \gamma, \varsigma) \in \Pi \times \Gamma \times \Sigma \\ (\pi, \gamma, \varsigma) \neq (\pi_j, \gamma_j, \varsigma_j)}} (\pi \wedge \gamma \wedge \varsigma)$$

Case $\psi = \varphi_1 \vee \varphi_2$: since the normal form itself is a disjunction, this case is straightforward.

Case $\psi = \diamond_i \varphi$: similar to $\psi = M_i \varphi$.

Case $\psi = M_i \varphi$: by induction hypothesis

$$\vdash \psi \leftrightarrow M_i \bigvee_{j \leq k} (\pi_j \wedge \gamma_j \wedge \varsigma_j)$$

By modal logic

$$\vdash \psi \leftrightarrow \bigvee_{j \leq k} M_i (\pi_j \wedge \gamma_j \wedge \varsigma_j)$$

By Lemma 2.10 and Lemma 2.11

$$\vdash \psi \leftrightarrow \bigvee_{j \leq k} M_i (\pi_j \wedge K_i \Box_N \gamma_j \wedge K_i \Box_N \varsigma_j)$$

By $S5(K)$

$$\vdash \psi \leftrightarrow \bigvee_{j \leq k} (M_i \pi_j \wedge K_i \Box_N \gamma_j \wedge K_i \Box_N \varsigma_j)$$

Applying our notation and Lemma 2.11 and Lemma 2.10

$$\vdash \psi \leftrightarrow \bigvee_{j \leq k} (M_i (\hat{\pi}_j^i \wedge \hat{\pi}_j^i) \wedge \gamma_j \wedge \varsigma_j)$$

By Lemma 2.8

$$\vdash \psi \leftrightarrow \bigvee_{j \leq k} (\hat{\pi}_j^i \wedge \gamma_j \wedge \varsigma_j)$$

Finally,

$$\vdash \psi \leftrightarrow \bigvee_{j \leq k} \bigvee_{\pi_j \in \Pi(\hat{\pi}_j^i)} ((\hat{\pi}_j^i \wedge \pi_j) \wedge \gamma_j \wedge \varsigma_j)$$

where $\Pi(\hat{\pi}_j^i)$ is the set of propositional descriptions restricted to the set of atoms occurring in $\hat{\pi}_j^i$, that is, that are not seen by i . \square

We require some subsidiary definitions. We begin by defining an alternative, possible worlds semantics for ECL-PC(PO). Given a frame F , a *Kripke model for ECL-PC(PO)* is a structure

$$K = \langle W, R_1^\diamond, \dots, R_n^\diamond, R_1^K, \dots, R_n^K, \pi \rangle$$

where $W = \Theta$ is a *set of worlds*, which correspond to possible valuations to \mathbb{A} , $R_i^\diamond \subseteq W \times W$, and $R_i^K \subseteq W \times W$, where these latter relations are defined as:

$$R_i^\diamond(w, w') \text{ iff } w \equiv_{\mathbb{A} \setminus \mathbb{A}_i} w', \text{ and } R_i^K(w, w') \text{ iff } w \equiv_{v_i} w'.$$

Finally, $\pi : W \rightarrow 2^{\mathbb{A}}$ gives the set of Boolean variables true at each world. The key clauses for \models^k (‘Kripke semantics’) are then as follows:

$$\begin{aligned} K, w \models^k p & \text{ iff } p \in \pi(w) & (p \in \mathbb{A}) \\ K, w \models^k \diamond_i \varphi & \text{ iff } \exists w' \in W \text{ s.t. } R_i^\diamond(w, w') \text{ and } K, w' \models^k \varphi \\ K, w \models^k K_i \varphi & \text{ iff } \forall w' \in W \text{ s.t. } R_i^K(w, w') \text{ and } K, w' \models^k \varphi \end{aligned}$$

LEMMA 3. *Let F, θ be an ECL-PC(PO) frame and associated valuation, let K, w be the corresponding Kripke model and world, and let φ be an arbitrary ECL-PC(PO) formula. Then:*

$$F, \theta \models^d \varphi \text{ iff } K, w \models^k \varphi.$$

We assume the standard definitions of maximally consistent sets and their existence via Lindenbaum’s lemma (see, e.g., [4, p.196]). We proceed to construct a canonical model

$$\hat{K} = \langle \hat{W}, \hat{R}_1^\diamond, \dots, \hat{R}_n^\diamond, \hat{R}_1^K, \dots, \hat{R}_n^K, \hat{\pi} \rangle, \text{ where:}$$

- \hat{W} is the set of all Λ_1 maximally consistent sets;
- $\hat{R}_i^\diamond(w, w')$ iff $\varphi \in w'$ implies $\diamond_i \varphi \in w$;
- $\hat{R}_i^K(w, w')$ iff $\varphi \in w'$ implies $M_i \varphi \in w$; and
- $\hat{\pi}(w) = \mathbb{A} \cap w$.

The following is a standard result for canonical models:

LEMMA 4 (TRUTH LEMMA.). *Let $\hat{K} = \langle \hat{W}, \hat{R}_1^\diamond, \dots, \hat{R}_n^\diamond, \hat{R}_1^K, \dots, \hat{R}_n^K, \hat{\pi} \rangle$ be a canonical model, $w \in \hat{W}$ be a world in \hat{K} , and φ be an arbitrary ECL-PC(PO) formula. Then:*

$$\hat{K}, w \models^k \varphi \text{ iff } \varphi \in w.$$

The truth lemma above gives rise to completeness wrt. a set of models, but it is not the kind of models we have associated with ECL-PC(PO). In the intended models, the modalities K_i and \diamond_i are defined with respect to valuations that are ‘similar’ with respect to the appropriate sets of atoms, while in the canonical model, those modal operators are defined as necessity operators with respect to a relation between maximal consistent sets that is defined in terms of membership of formulas in these sets. We now have to show that, in the canonical model, these two ways of looking at the modalities coincide. For this, our normal form Theorem 2 will be crucial.

But first we restrict ourselves to a generated submodel of \hat{K} . To be more precise, for the canonical model \hat{K} just obtained, and $w \in \hat{W}$, let $\hat{K}_{\bar{w}}$ be the model generated by w in the following sense. Let \hat{R}_N^\diamond be $\hat{R}_1^\diamond \cup \dots \cup \hat{R}_n^\diamond$. Then, define $\hat{W}_{\bar{w}} = \{v \mid \hat{R}_N^\diamond(w, v)\}$, and all relations $\hat{R}_{\bar{w}_i}^\diamond$ and $\hat{R}_{\bar{w}}^K$ and valuation $\hat{\pi}_{\bar{w}}$ are the old relations and valuation restricted to the new set $\hat{W}_{\bar{w}}$. The following is a known result about generated submodels:

$$\forall \varphi \forall v \in \hat{W}_{\bar{w}} \hat{K}, v \models \varphi \text{ iff } \hat{K}_{\bar{w}}, v \models \varphi$$

THEOREM 3 ($\hat{K}_{\bar{w}}$ SIMULATES AN ECL-PC(PO) FRAME.). *Let \hat{K} be as defined above, and take $w \in \hat{W}$. Consider the model $\hat{K}_{\bar{w}}$. Define, for every $i \in N$, the sets $\mathbb{A}_i = \{p \mid \text{ctrls}(i, p) \in w\}$, and $V_i = \{p \mid \text{sees}(i, p) \in w\}$. Then, in $\hat{K}_{\bar{w}}$, the accessibility relations satisfy the following properties:*

1. $\hat{R}_{\bar{w}_i}^\diamond(v, v')$ iff $\pi_{\bar{w}}(v) \equiv_{\mathbb{A} \setminus \mathbb{A}_i} \pi_{\bar{w}}(v')$.
2. $\hat{R}_{\bar{w}_i}^K(v, v')$ iff $\pi_{\bar{w}}(v) \equiv_{V_i} \pi_{\bar{w}}(v')$.

PROOF. Consider the first item. Suppose $\hat{R}_{\bar{w}_i}^\diamond(v, v')$, which means that for any φ , $\varphi \in v' \Rightarrow \diamond_i \varphi \in v$. Take any $p \in \mathbb{A} \setminus \mathbb{A}_i$. We show that $p \in v$ iff $p \in v'$. Suppose $p \in v$. By definition of \mathbb{A}_i , we have $\text{ctrls}(i, p) \notin w$, and, since w is a maximal consistent set, $\neg \text{ctrls}(i, p) \in w$. By Lemma 2, item 4 (take $\ell(\text{ctrls}(i, p)) = \text{ctrls}(i, p)$) we have $\Box_N \neg \text{ctrls}(i, p) \in w$, and, since v is \hat{R}_N^\diamond -reachable from w , we have $\neg \text{ctrls}(i, p) \in v$. This gives $(\neg \text{ctrls}(i, p) \wedge p) \in v$, which, by Lemma 2, item 1 gives us $\Box_i p \in v$. Now for contradiction, if $p \notin v'$, we would have $\neg p \in v'$, and by definition, $\diamond_i \neg p \in v$, which contradicts $\Box_i p \in v$. The reasoning for $p \notin v$ goes similar.

For the converse, suppose $\pi_{\bar{w}}(v) \equiv_{\mathbb{A} \setminus \mathbb{A}_i} \pi_{\bar{w}}(v')$, i.e., $v \cap (\mathbb{A} \setminus \mathbb{A}_i) = v' \cap (\mathbb{A} \setminus \mathbb{A}_i)$. Take an arbitrary $\varphi \in v'$, we have to show that $\diamond_i \varphi \in v$. By Theorem 2, we know that φ is equivalent to a disjunction as specified in (1), and since v' is a maximal consistent set, there must be (uniquely) a propositional description π , a control description γ and a visibility description ς such that $(\pi \wedge \Box_N \gamma \wedge \Box_N \varsigma) \in v'$. Since v and v' are both reachable from the same generating world w , we have $(\Box_N \gamma \wedge \Box_N \varsigma) \in v$ and hence, by (*comp*) and (*T*(\Box))

$$(\Box_i \gamma \wedge \Box_i \varsigma) \in v \quad (2)$$

Let us decompose π into $\pi_1 \wedge \pi_2$, where π_1 uses all the atoms p from \mathbb{A}_i , and π_2 uses all the atoms from $\mathbb{A} \setminus \mathbb{A}_i$. By Lemma 2, item 5, we have

$$\diamond_i \pi_1 \in v \quad (3)$$

Moreover $\pi \in v'$ implies that $\pi_2 \in v'$. Moreover by assumption $v \cap (\mathbb{A} \setminus \mathbb{A}_i) = v' \cap (\mathbb{A} \setminus \mathbb{A}_i)$. Hence, $\pi_2 \in v$. By Lemma 2, item 5, we then have that

$$\Box_i \pi_2 \in v \quad (4)$$

Collecting equations (2), (3) and (4), and using the modal validity $\vdash (\Box \alpha \wedge \diamond \beta) \rightarrow \diamond(\alpha \wedge \beta)$, we obtain $\diamond_i(\pi_1 \wedge \pi_2 \wedge \gamma \wedge \varsigma) \in v$. By Lemma 2.11, we conclude $\diamond_i(\pi_1 \wedge \pi_2 \wedge \Box_N \gamma \wedge \Box_N \varsigma) \in v$ which means that $\diamond_i \varphi \in v$.

We now prove the second item. Suppose $\hat{R}_{\bar{w}_i}^K(v, v')$, which means that for any φ , $\varphi \in v' \Rightarrow M_i \varphi \in v$. Take any $p \in V_i$. We show that $p \in v$ iff $p \in v'$. Suppose $p \in v$. By definition of V_i , we have $\text{sees}(i, p) \in w$. By Lemma 2, item 4 we have $\Box_N \text{sees}(i, p) \in w$, and, since v is \hat{R}_N^\diamond -reachable from w , we have $\text{sees}(i, p) \in v$. This gives $(\text{sees}(i, p) \wedge p) \in v$, which, by Lemma 2, item 2 gives us $K_i p \in v$. Now for contradiction, if $p \notin v'$, we would have $\neg p \in v'$, and by definition, $M_i \neg p \in v$, which contradicts $K_i p \in v$. The reasoning for $p \notin v$ goes similar.

For the converse, suppose $\pi_{\bar{w}}(v) \equiv_{V_i} \pi_{\bar{w}}(v')$, i.e., $v \cap (V_i) = v' \cap (V_i)$. Take an arbitrary $\varphi \in v'$, we have to show that $M_i \varphi \in v$. By Theorem 2, we know that φ is equivalent to a disjunction as specified in (1), and since v' is a maximal consistent set, there must be (uniquely) a propositional description π , a control description γ and a visibility description ς such that $(\pi \wedge \Box_N \gamma \wedge \Box_N \varsigma) \in v'$. Since v and v' are both reachable from the same generating world w , we have $(\Box_N \gamma \wedge \Box_N \varsigma) \in v$ and hence, by (*incl*)

$$(K_i \gamma \wedge K_i \varsigma) \in v \quad (5)$$

Let us decompose π into $\pi_1 \wedge \pi_2$, where π_1 uses all the atoms p from $\mathbb{A} \setminus V_i$, and π_2 uses all the atoms from V_i . By Lemma 2, item 7, we have

$$M_i \pi_1 \in v \quad (6)$$

Moreover $\pi \in v'$ means trivially that $\pi_2 \in v'$. Moreover by assumption $v \cap (V_i) = v' \cap (V_i)$. Hence, $\pi_2 \in v$. By Lemma 2, item 7, we then have that

$$K_i \pi_2 \in v \quad (7)$$

Collecting equations (5), (6) and (7), and using the modal validity $\vdash (\Box \alpha \wedge \diamond \beta) \rightarrow \diamond(\alpha \wedge \beta)$, we obtain $M_i(\pi_1 \wedge \pi_2 \wedge \gamma \wedge \varsigma) \in v$. By Lemma 2.11, we conclude $M_i(\pi_1 \wedge \pi_2 \wedge \Box_N \gamma \wedge \Box_N \varsigma) \in v$ which means that $M_i \varphi \in v$. \square

THEOREM 4 (COMPLETENESS OF Λ_1 .). Λ_1 is sound and complete with respect to the class of ECL-PC(PO) frames.

PROOF. Soundness is observed in Lemma 1. For completeness, take a Λ_1 -consistent formula φ . Consider a maximal consistent set w with $\varphi \in w$. We know that $\hat{K}, w \models \varphi$. Take the generated model $\hat{K}_{\bar{w}}$. We know that again $\hat{K}_{\bar{w}}, w \models \varphi$, and moreover, by Theorem 3, $\hat{K}_{\bar{w}}$ simulates an ECL-PC(PO) frame. \square

3. UNCERTAINTY ABOUT OWNERSHIP

The next type of uncertainty we consider relates to which agents control which variables. We refer to the logic we develop to capture such situations as the ECL-PC(UO), where ‘‘UO’’ stands for ‘‘uncertainty of ownership’’. The syntax of ECL-PC(UO) is identical to that of ECL-PC(PO), and so we will not present the syntax again here. In the semantics however, we substitute for every agent the set of propositions that it can see the value of, with a set of propositions which it sees the ownership of.

Given a set of agents N , atomic variables \mathbb{A} , and control partition $\mathbb{A}_1, \dots, \mathbb{A}_n$, a *controls observation* for agent i is as set $\Omega_i \subseteq \mathbb{A}$. The interpretation of Ω_i is that $p \in \Omega_i$ means that agent i *knows who has control over the variable p* , that is, the agent $j \in N$ such that $p \in \mathbb{A}_j$. Given this, we define a frame F for ECL-PC(UO) as:

$$F = \langle N, \mathbb{A}_1, \dots, \mathbb{A}_n, \Omega_1, \dots, \Omega_n \rangle \text{ where:}$$

- N and $\mathbb{A}_i \subseteq \mathbb{A}$ are as before, and
- Ω_i is the controls observation for agent i .

We now define a relation on *frames*, which will be used to give a semantics to our epistemic modalities. Let

$$F = \langle N, \mathbb{A}_1, \dots, \mathbb{A}_n, \Omega_1, \dots, \Omega_i, \dots, \Omega_n \rangle, \text{ and}$$

$$F' = \langle N, \mathbb{A}'_1, \dots, \mathbb{A}'_n, \Omega'_1, \dots, \Omega'_i, \dots, \Omega'_n \rangle$$

be two frames that contain the same agents and the same base set of propositional variables. Then we write $F \simeq_i F'$ to mean that (1) $\Omega_i = \Omega'_i$ and (2) for all $p \in \Omega_i$ and for all $j \in N$ we have $\mathbb{A}_j \cap \Omega_i = \mathbb{A}'_j \cap \Omega_i$. Thus, roughly, $F \simeq_i F'$ means that F' and F agree on the variables that i can see the ownership of, and moreover, for each of those variables, the control is assigned to the same agents in both frames.

Formally, the key steps in the semantics are defined as follows:

$$\begin{aligned} F, \theta \models^d p & \text{ iff } \theta(p) = \text{true} & (p \in \mathbb{A}) \\ F, \theta \models^d \diamond_i \varphi & \text{ iff } \exists \theta' \in \Theta : \theta' \equiv_{\mathbb{A} \setminus \mathbb{A}_i} \theta \text{ s.t. } M, \theta' \models^d \varphi \\ F, \theta \models^d K_i \varphi & \text{ iff } \forall F' : F' \simeq_i F \implies F', \theta \models^d \varphi \end{aligned}$$

EXAMPLE 2. Suppose we have a frame F in which $N = \{1, 2\}$, $\mathbb{A}_1 = \{p\}$, $\mathbb{A}_2 = \{q\}$, $\Omega_1 = \emptyset$, $\Omega_2 = \{p, q\}$. In this case, agent 1 has no information at all about which agent controls which variable: As far as this agent is concerned, any partition of controlled variables to agents is possible. Let $\theta(p) = \theta(q) = \text{true}$. We have:

- $F, \theta \models^d K_1(p \wedge q) \wedge K_2(p \wedge q)$

Unlike ECL-PC(PO), agents have no uncertainty about the actual value of variables. Thus both agents know that both variables are true in the valuation θ .

- $F, \theta \models^d \diamond_1(\neg p \wedge q) \wedge \neg K_1 \diamond_1(\neg p \wedge q)$

In fact, agent 1 can bring about $\neg p \wedge q$: he controls the variable p and he can choose $\neg p \wedge q$. However, because he is uncertain about whether he controls p , he does not know that he has the ability to choose $\neg p \wedge q$.

- $F, \theta \models^d \diamond_2(p \wedge \neg q) \wedge K_2 \diamond_2(p \wedge \neg q)$

Agent 2 can choose a value for q so as to bring about $p \wedge \neg q$ (assuming agent 1 leaves p unchanged). Moreover, since 2 knows that she controls q , she knows that she can choose $p \wedge \neg q$.

- $F, \theta \models^d K_2((p \wedge q) \wedge \diamond_1(\neg p \wedge \diamond_2 \neg q))$

Agent 2 knows that actually $p \wedge q$ holds, and that 1 can choose a situation where p is false and in which agent 2 furthermore can set q to false.

- $F, \theta \models^d K_1 \square_{\{1,2\}} \diamond_{\{1,2\}}(p \leftrightarrow \neg q) \wedge K_2 \square_1 \diamond_2(p \leftrightarrow \neg q)$

Agent 1 knows that together, the agents can always make the values of p and q different, but agent 2 even knows that, no matter which values 1 chooses for his variables, 2 can achieve a situation such that p and q are different.

Note that, by the same arguments as given for ECL-PC(PO), we may conclude that:

THEOREM 5. The model checking and satisfiability problems for ECL-PC(UO) are both PSPACE-complete.

We give an axiomatization for ECL-PC(UO) in Figure 2. Derivability \vdash in this section refers to that axiomatization. The following definitions and notations are useful.

DEFINITION 2. Define $\text{seeswho}(i, p)$ as $\bigvee_{j \in N} K_i \text{ctrls}(j, p)$. Let $SW = \{\text{seeswho}(j, p) \mid j \in N, p \in \mathbb{A}\}$. The elements of \mathbb{A} , CTRL and SW are our new basic propositions. A controls observation description ω is a full conjunction over SW. We note Ω the set of such controls observation descriptions. A new full description is a conjunction $\pi \wedge \gamma \wedge \omega$, where π, γ and ω are as explained above.

Let $P \subseteq \mathbb{A}$. We define $\text{CTRL}(P) = \{\bigwedge \text{ctrls}(i, p) \mid i \in N, p \in P, \text{ every } p \text{ appears only once}\}$. Finally let $\hat{\omega}^i$ be of the form $\bigwedge_{p \in \mathbb{A}} \ell(\text{seeswho}(i, p))$ and let the formula $\tilde{\omega}^i$ be of the form $\bigwedge_{p \in \text{At}, j \neq i} \ell(\text{seeswho}(j, p))$ such that $\hat{\omega}^i \wedge \tilde{\omega}^i$ is a controls observation description.

As with ECL-PC(PO), a full description $(\pi \wedge \gamma \wedge \omega)$ fully characterises a situation: it specifies which atoms are true and which are false (this is π), it specifies which agents control which variables (through γ) and it specifies exactly which agent is aware of who owns which variables (through ω). So semantically, it is immediately clear that any formula will be a disjunction of such full descriptions (namely, descriptions of those situations where φ is true), but our task is now to show that this is derivable in the logic.

CLPC Knowledge	φ	where φ is a CLPC tautology
(K(K))	$K_i(\varphi \rightarrow \psi) \rightarrow (K_i\varphi \rightarrow K_i\psi)$	
(T(K))	$K_i\varphi \rightarrow \varphi$	
(B(K))	$\varphi \rightarrow K_i M_i \varphi$	
(4(K))	$K_i\varphi \rightarrow K_i K_i \varphi$	
Ax1	$\psi \rightarrow K_i \psi$	when ψ objective
Ax2	$\diamond_N \psi \rightarrow K_i \diamond_N \psi$	when ψ objective
Ax3	$\ell(\text{seeswho}(i, p)) \rightarrow K_i \ell(\text{seeswho}(i, p))$	
Ax4	$\text{seeswho}(i, p) \wedge \ell(\text{ctrls}(j, p)) \rightarrow K_i \ell(\text{ctrls}(j, p))$	
Ax5	$\bigwedge_{p \in P} \neg \text{seeswho}(i, p) \rightarrow M_i(\gamma \wedge \tilde{\omega}^i)$	
Ax6	$\bigwedge_{p \in P} \text{seeswho}(i, p) \rightarrow (\gamma \rightarrow K_i \gamma)$	
Ax7	$M_i \hat{\omega}^i \wedge K_i \tilde{\omega}^i$	
Rules		
(MP)	from $\vdash \varphi \rightarrow \psi$ and $\vdash \varphi$ infer $\vdash \psi$	
(Nec(\square))	from $\vdash \varphi$ infer $\vdash \square_i \varphi$	
(Nec(K_i))	from $\vdash \varphi$ infer $\vdash K_i \varphi$	

Figure 2: Axiomatics of Λ_2 . The meta-variable i ranges over N , φ represents an arbitrary formula of ECL-PC(UO), p ranges over \mathbb{A} . Finally, $\hat{\omega}^i$, and $\tilde{\omega}^i$ are as specified in Definition 2, and $\gamma \in \text{CTRL}(P)$. Objective formulas have no modal operators.

LEMMA 5. The axiomatization for Λ_2 in Figure 2 is sound.

We now prove that this axiomatization is complete.

THEOREM 6 (NORMAL FORM). Every formula φ is provably equivalent to a disjunction of full descriptions, i.e., for every φ there exists a k and π_j, γ_j and ω_j ($1 \leq j \leq k$) such that

$$\vdash \varphi \leftrightarrow \bigvee_{1 \leq j \leq k} \pi_j \wedge \gamma_j \wedge \omega_j$$

The proof of Theorem 6 is omitted for reasons of space. We now define an alternative, possible worlds semantics for ECL-PC(UO). Given a frame $F = \langle N, \mathbb{A}_1, \dots, \mathbb{A}_n, \Omega_1, \dots, \Omega_i, \dots, \Omega_n \rangle$, a corresponding pointed Kripke model for ECL-PC(UO) is a structure

$$K, w_{(F, \theta)} = \langle W, R_1^\diamond, \dots, R_n^\diamond, R_1^K, \dots, R_n^K, \pi, w_{(F, \theta)} \rangle$$

where $W = \Pi \times \Gamma \times \Omega$ is a set of worlds that correspond to a frame and a propositional valuation. For every $w \in W$, we note $w(\pi)$ the propositional description it contains, $w(\gamma)$ the control description, and $w(\omega)$ the controls observation description. Given two states w and w' , a set of propositions X , we have already defined $w(\pi) \equiv_X w'(\pi)$. We define $w(\gamma) \equiv_X^i w'(\gamma)$ to mean that for every $p \in X$, $w(\gamma) \vdash \text{ctrls}(i, p)$ iff $w'(\gamma) \vdash \text{ctrls}(i, p)$. Similarly, we define $w(\omega) \equiv_X^i w'(\omega)$ to mean that for every $p \in X$, $w(\omega) \vdash \text{seeswho}(i, p)$ iff $w'(\omega) \vdash \text{seeswho}(i, p)$. Finally, the world $w_{(\theta, F)}$ is such that $w_{(\theta, F)}(\pi)$ describes θ , $w_{(\theta, F)}(\gamma)$ describes $\mathbb{A}_1, \dots, \mathbb{A}_n$ and $w_{(\theta, F)}(\omega)$ describes $\Omega_1, \dots, \Omega_n$.

The relations $R_i^\diamond \subseteq W \times W$, and $R_i^K \subseteq W \times W$, are defined as follows:

$$R_i^\diamond(w, w') \text{ iff } \begin{cases} w(\pi) \equiv_{\mathbb{A} \setminus \mathbb{A}_i} w'(\pi) \\ w(\omega) = w'(\omega) \\ w(\gamma) = w'(\gamma) \end{cases}$$

and

$$R_i^K(w, w') \text{ iff } \begin{cases} w(\pi) \equiv_{\mathbb{A}} w'(\pi) \\ w(\omega) \equiv_{\Omega_i}^i w'(\omega) \\ w(\gamma) \equiv_{\mathbb{A}_j \cap \Omega_i}^i w'(\gamma) \text{ for all } j \in N \end{cases}.$$

Finally, $\pi : W \rightarrow 2^{\mathbb{A}}$ gives the set of Boolean variables true at each world. We can then define a Kripke semantics for our language,

with the key clauses defined via the satisfiability relation \models^k as follows:

$$\begin{aligned} K, w \models^k p & \text{ iff } p \in \pi(w) & (p \in \mathbb{A}) \\ K, w \models^k \diamond_i \varphi & \text{ iff } \exists w' \in W \text{ s.t. } R_i^\diamond(w, w') \text{ and } K, w' \models^k \varphi \\ K, w \models^k K_i \varphi & \text{ iff } \forall w' \in W \text{ s.t. } R_i^K(w, w') \text{ and } K, w' \models^k \varphi \end{aligned}$$

The following is immediate.

LEMMA 6. *Let F, θ be an ECL-PC(UO) frame and associated valuation, let $K, w_{(\theta, F)}$ be the corresponding Kripke model and world, and let φ be an arbitrary ECL-PC(UO) formula. Then:*

$$F, \theta \models^d \varphi \text{ iff } K, w_{(\theta, F)} \models^k \varphi.$$

The definition of a canonical model \hat{K} for the logic is as before (although the model of course will be different, since the axioms are different!), and the truth lemma holds for this language as well. But in this case, we do not need to restrict ourselves to a generated submodel.

THEOREM 7 (\hat{K} SIMULATES AN ECL-PC(UO) FRAME.).

Let \hat{K} be as defined as above. Define, for every $i \in N$ and $v \in \hat{W}$, the sets $\mathbb{A}_{v_i} = \{p \mid \text{ctrls}(i, p) \in v\}$, and $\Omega_{v_i} = \{p \mid \exists j \in N, K_i \text{ctrls}(j, p) \in v\}$. Then, in \hat{K} , the accessibility relations satisfy the following properties:

1. $\hat{R}_i^\diamond(v, v')$ iff $\begin{cases} \pi(v) \equiv_{\mathbb{A}_{v_i}} \pi(v') \\ v(\omega) = v'(\omega) \\ v(\gamma) = v'(\gamma) \end{cases}$
2. $\hat{R}_i^K(v, v')$ iff $\begin{cases} \pi(v) \equiv_{\mathbb{A}} \pi(v') \\ v(\omega) \equiv_{\Omega_{v_i}}^i v'(\omega) \\ v(\gamma) \equiv_{\mathbb{A}_{v_j} \cap \Omega_{v_j}}^i v'(\gamma) \text{ for all } j \in N \end{cases}$

PROOF. We prove the second item. Suppose that $\hat{R}_i^K(v, v')$. By definition, it means that for all φ , $\varphi \in v'$ implies $M_i \varphi \in v$. We now prove the three properties of the right side of the item. We first show that $p \in v$ iff $p \in v'$. Suppose that $p \in v'$. Then $K_i p \in v'$ by Ax1. By hypothesis we obtain $M_i K_i p \in v$, which by S5 yields $p \in v$. The case $p \notin v'$ is similar.

We now show that $K_i \text{ctrls}(j, p) \in v$ iff $K_i \text{ctrls}(j, p) \in v'$. First, suppose that $K_i \text{ctrls}(j, p) \in v$. Then by hypothesis we have $M_i K_i \text{ctrls}(j, p) \in v$ and $K_i \text{ctrls}(j, p) \in v$ by S5. Second, suppose that $K_i \text{ctrls}(j, p) \notin v'$. Since v' is a m.c. set, $\neg K_i \text{ctrls}(j, p) \in v'$. Then, $M_i M_i \neg \text{ctrls}(j, p) \in v$ which by S5 is equivalent to $M_i \neg \text{ctrls}(j, p) \in v$ and $\neg K_i \text{ctrls}(j, p) \in v$. And since v is a m.c. set, we have $K_i \text{ctrls}(j, p) \notin v$.

Now, take any $j \in N$ and any $p \in \mathbb{A}_{v_j'} \cap \Omega_{v_j'}$. We show that $\text{ctrls}(j, p) \in v$ iff $\text{ctrls}(j, p) \in v'$. First, suppose that $\text{ctrls}(j, p) \in v'$. By definition of $\Omega_{v_j'}$, we have $K_j \text{ctrls}(j, p) \in v'$. By hypothesis, we have $M_j K_j \text{ctrls}(j, p) \in v$ which in S5 is equivalent to $\text{ctrls}(j, p) \in v$. Second, suppose that $\text{ctrls}(j, p) \notin v'$. Since v' is an m.c. set, $\neg \text{ctrls}(j, p) \in v'$. Also, by definition of $\Omega_{v_j'}$, we have $\text{seeswho}(i, p) \in v'$. Hence, by Axiom Ax4 we have $K_i \neg \text{ctrls}(j, p) \in v'$. Hence, we have $M_i K_i \neg \text{ctrls}(j, p) \in v$ which in S5 is equivalent to $\neg \text{ctrls}(j, p) \in v$, and since v is a m.c. set we obtain $\text{ctrls}(j, p) \notin v$.

We now prove the right to left direction of item 2. To do so, suppose that (hyp1) $\pi(v) \equiv_{\mathbb{A}} \pi(v')$, (hyp2) $v(\omega) \equiv_{\Omega_{v_i}}^i v'(\omega)$ and (hyp3) $v(\gamma) \equiv_{\mathbb{A}_{v_j} \cap \Omega_{v_j}}^i v'(\gamma)$ for all $j \in N$. We need to show that $\hat{R}_i^K(v, v')$, that is, for all φ we have $\varphi \in v'$ implies $M_i \varphi \in v$.

Take an arbitrary $\varphi \in v'$. By Theorem 6, we assume w.l.o.g. that for some k we have $\varphi \leftrightarrow \bigvee_{1 \leq j \leq k} (\pi_j \wedge \gamma_j \wedge \omega_j)$.

Since v' is an m.c. set, there is (uniquely) a full description $\pi \wedge \gamma \wedge \omega$ such that $(\pi \wedge \gamma \wedge \omega) \in v'$.

From (hyp1) we have $\pi \in v$ and by Ax1 we obtain

$$K_i \pi \in v \quad (8)$$

Let us write ω as $\omega_1 \wedge \omega_2$ such that ω_1 contains the $\ell(\text{seeswho}(i, p))$ literals (those concerning i 's observations) and ω_2 contains all the other literals in ω . Since by (hyp2) we have $v(\omega) \equiv_{\Omega_{v_i}}^i v'(\omega)$, we have $\omega_1 \in v$ and by Axiom Ax3 we get $K_i \omega_1 \in v$. Hence

$$K_i \omega_1 \in v \quad (9)$$

Let us now decompose γ into $\gamma_1 \wedge \gamma_2$ such that γ_1 contains all the $\text{ctrls}(j, p)$ appearing in γ such that $p \in \Omega_{v_j'}$ and ω_2 contains all the other control atoms appearing in γ .

From (hyp3) we know that for all $j \in N$ we have $v(\gamma) \equiv_{\mathbb{A}_{v_j'} \cap \Omega_{v_j'}}^i v'(\gamma)$. Then for all $p \in \mathbb{A}_{v_j'} \cap \Omega_{v_j'}$ and all $j \in N$, we have that $\text{ctrls}(j, p) \in v$ iff $\text{ctrls}(j, p) \in v'$.

Then we have $\gamma_1 \in v$ and by Axiom Ax6 we obtain

$$K_i \gamma_1 \in v \quad (10)$$

Finally, using Axiom Ax5 we obtain

$$M_i(\omega_2 \wedge \gamma_2) \in v \quad (11)$$

Combining (8), (9), (10), and (11) we then obtain $M_i(\pi \wedge \omega \wedge \gamma) \in v$, i.e., $M_i \varphi \in v$. \square

THEOREM 8 (COMPLETENESS OF Λ_2). Λ_2 is sound and complete with respect to the class of ECL-PC(UO) frames.

Let us finally sketch a general setup, in which:

1. not every atom $p \in \mathbb{A}$ needs to be in control of an agent;
2. agent i does not necessarily know what j sees (if $i \neq j$) and does not have complete ignorance either;
3. agent i does not necessarily know what j knows about control (if $i \neq j$) and does not have complete ignorance either.

To cater for this, let $\Upsilon_i = \langle \Omega_i, V_i \rangle$, where $\Omega_i \subseteq \mathbb{A}$ and $V_i \subseteq \mathbb{A}$. The idea is that for every atom in Ω_i , agent i knows who controls it, and for every atom in V_i , agent i knows what its truth value is. Now, a model M is of the form

$$M = \langle N, S, R^\Delta, \simeq \rangle, \text{ where}$$

1. S is a set of states $\langle \mathbb{A}_1, \dots, \mathbb{A}_n, \Upsilon_1, \dots, \Upsilon_n, \theta \rangle$;
 - (a) $\cup_{i \in N} \mathbb{A}_i \subseteq \mathbb{A}$ and $\mathbb{A}_i \cap \mathbb{A}_j \neq \emptyset$
 - (b) $\Upsilon_i = \langle \Omega_i, V_i \rangle$ with $\Omega_i, V_i \subseteq \mathbb{A}$
2. $R^\Delta : N \rightarrow S \times S$ is a binary relation. This relation satisfies the following: for every $\langle \mathbb{A}_1, \dots, \mathbb{A}_n, \Upsilon_1, \dots, \Upsilon_n, \theta \rangle \in S$, and every $\theta' \text{ such that } \theta \equiv_{\mathbb{A} \setminus \mathbb{A}_{i_s}} \theta'$, there is a state $t = \langle \mathbb{A}_1, \dots, \mathbb{A}_n, \Upsilon_1, \dots, \Upsilon_n, \theta' \rangle$;
3. Given two states $s = \langle \mathbb{A}_1, \dots, \mathbb{A}_n, \Upsilon_1, \dots, \Upsilon_n, \theta \rangle$ and $s' = \langle \mathbb{A}'_1, \dots, \mathbb{A}'_n, \Upsilon'_1, \dots, \Upsilon'_n, \theta' \rangle$, define

$$s \simeq s' \text{ iff } \begin{cases} \Upsilon_i = \Upsilon'_i \\ \forall p \in V_i \theta(p) = \theta'(p) \\ \forall p \in \Omega_i \forall j \in N (p \in \mathbb{A}_j \text{ iff } p \in \mathbb{A}'_j) \end{cases}$$

The semantics is very general and allows for a number of specialisations. Examples of such specialisations are:

1. For all states s and every agent i , $\Omega_{i_s} = \mathbb{A}$ (complete knowledge about control)
2. For all states s and t , and every agent i , the components Ω_{i_s} and Ω_{i_t} are the same.
3. For all states s and t , and every agent i , the components V_{i_s} and V_{i_t} are the same.

These properties entail some validities:

1. $\models ctrls(j, p) \leftrightarrow K_i ctrls(j, p)$
2. $\models K_i ctrls(j, p) \leftrightarrow (K_h K_i ctrls(j, p) \wedge \Box_N K_i ctrls(j, p))$
3. $\models sees(i, p) \leftrightarrow (K_h sees(i, p) \wedge \Box_N sees(i, p))$

In fact, all those specialisations apply to ECL-PC(PO). Other natural assumptions would be that for instance $\mathbb{A}_i \subseteq \Omega_i$ (corresponding to $ctrls(i, p) \rightarrow K_i ctrls(i, p)$) and $\mathbb{A}_i \subseteq V_i$ (corresponding to $ctrls(i, p) \rightarrow sees(i, p)$).

We give one simple scenario that can be modelled in this set-up, that of *Voting*. All agents either desire something (p_i) or not. They can reveal their preference through q_i : if $p_i \leftrightarrow q_i$, agent i is truthful, otherwise it lies. Here, $\mathbb{A}_i = \{q_i\}$, $\Omega_i = \{q_j \mid j \in N\}$ and $V_i = \{p_i\} \cup \{q_j \mid j \in N\}$. In other words, we assume agents cannot control what they prefer, although what they can do is choose their vote. We have here

$$\ell(p_i) \rightarrow K_i(\Diamond_i(\ell(p_i) \wedge q_i) \wedge \Diamond_i(\ell(p_i) \wedge \neg q_i))$$

i.e., i knows that it can vote truthfully but it can also lie. We also get $K_i q_j \rightarrow \neg(K_i p_j \vee K_i \neg p_j)$: even if i knows j 's vote, it does not know j 's real preference. Note that the information about what agents see and what they know about controls is still *global*, we have e.g. $K_i K_j ctrls(h, q_h)$.

4. CONCLUSION

As noted before, we added an information component to the logic of propositional control CL-PC ([14]). From a technical perspective, like in [7], our logic ECL-PC(PO), even if we would require that all agents see all propositional variables, is an extension of CL-PC, since as presented in [14], the distribution of propositional variables \mathbb{A} over agents is assumed as *given*. In ECL-PC(PO), it is not given, but it is *fixed*, implying that a specification φ may leave room for different distributions of the atoms, but once it is chosen, there is no way to refer to other distributions, not in terms of what agents can imagine, nor in terms of what they can achieve.

There are close connections between propositional logics of control and other logics that facilitate reasoning about the powers of coalitions, like Coalition Logic [11] and ATL [2]. In fact, CL-PC was partially motivated by the way the model checking system MOCHA for ATL [3] is designed, in which the system is divided in a number of modules (agents, in our terminology), each controlling its own set of Boolean variables. And indeed, there have been several attempts to add an epistemic component to ATL [13, 8, 1]. However, what those extensions all have in common is that the uncertainty of the agents is specified in an abstract way: in the Kripke models for the logics for cooperation and knowledge, the accessibility relations corresponding to knowledge are just given, abstract, equivalence relations. In our logic CL-PC(PO) the knowledge is determined by the variables of which the agent can see the truth value, and in ECL-PC(UO) this accessibility relation is determined by the variable of which the agent can see the ownership. In

this sense, we provide a *computationally grounded semantics* [16] for knowledge, which brings our approach closer to the interpreted systems approach to epistemic logic [5, 6]. Interestingly enough, the key idea of interpreted systems (two states are the same for agent i if the atoms that it sees have the same value) does not only apply to the epistemic dimension in our logics, but also to the control dimension: two states are reachable in terms of i 's control, if the values of the atoms not in i 's control is the same.

Future work should study how to combine our two approaches, as suggested at the end of Section 3, and to weaken some of the underlying assumptions regarding the agents' knowledge. Related to this, we would like to provide a completeness proof for our systems that does not rely on a normal form (and on the assumption that the number of propositional atoms is finite). Doing this, one needs to find a way of juggling with the two types of definitions of 'access' we are dealing with here: on the one hand, the canonical model in modal logic defines this in terms of membership of formulas in the states, whereas the interpreted systems approach would do this in terms of 'similarity' of the states. We hope that work of Lomuscio [9], connecting general *S5* semantics with that of interpreted systems, may give some first steps in this search. Another natural direction to be explored is to emphasize the group aspect of both dimensions: when forming a coalition C to bring about φ , i.e., $\Diamond_C \varphi$ gives rise to interesting questions from cooperative game theory, and epistemic logic provides the tools and results to combine this with interesting notions of *group knowledge*.

5. REFERENCES

- [1] T. Ågotnes. Action and knowledge in alternating-time temporal logic. *Synthese*, 149(2):377–409, 2006.
- [2] R. Alur, T. A. Henzinger, and O. Kupferman. Alternating-time temporal logic. *JACM*, 49(5):672–713, 2002.
- [3] R. Alur *et al.* Mocha: Modularity in model checking. In *CAV*, LNCS 1427, pp. 521–525, 1998.
- [4] P. Blackburn, M. de Rijke, and Y. Venema. *Modal Logic*. Cambridge University Press, 2001.
- [5] R. Fagin and J. Y. Halpern. Reasoning about knowledge and probability. In M. Y. Vardi, (ed.), *TARK*, 1988.
- [6] R. Fagin, J. Y. Halpern, Y. Moses, and M. Y. Vardi. *Reasoning About Knowledge*. MIT Press, 1995.
- [7] J. Gerbrandy. Logics of propositional control. In *AAMAS06*, pp. 193–200, 2006.
- [8] W. Jamroga and W. van der Hoek. Agents that know how to play. *Fundamenta Informaticae*, 63(2-3):185–219, 2004.
- [9] A. Lomuscio. *Knowledge Sharing among Ideal Agents*. PhD thesis, University of Birmingham, UK, 1999.
- [10] R. C. Moore. Reasoning about knowledge and action. In *IJCAI-77*, Cambridge, MA, 1977.
- [11] M. Pauly. *Logic for Social Software*. PhD thesis, University of Amsterdam, 2001. ILLC Dissertation Series 2001-10.
- [12] W. van der Hoek, D. Walther, and M. Wooldridge. Reasoning about the transfer of control. *JAIR*, 37:437–477, 2010.
- [13] W. van der Hoek and M. Wooldridge. Time, knowledge, and cooperation. *Studia Logica*, 75(1):125–157, 2003.
- [14] W. van der Hoek and M. Wooldridge. On the logic of cooperation and propositional control. *AI*, 164:81–119, 2005.
- [15] Y. Venema. A modal logic of substitution and quantification. *Logic Colloquium '92*, CSLI Pub., pp 293–309, 1996.
- [16] M. Wooldridge and A. Lomuscio. A computationally grounded logic of visibility, perception, and knowledge. *IGPL*, 9(2):273–288, 2001.