

## Expressing permission\*

William B. Starr  
Cornell University

**Abstract** This paper proposes a semantics for free choice permission that explains both the non-classical behavior of modals and disjunction in sentences used to grant permission, and their classical behavior under negation. It also explains why permissions can expire when new information comes in and why free choice arises even when modals scope *under* disjunction. On the proposed approach, deontic modals update preference orderings, and connectives operate on these updates rather than propositions. The success of this approach stems from its capacity to capture the difference between *expressing* the preferences that give rise to permissions and conveying propositions about those preferences.

**Keywords:** Free choice permission, modality, dynamic semantics, expressive meaning

### 1 Free choice permission and allied phenomena

While I will focus on deontic modals here, free choice effects arise in non-deontic and even non-modal contexts too (Fox 2007). As I will discuss in §3, once my analysis is stated it will be possible to articulate structural parallels across this wider range of data. In this section I will lay out the data to be explained, adding a few novel observations and examples, and saying why existing analyses are not fully satisfactory. §2 presents a new dynamic analysis, fully formalized in Appendix A.

#### 1.1 Free choices, hard choices

Using ‘ $\Rightarrow$ ’ and ‘implication’ to neutrally describe inferences that may be semantic or pragmatic in nature, the basic problem of free choice permission centers on three implications that I will call *Narrow Free Choice*, *Wide Free Choice* and *Double Prohibition*. For context, envision a perfectly informed labor representative  $X$  telling her constituents how to vote in an election. If  $X$  says (1a),  $X$  is intuitively committed to (1b) (Kamp 1973; von Wright 1968: 4–5).

---

\* Conversations with Chris Barker got me hooked on this topic, and helped set the empirical landscape for me. I thank Melissa Fusco, Hans Kamp, Sarah Murray, Jacopo Romoli, Malte Willer, the Cornell Semantics Group, six anonymous SALT reviewers and numerous SALT participants for very valuable feedback. Thanks also to Todd Snider and Jacob Collard for their help in preparing the paper.

**Narrow Free Choice (NFC)**

$$\text{May}(A \vee B) \Rightarrow \text{May}A \wedge \text{May}B$$

- (1) a. Members may vote for Anderson or Brady.  
 b. Members may vote for Anderson and members may vote for Brady.

Quite curiously, this implication also arises when *may* scopes under *or* (Kamp 1978: 273, Zimmermann 2000, Geurts 2005, Simons 2005).

**Wide Free Choice (WFC)**

$$\text{May}A \vee \text{May}B \Rightarrow \text{May}A \wedge \text{May}B$$

- (2) a. Members may vote for Anderson or you may vote Brady.  
 b. Members may vote for Anderson and you may vote for Brady.

Neither **NFC** nor **WFC** are valid in standard modal logic when *may* is treated as a possibility modal and these implications do not meet the standard cancellation test for implicatures (Simons 2005; Barker 2010). It is not felicitous for *X* to follow up her disjunctive permission statement by denying one of the disjuncts.

- (3) a. Members may vote for Anderson or Brady.  
 b. #But members may not vote for  $\left\{ \begin{array}{l} \text{Anderson} \\ \text{Brady} \end{array} \right\}$ .

As is appropriate for a political context, free choice implications can be ‘defeated’ by ignorance (Kamp 1978: 271) or uncooperativeness (Simons 2005: 273).

- (4) a. Members may vote for Anderson or Brady, but I don’t know which.  
 b. # Members may vote for  $\left\{ \begin{array}{l} \text{Anderson} \\ \text{Brady} \end{array} \right\}$ .
- (5) a. Members may vote for Anderson or Brady, but I won’t tell you which.  
 b. # Members may vote for  $\left\{ \begin{array}{l} \text{Anderson} \\ \text{Brady} \end{array} \right\}$ .

This is an important piece of data, but does little to determine whether free choice is a pragmatic or semantic effect. Some authors have taken it to be a kind of disambiguation (Simons 2005; Aloni 2007), while others have understood the *but* clauses as undercutting premises in a pragmatic inference (e.g., Zimmermann 2000).

The data so far tempt a non-classical semantics for disjunction or modals which predicts them as entailments. But that makes (1a) and (1b) equivalent, in which case it is difficult to predict their classical behavior under negation: negating (1a) intuitively means that *both* disjuncts are prohibited (Alonso-Ovalle 2006; Fox 2007).

**Double Prohibition (DP)**

$$\neg \text{May}(A \vee B) \Rightarrow \neg \text{May}A \wedge \neg \text{May}B$$

- (6) a. Members may not vote for Anderson or Brady.
- b. Members may not vote for Anderson and members may not vote for Brady.

A non-classical approach would seem to incorrectly predict a weak meaning for (6a):  $\neg(\text{May}A \wedge \text{May}B)$ . But a classical account gives exactly what's needed: if there's not a world where  $A \vee B$  is true, then there's not a world where  $A$  is true and there's not a world where  $B$  is true.

This twist seems to favor a pragmatic analysis that treats free choice implications as implicatures. But articulating an adequate pragmatic analysis has pushed the frontiers of pragmatics itself. Some theorists have rejected the Gricean axiom that implicatures are computed globally in terms of the whole utterance. They propose to treat implicatures locally, arising sub-sententially at the level of clauses, and postulate unpronounced operators to integrate with this process (e.g., [Chierchia 2006](#); [Fox 2007](#)). [Franke \(2009, 2011\)](#) instead suggests that these implicatures are the result of more interactive, iterated reasoning for which game-theoretic tools are needed.<sup>1</sup> While both of these sophisticated pragmatic approaches treat free choice effects as scalar implicatures, recent processing ([Chemla & Bott 2014](#)) and acquisition ([Tieu, Romoli, Zhou & Crain 2016](#)) studies demonstrate significant differences between free choice effects and scalar implicatures. This leaves open the possibility of a yet more subtle pragmatic approach ([Tieu et al. 2016](#)). But it also opens the door for a non-classical semantics that could somehow predict both free choice effects and DP. I will formulate such a theory in §2, but this alone does not distinguish that semantic account from others.

Among semantic theories of free choice effects, only [Aloni 2007](#), [Barker 2010](#), [Aher 2012](#) and [Willer 2015](#) offer some account of DP.<sup>2</sup> However, none of these theories offer compelling accounts of WFC. They all appeal to [Simons's \(2005: 281-2\)](#) proposal that across-the-board LF movement can transform (2a) to (1a) at LF. This would reduce the problem of predicting WFC to that of NFC. But it has not been observed in the literature that this approach faces a difficult over-generation

1 See [van Rooij 2010](#) for a helpful comparison of this approach with Neo-Gricean and localist ones. See [Schulz 2005](#) for a more traditional Neo-Gricean account.

2 The accounts of DP in [Aloni 2007](#) and [Barker 2010](#) are not fully satisfying. [Aloni 2007: 80](#) can predict DP with a particular selection of  $A \vee B$ 's alternatives, but does not offer a systematic account of how this selection is made. [Barker \(2010: §5\)](#) treats DP as an implicature, based on uncooperative or uninformed speakers blocking the implication. But this kind of data would also speak equally against a semantic explanation of NFC given (4) and (5). The analysis in §2 predicts DP as an entailment without further assumptions and offers a different account of (4) and (5).

problem. LF movement is a type-driven process, which makes it hard and *ad hoc* to limit it to particular modals and connectives of the same type. Yet, (7a) does not have a reading on which it means (7b).

- (7) a. Members may vote for Anderson and members may vote for Brady.  
 b. #Members may vote for Anderson and Brady.

May A  $\wedge$  May B doesn't transform to May (A  $\wedge$  B), despite being formally parallel to the alleged transformation of May A  $\vee$  May B into May (A  $\vee$  B). It is also worth noting that none of the existing pragmatic accounts explain WFC either (van Rooij 2010: 24). The analysis proposed in §2 will semantically predict WFC without appeal to movement. In developing that analysis, I will first look to allied phenomena which reveal free choice as part of a broader pattern of *resource sensitivity*.

## 1.2 Resource sensitivity and strong permission

Simons 2005 and Barker 2010 also stress the non-implications in (8), noting that when the disjuncts are not contextually exclusive — suppose Anderson and Brady are in a runoff election where members get to vote for two candidates — it cannot be inferred that one may not choose both disjuncts. Barker 2010 diagnoses this as an effect of permission being a *discrete resource* and embraces a non-classical logic to suit. On this theme I highlight (9), which shows that a hearer can't assume permission persists after one option has been chosen (Asher & Bonevac 2005: 304).

### **Resource Sensitivity (RS)**

1. May (A  $\vee$  B)  $\not\Rightarrow$  May (A  $\wedge$  B)
2. May (A  $\vee$  B)  $\not\Rightarrow$   $\neg$ May (A  $\wedge$  B)
3. May (A<sub>1</sub>  $\vee$  A<sub>2</sub>), A<sub>i</sub>  $\not\Rightarrow$  May A<sub>j</sub>; i, j  $\in$  {1, 2}

- (8) a. You may vote for Anderson or Brady.  
 b. #You may vote for both Anderson and Brady.  
 c. #You may not vote for both Anderson and Brady.
- (9) a. You may vote for Anderson or Brady.  
 b. You did vote for Anderson.  
 c. #You may (still) vote for Brady.

While RS1 and RS2 are non-entailments in standard modal logic, the conclusions follow as implicatures on most pragmatic approaches — see Barker 2010: §6.1. Similarly for RS3, which is predicted by standard modal logic but the conclusion

still follows as an implicature on most pragmatic approaches. Semantic analyses like [Barker 2010](#) and [Aloni 2007](#) also fail to predict RS3.<sup>3</sup>

Resource sensitivity appears to go even deeper, preventing an inference from  $\text{May } \phi$  to  $\text{May } \psi$  even when  $\phi$  entails  $\psi$ . This is clear with disjunction, where *You may vote for Anderson or Brady* does not follow from *You may vote for Anderson*. Many semantic theories predict this, but rely crucially on the semantics of disjunction to do so. However an example from [Starr 2016b](#): §2.3 suggests this phenomenon is more general. Arnie, a mobster, enlists his ruthless minion Monica to interrogate a known snitch, Jimmy. Arnie doesn't care for serious torture, so he's developed another method. He will have Monica inject Jimmy with both a deadly poison and its antidote. When the two solutions are administered simultaneously they produce only minor cramps and shortness of breath. Monica is to tell Jimmy that he will be administered the antidote if he confesses and informs on other snitches, relying on the minor symptoms to persuade. Here, (10b) does not intuitively follow from (10a).

- (10) a. Monica may inject Jimmy with poison and antidote.  
b. #Monica may inject Jimmy with poison.

Indeed, Arnie may call Monica en route to clarify — perhaps knowing she would love to give Jimmy just the poison and see him die — to say (11).

- (11) You may not inject Jimmy with poison. Inject him with the solution of poison and antidote.

So we have a failure of the inference from  $\text{May}(P \wedge A)$  to  $\text{May}P$  and from  $\neg\text{May}A$  to  $\neg\text{May}(P \wedge A)$ . While permission has been granted for  $P \wedge A$ , that resource cannot be used to generate permission for a related, but more general resource:  $P$ .

This line of investigation leads us back to the very passage where [von Wright 1968](#): 4–5 first observed the puzzle of free choice permission. It is claimed there that the kind of permission that gives rise to free choice is *strong permission*. Suppose I say *You may eat apples* and say *You may not eat bananas*. Bananas are forbidden, apples are not forbidden, but cherries also aren't forbidden. This distinction between the status of apples and cherries is exactly the distinction between strong and weak permission. As [Barker 2010](#): §3 highlights, this is crucial to capturing RS1 and RS2. By RS2,  $\text{May}(A \vee B)$  does not forbid  $A \wedge B$ . But by RS1, it also doesn't entail  $\text{May}(A \wedge B)$ . So *may* here must express strong permission. Note that this context does not support *You may eat cherries*, but it does support (12a) and (12b).

<sup>3</sup> [Fusco 2015](#) and [Asher & Bonevac 2005](#) are the only semantic analyses I know of which capture RS3. Unfortunately, neither captures WFC or DP, which are my primary focus in this paper. However Fusco (p.c.) informs me that her system affords a different semantics of negation that captures DP.

**Resource Sensitivity (RS) Cont.**4.  $\text{May}(A_1 \wedge A_2) \not\Rightarrow \text{May}A_i; i \in \{1, 2\}$ 5.  $\neg\neg\text{May}\phi \not\Rightarrow \text{May}\phi$ 6.  $\neg\text{Must}\neg\phi \not\Rightarrow \text{May}\phi$ 

- (12) a. It's not the case that you may not eat cherries.  
 b. It's not the case that you must not eat cherries.

If this is right, then the logic of *may* follows **RS5** and the logic of *must* follows **RS6**. The failure of duality in **RS6** is old news.<sup>4</sup> But the failure of double-negation elimination in modal contexts has not been observed. It requires not just a non-classical semantics for modals, but a non-classical semantics for negation.<sup>5</sup>

Weak permission is naturally suited to classical modal logic. If a set of worlds  $R$  models what's required, any proposition classically consistent with  $R$  is weakly permitted. To capture **RS** one must abandon an analysis of permission as weak permission, and classical logic with it. But this makes it harder, not easier, to capture the classical pattern of **DP**. **RS6** shows the need for a non-classical semantics of negation and this suggests a way forward. In the following section I will distinguish requirements and strong permission *dynamically*: they involve different ways of expressing preferences. Crucially, expressing preferences will not be equated with eliminating worlds where certain preferences are held. Instead, it will be analyzed as directly modifying a preference relation. Negation will then have two functions, one when it modifies a sentence that eliminates worlds — an informational, descriptive sentence — and one when it modifies preferences. Surprisingly, it is possible to do this without lexical ambiguity. This semantics of negation, when combined with a particular dynamic account of disjunction and strong permission, will capture all of the patterns discussed in this section.

**2 Expressing permission, dynamically**

The analysis developed in this section is motivated by a simple idea: expressing permission involves incrementally building a partial map of what can be done, rather than describing what the fully precise permission facts in some world are. Articulating this idea requires making precise this contrast between incremental, partial expression of permission and describing precise permission facts that hold in a world. I will do this by first presenting in §2.1 the simple model of informational dynamics from [Veltman 1996](#), and then contrasting it in §2.2 with the model of deontic dynamics proposed here. Section 2.3 will use this model to sketch a semantics


<sup>4</sup> von Wright (1968: 4–5). More recently: [Kratzer 1981](#): §4, [McNamara 2010](#) and [Cariani 2013](#): §5.

<sup>5</sup> I thank Malte Willer for encouraging me to think about double-negation.

and logic for *may* that captures free choice effects. Negation and DP are treated in §2.4. Section 2.5 uses these tools to explain RS1–6 and §2.6 returns to the issue of how ignorance and uncooperativeness can defeat free choice effects.

## 2.1 Information dynamics

The dynamics of information is simple. Information says the world is some of these ways, and none of those. This is captured by taking a state of information  $s$  to be a set of worlds. If one assumes a sentence’s only job is to provide information about the world, then a sentence  $\phi$ ’s meaning  $[ \phi ]$  can be thought of as a function from one state of information  $s$  to another  $s'$ . In Veltman’s (1996) terminology,  $s[ \phi ]$  is the result of updating  $s$  with  $\phi$ . On this dynamic approach, an atomic sentence  $A$  serves to eliminate worlds from  $s$  where  $A$  is false.  $\neg\phi$  removes worlds that would survive



**Figure 1** Atomic Update (Uppercase = True, Lowercase = False)

an update with  $\phi$ , while  $\phi \wedge \psi$  sequences the effects of its conjuncts.  $\phi \vee \psi$  unions the effects of its conjuncts. More formally:

### Informational Update Semantics (Veltman 1996)

1.  $s[A] = \{w \in s \mid w(A) = 1\}$
2.  $s[\neg\phi] = s - s[\phi]$
3.  $s[\phi \wedge \psi] = (s[\phi])[ \psi ]$
4.  $s[\phi \vee \psi] = s[\phi] \cup s[\psi]$

Adding deontic modals to this framework presents a choice: do deontic modals provide information about the world, and so update  $s$ , or do they have a different kind of effect entirely? The traditional approach in modal logic has been to assume that deontic modals provide information about the world.  $\diamond\phi$  eliminates any world  $w$  relative to which there is not an accessible  $\phi$ -world. On this analysis  $\diamond\phi$  provides information about the world, namely which worlds are accessible from our world. Veltman 1996 offers a slightly different approach to Might  $\phi$ , where it does not point-wise eliminate worlds based on their properties, but places a global test on the information state itself. The result of the test is  $s$  or  $\emptyset$ .

**Test Semantics for Might**  $s[\text{Might } \phi] = \{w \in s \mid s[\phi] \neq \emptyset\}$



I will explore an even further departure from the classical semantics. Deontic modals don't describe worlds or even test information states, they test and update what I will call a *deontic frame*  $\pi$ . A deontic frame will be modeled using preference relations between worlds. After all, like preferences, deontic modals serve to *motivate* agents to do things. On a traditional descriptive semantics the best one can do to capture this connection between motivation and deontic modals is to have deontic propositions *describe* the preferences that hold in a world e.g., *May A* is true in  $w$  if the most preferred worlds in  $w$  are consistent with  $A$ . The account developed here will allow one to model language which directly influences preferences, without recourse to propositions that passively describe those preferences.<sup>6</sup>

## 2.2 Deontic dynamics

Following [Kamp \(1973, 1978\)](#), [Lewis \(1979\)](#) and [van Rooij \(2000\)](#), I will analyze *May*  $\phi$  dynamically in terms of how it updates requirements/permissions  $\pi$ , rather than information  $s$  (a set of worlds). But the dynamic analysis I will propose has two key differences. One difference will be discussed later in §2.3. The difference I will focus on here is that  $\pi$  distinguishes weak permission and strong permission by having two separate 'preference frames' for what's required and what's strongly permitted — the motivation for modeling preference frames in terms of a strict preference ordering *and* an indifference ordering is discussed further in [Appendix A, Remark 1](#). The basic idea is that making requirements and providing permissions both involve presenting preferences, just in different ways.

### Practical Frames


$\pi := \langle R_\pi, P_\pi \rangle$  consists of **requirements**  $R_\pi$  and **strong permissions**  $P_\pi$

1. **Requirement Frame:**  $R_\pi := \langle r_\pi, \sim_\pi \rangle$ 
  - $r_\pi(w_1, w_2)$ :  $w_1$  is strictly preferable to  $w_2$
  - $w_1 \sim_\pi w_2$ :  $w_1$  is just as preferable as  $w_2$
2. **Permission Frame:**  $P_\pi := \langle p_\pi, \approx_\pi \rangle$  same as  $R_\pi$ .

This is best illustrated by considering the *indifferent practical frame I*, depicted in [Figure 2](#). The graph on the left represents the requirements, and that on the right the permissions. Wavy lines depict the indifference relation, while straight lines will be used for strict preferences (reflexive wavy lines are omitted in all diagrams for readability). This simple practical frame distinguishes weak and strong permission.


<sup>6</sup> For more on the philosophical motivations of this approach see [Starr 2016a](#).





**Figure 2** Indifferent Practical Frame: no requirements, no strong permissions

The requirements  $R_I$  are completely indifferent about which world is realized, so everything is weakly permitted. Yet, nothing is strongly permitted since  $P_I$  does not promote any worlds as better, or even equally good as, any other. As discourse unfolds, strict preferences are introduced to  $R_I$  and  $P_I$ , and indifference fills in between worlds that are not strictly preferred to one another. Figure 3 depicts this process by first introducing a requirement and then introducing a permission.



**Figure 3** Making A Required, Then B Permitted

First, A is made to be required in  $I$  —  $Req_A(I)$  — which involves making A-worlds preferred in both  $R_I$  and  $P_I$  — the notation of  $A(\cdot)$  is used for this operation. This just means that introducing an explicit requirement entails strong permission. Next, to make B permitted in the resulting state, one creates a permission ordering from the requirement ordering. This is done by adding a preference for B-worlds to the existing preferences. There are two crucial things to note here. First, only a preference for the top-ranked B-world is introduced. When a permission is introduced, it must be integrated with the existing requirements. Permission to do B, after A has been required, can only be permission to do  $A \wedge B$ . Second, this process

would appear to overwrite any prior permissions. How would one capture two distinct strong permissions to do  $A$  and  $\neg A$ ? This issue, along with the interaction of permission and information, requires augmenting the basic model of requirements and permissions sketched here.

### 2.3 Semantics and logic


Integrating information and deontic frames seems simple enough: let sentences update  $\langle s, \pi \rangle$ . But there is another crucial twist here that differentiates this analysis from others. Sentences will update a *set* of such pairs, as there can be many  $\pi$ 's and  $s$ 's at play in discourse:

**States**  $S$  is a set of **substates**:  $S = \{s_1^{\pi_1}, \dots, s_n^{\pi_n}\}$

- Each **substate**  $s^\pi$  consists of an information state  $s$  and a **practical frame**  $\pi$ :  $s^\pi := \langle s, \pi \rangle$ .

The *initial state*  $\mathbf{0}$  has a single substate with the set of all worlds  $W$  as its information state, and  $\mathbf{I}$  as its practical frame:  $\mathbf{0} := W^{\mathbf{I}}$ . In a state  $S$  where there are multiple substates each  $s^\pi \in S$  is competing for control over the agent's actions and beliefs. This is not to say that agents are uncertain about which unique  $s^\pi$  obtains, or that the discourse leaves a particular  $s^\pi$  underdetermined. Instead, the agents are allowing a range of  $s^\pi$  to remain in play to explore a wider range of options without needing to decide between them. As it turns out, the use of states rather than just substates provides a crucial resource for analyzing permission and disjunction.

$\text{May}A$  and  $\text{May}\neg A$  are intuitively consistent, but there is no single coherent  $P_\pi$  which both ranks  $A$ -worlds over  $\neg A$ -worlds and ranks  $\neg A$ -worlds over  $A$ -worlds. Substates solve this problem by allowing  $\text{May}A$  to create a new substate where there is strong permission for  $A$ , but also leave prior substates intact. On reflection, this makes sense: granting permission to do  $A$  allows the hearer to act in accord with a background  $\pi$  where  $A$  may not be preferred, but it also allows the hearer to act in accord with an ordering just like  $\pi$  except a permissive preference for  $A$  has been added — call it  $\text{Per}_A(\pi)$ . So a successful update of  $\mathbf{0}$  with  $\text{May}A$  will contain two substates, one with  $\mathbf{I}$  as its practical frame, and one with  $\text{Per}_A(\mathbf{I})$  as its practical frame. This is depicted in Figure 4 using the same basic conventions as before, only now the worlds pictured are from the relevant information state, boxes delimit substates and a bold box delimits the whole state.


**Figure 4**  $\mathbf{0}[\text{May } A]$ 

This semantics for *may* can be stated as the following recipe.

### Semantics for May

$S[\text{May } A]$ : Is A is weakly permitted by all  $\pi$ ? If *yes* do (a), if *no* do (b).

- a. Add strong permission for A to each  $\pi$ , put each augmented  $\pi$ ,  $Per_A(\pi)$ , in play as well as each original  $\pi$ .
  - Map  $S = \{s_1^{\pi_1}, \dots, s_n^{\pi_n}\}$  to  $S' = \{s_1^{\pi_1}, \dots, s_n^{\pi_n}, s_1^{Per_A(\pi_1)}, \dots, s_n^{Per_A(\pi_n)}\}$
- b. Reduce each  $s$  to  $\emptyset$ :  $\{\emptyset^{\pi_1}, \dots, \emptyset^{\pi_n}\}$

(See Appendix A, Definition 9 for full formalization)


As discussed above,  $Per_A(\pi)$  simply overwrites  $P_\pi$  with  $A(R_\pi)$  i.e., it creates a new permissive ordering from  $\pi$ 's requirements with an added preference for A-worlds. This statement of the semantics and Figure 4 make clear a crucial feature of the analysis: *May A* creates substates. This is crucial because disjunction also creates substates, predicting a special connection between the two.

The semantics for conjunction and disjunction is unchanged from above.

**Connective Semantics** 1.  $S[\phi \wedge \psi] = (S[\phi])[\psi]$ ; 2.  $S[\phi \vee \psi] = S[\phi] \cup S[\psi]$

But this semantics now predicts that disjunctions will create substates. For example,  $\mathbf{0}[A \vee B]$  will return  $\{\{w_{AB}, w_{Ab}\}^I, \{w_{AB}, w_{aB}\}^I\}$ . The fact that disjunction creates substates interacts in an important way with the semantics for *May*. Updating  $\mathbf{0}$  with *May A*  $\vee$  *May B* will result in a state just like  $\mathbf{0}[\text{May } A]$  in Figure 4, except there will another substate  $W^{Per_B(I)}$  where B-worlds are preferred. More generally:

$$\{s_1^{\pi_1}, \dots, s_n^{\pi_n}\}[\text{May } A \vee \text{May } B] = \{s_1^{\pi_1}, \dots, s_n^{\pi_n}, s_1^{Per_A(\pi_1)}, \dots, s_n^{Per_A(\pi_n)}, s_1^{Per_B(\pi_1)}, \dots, s_n^{Per_B(\pi_n)}\}$$



**Figure 5**  $\mathbf{0}[\text{May } A \vee \text{May } B]$

As Figure 5 shows, subsequently updating this state with either May A or May B will have no effect. May A always leaves incoming substates in the output state, so it could only add substates. But May A adds substates by overwriting their permissions. This means it will turn each  $s_i^{\pi_i} \in \{s_1^{\pi_1}, \dots, s_n^{\pi_n}\}[\text{May } A \vee \text{May } B]$  into  $s_i^{\text{Per}_A(\pi_i)}$ . Since each of those is in  $\{s_1^{\pi_1}, \dots, s_n^{\pi_n}\}[\text{May } A \vee \text{May } B]$  already, updating with May A after updating with May A  $\vee$  May B will not produce any change to the deontic frames. This is precisely what is required for practical consequence and for predicting **WFC**.  $S$  p-supports  $\phi$  when it doesn't change any of the  $\pi$ 's at play in  $S$ , and p-consequence is just p-support in any state that has been updated with the premises.

**P-Support**  $S \models \phi : S \models \phi \iff \Pi_S = \Pi_{S[\phi]}$ , where  $\Pi_S := \{\pi \mid s^\pi \in S \ \& \ s \neq \emptyset\}$

**P-Consequence**  $\phi_1, \dots, \phi_n \models \psi \iff \forall S: S[\phi_1] \dots [\phi_n] \models \psi$

This much explains **WFC**. Predicting **NFC** hinges on further details.

As noted above,  $A \vee B$  creates a substate for each disjunct. In this sense,  $\phi$ 's dynamic meaning determines its alternatives in  $S$ :

**Alternatives**  $\text{alt}_S(\phi) := \{a \mid \exists \pi: a^\pi \in S[\phi]\}$

As in [Simons 2005](#) and [Aloni 2007](#), one can formulate the semantics of  $\text{May } \phi$  so will operate on each of  $\phi$ 's alternatives:  $\text{May } \phi$  takes each  $a \in \text{alt}_S(\phi)$  and each input  $\pi$ , and tests whether  $a$  is consistent with what's required by  $\pi$ . If so, a substate featuring  $\text{Per}_a(\pi)$  is added to  $S$  — see [Definition 9](#) in [Appendix A](#). This predicts that  $S[\text{May}(A \vee B)] = S[\text{May}A \vee \text{May}B]$ . So **NFC** is valid, just as **WFC** is. It is worth noting that  $S[\text{May}(A \vee B)]$  does not in general support  $\text{May}(A \wedge B)$  (**RS1**). Conjunctive permission would add a substate where only  $w_{AB}$  is strictly preferred to every other world. To predict the other **RS** patterns and **DP**, one must formulate a semantics of negation which not only operates on information, but also preferences.

## 2.4 Negation and double prohibition


$\neg\phi$  will remove worlds that would survive an update with  $\phi$ , as in the semantics for negation from [§2.1](#). But, it also removes preferences that would result from an update with  $\phi$  — see [Definition 15](#), [Appendix A](#).

**Negation**  $S[\neg\phi]$ :

1. Remove information that would survive update with  $\phi$
2. Retract  $\phi$  preferences from each  $\pi$ , (notation:  $\pi \downarrow \phi$ )
  - a. Remove strict permissive preferences that  $\phi$  would add to  $\mathbf{0}$ , reverse them and make them both requirement and permissive preferences
  - b. Remove strict requirement preferences that  $\phi$  would add to  $\mathbf{0}$
  - c. If a strict preference relating  $w$  and  $w'$  was removed and not reversed, introduce indifference between  $w$  and  $w'$

The various clauses are best explained with two kinds of examples. One where  $S$  p-supports  $\neg\text{May}A$ , one where  $S$  does not p-support  $\neg\text{May}A$ .

To find a state that supports  $\neg\text{May}A$  one first has to find a state where  $A$  is inconsistent with what's required. By [Clause 1](#), the test with  $\text{May}A$  needs to fail, or else the information of the state will be reduced to  $\emptyset$ . [Figure 6](#) depicts  $R_{\pi_1}$  in  $S_1$ , which is an example of a state where the test imposed by  $\text{May}A$  will fail. But, to support  $\neg\text{May}A$ ,  $S_1$  must also already contain the preferences  $\neg\text{May}A$  would add, and lack the preferences it would remove. In particular, [Clause 2a](#) tells us that the state must *not* have a preference for  $A$ -worlds over  $\neg A$ -worlds in  $P_\pi$ , and the state must have  $\neg A$ -worlds preferred to  $A$ -worlds in  $P_\pi$  and  $R_\pi$ . ([Clause 2b](#) doesn't apply here, as  $\text{May}A$  does not change the requirements.) [Clause 2c](#) ensures that any worlds that are not related by strict preference are related by indifference.



**Figure 6**  $S_1 \models \neg\text{May} A$  and  $S_1 \not\models \neg\text{May}(A \vee B)$

State  $S_1$  p-supports  $\neg\text{May} A$ , but it does not p-support  $\neg\text{May}(A \vee B)$ . When one retracts  $\text{May}(A \vee B)$  from  $\pi_1$ , one must reverse any strict permissive preference that exists in  $\mathbf{0}[\text{May}(A \vee B)]$ . This means one must reverse all the preferences in  $A(P_1)$  and in  $B(P_1)$ , and put them together into both  $R_\pi$  and  $P_\pi$  of one practical frame  $\pi$ . Looking back at Figure 5, it should be clear that the result is the state depicted on the right in Figure 6. Here,  $w_{ab}$  is the only rational choice. This makes clear that  $\neg\text{May}(A \vee B)$  does not have a weak reading akin to  $\neg\text{May} A \vee \neg\text{May} B$ , despite the fact that the semantics validates **WFC** and **NFC**. Figure 6 tells one enough to see how **DP** ends up valid.  $S_1[\neg\text{May}(A \vee B)]$  is the minimal state that would p-support  $\neg\text{May}(A \vee B)$ . Indeed, the same state would have resulted from  $\mathbf{0}[\neg\text{May}(A \vee B)]$ . As discussed,  $S_1$  p-supports  $\neg\text{May} A$  but as the graphs make clear, all of the effects produced by  $\neg\text{May} A$  are already in place in  $S_1[\neg\text{May}(A \vee B)]$ . The same goes for  $\neg\text{May} B$ . In sum, this semantics somewhat miraculously makes  $\text{May}(A \vee B)$  behave non-classically when unembedded, but classically when embedded under negation. The key was a semantics for negation which operates not just on information, but on practical frames as well. This semantics for negation may look complex. But, conceptually, it is a simple and familiar idea:  $\neg\phi$  works by removing structures that would persist in a hypothetical update with  $\phi$ .

## 2.5 Resource sensitivity

### Resource Sensitivity (RS)

1.  $\text{May}(A \vee B) \not\Rightarrow \text{May}(A \wedge B)$
2.  $\text{May}(A \vee B) \not\Rightarrow \neg\text{May}(A \wedge B)$
3.  $\text{May}(A_1 \vee A_2), A_i \not\Rightarrow \text{May} A_j; i, j \in \{1, 2\}$
4.  $\text{May}(A_1 \wedge A_2) \not\Rightarrow \text{May} A_i; i \in \{1, 2\}$
5.  $\neg\neg\text{May} \phi \not\Rightarrow \text{May} \phi$
6.  $\neg\text{Must} \neg\phi \not\Rightarrow \text{May} \phi$

Of the above, only **RS1** has been explained. But when the semantics for negation is considered alongside Figure 5, it should be fairly clear how **RS2** is predicted. Updating  $\mathbf{0}[\text{May}A \vee \text{May}B]$  with  $\neg\text{May}(A \wedge B)$  would remove the preference for  $w_{AB}$  over  $w_{ab}$  in both  $A(R_I)$  and  $B(R_I)$ . So it cannot be that  $\neg\text{May}(A \wedge B)$  is a p-consequence of  $\text{May}A \vee \text{May}B$ . While on the topic of negation, **RS5** and **RS6** deserve attention.

The failures of double-negation elimination behind **RS5** are very specific, as suggested by the natural language data considered in §1.2. They are exactly in those states where there is a difference between what's weakly permitted and what's strongly permitted.  $\mathbf{0}$  is just such a state. Consider  $\mathbf{0}[\neg\neg\text{May}A]$ . This will remove from  $\mathbf{0}$  the permissive and requirement preferences that  $\neg\text{May}A$  would add to  $\mathbf{0}$ . Looking back at Figure 6, these will be any strict preferences for  $\neg A$ -worlds. There are no such preferences in  $\mathbf{0}$ , so  $\mathbf{0} \models \neg\neg\text{May}A$ . But clearly  $\mathbf{0} \not\models \text{May}A$ , since  $\text{May}A$  adds strong permission for  $A$ . It is worth noting that in states like  $S_1$  from Figure 6 this mismatch between weak and strong permission does not hold, and those contexts do not provide counterexamples to double-negation. It is therefore possible to formally specify a restricted version of double-negation, should one want to explain why it often sounds like a good inference.

Basically the same reasoning is behind **RS6**, although this requires specifying a semantics for *must*. Here I adopt the semantics developed in Starr 2016a — see Definition 13 in Appendix A.  $\text{May}\phi$  tests whether  $\phi$  is consistent with the worlds best according to each input  $R_\pi$ . If so, preferences for the best  $\phi$ -worlds are added to each  $R_\pi$  and  $P_\pi$ . If not, each substate is reduced to  $\emptyset^\pi$ . In  $\mathbf{0}$ ,  $\neg\text{Must}\neg A$  will idle since  $\mathbf{0}$  has no preferences to remove in the first place. But  $\text{May}A$  will clearly change  $\mathbf{0}$ : it will add strong permission for  $A$ . As with **RS5**, this is a limited failure of the classical pattern. It is only in very specific kinds of states that it will fail.

To see how **RS3** is predicted, consider updating  $\mathbf{0}[\text{May}A \vee \text{May}B]$  from Figure 5 with  $B$ . This simply trims out the  $\neg B$ -worlds, depicted below in Figure 7.<sup>7</sup> Subsequently updating with  $\text{May}A$  would not change this particular state, since it would turn all input practical frames into  $A(R_I)$ , and union them back into the state above.  $A(R_I)$  is already there. However, recall from §1.2 that in the natural language examples used to support **RS3**,  $w_{AB}$  was prohibited. In the state  $\mathbf{0}[\text{May}A \vee \text{May}B][\neg\text{May}(A \wedge B)][B]$ , only  $B^{\text{Per}_B(I)}$  and  $B^I$  will persist. But when  $\text{May}A$  transforms them into  $A(R_I)$  and unions it back into the state, a change occurs.

**RS4** follows from the fact that  $\text{May}(A \wedge B)$  will prefer  $w_{AB}$  to every world, and will not prefer  $w_{Ab}$  to  $w_{ab}$ . Since  $\text{May}A$  will add a substate where  $w_{AB}$  is preferred to  $w_{ab}$ , it cannot be a p-consequence of  $\text{May}(A \wedge B)$ .

<sup>7</sup> It is worth clarifying that  $B$  does not change the orderings, only the space of worlds. However, *may* and *must* are only concerned with  $c_S$ , so one can pretend as if they do. This difference will matter if the system is extended to deontics like *should* and *ought* which range over a wider class of worlds.






Figure 7  $\mathbf{0}[\text{May } A \vee \text{May } B][B]$

## 2.6 Coping with the ignorant and rude

Ignorance, as in (4), and uncooperativity, as in (5), cancel free choice effects.

- (4)
  - a. Members may vote for Anderson or Brady, but I don't know which.
  - b. # Members may vote for  $\left\{ \begin{array}{l} \text{Anderson} \\ \text{Brady} \end{array} \right\}$ .
- (5)
  - a. Members may vote for Anderson or Brady, but I won't tell you which.
  - b. # Members may vote for  $\left\{ \begin{array}{l} \text{Anderson} \\ \text{Brady} \end{array} \right\}$ .

On the analysis above,  $\text{May } A \vee \text{May } A$  and  $\text{May } (A \vee B)$  are equivalent. So (4) and (5) cannot involve disambiguating between narrow and wide-scope readings of the modal. Instead, I propose that both (4) and (5) cancel free choice effects by introducing uncertainty about which of two states to adopt.

Recall that if  $s_1^{\pi_1}, \dots, s_n^{\pi_n} \in S$ , then:

$$S[\text{May } (A \vee B)] = \{s_1^{\pi_1}, \dots, s_n^{\pi_n}, s_1^{A(\pi_1)}, \dots, s_n^{A(\pi_n)}, s_1^{B(\pi_1)}, \dots, s_n^{B(\pi_n)}\}$$

Both (4) and (5) result in higher-order uncertainty over which of two states should be adopted:  $S_A = \{s_1^{\pi_1}, \dots, s_n^{\pi_n}, s_1^{A(\pi_1)}, \dots, s_n^{A(\pi_n)}\}$  or  $S_B = \{s_1^{\pi_1}, \dots, s_n^{\pi_n}, s_1^{B(\pi_1)}, \dots, s_n^{B(\pi_n)}\}$ . Adapting the supervaluationist ideas of [Van Fraassen 1966](#) and [Stalnaker 1981](#), a sentence is supported despite such uncertainty only if it is supported by all resolutions of that uncertainty. Since only one resolution p-supports *MayA* —  $S_A$  — and only one p-supports *MayB* —  $S_B$  — neither permission claim is supported. This also clarifies the importance of interpreting substates as competing for control over actions and beliefs rather than uncertainty about what state one is in. Uncertainty involves deliberation and a forced choice between two or more options, while the former allows the agents to leave this choice unmade. But how exactly does higher-order uncertainty arise from the compositional semantics of (4) and (5)?

Consider first some fully explicit versions of the *but*-phrases.

(13) I don't know which of the two candidates you may vote for.

(14) I won't tell you which of the two candidates you may vote for.

Both of them convey, whether by presupposition, assertion or implicature, that one, and only one, of the two candidates may be voted for. Additionally, (13) asserts that the speaker is uncertain whether they should be in a state of mind represented by  $S_A$  or one represented by  $S_B$ . So (13)'s total contribution is that the speaker is uncertain about whether to adopt a state of mind represented by  $S_A$  or one represented by  $S_B$ , and that only one of these representations is correct. On the assumption that the speaker is more authoritative than the hearer about permissions, it follows that the state representing the hearer's state of mind should follow suit. That is, they should also be uncertain about which of the two states to adopt, and require a choice between them. The assertion of (14) is certainly different, as it entails that the speaker's state of mind is definitely represented by either  $S_A$  or by  $S_B$ . (14)'s total impact combines this with the information that only one state is correct. Here's how. Given the speaker's authority, the hearer should bring their state of mind to match the speaker's. As a result they are forced to choose between adopting  $S_A$  and  $S_B$ , but uncertain which one to choose. Combined with reasonable assumptions, this is enough to predict how free choice effects are blocked in (4) and (5).

It is plausible to assume that the elliptical *but*-phrases in (4) and (5) mean something like (13) and (14), respectively. It is also plausible to assume that *but* is a species of conjunction, and so sequentially updates the state. This means that the first conjuncts of (4) and (5) will license free choice inferences, but once the state is further updated by the second conjunct and its pragmatic implications inferred, the conversation enters a state that no longer licenses those free choice inferences. This kind of non-monotonicity was also key to explaining [RS3](#) where additional information blocked a free choice inference. It is only possible to give

this kind of analysis of ignorance and uncooperativity because the semantics builds non-monotonicity into the account of permission. While a more complete formal implementation of this analysis is needed, the sketch above shows that it will likely work out in a motivated and plausible way.

### 3 Conclusion

The semantics presented here covers more of the permission data in a more compelling way than the competing semantic analyses discussed in §1.1. The insight driving this semantics is that permission statements express incremental changes directly to preferences rather than describing fully precise permission facts. But I have said nothing about free choice effects that arise in contexts where permission is not involved, and it is not plausible to say that preferences are being expressed rather than described there. For example, it is well-known that in epistemic contexts *might* and disjunction lead to free choice effects, and similarly for disjunctions in the antecedents of conditionals. Further, [Klinedinst 2006](#), [Eckardt 2007](#) and [Fox 2007](#) observe that existential quantifiers with non-plural restrictors produce free choice inferences when they interact with disjunction.

- (15) a. There is beer in the fridge or the ice-bucket.  
 b.  $\Rightarrow$  There is beer in the fridge.  
 c.  $\Rightarrow$  There is beer in the ice-bucket.

Much further work is needed to say whether all of these other free choice effects also give rise to the kinds of resource sensitive reasoning detailed in §1.2. But even at this preliminary stage, it is crucial to clarify that the general style of semantics given here is not applicable only to the particulars of preferences and permission.

The crucial feature of the semantics is that it avoids fully precise descriptions of a particular semantic object — e.g., modal orderings — by instead incrementally building a partial map of that domain which exploits the way language users mentally represent that domain. Permission draws on preferences, which are incrementally constructed in a way that is sensitive to how humans represent them. The same kind of model for representing uncertainty and counterfactuals already exists ([Sloman 2005](#); [Pearl 2009](#)). Work on quantifiers, pluralities and discourse reference in dynamic semantics suggests similar resources for that phenomenon ([van den Berg 1996](#); [Nouwen 2003](#); [Brasoveanu 2008](#)). These approaches motivate rethinking our semantics of logical connectives in terms of how they incrementally modify partial representations of a domain. While it will not be possible to explore these connections here, there are enough structural parallels to make this a worthwhile direction for future research.

Expressing permission

## A Expressive deontic logic (EDL)

### Definition 1 (Syntax)

1.  $Wff_0 ::= At \mid (\neg Wff_0) \mid (Wff_0 \vee Wff_0) \mid (Wff_0 \wedge Wff_0)$
2.  $Wff ::= Wff_0 \mid (\text{May } Wff_0) \mid (\text{Must } Wff_0) \mid (\neg Wff) \mid (Wff \vee Wff) \mid (Wff \wedge Wff)$

**Definition 2 (Worlds  $W$ , Information States  $s$ )**  $W: At \mapsto \{0, 1\}$ ;  $s \subseteq W$

### Definition 3 (Practical Frames $\pi$ )

$\pi := \langle R_\pi, P_\pi \rangle$ , where  $R_\pi$  are *requirements* and  $P_\pi$  are *strong permissions*

1.  $R_\pi := \langle r_\pi, \sim_\pi \rangle$ 
  - $r_\pi(w, w')$ : ‘ $w$  is strictly preferable to  $w'$ ’
  - $w \sim_\pi w'$ : ‘ $w$  is just as preferable as  $w'$ ’
  - $w \not\sim_\pi w'$  iff  $r_\pi(w, w')$  and  $w \neq w'$
2.  $P_\pi := \langle p_\pi, \approx_\pi \rangle$ ; interpretation parallel to  $R_\pi$

**Remark 1** The need for both  $r_\pi$  and  $\sim_\pi$  comes from wanting to distinguish an agent who has irrational strict preferences i.e.,  $r_\pi = \{\langle w_1, w_2 \rangle, \langle w_2, w_1 \rangle\}$ , from an agent who takes  $w_1$  and  $w_2$  to be just as preferable. The former state of preference is expressed by  $\text{Must } A \wedge \text{Must } \neg A$  while the latter would support  $\neg \text{Must } A$  but neither  $\text{Must } A$  nor  $\text{Must } \neg A$ . Capturing these differences is essential to developing a thoroughly non-representational approach to deontic modality (Starr 2016a).

**Definition 4 (Indifferent Practical Frame)**  $\mathbf{I} := \langle \langle \emptyset, W^2 \rangle, \langle \emptyset, \emptyset \rangle \rangle$

### Definition 5 (States $S$ , Substates $s^\pi$ )

1. A **state**  $S$  is a set of **substates**:  $S = \{s_1^{\pi_1}, \dots, s_n^{\pi_n}\}$
2. A **substate**  $s^\pi$  is an information state  $s$  and a practical frame  $\pi$ :  $s^\pi := \langle s, \pi \rangle$

**Definition 6 (Initial State)**  $\mathbf{0} := \{W^1\}$  i.e., no information, practically indifferent

### Definition 7 (Conjunction, Disjunction)

1.  $S[\phi \wedge \psi] = (S[\phi])[ \psi ]$
2.  $S[\phi \vee \psi] = S[\phi] \cup S[\psi]$

**Definition 8 (Choice)**

$$Ch_s(R_\pi) := \{w_1 \in S \mid \nexists w_2 \in S: r_\pi(w_2, w_1) \ \& \ \exists w_2 \in S: w_1 \sim_\pi w_2 \ \text{or} \ r_\pi(w_1, w_2)\}$$

**Remark 2** Choice worlds are not dispreferred to any world and are either preferred to or just as preferable as at least one world. This second clause is necessary to ensure that a completely empty ordering does not make everything choosable. Intuitively, if you have absolutely no preferences no choice is good because none of them have anything going for them. This is relevant when considering  $Ch_W(P_I)$  which should be  $\emptyset$  rather than  $W$ . This captures the fact that everything is weakly permitted in  $I$  but nothing is strongly permitted.

**Definition 9 (May)**

$$S[\text{May } \phi] = \begin{cases} Per_\phi(S) & \text{if } \forall s^\pi \in S, \forall a \in alt_S(\phi): Ch_s(R_\pi) \cap a \neq \emptyset \\ \{\emptyset^\pi \mid s^\pi \in S\} & \text{otherwise} \end{cases}$$

**Remark 3** May  $\phi$  performs a test and then shifts the state depending on its outcome. It tests that for every substate and each of  $\phi$ 's alternatives  $a$ , the Choice worlds in that substate are consistent with  $a$ . In other words, it tests that each of  $\phi$ 's alternatives is weakly permitted in  $S$ . If the test is failed, each substate's information is reduced to  $\emptyset$ . If the test is passed,  $\phi$  becomes strongly permitted:  $Per_\phi(S)$ . This is done in two steps. First, one creates a new  $\pi$  for each of  $\phi$ 's alternatives  $a$ , notated  $Per_a(\pi)$ . As Definition 11.2 below states, this involves copying  $\pi$ 's requirements into the permission slot of  $\pi$ , and making  $a$  preferred in this new permission ordering. Definition 12 says that this is done by strictly preferring each  $a$  world in  $s$  over each non- $a$  world and making sure that  $a$  and non- $a$  worlds are not equally preferable. Second, one takes all such  $s^{Per_a(\pi)}$  and unions them together with  $S$  (Definition 11.1). This reflects the fact that permissions are not combined, but allowed to 'live alongside' one another. After all, May  $\phi$  and May  $\neg\phi$  are consistent.

**Remark 4** It is reasonable to wonder why substates and alternatives are universally quantified over in Definition 13. The universal quantification over substates predicts that *Ella is in her study or the parlor, you may not disturb her* can be supported by a state where visiting Ella is only problematic if she is in her study. The universal quantification over alternatives is required to make sure that May  $(A \vee B)$  requires both A and B to be weakly permitted.

**Remark 5** This semantics has May influencing both  $\pi$ , when the test is successful, and  $s$  when the test fails. This behavior is important when May  $\phi$  is negated. As Definition 15 details,  $\neg\psi$  eliminates preferences that  $\psi$  would add (reversing permissive preferences and making them requirements), and information  $\psi$  would add. So when the test imposed by May  $\phi$  fails,  $\neg\text{May } \phi$  will have no effect on the

information, since it takes  $s - \emptyset$ . But it will still have an effect on the preferences: it takes permissive preferences for  $\phi$ -worlds over  $\neg\phi$ -worlds, reverses them and adds them to the requirements. This correctly predicts that  $\text{Must}\neg\phi$  will be a practical consequence of  $\neg\text{May}\phi$ . This operation also predicts DP.  $\neg\text{May}(A \vee B)$  will end up adding to the requirements a preference for  $\neg A$ -worlds over  $A$ -worlds and  $\neg B$ -worlds over  $B$ -worlds, since it reverses each of the permissive preferences that  $\text{May}(A \vee B)$  would add and adds the inverse of this preference to *all* substates. The resulting state will therefore support both  $\neg\text{May}A$  and  $\neg\text{May}B$ .

**Definition 10 (Alternatives for  $\phi$  given  $S$ )**  $alt_S(\phi) := \{a \mid \exists\pi: a^\pi \in S[\phi]\}$

**Definition 11 (Permitting  $\phi$  in  $S$ ,  $a$  in  $\pi$ )**

1.  $Per_\phi(S) := S \cup \{s^{Per_a(\pi)} \mid s^\pi \in S \ \& \ a \in alt_S(\phi)\}$
2.  $Per_a(\pi) := \langle R_\pi, a \langle R_\pi \rangle \rangle$

**Definition 12 (Preferring  $a$  in  $R_\pi/P_\pi$ )**

1.  $a \langle R_\pi \rangle := \langle a \langle r_\pi \rangle, a \langle \sim r_\pi \rangle \rangle$
2.  $a \langle r_\pi \rangle := r_\pi \cup \{ \langle w, w' \rangle \in s^2 \mid w \in Ch_s(R_\pi) \cap a \ \& \ w' \notin Ch_s(R_\pi) \cap a \}$
3.  $a \langle \sim r_\pi \rangle := s^2 - a \langle r_\pi \rangle^2$

**Definition 13 (Must)**

$$S[\text{Must}(\phi)] = \begin{cases} Req_\phi(S) & \text{if } \forall s^\pi \in Req_\phi(S), \forall a \in alt_S(\phi): Ch_s(R_\pi) \subseteq a \\ \{\emptyset^\pi \mid s^\pi \in S\} & \text{otherwise} \end{cases}$$

**Remark 6**  $\text{Must}\phi$  first performs a shift — the *if*-clause of Definition 13 quantifies over  $s^\pi \in Req_\phi(S)$  rather than  $s^\pi \in S$  — and then performs a test on this shifted state. It shifts to a state where  $\phi$  is required, and tests that for every substate, all of  $\phi$ 's alternatives are entailed by the Choice worlds in that substate. If the test is failed, each substate's information is reduced to  $\emptyset$ . If the test is passed,  $\phi$  becomes required by making each of its alternatives preferred in  $R_\pi$  and  $P_\pi$  (Definition 14).

**Definition 14 (Requiring  $\phi$  in  $S$ ,  $a$  in  $\pi$ )**

1.  $Req_\phi(S) := \{s^{Req_a(\pi)} \mid s^\pi \in S \ \& \ a \in alt_S(\phi)\}$
2.  $Req_a(\pi) := \langle a \langle R_\pi \rangle, a \langle P_\pi \rangle \rangle$

**Definition 15 (Negation)** Reading  $s_i^{\pi_i} - s_j$  as  $(s_i - s_j)^{\pi_i}$ :  
 $S[-\phi] = \{s^{\pi} \downarrow \phi - \cup alt_{\{s^{\pi}\}}(\phi) \mid s^{\pi} \in S\}$

1.  $\pi \downarrow \phi := \langle R_{\pi} \uparrow \phi, P_{\pi} \downarrow \phi \rangle$ 
  - a.  $R_{\pi} \uparrow \phi := \langle (r_{\pi} - r(\phi)) \cup p(\phi)^{-1}, (\sim_{\pi} \cup r(\phi) \cup r(\phi)^{-1}) - (p(\phi) \cup p(\phi)^{-1}) \rangle$
  - b.  $P_{\pi} \downarrow \phi := \langle (p_{\pi} - p(\phi)) \cup p(\phi)^{-1}, \approx_{\pi} \rangle$
2.  $r(\phi) := \{ \langle w, w' \rangle \in r_{\pi_i} \mid s^{\pi_i} \in \mathbf{0}[\phi] \}$
3.  $p(\phi) := \{ \langle w, w' \rangle \in p_{\pi_i} \mid s^{\pi_i} \in \mathbf{0}[\phi] \}$

**Remark 7** The appearance of this definition belies its simplicity. Subtracting  $\cup alt_{\{s^{\pi}\}}(\phi)$  from  $s$  recreates the familiar effect of removing the  $\phi$ -worlds from  $s$ .  $\pi \downarrow \phi$  is the result of removing  $\phi$ -preferences from  $\pi$ . This is done in clauses 1a and 1b in slightly different ways for requirements and permissions. For requirements, one removes any requirement preferences  $\phi$  would add to  $\mathbf{0}$  i.e.,  $r(\phi)$ . One has to restore relations of indifference between these worlds, which is what  $(\sim_{\pi} \cup r(\phi) \cup r(\phi)^{-1})$  accomplishes. Additionally, one must add to the requirements the inverse of any preferences  $\phi$  would add to  $\mathbf{0}$  i.e.,  $p(\phi)$ , and remove relations of indifference between these worlds. This is needed to ensure that  $\neg \text{May } \phi \models \text{Must } \neg \phi$ . Removing preferences from the permissions proceeds similarly in clause 1b, but does not have the added complexity since  $\neg \text{Must } \phi$  does not need to p-entail  $\text{May } \neg \phi$ .

**Remark 8** Why does  $\neg \phi$  remove preferences  $\phi$  would add to  $\mathbf{0}$ , rather than preferences  $\phi$  would add to the input state  $S$ ? When  $S = \mathbf{0}[\text{May } A]$  and one considers  $S[-\text{May } A]$  it is clear that  $\text{May } A$  won't add any preferences to  $S$ . Thus negation wouldn't have any preferences to remove, and  $\mathbf{0}[\text{May } A]$  would counterintuitively p-support  $\neg \text{May } A$ .

**Definition 16 (Informational Support, Consequence)**

1.  $S \models \phi \iff c_S = c_{S[\phi]}$ , where  $c_S := \cup \{s \mid s^{\pi} \in S\}$
2.  $\phi_1, \dots, \phi_n \models \psi \iff \forall S: c_S[\phi_1] \dots [\phi_n] \models \psi$

**Definition 17 (Informational Consistency)**  $\exists S: S \models \phi_1, \dots, S \models \phi_n \ \& \ c_S \neq \emptyset$

**Definition 18 (Practical Support, Consequence)**

1.  $S \models \phi \iff \prod_S = \prod_{S[\phi]}$ , where  $\prod_S := \{ \pi \mid s^{\pi} \in S \ \& \ s \neq \emptyset \}$
2.  $\phi_1, \dots, \phi_n \models \psi: \forall S: S[\phi_1] \dots [\phi_n] \models \psi$



**Remark 9** The definition excludes  $\pi$ 's that feature only in substates with the empty information state. This allows the logic to validate disjunctive syllogism when one of the disjuncts is a deontic modal e.g.,  $\text{May} A \vee B$  and  $\neg B$  so  $\text{May} A$ .

**Definition 19 (Practical Consistency)**

$\exists S: S \models \phi_1, \dots, S \models \phi_n \ \& \ \cup \{Ch_s(R_\pi) \mid s^\pi \in S\} \neq \emptyset$

**Remark 10** Just as Definition 17 rules out the irrational information state  $\emptyset$ , Definition 19 rules out practically irrational states, namely ones which do not have at least one  $\pi$  that motivates the agents to choose at least one world.

**References**

- Aher, Martin. 2012. Free choice in deontic inquisitive semantics (DIS). In M. Aloni, V. Kimmelman, F. Roelofsen, G.W. Sassoon, K. Schulz & M. Westera (eds.), *Logic, Language and Meaning: 18th Amsterdam Colloquium 7218 Lecture Notes in Computer Science*, 22–31. Berlin: Springer. doi:10.1007/978-3-642-31482-7\_3.
- Aloni, Maria. 2007. Free choice, modals and imperatives. *Natural Language Semantics* 15(1). 65–94. doi:10.1007/s11050-007-9010-2.
- Alonso-Ovalle, Luis. 2006. *Disjunction in alternative semantics*. Amherst, MA: UMass Amherst PhD dissertation. <http://semanticsarchive.net/Archive/TVkY2ZIM/>.
- Asher, Nicholas & Daniel Bonevac. 2005. Free choice permission is strong permission. *Synthese* 145(3). 303–323. <http://www.jstor.org/stable/20118599>.
- Barker, Chris. 2010. Free choice permission as resource-sensitive reasoning. *Semantics and Pragmatics* 3(10). 1–38. doi:10.3765/sp.3.10.
- van den Berg, Martin. 1996. *Some aspects of the internal structure of discourse: The dynamics of nominal anaphora*. Amsterdam: ILLC, University of Amsterdam PhD dissertation.
- Brasoveanu, Adrian. 2008. Donkey pluralities: Plural information states versus non-atomic individuals. *Linguistics and Philosophy* 31. 129–209. doi:10.1007/s10988-008-9035-0.
- Cariani, Fabrizio. 2013. ‘Ought’ and resolution semantics. *Noûs* 47(3). 534–558. doi:10.1111/j.1468-0068.2011.00839.x.
- Chemla, Emmanuel & Lewis Bott. 2014. Processing inferences at the semantics/pragmatics frontier: Disjunctions and free choice. *Cognition* 130(3). 380–396. doi:10.1016/j.cognition.2013.11.013.
- Chierchia, Gennaro. 2006. Broaden your views: Implicatures of domain widening and the ‘logicality’ of language. *Linguistic Inquiry* 37(4). 535–590. doi:10.1162/ling.2006.37.4.535.

- Eckardt, Regine. 2007. Licensing *or*. In Uli Sauerland & Penka Stateva (eds.), *Presupposition and Implicature in Compositional Semantics*, 34–70. London: Palgrave Macmillan. doi:10.1057/9780230210752\_3.
- Fox, Danny. 2007. Free choice disjunction and the theory of scalar implicature. In Uli Sauerland & Penka Stateva (eds.), *Presupposition and Implicature in Compositional Semantics*, 71–120. New York: Palgrave Macmillan. doi:10.1057/9780230210752\_4.
- Franke, Michael. 2009. *Signal to act: Game theory in pragmatics*. Amsterdam: ILLC, University of Amsterdam PhD dissertation.
- Franke, Michael. 2011. Quantity implicatures, exhaustive interpretation, and rational conversation. *Semantics and Pragmatics* 4(1). 1–82. doi:10.3765/sp.4.1.
- Fusco, Melissa. 2015. Deontic modality and the semantics of choice. *Philosophers' Imprint* 15(28). 1–27. <http://hdl.handle.net/2027/spo.3521354.0015.028>.
- Geurts, Bart. 2005. Entertaining alternatives: Disjunctions as modals. *Natural Language Semantics* 13(4). 383–410. doi:10.1007/s11050-005-2052-4.
- Kamp, Hans. 1973. Free choice permission. *Proceedings of the Aristotelian Society* 74. 57–74. doi:10.1093/aristotelian/74.1.57.
- Kamp, Hans. 1978. Semantics versus pragmatics. In Franz Guenther & S.J. Schmidt (eds.), *Formal Semantics and Pragmatics for Natural Languages*, 255–287. Dordrecht: D. Reidel Pub. Co.
- Klinedinst, Nathan. 2006. *Plurality and possibility*: UCLA PhD dissertation.
- Kratzer, Angelika. 1981. The notional category of modality. In Hans-Jürgen Eikmeyer & Hannes Rieser (eds.), *Words, Worlds and Contexts*, 38–74. Berlin: Walter de Gruyter.
- Lewis, David K. 1979. A problem about permission. In Esa Saarinen, Risto Hilpinen, Ilkka Niiniluoto & Merrill Province Hintikka (eds.), *Essays in Honour of Jaakko Hintikka*, 163–175. Dordrecht: D. Reidel Pub. Co.
- McNamara, Paul. 2010. Deontic logic. In Edward N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy*, Summer 2010 edn. <http://plato.stanford.edu/archives/sum2010/entries/logic-deontic/>.
- Nouwen, Rick. 2003. *Plural pronominal anaphora in context*. Utrecht, The Netherlands: University of Utrecht PhD dissertation. <http://www.let.uu.nl/~Rick.Nouwen/personal/>.
- Pearl, Judea. 2009. *Causality: Models, Reasoning, and Inference*. Cambridge, England: Cambridge University Press 2nd edn.
- van Rooij, Robert. 2000. Permission to change. *Journal of Semantics* 17(2). 119–143. doi:10.1093/jos/17.2.119.
- van Rooij, Robert. 2010. Conjunctive interpretation of disjunction. *Semantics and Pragmatics* 3(11). 1–28. doi:10.3765/sp.3.11.
- Schulz, Katrin. 2005. A pragmatic solution for the paradox of free choice permission.

Expressing permission

- Synthese* 147(2). 343–377. doi:10.1007/s11229-005-1353-y.
- Simons, Mandy. 2005. Dividing things up: The semantics of or and the modal/or interaction. *Natural Language Semantics* 13(3). 271–316. doi:10.1007/s11050-004-2900-7.
- Sloman, Steven. 2005. *Causal Models: How People Think About the World and Its Alternatives*. New York: OUP.
- Stalnaker, Robert C. 1981. A defense of conditional excluded middle. In William L. Harper, Robert Stalnaker & Glenn Pearce (eds.), *Ifs: Conditionals, Belief, Decision, Chance, and Time*, 87–104. Dordrecht: D. Reidel Publishing Co.
- Starr, William B. 2016a. Dynamic expressivism about deontic modality. In N. Charlow & M. Chrisman (eds.), *Deontic Modality*, 355–394. New York: Oxford University Press. doi:10.1093/acprof:oso/9780198717928.003.0013.
- Starr, William B. 2016b. A preference semantics for imperatives. Under Review at *Semantics & Pragmatics*.
- Tieu, Lyn, Jacopo Romoli, Peng Zhou & Stephen Crain. 2016. Children’s knowledge of free choice inferences and scalar implicatures. *Journal of Semantics* 33(2). 269–298. doi:10.1093/jos/ffv001.
- Van Fraassen, Bas C. 1966. Singular terms, truth-value gaps and free logic. *Journal of Philosophy* 3. 481–495. doi:10.2307/2024549.
- Veltman, Frank. 1996. Defaults in update semantics. *Journal of Philosophical Logic* 25(3). 221–261. doi:10.1007/BF00248150.
- von Wright, Georg Henrik. 1968. Deontic logic and the theory of conditions. *Crítica: Revista Hispanoamericana de Filosofía* 2(6). 3–31. <http://www.jstor.org/stable/40103910>.
- Willer, Malte. 2015. Simplifying counterfactuals. In Thomas Brochhagen, Floris Roelofsen & Nadine Theiler (eds.), *20th Amsterdam Colloquium*, 428–437. Amsterdam: ILLC. <http://semanticsarchive.net/Archive/mVkOTk2N/AC2015-proceedings.pdf>.
- Zimmermann, Thomas Ede. 2000. Free choice disjunction and epistemic possibility. *Natural Language Semantics* 8(4). 255–290. doi:10.1023/A:1011255819284.

William B. Starr  
218 Goldwin Smith Hall  
Cornell University  
Ithaca, NY 14853  
[will.starr@cornell.edu](mailto:will.starr@cornell.edu)