

$eq(G)$	$\hat{=}$	$subseteq(G) \wedge supseteq(G)$
$subset(G)$	$\hat{=}$	$subseteq(G) \wedge \neg eq(G)$
$supseteq(G)$	$\hat{=}$	$supseteq(G) \wedge \neg eq(G)$
$incl(i)$	$\hat{=}$	$supseteq(\{i\})$
$excl(i)$	$\hat{=}$	$\neg incl(i)$
any	$\hat{=}$	$supseteq(\emptyset)$
$nei(G)$	$\hat{=}$	$\bigvee_{i \in G} incl(i)$
$ei(G)$	$\hat{=}$	$\neg nei(G)$
$gt(n)$	$\hat{=}$	$geq(n+1)$
$lt(n)$	$\hat{=}$	$\neg geq(n)$
$leq(n)$	$\hat{=}$	$lt(n+1)$
$maj(n)$	$\hat{=}$	$geq(\lceil (n+1)/2 \rceil)$
$ceq(n)$	$\hat{=}$	$(geq(n) \wedge leq(n))$

Table 1: Derived coalition predicates.

as $S5_n^{C,D}$ it is exponentially more succinct. We also provide a sound and complete axiomatisation of the logic. In Section 5 we present some detailed case studies, showing how the logic can be used in the specification and analysis of systems, while in section 6 we study and give a complete characterisation of the computational complexity of the model checking problem. We conclude in Section 7.

2. COALITION PREDICATES

We first introduce the language of coalition predicates (from [1]). In what follows we assume a set $Ag = \{1, \dots, n\}$ of agents. Syntactically, the language of coalition predicates is built from three atomic predicates $subseteq$, $supseteq$, geq and we derive a stock of other predicate forms from these¹. Formally, the syntax of coalition predicates is given by the following grammar:

$$P ::= subseteq(G) \mid supseteq(G) \mid geq(n) \mid \neg P \mid P \vee P$$

where $G \subseteq Ag$ is a set of agents and $n \in \mathbb{N}$ is a natural number.

The circumstances under which a coalition $G_0 \subseteq Ag$ satisfies a coalition predicate P are specified by the satisfaction relation “ \models_{cp} ”, defined by the following rules:

$$\begin{aligned} G_0 &\models_{cp} subseteq(G) \text{ iff } G_0 \subseteq G \\ G_0 &\models_{cp} supseteq(G) \text{ iff } G_0 \supseteq G \\ G_0 &\models_{cp} geq(n) \text{ iff } |G_0| \geq n \\ G_0 &\models_{cp} \neg P \text{ iff not } G_0 \models_{cp} P \\ G_0 &\models_{cp} P_1 \vee P_2 \text{ iff } G_0 \models_{cp} P_1 \text{ or } G_0 \models_{cp} P_2 \end{aligned}$$

We assume the conventional definitions of implication (\rightarrow), biconditional (\leftrightarrow), conjunction (\wedge), and exclusive-or (∇) in terms of \neg and \vee . We also find it convenient to make use of the derived predicates defined in Table 1.

3. ELQC

We now introduce our Epistemic Logic with Quantification over Coalitions – ELQC. We begin with some intuition about the logic. Recall that conventional $S5_n^{C,D}$ contains individual epistemic operators K_i (“agent i knows...”), as well as group knowledge operators C_G (“it is common knowledge in G that...”), D_G (“it is

¹We could work with a smaller base of predicates, deriving the remaining predicates from these, but the definitions would not be succinct; see the discussion in [1].

distributed knowledge in G that...”) and E_G (“everybody in G knows that...”). Although we can (and do) define these operators in ELQC, we start from a rather different looking operator base. For each of the modes of group knowledge $X \in \{C, D, E\}$, we introduce operators $\langle P \rangle_X$ and $[P]_X$, where P is a coalition predicate, as defined above. Then $\langle P \rangle_X \varphi$ will mean “there exists a group G such that G satisfies P and G has X -knowledge that φ ”, while $[P]_X \varphi$ means “in every group G satisfying P it is X -knowledge that φ ”. Notice that we introduce both the existential *and* universal operators as primitives, which may at first sight seem unnecessary; we return to this point later.

The following formula expresses the property described in the introduction to this paper:

$$\langle geq(2) \rangle_E \neg \langle gt(3) \rangle_E \varphi.$$

3.1 Syntax

The syntax of ELQC is defined by the grammar of Figure 1, with respect to a set Φ of atomic propositions, and Ag of agents. As usual, we use parentheses to disambiguate formulae, and define the remaining connectives of classical logic as abbreviations: $\perp \hat{=} \neg \top$, $\varphi \rightarrow \psi \hat{=} (\neg \varphi) \vee \psi$, $\varphi \leftrightarrow \psi \hat{=} (\varphi \rightarrow \psi) \wedge (\psi \rightarrow \varphi)$, and $\varphi \nabla \psi \hat{=} (\varphi \vee \psi) \wedge \neg(\varphi \wedge \psi)$. q

3.2 Semantics

Our semantics very closely follow the conventional $S5_n^{C,D}$ semantics for knowledge, and in particular, the models over which we interpret formulae of ELQC are the same as those used to interpret formulae of $S5_n^{C,D}$. Since the language is parameterised by the language of coalition predicates, the semantics are parameterised by the semantics \models_{cp} for coalition predicates, as defined in section 2. Moreover, models are implicitly parameterised by a set $Ag = \{1, \dots, n\}$ of agents, and a set Φ of Boolean variables.

Formally, a *model*, M , (over Ag, Φ) is an $(n+2)$ -tuple [3, p.17]:

$$M = \langle S, \sim_1, \dots, \sim_n, \pi \rangle,$$

where:

- S is a finite, non-empty set of *states*;
- $\sim_i \subseteq S \times S$ is an *epistemic accessibility relation* for each agent $i \in Ag$ — we require that each \sim_i is an equivalence relation; and
- $\pi : S \rightarrow 2^\Phi$ is a Kripke valuation function, which gives the set of primitive propositions satisfied in each state.

If $G \subseteq Ag$, we denote the union of G 's accessibility relations by \sim_G^E , so $\sim_G^E = (\bigcup_{i \in G} \sim_i)$. We use \sim_G^C to denote the transitive closure of \sim_G^E . Finally, \sim_G^D denotes the intersection of G 's accessibility relations. We use these relations to give a semantics to group knowledge modalities (cf. [3, p.66–70]).

A *pointed structure* is a pair M, s , where M is a model and s is a state in M . We interpret formulae of ELQC with respect to pointed structures, via the following rules (note that the rules make use of the semantic satisfaction relation \models_{cp} for coalition predicates, and implicitly assume a set Ag of agents; also, in what follows, we let X denote one of the three modes of knowledge $\{E, D, C\}$).

- $M, s \models \top$
- $M, s \models p$ iff $p \in \pi(s)$ (where $p \in \Phi$)

$\varphi ::=$	\top	/* truth constant */
	p	/* Boolean variables */
	$\neg\varphi$	/* negation */
	$\varphi \vee \psi$	/* disjunction */
	$\langle P \rangle_E \varphi$	/* there exists a P coalition in which every agent knows φ */
	$\langle P \rangle_D \varphi$	/* there exists a P coalition in which it is distributed knowledge that φ */
	$\langle P \rangle_C \varphi$	/* there exists a P coalition in which it is common knowledge that φ */
	$[P]_E \varphi$	/* in every P coalition, every agent knows φ */
	$[P]_D \varphi$	/* in every P coalition, it is distributed knowledge that φ */
	$[P]_C \varphi$	/* in every P coalition, it is common knowledge that φ */

Figure 1: Syntax of ELQC – P is a coalition predicate over Ag , $p \in \Phi$ is a Boolean variable.

- $M, s \models \neg\varphi$ iff
 $M, s \not\models \varphi$
- $M, s \models \varphi \vee \psi$ iff
 $M, s \models \varphi$ or $M, s \models \psi$
- $M, s \models \langle P \rangle_X \varphi$ iff
 $\exists G \subseteq Ag, G \models_{cp} P$ and $\forall s' \in S$ s.t. $s \sim_G^X s'$, we have
 $M, s' \models \varphi$
- $M, s \models [P]_X \varphi$ iff
 $\forall G \subseteq Ag$, if $G \models_{cp} P$ then $\forall s' \in S$ s.t. $s \sim_G^X s'$, we have
 $M, s' \models \varphi$.

As usual, we write $M \models \varphi$ if $M, s \models \varphi$ for all s in M , and $\models \varphi$ if $M \models \varphi$ for all M ; in this latter case, we say that φ is *valid*. Let us now consider why we needed to introduce both $\langle \dots \rangle_X$ and $[\dots]_X$ as primitives, rather than taking the usual modal logic route of defining one as the dual of the other. Taking duals of $\langle P \rangle_X$ and $[P]_X$ gives the following semantics:

- $M, s \models \neg\langle P \rangle_X \neg\varphi$ iff
 $\forall G \subseteq Ag$, if $G \models_{cp} P$ then $\exists s' \in S$ s.t. $s \sim_G^X s'$ and
 $M, s' \models \varphi$.
- $M, s \models \neg[P]_X \neg\varphi$ iff
 $\exists G \subseteq Ag, G \models_{cp} P$ and $\exists s' \in S$ s.t. $s \sim_G^X s'$, we have
 $M, s' \models \varphi$

Thus $\neg\langle P \rangle_X \neg\varphi$ holds if all P -coalitions considers φ to be X -possible, and $\neg[P]_X \neg\varphi$ holds if there is a P -coalition which considers φ to be X -possible. In particular, the dual of $\langle P \rangle_X$ is *not* the same as $[P]_X$, and vice versa. Hence the need to introduce both as atomic.

Note that if P is an Ag -inconsistent predicate, then $\models [P]_X \varphi$ for all modes X and formulae φ , and, similarly $\models \neg\langle P \rangle_X \varphi$. If P is a predicate denoting the empty set of agents, we have for all X that $\models [P]_X \varphi \leftrightarrow \langle P \rangle_X \varphi$, and also $\models [P]_E \varphi \wedge [P]_C \varphi \wedge \neg[P]_D \varphi$, for all formulae φ .

3.3 Some Definitions

The conventional epistemic operator K_i (“agent i knows...”) can be recovered from ELQC via the following definition:

$$K_i \varphi \doteq \langle eq(\{i\}) \rangle_E \varphi$$

Similarly, the standard common knowledge, distributed knowledge, and everyone knows operators (C_G , D_G , and E_G , respectively) can be recovered as follows:

$$\begin{aligned} C_G \varphi &\doteq \langle eq(G) \rangle_C \varphi \\ D_G \varphi &\doteq \langle eq(G) \rangle_D \varphi \\ E_G \varphi &\doteq \langle eq(G) \rangle_E \varphi \end{aligned}$$

Thus, ELQC includes as a fragment the conventional $S5_n^{C,D}$ logic of multi-agent knowledge, with common knowledge, distributed knowledge, and everyone-knows operators [3, 4]. We address the question of its exact relationship to $S5_n^{C,D}$ in the following section. In what follows, we will make frequent use of these abbreviations.

Before proceeding further, note that D_A is monotone in the coalition A , whereas E_A and C_A are anti-monotone. That is, suppose $A \subseteq B$; then, for all formulae φ , we have:

$$\begin{aligned} &\models D_A \varphi \rightarrow D_B \varphi \\ &\models E_B \varphi \rightarrow E_A \varphi \\ &\models C_B \varphi \rightarrow C_A \varphi. \end{aligned}$$

Since D -knowledge is monotonic, it may happen that taking an agent away from a coalition poses a threat to some distributedly known fact in the coalition. The following notion is analogous to one in [1], where a *weak veto player* for φ is an agent that must be present in any coalition that has the ability to achieve φ . We define a *weak veto knower*:

$$WVETO(i, D, \varphi) \doteq \neg \langle excl(i) \rangle_D \varphi$$

This says that no coalition without i has distributed knowledge that φ . This still does not imply that agent i would make φ distributed knowledge, hence we also define a notion of *strong veto knower*, which is an agent that is both necessary to form distributed knowledge of φ , and at least in one case sufficient:

$$SVETO(i, D, \varphi) \doteq WVETO(i, \varphi) \wedge \langle \text{supset}(\{i\}) \rangle_D \varphi$$

Finally, we might consider a general notion of a veto knower as an agent that is both necessary and sufficient for distributed knowledge: if you have this agent in your group, you will have distributed knowledge of φ , and without him, you will not have distributed knowledge of φ .

$$VETO(i, D, \varphi) \doteq WVETO(i, D, \varphi) \wedge \langle \text{supset}(\{i\}) \rangle_D \varphi.$$

Notice that we do not require that $\langle incl(\{i\}) \rangle_D \varphi$, since this would imply $K_i \varphi$. Assuming at least two agents in the system (i and one other), for all X we have $\models \langle incl(\{i\}) \rangle_X \varphi \rightarrow \langle incl(\{i\}) \rangle_X \varphi$. Thus, we have:

$$\begin{aligned} &\models VETO(i, D, \varphi) \rightarrow SVETO(i, D, \varphi) \\ &\models SVETO(i, D, \varphi) \rightarrow WVETO(i, D, \varphi). \end{aligned}$$

EXAMPLE 1. *The notion of a weak veto knower is not very interesting on its own – we have for instance that for every agent i , $WVETO(i, D, \perp)$: everybody is a veto knower for a group to distributedly know a contradiction. For the concept of a strong veto knower, consider a clerk F who is an expert in legislation about registering companies and organisations. For one thing, he knows that if an organisation has been paying tax for three years (t), or is registered with the chamber of commerce (c), it can call itself an official association (a). So, $K_F((t \vee c) \rightarrow a)$. Now, a particular organisation O is very keen on getting status a . Its treasurer A knows that t , and its secretary B that c . Then our clerk is a strong veto player in $\{A, B, F\}$ for distributedly knowing a , i.e., we have $SVETO(F, D, a)$. Everybody is a veto knower in this case, because the minimal coalitions for distributedly knowing a are $\{A, F\}$ and $\{B, F\}$. Would the clerk instead know the stronger rule $(t \wedge c) \rightarrow a$, then every agent in $\{A, B, F\}$ would have been a strong veto player wrt a .*

We have the following properties:

$$\begin{aligned} \models \varphi \rightarrow \psi &\Rightarrow \models WVETO(i, D, \psi) \rightarrow WVETO(i, D, \varphi) \\ \models WVETO(i, D, \varphi) &\leftrightarrow \neg \langle eq(Ag \setminus \{i\}) \rangle_{D\varphi} \\ \models SVETO(i, D, \varphi) &\leftrightarrow (\neg \langle eq(Ag \setminus \{i\}) \rangle_{D\varphi} \wedge \langle eq(Ag) \rangle_{D\varphi}) \end{aligned}$$

From the latter properties we derive an easy characterisation for strong veto knowers:

$M, w \models SVETO(i, D, \varphi)$ iff for all v such that $w \sim_{Ag}^D v$ we have $M, v \models \varphi$, but for some v such that $w \sim_{Ag \setminus \{i\}}^D v$, we have $M, v \models \neg\varphi$. In words: look at all the worlds that are accessible according to all of the agents simultaneously: they must satisfy φ . However, at the same time, there must be a world considered possible by all agents except i , in which $\neg\varphi$ holds.

In fact, we can lift the *VETO* predicates from individuals to coalitions A in several ways, on the positive side saying that a coalition $B \supseteq A$ has distributed knowledge of φ , but on the negative side varying from requiring that no coalition not including A distributedly knows φ , or that no coalition with some intersection with A has this property. Let $\bar{A} \equiv Ag \setminus A$.

$$WVETO_1(A, D, \varphi) \hat{=} \neg \langle subseteq(\bar{A}) \rangle_{D\varphi}$$

That is, no coalition leaving out *all* members of A distributedly knows φ . We have that

$$\models WVETO_1(A, D, \varphi) \leftrightarrow \neg \langle eq(\bar{A}) \rangle_{D\varphi}.$$

Another variant is the following.

$$WVETO_2(A, D, \varphi) \hat{=} \bigwedge_{i \in A} \neg \langle excl(i) \rangle_{D\varphi}$$

This says that every agent in A is necessary for distributedly knowing φ .

A coalition A is *weakly minimal* for $D\varphi$ if no subset of A distributedly knows φ .

$$WMIN_D(A, \varphi) \hat{=} \neg \langle subset(A) \rangle_{D\varphi}$$

A is *minimal* for $D\varphi$ if it in addition in fact has distributed knowledge of φ .

$$MIN_D(A, \varphi) \hat{=} D_A\varphi \wedge WMIN_D(A, \varphi)$$

Now, for common knowledge and everybody's knowledge, one can define similar notions, but rather than veto knowers and minimal coalitions, one has spoilers and maximal coalitions. A spoiler for φ is an agent that, would he enter a group A , would ensure that φ would not be known by everybody in A any longer, or φ would not be common knowledge any more. Similarly, a maximal coalition for $E\varphi$ or $C\varphi$ is a coalition that cannot absorb any new member without giving up the E (C)-knowledge of φ .

These predicates might be useful in the analysis of communication. For example, let $MIN_D(A, \varphi)$, and let i and j be two agents in A . Then a phone call between i and j , if φ is a purely propositional formula (no epistemic operators), would have the effect that $MIN_D(A \setminus \{i\}, \varphi)$ and $MIN_D(A \setminus \{j\}, \varphi)$: the weakly minimal set for $D\varphi$ does not need both i and j any longer!

3.4 Example Validities

Before we provide a sound and complete axiomatisation of ELQC in the next section, let us consider some examples of logical validities.

The following express the monotonicity properties of the three knowledge modes (see also Section 3.3).

$$\begin{aligned} \langle subseteq(G) \rangle_{D\varphi} &\rightarrow [supseteq(G)]_{D\varphi} \\ \langle subseteq(G) \rangle_{E\varphi} &\rightarrow [subseteq(G)]_{E\varphi} \\ \langle subseteq(G) \rangle_{C\varphi} &\rightarrow [subseteq(G)]_{C\varphi} \end{aligned}$$

Next, consider the interaction between the “diamonds” and the “boxes”. Let $X \in \{D, E, C\}$ be a mode of knowledge and let P be an arbitrary coalition predicate. The relationship between the operators and their duals (see Section 3.2) is as follows:

$$\begin{aligned} \langle P \rangle_X\varphi &\rightarrow \neg [P]_X\neg\varphi \\ [P]_X\varphi &\rightarrow \neg \langle P \rangle_X\neg\varphi \end{aligned}$$

Neither of these hold in the other direction.

We have that

$$\langle eq(G) \rangle_X\varphi \leftrightarrow [eq(G)]_X\varphi$$

– the two operators coincide for uniquely satisfiable predicates. In fact, we have the more general property ([1])

$$[P]_X\varphi \leftrightarrow \bigwedge_{G \models_{ep} P} \langle eq(G) \rangle_X\varphi$$

which shows that $[P]_X$ is definable in terms of $\langle P \rangle_X$. The reason that we still include it in the language is succinctness, as discussed in the next section.

Finally, let us look at some interaction properties between the modes of knowledge. For distributed knowledge, we have that:

$$[subseteq(G)]_{D\varphi} \leftrightarrow [\bigvee_{i \in G} eq(\{i\})]_{E\varphi}$$

and for common knowledge we have that:

$$[supseteq(G)]_{C\varphi} \leftrightarrow C_{Ag}\varphi$$

for arbitrary G .

4. EXPRESSIVE POWER

We know from the discussion above that ELQC contains $S5_n^{C,D}$ as a fragment; this naturally raises the question of exactly how the two logics are related. There are two obvious questions to ask. The first is whether ELQC is *strictly more expressive* than $S5_n^{C,D}$, i.e., whether there is some property of models that can be captured via a formula of ELQC that cannot be captured with $S5_n^{C,D}$. The second is whether ELQC is *strictly more succinct* than $S5_n^{C,D}$: that

$P0a$	$\vdash_{cp} \text{supseteq}(\emptyset)$
$P0b$	$\vdash_{cp} \text{subseq}(Ag)$
$P1$	$\vdash_{cp} \text{supseteq}(G) \wedge \text{subseq}(G') \leftrightarrow \text{subseq}(G \cup G')$
$P2$	$\vdash_{cp} \text{supseteq}(G) \rightarrow \neg \text{subseq}(G')$
$P3$	$\vdash_{cp} \text{subseq}(G \cup \{a\}) \wedge \neg \text{subseq}(\{a\}) \rightarrow \text{subseq}(G)$
$P4$	$\vdash_{cp} \text{subseq}(G) \rightarrow \text{subseq}(G')$
GEQ	$\vdash_{cp} \text{geq}(n) \leftrightarrow \bigvee_{G \subseteq Ag, G \geq n} \text{supseteq}(G)$
$Prop$	$\vdash_{cp} \psi$
MP	$\vdash_{cp} \varphi \rightarrow \psi, \vdash_{cp} \varphi \Rightarrow \vdash_{cp} \psi$
δAx	$\vdash_{ELQC} \delta(Ax)$
$\delta \langle \rangle$	$\vdash_{ELQC} \langle P \rangle_X \varphi \leftrightarrow \bigvee_{\{G \mid \vdash_{cp} \text{eq}(G) \rightarrow P\}} \langle \text{eq}(G) \rangle_X \varphi$
$\delta []$	$\vdash_{ELQC} [P]_X \varphi \leftrightarrow \bigwedge_{\{G \mid \vdash_{cp} \text{eq}(G) \rightarrow P\}} \langle \text{eq}(G) \rangle_X \varphi$
δR	$\delta(R)$

Table 2: Axioms and Rules for Epistemic Logic with Quantification over Coalitions. The condition of $P2$ is $G \not\subseteq G'$, for $P4$ it is $G \subseteq G'$, ψ in $Prop$ is a propositional tautology; Ax in δAx is any $S5_n^{C,D}$ -axiom, R in δR is any $S5_n^{C,D}$ -rule, X ranges over $\{C, D, E\}$.

is, whether there is a class of formulae of ELQC that can only be expressed in $S5_n^{C,D}$ with an unreasonable blow-up in the size of formulae. As an added bonus of our consideration of these issues, we will see that we get an axiomatisation of ELQC “for free” as a consequence of considering the relative expressive power of the logics. The methodology used here follows the pattern used in [1].

4.1 Absolute Expressive Power

To begin, consider the following translation τ from ELQC formulae to $S5_n^{C,D}$ formulae. For atoms p and \top , τ is the identity, and it distributes over negation and disjunction, and moreover:

$$\begin{aligned} \tau(\langle P \rangle_X \varphi) &= \bigvee_{\{G \mid \vdash_{cp} P\}} X_G(\tau(\varphi)) \\ \tau([P]_X \varphi) &= \bigwedge_{\{G \mid \vdash_{cp} P\}} X_G(\tau(\varphi)) \end{aligned}$$

We obviously have a translation in the other direction: let us call it δ , with defining clause $\delta(X_G \varphi) = \langle \text{eq}(G) \rangle_X \delta(\varphi)$. Hence, one can think of $\delta(\tau(\varphi))$ as a normal form for φ , where the only coalition predicate in φ is eq . That ELQC and $S5_n^{C,D}$ are equal with respect to absolute expressive power follows from the fact that the two translations preserve truth: the following is readily established.

THEOREM 1. *Let M, s be a pointed structure, φ be an ELQC formula, and ψ be an $S5_n^{C,D}$ formula. Then:*

1. $M, s \models_{ELQC} \varphi$ iff $M, s \models_{S5_n^{C,D}} \tau(\varphi)$
2. $M, s \models_{S5_n^{C,D}} \psi$ iff $M, s \models_{ELQC} \delta(\psi)$.

4.2 Axiomatisation

The translations introduced above provide the key to a complete axiomatisation of ELQC, which can be derived from axiomatisations of coalition predicates and the underlying logic $S5_n^{C,D}$. First, ELQC includes the δ translation of all the $S5_n^{C,D}$ axioms and rules, and axioms that state that the δ -translation is correct: see the lower part of Table 2. On top of that, ELQC is parametrised by an inference relation \vdash_{cp} for coalition predicates. The axioms for this in Table 2 are taken from [1]².

²Note, again, that the basic predicates are not independent, but we include all of them for succinctness.

THEOREM 2.

1. \vdash_{cp} is sound and complete: for any P , $\models_{cp} P \Leftrightarrow \vdash_{cp} P$ [1]
2. For any $S5_n^{C,D}$ formula φ , $\vdash_{S5_n^{C,D}} \varphi \Rightarrow \vdash_{ELQC} \delta(\varphi)$
3. Let φ be any ELQC formula. Then $\vdash_{ELQC} \varphi \leftrightarrow \delta(\tau(\varphi))$ and, in particular, $\vdash_{ELQC} \varphi$ iff $\vdash_{ELQC} \delta(\tau(\varphi))$.

The following is now immediate.

THEOREM 3 (COMPLETENESS AND SOUNDNESS). *Let φ be an arbitrary ELQC-formula. Then $\vdash_{ELQC} \varphi$ iff $\models_{ELQC} \varphi$.*

4.3 Succinctness

So, ELQC and $S5_n^{C,D}$ are equal with respect to absolute expressive power: but, as we now show, they are not equivalent with respect to succinctness. Define the length $\ell(\varphi)$ of both ELQC and $S5_n^{C,D}$ formulae φ , as follows, where X ranges over the three modes of knowledge $\{C, D, E\}$, Y over $\{C, D\}$ and i over Ag :

$$\begin{aligned} \ell(\top) = \ell(p) &= 1 \\ \ell(\varphi_1 \vee \varphi_2) &= \ell(\varphi_1) + \ell(\varphi_2) + 1 \\ \ell(\neg\varphi) &= \ell(\varphi) + 1 \\ \ell(\langle P \rangle_X \varphi) = \ell([P]_X \varphi) &= \text{prsize}(P) + \ell(\varphi) + 1 \\ \ell(Y_G \varphi) &= 1 + \text{coalsize}(G) + \ell(\varphi) \\ \ell(K_i \varphi) &= 2 + \ell(\varphi) \end{aligned}$$

where

$$\begin{aligned} \text{prsize}(\text{subseq}(G)) &= \text{coalsize}(G) + 1 \\ \text{prsize}(\text{supseteq}(G)) &= \text{coalsize}(G) + 1 \\ \text{prsize}(\neg P) &= \text{prsize}(P) + 1 \\ \text{prsize}(P_1 \vee P_2) &= \text{prsize}(P_1) + \text{prsize}(P_2) + 1 \\ \text{coalsize}(G) &= |G| \end{aligned}$$

Let φ and ψ be L_1 and L_2 formulae, respectively, where L_1 and L_2 both range over ELQC and $S5_n^{C,D}$. Then we say that they are equivalent with respect to some class of models if they have the same satisfying pairs M, s , that is, for each M, s in the class of models it is the case that $M, s \models_{L_1} \varphi$ iff $M, s \models_{L_2} \psi$. In the following theorem we show that ELQC is *exponentially more succinct* than $S5_n^{C,D}$. This notion of relative succinctness is taken from [7], who demonstrates that public announcement logic is exponentially more succinct than epistemic logic, and is also used in [1].

THEOREM 4. *There is an infinite sequence of distinct ELQC formulae $\varphi_0, \varphi_1, \dots$ such that, not only is the $S5_n^{C,D}$ formula $\tau(\varphi_i)$ equivalent to φ_i for every $i \geq 0$, but every $S5_n^{C,D}$ formula ψ_i that is equivalent to φ_i has the property $\ell(\psi_i) \geq 2^{|\varphi_i|}$.*

5. TWO CASE STUDIES

In this section, we give two case studies, to illustrate how ELQC can be used in the analysis and specification of systems.

5.1 Voting

Voting mechanisms are of great interest for multi-agent systems, and have been studied extensively in social choice theory. A particular focus has been on aspects related to strategy proofness, often under the assumption that the preferences of the voters are common knowledge. In some cases, however, voters may not want their preferences and votes to become public. Thus, it might be desirable that the voting mechanism is designed in such a way that the information it reveals about the votes of the individuals is limited (under the assumption that the result of the voting is made public). Of course, it is difficult to demand a “zero-knowledge” voting

mechanism where the publication of the result does not reveal anything about the individual votes, because the result will usually tell us something we did not already know. On the other hand, it is not only the case where a mechanism reveals the actual votes of particular individuals that may be problematic. It might, e.g., also be undesirable that two voters could pool their knowledge about their own respective votes together with the result of the voting (and common knowledge about the mechanism) and deduce the vote of a third voter. ELQC allows a fine grained epistemic analysis of such aspects of voting under incomplete information, as the following examples illustrate.

EXAMPLE 2. Assume that a committee consisting of Ann, Bill, Cath and Dave vote for who should be the leader of the committee (it is possible to vote for oneself). The winner is decided by majority voting (majority means at least three votes, if there is no majority there is no winner). After the secret voting, the winner is announced to be Ann. What do the voters know after this announcement?

We can model this situation formally by considering a Kripke structure with states AAAA, AAAB, AAAC, . . . , where AAAA is the state where all agents voted for Ann, AAAB the state where Ann, Bill and Cath voted for Ann and Dave voted for Bill, and so on. We model the situation after the announcement that Ann is the winner, so we only consider states $xyzw$ with at least three As. The accessibility relation is defined in the straightforward way, e.g., $xyzw \sim_A x'y'z'w'$ iff $x = x'$. We can use atomic propositions to express the votes in each state: a_b means that Ann votes for Bill, and is true in all states $Bxyz$ for any x, y and z . Let $Votes = \{x_f \wedge y_g \wedge z_h \wedge w_i \mid x, y, z, w, f, g, h, i \in Ag\}$ be the set of formulae representing complete votes. In any state of this model we have the following:

- $\bigwedge_{vote \in Votes} (vote \rightarrow \neg(gt(1))_E vote)$. At most one agent knows the complete voting.
- $\bigwedge_{vote \in Votes} (vote \rightarrow C_{Ag} \neg(gt(1))_E vote)$. The above is common knowledge among all agents.
- $\bigwedge_{vote \in Votes} (vote \rightarrow \bigwedge_i \neg K_i \neg(geq(1))_E vote)$. Every agent considers it possible that some agent knows the complete voting.
- $\bigwedge_{vote \in Votes} ((vote \wedge \neg(a_a \wedge b_a \wedge c_a \wedge d_a)) \rightarrow (geq(1))_E vote)$. If the voting is not unanimous, then there is some agent who knows the complete voting.
- $\bigwedge_{vote \in Votes} (vote \rightarrow [geq(4)]_D vote)$. The complete voting is distributed knowledge.

The third property above can be generalised, using distributed knowledge, for majority voting for a general number of agents n :

$$\bigwedge_{vote \in Votes} (vote \rightarrow \bigwedge_i \neg K_i \neg[gt(\lfloor \frac{n-1}{2} \rfloor)]_D vote)$$

– every agent considers it possible that any group consisting of at least (approximately) half of the agents have distributed knowledge of the complete voting.

EXAMPLE 3. Consider the same situation as in Example 2, except that the winner is not announced. Let proposition a mean that A wins (gets at least three votes) and proposition una_a mean that A wins unanimously. The following hold (in any state of the model):

- $\neg a \rightarrow (geq(2))_D \neg(geq(3))_E (\neg una_B \wedge \neg una_C \wedge \neg una_D)$. If A does not win, there is a group of at least two agents who distributively know that at most two agents know that neither B nor C nor D wins unanimously.

- $\neg a \rightarrow (geq(2))_E \neg(geq(4))_E (\neg una_B \wedge \neg una_C \wedge \neg una_D)$. If A does not win, at least two agents know that at most three agents know that neither B nor C nor D wins unanimously.

5.2 Gossiping

Consider the following situation.

Four friends each know a secret. They call each other. In each call they exchange all the secrets that they currently know of. Which phone calls should take place in order to spread all the secrets?

Let us call the friends $F = \{1, 2, 3, 4\}$ and the secrets they know $S = \{s_1, s_2, s_3, s_4\}$ (1 knows s_1 , etc.). Let σ be $s_1 \wedge s_2 \wedge s_3 \wedge s_4$. The aim of the communication protocol is that $(eq(F))_E \sigma$: each of the friends knows σ .

In order to find out which telephone calls should be made, we could analyse the individual knowledge of the agents in detail. For example, let first $X_i T$ mean that i exactly knows the elements of $T \subseteq S$:

$$X_i T = \bigwedge_{t \in T} K_i t \wedge \bigwedge_{s \in S \setminus T} (\neg K_i s \wedge \neg K_i \neg s)$$

Consider now a call between agents i and j . If

$$X_k T_k \wedge K_i T_i \wedge K_j T_j$$

holds before the call, where k is an agent different from i and j , then the following holds after the call:

$$X_k T_k \wedge K_i (T_i \cup T_j) \wedge K_j (T_j \cup T_i)$$

An alternative, which ELQC lends itself to, is to reason about distributed knowledge of coalitions. A precondition and postconditions of the scenario written in ELQC terms are

$$Pre = MIN_D(F, \sigma) \quad Post = \bigwedge_{f \in F} MIN_D(f, \sigma)$$

In words: before any phone call is being made, the smallest set that has distributed knowledge of the combination of all the secrets is the set of all friends, whereas the postcondition stipulates that afterwards, every individual has (distributed) knowledge of this combination. So the idea would be that at every call, some A such that $MIN_D(A, \sigma)$ holds should decrease.

We are interested in the minimal sets that have distributed knowledge of σ . For coalitions G_1, \dots, G_k , let $AM_\sigma(G_1, \dots, G_k)$ denote the fact that exactly G_1, \dots, G_k are the minimal coalitions with distributed knowledge of σ :

$$AM_\sigma(G_1, \dots, G_k) = \bigwedge_{1 \leq i \leq k} MIN_D(G_i, \sigma) \wedge \neg(\bigvee_{1 \leq i \leq k} \neg supseteq(G_i))_D \sigma$$

For instance, $AM_\sigma(123, 35) = MIN_D(123, \sigma) \wedge MIN_D(35, \sigma) \wedge \neg(\neg supseteq(123) \vee \neg supseteq(35))_D \sigma$ (here and in the following we use an abbreviated set notation for simplicity).

Starting with the precondition $AM_\sigma(1234)$, we must choose a first phone call. As the situation is symmetric for all agents the choice does not matter, and we chose a call between agents 1 and 2. After this call, it holds that $AM_\sigma(134, 234)$. Now we must consider the next call. We have that:

- A call between 1 and 3 gives $AM_\sigma(14, 34, 234)$
- A call between 3 and 4 gives $AM_\sigma(13, 14, 23, 24)$
- A call between 2 and 3 gives $AM_\sigma(134, 24, 34)$

As the call between 3 and 4 breaks up *both* coalitions 134 and 234, we choose that call. We thus have that $AM_\sigma(13, 14, 23, 24)$, and in this situation we get the following consequences of a call:

- A call between 1 and 3 gives $AM_\sigma(1, 3, 24)$
- A call between 1 and 4 gives $AM_\sigma(1, 4, 23)$
- A call between 2 and 3 gives $AM_\sigma(2, 3, 14)$
- A call between 2 and 4 gives $AM_\sigma(2, 4, 13)$

In this case all calls are equally informative when it comes to the size of minimal coalitions, so we chose the first one. Finally, a call between 2 and 4 is the only informative call, resulting in the post condition $AM_\sigma(1, 2, 3, 4)$. In summary, the following sequence of actions were performed, with associated pre- and post- conditions:

$$\begin{array}{lll} AM_\sigma(1234) & call(1, 2) & AM_\sigma(134, 234) \\ & call(3, 4) & AM_\sigma(13, 14, 23, 24) \\ & call(1, 3) & AM_\sigma(1, 3, 24) \\ & call(2, 4) & AM_\sigma(1, 3, 2, 4) \end{array}$$

Rather than analysing this problem by stating who knows that, we have used ELQC to reason about coalitions: in this case minimal coalitions that have minimal knowledge about the secret. The focus here has been on capturing pre- and post-conditions of the communication actions succinctly in ELQC, rather than on formalising the how models are updated as a result of actions. Combining the epistemic operators of ELQC with update operators of dynamic epistemic logics is an interesting idea for future work.

6. MODEL CHECKING

Model checking is a widely used approach to verifying, automatically and formally, that a given system has a certain property [2, 10]. The model checking problem for our logic can be understood as follows: Given a pointed structure M , s and a formula φ , is it the case that $M, s \models \varphi$? Now, we know that the model checking problem for the conventional ($S5_n$) logic of knowledge can be solved in time polynomial in the size of M , s , φ , and since this logic is a fragment of ELQC, we know that model checking for this fragment of ELQC is not going to be any harder. But, the obvious question arises, is model checking any harder for ELQC *in general*. The answer to this question is “yes – a lot harder”. The following theorem establishes that, in fact, ELQC model checking is as hard as any problem which requires solving a polynomial number of NP-hard problems [9, p.425].

THEOREM 5. *Model checking for ELQC is Δ_2^P -complete.*

PROOF. We prove hardness first. We reduce the *sequentially nested satisfiability* problem (SNSAT), introduced in [6]. An instance of SNSAT is given by a series of equations of the form

$$\begin{array}{ll} z_1 & = \exists X_1. \chi_1(X_1) \\ z_2 & = \exists X_2. \chi_2(X_2, z_1) \\ z_3 & = \exists X_3. \chi_3(X_3, z_1, z_2) \\ & \dots \\ z_k & = \exists X_k. \chi_k(X_k, z_1, \dots, z_{k-1}) \end{array}$$

where $X_1, \dots, X_k, Z = \{z_1, \dots, z_k\}$ are mutually disjoint sets of Boolean variables, and each $\chi_i(Y)$ is a propositional logic formula over the variables Y ; the idea is we first check whether $\chi_1(X_1)$ is satisfiable, and if it is, we assign z_1 the value true, otherwise assign it false; we then check whether χ_2 is satisfiable under the assumption that z_1 takes the value just derived, and so on. Thus the

result of each equation depends on the value of the previous one. The goal is to determine whether z_k is true. Let the input instance be as described above; we will assume w.l.o.g. that each χ_i is of Conjunctive Normal Form (i.e., a conjunction of clauses), so that negations are only applied to Boolean variables. The reduction is as follows. For each of the variable $v \in Z \cup X_1 \cup \dots \cup X_k$, we create two variables v^t and v^f , and two agents a_v^t and a_v^f . We also create an additional variable s_0 . The overall formula to be model checked has the following structure:

$$\langle \alpha \rangle_D \beta$$

where the coalition predicate α and ELQC formula β are defined as follows. Informally, we use the α formula to “guess” a valuation for the z_i variables in the input instance, while β checks that the values “guessed” in this way correctly describe the satisfiability relationship to the χ_i formulae.

First, we define α , as follows:

$$\alpha \hat{=} \bigwedge_{v \in Z} (incl(a_v^t) \nabla incl(a_v^f))$$

(Recall that ∇ is the exclusive-or operator.) Now, any coalition G such that $G \models_{cp} \alpha$ will correspond to a valuation to the variables $Z = \{z_1, \dots, z_k\}$, defined by $z_i = \top$ if $a_i^t \in G$ and $z_i = \perp$ if $a_i^f \in G$; from construction, for every z_i we must have either $a_{z_i}^t \in G$ or $a_{z_i}^f \in G$ but not both.

Next, we define β . First, we define a transformation on χ_i formulae, to obtain coalition predicates, as follows. Suppose the variables of χ_i are v_1, \dots, v_m . Then we let χ_i^* be the coalition predicate:

$$\chi_i^{**} \wedge \bigwedge_{j=1}^m (incl(a_{v_j}^t) \nabla incl(a_{v_j}^f))$$

where χ_i^{**} is the formula obtained from χ_i by systematically substituting $incl(v^f)$ for every negated instance of v , and $incl(v^t)$ for every un-negated occurrence of v . (Recall that we assume φ is in Conjunctive Normal Form.)

Next, we define a sequence Ψ_0, \dots, Ψ_k of formulae, as follows:

$$\Psi_i \hat{=} \begin{cases} \top & \text{if } i = 0 \\ (z_i^t \leftrightarrow C_{Ag}(s_0 \rightarrow \langle \chi_i^* \rangle_D \bigwedge_{0 \leq j \leq i-1} \Psi_j)) \wedge \\ (z_i^f \leftrightarrow C_{Ag}(s_0 \rightarrow \neg \langle \chi_i^* \rangle_D \bigwedge_{0 \leq j \leq i-1} \Psi_j)) & \text{for } i > 0. \end{cases}$$

We define β as:

$$\beta \equiv z_k^t \wedge \left(\bigwedge_{1 \leq i \leq k} \Psi_i \right)$$

We must now define the Kripke structure against which $\langle \alpha \rangle_D \beta$ is to be checked; the structure contains $1 + 2|Z \cup X_1 \cup \dots \cup X_k|$ states, and is illustrated in Figure 2. We have one initial state s_0 in which all propositions are true, and then for every variable $v \in Z \cup X_1 \cup \dots \cup X_k$, we have two states, one corresponding to the truth of v (v^t will be true in this state, but v^f will not), while the other corresponds to v being false (v^f will be true in this state, while v^t will not). For the epistemic accessibility relations, every state obviously has a self loop, the construction is shown in Figure 2.

We now claim that $M, s_0 \models \langle \alpha \rangle_D \beta$ iff $z_k = \top$ in the input instance of SNSAT. The correctness of the reduction is from construction: roughly, the outer part of the formula guesses a coalition, corresponding to a valuation for the Z variables, while the β part of the construction verifies the correctness of the assignment.

We now turn to membership of Δ_2^P . We sketch a polynomial time dynamic programming algorithm that decides the problem, making

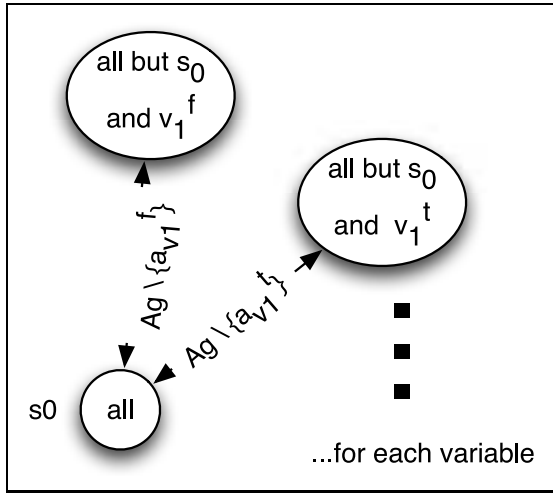


Figure 2: Illustrating the model construction for Theorem 5. (We omit reflexive, transitive, and Euclidean links.)

use of an NP oracle. The algorithm is closely related to conventional model checking algorithms for epistemic logics and modal logics in general [4, p.337]. In what follows, let $sf : ELQC \rightarrow 2^{ELQC}$ be a function that, given a formula φ of ELQC, returns the set of ELQC sub-formulae of φ [4, p.323]. Given input formula φ and pointed structure M, s , the algorithm constructs, via dynamic programming, a function $L : sf(\varphi) \rightarrow 2^S$, such that

$$\forall \psi \in sf(\varphi) : L(\psi) = \{s' \mid M, s' \models \psi\}.$$

The algorithm is as follows:

1. Generate $sf(\varphi)$ and put them in a sequence $\sigma = \varphi_1, \dots, \varphi_l$, in order of length, shortest first through to longest last, with ties broken arbitrarily. Let l be the length of σ , so $\sigma[l] = \varphi$.
2. For $i = 1$ to l do:
 - (a) if $\sigma[i] \in \Phi$ then $L(\sigma[i]) = \pi(\sigma[i])$;
 - (b) if $\sigma[i] = \neg\psi$ then $L(\sigma[i]) = S \setminus L(\psi)$;
 - (c) if $\sigma[i] = \psi \vee \chi$ then $L(\sigma[i]) = L(\psi) \cup L(\chi)$;
 - (d) if $\sigma[i]$ is of the form $\langle P \rangle_X \psi$ or $[P]_X \psi$, then:
 - i. set $L(\sigma[i]) = \emptyset$
 - ii. for each $s' \in S$, invoke the oracle to check whether $M, s' \models \sigma[i]$, and if so, set $L(\sigma[i]) = L(\sigma[i]) \cup \{s'\}$.
3. If $s \in L(\varphi)$ return “yes”, otherwise return “no”.

With respect to correctness of the approach, the only non-obvious step is of course the evaluation of $\langle P \rangle_X$ and $[P]_X$ operators; in particular, we need to show that the evaluation step for these operators can indeed be done with an NP oracle. In fact, this is straightforward. For example, suppose $\sigma[i]$ is of the form $\langle P \rangle_X \psi$. Then we simply guess a $G \subseteq Ag$, and verify that both $G \models_{cp} P$ and $M, s' \models X_G \psi$. The former verification step can obviously be done in polynomial time. To see that the latter step can also, observe that ψ is a strict sub-formula of $\sigma[i]$, and so $L(\psi)$ will be defined by the time we evaluate $\sigma[i]$, and so checking $M, s' \models X_G \psi$ simply involves checking that $\{s'' \mid s' \sim_G^X s''\} \subseteq L(\psi)$. This can clearly be done in polynomial time. \square

7. CONCLUSIONS

By adding a limited form of quantification, we demonstrated how epistemic group logic can become more succinct. In one sense we have been unnecessarily restrictive in our quantification: for instance, in our language, one cannot succinctly express “there is a coalition of size at least 3, which had distributed but not common knowledge that φ ”. It would be possible to allow for Boolean operators between the “type of knowledge” claims, with which the example above could be represented as $\langle geq(3) \rangle_{D \wedge \neg C} \varphi$.

Where our way of quantification is very careful and limited, and did not increase the expressive power of the epistemic logic, an approach at the other end of the spectrum is taken in [5], where the subject of knowledge can be very general terms (for instance, if you receive a mailing announcing you are one of the five lucky winners of some lottery, it is common knowledge among those five that there was such a lottery, even if none of the winners knows any of the others). The logic presented in [5] is not finitely axiomatisable however, and even its monadic fragment is undecidable ([5, p. 177]). It will be interesting to see how rich the coalition predicate logic can become before it increases the expressivity of the epistemic logic, and also what “natural” cut-offs points for such a coalition predicate logic is before the overall logic becomes undecidable.

Another interesting direction for future work is to incorporate the epistemic coalition predicates into dynamic epistemic logics [11]. Such a combination could be useful for analysing scenarios involving both dynamic epistemics and reasoning about the knowledge of coalitions of different sizes.

8. REFERENCES

- [1] T. Ågotnes, W. van der Hoek, and M. Wooldridge. Quantified coalition logic. In *Proc. IJCAI-07*, 2007.
- [2] E. M. Clarke, O. Grumberg, and D. A. Peled. *Model Checking*. MIT Press, 2000.
- [3] R. Fagin, J. Y. Halpern, Y. Moses, and M. Y. Vardi. *Reasoning About Knowledge*. MIT Press, 1995.
- [4] J. Y. Halpern and Y. Moses. A guide to completeness and complexity for modal logics of knowledge and belief. *Artif. Intell.*, 54:319–379, 1992.
- [5] B. Kooi. Dynamic term-modal logic. In *A Meeting of the Minds*, pages 173–186, College Publications, 2007.
- [6] F. Laroussinie, N. Markey, and Ph. Schnoebelen. Model checking CTL+ and FCTL is hard. In *Proc. FoSSaCS '01*, pages 318–331. Springer-Verlag, 2001.
- [7] C. Lutz. Complexity and succinctness of public announcement logic. In *Proc. AAMAS-2006*, 2006.
- [8] J.-J. Ch. Meyer and W. van der Hoek. *Epistemic Logic for AI and Computer Science*. Cambridge UP, 1995.
- [9] C. H. Papadimitriou. *Computational Complexity*. Addison-Wesley, 1994.
- [10] F. Raimondi and A. Lomuscio. Automatic verification of multi-agent systems by model checking via ordered binary decision diagrams. *Jnl of Appl. Logic*, 5:235–251, 2007.
- [11] H. van Ditmarsch, W. van der Hoek, and B. Kooi. *Dynamic Epistemic Logic*, volume 337 of *Synthese Library*. Springer-Verlag, 2007.