

The *Ceteris Paribus* Structure of Logics of Game Forms (Extended Abstract)*

Davide Grossi
University of Liverpool
D.Grossi@liverpool.ac.uk

Emiliano Lorini
IRIT-CNRS, Toulouse University
lorini@irit.fr

François Schwarzentruber
ENS Rennes
francois.schwarzentruber@ens-rennes.fr

Abstract

We present a simple *ceteris paribus* logic based on propositional equivalence, which we call CP. We show that CP provides a yardstick for comparing and unifying existing logics STIT, CL–PC and DL–PA.

1 Introduction

The logical analysis of agency and games—for an expository introduction to the field see van der Hoek and Pauly’s overview paper 2007—has boomed in the last two decades giving rise to a plethora of different logics in particular within the multi-agent systems field.¹ At the heart of these logics are always representations of the possible *choices* (or actions) of groups of players (or agents) and their *powers* to force specific outcomes of the game. Some logics take the former as primitives, like STIT (the logic of *seeing to it that*, [Belnap *et al.*, 2001; Horty, 2001]), some take the latter like CL (*coalition logic*, [Pauly, 2002; Goranko *et al.*, 2013]) and ATL (*alternating-time temporal logic*, [Alur *et al.*, 2002]).

In these formalisms the power of players is modeled in terms of the notion of effectivity. In a strategic game, the α -effectivity of a group of players consists of those sets of outcomes of the game for which the players have some collective action which forces the outcome of the game to end up in that set, no matter what the other players do [Moulin and Peleg, 1982]. So, if a set of outcomes X belongs to the α -effectivity of a set of players J , *there exists* an individual action for each agent in J such that, *for all* actions of the other players, the outcome of the game will be contained in X . If we keep the actions of the other agents fixed, then the selection of an individual action for each agent in J corresponds to a choice of J under the assumption that the other agents stick to their choices.

Scientific Context It was already observed by [van Benthem *et al.*, 2009] that this formalization of choice and power

*The full version of this paper is published as *The Ceteris Paribus Structure in Logics of Game Forms*, Journal of Artificial Intelligence Research, volume 53 (2015) pp. 91–126.

¹The richness of this logical landscape was the object of the IJCAI’13 invited talk by A. Herzig *Logics for Multi-Agent Systems: a Critical Overview*.

in games is of an ‘all other things being equal’, or *ceteris paribus*, nature. Considering which outcomes of a game are possible for a set of players J once the other players have fixed their actions, amounts to considering what may be the case under the *ceteris paribus* condition ‘*all actions of the agents not in J being equal (to their current ones)*’. In the current paper we leverage this intuition further and show how it can provide a novel systematization of many of the most influential formalisms in the field of logic and games.

Formal relationships linking the logics (or fragments thereof) we mentioned above have been object of several publications. Notable examples are: the embedding of CL into the next-time fragment of ATL [Goranko, 2001] and the embedding of CL into NCL (*normal coalition logic*, [Broersen *et al.*, 2007; Balbiani *et al.*, 2008a]), the embedding of CL and ATL into STIT [Broersen *et al.*, 2005; 2006]. Earlier contributions have also attempted more comprehensive systematizations of the field of logic and games. Two in particular are worth mentioning: Goranko and Jamroga’s work 2004, which compares game logics based on the computation tree abstraction like ATL and its variants; and Herzig’s work 2014, which provides a conceptual and syntax-based—while we favor here semantic methods—comparison of all the main formalisms in the literature.

We focus on the components of the semantics of those logics that have directly to do with the representation of choice and power, and we abstract away from the representation of time and repeated interaction. So the logics we will be working with are: the atemporal fragment of STIT, logic CL–PC (*coalition logic of propositional control*, [van der Hoek and Wooldridge, 2005]) and the starless fragment of DL–PA (*dynamic logic of propositional assignments*, [van Eijck, 2000; Balbiani *et al.*, 2013]). These logics cover, arguably, a large spectrum of the most influential existing formalisms. Logic STIT is often considered a standard in the literature, as it embeds both CL and ATL [Broersen *et al.*, 2005; 2006], so we use it as a natural starting point. Logic CL–PC is an influential extension of CL. Finally, logic DL–PA is an extension of PDL (*propositional dynamic logic*, [Harel *et al.*, 2000]), recently proposed in [Herzig *et al.*, 2011; Balbiani *et al.*, 2013].

Contribution To articulate our analysis, whose main technical tool consists of satisfiability-preserving embeddings,




Figure 1: Summary of the embeddings established in the paper and known from the literature

the paper introduces and studies—in its axiomatization and complexity—a simple ceteris paribus logic based on propositional equivalence, which we call CP. Such logic is the yardstick allowing us to compare and unify STIT, CL-PC and DL-PA. Figure 1 gives a graphical presentation of the embeddings established in the paper—as well as relevant ones already known in the literature.

An arrow indicates that each formula of the source logic is satisfiable if and only if a suitable translation of that formula is satisfiable in the target logic. DL-PA denotes the starless version of dynamic logic of propositional assignments, NCL and STIT denote the atemporal version of, respectively, normal coalition logic and the seeing-to-it logic. S5 denotes the normal modal logic of equivalence relations. Dotted lines indicate embeddings known in the literature: from CL-PC to DL-PA [Balbiani *et al.*, 2013] and from STIT to NCL (and vice versa) with respect to fragments of the respective languages [Lorini and Schwarzenrüber, 2011]. The embedding from STIT to CP assumes a bound on the STIT-models. All embeddings are polynomial except for the one from CP to S5.

2 A Simple Ceteris Paribus Logic

2.1 Equivalence Modulo a Set of Atoms

Consider a structure (W, V) where W is a set of states, and $V : \mathbf{P} \rightarrow 2^W$ a valuation function from a countable set of atomic propositions \mathbf{P} to subsets of W .

Definition 1. (Equivalence modulo X) Given a pair (W, V) , $X \subseteq \mathbf{P}$ and $|X| < \omega$, the relation $\sim_X^V \subseteq W^2$ is defined as:

$$w \sim_X^V w' \iff \forall p \in X : (w \in V(p) \iff w' \in V(p))$$

When X is a singleton (e.g. p), we will often write \sim_p^V instead of $\sim_{\{p\}}^V$. Also, in order to avoid clutter, we will often drop the reference to V in \sim_X^V when clear from the context.

Intuitively, two states w and w' are equivalent up to set X , or X -equivalent, if and only if they satisfy the same atoms in X (according to a given valuation V). The finiteness of X is clearly not essential in the definition. It is assumed because, as we will see, each set X will be taken to model a set of actions of some agent in a game form and sets of actions are always assumed to be finite.

Fact 1. (Properties of \sim_X) The following holds for any set of states W , valuation $V : \mathbf{P} \rightarrow 2^W$ and finite sets $X, Y \subseteq \mathbf{P}$:

- (i) \sim_X is reflexive, transitive and symmetric;
- (ii) if $X \subseteq Y$ then $\sim_Y \subseteq \sim_X$;
- (iii) if X is a singleton, \sim_X induces a partition of W with at most 2 cells;
- (iv) $\sim_X \cap \sim_Y = \sim_{X \cup Y}$;
- (v) $\sim_\emptyset = W^2$.

2.2 A Modal Logic of the \sim_X Relation

In this section we consider a simple modal language interpreted on relations \sim_X and axiomatize its logic on the class of structures (W, V) . The key modal operator of the language will be $\langle X \rangle$, whose intuitive meaning is ‘ φ is the case in some state which is X -equivalent to the current one’ or, to stress a ceteris paribus reading, ‘ φ is possible *all things expressed in X being equal*’. We call the resulting logic *propositional ceteris paribus logic*, CP in short.

Let \mathbf{P} be a countable set of atomic propositions. The language $\mathcal{L}_{\text{CP}}(\mathbf{P})$ is defined by the following BNF:

$$\mathcal{L}_{\text{CP}}(\mathbf{P}) : \varphi ::= p \mid \neg\varphi \mid (\varphi \wedge \psi) \mid \langle X \rangle\varphi$$

where p ranges over \mathbf{P} and X is a finite subset of atomic propositions ($X \subseteq \mathbf{P}$ and X finite). Note that as the set of finite subsets of atomic propositions is countable, the language $\mathcal{L}_{\text{CP}}(\mathbf{P})$ is also countable. The Boolean connectives $\top, \vee, \rightarrow, \leftrightarrow$ and the dual operators $[X]$ are defined as usual. Although we have taken diamond operators as primitive, we will for convenience also make use of box operators to state some results in later sections. The set $SF(\varphi)$ of subformulas of a formula φ is defined inductively in the usual way.

This is the class of models we will be working with:

Definition 2. (CP-models) Given a countable set \mathbf{P} , a CP-model for $\mathcal{L}_{\text{CP}}(\mathbf{P})$ is a tuple $\mathcal{M} = (W, V)$ where: W is a non-empty set of states; $V : \mathbf{P} \rightarrow 2^W$ is a valuation function. A CP-model is called universal if $W = 2^{\mathbf{P}}$ and V is s.t. $V(p) = \{w \mid p \in w\}$. It is called non-redundant if $\sim_{\mathbf{P}}$ is the identity relation in W^2 .

Intuitively, a CP-model consists just of a state-space and a valuation function for a given set of atoms.

Definition 3. (Satisfaction for CP-models) Let $\mathcal{M} = (W, V)$ be an CP-model for $\mathcal{L}_{\text{CP}}(\mathbf{P})$, $w \in W$ and $\varphi, \psi \in \mathcal{L}_{\text{CP}}(\mathbf{P})$ (Boolean clauses omitted):

$$\mathcal{M}, w \models_{\text{CP}} \langle X \rangle\varphi \iff \exists w' : w \sim_X^V w' \text{ AND } \mathcal{M}, w' \models_{\text{CP}} \varphi$$

Formula φ is CP-satisfiable, if and only if there exists a model \mathcal{M} and a state w such that $\mathcal{M}, w \models_{\text{CP}} \varphi$. Formula φ is valid in \mathcal{M} , noted $\mathcal{M} \models_{\text{CP}} \varphi$, if and only if for all $w \in W$, $\mathcal{M}, w \models_{\text{CP}} \varphi$. Validity and consequence are defined as usual.

So, modal operators are interpreted on the equivalence relations \sim_X induced by the valuation of the model. It is worth observing that the logic of this class of models is not closed under uniform substitution,² that is, logic CP is not uniform.³ To witness that, notice that formula $\{\{p\}\}p \vee \{\{p\}\}\neg p$ is valid, whereas $\{\{p\}\}\varphi \vee \{\{p\}\}\neg\varphi$ is not.

²For the definition of uniform substitution the reader is referred to the textbook by [Blackburn *et al.*, 2001, Def. 1.18].

³The terminology comes from Goldblatt’s work 1992.

Lemma 1. Let $\mathcal{L}_{CP}^0(\mathbf{P})$ the set of formulae $\varphi \in \mathcal{L}_{CP}(\mathbf{P})$ containing only $\langle \emptyset \rangle$ operators. The set of formulae of $\mathcal{L}_{CP}^0(\mathbf{P})$ which are CP-valid is the modal logic of Kripke frames (W, W^2) , i.e., logic S5.

In other words, the $\langle \emptyset \rangle$ operator of \mathcal{L}_{CP} is nothing but the global modality [Blackburn *et al.*, 2001, pp. 367–370]. The next lemma states that CP is actually the logic of the class of *relevant* CP-models.

Lemma 2. Every satisfiable CP-formula is satisfiable on a non-redundant model.

One can obtain a sound and strongly complete axiom system for CP by a standard reduction technique exploiting Lemma 1. It can also be shown that CP can be embedded in S5 by an exponential satisfiability-preserving translation, and therefore that the satisfiability problem for CP is decidable and in NEXPTIME. We refer to the full paper for the details.

3 The Ceteris Paribus Structure of STIT

3.1 Atemporal Group STIT

First let us recall the syntax and the semantics of atemporal group STIT. The language of this logic is built from a countable set of atomic propositions \mathbf{P} and a finite set of agents $AGT = \{1, \dots, n\}$ and is defined by the following BNF:

$$\mathcal{L}_{G-STIT}(\mathbf{P}, AGT) : \varphi ::= p \mid \neg\varphi \mid (\varphi \wedge \varphi) \mid [J : stit]\varphi$$

where p ranges over \mathbf{P} and J ranges over 2^{AGT} . The construction $[J : stit]\varphi$ is read “group J sees to it that φ is true regardless of what the other agents choose”. We define the dual operator $\langle J : stit \rangle \varphi \stackrel{\text{def}}{=} \neg[J : stit]\neg\varphi$. When $J = \emptyset$, the construction $[\emptyset : stit]\varphi$ is read “ φ is true regardless of what every agent chooses” or simply “ φ is necessarily true”.

Definition 4 (STIT-Kripke model, [Herzig and Schwarzenruber, 2008]). A STIT-Kripke model $\mathcal{M} = (W, \{R_J\}_{J \subseteq AGT}, V)$ is a 3-tuple where:

- W is a non-empty set of worlds;
- for all $J \subseteq AGT$, R_J is an equivalence relation such that: (i) $R_J \subseteq R_\emptyset$; (ii) $R_J = \bigcap_{j \in J} R_{\{j\}}$; (iii) for all $w \in W$ and $(w_1, \dots, w_n) \in W^n$, if $u_1 \in R_{\{1\}}(w), \dots, u_n \in R_{\{n\}}(w)$ then $\bigcap_{1 \leq j \leq n} R_{\{j\}}(u_j) \neq \emptyset$;
- $V : \mathbf{P} \rightarrow 2^W$ is a valuation function for atomic propositions;

with $R_J(w) = \{u \in W : (w, u) \in R_J\}$ for any $J \in 2^{AGT}$.

The partition induced by the equivalence relation R_J is the set of possible choices of the group J . Indeed, in STIT a choice of a group J at a given world w is identified with the set of possible worlds $R_J(w)$. We call $R_J(w)$ the set of possible outcomes of group J 's choice at world w , in the sense that group J 's current choice at w forces the possible worlds to be in $R_J(w)$. The set $R_\emptyset(w)$ is simply the set of possible outcomes at w , or said differently, the set of outcomes of the current game at w . According to Condition (i), the

set of possible outcomes of a group J 's choice is a subset of the set of possible outcomes. Condition (ii), called *additivity*, means that the choices of the agents in a group J is made up of the choices of each individual agent and no more. Condition (iii) corresponds to the property of *independence of agents*: whatever each agent decides to do, the set of outcomes corresponding to the joint action of all agents is non-empty. More intuitively, this means that agents can never be deprived of choices due to the choices made by other agents. In Lorini and Schwarzenruber's work 2011 determinism for the group AGT was assumed. That is to say that the set of outcomes corresponding to a joint action of all agents is a singleton. Horty's group STIT logic [Horty, 2001] does not suppose this. Here we deal with Horty's version of STIT. So a STIT model is a game form in which a joint action of all agents might determine more than one outcome.

3.2 Embedding Atemporal STIT into CP

That STIT is not embeddable into CP can be inferred from various reasons, for instance: STIT satisfiability problem is undecidable if there are more than 3 agents [Herzig and Schwarzenruber, 2008], unlike CP; STIT does not have the finite model property [Herzig and Schwarzenruber, 2008], unlike CP, which we have seen is embeddable S5.

We can nevertheless embed a variant of group STIT under the assumption that every agent has a finite and bounded number of actions in his repertoire. The embedding uses a special formula $GRID_m$ which enforces a CP model to encode a game form, and a suitably defined translation tr_1 .

Theorem 1. Let us consider a group STIT formula φ . Let m be an integer. Then the following items are equivalent: 1. φ is STIT-satisfiable in a STIT-model where each agent has at most 2^m actions; 2. φ is STIT-satisfiable in a STIT-model where each agent has exactly 2^m actions; 3. $GRID_m \wedge tr_1(\varphi)$ is CP-satisfiable.

3.3 Atemporal Individual STIT

In this subsection, we consider the following fragment of STIT called atemporal individual STIT:

$$\mathcal{L}_{I-STIT}(\mathbf{P}, AGT) : \varphi ::= p \mid \neg\varphi \mid (\varphi \wedge \varphi) \mid [\{j\} : stit]\varphi$$

where p ranges over \mathbf{P} and j ranges over AGT .

This fragment of STIT, axiomatized by [Xu, 1998], has the exponential finite model property (see Lemma 7 in [Balbiani *et al.*, 2008b]). Moreover, as the following theorem highlights, it can be embedded in logic CP, through the above translation.

Theorem 2. Let us consider a STIT formula φ of the individual STIT fragment. Let m be the length of φ . Then the following three items are equivalent: 1. φ is STIT-satisfiable 2. φ is STIT-satisfiable in a model where each agent has at most 2^m actions; 3. $GRID_m \wedge tr_1(\varphi)$ is CP-satisfiable.

Thanks to Theorem 2, we reduce the NEXPTIME-complete satisfiability problem of individual STIT [Balbiani *et al.*, 2008b] to the CP-satisfiability problem. As the reduction is polynomial, we obtain the following lower bound complexity result for the CP-satisfiability problem. We thus obtain also a lower bound result of NEXPTIME-hardness for the complexity of the satisfiability problem of CP.

4 The *Ceteris Paribus* Structure of Coalition Logic of Propositional Control

4.1 Logic CL–PC

CL–PC was introduced by [van der Hoek and Wooldridge, 2005] as a formal language for reasoning about capabilities of agents and coalitions in multiagent environments. In this logic the notion of capability is modeled by means of the concept of *control*. In particular, it is assumed that each agent i is associated with a specific finite subset \mathbf{P}_i of the finite set of all propositions \mathbf{P} . \mathbf{P}_i is the set of propositions *controlled* by the agent i . That is, the agent i has the ability to assign a (truth) value to each proposition \mathbf{P}_i but cannot affect the truth values of the propositions in $\mathbf{P} \setminus \mathbf{P}_i$. In the variant of CL–PC studied by [van der Hoek and Wooldridge, 2005] it is also assumed that control over propositions is exclusive, that is, two agents cannot control the same proposition. Moreover, it is assumed that control over propositions is complete, that is, every proposition is controlled by at least one agent.

The *language of CL–PC* is built from a *finite* set of atomic propositions \mathbf{P} and a finite set of agents $AGT = \{1, \dots, n\}$, and is defined by the following BNF:

$$\mathcal{L}_{\text{CL-PC}}(\mathbf{P}, AGT) : \varphi ::= p \mid \neg\varphi \mid (\varphi \wedge \varphi) \mid \diamond_J \varphi$$

where p ranges over \mathbf{P} and J ranges over 2^{AGT} . Operator \diamond_J is called *cooperation modality*, and the construction $\diamond_J \varphi$ means that “group J has the contingent ability to achieve φ ”.

Definition 5 (CL–PC model). *A model for CL–PC is a tuple $\mathcal{M} = (\mathbf{P}_1, \dots, \mathbf{P}_n, X)$ where: (i) $\mathbf{P}_1, \dots, \mathbf{P}_n$ is a partition of \mathbf{P} among the agents in AGT ; (ii) $X \subseteq \mathbf{P}$ is the set of propositions which are true in the initial state.*

For every group of agents $J \subseteq AGT$, let $\mathbf{P}_J = \bigcup_{i \in J} \mathbf{P}_i$ be the set of atomic propositions controlled by the group J . Moreover, for every group $J \subseteq AGT$ and for every set of atomic propositions $X \subseteq \mathbf{P}$, let $X_J = X \cap \mathbf{P}_J$ be the set of atomic propositions in X controlled by the group J . Sets X_J are called J -valuations.

Given a CL–PC model $\mathcal{M} = (\mathbf{P}_1, \dots, \mathbf{P}_n, X)$, the truth condition for the CL–PC modal operators is:

$$\mathcal{M} \models_{\text{CL-PC}} \diamond_J \varphi \iff \exists X'_J \subseteq \mathbf{P}_J : \mathcal{M} \bigoplus X'_J \models_{\text{CL-PC}} \varphi$$

where $\mathcal{M} \bigoplus X'_J$ is the CL–PC model $(\mathbf{P}_1, \dots, \mathbf{P}_n, X'')$ such that: $X''_{AGT \setminus J} = X_{AGT \setminus J}$ and $X''_J = X'_J$. That is, $\diamond_J \varphi$ is true at a given model \mathcal{M} if and only if, the coalition J can change the truth values of the atoms that it controls in such a way that φ will be true afterwards (i.e., given the actual truth-value combination of the atoms which are not controlled by J , there exists a truth-value combination of the atoms controlled by J which ensures φ).

4.2 Embedding CL–PC into STIT and CP

The aim of this section is to provide an embedding of CL–PC into the variant of atemporal group STIT with bounded choices (atemporal ‘bounded’ group STIT) that have been presented in Section 3.1. This is achieved by making use of a suitably defined translation tr_2 and dedicated formulae enforcing a CL–PC-like structure on STIT models: formulae

EXC^+ and EXC^- enforce that control over atomic propositions in \mathbf{P} is exclusive, formula $COMPL$ enforces that exercise of control over atomic propositions in \mathbf{P} is complete; formula $GRID^*$ enforces that all the possible truth-value combinations of the atomic propositions in \mathbf{P} are possible.

Theorem 3. *Let $m = |\mathbf{P}|$. Then, a CL–PC formula φ is CL–PC-satisfiable if and only if $(EXC^+ \wedge EXC^- \wedge COMPL \wedge GRID^*) \wedge tr_2(\varphi)$ is satisfiable in a STIT model where each agent has at most 2^m actions.*

As CP embeds atemporal ‘bounded’ group STIT (Theorem 1 in Section 3.1), from Theorem 3 it follows that CP also embeds CL–PC.

5 The *Ceteris Paribus* Structure of Dynamic Logic of Propositional Assignments

5.1 Logic DL–PA

The dynamic logic of propositional assignments (DL–PA) is the *concrete* variant of propositional dynamic logic (PDL) [Harel *et al.*, 2000] in which atomic programs are assignments of propositional variables to true or to false. The complexities of the model checking and of the satisfiability problem for DL–PA have been recently studied by [Balbiani *et al.*, 2013]. The *starless* version of DL–PA was previously studied by [van Eijck, 2000] and recently put to use by [Herzig *et al.*, 2011], who have shown that it embeds CL–PC.

The *language of DL–PA* is built from a *finite* set of atomic propositions \mathbf{P} and is defined by the following BNF:

$$\begin{aligned} \pi & ::= +p \mid -p \mid \pi; \pi \mid \pi \cup \pi \mid \pi^* \mid \varphi? \\ \mathcal{L}_{\text{DL-PA}}(\mathbf{P}) : \varphi & ::= p \mid \neg\varphi \mid (\varphi \wedge \varphi) \mid \langle \pi \rangle \varphi \end{aligned}$$

We will use $\pm p$ to denote $(+p \cup -p)$. A DL–PA-model is a set $X \subseteq \mathbf{P}$, that is, a propositional valuation. We provide the semantics for the key operators, the rest is standard: $\llbracket +p \rrbracket = \{(X, X') \mid X' = X \cup \{p\}\}$; $\llbracket -p \rrbracket = \{(X, X') \mid X' = X \setminus \{p\}\}$; $\llbracket \langle \pi \rangle \varphi \rrbracket = \{X \mid \text{there exists } X' \text{ s.th. } (X, X') \in \llbracket \pi \rrbracket \text{ and } X' \in \llbracket \varphi \rrbracket\}$.

We write $X \models_{\text{DL-PA}} \varphi$ for $X \in \llbracket \varphi \rrbracket$. We will refer to the fragment of DL–PA without $*$ operator as *starless* DL–PA.

5.2 Embeddings between Starless DL–PA and CP

A DL–PA model X can be translated directly to the pointed CP model $\mathcal{M}_{\text{CP}}, X$. The desired embedding is obtained through a suitable translation tr_3

Theorem 4. *Let X be a DL–PA model and φ belong to the language of starless DL–PA:*

$$X \models_{\text{DL-PA}} \varphi \iff \mathcal{M}_{\text{CP}}, X \models_{\text{CP}} tr_3(\varphi)$$

An embedding in the opposite direction can also be obtained, through a suitably defined translation tr_4 :

Theorem 5. *Let \mathcal{M} be a CP-model and $\varphi \in \mathcal{L}_{\text{CP}}(\mathbf{P})$:*

$$\mathcal{M}, w \models_{\text{CP}} \varphi \iff w \models_{\text{DL-PA}} tr_4(\varphi)$$

From Theorem 5 we can finally obtain as corollary a satisfiability-preserving embedding of CP into DL–PA exploiting a formula $GRID^{**}$ forcing a CP-model to contain all propositional valuations from \mathbf{P} . Then, $tr_4(\varphi)$ is DL–PA satisfiable if and only if $GRID^{**} \wedge \varphi$ is CP satisfiable.

References

- [Alur *et al.*, 2002] R. Alur, T. Henzinger, and O. Kupferman. Alternating-time temporal logic. *Journal of the ACM*, 49:672–713, 2002.
- [Balbiani *et al.*, 2008a] P. Balbiani, O. Gasquet, A. Herzig, F. Schwarzentruber, and N. Troquard. Coalition games over kripke semantics. In C. Dégremon, L. Keiff, and H. Rückert, editors, *Festschrift in Honour of Shahid Rahman*, pages 1–12. College Publications, 2008.
- [Balbiani *et al.*, 2008b] P. Balbiani, A. Herzig, and N. Troquard. Alternative axiomatics and complexity of deliberative stit theories. *Journal of Philosophical Logic*, 37(4):387–406, 2008.
- [Balbiani *et al.*, 2013] P. Balbiani, A. Herzig, and N. Troquard. Dynamic logic of propositional assignments: A well-behaved variant of PDL. In *Proceedings of the 28th ACM/IEEE Symposium on Logic in Computer Science (LICS 2013)*, pages 143–152. IEEE Computer Society, 2013.
- [Belnap *et al.*, 2001] N.D. Belnap, M. Perloff, and M. Xu. *Facing the future: agents and choices in our indeterminist world*. Oxford University Press, USA, 2001.
- [Blackburn *et al.*, 2001] P. Blackburn, M. de Rijke, and Y. Venema. *Modal Logic*. Cambridge University Press, Cambridge, 2001.
- [Broersen *et al.*, 2005] J. Broersen, A. Herzig, and N. Troquard. From coalition logic to STIT. In A. Lomuscio, E. de Vink, and M. Wooldridge, editors, *Proceedings of the Third International Workshop on Logic and Communication in Multi-agent Systems (LCMAS’05)*, pages 23–35, 2005.
- [Broersen *et al.*, 2006] J. Broersen, A. Herzig, and N. Troquard. Embedding alternating-time temporal logic in strategic STIT logic of agency. *Journal of Logic and Computation*, 16(5):559–578, 2006.
- [Broersen *et al.*, 2007] J. Broersen, A. Herzig, and N. Troquard. A normal simulation of coalition logic and an epistemic extension. In D. Samet, editor, *Proceedings TARK’07*, pages 92–101. ACM Press, 2007.
- [Goldblatt, 1992] R. Goldblatt. *Logics of Time and Computation*. CSLI, 1992.
- [Goranko and Jamroga, 2004] V. Goranko and W. Jamroga. Comparing semantics of logics for multi-agent systems. *Synthese*, 139:241–280, 2004.
- [Goranko *et al.*, 2013] V. Goranko, W. Jamroga, and P. Turcini. Strategic games and truly playable effectivity functions. *Journal of Autonomous Agents and Multi-Agent Systems*, (26):288–314, 2013.
- [Goranko, 2001] V. Goranko. Coalition games and alternating temporal logics. In *Proceedings of the 8th conference on theoretical aspects of rationality and knowledge (TARK’01)*, pages 259–272, 2001.
- [Harel *et al.*, 2000] D. Harel, D. Kozen, and J. Tiuryn. *Dynamic Logic*. MIT Press, 2000.
- [Herzig and Schwarzentruber, 2008] A. Herzig and F. Schwarzentruber. Properties of logics of individual and group agency. *Advances in modal logic*, 7:133–149, 2008.
- [Herzig *et al.*, 2011] A. Herzig, E. Lorini, F. Moisan, and N. Troquard. A dynamic logic of normative systems. In T. Walsh, editor, *Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence (IJCAI 2011)*, pages 228–233. AAAI Press, 2011.
- [Herzig, 2014] A. Herzig. Logics of knowledge and action: A critical analysis and challenges. *Journal of Autonomous Agents and Multi-Agent Systems*, DOI: 10.1007/s10458-014-9267-z, 2014.
- [Horty, 2001] J. F. Horty. *Agency and Deontic Logic*. Oxford University Press, Oxford, 2001.
- [Lorini and Schwarzentruber, 2011] E. Lorini and F. Schwarzentruber. A logic for reasoning about counterfactual emotions. *Artificial Intelligence*, 175(3-4):814–847, 2011.
- [Moulin and Peleg, 1982] H. Moulin and B. Peleg. Cores of effectivity functions and implementation theory. *Journal of Mathematical Economics*, 10:115–145, 1982.
- [Pauly, 2002] M. Pauly. A modal logic for coalitional power in games. *Journal of Logic and Computation*, 12(1):149–166, 2002.
- [van Benthem *et al.*, 2009] J. van Benthem, P. Girard, and O. Roy. Everything else being equal: A modal logic for ceteris paribus preferences. *Journal of Philosophical Logic*, 38:83–125, 2009.
- [van der Hoek and Pauly, 2007] W. van der Hoek and M. Pauly. Modal logic for games and information. In P. Blackburn, J. van Benthem, and F. Wolter, editors, *Handbook of Modal Logic*, pages 1077–1146. Elsevier, 2007.
- [van der Hoek and Wooldridge, 2005] W. van der Hoek and M. Wooldridge. On the logic of cooperation and propositional control. *Artificial Intelligence*, 164:81–119, 2005.
- [van Eijck, 2000] J. van Eijck. Making things happen. *Studia logica*, 66(1):41–58, 2000.
- [Xu, 1998] M. Xu. Axioms for deliberative STIT. *Journal of Philosophical Logic*, 27:505–552, 1998.