THE LOGIC OF JOINT ABILITY IN TWO-PLAYER TACIT GAMES

PETER HAWKE

Stanford University, Department of Philosophy

Abstract. Logics of joint strategic ability have recently received attention, with arguably the most influential being those in a family that includes Coalition Logic (CL) and Alternating-time Temporal Logic (ATL). Notably, both CL and ATL bypass the epistemic issues that underpin Schelling-type *coordination problems*, by apparently relying on the meta-level assumption of (perfectly reliable) communication between cooperating rational agents. Yet such epistemic issues arise naturally in settings relevant to ATL and CL: these logics are standardly interpreted on structures where agents move simultaneously, opening the possibility that an agent cannot foresee the concurrent choices of other agents. In this paper we introduce a variant of CL we call *Two-Player Strategic Coordination Logic* (SCL^2). The key novelty of this framework is an operator for capturing coalitional ability when the cooperating agents cannot share strategic information. We identify significant differences in the expressive power and validities of SCL^2 and CL^2, and present a sound and complete axiomatization for SCL^2. We briefly address conceptual challenges when shifting attention to games with more than two players and stronger notions of rationality.

§1. Introduction: information and joint ability.

1.1. Overview. Ability is a fundamental concept in the study of agency. An agent is predominantly an entity that can act, and talk of power or ability captures what that agent can *accomplish* through action. Thus, ability finds interest in the theory of planning agency; moral responsibility; know how; freedom; state and institutional power; voting and social choice; game theory; and the multi-agent systems paradigm in computer science. Recently, the tools of logic have been used to study *joint strategic ability*, largely in the tradition of the closely related *Coalition Logic* (CL) of Pauly (2001) and *Alternating-time Temporal Logic* (ATL) of Alur, Henzinger, & Kupferman (2002).[1]

In this paper, we introduce a novel variation of CL, so as to explore a specific aspect of the interaction between joint ability and the epistemic status of a coalition. We broadly locate our study with two remarks. First, as is typical in the logical tradition, we focus on a strong sense of 'ability': the ability to *guarantee* an outcome, a notion closely aligned with that of a 'winning strategy'.[2] This notion sometimes travels under the name of agential *power* or *forcing* capacity.[3]

Secondly, we are interested in the semantics and logic of joint ability claims in situations where agents face uncertainty generated by the *simultaneous moves of other agents*.

[1] The STIT family of logics can also convincingly handle ability claims (Horty & Belnap, 1995), (Horty, 2001). For the relationship between STIT and CL/ATL, see Broersen, Herzig, & Troquard (2006). For a further treatment of ability claims, see Pacuit & Simon (2011).

[2] And, plausibly, sincere promise-making: compare Mele (2003).

[3] *Cf.* Part III of van Benthem (2014).

Following a suggestion in the literature, we call such situations games of *almost perfect information* (van der Hoek & Pauly, 2007).[4] In some such games, a winning strategy is unavailable because the agents in the coalition cannot coordinate their individual actions. Our aim is to explore this threat.

In the present section, we discuss the conceptual motivations for the logic that interests us. Things take a technical turn in later sections, as we propose and systematically study a particular logical system.

1.2. Conceptual motivation. CL enriches classical propositional logic with *coalitional modalities*, with the expression $\langle\langle A \rangle\rangle \psi$ intended to mean that the coalition A has the joint ability to secure outcome ψ in the next move. Attempts to establish precise semantics for these expressions have highlighted an ambiguity in the intuitive notion of (joint) ability: while ability *may* be understood as completely determined by physical facts or conventions that bear on the availability of actions, such a reading overlooks the subtle *epistemic* dimensions of having a winning strategy.[5] To illustrate, consider some variations on the well-known *coordinated attack problem*.[6]

EXAMPLE 1.1 (Almost perfect information). *Two armies, respectively commanded by generals* **a** *and* **b***, are positioned separately in the hills overlooking the fortress of their mutual enemy. Neither army is strong enough to defeat the enemy individually, but a simultaneous attack will guarantee victory. It is common knowledge among the generals that* (i) *the other army is stationed nearby,* (ii) *they share the joint goal of defeating the enemy and* (iii) *that this can only be accomplished via coordination. However, the generals have no means for immediate communication and no predetermined agreement to attack at a certain time. Conclusion*: *the generals cannot enforce victory* (*though they could win as a matter of luck*).

EXAMPLE 1.2 (Imperfect information). *Imagine a similar scenario to Example 1.1 with two variations: suppose first that the generals have a perfectly reliable communication channel and can therefore coordinate on a plan of action. Second, the enemy's strength is such that the two armies, even if working together, can only win with the element of surprise. In fact, victory is only assured if they attack the enemy when the changing of the guard occurs. Neither general knows what times guard change, however. Conclusion: the generals cannot enforce a winning outcome.*

EXAMPLE 1.3 (Incomplete information). *Suppose that the generals are able to reliably communicate, and can therefore coordinate. However, general* **a** *has some mistaken beliefs about the possible outcomes of the strategic interaction*: *arrogantly, she believes* (*falsely*) *that general* **b***'s army is incompetent and will hinder an attack. She also believes* (*falsely*) *that her army is strong enough to win single-handedly. Conclusion*: *the generals cannot enforce a winning outcome.*[7]

---

[4] These are familiar to the game theorist. For instance, the prisoner's dilemma is usually understood as a game where the agents move simultaneously without any prior communication or convention.

[5] The interaction between strategic interaction and epistemic state is a long-standing research programme in game theory. For an overview, see Chapters 5 and 7 of Leyton-Brown & Shoham (2008) and Pacuit & Roy (2016).

[6] For more on the coordinated attack problem, see Fagin, Halpern, & Moses (1995).

[7] Is it more accurate to say that the generals do not have the ability to enforce the outcome, or to say that they have the ability, but they don't know *how to use* that ability? To our mind, it makes little difference. The latter amounts to saying that the generals are not able to make use of

Thus, a lack information can disable a coalition in different ways. Using standard terminology from game theory, Example 1.2 is a game of *imperfect information*, where players are uncertain of the current state of the game due to limitations on memory or observational powers.[8] Example 1.3 is a game of *incomplete information*, with players that are mistaken or uncertain about the *structural features* of the game, including possibly the nature of the players.[9] Finally, of special concern, Example 1.1 is a game of *almost perfect information*, following van der Hoek & Pauly (2007). In such games, players move simultaneously, and may thus be unsure "about the actions the other players are simultaneously taking" (van der Hoek & Pauly, 2007, p. 1085). The notion of almost perfect information is therefore closely tied to those of communication and binding agreement. We use Harsanyi's more compact terminology (Harsanyi, 1977, pp. 133–134) and refer to games of almost perfect information as *tacit games* when there is *no* reliable communication within coalitions. This contrasts with the other extreme of *vocal games*, where players communicate at will.

Epistemic considerations have inspired logics that marry strategic ability and information, such as Alternating-time Temporal Epistemic Logic (ATEL) and its variations.[10] However, these logics are generally aimed at discussing ability in the context of imperfect information. Thus, incomplete information and almost perfect information are not addressed directly (though perhaps imperfect information can sometimes *simulate* other modes of uncertainty).[11] In particular, epistemic issues related to *coordination* have received limited attention. This highlights a tension in the CL framework: according to this semantics, $\langle\langle A \rangle\rangle \psi$ is true if there exists an action for each agent in the coalition such that simultaneous performance guarantees ψ. This, however, ignores the possibility that limited information-exchange could hinder predictions about the choices of other agents, impeding coordination. Example 1.1 illustrates this: here, there are several *incompatible* winning joint strategies (namely, to jointly attack at time t), and so each agent must select her contributing action so as to match the choice of her comrade. Harsanyi labels this general problem the *strategy-coordination problem*, judging it to be one of the four central problem facing players in a game (Harsanyi, 1977, chap. 7). Altogether, the notion of ability at work in CL (and ATL and ATEL) indicates a meta-level assumption that agents are informed about the intended moves of all (and only) the members of the coalition to which they belong. Essentially, it is assumed that they are playing a vocal game.

Can satisfactory semantics be provided for a coalitional ability operator for tacit games? In this paper, we take up this challenge. We introduce a new variant of CL which we call, generically, *Strategic Coordination Logic* (SCL). The leading idea is to approach the joint actions available at a game state as a universe of possibilities upon which to interpret a language for epistemic logic, thereby capturing an agent's information (relative to that universe) when selecting her individual action. Ability claims are then interpreted so as to be sensitive to this information.

In particular, we explore a two-player logic we call Two-player Strategic Coordination Logic (SCL^2). SCL^2 has two types of coalitional modality: $\langle\langle A \rangle\rangle$ and $((A))$. A formula

---

their ability to win. However, denying that one has the ability to use an ability to guarantee an outcome seems essentially equivalent to merely denying that one has the ability to guarantee that outcome.

[8] Cf. the notion of *extensive game with imperfect information* in Chapter 5 of Leyton-Brown & Shoham (2008).

[9] Cf. the notion of *Bayesian game* in Chapter 7 of Leyton-Brown & Shoham (2008).

[10] See van der Hoek & Wooldridge (2003) and, for instance, van der Hoek & Jamroga (2004).

[11] For exceptions to this rule - where epistemic issues beyond imperfect information are considered explicitly - see Ghaderi, Lesperance, & Levesque (2007) and Goranko & Hawke (2010).

$\langle\langle A\rangle\rangle\varphi$ informally means that coalition A can jointly achieve φ under the assumption of perfectly reliable communication. Thus the operator $\langle\langle A\rangle\rangle$ is intended to coincide with the strategic ability operator $\langle\langle A\rangle\rangle$ in CL. Using the semantics introduced in this paper, this claim can be made precise. More significantly, a formula of the form $((A))\varphi$ informally means that A can jointly achieve φ even when the members of A *cannot* communicate. The most significant contribution of the present paper is to supply precise semantics for this operator. According to these semantics, one case in which $((\mathbf{a}, \mathbf{b}))\varphi$ holds is if one of agents $\mathbf{a}$ and $\mathbf{b}$ can achieve φ single-handedly. More subtly, however, $((\mathbf{a}, \mathbf{b}))\varphi$ holds when one of $\mathbf{a}$ and $\mathbf{b}$ has an individual action she knows can guarantee φ only because the other agent can be expected *not* to choose certain individual actions—namely, not to choose an individual action he knows would guarantee $\neg\varphi$.

We present two versions of our semantics. §3.3 offers a compact version that does not explicitly utilize epistemic notions. §3.5 offers an alternative that uses an influential epistemic logic, *Public Announcement Logic*, to flesh out the epistemic dimension. In §3.6, we prove their equivalence.

A notable conclusion from our study will be that the $\langle\langle A\rangle\rangle$ and $((A))$ operators show marked *logical* differences. In particular, while $\langle\langle A\rangle\rangle\varphi$ is closed under logical implication, $((A))\varphi$ is not (i.e., the latter operator is not *monotone*) and while distinct agents can always *combine* their abilities under the $\langle\langle A\rangle\rangle$ operator, the same cannot be said for the $((A))$ operator. It can, therefore, be misleading to treat the logic of the $\langle\langle A\rangle\rangle$ operator as paradigmatic for joint ability.

1.3.   Scope.   We deliberately constrain our scope. The interaction between joint ability, coordination, rationality, and information is complex, and we only scratch the surface. We make several assumptions about our agents.

- **Solidarity:** it is common knowledge among agents in a coalition which agents are in the coalition and that each such agent will act as if each joint goal of the coalition is an individual goal for that agent.
- **Minimal rationality:** it is common knowledge among agents in a coalition that every agent in that coalition always selects, if one is available, an action that guarantees her individual goals are met and no agent in the coalition voluntarily performs actions that guarantee that their individual goals are not met.
- **Complete information:** the structure of the game is common knowledge among the agents, including the number of states, agents or actions, and the precise outcomes of joint actions.
- **Perfect information:** at any point in play, the current state of the game is common knowledge among the agents.

We count these as informal *meta-level* assumptions, for we neither explicitly represent these assumptions in our models, nor provide resources in our logical language for expressing them. Rather, they inform the intuitive plausibility of our semantic clauses. This approach reflects the norm in the literature on the logic of ability. For instance, the appeal of CL and ATL hinges on (something like) the above meta-level assumptions, along with, in our view, an additional one: that the agents in a coalition can reliably communicate. Thus we envisage our study as largely agreeing with the meta-level assumptions of CL and ATL, while dropping the assumption of communication.[12]

---

[12]   It is of course of interest to develop a logical system that allows one to both represent the above assumptions (or their failure) in a model and express these assumptions (or their denial) in a

Our assumptions guide the informal interpretation of $((A))\varphi$ and $\langle\langle A\rangle\rangle\varphi$. We intend the former to *both* mean that A can enforce φ even if they cannot communicate *and* that A can enforce φ even if they cannot coordinate. Likewise, we intend the latter to *both* mean that A can enforce φ if they can communicate *and* that A can enforce φ if they can coordinate. Note that, in general, being in a position to coordinate cannot be identified with being in a position to communicate. Consider the prisoner's dilemma: plausibly, this scenario would *not* lead to coordination/cooperation of the agents if they were merely able to communicate. What is further required is that the agents can form *binding commitments* with respect to *common goals/preferences*. However, since our assumption of solidarity ensures commitment within a coalition, our discussion treats communication and coordination as interchangeable conditions.[13]

Two further aspects of our approach bear emphasis: we (i) limit discussion to two-player games and (ii) define the semantics of SCL^2 on the same class of structures used to interpret two-player CL. Aspect (i) is a limitation, hopefully to be remedied in later work. Nevertheless, (i) and (ii) serve crucial roles in the current discussion, for we have the specific goals of (a) motivating a version of SCL on the basis of compelling intuitions concerning simple games, and our intuitions are most clearly stoked by certain two-player scenarios, and (b) we aim to directly compare the expressive power and proof system of this version of SCL with CL, a task most effectively carried out if these systems are interpreted on common structures. Note that a focus on two-player games generates *general* lessons: for instance, differences in *expressivity* and *invalidity*[14] between CL^2 and SCL^2 remain relevant however we decide to extend SCL to deal with more agents. Finally, two-player games are a justifiable starting point when one realizes that extension beyond this simple (though educational) setting is not straightforward: in 5.1, we present an intuitive counter-example (in a three-agent setting) to a *naive* extension of our two-player semantics. We tentatively propose a more sophisticated extension.

Finally we assume, for simplicity, that a communication channel between two agents is always symmetric and that communication takes zero time units. The latter assumption allows us to ignore situations where the agents can communicate *in principle* but do not have time to formulate an agreement.

1.4. Plan.   In §2, we provide, as necessary background, the syntax and semantics of CL and Public Announcement Logic (PAL), where the latter is a type of dynamic epistemic logic that will be crucial in our definitions of the strategic ability operators in SCL^2. In §3, we provide some simple games that test our intuitions concerning joint ability (and coordination) against the notion of ability at work in standard CL. We use this to motivate the syntax and semantics of SCL^2 and then discuss the success of this logic in dealing with the test cases. In §4, we present technical results: we compare the expressive power of SCL^2 and CL; compare the validities of CL and SCL^2; and state and prove a sound and complete axiomatization for SCL^2. In §5, we briefly discuss the prospects for extending SCL to settings with more than two agents, and for making stronger assumptions about the rationality of our agents.

---

precise language. This task must be left to other papers, however. For instance the logic ATEL offers such an approach with respect to the question of perfect information. See van der Hoek & Wooldridge (2003).

[13]   Thanks to an anonymous referee for pressing us to clarify these points.

[14]   Note our meaning: φ is an invalidity just in case φ is not a validity.

1.5. Contrasts and connections to related work. Though coordination problems drive our discussion, we steer away from the diagnosis that has chiefly animated the debate since the pioneering work of Schelling (1960): that real (rational) agents overcome coordination problems through psychological devices such as *focus* (Schelling, 1960) or *framing* (Bacharach, 2006), which transform the situation in question into a more tractable kind of game (from the agents' perspective). To marry formal treatments along this line with a logic of ability seems to us an intriguing topic. However, in this paper, we assume that a coordination problem faced by our agents cannot be dissolved through mere redescription. Second, we note that the notions of *signal* and *intention* relate to the issues of this paper. However, we here background this complexity. A discussion closely related to our own can be found in Ghaderi *et al.*, (2007), where the epistemic issues of coordination are identified, discussed, and a logical theory introduced. There are significant differences between this approach and ours, however, both in emphasis and formal techniques. Notably, the logical theory of Ghaderi *et al.*, (2007) is based on the extremely expressive formalism known as the *situation calculus*, while ours is in the tradition of CL/ATL. In another related work, (van Benthem, 2007) makes similar use of the machinery of dynamic epistemic logic to study strategic interactions between agents.[15] An important difference between the current work and (van Benthem, 2007), however, is that the former is less concerned with sophisticated solution concepts from noncooperative game theory (which rely crucially on information about the preferences of agents), and rather with coalitional *ability*, which we understand as abstracting away from actual agent preference. On a similar note, we register interesting connections between the current work and the study of the knowledge-theoretic aspects of strategic voting, as in Chopra, Pacuit, & Parikh (2004) and van Ditmarsch, Lang, & Saffidine (2012).

§2. Technical background.

2.1. Concurrent game structures. We interpret CL and SCL on *concurrent game structures* (CGSs). A generalization of labelled transition systems, these structures provide sufficient flexibility to represent various situations of interest to the game theorist. In particular, each state of a CGS can be associated with a game in strategic form.[16]

DEFINITION 2.1 (Concurrent game structure). *Relative to a set of atomic formulae* Π, *a* concurrent game structure *is a tuple* $\langle k, Q, d, \delta, \pi \rangle$, *where*:

- k *is a number of agents. We represent the set of agents with* $\mathcal{A} = \{1, 2, \ldots, k\}$
- Q *is a nonempty set of states*
- d *is a function* $d : Q \times \mathcal{A} \to \mathbb{N}$, *where* $d(q, a)$ *represents the number of actions available to agent* a *at state* q. *We thus represent the actions available to* a *at* q *with* $\mathrm{D}_a(q) = \{1, 2, \ldots, d(q, a)\}$. *A joint action at state* q (*denoted by* σ_q *or sometimes just* σ, *when the context is clear*) *is a tuple* $\langle j_1, j_2, \ldots, j_k \rangle$, *where* $j_i \leq d_i(q)$ *for every* $i \leq k$. *Given* $\sigma = \langle j_1, j_2, \ldots, j_k \rangle$, *we sometimes write* σ^i *to indicate* j_i. *We use* $\mathrm{D}(q)$ *to denote the set of joint actions at* q.
- δ, *the transition function, is a function that maps a state* q *and a joint action at* q *to a state in* Q.
- $\pi : \Pi \to \mathcal{P}(Q)$ *is a valuation function that assigns a set of states in* Q *to each atomic formula.*

---

[15] For more on this line, see Chapter 12 of van Benthem (2014).

[16] Again, see Leyton-Brown & Shoham (2008) or Pacuit & Roy (2016) for an overview of standard game theoretic terminology.

We use '2-CGS' to indicate a CGS with $\mathcal{A} = \{1, 2\}$. We use the name **a** to refer to agent 1 and **b** to refer to agent 2. We use 'a' as a meta-variable for agents.

2.2. Syntax and semantics of CL. Two-player coalition logic (CL^2) is the benchmark against which we compare SCL^2. We now outline syntax and semantics for CL. This semantics differs from the original semantics of (Pauly, 2001) based on coalition effectivity models. Instead, we use semantics based on concurrent game structures, essentially that of (Alur *et al*., 2002). Technically, nothing is lost in this move: the semantics based on coalition models is equivalent to that based on concurrent game structures, as per Goranko & Jamroga (2004) and Goranko, Jamroga, & Turrini (2013).[17]

DEFINITION 2.2 (Syntax of CL). *We denote the language of CL, relative to a finite set of agents $\mathcal{A}$ and countable set of atomic formulae Π, by $\mathcal{L}_{CL}(\mathcal{A}, \Pi)$ (or just $\mathcal{L}_{CL}$ when the context is clear). The recursive definition of the formulae of this language is as follows*:

$$\varphi ::= P \mid \neg\varphi \mid \varphi \wedge \varphi \mid \varphi \vee \varphi \mid \langle\langle A \rangle\rangle\varphi,$$

where A denotes a subset of $\mathcal{A}$ and $P \in \Pi$.

The intended informal interpretation of the $\langle\langle A \rangle\rangle\psi$ is "coalition A has the joint ability to guarantee the outcome ψ in the next move". We argue in due course that this can be appropriately qualified with "on the assumption of perfectly reliable communication within the coalition".

$\mathcal{L}_{CL}(\mathcal{A}, \Pi)$ is interpreted on the class of CGSs in which the set of agents in the structure match that associated with the language, and the valuation function π assigns sets of states to the formulae in Π. We require a number of auxiliary notions.

DEFINITION 2.3 (Successor). *Given a CGS $\mathcal{S}$, we say that state q' is a* successor *of state q if there is a joint action σ such that $\delta(q, \sigma) = q'$. We denote the set of successors to state q by $succ(q)$.*

DEFINITION 2.4 (*A*-move, contribution, reach). *Given a CGS $\mathcal{S}$ and a coalition of agents $A \subseteq \mathcal{A}$, an A-move on $\mathcal{S}$ at q is a tuple of actions σ^A with a position for an action at q for each agent in A. We say that agent a* contributes α_a to σ^A. *Every A-move that the coalition A can select at state q induces a set of q-successors on Q (the* outcomes *of that move) which we denote by $succ(q, \sigma^A)$. If $q' \in succ(q, \sigma^A)$, we say that q' can be* reached *from q using σ^A. We write $succ(q, \alpha_1, \ldots, \alpha_n)$ to mean $succ(q, \langle\alpha_1, \ldots, \alpha_n\rangle)$.*

DEFINITION 2.5 (Semantics of CL). *Let $\mathcal{S}$ denote a CGS and q any state in $\mathcal{S}$. Then we interpret $\mathcal{L}_{CL}$ formulae in the following way*:

- $\mathcal{S}, q \vDash_{CL} P$, where $P \in \Pi$, *iff $P \in \pi(q)$.*
- $\mathcal{S}, q \vDash_{CL} \neg\varphi$ *iff $\mathcal{S}, q \nvDash \varphi$.*

---

[17] More precisely, the class of *truly playable* effectivity functions is identical to the class of α-effectivity functions (i.e., the effectivity functions generated by the class of strategic games). According to a refinement of Pauly's semantics offered by Goranko *et al*., (2013), the language of CL is interpreted on the class of truly playable coalition models, which are comprised of a set of states, a truly playable effectivity function at each world and a valuation. On this approach, the ability modality is interpreted at a state in terms of the outcomes for which a coalition is effective at that state. Meanwhile, a CGS-based semantics for CL effectively associates a strategic game with each state, and interprets the ability modality in a way that tracks the α-effectivity function generated by the strategic game at that state.

- $\mathcal{S}, q \vDash_{CL} \varphi \wedge \psi$ iff $\mathcal{S}, q \vDash \varphi$ and $\mathcal{S}, q \vDash \psi$.
- $\mathcal{S}, q \vDash_{CL} \varphi \vee \psi$ iff $\mathcal{S}, q \vDash \varphi$ or $\mathcal{S}, q \vDash \psi$.
- $\mathcal{S}, q \vDash_{CL} \langle\langle A \rangle\rangle\varphi$ iff there exists an A-move σ^A such that if $q' \in succ(q, \sigma^A)$ then $\mathcal{S}, q' \vDash \varphi$.

CL is just the next-time fragment of ATL. We exploit this connection to import useful results or definitions originally framed in the context of ATL.

2.3. Syntax and semantics of PAL.   The logic PAL will find utility in our discussion of SCL^2. Roughly, the idea behind Public Announcement Logic (PAL) is to provide logical tools for reasoning about how *public announcements* influence the epistemics of a group of agents: the public announcement of (true) proposition φ *updates* the epistemic situation for the agents, by eliminating from their consideration all possible states in which φ is not true. See van Ditmarsch, van der Hoek, & Kooi (2008) for a full discussion.

"Announcement" need not be interpreted as something literally emanating from a loudspeaker. Generally, an announcement may be understood as an event in which certain information becomes publicly available, by whatever source. For the applications in this paper, a PAL announcement is best interpreted as a *common inference* that is made by a group of agents on the strength of common knowledge of the structure of the game, nature of the players, and so forth, *not* as an act of communication.

DEFINITION 2.6 (Syntax of PAL). *Let a finite set of agents $\mathcal{A}$ and a countable set of atomic formulae Π be given. We inductively define the language of PAL (*with distributed knowledge operators*), denoted by $\mathcal{L}_{PAL}(\mathcal{A}, \Pi)$, with the following BNF:*

$$\varphi ::= P \mid \neg\varphi \mid \varphi \wedge \varphi \mid \varphi \vee \varphi \mid K_a\varphi \mid D_A\varphi \mid \langle\varphi\rangle\,\varphi,$$

where a denotes an agent in $\mathcal{A}$, A denotes a subset of $\mathcal{A}$ and $P \in \Pi$. Given operator K_a, we define the dual operator $\hat{K}_a$ by $\hat{K}_a\varphi := \neg K_a\neg\varphi$.

Informally, the intended interpretation of an expression of the form $K_a\varphi$ is "agent a knows that φ". The intended interpretation of $D_A\varphi$ is "it is distributed knowledge amongst the members of A that φ holds". Finally, the intended interpretation of $\langle\varphi\rangle\,\psi$ is "after announcement of φ, it is true that ψ".

Formulae of this language are interpreted on the following structures.

DEFINITION 2.7 (Multi-agent epistemic structure). *Given a finite set of agents $\mathcal{A} = \{1, \ldots, k\}$ and a set of atomic formulae Π, a* multi-agent epistemic structure *is a tuple $\mathcal{M} = \langle S, \{\sim_a\}_{a \in \mathcal{A}}, V\rangle$, where*

- *S is a nonempty set of possible states.*
- *For each agent $a \in \mathcal{A}$ there is an equivalence relation $\sim_a \subseteq S \times S$ (that is, each relation is reflexive, transitive and symmetric).*
- *$V : \Pi \to \mathcal{P}(S)$ is a valuation function.*

DEFINITION 2.8 (Semantics of PAL). *For a given a multi-agent epistemic structure $\mathcal{M}$ and state $s \in S$, let $\sim_A$ denote the intersection of the equivalence relations associated with the members of A i.e., the relation $\bigcap_{a \in A} \sim_a$. Then:*

- *$\mathcal{M}, s \vDash_{PAL} P$, for $P \in \Pi$, iff $s \in V(P)$,*
- *$\mathcal{M}, s \vDash_{PAL} \neg\varphi$ iff $\mathcal{M}, s \nvDash p$,*
- *$\mathcal{M}, s \vDash_{PAL} \varphi \wedge \psi$ iff $\mathcal{M}, s \vDash \varphi$ and $\mathcal{M}, s \vDash \psi$,*

- $\mathcal{M}, s \vDash_{PAL} \varphi \vee \psi$ iff $\mathcal{M}, s \vDash \varphi$ or $\mathcal{M}, s \vDash \psi$,
- $\mathcal{M}, s \vDash_{PAL} K_a \varphi$ iff for all $t \in S : s \sim_a t$ implies $\mathcal{M}, t \vDash \varphi$,
- $\mathcal{M}, s \vDash_{PAL} D_A \varphi$ iff for all $t \in S : s \sim_A t$ implies $\mathcal{M}, t \vDash \varphi$,
- $\mathcal{M}, s \vDash_{PAL} \langle \varphi \rangle \psi$ iff $\mathcal{M}, s \vDash \varphi$ and $\mathcal{M}|_\varphi, s \vDash \psi$,

where $\mathcal{M}|_\varphi$, the update of $\mathcal{M}$ with respect to φ, is the multi-agent epistemic structure $\langle S', \{\sim'_a\}_{a \in \mathcal{A}}, V' \rangle$ with

- $S' = \varphi^{\mathcal{M}}$, where $\varphi^{\mathcal{M}} = \{s \in S \mid \mathcal{M}, s \vDash \varphi\}$,
- $\sim'_a = \sim_a \cap (\varphi^{\mathcal{M}} \times \varphi^{\mathcal{M}})$,
- $V(P)' = V(P) \cap \varphi^{\mathcal{M}}$.

§3. Introducing two-player strategic coordination logic. In this section, we develop a semantics for SCL^2. We formulate two versions. After identifying various intuitive examples as test data, we offer a compact definition of the satisfaction relation $\vDash_{SCL}$. Then, we define a second satisfaction relation $\vDash^e_{SCL}$ that is characterized in explicitly epistemic terms. This second formulation is more involved but is also, in our view, conceptually richer and more flexible (to appreciate this flexibility, consider the natural ways of extending it in §5). In 3.6, we prove that the two semantics are equivalent. Finally, we confirm that our semantics accounts for the test data.

3.1. Motivating examples. What impact might almost perfect information have on joint ability? We gather data for the semantics of a new ability operator $((A))$ from five (counter-)examples. Each describes a CGS involving two agents **a** and **b**. We are interested in the coalitional ability of these agents at the start state in each structure (the node with a double border in the accompanying figures). Each agent chooses between action 0 and 1. For each example, we offer an intuitive judgement about the joint ability, first under the assumption of reliable communication between the agents, then under the assumption of a lack of reliable communication.

EXAMPLE 3.1. *We formalize Example 1.1 as a pure coordination game $\mathcal{S}_1$, depicted in Figure 1. Suppose that **a** and **b** are in a coalition with goal ψ. They achieve this goal only if both choose action 0 or both choose 1.*

Can our coalition enforce ψ? A common sense answer points out an ambiguity in our presentation: can the agents share information? Or not? The intuitive assessment is then as follows: if the agents can share information, then they are able to guarantee ψ. If the agents are not positioned to share information, then clearly they cannot guarantee ψ.

What assessment follows from the semantics for CL? Denote the start state by q. Then: $\mathcal{S}, q \vDash_{CL} \langle\langle \mathbf{a}, \mathbf{b} \rangle\rangle \psi$. Thus, CL gives a satisfactory assessment only if we assume that the agents can communicate (and form binding agreements).
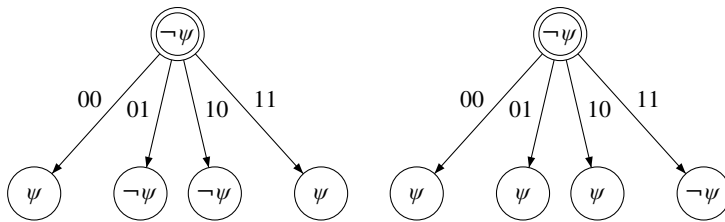

Fig. 1. On the left, the concurrent game structure $\mathcal{S}_1$, as in Example 3.1; on the right, $\mathcal{S}_2$ from Example 3.2.

EXAMPLE 3.2. *Consider the game scenario S_2 depicted in Figure 1. Here, the agents only fail to achieve ψ if they both select action 1.*

Intuitively, our (rational) agents can coordinate in this game, whether or not they can communicate. For if the agents cannot communicate (but share solidarity), both will select action 0, since, for either agent, if that agent plays action 0 then she guarantees that the coalition's goal ψ is accomplished, whatever the other agent does. (Recall that solidarity dictates that a rational agent will always select an individual action that they know guarantees the success of the coalition, if one exists.)

EXAMPLE 3.3. *Consider game S_3, on the left in Figure 2. Here, there is only one way in which the agents can achieve ψ: both select 0.*

Intuitively, can the agents enforce ψ? The answer, we propose, is "yes", no matter whether the agents can communicate or not (again presuming solidarity and rationality). In this case, the semantics of CL provide a satisfactory assessment. However, it is instructive to note that the exact reasoning behind this intuitive answer depends on whether or not communication is assumed. If communication is assumed, then the agents will settle on the joint action $\langle 0, 0 \rangle$. If the agents cannot communicate, then it seems they can mutually rely on each other to not choose an action (namely, action 1) that would sabotage their success.

EXAMPLE 3.4. *Consider the game S_4, on the right in Figure 2. Here we have two formulae of interest, ψ_1 and ψ_2, effecting a blend of the considerations from previous examples. Assume that **a** and **b** cannot communicate. Can they enforce ψ_1? The answer, intuitively, is 'yes', for similar considerations to Example 3.2. Similarly, they can achieve ψ_2. However, can they enforce both ψ_1 and ψ_2 simultaneously (i.e., $\psi_1 \wedge \psi_2$)? Without communication it is intuitively clear that they cannot. To see this, label the states in the figure with $\psi_1 \wedge \psi_2$ and $\neg(\psi_1 \wedge \psi_2)$ where appropriate. It is striking that we thereby recreate the situation in Example 3.1 where the agents cannot coordinate. Thus, just because disjoint sub-coalitions in a coalition can enforce certain goals, it does not follow, under the assumption of no communication, that the coalition as a whole can enforce those sub-goals simultaneously. This contrasts with the ability operator of CL.*

EXAMPLE 3.5. *Consider the game S_5 in Figure 3. We note, drawing on earlier examples, that the coalition can enforce ψ_1 but cannot enforce ψ_2 if we assume a lack of reliable communication.*

Now consider the goal $\psi_1 \vee \psi_2$ (suppose the coalition is indifferent between ψ_1 and ψ_2). According to our earlier assessments, the coalition cannot enforce $\psi_1 \vee \psi_2$ without communication (cf. Example 3.1), despite $\psi_1 \rightarrow (\psi_1 \vee \psi_2)$ being a propositional validity. In a similar vein, if communication is not possible, the coalition can enforce $\psi_1 \wedge \psi_2$ (cf. Example 3.3) but cannot enforce ψ_2 (cf. Example 3.1). We conclude: though a coalition can guarantee θ and θ logically implies φ, it does not in general follow that the coalition
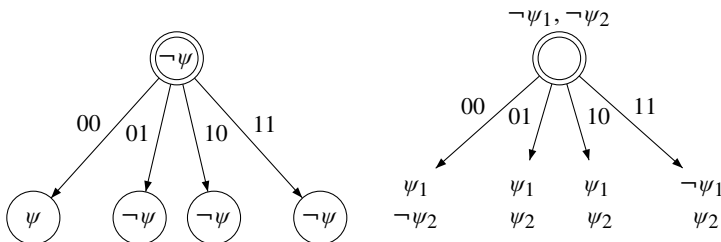

Fig. 2. On the left, the concurrent game structure S_3, as in Example 3.3; on the right, S_4, as in Example 3.4.
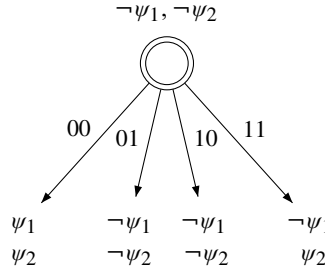
Fig. 3. $\mathcal{S}_5$ as in Example 3.5.

can enforce φ - at least not without communication. The ability-with-no-communication operator is not monotonic.

3.2. Syntax.

DEFINITION 3.6 (Syntax of SCL). *Given a set of agents $\mathcal{A}$ and a set of propositional atoms Π, the language $\mathcal{L}_{SCL}(\mathcal{A}, \Pi)$ (or $\mathcal{L}_{SCL}$ where the context is clear) is given recursively by:*

$$\varphi ::= P \mid \neg\varphi \mid \varphi \wedge \varphi \mid \varphi \vee \varphi \mid \langle\langle A \rangle\rangle\varphi \mid ((A))\varphi,$$

where $P \in \Pi$ and A denotes a subset of $\mathcal{A}$. A propositional formula *is a boolean combination of atoms. A* coalitional atom *is a formula of the form $\langle\langle A \rangle\rangle\varphi$ or $((A))\varphi$.*

$\varphi \rightarrow \psi$ is treated as an abbreviation of $\neg\varphi \vee \psi$. The intended interpretation of $\langle\langle A \rangle\rangle\varphi$ is "*A* can guarantee φ after the next move, on the assumption that all members of *A* can reliably communicate with one another". Again, this operator is intended to be equivalent to the corresponding coalitional operator of CL. The intended interpretation of $((A))\varphi$ is "*A* can guarantee φ after the next move, even if the members of *A* cannot communicate with one another".

3.3. Semantics, nonepistemic version.   From here, we work with

$$\mathcal{L}_{SCL}(\{1, 2\}, \Pi)$$

and use **a** and **b** to respectively name the two agents. Thus, our modalities are: $\langle\langle\rangle\rangle$, $(())$, $\langle\langle\mathbf{a}\rangle\rangle$, $\langle\langle\mathbf{b}\rangle\rangle$, $((\mathbf{a}))$, $((\mathbf{b}))$, $\langle\langle\mathbf{a}, \mathbf{b}\rangle\rangle = \langle\langle\mathcal{A}\rangle\rangle$, and $((\mathbf{a}, \mathbf{b})) = ((\mathcal{A}))$.

DEFINITION 3.7 (SCL^2 semantics V.1). *Let $\mathcal{S}$ denote a 2-CGS and q any state in $\mathcal{S}$. We define the relation $\vDash_{SCL}$ as follows. The clauses for $\mathcal{L}_{SCL^2}$ formulae are exactly as in the semantics of CL, except:[18]*

> $\mathcal{S}, q \vDash_{SCL} ((A))\varphi$ iff $\mathcal{S}, q \vDash_{SCL} \langle\langle A \rangle\rangle\varphi$ and there is an agent $a \in A$ and action $\alpha \in \mathrm{D}_a(q)$ such that for every φ-solidary joint action σ we have $\mathcal{S}, \delta(q, \sigma^*) \vDash \varphi$,

where:

- *a's contribution to σ^* is action α, and otherwise the components of σ^* match σ;*
- *an action $\alpha \in \mathrm{D}_{\mathbf{a}}(q)$ is a φ-solidary action for **a** at q just in case there is an action $\beta \in \mathrm{D}_{\mathbf{b}}(q)$ such that $\mathcal{S}, \delta(q, \alpha, \beta) \vDash_{SCL} \varphi$ (φ-solidary actions for **b** at q are similarly defined, with **a** and **b** swapped);*
- *σ is a φ-solidary joint action just in case every agent in A contributes a φ-solidary action to σ.*

---

[18] Thanks to an anonymous referee for suggestions that improved the readability of this definition.

The intuitive rationale for the above clause reflects the meta-level assumptions outlined in Section 1.3. Consider goal φ and coalition A. Assume that there is common knowledge of solidarity between the members of A (and so each treats the goal of the coalition as if it were an individual goal); that it is commonly known that the members of A are minimally rational and so will not choose an action that sabotages their individual goal; and that there is no communication between the members of A. In this case, when is it true that A can enforce φ? Answer: if there is an agent a in A that can enforce φ so long as no agent in A chooses an action that enforces $\neg\varphi$ (i.e., no agent in A irrationally sabotages the pursuit of φ).

3.4. Generated epistemic structures.   To unpack the epistemic dimension of our semantics, we introduce the auxiliary notion of a *generated epistemic structure*. Notice, given some CGS $\mathcal{S}$ and state q in $\mathcal{S}$, that the set $D(q)$ of joint actions available at q forms a space of *possibilities* (an *action model* in the terminology of (van Ditmarsch *et al.*, 2008)). The agents jointly choose which possibility to actualize. In the context of the issues under discussion, we desire a way to represent the information a given agent has when making this choice. To this end, we follow (van Benthem, 2007) in noticing that the space of possible joint actions is naturally endowed with agent-relative indistinguishability relations, thereby generating a multi-agent epistemic structure. Associate with each possible joint action the propositions which result from executing that action. Then, on the supposition that each agent selects his individual action but not those of other agents, we relate, for each agent, two joint possible actions just in case that agent performs the same individual action in those two joint actions. An agent knows (that is, can accurately predict) that an individual action of theirs will bring about a certain outcome just in case that outcome is brought about by every joint action in which the agent chooses that individual action. This provides for individual knowledge of outcomes under the supposition of no communication. Then, the individual knowledge of the agents, in the wake of communication, coincides with the *distributed* knowledge relative to our indistinguishability relations.

A precise definition of generated epistemic structure must be offered simultaneously as that of $\vDash^e_{SCL}$, as the definitions are mutually recursive. For ease of exposition, we present the former here and the latter in the next section. Notice then that the relation $\vDash^e_{SCL}$ appears in the next definition.

DEFINITION 3.8 (Generated epistemic structure). *Given a CGS*
$\mathcal{S} = \langle k, Q, d, \delta, \Pi, \pi \rangle$ *and a state q in $\mathcal{S}$, the* epistemic structure generated by $\langle \mathcal{S}, q \rangle$
$\mathcal{M}_{\mathcal{S}}(q)$ *(or $\mathcal{M}(q)$ or $\mathcal{M}$ where the context is clear) is the epistemic structure*

$$\langle S, \{\sim_a\}_{a \in \mathcal{A}}, V \rangle$$

relative to language $\mathcal{L}_{PAL}(\mathcal{A}, \Pi^+)$, where

- $S = D(q)$
- *For each $a \in \mathcal{A}$, the relation $\sim_a$ is defined by: $\sigma_1 \sim_a \sigma_2$ iff $\sigma_1^a = \sigma_2^a$, for every $\sigma_1, \sigma_2 \in D(q)$*
- $\Pi^+ = \Pi \cup \{\varphi \in \mathcal{L}_{SCL} \mid \varphi$ *is a coalitional formula*$\}$
- *For $\varphi \in \Pi^+$,*

$$V(\varphi) = \left\{ \sigma \in D(q) \mid \mathcal{S}, \delta(q, \sigma) \vDash^e_{SCL} \varphi \right\}$$

Notice that if $\varphi \in \Pi$, then this amounts to:

$$V(\varphi) = \{ \sigma \in D(q) \mid \delta(q, \sigma) \in \pi(\varphi) \}.$$

As an example, Figure 4 gives the epistemic structure generated by the coordination game at q in Example 3.1 (note that the ψ in this depiction is not necessarily a member of Π^+, but perhaps a boolean combination of these atoms).

Given $\mathcal{M}(q)$, an epistemic structure generated by CGS $\mathcal{S}$ at point q, the formulae of $\mathcal{L}_{PAL}$ can be interpreted using the standard semantics. We intend the informal interpretation of such formulae to be slightly nonstandard. If φ is a propositional atom, we read $\mathcal{M}(q), \sigma \vDash \varphi$ as "if σ is jointly chosen, then the joint choice of action to be played results in φ". The statement $\mathcal{M}(q), \sigma \vDash K_a\varphi$ is read as "if σ is jointly chosen to be played, then a has enough information to accurately predict (i.e., has "knowledge") that (the individual action they contribute to) the joint choice of action to be played will result in φ". With this in mind, the statement $\mathcal{M}(q), \sigma \vDash D_A\varphi$ is read as one would expect. The statement $\mathcal{M}(q), \sigma \vDash \langle\psi\rangle K_a\varphi$ is read as "if σ is jointly chosen, then after the announcement that ψ (i.e., after the elimination of all joint actions at which $\neg\psi$ holds), agent a knows that (the individual action they contribute to) the joint choice of action to be played will result in φ".

3.5. Semantics, epistemic version.

DEFINITION 3.9 (SCL^2 semantics V.2). *Let $\mathcal{S}$ denote a 2-CGS and q any state in $\mathcal{S}$. We define $\vDash^e_{SCL}$ as follows. For propositional atoms and boolean compositions, the interpretation is as in the semantics of CL. For coalitional atoms*:

- $\mathcal{S}, q \vDash^e_{SCL} \langle\langle A\rangle\rangle\varphi$ *iff there exists a joint move* $\sigma \in D(q)$ *such that*

$$\mathcal{M}(q), \sigma \vDash_{PAL} D_A\varphi$$

- $\mathcal{S}, q \vDash^e_{SCL} ((A))\varphi$ *iff there exists a joint move* $\sigma \in D(q)$ *such that*

$$\mathcal{M}(q), \sigma \vDash_{PAL} \langle \bigwedge_{a \in A} \hat{K}_a\varphi \rangle \bigvee_{a \in A} K_a\varphi.$$

The following is easy to check.

PROPOSITION 3.10.

(a) $\mathcal{S}, q \vDash^e_{SCL} \langle\langle a\rangle\rangle\varphi$ *iff there is a joint move* $\sigma \in D(q)$ *s.t.* $\mathcal{M}(q), \sigma \vDash_{PAL} K_a\varphi$
(b) $\mathcal{S}, q \vDash^e_{SCL} \langle\langle a\rangle\rangle\varphi$ *iff* $\mathcal{S}, q \vDash^e_{SCL} ((a))\varphi$
(c) $\mathcal{S}, q \vDash^e_{SCL} \langle\langle \mathcal{A}\rangle\rangle\varphi$ *iff there is a joint move* $\sigma \in D(q)$ *s.t.* $\mathcal{M}(q), \sigma \vDash_{PAL} \varphi$.

Item (c) is a consequence of the fact that $D_\mathcal{A}\varphi$ is equivalent to φ on generated epistemic models, since these models have a *hyper-cube* structure.

The rationale for the $\langle\langle A\rangle\rangle$ clause is hopefully transparent: a coalition can (coordinate to) guarantee an outcome ψ, under the assumption of communication, just in case there is a joint action such that it is distributed knowledge among the coalition that their combined contribution guarantees ψ, no matter what the agents outside the coalition do.
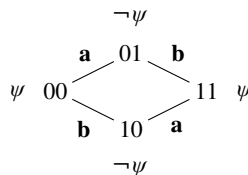

Fig. 4. The generated epistemic structure $\mathcal{M}_{\mathcal{S}_1}(q)$.

The interpretation of the $((A))$ clause requires more explanation. Assuming solidarity, it may be *commonly inferred* by members of A (i.e., "announced") that no member of A will choose an individual action that that member knows will guarantee that A will *not* achieve its aim. More precisely, if we are considering whether A can jointly achieve ψ, the generated epistemic structure may be updated by eliminating every state at which $\bigvee_{a \in A} K_a \neg \psi$ holds. Equivalently, it may be "announced" that $\bigwedge_{a \in A} \hat{K}_a \psi$ holds. After this update, if any agent in the coalition has an individual action that she knows will guarantee success, then it is assured that the coalition can enforce their goal. More precisely, if there exists a possible joint action, in the updated generated epistemic structure, of which it is true that $\bigvee_{a \in A} K_a \psi$, then A can enforce ψ.

3.6. Equivalence.

THEOREM 3.11. *For every 2-CGS $\mathcal{S}$, state q and $\varphi \in \mathcal{L}_{\mathrm{SCL}^2}$, we have*:

$$\mathcal{S}, q \vDash_{SCL} \varphi \text{ iff } \mathcal{S}, q \vDash^e_{SCL} \varphi.$$

Proof. See Appendix 7. □

From here out we switch between the two versions of our semantics whenever it is convenient, and no longer distinguish $\vDash_{SCL}$ and $\vDash^e_{SCL}$.

3.7. Application to examples.

Consider $\mathcal{S}_1$ from Example 3.1, with q denoting the start state. The generated epistemic structure $\mathcal{M}_{\mathcal{S}_1}(q)$ is, again, represented in Figure 4. This structure is unchanged after announcement of $\bigwedge_{a \in A} \hat{K}_a \psi$, since there is no possible action at which either $K_{\mathbf{a}} \neg \psi$ or $K_{\mathbf{b}} \neg \psi$. Further, there is no possible action at which either $K_{\mathbf{a}} \psi$ or $K_{\mathbf{b}} \psi$. In total, we have that for any possible action $\sigma \in \mathrm{D}(q)$, it is false that $\mathcal{M}_{\mathcal{S}_1}(q), \sigma \vDash_{PAL} \langle \bigwedge_{a \in \mathcal{A}} \hat{K}_a \psi \rangle \bigvee_{a \in \mathcal{A}} K_a \psi$. Thus, it is false that $\mathcal{S}_1, q \vDash ((A)) \psi$, as desired. Nevertheless, it is true that $\mathcal{S}_1, q \vDash \langle\langle A \rangle\rangle \psi$.

The case of $\mathcal{S}_2$, from Example 3.2, is easy to evaluate. The announcement of $\bigwedge_{a \in \mathcal{A}} \hat{K}_a \psi$ leaves the generated epistemic structure $\mathcal{M}_{\mathcal{S}_2}(q)$ unchanged (see Figure 5). Notice, however, that $K_{\mathbf{a}} \psi$ is true of $\langle 0, 0 \rangle$, so $\mathcal{S}_2, q \vDash ((A)) \psi$, as desired.

Now consider $\mathcal{S}3$ from Example 3.3, with q again denoting the state at which the agents choose an action. The generated epistemic structure $\mathcal{M}_{\mathcal{S}_3}(q)$ is represented in Figure 5, along with the updated structure after the announcement of $\bigwedge_{a \in \mathcal{A}} \hat{K}_a \psi$. Since at every possible action other than $\langle 0, 0 \rangle$ either $K_{\mathbf{a}} \neg \psi$ or $K_{\mathbf{b}} \neg \psi$, the structure is reduced to one possible action after this announcement. Since it is then true of $\langle 0, 0 \rangle$ that $K_{\mathbf{a}} \psi$, we have a possible action $\sigma \in \mathrm{D}(q)$ such that $\mathcal{M}_{\mathcal{S}_3}(q), \sigma \vDash_{PAL} \langle \bigwedge_{a \in \mathcal{A}} \hat{K}_a \psi \rangle \bigvee_{a \in \mathcal{A}} K_a \psi$. Thus, it is true that $\mathcal{S}_3, q \vDash ((A)) \psi$, as desired (not to mention it is true that $\mathcal{S}_3, q \vDash \langle\langle A \rangle\rangle \psi$).

The reader can verify that the semantics of SCL^2 matches our intuitive assessment of Examples 3.4 and 3.5.
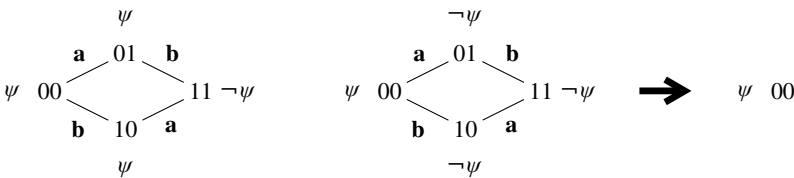

Fig. 5. On the left, the generated epistemic structure $\mathcal{M}_{\mathcal{S}_2}(q)$. On the right, the simple generated epistemic structure $\mathcal{M}_{\mathcal{S}_3}(q)$, with update after the announcement of $\bigwedge_{a \in \mathcal{A}} \hat{K}_a \psi$.

§4. Technical results: CL versus SCL.

4.1. Expressivity. In this section, we study the relative expressivity of SCL^2 and CL. Is the logic SCL^2 able to express, with the $((A))$ operator, a property that cannot be expressed in CL? To settle this, we first introduce some useful machinery (Ågotnes, Goranko, & Jamroga, 2007).

DEFINITION 4.1 (Equivalence of formulae). *Suppose that $\mathcal{L}_1$ and $\mathcal{L}_2$ are interpreted on the same class of models $\mathbb{M}$ (in accordance with satisfaction relations $\vDash_1$ and $\vDash_2$, respectively). Consider formulae $\varphi_1 \in \mathcal{L}_1$ and $\varphi_2 \in \mathcal{L}_2$. Then φ_1 and φ_2 are* equivalent *just in case they are true in the same states (that is, for any $\mathcal{M} \in \mathbb{M}$ and $q \in \mathcal{M}$, we have that $\mathcal{M}, q \vDash_1 \varphi_1$ iff $\mathcal{M}, q \vDash_2 \varphi_2$). We denote this by $\varphi_1 \equiv \varphi_2$.*

DEFINITION 4.2 (Expressive power). *Suppose that $\mathcal{L}_1$ and $\mathcal{L}_2$ are interpreted on the same class of models. $\mathcal{L}_2$ is at least as expressive as $\mathcal{L}_1$ iff for every $\varphi_1 \in \mathcal{L}_1$ there is $\varphi_2 \in \mathcal{L}_2$ such that $\varphi_1 \equiv \varphi_2$. We write $\mathcal{L}_1 \preceq \mathcal{L}_2$. Further, if $\mathcal{L}_1 \preceq \mathcal{L}_2$ and $\mathcal{L}_2 \npreceq \mathcal{L}_1$ we write $\mathcal{L}_1 \prec \mathcal{L}_2$, and say that $\mathcal{L}_2$ is* more expressive *than $\mathcal{L}_1$.*

We present a notion of bisimulation for CGSs, following (Ågotnes *et al.*, 2007).

DEFINITION 4.3 (Bisimulation for CGSs). *Let CGSs $\mathcal{S}_1 = \langle k, Q_1, d_1, \delta_1, \Pi_1, \pi_1 \rangle$ and $\mathcal{S}_2 = \langle k, Q_2, d_2, \delta_2, \Pi_2, \pi_2 \rangle$ be given, with $\mathcal{A} = \{1, 2, \ldots, k\}$.*

1. *Let a set of agents $A \subseteq \mathcal{A}$ be given. A relation $\beta \subseteq Q_1 \times Q_2$ is a (global) A-bisimulation between $\mathcal{S}_1$ and $\mathcal{S}_2$, denoted $\mathcal{S}_1 \rightleftarrows_\beta^A \mathcal{S}_2$, iff for any $q_1 \in Q_1$ and $q_2 \in Q_2$, $q_1 \beta q_2$ implies that*

   Local harmony *$\pi_1(q_1) = \pi_2(q_2)$*
   Forth *For any A-move σ_1^A at q_1, there exists an A-move σ_2^A at q_2 such that for every successor state $q_2^* \in succ(\mathcal{S}_2, q_2, \sigma_2^A)$ there exists a successor state $q_1^* \in succ(\mathcal{S}_1, q_1, \sigma_1^A)$ such that $q_1^* \beta q_2^*$*
   Back *Likewise, for 1 and 2 swapped*

3. *If β is a A-bisimulation between $\mathcal{S}_1$ and $\mathcal{S}_2$ for every $A \subseteq \mathcal{A}$, we call it a (full global) bisimulation between $\mathcal{S}_1$ and $\mathcal{S}_2$, denoted $\mathcal{S}_1 \rightleftarrows_\beta \mathcal{S}_2$*

THEOREM 4.4. *If $\mathcal{S}_1 \rightleftarrows_\beta \mathcal{S}_2$ and $q_1 \beta q_2$, then, for every formula $\varphi \in \mathcal{L}_{CL}(\Pi, \mathcal{A})$, we have that $\mathcal{S}_1, q_1 \vDash_{CL} \varphi$ iff $\mathcal{S}_2, q_2 \vDash_{CL} \varphi$.*

Proof. See the appendix of Ågotnes *et al.*, (2007). □

An appropriate notion of bisimulation in hand, a standard strategy for showing that $\mathcal{L}$ has expressive power beyond $\mathcal{L}_{CL}$ presents itself: if we can find two models such that 1) the models are bisimilar to each other (and so cannot be distinguished by $\mathcal{L}_{CL}$) and 2) there is some formula from $\mathcal{L}$ that holds on the one model but not the other, then we conclude that $\mathcal{L}$ can express distinctions that CL cannot.

THEOREM 4.5. *$\mathcal{L}_{CL^2} \prec \mathcal{L}_{SCL^2}$.*

Proof. $\mathcal{L}_{CL^2} \preceq \mathcal{L}_{SCL^2}$ is immediate. Of more interest: $\mathcal{L}_{SCL^2} \npreceq \mathcal{L}_{CL^2}$. Consider the 2-CGS $\mathcal{S}_6$, as depicted in Figure 6 and the 2-CGS $\mathcal{S}_3$ in Example 3.3, and depicted (again) in Figure 6. In $\mathcal{S}_6$, each agent has three moves from which to choose at the start state—namely, 0, 1, and 2. For readability, we label the states with the joint action that leads to that state.
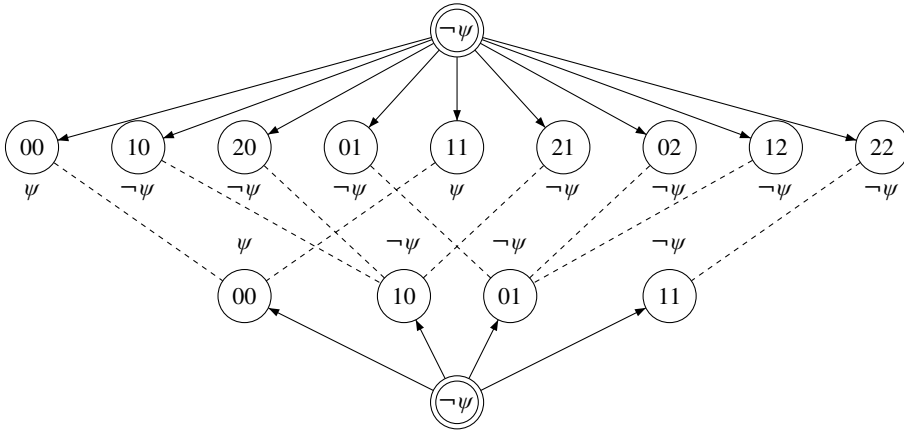
Fig. 6. The structure $\mathcal{S}_6$ (above), the structure $\mathcal{S}_3$ (below), and the bisimulation between them (the broken lines between states).

We know from the discussion of Example 3.3 that $\mathcal{S}_3, q \vDash ((\mathcal{A}))\psi$, where q refers to the start state in $\mathcal{S}_3$ (for convenience, we refer to the start state in both structures as q). However, it is straightforward to check that $\mathcal{S}_6, q \nvDash ((\mathcal{A}))\psi$: the generated epistemic model for $\mathcal{S}_6$ at q is updated to eliminate all joint actions in which some agent chooses action 2, but this just leaves us with the epistemic structure generated by $\mathcal{S}_1$ in Example 3.1.

We claim that a full global bisimulation exists between $\mathcal{S}_3$ and $\mathcal{S}_6$. We depict this bisimulation with the broken lines between states in the respective structures in Figure 6 (the start states in $\mathcal{S}_3$ and $\mathcal{S}_6$ are also related by this relation, but we omit this in the figure for readability). Thus, no CL^2 formula can distinguish $\mathcal{S}_3$ from $\mathcal{S}_6$. Since the SCL-formula $((\mathcal{A}))\psi$ distinguishes the two structures, there is no equivalent formula in CL^2 for the formula $((\mathcal{A}))\psi$ in SCL^2.   □

4.2. Validities. We now consider some significant validities (and invalidities) of SCL^2. In particular, we present a sound and complete axiomatization. For comparison, we start with a complete axiomatization for CL.

PROPOSITION 4.6 (Axiomatization for CL). *The following schemata constitute a sound and complete proof system for CL, including* CL^2:

- *All propositional tautologies*
- $\bot_{\langle\langle\rangle\rangle} : \neg\langle\langle A\rangle\rangle\bot$
- $\top_{\langle\langle\rangle\rangle} : \langle\langle A\rangle\rangle\top$
- $\mathcal{A}$**-duality** $: \neg\langle\langle\emptyset\rangle\rangle\neg\psi \to \langle\langle\mathcal{A}\rangle\rangle\psi$
- **Superadditivity**$_{\langle\langle\rangle\rangle} : \langle\langle A\rangle\rangle\psi_1 \wedge \langle\langle B\rangle\rangle\psi_2 \to \langle\langle A \cup B\rangle\rangle(\psi_1 \wedge \psi_2)$, *if* $A \cap B = \emptyset$
- **Modus ponens** $:$ *from* ψ_1 *and* $\psi_1 \to \psi_2$, *infer* ψ_2
- $\langle\langle A\rangle\rangle$**-monotonicity** $:$ *from* $\psi_1 \to \psi_2$, *infer* $\langle\langle A\rangle\rangle\psi_1 \to \langle\langle A\rangle\rangle\psi_2$

Proof. See Pauly (2001), Goranko *et al.*, (2013), and van Drimmelen & Goranko (2006).   □

DEFINITION 4.7 (Axiomatization for SCL^2). *The set of validities for* SCL^2 *includes all those based on the following schemata* (*with* $A, B \subseteq \mathcal{A}$ *and* $a \in \mathcal{A}$):

- All propositional tautologies
- $\langle\langle\rangle\rangle$-*axioms*: *as in the axiomatization for* CL^2

- $(())$-*axioms*:

  $\emptyset : \neg ((\emptyset)) \top$

  Coalition monontonicity$_{(())}$ $: ((A))\psi \to ((B))\psi$ *when* $A \subseteq B$
- Interaction axioms:

  Int1 $: ((A))\psi \to \langle\langle A \rangle\rangle \psi$

  Int2 $: \langle\langle a \rangle\rangle \psi \to ((a))\psi$

  Int3 $: \left(\bigwedge_{a \in A} \neg \langle\langle a \rangle\rangle \psi \right) \wedge \left(\bigwedge_{a \in A} \neg \langle\langle a \rangle\rangle \neg \psi \right) \to \neg ((A))\psi$
- Rules of inference:

  Modus ponens

  $\langle\langle A \rangle\rangle$**-monotonicity**

  $((A))$**-equivalence**: *from* $\psi_1 \leftrightarrow \psi_2$ *infer* $((A))\psi_1 \leftrightarrow ((A))\psi_2$.

Coalition monotonicity says that a joint ability cannot be lost with the addition of more agents to the coalition. **Superadditivity**$_{(())}$ says that disjoint coalitions can pool their abilities—so long as they can communicate. The **Int1**-axiom says that if a coalition can enforce something without communication, then they can enforce it with communication. The **Int2**-axiom says that the ability of individual agents does not depend on communication. The **Int3**-axiom says that, for a coalition that can't communicate, if none of the agents can individually ensure φ and none have individual actions that are rationally avoided if the goal is φ, then the coalition cannot enforce φ.

We mention three useful and illuminating derivable validities. First: $\langle\langle a \rangle\rangle\varphi \to ((A))\varphi$. Second: $\langle\langle a \rangle\rangle\varphi \wedge \langle\langle b \rangle\rangle\neg\varphi \to \bot$. Third: **Coalition monotonicity**$_{(())}$—that is, $\langle\langle A \rangle\rangle\psi \to \langle\langle B \rangle\rangle\psi$ when $A \subseteq B$—follows from **Superadditivity**$_{(())}$.

THEOREM 4.8 (Soundness). *The above axioms are valid and the inference rules preserve validity.*

Proof. The proof is routine. We prove two of the results to give a feel for the reasoning. The rest is left as an exercise. Throughout, let $\mathcal{S}$ refer to an arbitrary 2-CGS, q an arbitrary state in $\mathcal{S}$ and $\mathcal{M}$ the generated epistemic structure $\mathcal{M}_{\mathcal{S}}(q)$.

Superadditivity$_{(())}$: suppose that

$$\mathcal{S}, q \vDash \langle\langle A \rangle\rangle\psi_1 \wedge \langle\langle B \rangle\rangle\psi_2,$$

where A and B are disjoint. By our semantics, there exist two joint actions σ_1 and σ_2 such that $\mathcal{M}, \sigma_1 \vDash D_A \psi_1$ and $\mathcal{M}, \sigma_2 \vDash D_B \psi_2$. Now consider any joint action σ at q in which each agent $a \in A$ performs the action σ_1^a and each agent $b \in B$ performs the action σ_2^b. Then

$$\mathcal{M}, \sigma \vDash \psi_1 \wedge \psi_2.$$

It follows that $D_{A \cup B}(\psi_1 \wedge \psi_2)$ holds at every such σ, and so

$$\mathcal{S}, q \vDash \langle\langle A \cup B \rangle\rangle(\psi_1 \wedge \psi_2).$$

Int3: suppose that

$$\mathcal{S}, q \vDash \bigwedge_{a \in A} \neg \langle\langle a \rangle\rangle \psi$$

and that

$$\mathcal{S}, q \vDash \bigwedge_{a \in A} \neg \langle\langle a \rangle\rangle \neg \psi.$$

Now, by definition, there exists $\sigma \in D(q)$ such that $\mathcal{M}, \sigma \vDash K_a \neg \psi$ just in case $\mathcal{S}, q \vDash \langle\langle a \rangle\rangle \neg \psi$. Hence, for all $\sigma \in D(q)$ and $a \in A$, we have that $\mathcal{M}, \sigma \nvDash K_a \neg \psi$. Thus, the update following the announcement of $\bigwedge_{a \in A} \hat{K}_a \psi$ leaves the model unchanged. This means that $\mathcal{S}, q \vDash ((A))\psi$ can hold only if there exists $\sigma \in D(q)$ such that $\mathcal{M}, \sigma \vDash K_a \psi$ for some $a \in A$. However, there exists $\sigma \in D(q)$ such that $\mathcal{M}, \sigma \vDash K_a \psi$ just in case $\mathcal{S}, q \vDash \langle\langle a \rangle\rangle \psi$, which is assumed to be false for all $a \in A$. So, for all $\sigma \in D(q)$ and $a \in A$, it follows that $\mathcal{M}, \sigma \nvDash K_a \psi$, and so $\mathcal{S}, q \nvDash ((A))\psi$. $\square$

THEOREM 4.9 (Completeness). *The above axiom system is complete.*

A sketch of the proof is left to Appendix 8. The proof also establishes the following theorem, since it is shown that every consistent SCL^2 formula is satisfiable by a *finite* CGS.

THEOREM 4.10 (Finite model property). SCL^2 *has the finite model property.*

COROLLARY 4.11 (Decidability). *The satisfiability problem for* SCL^2 *is decidable.*

4.3. Invalidities.   Finally, we state some *invalidities* of SCL^2, illustrating significant logical differences between $\langle\langle A \rangle\rangle$ and $((A))$.

PROPOSITION 4.12 (Invalidities). *The following are not valid for* SCL^2:

(1) $\langle\langle \mathbf{a} \rangle\rangle \psi_1 \wedge \langle\langle \mathbf{b} \rangle\rangle \psi_2 \rightarrow ((\mathcal{A}))(\psi_1 \wedge \psi_2)$,
(2) *from* $\psi_1 \rightarrow \psi_2$, *conclude* $((A))\,\psi_1 \rightarrow ((A))\,\psi_2$.

Proof.

(1) Example 3.4 provides a counter-example.
(2) Example 3.5 provides a counter-example: $\psi_1 \rightarrow (\psi_1 \vee \psi_2)$ is valid, yet for $\mathcal{S}_5$ in Example 3.5 it holds that $\mathcal{S}_5, q \vDash ((A))\,\psi_1 \wedge \neg ((A))\,(\psi_1 \vee \psi_2)$. $\square$

§5. Further directions.

5.1. Extension to arbitrary coalitions.   CL and ATL allow for the study of joint ability in settings with more than two agents. Can we extend SCL^2 to similar effect, in a natural manner that continues to do justice to our important test cases? We believe this is likely, but the difficulties should not be under-estimated. To illustrate, we show that a naive attempt to extend the semantics of SCL^2 runs into intuitive counter-examples. Let $\mathcal{S}$ be a three-player CGS. Suppose that we essentially leave the semantics for SCL^2 untouched for this novel setting:

- $\mathcal{S}, q \vDash \langle\langle A \rangle\rangle \varphi$ iff there exists a joint move $\sigma \in D(q)$ such that

$$\mathcal{M}(q), \sigma \vDash_{PAL} D_A \varphi.$$

- $\mathcal{S}, q \vDash ((A))\,\varphi$ iff there exists a joint move $\sigma \in D(q)$ such that

$$\mathcal{M}(q), \sigma \vDash_{PAL} \langle \bigwedge_{a \in A} \hat{K}_a \varphi \rangle \bigvee_{a \in A} K_a \varphi.$$

We now identify a game for which, we claim, sound informal reasoning indicates that a certain coalition can enforce a certain outcome, yet the above semantics delivers the opposite conclusion.

EXAMPLE 5.1. *Consider the interaction between* $\mathbf{a}$, $\mathbf{b}$, *and* $\mathbf{c}$ *in Figure 7. Each agent has two actions, 0 and 1. For readability, we label the outcome states with the joint action that*
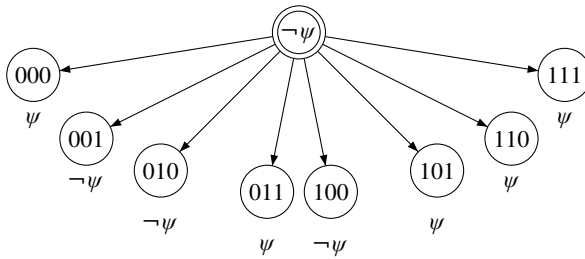
Fig. 7. Structure $\mathcal{S}_7$ from Example 5.1.

leads to that outcome. Suppose that **a** *and* **b** *are in a coalition with goal* ψ. **c** *is outside of the coalition, and so a potential source for the derailment of any strategy carried out by* **a** *and* **b**. *Suppose further that* **a** *and* **b** *cannot communicate* (*though there is common knowledge of solidarity between them*). *That is, this is again a tacit game. Reasoning informally, can* **a** *and* **b** *enforce* ψ?

It is easier, perhaps, to make a judgment when viewing the generated epistemic structure for this CGS. We depict this structure in Figure 8. For readability, we represent the epistemic relations for **a** *and* **b** *in the following manner* (*we ignore the relations for* **c**): *if only* **a** *is unable to distinguish two joint actions, then this is represented with a broken line composed of dashes; if only* **b** *is unable to distinguish two joint actions, then this is represented with a broken line composed of dots; if both* **a** *and* **b** *are unable to distinguish two joint actions, we represent this with an unbroken line. We may think of each player as choosing between two facets of the cube in Figure 8:* **a** *chooses between left and right;* **b** *chooses between front and back;* **c** *chooses between top and bottom. Thus, to ask if* **a** *and* **b** *can enforce* ψ *is to ask if each can choose a respective facet so that the overlap is comprised only of joint actions that bring about* ψ.

So, can **a** *and* **b** *enforce* ψ? *Intuitively, yes: if suitably rational,* **a** *and* **b** *both play action 1. Informally, the reasoning is as follows: to select 0 is to land the coalition with a strategy that can be thwarted by* **c**, *no matter what the other agent does. On the other hand, playing 1 leaves open the possibility of* both *playing 1, the* only *joint strategy that carries no threat from* **c**. *Action 1 "dominates" 0, in this sense. Further, each player can rely on the other to act on this reasoning due to common knowledge of solidarity.*

Yet according to the proposed semantics: $\mathcal{S}_7, q \nvDash ((\{\mathbf{a}, \mathbf{b}\}))\psi$, *since for every joint action* $\sigma \in \mathrm{D}(\sigma)$ *we have* $\mathcal{M}_{\mathcal{S}_7}, \sigma \vDash \hat{K}_\mathbf{a}\psi \wedge \hat{K}_\mathbf{b}\psi \wedge (\neg K_\mathbf{a}\psi) \wedge (\neg K_\mathbf{b}\psi)$.
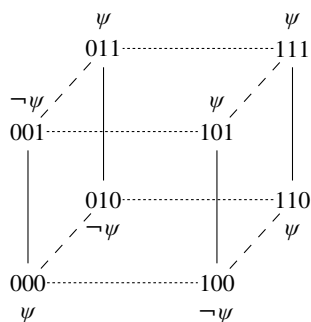

Fig. 8. The generated epistemic structure $\mathcal{M}_{\mathcal{S}_7}(q)$.

There is a less obvious extension of the SCL^2 semantics that better matches our intuitive assessment of the above example: $\mathcal{S}, q \vDash ((A))\,\varphi$ iff there exists a joint move $\sigma \in \mathrm{D}(q)$ such that

$$\mathcal{M}(q), \sigma \vDash_{PAL} \langle \bigwedge_{a \in A} \hat{K}_a D_A \varphi \rangle \bigvee_{a \in A} K_a \varphi.$$

Recall that it is reasonable to here read $D_A \varphi$ as expressing that the joint contribution of A to σ guarantees the outcome φ, no matter what the agents outside of A do. The clause may therefore be interpreted as follows: $((A))\varphi$ holds at q just in case were the agents in A to ignore all actions that satisfy the condition "the agent for which that action is available knows that that action does not form part of a winning joint action for the coalition", then at least one agent has an action she knows ensures success. The reader can verify two promising facts: (i) the proposed clause delivers the intuitive assessment of our last game: **a** and **b** can enforce ψ; (ii) when applied to the two-player setting, the clause is equivalent to that for SCL^2. That is, our proposal provides a *genuine extension* of the semantics for SCL^2.[19]

This proposal is only tentative. We predict that many subtle phenomena emerge for joint ability in tacit games with more than two players.

5.2. Complex communication networks.   We have assumed thus far that communication channels always run in both directions. It is also of interest to study situations where communication runs in only one direction. Further, consideration of more than two agents opens the possibility of disjoint communication networks within a coalition. For instance, **a** and **b** may be able to communicate, while **c** can communicate with neither. These more sophisticated settings can be studied by enriching our structures with *communication relations* (cf. Chopra *et al*., (2004)). However, it is in the spirit of the current paper to instead exploit the existing resources in a CGS, by introducing a more sophisticated logical language. Suppose that we enrich a propositional language with coalitional modalities of the following form: $\langle\langle \mathbf{A} \rangle\rangle$, where $\mathbf{A}$ is a *set* of sets of agents. The intended interpretation of the expression $\langle\langle \mathbf{A} \rangle\rangle\varphi$ is then: where $\mathbf{A} = \{A_1, A_2, \ldots, A_n\}$, the coalition $\bigcup_{i \leq n} \mathbf{A}$ can enforce φ if each subcoalition A_i can communicate among themselves (and even if there is no communication *between* any distinct subcoalitions A_i and A_j). Note that such a language is strictly more expressive than that for SCL^2, for $\langle\langle A \rangle\rangle\varphi$ can be defined as $\langle\langle \{A\} \rangle\rangle\varphi$ and $((A))\varphi$ can be defined as $\langle\langle \{\{a_1\}, \{a_2\}, \ldots, \{a_k\}\} \rangle\rangle\varphi$ where $A = \{a_1, a_2, \ldots, a_k\}$. Our semantics for SCL^2 can be extended to this setting in a natural way: $\mathcal{S}, q \vDash_{SCL} \langle\langle \mathbf{A} \rangle\rangle\varphi$ iff there exists a joint move $\sigma \in \mathrm{D}(q)$ such that

$$\mathcal{M}(q), \sigma \vDash_{PAL} \langle \bigwedge_{A \in \mathbf{A}} \hat{D}_A \varphi \rangle \bigvee_{A \in \mathbf{A}} D_A \varphi.$$

5.3. Stronger notions of rationality.   We have thus far assumed that our agents have (what we call) minimal rationality: they select from actions that necessitate that their goals are met, if available, and will not voluntarily perform actions that necessitate that their goals are not met. This assumption can be strengthened. For instance, one might consider agents that are (what we call) *robustly rational*: such an agent does not voluntarily select actions that are dominated, in a sense familiar to game theorists.

---

[19] The key points: $D_{\mathcal{A}}\varphi$ is equivalent to φ, while $D_a \varphi$ is equivalent to $K_a \varphi$ and $\hat{K}_a K_a \varphi$ is equivalent to $\hat{K}_a \varphi$.

DEFINITION 5.2. *Consider agent a and an individual action α for a at state q. α is dominated relative to goal φ iff there is an action β for a at state q such that*:

- *for every joint action σ where a contributes α to σ: if σ leads to a φ state, then so does σ*, where σ* is identical with σ except that a contributes β,*
- *there exists a joint action σ where a contributes α to σ, σ leads to a ¬φ state and σ* (as above) leads to a φ state.*

This suggests a new coalitional modality $\lceil\lceil A\rceil\rceil\varphi$, where $\mathcal{S}, q \vDash_{SCL} \lceil\lceil A\rceil\rceil\varphi$ iff there is an agent a in A with an individual action that enforces φ in every outcome that is reachable by a joint action that is not eliminated by the iterated removal of dominated individual actions (for agents in coalition A, and relative to goal φ).

Unsurprisingly, ascribing (common knowledge of) powerful reasoning skills to our agents enhances their joint abilities (at the cost of wading deeper into idealization). To see this, consider the following game, played by agents **a** and **b**.

$$\begin{bmatrix} \neg\varphi & \varphi & \varphi \\ \varphi & \varphi & \neg\varphi \\ \varphi & \neg\varphi & \neg\varphi \end{bmatrix}$$

Each agent has three actions available (0,1,2). The rows in the above matrix representing the actions for **a** and the columns the actions for **b**. The entries in the matrix indicate the outcome (relative to φ) for each joint action. Though it is cumbersome to depict, this game can be captured by a CGS, played at state q. Consider the coalition A of **a** and **b**. Can they enforce φ, if the game is tacit? Note that no agent has an action that she knows (before any reasoning) will guarantee φ, nor does she have an action that she knows will guarantee $\neg\varphi$. This is enough to ensure $\neg((A))\varphi$ (cf. axiom **Int3**). However, note further that action 2 for **b** is dominated by action 1, and that if action 2 is "eliminated" by iterated dominance reasoning, then agent **a** has an action that she *then* knows will guarantee φ: namely, action 1. Thus, on the basis of our rough semantics for $\lceil\lceil A\rceil\rceil\varphi$, it is evident that $\lceil\lceil A\rceil\rceil\varphi$ holds at q.

§6. Conclusion and further work. We have introduced and motivated a new variation on Coalition Logic: Two-Player Strategic Coordination Logic. Its purpose is to deal more effectively with situations of almost perfect information, and the accompanying effects on the ability of coalitions to coordinate. We have put forward a precise syntax and semantics for SCL^2, and provided technical results concerning expressivity and validity that illuminate differences between CL and SCL^2.

There is scope for further research. Certain technical matters are not dealt with in this paper, such as the complexity of model-checking. Another route is to explore further refinements of Strategic Coordination Logic. The considerations of 5 show that various subtle extensions of SCL^2 are worth a spirited investigation.

§7. Appendix A: Proof of Theorem 3.11. By induction on the complexity of φ, with induction hypothesis: for every 2-CGS $\mathcal{S}$ and state q, we have $\mathcal{S}, q \vDash_{SCL} \varphi$ just in case $\mathcal{S}, q \vDash_{SCL}^e \varphi$. The nontrivial cases are provided by the coalitional atoms. For the $\langle\langle A\rangle\rangle\varphi$ clause, note the equivalence of the following.

  i. there is an A-move σ^A at q s.t. $\mathcal{S}, q' \vDash_{SCL} \varphi$ for every $q' \in succ(q, \sigma^A)$.
  ii. there exists a joint move $\sigma \in D(q)$ s.t. for all $\tau \in D(q)$, if $\sigma^a = \tau^a$ for all $a \in A$, then.

  iii. there exists a joint move $\sigma \in \mathrm{D}(q)$ s.t. for all $\tau \in \mathrm{D}(q)$, if $\sigma^a = \tau^a$ for all $a \in A$, then $\mathcal{S}, \delta(q, \tau) \vDash^e_{SCL} \varphi$.

  iv. there exists a joint move $\sigma \in \mathrm{D}(q)$ s.t. for all $\tau \in \mathrm{D}(q)$, if $\sigma^a = \tau^a$ for all $a \in A$, then $\mathcal{M}(q), \tau \vDash_{PAL} \varphi$.

  v. there exists a joint move $\sigma \in \mathrm{D}(q)$ s.t. for all $\tau \in \mathrm{D}(q)$, if $\sigma \sim_A \tau$ then $\mathcal{M}(q), \tau \vDash_{PAL} \varphi$.

  vi. there exists a joint move $\sigma \in \mathrm{D}(q)$ s.t. $\mathcal{M}(q), \sigma \vDash_{PAL} D_A \varphi$.

The equivalence between ii and iii is a consequence of our inductive hypothesis. The equivalence between iv and v is a consequence of the definition of $\sim$ in a generated epistemic structure.

For the $((A))\varphi$ clause: note first that, as a matter of the definition of $\mathcal{M}_\mathcal{S}(q)$, an action $\alpha \in \mathrm{D}_a(q)$ is φ-solidary in $\mathcal{S}$ at q just in case

$$\mathcal{M}_\mathcal{S}(q), \sigma \vDash_{PAL} \neg K_a \neg \varphi$$

for every σ where $\sigma^a = \alpha$. Thus, σ is a φ-solidary joint action exactly when $\mathcal{M}_\mathcal{S}(q), \sigma \vDash_{PAL} \bigwedge_{a \in A} \hat{K}_a \varphi$. Thus, the following are equivalent:

  i. $\mathcal{S}, q \vDash_{SCL} \langle\langle A \rangle\rangle \varphi$ and there is an agent $a \in A$ and action $\alpha \in \mathrm{D}_a(q)$ such that for every φ-solidary joint action σ we have $\mathcal{S}, \delta(q, \sigma^*) \vDash_{SCL} \varphi$ (where a's contribution to σ^* is action α, and otherwise the components of σ^* match σ).

  ii. $\mathcal{S}, q \vDash_{SCL} \langle\langle A \rangle\rangle \varphi$ and there is an agent $a \in A$ and action $\alpha \in \mathrm{D}_a(q)$ such that for every φ-solidary joint action σ we have $\mathcal{S}, \delta(q, \sigma^*) \vDash^e_{SCL} \varphi$.

  iii. there exists $a \in A$, an action $\alpha \in \mathrm{D}_a(q)$ and a joint move $\sigma \in \mathrm{D}(q)$ such that: a contributes α to σ and σ is φ-solidary. Further, for any joint action $\tau \in \mathrm{D}(q)$, if a contributes α to τ and τ is φ-solidary, then $\mathcal{S}, \delta(q, \tau) \vDash^e_{SCL} \varphi$.

  iv. there exists a joint action $\sigma \in \mathrm{D}(q)$ and $a \in A$ such that: first,

$$\mathcal{M}(q), \sigma \vDash_{PAL} \bigwedge_{b \in A} \hat{K}_b \varphi$$

and, second, if both $\sigma \sim_a \tau$ and

$$\mathcal{M}(q), \tau \vDash_{PAL} \bigwedge_{b \in A} \hat{K}_b \varphi$$

then $\mathcal{M}(q), \tau \vDash_{PAL} \varphi$ (note that φ is a formula in $\mathcal{L}_{PAL}$ that contains no epistemic operators)

  v. there exists a joint action $\sigma \in \mathrm{D}(q)$ s.t.

$$\mathcal{M}(q), \sigma \vDash_{PAL} \bigwedge_{b \in A} \hat{K}_b \varphi \text{ and } \mathcal{M}(q)|_{\bigwedge_{b \in A} \hat{K}_b \varphi}, \sigma \vDash_{PAL} \bigvee_{b \in A} K_b \varphi.$$

  vi. there exists a joint move $\sigma \in \mathrm{D}(q)$ s.t.

$$\mathcal{M}(q), \sigma \vDash_{PAL} \langle \bigwedge_{b \in A} \hat{K}_b \varphi \rangle \bigvee_{b \in A} K_b \varphi.$$

§8. Appendix B: Completeness proof sketch. Though we do not have the space for every detail, we sketch a proof for the following result: if a $\mathcal{L}_{SCL^2}$ formula ψ is consistent

with respect to the proof system of 4.2 (i.e., $\nvdash \neg\psi$, where $\vdash$ is the provability relation associated with our proof system), then there exists a (finite) 2-CGS $\mathcal{S}_\psi$ and a state q_ψ in $\mathcal{S}_\psi$ such that: $\mathcal{S}_\psi, q_\psi \vDash \psi$. The technique of our proof is heavily influenced by that in (van Drimmelen & Goranko, 2006).

We assume that every formula, including ψ, is in a convenient *normal form*.

DEFINITION 8.1 (Normal form). *A sentence $\varphi \in \mathcal{L}_{SCL^2}$ is in* normal form *just in case negations appear only in front of propositional atoms or coalitional atoms.*

PROPOSITION 8.2. *For every sentence $\varphi \in \mathcal{L}_{SCL^2}$ there is a provably equivalent sentence $\varphi^* \in \mathcal{L}_{SCL^2}$ that is in normal form.*

Proof. Routine. □

One instance in which we are sloppy in our forthcoming discussion is that we write $\neg\varphi$ when we generally mean the *normal form* of this sentence.

We work with a finite set of formulae called the *closure* of ψ. The states of $\mathcal{S}_\psi$ will be maximally consistent subsets of the closure of ψ.

DEFINITION 8.3. *The* closure *for formula φ, denoted $cl(\varphi)$, is the smallest set of formulae where*:

- *if ψ is a subformula of φ then $\psi \in cl(\varphi)$,*
- *if $\psi \in cl(\varphi)$ then $\neg\psi \in cl(\varphi)$,*
- *if $\psi_1, \psi_2 \in cl(\varphi)$ then $\psi_1 \lor \psi_2 \in cl(\varphi)$,*
- *if a coalition formula $\langle\langle \mathbf{a}, \mathbf{b} \rangle\rangle\psi$ or $((\mathbf{a}, \mathbf{b}))\psi$ or $\langle\langle a \rangle\rangle\psi$ is in $cl(\varphi)$, then so is all of $\langle\langle \mathbf{a}, \mathbf{b} \rangle\rangle\psi$, $((\mathbf{a}, \mathbf{b}))\psi$, $\langle\langle a \rangle\rangle\psi$, $\langle\langle \mathbf{a}, \mathbf{b} \rangle\rangle\neg\psi$, $((\mathbf{a}, \mathbf{b}))\neg\psi$ and $\langle\langle a \rangle\rangle\neg\psi$.*

Note then that $cl(\varphi)$ is finite for any φ.

We now name four sets of formulae that will play a role in the definition of the actions available to the agents in $\mathcal{S}_\psi$.

DEFINITION 8.4 (Coalition formulae). *Let Δ be a set of $\mathcal{L}_{SCL^2}$ formulae.*

- *$\Psi_{pos\langle\langle\rangle\rangle}(\Delta)$ is the set of positive communication formulae in Δ, formulae of the form $\langle\langle A \rangle\rangle\varphi$ in Δ.*
- *$\Psi_{neg\langle\langle\rangle\rangle}(\Delta)$ is the set of negative communication formulae in Δ. formulae of the form $\neg\langle\langle A \rangle\rangle\varphi$ in Δ.*
- *$\Psi_{pos(())}(\Delta)$ is the set of positive coordination formulae in Δ, formulae of the form $((A))\varphi$ in Δ.*
- *$\Psi_{neg(())}(\Delta)$ is the set of negative coordination formulae in Δ, formulae of the form $\neg((A))\varphi$ in Δ.*

Ψ is the ordered union of $\Psi_{neg\langle\langle\rangle\rangle}$, $\Psi_{pos\langle\langle\rangle\rangle}$, $\Psi_{pos(())}$, and $\Psi_{neg(())}$ (with negative communication formulae first in the order, and so forth), with Δ withheld if the context is unambiguous. Thus a natural number may be associated with each coalitional atom of interest, called its *index*, ranging from 1 to $l+m+n+o$. If j is the index associated with, for instance, $\langle\langle A \rangle\rangle\varphi$ in Ψ, we sometimes refer to φ as φ_j, and similarly for other coalitional atoms.

We now define the model $\mathcal{S}_\psi$ for ψ. A word on our general approach: $\mathcal{S}_\psi$ is tree-like, of a depth corresponding to the modal depth of ψ. Each state/node s is a (finite) set of formulae: a maximally consistent subset of $cl(\psi)$ generated by a consistent subset Δ_s determined by the coalitional atoms that appear in the parent node. In particular, we associate with the root state a set Γ where $\psi \in \Gamma$. A key step in our proof will be to show that each Δ_s is indeed

consistent. We then prove a truth lemma to the effect that $\varphi \in s$ just in case $\mathcal{S}_\psi, s \vDash \varphi$. In particular, this holds for the root node - and so ψ holds at the root.

DEFINITION 8.5 (Tree model for ψ). *Consider consistent formula ψ and Γ_ψ, where the latter is a maximally consistent subset of $cl(\psi)$ such that $\psi \in \Gamma_\psi$. A tree model for ψ is a 2-CGS $\mathcal{S}_\psi$ that meets the following constraints*:

- $\mathcal{A} = \{1, 2\}$
- Q *is a finite set of maximally consistent subsets of $cl(\psi)$, d an action function and δ a transition function that together satisfy the following constraints*:

    1. *the root state is Γ_ψ.*
    2. *if $\Delta \in Q$, then $d(\Delta, a) = l + m + n + o$, for agent a, where*:

        — $|\Psi_{neg\langle\langle\rangle\rangle}(\Delta)| = l$,
        — $|\Psi_{pos\langle\langle\rangle\rangle}(\Delta)| = m$,
        — $|\Psi_{pos(())}(\Delta)| = n$,
        — $|\Psi_{neg(())}(\Delta)| = o$.

        Thus, we associate with each coalition formula in Δ an action i, available to each agent at Δ, called a vote *for that formula.*
    3. *if joint actions σ_1 and σ_2 at state Δ are not identical, then $\delta(\Delta, \sigma_1) \neq \delta(\Delta, \sigma_2)$.*
    4. *if $\neg\langle\langle \mathbf{a}, \mathbf{b}\rangle\rangle\varphi_i \in \Delta$ (or, equivalently, by $\mathcal{A}$-**duality**: $\langle\langle\emptyset\rangle\rangle\neg\varphi_i \in \Delta$) then $\neg\varphi_i \in \delta(\Delta, \sigma)$ for every joint action σ at state Δ.*
    5. *if $\neg\langle\langle\mathbf{a}\rangle\rangle\varphi_i \in \Delta$ and either $((\mathbf{a}, \mathbf{b}))\varphi_i \notin \Delta$ or $\neg\langle\langle\mathbf{b}\rangle\rangle\varphi_i \notin \Delta$ or $\langle\langle\mathbf{b}\rangle\rangle\neg\varphi_i \notin \Delta$, then $\neg\varphi_i \in \delta(\Delta, \sigma)$ for every joint action σ such that: at least one agent votes for a negative communication formula and*

    $$i = [\sum_{a \in N} \sigma^a] mod \, l$$

    where N is the subset of agents that vote for a negative communication formulae in σ i.e., the negatively voting agents "jointly" vote to block $\langle\langle\mathbf{a}\rangle\rangle\varphi_i$. Likewise with $\mathbf{a}$ and $\mathbf{b}$ reversed.
    6. *if $\langle\langle\mathbf{a}, \mathbf{b}\rangle\rangle\varphi_i \in \Delta$, then $\varphi \in \delta(\Delta, \langle i, i\rangle)$ i.e., if both agents vote for φ_i, then φ_i is achieved.*
    7. *if $\langle\langle\mathbf{a}\rangle\rangle\varphi_i \in \Delta$, then $\varphi \in \delta(\Delta, \langle i, j\rangle)$ for every action j for $\mathbf{b}$ i.e., if agent $\mathbf{a}$ votes for φ, then φ is achieved. Likewise with $\mathbf{a}$ and $\mathbf{b}$ reversed.*
    8. *if $((\mathbf{a}, \mathbf{b}))\varphi_i \in \Delta$ and $\neg\langle\langle\mathbf{a}\rangle\rangle\varphi_i \in \Delta$ and $\neg\langle\langle\mathbf{b}\rangle\rangle\varphi_i \in \Delta$, then $\varphi_i \in \delta(\Delta, \langle i, j\rangle)$ for every j except when j is a vote for a formula $\langle\langle\mathbf{b}\rangle\rangle\chi_j$ where $\vdash \chi_j \rightarrow \neg\varphi_i$; and $\varphi_i \in \delta(\Delta, \langle j, i\rangle)$ for every j except when j is a vote for a formula $\langle\langle\mathbf{a}\rangle\rangle\chi_j$ where $\vdash \chi_j \rightarrow \neg\varphi_i$.*
    9. *if $\neg((\mathbf{a}, \mathbf{b}))\varphi_i \in \Delta$ and $\langle\langle\mathbf{a}, \mathbf{b}\rangle\rangle\varphi_j \in \Delta$ with $\varphi_i = \varphi_j$, then $\varphi_i \in \delta(\Delta, \langle i, i\rangle)$, $\neg\varphi_i \in \delta(\Delta, \langle k, j\rangle)$ and $\neg\varphi_i \in \delta(\Delta, \langle j, k\rangle)$ for every action $k \neq j$.*

- π *assigns a state Δ to proposition letter p iff $p \in \Delta$.*

We next prove that a tree model exists for ψ. A preliminary lemma will be useful.

LEMMA 8.6 (Disjoint coalition consistency). *Let*

$$\{\langle\langle A_1\rangle\rangle\varphi_1, \ldots, \langle\langle A_n\rangle\rangle\varphi_n, \neg\langle\langle A_{n+1}\rangle\rangle\varphi_{n+1}\}$$

be a CL-consistent set of formulae, with $A_1, \ldots, A_n, \mathcal{A} \setminus A_{n+1}$ pairwise disjoint coalitions. Then

$$\{\varphi_1, \ldots, \varphi_n, \neg\varphi_{n+1}\}$$

is CL-consistent.

Proof. As in Lemma 31 of (van Drimmelen & Goranko, 2006). □

LEMMA 8.7 (Existence Lemma). *A tree model $\mathcal{S}_\psi$ exists for every consistent ψ.*

Proof. Given Γ_ψ, we need to show that the constraints placed on the states Q of $\mathcal{S}_\psi$ are realizable i.e., that items 4 through 9 in Definition 8.5 do not mutually entail that a SCL^2-inconsistent set of formulae is contained in a maximally consistent subset Δ of $cl(\psi)$. In other words, we must show that items 4 through 9 are mutually consistent. Since items 5 through 9 together can at most insist that two particular formulae appear in Δ, the proof can be carried out by cases, checking that these possible pairs are not inconsistent with each other or formulae imposed by item 4. For every case, Lemma 8.6 and our axiom system find application. We leave most cases as an exercise for the reader, but illustrate the techniques by showing that constraints 5 to 9 are consistent with constraint 4.

Consider the construction of sets of formulae Δ_σ - treating the constraints of Definition 8.5 as *rules* for adding formulae to each Δ_σ - where Δ is a given SCL^2-consistent set and σ a joint action determined by the contents of Δ, as in Definition 8.5. We argue that the application of rule 4 cannot in conjunction with another rule result in the inconsistency of Δ_σ.

Suppose that $\neg\varphi$ has been added to (each) Δ_σ because $\neg\langle\langle \mathbf{a}, \mathbf{b}\rangle\rangle\varphi \in \Psi_{neg\langle\langle\rangle\rangle}(\Delta)$ (by rule 4 of Definition 8.5). Consider φ^* such that $\vdash \varphi^* \rightarrow \varphi$ (i.e., φ^* and $\neg\varphi$ are inconsistent).

First, for reductio, assume that $\neg\langle\langle \mathbf{a}, \mathbf{b}\rangle\rangle\neg\varphi^* \in \Psi_{neg\langle\langle\rangle\rangle}(\Delta)$. Using the $\mathcal{A}$-**duality** axiom, it follows that $\langle\langle\emptyset\rangle\rangle\varphi^* \in \Delta$. Since Δ is consistent, it follows that

$$\{\neg\langle\langle \mathbf{a}, \mathbf{b}\rangle\rangle\varphi, \langle\langle\emptyset\rangle\rangle\varphi^*\}$$

is consistent. Using Lemma 8.6, it follows that $\{\neg\varphi, \varphi^*\}$ is consistent, a contradiction. So no application of rule 4 can introduce φ into Δ_σ. Second, for reductio, assume that $\langle\langle \mathbf{a}, \mathbf{b}\rangle\rangle\varphi^* \in \Psi_{pos\langle\langle\rangle\rangle}(\Delta)$. Then the $\langle\langle A\rangle\rangle$-monotonicity rule ensures that $\langle\langle \mathbf{a}, \mathbf{b}\rangle\rangle\varphi \in \Delta$, contradicting the consistency of Δ. Hence, φ^* cannot be added by rule 6. Due to axiom **Int1**, the consistency of Δ further ensures that $((\mathbf{a}, \mathbf{b}))\varphi^* \notin \Psi_{pos(())}(\Delta)$. Hence, φ^* cannot be added by rule 8. Next, note that together axioms **Coalition Monotonicity**$_{\langle\langle\rangle\rangle}$, **Int1** and **Int2** give $\langle\langle a\rangle\rangle\phi \rightarrow \langle\langle \mathbf{a}, \mathbf{b}\rangle\rangle\phi$ for every ϕ and a, thus the consistency of Δ ensures that $\langle\langle a\rangle\rangle\varphi^* \notin \Psi_{pos\langle\langle\rangle\rangle}(\Delta)$ for every a. Hence φ^* cannot be added by rule 7. Next, since axiom $\mathcal{A}$-**duality** gives us that $\neg\langle\langle \mathbf{a}, \mathbf{b}\rangle\rangle\varphi \rightarrow \langle\langle\emptyset\rangle\rangle\neg\varphi \in \Delta$, we have that $\langle\langle\emptyset\rangle\rangle\neg\varphi \in \Delta$. Combined with Lemma 8.6, it follows that $\neg\langle\langle a\rangle\rangle\neg\varphi^* \notin \Delta$ for every a. Hence, φ^* cannot be added to Δ_σ using rule 5. Finally, note that rule 9 is inapplicable for the case of $\neg((\mathbf{a}, \mathbf{b}))\varphi^*$ and $\langle\langle \mathbf{a}, \mathbf{b}\rangle\rangle\varphi$, since $\neg\langle\langle \mathbf{a}, \mathbf{b}\rangle\rangle\varphi^* \notin \Delta$; and rule 9 is inapplicable for the case of $\neg((\mathbf{a}, \mathbf{b}))\neg\varphi^*$ and $\langle\langle \mathbf{a}, \mathbf{b}\rangle\rangle\neg\varphi$, since **Coalition** axiom **Monotonicity**$_{\langle\langle\rangle\rangle}$ and $\langle\langle\emptyset\rangle\rangle\neg\varphi^*$ entails that $\langle\langle a\rangle\rangle\neg\varphi^* \in \Delta$, which in turn entails that $((\mathbf{a}, \mathbf{b}))\neg\varphi^* \in \Delta$ by **Int1** and **Coalition Monotonicity**$_{(())}$. Altogether, there is no rule that along with the formulae in $\Psi(\Delta)$ determines that φ^* is added to Δ_σ. □

LEMMA 8.8 (Truth Lemma). *Given tree model $\mathcal{S}$, we have, for every formula φ:*

$$\mathcal{S}, \Delta \vDash \varphi \text{ if } \varphi \in \Delta.$$

Proof. By induction on the structure of φ.

- φ is a proposition letter. Then the constraints on $\mathcal{S}$ give: if $p \in \Delta$ then $\Delta \in \pi(p)$.
- φ is the negation of a proposition letter. Then the constraints on $\mathcal{S}$ give: if $\neg p \in \Delta$ then $\Delta \notin \pi(p)$.
- φ is of the form $\psi_1 \wedge \psi_2$. If $\varphi \in \Delta$ then both $\psi_1 \in \Delta$ and $\psi_2 \in \Delta$, since $\psi_1, \psi_2 \in cl(\varphi)$ and Δ is maximally consistent.
- φ is of the form $\psi_1 \vee \psi_2$. If $\varphi \in \Delta$ then either $\psi_1 \in \Delta$ or $\psi_2 \in \Delta$, since $\psi_1, \psi_2 \in cl(\varphi)$ and Δ is maximally consistent.
- φ is of the form $\neg\langle\langle \mathbf{a}, \mathbf{b} \rangle\rangle\psi$. Then item 4 of Definition 8.5 ensures that $\neg\psi$ holds at every state reachable from Δ by some joint action, and so $\mathcal{S}, \Delta \vDash \varphi$.
- φ is of the form $\langle\langle \emptyset \rangle\rangle\psi$. Then $\mathcal{A}$-**duality** guarantees that $\neg\langle\langle \mathbf{a}, \mathbf{b} \rangle\rangle\neg\psi \in \Delta$, and so item 4 of Definition 8.5 again guarantees that $\mathcal{S}, \Delta \vDash \varphi$.
- φ is of the form $\neg\langle\langle \mathbf{a} \rangle\rangle\psi$ (*mutatis mutandis* for $\mathbf{a}$ replaced by $\mathbf{b}$). There are two cases.
  Case 1. suppose that either $((\mathbf{a}, \mathbf{b}))\psi \notin \Delta$ or $\neg\langle\langle \mathbf{b} \rangle\rangle\psi \notin \Delta$ or $\langle\langle \mathbf{b} \rangle\rangle\neg\psi \notin \Delta$. Then item 5 of Definition 8.5 guarantees that for every action i available to $\mathbf{a}$ at Δ, there is an action j available to $\mathbf{b}$ at Δ such that: $\neg\psi \in \delta(\Delta, \langle i, j \rangle)$ (thereby guaranteeing that $\mathcal{S}, \Delta \vDash \varphi$). If i is not the index for a negative communication formula, then the corresponding action j is the index for the negative communication formula $\neg\langle\langle \mathbf{a} \rangle\rangle\psi$. If i is the index for a negative communication formula then there exists some index j for a negative communication formula such that $[i + j]\mathrm{mod}l$ is equal to the index for the negative communication formula $\neg\langle\langle \mathbf{a} \rangle\rangle\psi$.
  Case 2. suppose that $((\mathbf{a}, \mathbf{b}))\psi \in \Delta$ and $\neg\langle\langle \mathbf{b} \rangle\rangle\psi \in \Delta$ and $\langle\langle \mathbf{b} \rangle\rangle\neg\psi \in \Delta$. Then item 8 of Definition 8.5 guarantees that for every action i for $\mathbf{a}$ there is an action j for $\mathbf{b}$—namely, the index for $\langle\langle \mathbf{b} \rangle\rangle\neg\psi$ - such that $\neg\psi \in \delta(\Delta, \langle\langle i, j \rangle\rangle)$, and so $\mathcal{S}, \Delta \vDash \varphi$.
- φ is of the form $\langle\langle \mathbf{a}, \mathbf{b} \rangle\rangle\psi$. Then item 6 of the Definition 8.5 guarantees the existence of a state reachable from Δ - namely, $\delta(\Delta, \langle i, i \rangle)$, where i is the index for φ—at which ψ holds.
- φ is of the form $\langle\langle a \rangle\rangle\psi$. Then item 7 of Definition 8.5 ensures that the index i for φ is an action that ensures that ψ holds for every joint action where a plays i.
- Since **Int1** and **Int2** give us that $((a))\psi$ is equivalent to $\langle\langle a \rangle\rangle\psi$, cases involving formula of the form $((a))\psi$ and $\neg((a))\psi$ may be understood as covered by previous cases.
- φ is of the form $((\mathbf{a}, \mathbf{b}))\psi$. There are two cases.
  Case 1. either $\langle\langle \mathbf{a} \rangle\rangle\psi \in \Delta$ or $\langle\langle \mathbf{b} \rangle\rangle\psi \in \Delta$. Item 7 of Definition 8.5 guarantees that such formula hold at Δ, and the validity of $\langle\langle a \rangle\rangle\psi \rightarrow ((\mathbf{a}, \mathbf{b}))\psi$ then guarantees that $((\mathbf{a}, \mathbf{b}))\psi$ holds at Δ.
  Case 2. neither $\langle\langle \mathbf{a} \rangle\rangle\psi \in \Delta$ nor $\langle\langle \mathbf{b} \rangle\rangle\psi \in \Delta$. Then item 8 of Definition 8.5 guarantees that, for each agent a, the index i for ψ is an action such that: if every action that ensures $\neg\psi$ is deleted in the generated epistemic model at Δ, then the agent a knows that i ensures ψ in that updated model.
- φ is of the form $\neg((\mathbf{a}, \mathbf{b}))\psi$. There are two cases.
  Case 1. if $\neg\langle\langle A \rangle\rangle\psi \in \Delta$, then this formula is guaranteed to hold at Δ by item 4 of Definition 8.5. The validity of $\neg\langle\langle A \rangle\rangle\psi \rightarrow \neg((\mathbf{a}, \mathbf{b}))\psi$ then ensures that φ holds at Δ.
  Case 2. if $\langle\langle A \rangle\rangle\psi \in \Delta$, then by item 9 of Definition 8.5, we have that $\psi \in \delta(\Delta, \langle i, i \rangle)$, where i is the index for $\neg((\mathbf{a}, \mathbf{b}))\psi$. But now note that by item 6 of

Definition 8.5, we have that $\psi \in \delta(\Delta, \langle j, j \rangle)$, where j is the index for $\langle\langle A \rangle\rangle \psi$. So neither action i nor action j will be deleted from generated epistemic model at Δ, with the announcement that $\langle \hat{K}_{\mathbf{a}} \psi \wedge \hat{K}_{\mathbf{b}} \psi \rangle$. Thus, joint actions $\langle i, j \rangle$ and $\langle j, i \rangle$ appear in the updated generated epistemic model. But since item 9 guarantees that $\neg \psi$ holds at $\delta(\Delta, \langle k, j \rangle)$ and $\delta(\Delta, \langle j, k \rangle)$ for every $k \neq j$, we have that $\neg K_a \psi$ holds everywhere in the updated epistemic model. $\qquad \square$

Our main result follows as an immediate corollary.

THEOREM 8.9. *For every consistent formula ψ there exists a 2-CGS and a state in that structure at which ψ is satisfied.*

BIBLIOGRAPHY

Ågotnes, T., Goranko, V., & Jamroga, W. (2007). Alternating-time temporal logics with irrevocable strategies. In Samet, D., editor. *Proceedings of the 11th Conference on Theoretical Aspects of Rationality and Knowledge* (*TARK XI*). New York, NY, USA: Presses Universitaires de Louvain, pp. 15–24.

Alur, R., Henzinger, T. A., & Kupferman, O. (2002). Alternating-time temporal logic. *Journal of the ACM*, **49**, 672–713.

Bacharach, M. (2006). *Beyond Individual Choice*. Princeton, NJ, USA: Princeton University Press.

Broersen, J., Herzig, A., & Troquard, N. (2006). Embedding alternating-time temporal logic in strategic stit logic of agency. *Journal of Logic and Computation*, **16**(5), 559–578.

Chopra, S., Pacuit, E., & Parikh, R. (2004). Knowledge-theoretic properties of strategic voting. In Alferes, J. and Leite, J., editors. *Proceedings of Logics in Artificial Intelligence: 9th European Conference* (*JELIA*), LNCS, Vol. 3229. Berlin, Heidelberg: Springer, pp. 18–30.

Fagin, R., Halpern, J. Y., Moses, Y., & Vardi, M. Y. (1995). *Reasoning about Knowledge*. Massachusetts, USA: MIT Press.

Ghaderi, H., Lesperance, Y., & Levesque, H. (2007). A logicial theory of coordination and joint ability. In M. Huhns and O. Shehory, editors. *Proceedings of Twenty-Second Conference on Artifical Intelligence* (*AAAI07*). Vancouver, BC: ACM, pp. 421–426.

Goranko, V. & Hawke, P. (2010). On the dynamics of information and abilities of players in multi-player games. In *Proceedings of LOFT 10*. Paper presented at the 9th Conference on Logic and the Foundations of Game and Decision Theory (LOFT 9), University of Toulouse, France, July 2010.

Goranko, V. & Jamroga, W. (2004). Comparing semantics for logics of multi-agent systems. *Synthese*, **139**(2), 241–280.

Goranko, V., Jamroga, W., & Turrini, P. (2013). Strategic games and truly playable effectivity functions. *Autonomous Agents and Multi-Agent Systems*, **26**, 288–314.

Harsanyi, J. (1977). *Rational Behavior and Bargaining Equilibrium in Games and Social Situations*. New York, NY, USA: Cambridge University Press.

Horty, J. & Belnap, N. (1995). A study of action, omission, ability and obligation. *Journal of Philosophical Logic*, **24**, 583–644.

Horty, J. F. (2001). *Agency and Deontic Logic*. New York, NY, USA: Oxford University Press.

Leyton-Brown, K. & Shoham, Y. (2008). *Essentials of Game Theory: A Concise, Multidisciplinary Introduction*. Synthesis Lectures on Artificial Intelligence and Machine Learning. San Rafael, CA: Morgan and Claypool Publishers.

Mele, A. R. (2003). Agent's abilities. *Nous*, **37**(3), 447–470.

Pacuit, E. & Roy, O. (2016). Epistemic foundations of game theory. In Zalta, E. N., editor. The Stanford Encyclopedia of Philosophy. CA, USA: Metaphysics Research Lab, Stanford University. Available at: https://plato.stanford.edu/archives/win2016/entries/epistemic-game/.

Pacuit, E. & Simon, S. (2011). Reasoning with protocols under imperfect information. *The Review of Symbolic Logic*, **4**(3), 412–444.

Pauly, M. (2001). Logic for Social Software. Ph.D. Thesis, University of Amsterdam.

Schelling, T. (1960). *The Strategy of Conflict*. Massachusetts, USA: Harvard University Press.

van Benthem, J. (2007). Rational dynamics and epistemic logic in games. *International Game Theory Review*, **9**(1), 13–45.

van Benthem, J. (2014). *Logic in Games*. Massachusetts, USA: MIT Press.

van der Hoek, W. & Jamroga, W. (2004). Agents who know how to play. *Fundamenta Informaticae*, **62**, 1–35.

van der Hoek, W. & Pauly, M. (2007). Modal logic for games and information. In Blackburn, P., van Benthem, J., and Wolter, F., editors. *Handbook of Modal Logic*. Amsterdam, Netherlands: Elsevier, pp. 1077–1148.

van der Hoek, W., & Wooldridge, M. (2003). Cooperation, knowledge and time: Alternating-time temporal epistemic logic and its applications. *Studia Logica*, **75**, 125–127.

van Ditmarsch, H., Lang, J., & Saffidine, A. (2012). Strategic voting and the logic of knowledge. In *Proceedings of 14th TARK*. Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems - Volume 3, International Foundation for Autonomous Agents and Multi-agent Systems, Richland, SC, USA, pp. 1247–1248.

van Ditmarsch, H., van der Hoek, W., & Kooi, B. (2008). *Dynamic Epistemic Logic*. Dordrecht, The Netherlands: Springer.

van Drimmelen, G. & Goranko, V. (2006). Complete axiomatization and decidability of alternating-time temporal logic. *Theoretical Computer Science*, **353**, 93–117.

DEPARTMENT OF PHILOSOPHY
STANFORD UNIVERSITY
450 SERRA MALL
STANFORD, CA 94305, USA
E-mail: phawke@stanford.edu