Three Deontic Logics for Rational Agency in Games

Huimin Dong* and Olivier Roy

Department of Philosophy, University of Bayreuth, Germany
huimin.dong@uni-bayreuth.de

Abstract. In [13, 1] the authors have argued that, in games, obligations and permissions should be viewed, respectively, as giving necessary and sufficient conditions for rationality. This gives rise to a specific deontic logic where, for instance, O and P are not dual notions and P becomes a "free choice" permission operator. Similar deontic logics have been proposed in the literature, as early as van Benthem's [15], and more recently in [9]. In this paper we study the relation between these deontic logics for rational agency in games. We compare their deductive power, provide translation results, and emphasize the different views they take on what players ought to, or may do.

This paper studies a family of deontic logics that diverge from Standard Deontic Logic (SDL) [10] in that O and P are not dual, and P validates the infamous "free choice" principle:

$$P(A \vee B) \rightarrow PA \wedge PB \qquad \text{(FCP)}$$

In [13, 1], the authors have argued that such deontic logics are well-suited to capturing rational obligations and permissions in games, i.e. what the players ought to, and may do according to particular solution concepts. The similarity between the logic proposed by the authors and a number of other deontic systems has been observed in [1]. But the precise comparison remained to be made. This is the main contribution of the present paper.

This contribution should be of interest to philosophical logicians working on game theory for two reasons. It shows first that three independent proposals in deontic logic are well-suited to describe the rational obligations and permissions that bear on players in games, even though these systems might not have been originally devised for that purpose. This is a conceptual contribution. Second, on the formal side, it provides a systematic comparison of the deductive power of the first two systems studied here, and shows that the third can be embedded in the first.

Section 1 reviews the normative interpretation of solution concepts in game theory, and the argument given in [13, 1] for the particular structure of obligations and permissions that they give rise to. Section 2 provides the first comparison, between van Benthem's "Minimal Deontic Logic" [15] and Anglberger *et al.*

* Supported by the MOE Project of Key Research Institute of Humanities and Social Sciences in Universities (No. 12JJD720005), and by the China Scholarship Council grant (CSC No. 201306380078).

"Obligations as Weakest Permissions" [12]. Section 3 compares van Benthem's system with Trupuz and Kulicki's "Deontic Boolean Action Logic" [9]. Section 4 concludes.

1 The deontic logic of rational recommendations in games

By rational recommendations in games we mean obligations and permissions stemming from classical game-theoretic solution concepts. The goal of this section is to argue that such recommendations give rise to a specific deontic logic, one that differs from SDL. On the way there we give a brief, informal introduction to the game-theoretic solution concepts that we have in mind (Section 1.1).

1.1 Normative interpretation of solution concepts

In a game a number of self-interested players interact in what Schelling called "interdependent decision" [refs]. The result of each player's decision depends on what all the other players do. In this paper we will be looking only at so-called games in strategic form. The formal definition is this:

Definition 1. *A game in strategic form* $\mathbb{G}$ *is a tuple* $\langle I, \{S_i\}_{i \in I}, \pi \rangle$*.* I *is a finite set of agents or players.* S_i *is a finite set of actions or strategies for each player i. A strategy profile* σ *is a combination of strategies, one for each player. We write* σ_i *for agent i's strategy in* σ*, and* σ_{-i} *for the strategies of all agents except i in* σ*. The set of all strategy profiles is denoted S. The payoff function* $\pi : S \to \mathbb{R}^I$ *assigns to each strategy profile a vector of real-valued payoffs for the players. We use* π_i *to denote i's component in that vector.*

Let us consider a concrete example: the game "Guess 2/3 of the Average," a classic for introductory game-theory courses and experiments [1]. A number n of players have to choose a natural number between 0 and 100. They do so simultaneously, without knowing what the others do. The winner is the player whose choice is closest to 2/3 of the average number chosen. In case there is more than one "winner", the players split the prize. Players prefer having more of the prize than less. So each prefers to be the unique winner.

Let us consider a very simple, two-player version of that game, say between Ann and Bob. So I = {Ann, Bob}, and $S_{Ann} = \{0, 1, ..., 100\}$, and similarly for Bob. The payoff function is first defined player-wise, as follows, with a pair (k, l) representing Ann's choice k and Bob's choice l, and $a = 2/3\frac{k+l}{2}$.

$$\pi_{Ann}(k,l) = \begin{cases} 1 & |k-a| < |l-a| \\ 1/2 & |k-a| = |l-a| \\ 0 & \textit{otherwise} \end{cases}$$

[1] See `https://en.wikipedia.org/wiki/Guess_2/3_of_the_average` for an overview.

Bob's payoff function $\pi_{Bob}(k, l)$ is just defined as $1 - \pi_{Ann}(k, l)$, and $\pi(k, l)$ as $(\pi_{Ann}(k, l), \pi_{Bob}(k, l))$.

What should Ann and Bob do in this game? The standard solution concept to be applied here is that of *Nash equilibrium*, which is computed using the *best response dynamics*. Put yourself in Ann's position, and consider the case in which you and Bob choose 100. In that case you split the prize. Both of you are equally near to 2/3 of 100. But then, given that Bob chooses 100, you (Ann) could have done better by choosing any lower number, claiming then the prize all for yourself. In game-theoretic terms, choosing 100 is not a *best response* for Ann to Bob choosing 100. Technically, a best response function is an action that, *given the choice of the other players*, yields an outcome that is at least as good as the outcome yielded by any other action. Formally, for each $l \in S_{Bob}$, Ann's best response $br(Ann, l)$ is defined as $\{k : \pi_{Ann}(k, l) \geq \pi_{Ann}(k', l), \text{for all } k' \in S_{Ann}\}$.[2] Observe that for Ann this set is not a singleton. Best response need not be unique. Playing *anything* lower than 100 will make her the unique winner, given that Bob plays 100. But playing 100 is *not* a best response. The situation is of course entirely symmetric for Bob. Given that Ann plays 100, the best response for him is to play something lower.

Suppose Ann and Bob then play different numbers, each lower than 100, but higher or equal to 1. Say Ann plays the highest number. This is not a best response for her. She should play a lower number, either slightly above or slightly below Bob's, depending on how far he is from the 2/3 of the average point. But doing so will lower that point, now making even lower choices best replies for Bob.

This dynamics will go on until both Ann and Bob have chosen 0. There they play a *mutual best response*, a.k.a. a *Nash equilibrium*. Given that you choose 0, Bob has no incentive to choose anything else. In all other cases he forfeits the prize entirely to Ann. And vice-versa for her. The formal definition of a Nash equilibrium is as follows:

Definition 2. *Let $\mathbb{G}$ be a game in strategic form and br be a best response function for that game. Then σ is a Nash equilibrium iff $\sigma_i = br(i, \sigma_{-i})$ for all players $i \in I$.*

Best response and equilibrium play are two *solution concepts* for games. In their normative interpretation, they are intended to capture the idea of a rational action or a rational play. In this paper we will use best response as our running example. Ann and Bob should not, on pain of irrationality, play actions that are not best response to one another.[3] We will call *rational recommendations* the normative prescriptions that one gets from such solution concepts in games,

[2] The general definition of the best response set $br(i, \sigma_{-i})$ for player i to the choice σ_{-i} of the others is $\{s_i \in S_i : \pi_i(s_i, \sigma_{-i}) \geq \pi_{Ann}(s'_i, \sigma_{-i}), \text{for all } s'_i \in S_i\}$.

[3] Why? One way to answer is to go back to decision theory. There the standard of rationality for decision under risk is maximization of expected utility. Simply put, a player should choose actions for which the player has the strongest belief that that action will lead to a good outcome. Choosing otherwise can lead to practical incoherence. See [8] for an overview of the normative interpretation of decision theory, and [11] for an overview of its application to games.

for instance the recommendation that Ann should not play 100 if Bob plays that too. Different solution concepts will of course yield different (but not unrelated!) rational recommendations in games. And Nash equilibrium and best response are surely not the only solution concepts around. In recent years iterated elimination procedures, using either strict or weak dominance, have for instance attracted much attention from epistemic game theorist and logicians [refs]. Here, however, we will not look at the structure of specific solution concepts. There is already an extensive literature on logical characterizations of, say iterated strict dominance or equilibrium play [refs - Pauly and van der Hoek, Tamminga]. Rather, our point is that the abstract logic of rational recommendations in games, *whatever the underlying solution concept* should have a particular structure. This is what we will argue now.

1.2 The logical structure of rational recommendations

We now review the argument given in [13, 1] for the following claim: rational obligations and permissions in games should be seen, respectively, as giving necessary and sufficient conditions for rational play, and as a result the two notions should not be seen as dual to each other. This is a philosophical argument. If the argument is correct then this has important implications for the deontic logic of such rational obligations and permissions.

Rationality is the key normative notion underlying solution concepts. As we have seen, solution concepts pinpoint a subset of profiles that are intuitively deemed rational in a game, sometimes given additional information about the strategies that are at play or the beliefs of the players. Consider again Guess 2/3 of the Average. Here best response prescribes for both Ann and Bob to play a lower number, given that the other plays 100. A Nash equilibrium profile in that game is one where Ann and Bob play mutual best response to what the other is doing. The profile (0,0) is the unique Nash equilibrium in pure strategies of that game.[4]

Solution concepts, interpreted normatively, give recommendations to the players. But what kind? Our first claim is that they provide rational *permissions*, as opposed to obligations.

> **Rational Permissions** Solution concepts in games pinpoint rational *permissions*, not necessarily rational obligations.

The argument for this claim starts with the basic observation that there is in general no unique solution to a given game. Take again the recommendation of best response given that the other plays 100. Any number from 0 to 99 is a best response. The only non-best response is playing 100 oneself. In the face of such a plurality of solutions it doesn't make sense to say that the players *ought* to play *all* of these numbers. They simply cannot do that. These are mutually

[4] The situation is more complicated for more than two players. There everyone playing 1 can be an equilibrium in pure strategies. A unique deviation to 0 might not lower the average enough to ensure a win.

exclusive actions. So, if "ought implies can" then it is not the case that players are under a rational obligation to play every solution. What remains is that playing any solution of a game is rationally permissible. In Guess 2/3 of the Average, given that the other is playing 100, any number between 0 and 99 is rationally permissible. The situation is of course not particular to best response in this particular instance of that game. Non-unique solutions are ubiquitous in game theory. In the face of this, the adequate ways to understand rational recommendations from solution concepts is in terms of rational permissions. [5]

Our next claim is that rational permissions provide sufficient conditions for best response. Here we only illustrate it through our running example. The argument is developed in detail in [13, 1]. Given that Bob plays 100, it is rationally permissible for Ann to play any number, as long as it is lower than 100. Let us introduce some action-theoretic terminology, which we will make formal later on. Call an *action* or a *strategy type* just a set of action/strategies for one player, and similarly for strategy profiles. In Guess 2/3 of the Average, the type "playing a number lower than 100" is rationally permitted for Ann by best response, given that Bob plays 100. *If* she plays any strategy of that type *then* she plays a best response strategy. Playing less than 100 is sufficient for rationality. Observe, furthermore, that playing any number which is an instance of a logically stronger action type will also be best response for Ann against Bob playing 100.[6] So picking among the set of *even* numbers lower than 100, or just picking 0 for that matter, will imply playing a best response. From the perspective of best response to 100 alone, these are all on a par.

So if "playing a number lower than 100" is a rationally permitted type, and rational permission for the action type provides sufficient conditions for rationality, then any of these logically stronger types should be seen as permitted as well. Best response cannot distinguish them any further. In [13, 1] it is argued that this holds more generally, for any rational permission in games. In a nutshell, this gives us the following principle:

> **Strong Rational Permissions (SRP)** An action type φ is rationally permitted in game $\mathbb{G}$ if and only if playing a strategy of type φ implies playing a rational strategy.

If SRP is correct, this has important consequences for the logical analysis of rational recommendations. The most important one is that one should take

[5] Of course, when there is a unique profile that is rationally permitted, that profile becomes rationally obligatory. This is the case, for instance, for the profile (0,0), using rational recommendations from Nash equilibrium in the game above. This is not only what the players are rationally permitted to do. They ought to play (0,0). This will correspond to a logical principle connecting obligations and permissions, which we will encounter later on.

[6] A type φ is logically stronger than type ψ when φ is strictly contained in ψ. Playing a strategy of type φ implies playing strategy of type ψ.

"Free Choice Permission" on board:[7]

$$P(\varphi \vee \psi) \rightarrow P\varphi \wedge P\psi$$

Indeed, if $\varphi \vee \psi$ is viewed as a non-deterministic choice between playing φ or ψ, then if playing a strategy of either type is sufficient for rationality, then both playing any strategy of type φ and any strategy of type ψ is sufficient for rational play. So both are permitted by SRP.

If permissions provide sufficient conditions for rationality, the natural counterpart is the view of obligations as necessary conditions. In other words:

> **Weak Rational Obligations (WRO)** An action type φ is rationally obligatory, or *rationally required* in game $\mathbb{G}$ if and only if not playing a strategy of type φ implies not playing a rational strategy.

We call this principle "weak" because it suggests a form of closure of obligations under logically weaker types. If it is rationally required to do φ then playing rationally implies playing a strategy of type φ. But then this also implies playing any weaker type of strategy, and in particular the trivial type $\varphi \vee \neg\varphi$. So logically very weak types of action will turn out to be obligatory. We will see this in the concrete case of Guess 2/3 of the Average in Section 2.

An important consequence of accepting both SRP and WRO is that obligations and permissions are not necessarily duals anymore. In our example it is not the case that playing a number higher than 50 is permitted for Ann as best response to 100. This is not sufficient for best response to 100, because playing 100 herself is a strategy of that type. So by SRP this is not permitted. But not playing a number higher than 50 is just the same as playing a number lower or equal to 49. But this cannot be obligatory either, because not playing this does not entail not playing a best response. So rational permissions and rational requirements are not dual here.

We take these as the central features of rational recommendations in games: SRP, WRO— viz. obligations and permissions provide, respectively, necessary and sufficient conditions for rationality — and these two normative categories are not dual. In this section we have sketched the philosophical arguments for these claims. They are developed in more detail in [13, 1], where the authors present a deontic logic of "obligation as weakest permission" that has all these features. In the next sections we compare that logic with two very congenial

[7] This principle has a reputation for misbehavior in combination with SDL. If permissions are normal modalities in the technical sense, then $P\varphi \rightarrow P\psi$ becomes easily derivable using FCP. We will see an example of this in Section 2. So rational permissions should not be normal, and indeed they are not in any of the logical systems that we present below. Free Choice can also cause problems for non-normal modalities, as long as they are extensional. This is the by now well-known "vegetarian free lunch" example [7]. If ordering a vegetarian meal is permitted, then by FCP the logically stronger action type "ordering a vegetarian meal and not pay for it" must be also permitted, at least if the boolean constructors on action types are classical. See [2] for an answer to that criticism.

proposals, an earlier one by van Benthem and a more recent one by Trypuz and Kulicki. As we will see they all share the three central features, and as such can be seen as logic for rational recommendations in games, although they differ either in their philosophical commitment or their expressive power.

2 Minimal Deontic Logic and Weakest Permissions

We start by comparing the logic of "obligations as weakest permissions" [12] with van Benthem's "minimal deontic logic" [15]. They have much in common, both formally and in their analysis of rational recommendations. They differ, however, in their view of obligation. Applied to rational recommendations, van Benthem's deontic logic makes obligatory every necessary condition for rational play. In other words, playing any action type that rules out being rational is forbidden. This is not the case in the logic of obligations as weakest permissions. There the unique obligation bearing on the players is to play a rational strategy. The crux of this difference turns out to be the relation between obligation and permission in these two systems.

2.1 Common language

The two logics that we study now share the same language. They take both O and P as primitive, and use a universal modality $\Diamond$.

Definition 3. *Let p be any element of a given (countable) set Prop of atomic propositions. The language $\mathcal{L}$ is defined as follows:*

$$\varphi := p \mid \neg\varphi \mid \varphi \wedge \varphi \mid \Diamond\varphi \mid P\varphi \mid O\varphi$$

2.2 Propositions, action types

First a note on the interpretation of the structures used in the the semantics of our first two systems. These structures are familiar to modal logicians: binary relations or neighborhood functions defined on a set of objects. These objects, however, will be viewed from an action-theoretic perspective. Instead of thinking of them as possible worlds, we will take them to be atomic actions. When discussing concrete games, these will be either strategy profiles or strategies for individual players. While sets of states in standard Kripke semantics are propositions, here they will be taken as action types. In our running example, for instance, "playing an even number", "playing a number less than 100", or "playing 0" are all action types, with the latter just happening to be an atomic one. So the standard Boolean operations on propositions will correspond here to action type constructors, pretty much as in Propositional Dynamic Logic (PDL) or Boolean Modal Logic (BML) [3], and the resulting deontic logic will be one of "ought to do", as opposed to "ought to be."

2.3 Minimal Deontic Logic

In van Benthem's Minimal deontic logic the obligation O is a normal modality, and the permission operator P is a so-called window modality [3] defined on the set of normatively ideal action tokens. The presentation here is slightly different than in [15], to make it more in line with what comes later.

Definition 4 (Minimal Deontic Logic (MDL) Models). *A model* $M = \langle W, R_D, V \rangle$ *for the minimal deontic logic is defined as follows:*

- $W \neq \emptyset$ *is a set of atomic action.*
- $R_D \subseteq W \times W$.
- $V : Prop \to \wp(W)$ *is a valuation function*

This is just a standard Kripke model for deontic logic. The relation R_D pinpoints the normatively ideal action type, from the perspective of each atomic action or profile. In games, the normatively ideal actions will be those recommended by a specific solution concept. In principle what is rational or ideal in a MDL model may vary from action to action. This can be used to represent typical cases of interdependence between what one player does and what is rational for the others to do, as observed for instance in the best response dynamics that leads to the (0,0) equilibrium of Guess 2/3 of the average. In this section and the next, however, we will only consider *uniform* models. These are models where $R_D[w] = R_D[w']$ for all w, w', with $R_D[w] = \{v : R_D(w, v)\}$. The set of rational atomic actions is the same throughout the model.

The difference with SDL shows in the truth conditions for P.

$$M, w \models \Box\varphi \text{ iff } \forall v \in W.\ M, v \models \varphi$$
$$M, w \models O\varphi \text{ iff } \forall v \in W.\ (R_D wv \Rightarrow M, v \models \varphi)$$
$$M, w \models P\varphi \text{ iff } \forall v \in W.\ (M, v \models \varphi \Rightarrow R_D wv)$$

P is thus a "window modality" [3]. $P\varphi$ is true of all action types φ that are sub-types of the ideal type specified by R_D. In less technical terms, $P\varphi$ is true whenever playing a strategy of type φ ensures a rational play. So in this logical system permissions provide sufficient conditions for an action type to be "legal" or "licensed" by a given normative theory. The normative theory we are looking at right now is of course the rational recommendations stemming from a given solution concept in games. An action type is permitted, in this view, if playing that action type implies playing a rational strategy. So the logic embodies SRP.

Obligations, on the other hand, can be seen as providing necessary conditions for rationality in that system, and hence also to capture WRO. The core interaction principle behind this is the following, which we will encounter often later on in this paper:

$$(O\varphi \wedge P\psi) \to \Box(\psi \to \varphi)$$

In the context of games, this principle says that if one is not rational unless she plays a strategy of type φ, while playing a strategy of type ψ guarantees a rational play, then it must be that all type ψ strategies are φ strategies. Let Ra be the type of all rational strategies. Then combining SRP and WRO, we have:

$$\psi \to Ra \to \varphi$$

This connection between ψ and φ is expressed in the consequent of the previous formula, crucially using the universal modality. Observe that this logic also captures the third core features of logics for rational recommendations in games: O and P are not dual here.

Let us illustrate how this system would handle rational recommendations in Guess 2/3 of the Average, with Ann and Bob playing the game, and Bob playing 100. Then best response recommends Ann to play any number lower than 100. A natural way to represent this situation as a minimal deontic model is to take W to be the set of all pairs $(n, 100)$, with $0 \leq n \leq 100$, i.e. all possible choices for Ann given that Bob plays 100. The set of plays where Ann plays a best response against Bob playing 100, i.e. $br(Ann, 100)$, is then simply W minus the pair $(100, 100)$. Defining then $R_D[w]$ as $br(Ann, 100)$, for all w, we get the expected recommendations. Ann ought not to play 100, and playing anything below that is permitted, because it implies playing a best response.

So with van Benthem's minimal deontic logic we have our first logic for rational recommendations in games. It is worth considering its axiomatization before moving to our second logic. The full system is presented in Table 1. Since

(K-□)	$\Box(\varphi \to \psi) \to (\Box\varphi \to \Box\psi)$	(Incl)	$\Box\varphi \to O\varphi$
(K-O)	$O(\varphi \to \psi) \to (O\varphi \to O\psi)$	(CoIncl)	$\Box\neg\varphi \to P\varphi$
(OR)	$P\varphi \wedge P\psi \to P(\varphi \vee \psi)$	(WP)	$(O\varphi \wedge P\psi) \to \Box(\psi \to \varphi)$
(NEC)	$\dfrac{\varphi}{\Box\varphi}$	(Flip)	$\dfrac{\varphi \to \psi}{P\psi \to P\varphi}$

Table 1. The sound and complete axiom system for van Benthem's Minimal Deontic Logic. All propostional tautologies are also taken as axioms.

O is given a standard Kripke semantics, it validates K and the necessitation rule. The latter is readily derivable using (NEC) and (Incl). So O is a normal modality. Note in passing that this implies that in this logic some of the classical deontic paradoxes, for instance contrary to duty [5] or Rosss paradox, will be valid.

P, on the other hand, is not a normal modality. Necessitation fails for permissions in that logic. Permissions do validate Free Choice, which is directly derivable by (Flip). Note also that from this rule one can both derive the K axiom for P and the extensionality rule (E), familiar to axiomatization of non-normal modal logics, stating that if $\varphi \leftrightarrow \psi$ is a theorem of the logic, then one can substitute φ and ψ within the scope of a permission operator.

2.4 Obligation as Weakest Permission

The second logic we consider is the logic of obligations as weakest permission. The main difference with van Benthem's system is that obligations are no longer normal modalities here. An action type φ is obligatory, in this logic, whenever it is *exactly* the normatively ideal action type, i.e. a type of rational strategy. To put it bluntly, in this logic there is only one thing that agents ought to do: be rational!

Definition 5 (OWP models). *A model for the logic of obligations as weakest permission, or* OWP-*model for short, is a tuple* $M = \langle H, n_P, n_O, \|\cdot\| \rangle$ *where*

- *H is a set of atomic actions.*
- *$n_P : H \to \wp\wp(H)$ is a neighborhood function assigning sets of subsets of H to each $h \in H$ such that*
  - *If $X \cup Y \in N_P(h)$ then $X \in n_P(h)$ & $Y \in n_P(h)$*
- *$n_O : H \to \wp\wp(H)$ is a neighborhood function assigning a set of subsets of H to each $h \in H$ such that*
  - *$\emptyset \notin n_O(h)$.*
  - *(Ought-Perm) If $X \in n_O(h)$ then $X \in n_P(h)$*
  - *(Weakest-Perm) If $X \in n_O(h)$ then $Y \subseteq X$ for all $Y \in n_P(h)$.*
- *$\|\cdot\| : Prop \to \wp(H)$ is a valuation function.*

As usual, an OWP-frame is an OWP-model minus the valuation $\|\cdot\|$. We postpone the discussion of the frame conditions for a moment, observing only that putting (Ought-Perm) and (Weakest-Perm) together we obtain the following:

Observation 1 *If $n_O(h) \neq \emptyset$ then*[8]

$$n_O(h) = \{\bigcup n_P(h)\} \tag{C1}$$

In other words, if an action type is obligatory, then it is the *unique* action type that is obligatory, up to logical equivalence, and this is the logically weakest permission that the agent has.

The truth conditions for O and P in OWP-models are standard for neighborhood semantics, and validity is defined as usual. Let us abuse our notation and write $\|\varphi\|$ for $\{h : M, h \models \varphi\}$.

$$M, h \models \Box\varphi \text{ iff } \forall h' \in H.\ M, h' \models \varphi$$
$$M, h \models O\varphi \text{ iff } \|\varphi\| \in n_o(h)$$
$$M, h \models P\varphi \text{ iff } \|\varphi\| \in n_p(h)$$

The set of valid formulas is completely axiomatizable by the system in Table 2. Some of its features are worth highlighting. First, the E rule mentioned

[8] Keep in mind that $\emptyset \in n_O(h)$ and $n_O(h) = \emptyset$ are two very different conditions. Here the first would boil down to a violation of the "ought implies can" principle. It makes the impossible action type obligatory. In the second case *nothing* is obligatory.

earlier for P in MDL is derivable here as well, both for P, in the exact same way as above, and for O. The latter follows from necessitation for $\Box$ and (Univ). Maybe more surprisingly, K for O is derivable in this logic! Model-theoretically, this is a consequence of the uniqueness of obligations in that system. If both $\varphi \to \psi$ and φ are obligatory, then by Observation 1 they must have the same extension. Axiomatically, two applications of WP make $\Box((\varphi \to \psi) \leftrightarrow \varphi)$ derivable. But then a few steps of normal modal and propositional reasoning with $\Box$ gives us $\Box(\varphi \leftrightarrow \psi)$, from which one application of E for O outputs $O\psi$. Now deontic logicians might fear that paradoxes loom again in the presence of K. But the undesirable consequences of that theorem are limited by the absence of necessitation for O in that logic. None of Ross or the Contrary to Duty paradoxes hold here. See [12] for the details.

Rational recommendations in games work in OWP-models essentially as in van Benthem's minimal deontic logic. Take again our running example. Ann's obligations are defined exactly as before. Taking H to be the set of all pairs $(n, 100)$, $n_O(h)$ just contains the set $br(Ann, 100)$, for all h, and each $n_P(h)$ is defined by taking the closure under subsets of $br(Ann, 100)$. This gives us essentially the same result as before: Ann ought to play something less that 100, and any sub-type of that, for instance playing less than 10, or just 1, is permitted.

(Univ)	$\Box(\varphi \leftrightarrow \psi) \to (D\varphi \leftrightarrow D\psi)$	(WP)	$(O\varphi \wedge P\psi) \to \Box(\psi \to \varphi)$
(O-P)	$O\varphi \to P\varphi$	(Flip)	$\dfrac{\varphi \to \psi}{P\psi \to P\varphi}$
(O-Can)	$O\varphi \to \Diamond\varphi$		

Table 2. The sound and complete axiom system for OWP. All propositional tautologies, as well as the S5 axioms for $\Diamond$, are also assumed here. In (Univ) D is either O or P.

2.5 Comparison

Our first two logics for rational recommendations in games thus have many things in common. They share the main substantive principles for which we argued for in Section 1.2: O and P are not dual, P validates Free Choice by (Flip), and we have (WP) in both logics. Furthermore obligations validate K in both systems.

Some of the axiomatic divergence between the two systems reflect minute frame-theoretic differences that can be easily accommodated. There is for instance no "ought implies can" principle in MDL, because the semantics allow for "blind" atomic actions, i.e. actions from which no normatively ideal actions can be reached. But the principle can be added, forcing us into the class of serial MDL frames. On the other hand, both the obligation and the permission neighborhoods can be empty in OWP models, which explains the invalidity of

$\Box\neg\varphi \to P\varphi$ in that logic. If, however, the permission neighborhood is not empty, then the principle holds. This is reflected by the following theorem of OWP:

$$P\psi \to (\Box\neg\varphi \to P\varphi)$$

The derivation starts by factoring ψ into its logically equivalent $(\psi\wedge\varphi)\vee(\psi\wedge\neg\varphi)$. Using the derivable E rule and (Flip), one gets $P(\psi \wedge \varphi)$. But then since $\neg\varphi$ propositionally implies $(\psi \wedge \varphi) \leftrightarrow \varphi$, a standard bit of normal modal reasoning with $\Box$, and one last application of E gives us $P\varphi$, as required.

A similar argument explains the absence of (OR) in OWP. The principle is valid only in the class of OWP frames where the *obligation* neighborhood is not empty. Note that in OWP-frames this might happen even if the permission neighborhood isn't. This is mirrored, once again, by the following prefixed version of (Univ), which is also a theorem of OWP.[9]

$$O\chi \to (P\varphi \wedge P\psi \to P(\varphi \vee \psi))$$

This time the derivation starts with applying (WP) to $O\chi$ and $P\varphi$, to deliver $\Box(\varphi \to \chi)$, and similarly for ψ. Then after some steps of propositional and normal modal reasoning we get $\Box((\chi \wedge (\varphi \vee \psi)) \leftrightarrow (\varphi \vee \psi))$. The proof finishes by using (O-P) on $O\chi$ again, and working our way to $P(\varphi \vee \psi)$ using the same factorization and flipping routine as above.

The main point of divergence between the two systems, how they handle obligations, rests on the apparently innocuous (O-P) principle:

$$O\varphi \to P\varphi$$

This principle is not valid in van Benthem's MDL while, in OWP, it nails down the uniqueness of obligations.

The absence of (O-P) in MDL requires giving up some old thinking habits from Standard Deontic Logic. Consider again the example of Ann's best response to Bob playing 100 in Guess 2/3 of the Average. The result of the construction sketched in Section 1.1 is that Ann ought not only to "play any number lower than 100." She also ought to play any action type that is logically weaker than "playing any number lower than 100". So in particular the trivial action type $\top$ is rationally required of her. She ought to play *a* number, whatever that number is, simply because if she doesn't play any number then she won't play any best response number. But playing whatever number is *not* permitted for Ann, despite the fact that she is rationally required to do this. This is a particular case where obligation does not imply permission in MDL. This might feel counter-intuitive to the reader, probably because of the ease with which we have learned to derive permissions from obligations in Standard Deontic Logic. Against this one should keep in mind the interpretation of obligations and permissions in MDL as necessary and sufficient conditions for rationality. Necessary conditions need not to be sufficient, of course, so O should not imply P in that interpretation.

[9] We are grateful to Frederik van de Putte for drawing our attention to this fact.

In fact adding this principle to MDL results in the same deontic trivialization as when Free Choice is added to SDL (c.f. footnote 7 on p. 6). Everything gets permitted. Necessitation for O gives $O\top$, which then with (O-P) and (Flip) yields $P\varphi$ for any φ whatsoever, as anything implies the tautology. OWP avoids this trivialization because obligations are not closed under logical consequences. Although it satisfies K, this logic invalidates the so-called inheritance rule:

$$\frac{\varphi \rightarrow \psi}{O\varphi \rightarrow O\psi}$$

This can be illustrated again in our running example. Ann's best response to Bob playing 100 is to play any number lower than his. This logically implies that she plays *a* number. But, unlike in MDL, here it doesn't follow that she ought to play any number as well. What Ann ought to do, here, is *only* to play a best response. This is a direct consequence of the interplay between (O-P) and (WP).

More generally, by accepting that obligation implies permission, OWP is committed to the view that obligations pinpoint necessary *and* sufficient conditions for rationality. Hence the uniqueness of obligations, up to coextensionality. As mentioned earlier, the only types of strategy agents ought to play in that logic are rational strategies.

MDL and OWP do overlap, but precisely in the trivial cases where nothing but the trivial action type $\top$ is obligatory, and hence everything is permitted.[10] Indeed, any a MDL-model where $R_D = W \times W$ can be turned into an OWP model, by taking $n_O(w) = \{W\}$ for all w, and $n_P(w)$ the full power set of W, and conversely for starting from such a OWP-model. It should be clear that the two will satisfy exactly the same formulas. The converse is also true. For any MDL-frame where R_D is not the universal relation, taking the set of accessible atomic actions at each w to construct $n_O(w)$ will yield divergent obligations in OWP and MDL, at some w.

Let us summarize the findings of this first comparison. The most important point of agreement between MDL and OWP, and the way in which they differ most from Standard Deontic Logic, is that rational obligation and permission provide necessary and sufficient conditions for rational play in games. This is witnessed by their acceptance of the (Flip) rule and the (WP) axiom. Taking this view on board, however, raises a dilemma. On the one hand one can stay as close as possible to SRP and WRO. Then one is forced, on pain of trivialization, to give up the familiar "O implies P" principle. This leads to situations like Ann's described above, where a player ought to play an action type that this not rationally permitted. On the other hand, if one chooses to keep the implication from obligation to permission, then one restricts the former to one particular necessary condition for rationality, namely the necessary *and* sufficient one, letting go of the "only if" direction of WRO. In short, the main axiomatic difference between MDL and OWP reflects a difference in philosophical commitment, to

[10] Note that with (Flip) or Free Choice Permission in the system, $\top$ being the only obligatory action and everything being permitted are just two sides of the same coin.

the familiar "obligation implies permission" principle and to the main features of logics of rational recommendations in games.

3 Deontic Action Logics

We will now look a richer language to describe rational recommendations in games, Deontic Boolean Action Logic (DBAL), proposed in [9]. This logic differs from MDL and OWP in that it draws a sharper distinction between the Boolean construction of action types and the Boolean connectives applied to obligations and permissions. Unlike in the previous section, where we mainly looked at axiomatic differences between MDL and OWP, here the main contributions will be translation results, showing that DBAL is embeddable in MDL.

3.1 Deontic Boolean Action Logic

By interpreting the points in MDL and OWP models as actions, atomic and complex formulas in $\mathcal{L}$ were naturally interpreted as action types, constructed much as in Propositional Dynamic Logic (PDL). The logic we consider now explicitly uses a language similar to Boolean Modal Logic (BML), and by doing so distinguishes action type constructors and standard Boolean connectives for formulas. The result is the following two-sorted language:

Definition 6. *The language $\mathcal{L}^*$ for Deontic Boolean Action Logic (DBAL) is defined as follows:*

$$\varphi := \alpha \doteq \alpha \mid \neg\varphi \mid \varphi \to \varphi \mid P\alpha \mid F\alpha$$
$$\alpha := a \mid \mathbf{1} \mid \overline{\alpha} \mid \alpha \cup \alpha \mid \alpha \cap \alpha$$

where a is one of the elements of a finite set Act_0 of action generators.

Action types $\alpha \in Act$ are thus constructed out of primitive action types, taken from a given set of action generators Act_0, the trivial $\mathbf{1}$ action, and the usual BML connectives of complementation $\overline{\alpha}$, nondeterministic choice $\cup$ and parallel execution $\cap$. The impossible action $\mathbf{0}$ is defined as $\overline{\mathbf{1}}$.

The language itself has, strictly speaking, three types of atomic sentence. First, the language can describe equivalence of action types, using the $\doteq$ connective. Second comes the deontic statements, stating that certain action types are (rationally) permitted or forbidden. An obligation to do α, in that language, is defined as the complement $\overline{\alpha}$ of that action type being prohibited.

This syntax reveals the main difference between DBAL and the two systems presented in the previous section. In DBAL only action types are allowed within the scope of the deontic modalities. So we are still in the realm of logics for "ought to do." Here, however, no deontic modalities, or boolean connectives for that matter, can occur in the scope of an F or a P. Formulas of the form $PF\alpha$ or $F\neg\alpha$ are not well-formed here. The deontic modalities function as properties of action types, and deontic statements that such and such type is forbidden or permitted are (structured) atomic sentences in this language. The Boolean connectives $\neg$, $\rightarrow$, etc., are here to form complex statements about equivalence and deontic properties of action types.

The semantics offer tools to interpret the construction of complex action types and their deontic properties. Although there is a syntactic distinction between the action-theoretic and the rest of the language here, in the model we only work with atomic actions and deontic properties.

Definition 7. *A DBAL model M is a tuple $\langle E, LEG, ILL, I\rangle$ where:*

1. *E is a non-empty set of atomic actions*
2. *LEG and ILL are subsets of 2^E such that*
   (a) $X \in LEG(ILL)$ *and* $Y \subseteq X \Rightarrow Y \in LEG(ILL)$
   (b) $X \in LEG(ILL)$ *and* $Y \in LEG(ILL) \Rightarrow X \cup Y \in LEG(ILL)$
   (c) $LEG \cap ILL = \{\emptyset\}$
3. *$I : Act_0 \rightarrow \wp(E)$ is an interpretation function assigning to each action generator to a subset of actions, i.e., $I(a) \subseteq E$ where $a \in Act_0$*

Reflecting the syntactic primitives, DBAL models come equipped with two sets of action types, the legal and the illegal one. In its LEG instance, 2a gives us the by now familiar downward closure of the permissions, corresponding to Free Choice. For ILL this condition gives us, given the definition of obligation, the supersets closure condition that we have already encountered in MDL. To spell this out in detail we need first to extend the interpretation function to arbitrary action types.

- $I(\mathbf{1}) = E$
- $I(\overline{\alpha}) = E - I(\alpha)$
- $I(\alpha \cup \beta) = I(\alpha) \cup I(\beta)$
- $I(\alpha \cap \beta) = I(\alpha) \cap I(\beta)$

With this in hand the truth conditions for formulas of $\mathcal{L}^*$ become the following. Observe that the Boolean connectives for formulas are interpreted globally, instead of locally.

$$M \models \alpha \doteq \beta \text{ iff } I(\alpha) = I(\beta)$$
$$M \models \neg\varphi \text{ iff } M \not\models \varphi$$
$$M \models \varphi \wedge \psi \text{ iff } M \models \varphi \;\&\; M \models \psi$$
$$M \models P\alpha \text{ iff } I(\alpha) \in LEG$$
$$M \models F\alpha \text{ iff } I(\alpha) \in ILL$$

The sound and complete axiomatization of DBAL is very close to the one for van Benthem's MDL, modulo the additional apparatus to describe equivalence of action types, and the fact that prohibitions instead of obligations are primitive. See Table 3. One indeed recognizes Free Choice and its converse (OR) in the biconditional (OR-P). Observe that the following valid analogous to (WP) is expressible in this system:

$$O\alpha \wedge P\beta \rightarrow (\beta \cap \alpha) \doteq \beta$$

(OR-P) $P(\alpha \cup \beta) \leftrightarrow P\alpha \wedge P\beta$	(Incl) $P\alpha \wedge F\beta \rightarrow \alpha \cap \beta \doteq \mathbf{0}$
(OR-F) $F(\alpha \cup \beta) \leftrightarrow F\alpha \wedge F\beta$	(BA) $\dfrac{\alpha \doteq \beta, \varphi[\alpha]}{\varphi[\alpha/\beta]}$
($\mathbf{0}$-P/F) $\alpha \doteq \mathbf{0} \leftrightarrow F\alpha \wedge P\alpha$	

Table 3. The sound and complete axiom system for Deontic Boolean Action Logic. All propositional tautologies are also taken as axioms, as well as the usual axioms for Boolean algebras for boolean action terms α. Above $\varphi[\alpha/\beta]$ indicates to replace all occurrences of α in φ by β.

3.2 Equivalence of level-1 MDL, and DBAL

We now show that a DBAL corresponds to a simple fragment of MDL. They are inter-translatable. We start with the syntactic translation, from DBAL to MDL. Let Act_0 be our given set of action generators. We first create one atomic proposition per generator. So set $Prop_{Act_0} = \{p_a : a \in Act_0\}$. Call $\mathcal{L}_{Act_0}$ the language $\mathcal{L}$ generated by the rule on page 7 using $Prop_{Act_0}$. The translation function $T : \mathcal{L}^* \rightarrow \mathcal{L}_{Act_0}$ is then defined as follows:

1. $T(a) = p_a$
2. $T(\mathbf{1}) = \top$
3. $T(\overline{\alpha}) = \neg T(\alpha)$
4. $T(\alpha \cup \beta) = T(\alpha) \vee T(\beta)$
5. $T(\alpha \cap \beta) = T(\alpha) \wedge T(\beta)$
6. $T(\alpha \doteq \beta) = \Box(T(\alpha) \leftrightarrow T(\beta))$
7. $T(\neg\varphi) = \Box(\neg T(\varphi))$
8. $T(\varphi \rightarrow \psi) = \Box(T(\varphi) \rightarrow T(\psi))$
9. $T(P\alpha) = PT(\alpha)$
10. $T(F\alpha) = O\neg T(\alpha)$

Now let $M = \langle E, LEG, ILL, I\rangle$ be a DBAL model. We construct a uniform MDL model $M^T = \langle W, R_D, V\rangle$ as follows. The only subtle matter is in the construction of R_D, as Act_0 is finite but E needs not be, not all subsets of E need to be definable. We use definable ones here:

- $W = E$
- For all $w \in W$, $R_D[w] = I(\alpha)$ for α such that $M \vDash P\alpha$ and for all β such that $M \vDash P\beta$, $I(\beta) \subseteq I(\alpha)$.
- $V(p_a) = I(a)$

It should be clear that M^T is a MDL model.

Observation 2 *Let $M = \langle E, LEG, ILL, I\rangle$ be a DBAL model. Then M^T is a MDL model.*

Proof. The only non-trivial condition is the one for R_D. We need to check that the set $R_D[w]$ is well-defined. This follows directly from finiteness of Act_0 and the fact that if $P\alpha$ and $P\beta$ are true, then $P(\alpha \cup \beta)$ is as well.

With this in hand we can show the main result of this section, namely that the translation T is truth preserving.

Theorem 1. *For all formula φ of $\mathcal{L}^*$, DBAL model M, and state $w \in W$ for M^T:*

$$M \models \varphi \text{ iff } M^T, w \models T(\varphi)$$

Proof. The proof is by induction on the complexity of $\varphi \in \mathcal{L}^*$, which in turns requires sub-inductions on action types for atomic formulas. The former is straightforward. We focus on the latter cases.

There are three cases to consider. First $\varphi = \alpha \doteq \beta$. Suppose that $\alpha, \beta \in Act$. Then $M \models \alpha \doteq \beta$ iff $I(\alpha) = I(\beta)$ iff $V(p_\alpha) = V(p_\beta)$, by construction. But the latter happens iff $M^T, w \models \Box(T(\alpha) \leftrightarrow T(\beta))$. The inductive steps follow similarly, using the inductive hypothesis to go from $I(\alpha)$ and $I(\beta)$, for complex α and β, to the truth sets $\|T(\alpha)\|$ and $\|T(\beta)\|$, respectively.

Now consider the case where $\varphi = P\alpha$. Suppose first that $\alpha \in Act$. Then $M \models P\alpha$ iff $I(\alpha) \in LEG$. Now the key observation, which follows directly from our construction, is that the latter happens iff $V(p_a) \subseteq R_D(w)$ in M^T, for any w, and similarly for arbitrary action type α.

Finally, suppose that $\varphi = F\alpha$. Start with $\alpha \in Act$. We have that $M \models F\alpha$ iff $I(\alpha) \in ILL$. Now we know that $I(\alpha) \cap X = \emptyset$ for all $X \in LEG$. But by our construction of M^T this means that $R_D[w] \in LEG$, so $R_D[w] \cap V(p_\alpha) = \emptyset$, so $R_D[w] \subseteq W - V(p_\alpha) = \|\neg T(\alpha)\|$, so $M, w \models O\neg T(\alpha)$. The inductive step proceeds similarly.

Observe that the syntactic translation maps formulas of $\mathcal{L}^*$ into a simple fragment of MDL. First of all, the fragment without embedded deontic operators, simply by the syntactic restrictions on DBAL. No $\Box$ or $\Diamond$ occur in the scope of a deontic operator either. Furthermore, no atomic proposition occurs "free" in the translated formulas. They are always in the scope of either a universal modality, when translated from atoms $\alpha \doteq \beta$, or a deontic operator. Finally, since Act_0 is finite, this translation only uses finitely many atomic proposition. Call this simple fragment $\mathcal{L}_1$. We show now that $\mathcal{L}_1$ can be translated back into DBAL. The translation goes in two steps, first for action types and then for arbitrary formulas.

Let $\mathcal{L}_0$ be a fragment of MDL where no modal operator ($\Box$, O or P) occurs, defined over a given finite set of atomic propositions $Prop$. Define $Act_{Prop} = \{a_p : p \in Prop\}$. Then the action-translation $\tau : \mathcal{L}_0 \to \mathcal{L}$ is defined as follows:

$$\tau(p) = a_p$$
$$\tau(\neg\varphi) = \overline{\tau(\varphi)}$$
$$\tau(\varphi \land \psi) = \tau(\varphi) \cap \tau(\psi)$$

Now we are ready to define our translation from $\mathcal{L}_1$ to $\mathcal{L}^*$. Recall that no atomic proposition occurs "free" in $\mathcal{L}_1$. The translation $\rho : \mathcal{L}_1 \to \mathcal{L}^*$ is thus only defined for complex formula, and ultimately resorts on the action-translation τ

just defined.

$$\rho(\neg\varphi) = \neg\rho(\varphi)$$
$$\rho(\varphi \wedge \psi) = \rho(\varphi) \wedge \rho(\psi)$$
$$\rho(\Diamond\varphi) = \neg(\tau(\varphi) \doteq \mathbf{0})$$
$$\rho(P\varphi) = P(\tau(\varphi))$$
$$\rho(O\varphi) = F(\overline{\tau(\varphi)})$$

Now we show how to transform MDL models into DBAL ones. The construction works locally, generating one DBAL model for each point w in the original deontic model. Of course, in the special case of uniform MDL models, where the set of normatively ideal or rational atomic actions is the same at all w, this is not necessary.

Let $M = \langle W, R_D, V\rangle$ be a MDL model. We construct a model $M^{w^*} = \langle E, LEG, ILL, I\rangle$ with to a particular point $w^* \in W$ as follows:

1. $E = W$
2. $LEG = \{X \mid X \subseteq R_D[w^*]\}$
3. $ILL = \{X \mid R_D[w^*] \subseteq \overline{X}\}$
4. $I(\tau(p)) = \|p\|$ for each $p \in Prop_0$ in DAL-language

The resulting model is indeed a DBAL model.

Observation 3 *The model $M^{w^*} = \langle E, LEG, ILL, I\rangle$ w.r.t. particular point $w^* \in W$ constructed before is a DBAL model.*

Proof. The verification of LEG is easy to see. So we only verify those three conditions for ILL.

1. Suppose $X \in ILL$ and $Y \subseteq X$. Then we have $R_D[w^*] \subseteq \overline{X}$ and $\overline{X} \subseteq \overline{Y}$. So $R_D[w^*] \subseteq \overline{Y}$. So $Y \in ILL$.
2. Suppose $X \in ILL$ and $Y \in ILL$. Then $R_D[w^*] \subseteq \overline{X}$ and $R_D[w^*] \subseteq \overline{Y}$. Then it implies $R_D[w^*] \subseteq \overline{X} \cap \overline{Y}$. That is $R_D[w^*] \subseteq \overline{X \cup Y}$. Thus $X \cup Y \in ILL$.
3. If not, then there is some $X \in LEG$ and $X \in ILL$ such that $X \neq \emptyset$. So $X \subseteq R_D[w^*]$ and $R_D[w^*] \subseteq \overline{X}$. It follows $X \subseteq \overline{X}$. But that's not possible because $X \neq \emptyset$. So, $LEG \cap ILL = \{\emptyset\}$.

With this in hand we can now show the main result of this section:

Theorem 2. *Let $M = \langle W, R_D, V\rangle$ be a MDL model. Then for all formulas $\psi \in \mathcal{L}_0$ and $\varphi \in \mathcal{L}_1$,*

1. *$M, w \vDash \psi$ iff $w \in I(\tau(\psi))$ for every $w \in W$, i.e., $\|\psi\| = I(\tau(\psi))$.*
2. *$M, w^* \vDash \varphi$ iff $M^{w^*} \vDash \rho(\varphi)$*

Proof. 1 follows directly from the construction of M^{w^*} and the definition of the translation function τ. For 2 we only show the cases of $\Diamond\varphi, P\varphi$ and $O\varphi$. The cases for the Boolean connectives follow directly.

1. Case $\Diamond\varphi$: Suppose $M, w^* \vDash \Diamond\varphi$. Then there is some $w' \in W$ such that $M, w' \vDash \varphi$. By the first result and the construction, $w' \in I(\tau(\varphi))$ where $w' \in E$. That means $M^{w^*} \vDash \neg(\tau(\varphi) \doteq \mathbf{0})$.
   On the other hand, suppose $M^{w^*} \vDash \neg(\tau(\varphi) \doteq \mathbf{0})$. So there is some $w' \in I(\tau(\varphi))$. By the first result, $M, w' \vDash \varphi$. That is $M, w \vDash \Diamond\varphi$.
2. Case $P\varphi$: Suppose $M, w^* \vDash P\varphi$. That is $\|\varphi\| \subseteq R_D[w^*]$. Then $I(\tau(\varphi)) \in LEG$ by the first result and the construction. Thus $M^{w^*} \vDash P(\tau(\varphi))$.
   On the other hand, suppose $M^{w^*} \vDash P(\tau(\varphi))$. That is to say, $I(\tau(\varphi)) \in LEG$. By the construction, it implies $\|\varphi\| \subseteq R_D[w^*]$. So we have $M, w^* \vDash P\varphi$.
3. Case $O\varphi$: Suppose $M, w^* \vDash O\varphi$. That is $R_D[w^*] \subseteq \|\varphi\|$. Then $\overline{I(\tau(\varphi))} \in ILL$ by the first result and the construction. Namely $I(\overline{\tau(\varphi)}) \in ILL$. So we have $M, w^* \vDash F(\overline{\tau(\varphi)})$.
   On the other hand, suppose $M, w^* \vDash F(\overline{\tau(\varphi)})$. It implies that $\overline{I(\tau(\varphi))} \in ILL$. So $R_D[w^*] \subseteq \|\varphi\|$ by the first result and the construction. Thus $M, w^* \vDash O\varphi$.

Put together, these two translation results show that DBAL can be embedded in MDL. The upshot of the comparison between MDL and OWP was that the two logics embodied different philosophical commitments regarding the structure of rational obligations and permissions. This is not the case here. The main difference between MDL and DBAL is that the latter doesn't allow for embedding deontic modalities.

4 Conclusion

The goal of this paper was to provide an explicit comparison between three related logics which, we argued, are well-suited to study rational recommendations in games: van Benthem's "Minimal Deontic Logic", Anglberger et al's "Obligation as Weakest Permission", and Trypuz and Kulicki's "Deontic Boolean Action Logic." All three systems can be seen as endorsing the idea that obligations and permissions in games provide necessary and sufficient conditions for rationality. We argued that the first two differ on their view on the relation between obligations and permissions: on OWP the former imply the latter, but not in MDL. As a result of this, we argued, MDL stays closer to the core philosophical principles for rational recommendations in games, while OWP stays closer to the intuitive idea that an action type cannot be obligatory while at the same time not being permitted. We then showed that third logic can be embedded in MDL. The upshot of that result is that even though the former draws a sharper syntactic distinction between propositional connectives and constructors for complex action types, this distinction is blurred again at the semantic level. DBAL can be seen as a syntactically restricted MDL, but this needs not to be conceptually implausible. Given the action-theoretic interpretation we used for the semantics of MDL, formulas of the form $OP\varphi$ means that the "action type" $P\varphi$ is obligatory . We leave it to the reader to decide whether $P\varphi$ can plausibly support this action-theoretic reading.

The comparison in this paper left out a number of other deontic systems that share the core principles identified in Section 3, the most salient being [4, 6]. These logical systems go one step further than DBAL in distinguishing action-theoretic and Boolean connectives. They do so by introducing richer semantic structures, much as in propositional dynamic logic or labeled transition systems. This is certainly desirable, as it can lead to fruitful analysis of dynamic solutions of games, for instance to the computation of Nash equilibrium in our Guess 2/3 of the Average example. At the formal level, however, the richer syntax and semantics have so far blocked our attempts at an insightful translation and/or axiomatic comparison between these systems and the ones studied here. The main point there, which is also the final take home message of this paper, is that because they embody the same core principle relating obligations and permissions to necessary and sufficient conditions for rationality, they are part of larger family of logical systems that can be put to use in the analysis of what players ought to do in strategic interaction.

References

1. N. G. A.J.J. Anglberger and O. Roy. Obligation, free choice and the logic of weakest permission. *Review of Symbolic Logic*, 2015.
2. A. Anglberger, H. . Dong, and O. Roy. Open reading without free choice. In *DEON'14*, 2014.
3. P. Blackburn, M. De Rijke, and Y. Venema. *Modal Logic*, volume 53. Cambridge University Press, 2002.
4. P. Castro and T. Maibaum. A complete and compact propositional deontic logic. In C. Jones, Z. Liu, and J. Woodcock, editors, *Theoretical Aspects of Computing ICTAC 2007*, volume 4711 of *Lecture Notes in Computer Science*, pages 109–123. Springer Berlin Heidelberg, 2007.
5. R. M. Chisholm. Contrary-to-duty imperatives and deontic logic. *Analysis*, pages 33–36, 1963.
6. J. Czelakowski. Action and deontology. In *Freedom and Enforcement in Action*, pages 143–193. Springer, 2015.
7. S. Hansson. The varieties of permissions. In D. Gabbay, J. Horty, X. Parent, R. van der Meyden, and L. van der Torre, editors, *Handbook of Deontic Logic and Normative Systems*, volume 1. College Publication, 2013.
8. J. Joyce. Bayesianism. In Piers Rawling and Alfred R. Mele, editors, *The Oxford Handbook of Rationality*, pages 132–155. Oxford: Oxford University Press, 2004.
9. P. Kulicki and R. Trypuz. On deontic action logics based on boolean algebra. *Journal of Logic and Computation*, 2013.
10. P. McNamara. Deontic logic. In E. N. Zalta, editor, *Standford Encyclopedia of Philosophy*. 2006. Spring Edition.
11. E. Pacuit and O. Roy. Epistemic foundations of game theory. *Stanford encyclopedia of philosophy*, 2015.
12. O. Roy, A. Anglberger, and N. Gratzl. The logic of obligation as weakest permission. *Deontic Logic in Computer Science*, page 139–150, 2012.
13. O. Roy, A. Anglberger, and N. Gratzl. The logic of best action from a deontic perspective. In A. Baltag and S. Smets, editors, *Johan FAK van Benthem on Logical and Informational Dynamics*. Springer, 2014.

14. R. Trypuz and P. Kulicki. A systematics of deontic action logics based on boolean algebra. *Logic and Logical Philosophy*, 18(3-4):253–270, 2010.
15. J. Van Benthem. Minimal deontic logics. *Bulletin of the Section of Logic*, 8(1):36–42, 1979.

Appendix

Here are the syntactic proofs sketched in the paper.

Observation 4 *In MDL, the necessitation rule holds for O, i.e.,* $\dfrac{\varphi}{O\varphi}$.

Proof. Suppose φ is a theorem of MDL. Then by (NEC) we have $\Box\varphi$. So $O\varphi$ is derivable fro $\Box\varphi$ by using (Incl) and Modus Ponens.

Observation 5 *In MDL, Free choice, K axiom, and E rule hold for P, i.e.,*

1. $P(\varphi \vee \psi) \rightarrow (P\varphi \wedge P\psi)$,
2. $P(\varphi \rightarrow \psi) \rightarrow (P\varphi \rightarrow P\psi)$, *and*
3. $\dfrac{\varphi \leftrightarrow \psi}{P\varphi \leftrightarrow P\psi}$.

Proof. 1. Because $\varphi \vee \psi$ is derivable from both φ and ψ, we have $P(\varphi \vee \psi) \rightarrow P\varphi$ and $P(\varphi \vee \psi) \rightarrow P\psi$ by applying (Flip) twice. So $P(\varphi \vee \psi) \rightarrow (P\varphi \wedge P\psi)$ by using propositional logic.
2. Suppose $P(\varphi \rightarrow \psi)$ and $P\varphi$ are theorems of MDL. Then $P((\varphi \rightarrow \psi) \vee \varphi)$ is followed by (OR). On the other hand, we know that $\psi \rightarrow ((\varphi \rightarrow \psi) \vee \varphi)$ is a tautology by propositional logic. Then $P((\varphi \rightarrow \psi) \vee \varphi) \rightarrow P\psi$ is derivable by using (Flip). Thus, we have $P\psi$ by Modus Ponens.
3. Suppose $\varphi \leftrightarrow \psi$ is a theorem of MDL. So $\varphi \rightarrow \psi$ and $\psi \rightarrow \varphi$ are both theorems, as the knowledge of propositional logic. Then, by using (Flip) twice, we have $P\psi \rightarrow P\varphi$ and $P\varphi \rightarrow P\psi$. So $P\varphi \leftrightarrow P\psi$ is derivable by propositional logic.

Observation 6 *In OWP, E rule holds for P and O, i.e.,*

1. $\dfrac{\varphi \leftrightarrow \psi}{P\varphi \leftrightarrow P\psi}$, *and*
2. $\dfrac{\varphi \leftrightarrow \psi}{O\varphi \leftrightarrow O\psi}$.

Proof. 1. Suppose $\varphi \leftrightarrow \psi$ is a theorem of OWP. So, by using (Flip) twice, we have $P\varphi \leftrightarrow P\psi$.
2. Suppose $\varphi \leftrightarrow \psi$ is a theorem of OWP. By necessitation for $\Box$, we have $\Box(\varphi \leftrightarrow \psi)$. So, by using (Univ), it follows $O\varphi \leftrightarrow O\psi$.

Observation 7 *In OWP, the K axiom holds for O, i.e.,* $O(\varphi \rightarrow \psi) \rightarrow (O\varphi \rightarrow O\psi)$.

Proof. Suppose $O(\varphi \to \psi)$ and $O\varphi$ are theorems of OWP. Then, by applying (O-P) on $O\varphi$, it follows $P\varphi$. So, we have $\Box(\varphi \to (\varphi \to \psi))$ by applying (WP) on $O(\varphi \to \psi)$ and $P\varphi$. We have $\Box((\varphi \to \psi) \to \varphi)$ in a similar way as previously. By the axiom for $\Box$, it follows $\Box((\varphi \to \psi) \leftrightarrow \varphi)$. In addition, $(\varphi \to \psi) \leftrightarrow \varphi$ and $\varphi \leftrightarrow \psi$ are logically equivalent. So we know that $\Box(\varphi \leftrightarrow \psi)$ by using E rule for $\Box$. Then by using (Univ) it follows $O\varphi \leftrightarrow O\psi$. Thus $O\psi$ is derived as $O\varphi$ is a theorem of OWP.

Observation 8 $P\psi \to (\Box\neg\varphi \to P\varphi)$ *is a theorem of OWP.*

Proof. Suppose $P\psi$ and $\Box\neg\varphi$ are theorems of OWP. On the one hand, we know ψ is logically equivalent to $(\psi \wedge \varphi) \vee (\psi \wedge \neg\varphi)$. So, by using the E rule for P and (Flip), we have $P(\varphi \wedge \psi)$. On the other hand, $\neg\varphi$ logically implies $(\psi \wedge \varphi) \leftrightarrow \varphi$. By applying the K axiom and (NEC) for $\Box$ on $\Box\neg\varphi$, it follows $\Box((\psi \wedge \varphi) \leftrightarrow \varphi)$. By (Univ) and $P(\psi \wedge \varphi)$, we have $P\varphi$ as theorem of OWP.

Observation 9 $O\chi \to ((P\varphi \wedge P\psi) \to P(\varphi \vee \psi)$ *is a theorem of OWP.*

Proof. Suppose $O\chi$ and $P\varphi \wedge P\psi$ are theorems of OWP. Applying (WP) to $O\chi$ and $P\varphi$, it follows $\Box(\varphi \to \chi)$. Similarly, we also have $\Box(\psi \to \chi)$. So we have $\Box(\varphi \vee \psi \to \chi)$ as a result of normal modal logic for $\Box$. By using normal modal logic again, it follows that $\Box((\chi \wedge (\varphi \vee \psi)) \leftrightarrow (\varphi \vee \psi))$. By using (Univ), we get $P(\chi \wedge (\varphi \vee \psi)) \leftrightarrow P(\varphi \vee \psi)$. Applying (O-P) on $O\chi$, we have $P\chi$. By the propositional logic, E rule for P, and (Flip), it follows that $P(\chi \wedge (\varphi \vee \psi))$. Thus we get $P(\varphi \vee \psi)$.

Observation 10 $P\top$ *is derivable after adding up (O-P) into MDL.*

Proof. By the necessitation rule for O in MDL, we have $O\top$. Then $P\top$ is followed by (O-P). On the other hand, by using (Flip), $P\top \to P\varphi$ for any φ. Thus, $P\varphi$ is derivable, after adding up (O-P) into MDL.