

Expressivity Results for Deontic Logics of Collective Agency

Allard Tamminga* Hein Duijf† Frederik Van De Putte‡

Abstract

We use a deontic logic of collective agency to study reducibility questions about collective agency and collective obligations. The logic that is at the basis of our study is a multi-modal logic in the tradition of *stit* (‘sees to it that’) logics of agency. Our full formal language has constants for collective and individual deontic admissibility, modalities for collective and individual agency, and modalities for collective and individual obligations. We classify its twenty-seven sublanguages in terms of their expressive power. This classification enables us to investigate reducibility relations between collective deontic admissibility, collective agency, and collective obligations, on the one hand, and individual deontic admissibility, individual agency, and individual obligations, on the other.

Keywords: deontic logic; collective admissibility; collective agency; collective obligation; expressivity; bisimulation.

1 INTRODUCTION

In the philosophy of the social sciences, individualism is the methodological precept that any social phenomenon is ultimately to be explained in terms of the actions and interactions of individuals. One of the central questions in the debate on individualism is whether statements about collectives can be reduced to statements about individuals. Does it hold that for every statement about collectives there is a logically equivalent statement about individuals? In this paper, we use

*University of Groningen, Oude Boteringestraat 52, 9712 GL Groningen, The Netherlands.

†Vrije Universiteit Amsterdam, De Boelelaan 1105, 1081 HV Amsterdam, The Netherlands.

‡Ghent University, Blandijnberg 2, 9000 Gent, Belgium and University of Bayreuth, Universitätsstraße 30, 95447 Bayreuth, Germany.

techniques from modal logic to find answers to specific reducibility questions, including: Does it hold that for every statement about collective agency there is a logically equivalent statement about individual agency? Does it hold that for every statement about collective obligations there is a logically equivalent statement about collective agency and individual obligations?

By way of a formal study of expressivity relations, we aim to assess which sorts of reductionism about collective agency and collective obligations are tenable. Accordingly, we rephrase the reducibility questions as questions of expressivity: Can statements from a specific language that includes specific statements about collectives be *expressed* with statements from another language? We follow standard practice in modal logic and develop new notions of bisimulation to determine the answers to these questions of expressivity. Doing so enables us to answer our reducibility questions with logical precision.

To answer our expressivity questions, we use a deontic logic of collective agency. This multi-modal logic is a simplified and generalized version of Horty's (1996; 2001) deontic logic of agency.¹ Alongside the usual sentential connectives and the universal modality, our full formal language has constants for collective and individual deontic admissibility, modalities for collective and individual agency, and modalities for collective and individual obligations.² Truth-conditions for the formulas of the full formal language are specified in terms of *deontic game models* (Tamminga and Hindriks 2020, §3.2).

There are twenty-seven sublanguages of the full language, depending on (a) whether constants for collective and/or individual deontic admissibility are excluded, (b) whether modalities for collective and/or individual agency are excluded, and (c) whether modalities for collective and/or individual obligations are excluded. Because each of these three items gives rise to three possibilities, there are exactly twenty-seven sublanguages. The main technical contribution of

¹The deontic logic of collective agency is a logic in the tradition of *stit* ('sees to it that') logics of agency. Seminal works include Kanger (1957), Pörn (1970), von Kutschera (1986), and Horty and Belnap (1995). See Belnap, Perloff, and Xu (2001) for a textbook presentation of *stit* logic. There are other formal accounts of agency and normativity. Rather than using *stit* logic, Meyer (1988) and Segerberg (2012) use *dynamic* logic to analyse individual actions and obligations. Czelakowski (2020) uses *formal linguistics* as a basis for a deontic logic of consecutive individual actions.

²On admissibility, see Arrow (1951, p. 429), Luce and Raiffa (1957, pp. 287 and 307), Savage (1972, p. 21), and Kohlberg and Mertens (1986, §2.7.A). Horty (2001, p. 130) relies on collective deontic admissibility to define group obligations, and Tamminga and Duijf use collective deontic admissibility to analyse collective rationality (2017, pp. 200–201) and backward-looking collective moral responsibility (2017, §6).

this paper consists in charting this new territory and ordering these twenty-seven sublanguages in terms of their expressive power.

Our paper is organized as follows. In Section 2, we first define the full language and give truth-conditions for its formulas in terms of deontic game models. The twenty-seven sublanguages are then defined. Next, we define the concept of expressivity and give three criteria on expressivity that help to classify the twenty-seven sublanguages in terms of their expressive power (Section 3). We chart the sublanguages that are equally expressive in Section 4 and then give the full picture of the expressivity relations in Section 5. Subsequently, we introduce a class of bisimulation relations between pointed deontic game models and prove a Hennessy-Milner theorem for each of the twenty-seven sublanguages (Section 6). In the following three sections, we prove that the full picture of expressivity relations is accurate. We briefly discuss two assumptions on our models and languages in Section 10. A short summary and some philosophical considerations conclude the paper.

2 LANGUAGES AND SEMANTICS

Let \mathfrak{P} be a fixed countable set of atomic propositions and let \mathcal{N} be a fixed finite set of individual agents that contains *at least three elements*. (We return to this assumption in Section 10.) We use p and q as variables for atomic propositions, i , j , and k as variables for individual agents, and \mathcal{F} and \mathcal{G} as variables for non-empty sets of individual agents. We use $-\mathcal{G}$ to refer to the complement $\mathcal{N} - \mathcal{G}$. The *full language* $\mathcal{L}_{\star_{\mathcal{G}}, [\mathcal{G}], (\mathcal{G})}$ of our deontic logic of collective agency contains (i) *deontic admissibility constants* of the form “Group \mathcal{G} of agents performs a deontically admissible group action” (formalized as $\star_{\mathcal{G}}$) and “Agent i performs a deontically admissible individual action” (formalized as \star_i). Moreover, it contains (ii) *necessity statements* of the form “It is settled true that ϕ ” (formalized as $\Box\phi$), (iii) *agentive statements* of the form “Group \mathcal{G} of agents sees to it that ϕ ” (formalized as $[\mathcal{G}]\phi$) and “Agent i sees to it that ϕ ” (formalized as $[i]\phi$), and (iv) *deontic statements* of the form “Group \mathcal{G} of agents ought to see to it that ϕ ” (formalized as $(\mathcal{G})\phi$) and “Agent i ought to see to it that ϕ ” (formalized as $(i)\phi$). Its Backus-Naur form is the following:

$$\phi := p \mid \star_{\mathcal{G}} \mid \neg\phi \mid (\phi \wedge \phi) \mid \Box\phi \mid [\mathcal{G}]\phi \mid (\mathcal{G})\phi$$

where p ranges over \mathfrak{P} and \mathcal{G} ranges over non-empty subsets of \mathcal{N} .

The operators \rightarrow , \leftrightarrow , \diamond , and $\langle \mathcal{G} \rangle$ abbreviate the usual constructions. We leave out brackets and braces if the omission does not give rise to ambiguities. Accord-

ingly, the formulas \star_i , $[i]\phi$, and $(i)\phi$ are shorthand for the formulas $\star_{\{i\}}$, $[\{i\}]\phi$, and $(\{i\})\phi$, respectively.

2.1 DEONTIC GAME MODELS

Truth-conditions for the formulas of $\mathfrak{L}_{\star_{\mathcal{G}},[\mathcal{G}],(\mathcal{G})}$ are specified in terms of deontic game models. A deontic game model is a one-shot game in which each individual agent i is assigned a non-empty and finite set A_i of available individual actions. The set $A_{\mathcal{G}}$ of group actions that are available to a non-empty set \mathcal{G} of individual agents is given by $A_{\mathcal{G}} = \times_{i \in \mathcal{G}} A_i$. The set A of action profiles equals the set $A_{\mathcal{N}}$ of group actions that are available to the grand coalition \mathcal{N} of all individual agents and is given by $A = A_{\mathcal{N}} = \times_{i \in \mathcal{N}} A_i$. The set of action profiles plays the role of the set of worlds in possible-worlds models for standard deontic logic. A deontic ideality function d assigns to each action profile a in A a value $d(a)$ that is either 0 or 1. A valuation function assigns to each atomic proposition p in \mathfrak{P} a set $v(p)$ of action profiles where p is true.³ Formally,

Definition 1 (Deontic Game Model). *A deontic game model M is a quadruple $\langle \mathcal{N}, (A_i), d, v \rangle$ such that for each agent i in \mathcal{N} it holds that A_i is a non-empty and finite set of actions available to agent i , $d : A \rightarrow \{0, 1\}$ is a deontic ideality function such that there is at least one a in A with $d(a) = 1$, and $v : \mathfrak{P} \rightarrow \wp(A)$ is a valuation function.*

To specify truth-conditions for deontic admissibility constants and for deontic statements, we order the group actions that are available to any (possibly singleton) set \mathcal{G} of individual agents by way of a dominance relation:

Definition 2 (Weak Dominance). *Let $M = \langle \mathcal{N}, (A_i), d, v \rangle$ be a deontic game model. Let $\mathcal{G} \subseteq \mathcal{N}$ be a non-empty set of individual agents. Let $a_{\mathcal{G}}, b_{\mathcal{G}} \in A_{\mathcal{G}}$. Then*

$$a_{\mathcal{G}} \succeq_M b_{\mathcal{G}} \quad \text{iff} \quad \text{for all } c_{-\mathcal{G}} \in A_{-\mathcal{G}} \text{ it holds that } d(a_{\mathcal{G}}, c_{-\mathcal{G}}) \geq d(b_{\mathcal{G}}, c_{-\mathcal{G}}).$$

Strong dominance is defined in terms of weak dominance: $a_{\mathcal{G}} \succ_M b_{\mathcal{G}}$ if and only if $a_{\mathcal{G}} \succeq_M b_{\mathcal{G}}$ and $b_{\mathcal{G}} \not\prec_M a_{\mathcal{G}}$.

A group action that is available to any (possibly singleton) set \mathcal{G} of individual agents is *deontically admissible* if and only if it is *not* strongly dominated by any of \mathcal{G} 's available group actions. Accordingly,

³We adopt the notational conventions of Osborne and Rubinstein (1994, §1.7).

Definition 3 (Deontic Admissibility). *Let $M = \langle \mathcal{N}, (A_i), d, v \rangle$ be a deontic game model. Let $\mathcal{G} \subseteq \mathcal{N}$ be a non-empty set of individual agents. Then the set of \mathcal{G} 's deontically admissible actions in M , denoted by $\text{Adm}_M(\mathcal{G})$, is given by*

$$\text{Adm}_M(\mathcal{G}) = \{a_{\mathcal{G}} \in A_{\mathcal{G}} : \text{there is no } b_{\mathcal{G}} \in A_{\mathcal{G}} \text{ such that } b_{\mathcal{G}} \succ_M a_{\mathcal{G}}\}.$$

We can now specify the truth-conditions for the formulas of the full language $\mathfrak{L}_{\star_{\mathcal{G}}, [\mathcal{G}], (\mathcal{G})}$. (Note that by setting $\mathcal{G} = \{i\}$ we obtain the truth-conditions for the individualistic formulas \star_i , $[i]\phi$, and $(i)\phi$.)

Definition 4 (Truth-Conditions). *Let $M = \langle \mathcal{N}, (A_i), d, v \rangle$ be a deontic game model. Let $\mathcal{G} \subseteq \mathcal{N}$ be a non-empty set of individual agents. Let $a \in A$ be an action profile. Let $p \in \mathfrak{P}$ be an atomic formula and let $\phi, \psi \in \mathfrak{L}_{\star_{\mathcal{G}}, [\mathcal{G}], (\mathcal{G})}$ be arbitrary formulas. Then*

$$\begin{aligned} (M, a) \models p & \quad \text{iff} \quad a \in v(p) \\ (M, a) \models \star_{\mathcal{G}} & \quad \text{iff} \quad a_{\mathcal{G}} \in \text{Adm}_M(\mathcal{G}) \\ (M, a) \models \neg\phi & \quad \text{iff} \quad (M, a) \not\models \phi \\ (M, a) \models \phi \wedge \psi & \quad \text{iff} \quad (M, a) \models \phi \text{ and } (M, a) \models \psi \\ (M, a) \models \Box\phi & \quad \text{iff} \quad (M, b) \models \phi \text{ for all } b \in A \\ (M, a) \models [\mathcal{G}]\phi & \quad \text{iff} \quad (M, b) \models \phi \text{ for all } b \in A \text{ with } b_{\mathcal{G}} = a_{\mathcal{G}} \\ (M, a) \models (\mathcal{G})\phi & \quad \text{iff} \quad (M, b) \models \phi \text{ for all } b \in A \text{ with } b_{\mathcal{G}} \in \text{Adm}_M(\mathcal{G}). \end{aligned}$$

Given a deontic game model M , we write $M \models \phi$ if for all action profiles a in A it holds that $(M, a) \models \phi$. A formula ϕ is *valid* (notation: $\models \phi$) if for all deontic game models M it holds that $M \models \phi$.

2.2 TWENTY-SEVEN SUBLANGUAGES

The full language $\mathfrak{L}_{\star_{\mathcal{G}}, [\mathcal{G}], (\mathcal{G})}$ of the deontic logic of collective agency contains constants for collective and individual *deontic admissibility*, modalities for collective and individual *agency*, and modalities for collective and individual *obligations*. Our basic language is the standard modal language \mathfrak{L} built from the set \mathfrak{P} of atomic propositions, using negation (\neg), conjunction (\wedge), and the universal modality (\Box). (We return to the assumption of including the universal modality in Section 10.) The sublanguages of our full language are obtained from the basic language \mathfrak{L} by adding (a) either collective deontic admissibility constants, individual deontic admissibility constants, or no deontic admissibility constants; (b) collective agency modalities, individual agency modalities, or no agency modalities; and/or (c) collective obligation modalities, individual obligation modalities, or no obligation modalities. This gives us twenty-seven languages, ranging from

\mathfrak{L} to $\mathfrak{L}_{\star_{\mathcal{G}},[\mathcal{G}],(\mathcal{G})}$. Their names are systematized as follows: we use $\mathfrak{L}_{x,y,z}$ with $x \in \{\star_{\mathcal{G}}, \star_i, \epsilon\}$, $y \in \{[\mathcal{G}], [i], \epsilon\}$, and $z \in \{(\mathcal{G}), (i), \epsilon\}$, where ϵ represents the omission of the relevant constant or modality. Their Backus-Naur forms are read off from their names. For instance, if $x = \star_i$, $y = [\mathcal{G}]$, and $z = \epsilon$, we obtain the language $\mathfrak{L}_{\star_i,[\mathcal{G}]}$. Its Backus-Naur form is

$$\phi := p \mid \star_i \mid \neg\phi \mid (\phi \wedge \phi) \mid \Box\phi \mid [\mathcal{G}]\phi$$

where p ranges over \mathfrak{P} , i ranges over \mathcal{N} , and \mathcal{G} ranges over non-empty subsets of \mathcal{N} .

We use \mathbf{L} to refer to the set that consists of exactly all of our twenty-seven sublanguages. Truth-conditions for the formulas of any $\mathfrak{L}^* \in \mathbf{L}$ are specified by Definition 4. We now order the twenty-seven languages in \mathbf{L} in terms of their expressive power.⁴

3 EXPRESSIVITY

Let us first recall the concept of expressive power, as standardly used in modal logic. We say that a formula ψ from language \mathfrak{L}^{**} can be expressed in language \mathfrak{L}^* if and only if there is a ϕ in \mathfrak{L}^* such that ϕ and ψ are logically equivalent. If every formula from the language \mathfrak{L}^{**} can be expressed in the language \mathfrak{L}^* , we say that \mathfrak{L}^* is at least as expressive as \mathfrak{L}^{**} . Formally,

Definition 5. *Let $\mathfrak{L}^*, \mathfrak{L}^{**} \in \mathbf{L}$. Then \mathfrak{L}^* is at least as expressive as \mathfrak{L}^{**} (notation: $\mathfrak{L}^* \succeq \mathfrak{L}^{**}$) if and only if for every ψ in \mathfrak{L}^{**} there is a ϕ in \mathfrak{L}^* such that $\models \psi \leftrightarrow \phi$.*

As usual, $\mathfrak{L}^* \equiv \mathfrak{L}^{**}$ abbreviates the conjunction of $\mathfrak{L}^* \succeq \mathfrak{L}^{**}$ and $\mathfrak{L}^{**} \succeq \mathfrak{L}^*$. Moreover, $\mathfrak{L}^* \succ \mathfrak{L}^{**}$ abbreviates the conjunction of $\mathfrak{L}^* \succeq \mathfrak{L}^{**}$ and $\mathfrak{L}^{**} \not\succeq \mathfrak{L}^*$. Note that $\mathfrak{L}^{**} \not\succeq \mathfrak{L}^*$ if and only if there is a ϕ in \mathfrak{L}^* such that for every ψ in \mathfrak{L}^{**} it holds that $\not\models \psi \leftrightarrow \phi$. Accordingly, there is no formula in \mathfrak{L}^{**} that is logically equivalent to ϕ .

To chart the expressivity landscape of our twenty-seven languages, we introduce three criteria of expressive power. The first two criteria are straightforward. First, consider the obvious *inclusion criterion* for $\mathfrak{L}^* \succeq \mathfrak{L}^{**}$:

⁴Other metalogical investigations of these languages are scarce. Xu (1998), Wansing (2006), and Balbiani, Herzig, and Troquard (2008) prove completeness for $\mathfrak{L}_{[i]}$, and Murakami (2005) proves completeness for $\mathfrak{L}_{[i],(i)}$. Kooi and Tamminga (2008) characterize the possibility of moral conflicts in $\mathfrak{L}_{[\mathcal{G}],(\mathcal{G})}$. Balbiani, Herzig, and Troquard (2008) and Herzig and Schwarzentruher (2008) study the complexity of $\mathfrak{L}_{[i]}$, and the latter study also the complexity of $\mathfrak{L}_{[\mathcal{G}]}$. Duijf, Tamminga and Van De Putte (ms.) define a new and stronger concept of inexpressivity based on *nomological* equivalence and prove that there are no statements ϕ and ψ in $\mathfrak{L}_{\star_i,[i]}$ such that ϕ and $\star_{\mathcal{G}}$ are equivalent modulo ψ (where ψ specifies a set of individualistic background laws).

Lemma 1 (Inclusion Criterion). *Let $\mathfrak{L}^*, \mathfrak{L}^{**} \in \mathbf{L}$. Then*

If $\mathfrak{L}^ \supseteq \mathfrak{L}^{**}$, then $\mathfrak{L}^* \succeq \mathfrak{L}^{**}$.*

Secondly, consider the *transitivity criterion* for $\mathfrak{L}^* \not\preceq \mathfrak{L}^{**}$. It follows from the transitivity of \succeq :

Lemma 2 (Transitivity Criterion). *Let $\mathfrak{L}^*, \mathfrak{L}^{**}, \mathfrak{L}^{***}, \mathfrak{L}^{****} \in \mathbf{L}$. Then*

*If $\mathfrak{L}^{***} \not\preceq \mathfrak{L}^{****}$ and $\mathfrak{L}^{***} \succeq \mathfrak{L}^*$ and $\mathfrak{L}^{**} \succeq \mathfrak{L}^{****}$, then $\mathfrak{L}^* \not\preceq \mathfrak{L}^{**}$.*

The third criterion is a criterion for $\mathfrak{L}^* \not\preceq \mathfrak{L}^{**}$ that is defined in terms of pointed deontic game models. The criterion applies if two pointed deontic game models validate exactly the same formulas from \mathfrak{L}^* but give different truth-values to a formula from \mathfrak{L}^{**} . Any ordered pair (M, a) that consists of a deontic game model M and one of its action profiles a is a *pointed deontic game model*. Two pointed deontic game models are equivalent on \mathfrak{L}^* if and only if they validate exactly the same formulas from \mathfrak{L}^* :

Definition 6 (Equivalence on \mathfrak{L}^*). *Let $\mathfrak{L}^* \in \mathbf{L}$. Let (M, a) and (M', a') be pointed deontic game models. Then (M, a) and (M', a') are equivalent on \mathfrak{L}^* (notation: $(M, a) \equiv_{\mathfrak{L}^*} (M', a')$) iff for all $\phi \in \mathfrak{L}^*$ it holds that $(M, a) \models \phi$ if and only if $(M', a') \models \phi$.*

We can now formulate the *equivalence criterion* for $\mathfrak{L}^* \not\preceq \mathfrak{L}^{**}$:

Lemma 3 (Equivalence Criterion). *Let $\mathfrak{L}^*, \mathfrak{L}^{**} \in \mathbf{L}$. Let (M, a) and (M', a') be pointed deontic game models. Suppose $(M, a) \equiv_{\mathfrak{L}^*} (M', a')$. Then*

*If $(M, a) \models \psi$ and $(M', a') \not\models \psi$ for some $\psi \in \mathfrak{L}^{**}$, then $\mathfrak{L}^* \not\preceq \mathfrak{L}^{**}$.*

Proof. Assume $(M, a) \equiv_{\mathfrak{L}^*} (M', a')$. Suppose $(M, a) \models \psi$ and $(M', a') \not\models \psi$ for some $\psi \in \mathfrak{L}^{**}$. Suppose $\mathfrak{L}^* \succeq \mathfrak{L}^{**}$. Because $\psi \in \mathfrak{L}^{**}$, there is a ϕ in \mathfrak{L}^* such that $\models \psi \leftrightarrow \phi$. Then $(M, a) \models \phi$ and $(M', a') \not\models \phi$ and $\phi \in \mathfrak{L}^*$. Contradiction. Therefore, $\mathfrak{L}^* \not\preceq \mathfrak{L}^{**}$. \square

The *inclusion criterion*, the *transitivity criterion*, and the *equivalence criterion* help to classify the twenty-seven sublanguages of $\mathfrak{L}_{\times \mathcal{G}, [\mathcal{G}], (\mathcal{G})}$ in terms of their expressive power. We start with the sublanguages that have precisely the same expressive power.

4 EQUALLY EXPRESSIVE SUBLANGUAGES

Any obligation modality can be given an Andersonian-Kangerian definition in terms of the universal modality and the relevant deontic admissibility constant:

Lemma 4. *Let $\mathcal{G} \subseteq \mathcal{N}$ be a non-empty set of individual agents. Let $\phi \in \mathfrak{L}_{\star_{\mathcal{G}},[\mathcal{G}],(\mathcal{G})}$. Then*

$$\models (\mathcal{G})\phi \leftrightarrow \Box(\star_{\mathcal{G}} \rightarrow \phi).$$

The Andersonian-Kangerian definitions of obligation modalities are central to the proofs that some of the sublanguages of $\mathfrak{L}_{\star_{\mathcal{G}},[\mathcal{G}],(\mathcal{G})}$ are equally expressive.

Lemma 5. *Let $x \in \{\star_{\mathcal{G}}, \star_i, \epsilon\}$, $y \in \{[\mathcal{G}], [i], \epsilon\}$, and $z \in \{(\mathcal{G}), (i), \epsilon\}$. Then*

- (i) $\mathfrak{L}_{\star_{\mathcal{G}},y,z} \supseteq \mathfrak{L}_{x,y,(\mathcal{G})}$
- (ii) $\mathfrak{L}_{\star_{\mathcal{G}},y,z} \supseteq \mathfrak{L}_{x,y,(i)}$
- (iii) $\mathfrak{L}_{\star_i,y,z} \supseteq \mathfrak{L}_{\star_i,y,(i)}$
- (iv) $\mathfrak{L}_{\star_i,y,z} \supseteq \mathfrak{L}_{y,(i)}$.

Proof. (i) We show that for every $\psi \in \mathfrak{L}_{x,y,(\mathcal{G})}$ there is a $\phi \in \mathfrak{L}_{\star_{\mathcal{G}},y,z}$ such that $\models \psi \leftrightarrow \phi$ by structural induction on ψ .

Basis: we check cases $\psi = p$, $\psi = \star_i$ (if $x = \star_i$), and $\psi = \star_{\mathcal{G}}$ (if $x = \star_{\mathcal{G}}$). Because all of them are also elements of $\mathfrak{L}_{\star_{\mathcal{G}},y,z}$, there is a $\phi \in \mathfrak{L}_{\star_{\mathcal{G}},y,z}$ such that $\models \psi \leftrightarrow \phi$.

Induction Hypothesis: for all $\chi \in \mathfrak{L}_{x,y,(\mathcal{G})}$ with fewer operators than ψ there is a $\phi_{\chi} \in \mathfrak{L}_{\star_{\mathcal{G}},y,z}$ such that $\models \chi \leftrightarrow \phi_{\chi}$.

Induction Step: we check case $\psi = (\mathcal{G})\chi$. By Lemma 4, $\models (\mathcal{G})\chi \leftrightarrow \Box(\star_{\mathcal{G}} \rightarrow \chi)$. By the Induction Hypothesis, there is a $\phi_{\chi} \in \mathfrak{L}_{\star_{\mathcal{G}},y,z}$ such that $\models \chi \leftrightarrow \phi_{\chi}$. Hence, $\models (\mathcal{G})\chi \leftrightarrow \Box(\star_{\mathcal{G}} \rightarrow \phi_{\chi})$. Note that $\Box(\star_{\mathcal{G}} \rightarrow \phi_{\chi}) \in \mathfrak{L}_{\star_{\mathcal{G}},y,z}$. Hence, there is a $\phi \in \mathfrak{L}_{\star_{\mathcal{G}},y,z}$ such that $\models \psi \leftrightarrow \phi$.

The other cases are similar.

(ii), (iii), and (iv) are proved analogously. \square

Theorem 1. *Each of the following holds:*

- (i) $\mathfrak{L}_{\star_{\mathcal{G}},[\mathcal{G}]} \equiv \mathfrak{L}_{\star_{\mathcal{G}},[\mathcal{G}],(i)} \equiv \mathfrak{L}_{\star_{\mathcal{G}},[\mathcal{G}],(\mathcal{G})}$
- (ii) $\mathfrak{L}_{\star_{\mathcal{G}},[i]} \equiv \mathfrak{L}_{\star_{\mathcal{G}},[i],(i)} \equiv \mathfrak{L}_{\star_{\mathcal{G}},[i],(\mathcal{G})}$
- (iii) $\mathfrak{L}_{\star_{\mathcal{G}}} \equiv \mathfrak{L}_{\star_{\mathcal{G}},(i)} \equiv \mathfrak{L}_{\star_{\mathcal{G}},(\mathcal{G})}$
- (iv) $\mathfrak{L}_{\star_i,[\mathcal{G}]} \equiv \mathfrak{L}_{\star_i,[\mathcal{G}],(i)}$
- (v) $\mathfrak{L}_{\star_i,[i]} \equiv \mathfrak{L}_{\star_i,[i],(i)}$
- (vi) $\mathfrak{L}_{\star_i} \equiv \mathfrak{L}_{\star_i,(i)}$.

Proof. (i) By Lemma 5(ii), we have $\mathfrak{L}_{\star\mathcal{G},[\mathcal{G}]} \succeq \mathfrak{L}_{\star\mathcal{G},[\mathcal{G}],(i)}$. By the inclusion criterion, it holds that $\mathfrak{L}_{\star\mathcal{G},[\mathcal{G}],(i)} \succeq \mathfrak{L}_{\star\mathcal{G},[\mathcal{G}]}$. Hence, $\mathfrak{L}_{\star\mathcal{G},[\mathcal{G}]} \equiv \mathfrak{L}_{\star\mathcal{G},[\mathcal{G}],(i)}$. By Lemma 5(i), we have $\mathfrak{L}_{\star\mathcal{G},[\mathcal{G}]} \succeq \mathfrak{L}_{\star\mathcal{G},[\mathcal{G}],(\mathcal{G})}$. By the inclusion criterion, it holds that $\mathfrak{L}_{\star\mathcal{G},[\mathcal{G}],(\mathcal{G})} \succeq \mathfrak{L}_{\star\mathcal{G},[\mathcal{G}]}$. Hence, $\mathfrak{L}_{\star\mathcal{G},[\mathcal{G}]} \equiv \mathfrak{L}_{\star\mathcal{G},[\mathcal{G}],(\mathcal{G})}$. The transitivity of \succeq gives us $\mathfrak{L}_{\star\mathcal{G},[\mathcal{G}]} \equiv \mathfrak{L}_{\star\mathcal{G},[\mathcal{G}],(i)} \equiv \mathfrak{L}_{\star\mathcal{G},[\mathcal{G}],(\mathcal{G})}$.

Claims (ii) through (vi) are proved analogously. \square

Because our focus is on charting the expressivity relations between our twenty-seven sublanguages, we refer to any class of equally expressive languages by the language in that class with the shortest name. Hence, we use the six languages in the first column in Theorem 1 to refer to the six classes of equally expressive languages that were identified in this theorem. Accordingly, $\mathfrak{L}_{\star\mathcal{G},[\mathcal{G}]}$ refers to the class of the languages in \mathbf{L} that are equally expressive as the full language $\mathfrak{L}_{\star\mathcal{G},[\mathcal{G}],(\mathcal{G})}$.

5 THE FULL PICTURE OF THE EXPRESSIVITY RELATIONS

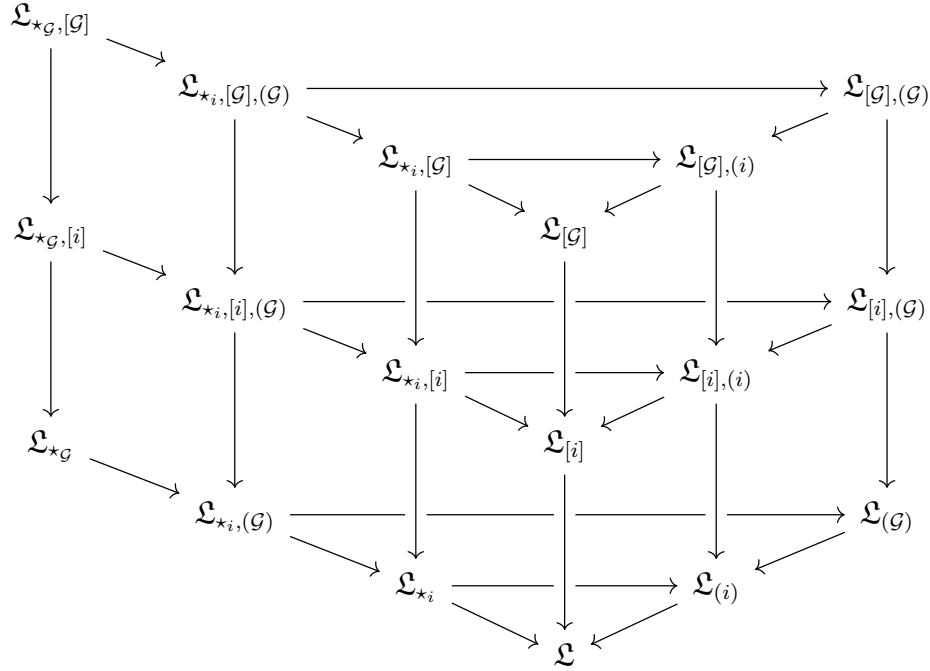


Figure 1: Expressivity Relations between the Eighteen Languages

The eighteen sublanguages that remain to be classified are the six languages of

the first column of Theorem 1 and the sublanguages in \mathbf{L} that were not mentioned in this theorem. No two of these eighteen sublanguages are equally expressive. If we represent $\mathfrak{L}^* \succ \mathfrak{L}^{**}$ by an arrow from \mathfrak{L}^* to \mathfrak{L}^{**} , the expressivity relations between the eighteen languages can be pictured as in Figure 1.

Accordingly, $\mathfrak{L}_{*\mathcal{G},[\mathcal{G}]}$ is the most expressive language in \mathbf{L} and \mathfrak{L} is the least expressive language in \mathbf{L} . To prove that this picture adequately represents the expressivity relations among the eighteen sublanguages, we think of the picture as a three-dimensional prism and decompose it in the *left front face* (Figure 2), the *right front face* (Figure 9), and the *horizontal* (Figure 10).

6 BISIMULATIONS AND HENNESSY-MILNER THEOREMS

To prove that some sublanguage \mathfrak{L}^* is *not* at least as expressive as another sublanguage \mathfrak{L}^{**} , we apply the equivalence criterion to specific pairs of pointed deontic game models. The equivalence criterion can only be applied if the pointed deontic game models in the pair are \mathfrak{L}^* -equivalent. To prove that they are \mathfrak{L}^* -equivalent, we develop a suitable concept of \mathfrak{L}^* -bisimulation and show that the two languages are \mathfrak{L}^* -bisimilar. All of this is standard practice in modal logic.⁵

An \mathfrak{L}^* -bisimulation between two pointed deontic game models (M, a) and (M', a') is a relation $R \subseteq A \times A'$ linking the models' action profiles that satisfies a specific set of structural conditions associated with the logical operators in \mathfrak{L}^* . For every language \mathfrak{L}^* in \mathbf{L} we present such a set of structural conditions on R that characterizes \mathfrak{L}^* -equivalence. Accordingly, for every language \mathfrak{L}^* in \mathbf{L} we show that two given pointed deontic game models are \mathfrak{L}^* -bisimilar if and only if they are \mathfrak{L}^* -equivalent. Such theorems are known as Hennessy-Milner theorems. We prove a general claim that summarizes the Hennessy-Milner theorems for all sublanguages in \mathbf{L} (Theorem 2). To establish this general claim, we use a technique from Van Benthem, Bezhianishvili, Enqvist, and Yu (2017, § 3) and first prove a general lemma on sublanguages in \mathbf{L} . (Note that this technique is also called upon in our proof of the general claim itself.)

Lemma 6. *Let $\mathfrak{L}^* \in \mathbf{L}$. Let $M = \langle \mathcal{N}, (A_i), d, v \rangle$ and $M' = \langle \mathcal{N}', (A'_i), d', v' \rangle$ be deontic game models. Let $b \in A$ and $b' \in A'$. If $(M, b) \equiv_{\mathfrak{L}^*} (M', b')$, then for every $c \in A$ there is a $c' \in A'$ such that $(M, c) \equiv_{\mathfrak{L}^*} (M', c')$.*

Proof. Assume $(M, b) \equiv_{\mathfrak{L}^*} (M', b')$. Take an arbitrary $c \in A$. For every $d \in A$, let $\phi_{c,d} = p \vee \neg p$ if $(M, c) \equiv_{\mathfrak{L}^*} (M, d)$; otherwise, let $\phi_{c,d} = \psi$ for some $\psi \in \mathfrak{L}^*$

⁵Blackburn, De Rijke, and Venema (2001, § 2.2) provide a textbook presentation of bisimulation in modal logic.

for which it holds that $(M, c) \models \psi$ and $(M, d) \not\models \psi$. Let $\phi_c = \bigwedge_{d \in A} \phi_{c,d}$. The finiteness of A ensures that ϕ_c is well defined. Note that (\dagger) for every $d \in A$ it holds that if $(M, d) \models \phi_c$, then $(M, c) \equiv_{\mathfrak{L}^*} (M, d)$.

Because $(M, c) \models \phi_c$, it holds that $(M, b) \models \diamond \phi_c$. By our assumption, $(M', b') \models \diamond \phi_c$. Then there is a $c' \in A'$ such that $(M', c') \models \phi_c$. Suppose $(M, c) \not\equiv_{\mathfrak{L}^*} (M', c')$. Then there is a $\chi \in \mathfrak{L}^*$ such that $(M, c) \models \chi$ and $(M', c') \not\models \chi$. Then $(M', c') \models \phi_c \wedge \neg \chi$ and hence $(M', b') \models \diamond(\phi_c \wedge \neg \chi)$. By our assumption, $(M, b) \models \diamond(\phi_c \wedge \neg \chi)$. Then there is a $d \in A$ such that $(M, d) \models \phi_c \wedge \neg \chi$. By (\dagger) , the first conjunct entails that $(M, c) \equiv_{\mathfrak{L}^*} (M, d)$. However, we have $(M, c) \models \chi$ and $(M, d) \not\models \chi$. Contradiction. Hence, $(M, c) \equiv_{\mathfrak{L}^*} (M', c')$. Therefore, because $c \in A$ was arbitrary, for every $c \in A$ there is a $c' \in A'$ such that $(M, c) \equiv_{\mathfrak{L}^*} (M', c')$. \square

We now present a list of nine structural conditions on relations R that link the action profiles of two given deontic game models. Each of the twenty-seven languages in \mathbf{L} is assigned a specific subset of these nine structural conditions. Because all of these subsets include at least three basic structural conditions, we set these three conditions apart. The relations that satisfy them are *basic bisimulations*:

Definition 7. Let $M = \langle \mathcal{N}, (A_i), d, v \rangle$ and $M' = \langle \mathcal{N}', (A'_i), d', v' \rangle$ be deontic game models. A relation $R \subseteq A \times A'$ is a *basic bisimulation* between M and M' if R satisfies each of the following clauses:

- (i) for all $(b, b') \in R$ and all $p \in \mathfrak{P}$ it holds that $b \in v(p)$ iff $b' \in v'(p)$
- (ii) for all $b \in A$ there is a $b' \in A'$ such that $(b, b') \in R$
- (iii) for all $b' \in A'$ there is a $b \in A$ such that $(b, b') \in R$.

The other six structural conditions on relations R are specified as follows:

Definition 8. Let $M = \langle \mathcal{N}, (A_i), d, v \rangle$ and $M' = \langle \mathcal{N}', (A'_i), d', v' \rangle$ be deontic game models. Let $R \subseteq A \times A'$ be a relation between M and M' . The six optional structural conditions on R are the following:

- $C_{\star \mathcal{G}}$: for all $(b, b') \in R$ and all non-empty $\mathcal{G} \subseteq \mathcal{N}$ it holds that $b_{\mathcal{G}} \in \text{Adm}_M(\mathcal{G})$ iff $b'_{\mathcal{G}} \in \text{Adm}_{M'}(\mathcal{G})$
- $C_{\star i}$: for all $(b, b') \in R$ and all $i \in \mathcal{N}$ it holds that $b_i \in \text{Adm}_M(i)$ iff $b'_i \in \text{Adm}_{M'}(i)$

- $C_{[\mathcal{G}]}$: (i) for all $(b, b') \in R$ and all $c \in A$ and all non-empty $\mathcal{G} \subseteq \mathcal{N}$ it holds that if $b_{\mathcal{G}} = c_{\mathcal{G}}$, then there is a $c' \in A'$ such that $b'_{\mathcal{G}} = c'_{\mathcal{G}}$ and $(c, c') \in R$
(ii) for all $(b, b') \in R$ and all $c' \in A'$ and all non-empty $\mathcal{G} \subseteq \mathcal{N}$ it holds that if $b'_{\mathcal{G}} = c'_{\mathcal{G}}$, then there is a $c \in A$ such that $b_{\mathcal{G}} = c_{\mathcal{G}}$ and $(c, c') \in R$
- $C_{[i]}$: (i) for all $(b, b') \in R$ and all $c \in A$ and all $i \in \mathcal{N}$ it holds that if $b_i = c_i$, then there is a $c' \in A'$ such that $b'_i = c'_i$ and $(c, c') \in R$
(ii) for all $(b, b') \in R$ and all $c' \in A'$ and all $i \in \mathcal{N}$ it holds that if $b'_i = c'_i$, then there is a $c \in A$ such that $b_i = c_i$ and $(c, c') \in R$
- $C_{(\mathcal{G})}$: (i) for all $c \in A$ and all non-empty $\mathcal{G} \subseteq \mathcal{N}$ it holds that if $c_{\mathcal{G}} \in \text{Adm}_M(\mathcal{G})$, then there is a $c' \in A'$ such that $c'_{\mathcal{G}} \in \text{Adm}_{M'}(\mathcal{G})$ and $(c, c') \in R$
(ii) for all $c' \in A'$ and all non-empty $\mathcal{G} \subseteq \mathcal{N}$ it holds that if $c'_{\mathcal{G}} \in \text{Adm}_{M'}(\mathcal{G})$, then there is a $c \in A$ such that $c_{\mathcal{G}} \in \text{Adm}_M(\mathcal{G})$ and $(c, c') \in R$
- $C_{(i)}$: (i) for all $c \in A$ and all $i \in \mathcal{N}$ it holds that if $c_i \in \text{Adm}_M(i)$, then there is a $c' \in A'$ such that $c'_i \in \text{Adm}_{M'}(i)$ and $(c, c') \in R$
(ii) for all $c' \in A'$ and all $i \in \mathcal{N}$ it holds that if $c'_i \in \text{Adm}_{M'}(i)$, then there is a $c \in A$ such that $c_i \in \text{Adm}_M(i)$ and $(c, c') \in R$.

Each of the twenty-seven sublanguages in \mathbf{L} is assigned a specific subset of these six structural conditions in the following way (C_{ϵ} is the empty condition that is satisfied trivially):

Definition 9. Let $M = \langle \mathcal{N}, (A_i), d, v \rangle$ and $M' = \langle \mathcal{N}, (A'_i), d', v' \rangle$ be deontic game models. Let $x \in \{\star_{\mathcal{G}}, \star_i, \epsilon\}$, $y \in \{[\mathcal{G}], [i], \epsilon\}$, and $z \in \{(\mathcal{G}), (i), \epsilon\}$. A relation $R \subseteq A \times A'$ is an $\mathfrak{L}_{x,y,z}$ -bisimulation between M and M' if R is a basic bisimulation that satisfies conditions C_x , C_y , and C_z .

Consequently, an \mathfrak{L} -bisimulation is just a basic bisimulation, and an $\mathfrak{L}_{\star_i, [\mathcal{G}]}$ -bisimulation is a basic bisimulation that also satisfies conditions C_{\star_i} and $C_{[\mathcal{G}]}$. We write $(M, a) \rightleftharpoons_{\mathfrak{L}_{x,y,z}} (M', a')$ if there is an $\mathfrak{L}_{x,y,z}$ -bisimulation R between M and M' such that $(a, a') \in R$.

We can now establish our general claim:

Theorem 2 (Hennessy-Milner). Let $\mathfrak{L}^* \in \mathbf{L}$. Then for all pointed deontic game models (M, a) and (M', a') it holds that $(M, a) \rightleftharpoons_{\mathfrak{L}^*} (M', a')$ if and only if $(M, a) \equiv_{\mathfrak{L}^*} (M', a')$.

Proof. Although the general claim is in fact a conjunction of claims on twenty-seven individual cases, the proofs of all these claims are very similar. Because the left-to-right direction is always proved by a straightforward structural induction on ϕ , we omit it here. As for the right-to-left direction, we will only give a full proof for the case where $\mathfrak{L}^* = \mathfrak{L}_{\star_{\mathcal{G}}, [\mathcal{G}], (\mathcal{G})}$. Suppose $(M, a) \equiv_{\mathfrak{L}_{\star_{\mathcal{G}}, [\mathcal{G}], (\mathcal{G})}} (M', a')$. Let $R = \{(b, b') \in A \times A' : (M, b) \equiv_{\mathfrak{L}_{\star_{\mathcal{G}}, [\mathcal{G}], (\mathcal{G})}} (M', b')\}$. We prove that R is an $\mathfrak{L}_{\star_{\mathcal{G}}, [\mathcal{G}], (\mathcal{G})}$ -bisimulation between M and M' . We first show that R is a basic bisimulation. Clause (i) of Definition 7 follows from the definition of R . Clauses (ii) and (iii) of Definition 7 follow from Lemma 6. Hence, R is a basic bisimulation. Next, we show that R satisfies conditions $C_{\star_{\mathcal{G}}}$, $C_{[\mathcal{G}]}$, and $C_{(\mathcal{G})}$.

$C_{\star_{\mathcal{G}}}$: Take an arbitrary $(b, b') \in R$ and an arbitrary non-empty $\mathcal{G} \subseteq \mathcal{N}$. Suppose $b_{\mathcal{G}} \in \text{Adm}_M(\mathcal{G})$. Then $(M, b) \models \star_{\mathcal{G}}$. By the definition of R , we have $(M', b') \models \star_{\mathcal{G}}$. Hence, $b'_{\mathcal{G}} \in \text{Adm}_{M'}(\mathcal{G})$. The proof of the converse is analogous.

$C_{[\mathcal{G}]}$: We prove clause (i) of condition $C_{[\mathcal{G}]}$. Take an arbitrary pair $(b, b') \in R$, an arbitrary $c \in A$, and an arbitrary non-empty $\mathcal{G} \subseteq \mathcal{N}$. Suppose $b_{\mathcal{G}} = c_{\mathcal{G}}$. Define ϕ_c as in the proof of Lemma 6. Then $(M, c) \models \phi_c$. Because $b_{\mathcal{G}} = c_{\mathcal{G}}$, it follows that $(M, b) \models \langle \mathcal{G} \rangle \phi_c$. By supposition and the definition of R , it must be that $(M', b') \models \langle \mathcal{G} \rangle \phi_c$. Then there is a $c' \in A'$ such that $b'_{\mathcal{G}} = c'_{\mathcal{G}}$ and $(M', c') \models \phi_c$. By the same reasoning as in the proof of Lemma 6, it must be that $(M, c) \equiv_{\mathfrak{L}_{\star_{\mathcal{G}}, [\mathcal{G}], (\mathcal{G})}} (M', c')$. Hence $(c, c') \in R$. Because (b, b') and c and \mathcal{G} were arbitrary, for all $(b, b') \in R$ and all $c \in A$ and all non-empty $\mathcal{G} \subseteq \mathcal{N}$ it holds that if $b_{\mathcal{G}} = c_{\mathcal{G}}$, then there is a $c' \in A'$ such that $b'_{\mathcal{G}} = c'_{\mathcal{G}}$ and $(c, c') \in R$. The proof of clause (ii) of condition $C_{[\mathcal{G}]}$ is analogous.

$C_{(\mathcal{G})}$: We prove clause (i) of condition $C_{(\mathcal{G})}$. Take an arbitrary $c \in A$ and an arbitrary non-empty $\mathcal{G} \subseteq \mathcal{N}$. Suppose $c_{\mathcal{G}} \in \text{Adm}_M(\mathcal{G})$. Define ϕ_c as in the proof of Lemma 6. Then $(M, c) \models \phi_c$. Hence, $(M, c) \not\models (\mathcal{G})\neg\phi_c$ and therefore $(M, a) \not\models (\mathcal{G})\neg\phi_c$. By supposition, it must be that $(M', a') \not\models (\mathcal{G})\neg\phi_c$. Then there is a $c' \in A'$ such that $c'_{\mathcal{G}} \in \text{Adm}_{M'}(\mathcal{G})$ and $(M', c') \models \phi_c$. By the same reasoning as in the proof of Lemma 6, it must be that $(M, c) \equiv_{\mathfrak{L}_{\star_{\mathcal{G}}, [\mathcal{G}], (\mathcal{G})}} (M', c')$. Hence, $(c, c') \in R$. Because c and \mathcal{G} were arbitrary, for all $c \in A$ and all $\mathcal{G} \in \mathcal{N}$ it holds that if $c_{\mathcal{G}} \in \text{Adm}_M(\mathcal{G})$, then there is a $c' \in A'$ such that $c'_{\mathcal{G}} \in \text{Adm}_{M'}(\mathcal{G})$ and $(c, c') \in R$. The proof of clause (ii) of condition $C_{(\mathcal{G})}$ is analogous.

Hence, R is an $\mathfrak{L}_{\star_{\mathcal{G}},[\mathcal{G}],(\mathcal{G})}$ -bisimulation between M and M' . Because $(a, a') \in R$, we conclude that $(M, a) \stackrel{\cdot}{\sim}_{\mathfrak{L}_{\star_{\mathcal{G}},[\mathcal{G}],(\mathcal{G})}} (M', a')$.

The proofs of the other twenty-six cases are simple variations of the above proof. Consider the language $\mathfrak{L}_{x,y,z}$. If $x = \star_i$, $y = [i]$, or $z = (i)$, then replace each occurrence of \mathcal{G} in the corresponding part of the above proof with i . If $x = \epsilon$, $y = \epsilon$, or $z = \epsilon$, then leave out the corresponding part of the above proof. \square

7 THE LEFT FRONT FACE




Figure 2: The Left Front Face Expressivity Relations

To chart the expressivity relations among the twelve sublanguages of the left front face of the full expressivity picture (Figure 2), it suffices to prove five theorems. Using these five theorems, the inclusion criterion, the transitivity criterion and Lemma 5(i), we can easily find the expressivity relations between every pair from the twelve sublanguages of Figure 2. The labels of the arrows in Figure 2 indicate the theorem that can be used to establish the corresponding expressivity relation. The five theorems are the following:

- (i) $\mathfrak{L}_{\star_{\mathcal{G}},[i]} \not\sim \mathfrak{L}_{[\mathcal{G}]}$ (Theorem 3)
- (ii) $\mathfrak{L}_{\star_{\mathcal{G}}} \not\sim \mathfrak{L}_{[i]}$ (Theorem 4)
- (iii) $\mathfrak{L}_{\star_i,([\mathcal{G}],(\mathcal{G}))} \not\sim \mathfrak{L}_{\star_{\mathcal{G}}}$ (Theorem 5)
- (iv) $\mathfrak{L}_{\star_i,([\mathcal{G}]} \not\sim \mathfrak{L}_{\star_i,(\mathcal{G})}$ (Theorem 6)
- (v) $\mathfrak{L}_{[\mathcal{G}]} \not\sim \mathfrak{L}_{\star_i}$ (Theorem 7).

Theorem 3. $\mathfrak{L}_{\star_{\mathcal{G}},[i]} \not\sim \mathfrak{L}_{[\mathcal{G}]}$.

Proof. Because \mathcal{N} contains at least three elements, there are $i, j, k \in \mathcal{N}$ such that $i \neq j$ and $j \neq k$ and $i \neq k$. We define two pointed deontic game models that are $\mathfrak{L}_{\star_{\mathcal{G}},[i]}$ -bisimilar (and hence validate exactly the same $\mathfrak{L}_{\star_{\mathcal{G}},[i]}$ -formulas) but give different truth-values to the $\mathfrak{L}_{[\mathcal{G}]}$ -formula $[i, k]p$.

Let $M_1 = \langle \mathcal{N}, (A_i), d, v \rangle$ be such that $A_i = \{a_i, b_i\}$ and $A_j = \{a_j, b_j\}$ and $A_k = \{a_k, b_k\}$ and $A_l = \{a_l\}$ for all $l \in \mathcal{N} - \{i, j, k\}$. The set A of action profiles is $\times_{i \in \mathcal{N}} A_i$. (Note that $A_{-\{i,j,k\}} = \{a_{-\{i,j,k\}}\}$.) Let $d(c) = 1$ for all $c \in A$. Lastly, let $v(p) = \{(a_i, a_j, a_k, a_{-\{i,j,k\}}), (a_i, a_j, b_k, a_{-\{i,j,k\}}), (a_i, b_j, a_k, a_{-\{i,j,k\}}), (b_i, a_j, a_k, a_{-\{i,j,k\}})\}$ and let $v(q) = \emptyset$ for all $q \in \mathfrak{P} - \{p\}$. Model M_1 can be pictured as in Figure 3 (the group action $a_{-\{i,j,k\}}$ is not represented).

| | | |
|-------|-------|-------|
| | a_j | b_j |
| a_i | $1/p$ | $1/p$ |
| b_i | $1/p$ | 1 |
| | a_k | |

| | | |
|-------|-------|-------|
| | a_j | b_j |
| a_i | $1/p$ | 1 |
| b_i | 1 | 1 |
| | b_k | |

Figure 3: Deontic game model M_1 .

Let $M'_1 = \langle \mathcal{N}, (A'_i), d', v' \rangle$ be such that $A'_i = \{a'_i, b'_i\}$ and $A'_j = \{a'_j, b'_j\}$ and $A'_k = \{a'_k\}$ and $A'_l = \{a'_l\}$ for all $l \in \mathcal{N} - \{i, j, k\}$. The set A' of action profiles is $\times_{i \in \mathcal{N}} A'_i$. (Note that $A'_{-\{i,j,k\}} = \{a'_{-\{i,j,k\}}\}$.) Let $d'(c') = 1$ for all $c' \in A'$. Lastly, let $v'(p) = \{(a'_i, a'_j, a'_k, a'_{-\{i,j,k\}}), (b'_i, b'_j, a'_k, a'_{-\{i,j,k\}})\}$ and $v'(q) = \emptyset$ for all $q \in \mathfrak{P} - \{p\}$. Model M'_1 can be pictured as in Figure 4 (the group action $a'_{-\{i,j,k\}}$ is not represented).

| | | |
|--------|--------|--------|
| | a'_j | b'_j |
| a'_i | $1/p$ | 1 |
| b'_i | 1 | $1/p$ |
| | a'_k | |

Figure 4: Deontic game model M'_1 .

Let $R = \{(c, c') \in A \times A' : c \in v(p) \text{ iff } c' \in v'(p)\}$. Because for all non-empty $\mathcal{G} \subseteq \mathcal{N}$ it holds that $M_1 \models \star_{\mathcal{G}}$ and $M'_1 \models \star_{\mathcal{G}}$, the relation R satisfies bisimulation condition $C_{\star_{\mathcal{G}}}$. For every $i \in \mathcal{N}$ and every $c \in A$ there are $c_{-i}^*, c_{-i}^{**} \in A_{-i}$ such that $(c_i, c_{-i}^*) \in v(p)$ and $(c_i, c_{-i}^{**}) \notin v(p)$. Likewise, for every $i \in \mathcal{N}$ and every $c' \in A'$ there are $c'_{-i}, c''_{-i} \in A'_{-i}$ such that $(c'_i, c'_{-i}) \in v(p)$ and $(c'_i, c''_{-i}) \notin v(p)$. From these observations, it follows that R satisfies bisimulation condition $C_{[i]}$. Therefore, R is an $\mathfrak{L}_{\star_{\mathcal{G}}, [i]}$ -bisimulation between M_1 and M'_1 . Hence, $(M_1, (a_i, a_j, a_k, a_{-\{i,j,k\}})) \rightleftharpoons_{\mathfrak{L}_{\star_{\mathcal{G}}, [i]}} (M'_1, (a'_i, a'_j, a'_k, a'_{-\{i,j,k\}}))$. By Theorem 2, we obtain $(M_1, (a_i, a_j, a_k, a_{-\{i,j,k\}})) \equiv_{\mathfrak{L}_{\star_{\mathcal{G}}, [i]}} (M'_1, (a'_i, a'_j, a'_k, a'_{-\{i,j,k\}}))$. However, note that it holds that $(M_1, (a_i, a_j, a_k, a_{-\{i,j,k\}})) \models [i, k]p$ and also that $(M'_1, (a'_i, a'_j, a'_k, a'_{-\{i,j,k\}})) \not\models [i, k]p$. By Lemma 3, $\mathfrak{L}_{\star_{\mathcal{G}}, [i]} \not\subseteq \mathfrak{L}_{[i]}$. \square

Theorem 4. $\mathfrak{L}_{\star_{\mathcal{G}}} \not\subseteq \mathfrak{L}_{[i]}$.

Proof. Because \mathcal{N} contains at least two elements, there are $i, j \in \mathcal{N}$ such that $i \neq j$. We define two pointed deontic game models that are $\mathfrak{L}_{\star_{\mathcal{G}}}$ -bisimilar (and hence validate exactly the same $\mathfrak{L}_{\star_{\mathcal{G}}}$ -formulas) but give different truth-values to the $\mathfrak{L}_{[i]}$ -formula $[i]p$.

Let $M_2 = \langle \mathcal{N}, (A_i), d, v \rangle$ be such that $A_i = \{a_i, b_i\}$ and $A_j = \{a_j\}$ and $A_k = \{a_k\}$ for all $k \in \mathcal{N} - \{i\}$. The set A of action profiles is $\times_{i \in \mathcal{N}} A_i$. (Note that $A_{-\{i,j\}} = \{a_{-\{i,j\}}\}$.) Let $d(c) = 1$ for all $c \in A$. Lastly, let $v(p) = \{(a_i, a_j, a_{-\{i,j\}})\}$ and $v(q) = \emptyset$ for all $q \in \mathfrak{P} - \{p\}$. Model M_2 can be pictured as in Figure 5 (left; the group action $a_{-\{i,j\}}$ is not represented).

Let $M'_2 = \langle \mathcal{N}, (A'_i), d', v' \rangle$ be such that $A'_i = \{a'_i, b'_i\}$ and $A'_j = \{a'_j, b'_j\}$ and $A'_k = \{a'_k\}$ for all $k \in \mathcal{N} - \{i, j\}$. The set A' of action profiles is $\times_{i \in \mathcal{N}} A'_i$. (Note that $A'_{-\{i,j\}} = \{a'_{-\{i,j\}}\}$.) Let $d'(c') = 1$ for all $c' \in A'$. Lastly, let $v'(p) = \{(a'_i, a'_j, a'_{-\{i,j\}}), (b'_i, b'_j, a'_{-\{i,j\}})\}$ and $v'(q) = \emptyset$ for all $q \in \mathfrak{P} - \{p\}$. Model M'_2 can be pictured as in Figure 5 (right; the group action $a'_{-\{i,j\}}$ is not represented).

| | |
|-------|-------|
| | a_j |
| a_i | $1/p$ |
| b_i | 1 |

| | | |
|--------|--------|--------|
| | a'_j | b'_j |
| a'_i | $1/p$ | 1 |
| b'_i | 1 | $1/p$ |

Figure 5: Deontic game models M_2 (left) and M'_2 (right).

Let $R = \{(c, c') \in A \times A' : c \in v(p) \text{ iff } c' \in v'(p)\}$. It is easy to check that R is an \mathfrak{L}_{\star_G} -bisimulation between M_2 and M'_2 . Hence, $(M_2, (a_i, a_j, a_{-\{i,j\}})) \rightleftarrows_{\mathfrak{L}_{\star_G}} (M'_2, (a'_i, a'_j, a'_{-\{i,j\}}))$. By Theorem 2, it must be that $(M_2, (a_i, a_j, a_{-\{i,j\}})) \equiv_{\mathfrak{L}_{\star_G}} (M'_2, (a'_i, a'_j, a'_{-\{i,j\}}))$. However, observe that $(M_2, (a_i, a_j, a_{-\{i,j\}})) \models [i]p$ and $(M'_2, (a'_i, a'_j, a'_{-\{i,j\}})) \not\models [i]p$. By Lemma 3, it must be that $\mathfrak{L}_{\star_G} \not\subseteq \mathfrak{L}_{[i]}$. \square

Theorem 5. $\mathfrak{L}_{\star_i, [\mathcal{G}], (\mathcal{G})} \not\subseteq \mathfrak{L}_{\star_G}$.

Proof. Because \mathcal{N} contains at least two elements, there are $i, j \in \mathcal{N}$ such that $i \neq j$. We define two pointed deontic game models that are $\mathfrak{L}_{\star_i, [\mathcal{G}], (\mathcal{G})}$ -bisimilar (and hence validate exactly the same $\mathfrak{L}_{\star_i, [\mathcal{G}], (\mathcal{G})}$ -formulas) but give different truth-values to the \mathfrak{L}_{\star_G} -formula $\star_{\{i,j\}}$.

Let $M_3 = \langle \mathcal{N}, (A_i), d, v \rangle$ be such that $A_i = \{a_i\}$ and $A_j = \{a_j\}$ and $A_k = \{a_k\}$ for all $k \in \mathcal{N} - \{i, j\}$. The set A of action profiles is $\times_{i \in \mathcal{N}} A_i$. (Note that $A_{-\{i,j\}} = \{a_{-\{i,j\}}\}$.) Let $d(a_i, a_j, a_{-\{i,j\}}) = 1$. Lastly, let $v(p) = \emptyset$ for all $p \in \mathfrak{P}$. Model M_3 can be pictured as in Figure 6 (left; the group action $a_{-\{i,j\}}$ is not represented).

Let $M'_3 = \langle \mathcal{N}, (A'_i), d', v \rangle$ be such that $A'_i = \{a'_i, b'_i\}$ and $A'_j = \{a'_j, b'_j\}$ and $A'_k = \{a'_k\}$ for all $k \in \mathcal{N} - \{i, j\}$. The set A' of action profiles is $\times_{i \in \mathcal{N}} A'_i$. (Note that $A'_{-\{i,j\}} = \{a'_{-\{i,j\}}\}$.) Let $d'(a'_i, b'_j, a'_{-\{i,j\}}) = d'(b'_i, a'_j, a'_{-\{i,j\}}) = 1$ and $d'(a'_i, a'_j, a'_{-\{i,j\}}) = d'(b'_i, b'_j, a'_{-\{i,j\}}) = 0$. Lastly, let $v(p) = \emptyset$ for all $p \in \mathfrak{P}$. Model M'_3 can be pictured as in Figure 6 (right; the group action $a'_{-\{i,j\}}$ is not represented).

| | |
|-------|-------|
| | a_j |
| a_i | 1 |

| | | |
|--------|--------|--------|
| | a'_j | b'_j |
| a'_i | 0 | 1 |
| b'_i | 1 | 0 |

Figure 6: Deontic game models M_3 (left) and M'_3 (right).

Let $R = A \times A'$. It is easy to check that R is an $\mathfrak{L}_{\star_i, [\mathcal{G}], (\mathcal{G})}$ -bisimulation between M_3 and M'_3 . Hence, $(M_3, (a_i, a_j, a_{-\{i,j\}})) \rightleftarrows_{\mathfrak{L}_{\star_i, [\mathcal{G}], (\mathcal{G})}} (M'_3, (a'_i, a'_j, a'_{-\{i,j\}}))$. By Theorem 2, it must be that $(M_3, (a_i, a_j, a_{-\{i,j\}})) \equiv_{\mathfrak{L}_{\star_i, [\mathcal{G}], (\mathcal{G})}} (M'_3, (a'_i, a'_j, a'_{-\{i,j\}}))$.

However, observe that $(M_3, (a_i, a_j, a_{-\{i,j\}})) \models \star_{\{i,j\}}$ and $(M'_3, (a'_i, a'_j, a'_{-\{i,j\}})) \not\models \star_{\{i,j\}}$. By Lemma 3, $\mathfrak{L}_{\star_i, [\mathcal{G}], (\mathcal{G})} \not\equiv \mathfrak{L}_{\star_{\mathcal{G}}}$. \square

Theorem 6. $\mathfrak{L}_{\star_i, [\mathcal{G}]} \not\equiv \mathfrak{L}_{\star_i, (\mathcal{G})}$.

Proof. Because \mathcal{N} contains at least two elements, there are $i, j \in \mathcal{N}$ such that $i \neq j$. We define two pointed deontic game models that are $\mathfrak{L}_{\star_i, [\mathcal{G}]}$ -bisimilar (and hence validate exactly the same $\mathfrak{L}_{\star_i, [\mathcal{G}]}$ -formulas) but give different truth-values to the $\mathfrak{L}_{\star_i, (\mathcal{G})}$ -formula $(i, j)p$.

Let $M_4 = \langle \mathcal{N}, (A_i), d, v \rangle$ be such that $A_i = \{a_i, b_i\}$ and $A_j = \{a_j, b_j\}$ and $A_k = \{a_k\}$ for all $k \in \mathcal{N} - \{i, j\}$. The set A of action profiles is $\times_{i \in \mathcal{N}} A_i$. (Note that $A_{-\{i,j\}} = \{a_{-\{i,j\}}\}$.) Let $d(a_i, a_j, a_{-\{i,j\}}) = d(b_i, b_j, a_{-\{i,j\}}) = 1$ and $d(a_i, b_j, a_{-\{i,j\}}) = d(b_i, a_j, a_{-\{i,j\}}) = 0$. Lastly, let $v(p) = \{(a_i, a_j, a_{-\{i,j\}}), (b_i, b_j, a_{-\{i,j\}})\}$ and $v(q) = \emptyset$ for all $q \in \mathfrak{P} - \{p\}$. Model M_4 can be pictured as in Figure 7 (left; the group action $a_{-\{i,j\}}$ is not represented).

Let $M'_4 = \langle \mathcal{N}, (A'_i), d', v' \rangle$ be such that $A'_i = \{a'_i, b'_i\}$ and $A'_j = \{a'_j, b'_j\}$ and $A'_k = \{a'_k\}$ for all $k \in \mathcal{N} - \{i, j\}$. The set A' of action profiles is $\times_{i \in \mathcal{N}} A'_i$. (Note that $A'_{-\{i,j\}} = \{a'_{-\{i,j\}}\}$.) Let $d'(a'_i, b'_j, a'_{-\{i,j\}}) = d'(b'_i, a'_j, a'_{-\{i,j\}}) = 1$ and $d'(a'_i, a'_j, a'_{-\{i,j\}}) = d'(b'_i, b'_j, a'_{-\{i,j\}}) = 0$. Lastly, let $v'(p) = \{(a'_i, a'_j, a'_{-\{i,j\}}), (b'_i, b'_j, a'_{-\{i,j\}})\}$ and $v'(q) = \emptyset$ for all $q \in \mathfrak{P} - \{p\}$. Model M'_4 can be pictured as in Figure 7 (right; the group action $a'_{-\{i,j\}}$ is not represented).

| | | |
|-------|-------|-------|
| | a_j | b_j |
| a_i | $1/p$ | 0 |
| b_i | 0 | $1/p$ |

| | | |
|--------|--------|--------|
| | a'_j | b'_j |
| a'_i | $0/p$ | 1 |
| b'_i | 1 | $0/p$ |

Figure 7: Deontic game models M_4 (left) and M'_4 (right).

Let $R = \{(c, c') \in A \times A' : c \in v(p) \text{ iff } c' \in v'(p)\}$. It is easy to check that R is an $\mathfrak{L}_{\star_i, [\mathcal{G}]}$ -bisimulation between M_4 and M'_4 . Hence, $(M_4, (a_i, a_j, a_{-\{i,j\}})) \rightleftarrows_{\mathfrak{L}_{\star_i, [\mathcal{G}]}} (M'_4, (a'_i, a'_j, a'_{-\{i,j\}}))$. By Theorem 2, it must be that $(M_4, (a_i, a_j, a_{-\{i,j\}})) \equiv_{\mathfrak{L}_{\star_i, [\mathcal{G}]}} (M'_4, (a'_i, a'_j, a'_{-\{i,j\}}))$. However, observe that $(M_4, (a_i, a_j, a_{-\{i,j\}})) \models (i, j)p$ and $(M'_4, (a'_i, a'_j, a'_{-\{i,j\}})) \not\models (i, j)p$. By Lemma 3, $\mathfrak{L}_{\star_i, [\mathcal{G}]} \not\equiv \mathfrak{L}_{\star_i, (\mathcal{G})}$. \square

Theorem 7. $\mathfrak{L}_{[\mathcal{G}]} \not\equiv \mathfrak{L}_{\star_i}$.

Proof. Because \mathcal{N} contains at least two elements, there are $i, j \in \mathcal{N}$ such that $i \neq j$. We define two pointed deontic game models that are $\mathfrak{L}_{[\mathcal{G}]}$ -bisimilar (and

hence validate exactly the same $\mathcal{L}_{[\mathcal{G}]}$ -formulas) but give different truth-values to the \mathcal{L}_{\star_i} -formula \star_i .

Let $M_5 = \langle \mathcal{N}, (A_i), d, v \rangle$ be such that $A_i = \{a_i, b_i\}$ and $A_j = \{a_j\}$ and $A_k = \{a_k\}$ for all $k \in \mathcal{N} - \{i, j\}$. The set A of action profiles is $\times_{i \in \mathcal{N}} A_i$. (Note that $A_{-\{i, j\}} = \{a_{-\{i, j\}}\}$.) Let $d(c) = 1$ for all $c \in A$. Lastly, let $v(p) = \{(b_i, a_j, a_{-\{i, j\}})\}$ and $v(q) = \emptyset$ for all $q \in \mathfrak{P} - \{p\}$. Model M_5 can be pictured as in Figure 8 (left; the group action $a_{-\{i, j\}}$ is not represented).

Let $M'_5 = \langle \mathcal{N}, (A'_i), d', v' \rangle$ be such that $A'_i = \{a'_i, b'_i\}$ and $A'_j = \{a'_j\}$ and $A'_k = \{a'_k\}$ for all $k \in \mathcal{N} - \{i, j\}$. The set A' of action profiles is $\times_{i \in \mathcal{N}} A'_i$. (Note that $A'_{-\{i, j\}} = \{a'_{-\{i, j\}}\}$.) Let $d'(a'_i, a'_j, a'_{-\{i, j\}}) = 0$ and $d'(b'_i, a'_j, a'_{-\{i, j\}}) = 1$. Lastly, let $v'(p) = \{(b'_i, a'_j, a'_{-\{i, j\}})\}$ and $v'(q) = \emptyset$ for all $q \in \mathfrak{P} - \{p\}$. Model M'_5 can be pictured as in Figure 8 (right; the group action $a'_{-\{i, j\}}$ is not represented).

$$\begin{array}{c}
 a_j \\
 \hline
 a_i \quad \boxed{1} \\
 \hline
 b_i \quad \boxed{1/p}
 \end{array}
 \qquad
 \begin{array}{c}
 a'_j \\
 \hline
 a'_i \quad \boxed{0} \\
 \hline
 b'_i \quad \boxed{1/p}
 \end{array}$$

Figure 8: Deontic game models M_5 (left) and M'_5 (right).

Let $R = \{(c, c') \in A \times A' : c \in v(p) \text{ iff } c' \in v'(p)\}$. It is easy to check that R is an $\mathcal{L}_{[\mathcal{G}]}$ -bisimulation between M_5 and M'_5 . Hence, $(M_5, (a_i, a_j, a_{-\{i, j\}})) \rightleftarrows_{\mathcal{L}_{[\mathcal{G}]}} (M'_5, (a'_i, a'_j, a'_{-\{i, j\}}))$. By Theorem 2, it must be that $(M_5, (a_i, a_j, a_{-\{i, j\}})) \equiv_{\mathcal{L}_{[\mathcal{G}]}} (M'_5, (a'_i, a'_j, a'_{-\{i, j\}}))$. However, observe that $(M_5, (a_i, a_j, a_{-\{i, j\}})) \models \star_i$ and $(M'_5, (a'_i, a'_j, a'_{-\{i, j\}})) \not\models \star_i$. By Lemma 3, $\mathcal{L}_{[\mathcal{G}]} \not\equiv \mathcal{L}_{\star_i}$. \square

8 THE RIGHT FRONT FACE

To chart the expressivity relations among the nine sublanguages of the right front face of the full expressivity picture (Figure 9), it suffices to prove four theorems. Using these four theorems, the inclusion criterion, and the transitivity criterion, we can easily find the expressivity relations between every pair from the nine sublanguages in Figure 9. The labels of the arrows in Figure 9 indicate which theorem must be used to establish the corresponding expressivity relation. The four theorems are the following:

- (i) $\mathcal{L}_{[i], (\mathcal{G})} \not\equiv \mathcal{L}_{[\mathcal{G}]}$ (Theorem 8)
- (ii) $\mathcal{L}_{(\mathcal{G})} \not\equiv \mathcal{L}_{[i]}$ (Theorem 9)
- (iii) $\mathcal{L}_{[\mathcal{G}]} \not\equiv \mathcal{L}_{(i)}$ (Theorem 10)
- (iv) $\mathcal{L}_{[\mathcal{G}], (i)} \not\equiv \mathcal{L}_{(\mathcal{G})}$ (Theorem 11).




Figure 9: The Right Front Face Expressivity Relations

Theorem 8. $\mathfrak{L}_{[i],(\mathcal{G})} \not\preceq \mathfrak{L}_{[\mathcal{G}]}$.

Proof. By Lemma 5(i), it holds that $\mathfrak{L}_{\star_{\mathcal{G}},[i]} \succeq \mathfrak{L}_{[i],(\mathcal{G})}$. Theorem 3 states that $\mathfrak{L}_{\star_{\mathcal{G}},[i]} \not\preceq \mathfrak{L}_{[\mathcal{G}]}$. By the transitivity criterion, $\mathfrak{L}_{[i],(\mathcal{G})} \not\preceq \mathfrak{L}_{[\mathcal{G}]}$. \square

Theorem 9. $\mathfrak{L}_{(\mathcal{G})} \not\preceq \mathfrak{L}_{[i]}$.

Proof. By Lemma 5(i), it holds that $\mathfrak{L}_{\star_{\mathcal{G}}} \succeq \mathfrak{L}_{(\mathcal{G})}$. Theorem 4 states that $\mathfrak{L}_{\star_{\mathcal{G}}} \not\preceq \mathfrak{L}_{[i]}$. By the transitivity criterion, $\mathfrak{L}_{(\mathcal{G})} \not\preceq \mathfrak{L}_{[i]}$. \square

Theorem 10. $\mathfrak{L}_{[\mathcal{G}]} \not\preceq \mathfrak{L}_{(i)}$.

Proof. Consider models M_5 and M'_5 from the proof of Theorem 7. Recall that $(M_5, (a_i, a_j, a_{-\{i,j\}})) \Leftrightarrow_{\mathfrak{L}_{[\mathcal{G}]}} (M'_5, (a'_i, a'_j, a'_{-\{i,j\}}))$. However, observe that $(M_5, (a_i, a_j, a_{-\{i,j\}})) \not\models (i)p$ and $(M'_5, (a'_i, a'_j, a'_{-\{i,j\}})) \models (i)p$. By Lemma 3, $\mathfrak{L}_{[\mathcal{G}]} \not\preceq \mathfrak{L}_{(i)}$. \square

Theorem 11. $\mathfrak{L}_{[\mathcal{G}],(i)} \not\preceq \mathfrak{L}_{(\mathcal{G})}$.

Proof. Consider models M_4 and M'_4 from the proof of Theorem 6. Let $R = \{(c, c') \in A \times A' : c \in v(p) \text{ iff } c' \in v'(p)\}$. It is easy to check that R is an $\mathfrak{L}_{[\mathcal{G}],(i)}$ -bisimulation between M_4 and M'_4 . Hence, $(M_4, (a_i, a_j, a_{-\{i,j\}})) \rightleftharpoons_{\mathfrak{L}_{[\mathcal{G}],(i)}} (M'_4, (a'_i, a'_j, a'_{-\{i,j\}}))$. By Theorem 2, it must be that $(M_4, (a_i, a_j, a_{-\{i,j\}})) \equiv_{\mathfrak{L}_{[\mathcal{G}],(i)}} (M'_4, (a'_i, a'_j, a'_{-\{i,j\}}))$. However, observe that $(M_4, (a_i, a_j, a_{-\{i,j\}})) \models (i, j)p$ and $(M'_4, (a'_i, a'_j, a'_{-\{i,j\}})) \not\models (i, j)p$. By Lemma 3, $\mathfrak{L}_{[\mathcal{G}],(i)} \not\equiv \mathfrak{L}_{(\mathcal{G})}$. \square

9 THE HORIZONTALS

$$\begin{array}{ccc}
\mathfrak{L}_{\star_i, [\mathcal{G}], (\mathcal{G})} & \xrightarrow{\hspace{10em}} & \mathfrak{L}_{[\mathcal{G}], (\mathcal{G})} \\
& & \mathfrak{L}_{\star_i, [\mathcal{G}]} \xrightarrow{\hspace{2em}} \mathfrak{L}_{[\mathcal{G}], (i)} \\
\mathfrak{L}_{\star_i, [i], (\mathcal{G})} & \xrightarrow{\hspace{10em}} & \mathfrak{L}_{[i], (\mathcal{G})} \\
& & \mathfrak{L}_{\star_i, [i]} \xrightarrow{\hspace{2em}} \mathfrak{L}_{[i], (i)} \\
\mathfrak{L}_{\star_i, (\mathcal{G})} & \xrightarrow{\hspace{10em}} & \mathfrak{L}_{(\mathcal{G})} \\
& & \mathfrak{L}_{\star_i} \xrightarrow{\hspace{2em}} \mathfrak{L}_{(i)}
\end{array}$$

Figure 10: The Horizontal Expressivity Relations

To chart the six horizontal expressivity relations among the twelve sublanguages of Figure 10, it suffices to prove one theorem. Using this theorem, the inclusion criterion, the transitivity criterion and Lemma 5(iv), we can easily prove the six horizontal expressivity relations. The theorem we need is the following:

Theorem 12. $\mathfrak{L}_{[\mathcal{G}], (\mathcal{G})} \not\equiv \mathfrak{L}_{\star_i}$.

Proof. Because \mathcal{N} contains at least two elements, there are $i, j \in \mathcal{N}$ such that $i \neq j$. We define two pointed deontic game models that are $\mathfrak{L}_{[\mathcal{G}], (\mathcal{G})}$ -bisimilar (and hence validate exactly the same $\mathfrak{L}_{[\mathcal{G}], (\mathcal{G})}$ -formulas) but give different truth-values to the $\mathfrak{L}_{\star_i, (\mathcal{G})}$ -formula \star_i .

Let $M_6 = \langle \mathcal{N}, (A_i), d, v \rangle$ be such that $A_i = \{a_i\}$ and $A_j = \{a_j\}$ and $A_k = \{a_k\}$ for all $k \in \mathcal{N} - \{i, j\}$. The set A of action profiles is $\times_{i \in \mathcal{N}} A_i$. (Note that $A_{-\{i,j\}} = \{a_{-\{i,j\}}\}$.) Let $d(a_i, a_j, a_{-\{i,j\}}) = 1$. Lastly, let $v(p) = \emptyset$ for all $p \in \mathfrak{P}$.

Model M_6 can be pictured as in Figure 11 (left; the group action $a_{-\{i,j\}}$ is not represented).

Let $M'_6 = \langle \mathcal{N}, (A'_i), d', v \rangle$ be such that $A'_i = \{a'_i, b'_i\}$ and $A'_j = \{a'_j, b'_j\}$ and $A'_k = \{a'_k\}$ for all $k \in \mathcal{N} - \{i, j\}$. The set A' of action profiles is $\times_{i \in \mathcal{N}} A'_i$. (Note that $A'_{-\{i,j\}} = \{a'_{-\{i,j\}}\}$.) Let $d'(a'_i, a'_j, a'_{-\{i,j\}}) = 1$ and $d'(a'_i, b'_j, a'_{-\{i,j\}}) = d'(b'_i, a'_j, a'_{-\{i,j\}}) = d'(b'_i, b'_j, a'_{-\{i,j\}}) = 0$. Lastly, let $v(p) = \emptyset$ for all $p \in \mathfrak{P}$. Model M'_6 can be pictured as in Figure 11 (right; the group action $a'_{-\{i,j\}}$ is not represented).

| | | | | | |
|-------|------------|--|--------|-------------|-------------|
| a_i | a_j 1 | | a'_i | a'_j 1 | b'_j 0 |
| | | | b'_i | 0 | 0 |

Figure 11: Deontic game models M_6 (left) and M'_6 (right).

Let $R = A \times A'$. It is easy to check that R is an $\mathfrak{L}_{(\mathcal{G}),[\mathcal{G}]}$ -bisimulation between M_6 and M'_6 . Hence, $(M_6, (a_i, a_j, a_{-\{i,j\}})) \xleftrightarrow{\mathfrak{L}_{(\mathcal{G}),[\mathcal{G}]}} (M'_6, (b'_i, b'_j, a'_{-\{i,j\}}))$. By Theorem 2, it must be that $(M_6, (a_i, a_j, a_{-\{i,j\}})) \equiv_{\mathfrak{L}_{(\mathcal{G}),[\mathcal{G}]}} (M'_6, (b'_i, b'_j, a'_{-\{i,j\}}))$. However, observe that $(M_6, (a_i, a_j, a_{-\{i,j\}})) \models \star_i$ and $(M'_6, (b'_i, b'_j, a'_{-\{i,j\}})) \not\models \star_i$. By Lemma 3, $\mathfrak{L}_{(\mathcal{G}),[\mathcal{G}]} \not\subseteq \mathfrak{L}_{\star_i}$. \square

10 WEAKER ASSUMPTIONS ON MODELS AND LANGUAGES

Before concluding the paper, we briefly discuss two assumptions on our models and languages: (a) the assumption that \mathcal{N} is a fixed finite set of individual agents that contains *at least three elements*, and (b) the assumption that each sublanguge in \mathbf{L} includes the *universal modality*. How would the expressivity relations be affected if we were to weaken these assumptions? (Because these assumptions are not central to the philosophical debate on reductionism in the social sciences, the following considerations are largely technical in nature.)

Throughout the paper, we assumed that \mathcal{N} is a fixed finite set of individual agents that contains *at least three elements*. Note that only Theorems 3 and 8 rely on the assumption that \mathcal{N} contains at least three agents. There are two cases to consider:

- (i) If \mathcal{N} were to contain exactly one agent, then the full language would not contain statements about (non-singleton) groups. Consequently, we would need to consider only *eight* sublanguages. The expressivity relations for these eight languages are given by the subfigure of Figure 1 that contains

only those languages in \mathbf{L} that exclude collective admissibility constants, modalities for collective agency, and modalities for collective obligations.

- (ii) If \mathcal{N} were to contain exactly two agents, then the only statements about (non-singleton) groups would be of the form $\star_{\mathcal{N}}\phi$, $[\mathcal{N}]\phi$, and $(\mathcal{N})\phi$. Note that our models validate $[\mathcal{N}]\phi \leftrightarrow \phi$. Consequently (and unlike the case where \mathcal{N} contains at least three agents), for every language \mathcal{L}^* in \mathbf{L} it holds that \mathcal{L}^* is equally expressive as the language \mathcal{L}^{**} that is obtained from \mathcal{L}^* by adding modalities for collective agency. Hence, under these circumstances, modalities for collective agency could be added without an increase of expressive power. Our results concerning the sublanguages that exclude modalities for collective agency rely only on the assumption that \mathcal{N} contains at least two agents, and hence the expressivity relations among these sublanguages are given by Figure 1.

Finally, each sublanguage in \mathbf{L} includes the universal modality. What can we say about the expressivity relations between languages that exclude the universal modality? Note that because \mathcal{N} contains at least two distinct individual agents i and j , our models validate $\Box\phi \leftrightarrow [i][j]\phi$. Accordingly, in each language in \mathbf{L} that has modalities for individual agency, the universal modality is definable in terms of modalities for individual agency. (Note that for every language \mathcal{L}^* in \mathbf{L} it holds that if \mathcal{L}^* has modalities for collective agency, then \mathcal{L}^* has also modalities for individual agency.) Hence, in each language in \mathbf{L} that has modalities for individual agency, the universal modality could be left out without losing expressive power. This does *not* hold true for those languages in \mathbf{L} that lack modalities for individual agency. These languages lose expressive power if the universal modality is left out. As a consequence, some expressivity relations no longer obtain: in languages in \mathbf{L} without modalities for individual agency and without the universal modality, modalities for obligations cannot be defined in terms of the universal modality and the relevant deontic admissibility constant.

11 CONCLUSION

We classified the twenty-seven sublanguages in \mathbf{L} in terms of their expressive power. To achieve this classification, we first proved twenty-seven Hennessy-Milner theorems: for every language \mathcal{L}^* in \mathbf{L} it holds that two pointed deontic game models are equivalent on \mathcal{L}^* if and only if there is a bisimulation relation between the two pointed models that satisfies a set of structural conditions that are specific to \mathcal{L}^* . The Hennessy-Milner theorems are central to our proofs that a given sublanguage is strictly more expressive than another.

The expressivity landscape we charted in this paper strongly suggests that the answers to the two philosophical questions in the introduction are negative. Because $\mathfrak{L}_{[G]} \succ \mathfrak{L}_{[i]}$ and $\mathfrak{L}_{(G)} \succ \mathfrak{L}_{[G],(i)}$, there are statements about collective agency that are not logically equivalent to any statement about individual agency, and there are statements about collective obligations that are not logically equivalent to any statement about collective agency and individual obligations. Naturally, these are not the only reducibility questions that are answered by our current study. For example, because $\mathfrak{L}_{(G)} \succ \mathfrak{L}_{[i]}$ and $\mathfrak{L}_{[G]} \succ \mathfrak{L}_{[i],(G)}$, there are statements about collective obligations that are not logically equivalent to any statement about individual agency, and there are statements about collective agency that are not logically equivalent to any statement about individual agency and collective obligations.

Have we refuted reductionism about collective agency and collective obligations? No, we have not. Our main philosophical contribution to the debate on reductionism about collective agency and collective obligations is methodological: our study indicates that logical methods can be used to find precise answers to the questions we raised in the introduction. To vindicate reductionism, the philosophical interpretation of our logical results must be contested. This can be done in various ways. For instance, the modelling considerations that lie at the basis of our semantics might be criticized. The reductionist might specify alternative models and/or truth-conditions for statements about collective and individual deontic admissibility, statements about collective and individual agency, and statements about collective and individual obligations. Then, she might prove that her alternative modelling supports reductionism about collective agency and collective obligations. Our methodological contribution therefore highlights the fruitfulness of logical methods for the study of reductionism in the philosophy of the social sciences.

ACKNOWLEDGEMENTS

Allard Tamminga and Hein Duijf gratefully acknowledge financial support from the ERC-2013-CoG project REINS, no. 616512. Hein Duijf gratefully acknowledges funding from the ERC-2017-CoG project SEA, no. 771074. Frederik Van De Putte's research is funded by a Marie Skłodowska-Curie Fellowship, no. 795329, and by the Research Foundation – Flanders (FWO-Vlaanderen). We thank the anonymous referees at this journal for their valuable suggestions and comments.

REFERENCES

- Arrow, K. J. (1951). Alternative approaches to the theory of choice in risk-taking situations. *Econometrica* 19: 404–437.

- Balbani, P., A. Herzig, and N. Troquard (2008). Alternative axiomatics and complexity of deliberative stit theories. *Journal of Philosophical Logic* 37: 387–406.
- Belnap, N., M. Perloff, and M. Xu (2001). *Facing the Future: Agents and Choices in Our Indeterminist World*. New York: Oxford University Press.
- Blackburn, P., M. de Rijke, and Y. Venema (2001). *Modal Logic*. Cambridge: Cambridge University Press.
- Czelakowski, J. (2020). Deontology of compound actions. *Studia Logica* 108: 5–47.
- Duijf, H., A. Tamminga and F. Van De Putte (ms.). An impossibility result on methodological individualism.
- Herzig, A. and F. Schwarzenrüber (2008). Properties of logics of individual and group agency. In: C. Areces and R. Goldblatt (eds.), *Advances in Modal Logic: Volume 7* (pp. 133–149). London: College Publications.
- Horty, J. F. (1996). Agency and obligation. *Synthese* 108: 269–307.
- Horty, J. F. (2001). *Agency and Deontic Logic*. New York: Oxford University Press.
- Horty, J. F. and N. Belnap (1995). The deliberative stit: a study of action, omission, ability, and obligation. *Journal of Philosophical Logic* 24: 583–644.
- Kanger, S. (1957). New foundations for ethical theory. Privately distributed pamphlet, first published in R. Hilpinen (ed.), *Deontic Logic* (pp. 36–58). Dordrecht 1971: D. Reidel Publishing Company.
- Kohlberg, E. and J.-F. Mertens (1986). On the strategic stability of equilibria. *Econometrica* 54: 1003–1037.
- Kooi, B. and A. Tamminga (2008). Moral conflicts between groups of agents. *Journal of Philosophical Logic* 37: 1–21.
- Luce, R. D. and H. Raiffa (1957). *Games and Decisions*. New York: John Wiley and Sons.
- Meyer, J.-J. Ch. (1988). A different approach to deontic logic: deontic logic viewed as a variant of dynamic logic. *Notre Dame Journal of Formal Logic* 29: 109–136.
- Murakami, Y. (2005). Utilitarian deontic logic. In: R. Schmidt, I. Pratt-Hartmann, M. Reynolds, and H. Wansing (eds.), *Advances in Modal Logic: Volume 5* (pp. 211–230). London: College Publications.
- Osborne, M. J. and A. Rubinstein (1994). *A Course in Game Theory*. Cambridge, MA: The MIT Press.
- Pörn, I. (1970). *The Logic of Power*. Oxford: Basil Blackwell.

- Savage, L. J. (1972). *The Foundations of Statistics*, second revised edition. New York: Dover Publications, Inc.
- Segeberg, K. (2012). D Δ L: a dynamic deontic logic. *Synthese* 185: 1–17.
- Tamminga, A. and H. Duijf (2017). Collective obligations, group plans and individual actions. *Economics and Philosophy* 33: 187–214.
- Tamminga, A. and F. Hindriks (2020). The irreducibility of collective obligations. *Philosophical Studies*, in press.
- van Benthem, J., N. Bezhianishvili, S. Enqvist, and J. Yu (2017). Instantial neighbourhood logic. *Review of Symbolic Logic* 10: 116–144.
- von Kutschera, F. (1986). Bewirken. *Erkenntnis* 24: 253–281.
- Wansing, H. (2006). Tableaux for multi-agent deliberative-stit logic. In: G. Governatori, I. Hodkinson and Y. Venema (eds.), *Advances in Modal Logic: Volume 6* (pp. 503–520). London: College Publications.
- Xu, M. (1998). Axioms for deliberative stit. *Journal of Philosophical Logic* 27: 505–552.