

# Deontic Modality

EDITED BY  
Nate Charlow and  
Matthew Chrisman

**OXFORD**  
UNIVERSITY PRESS

## OXFORD

UNIVERSITY PRESS

Great Clarendon Street, Oxford, OX2 6DP,  
United Kingdom

Oxford University Press is a department of the University of Oxford.  
It furthers the University's objective of excellence in research, scholarship,  
and education by publishing worldwide. Oxford is a registered trade mark of  
Oxford University Press in the UK and in certain other countries

© the several contributors 2016

The moral rights of the authors have been asserted

First Edition published in 2016

Impression: 1

All rights reserved. No part of this publication may be reproduced, stored in  
a retrieval system, or transmitted, in any form or by any means, without the  
prior permission in writing of Oxford University Press, or as expressly permitted  
by law, by licence or under terms agreed with the appropriate reprographics  
rights organization. Enquiries concerning reproduction outside the scope of the  
above should be sent to the Rights Department, Oxford University Press, at the  
address above

You must not circulate this work in any other form  
and you must impose this same condition on any acquirer

Published in the United States of America by Oxford University Press  
198 Madison Avenue, New York, NY 10016, United States of America

British Library Cataloguing in Publication Data

Data available

Library of Congress Control Number: 2015960289

ISBN 978-0-19-871792-8

Printed in Great Britain by  
Clays Ltd, St Ives plc

Links to third party websites are provided by Oxford in good faith and  
for information only. Oxford disclaims any responsibility for the materials  
contained in any third party website referenced in this work.

## Contents

<i>List of Contributors</i>	vii
Introduction	1
<i>Nate Charlow and Matthew Chrisman</i>	
1. Deontic Modals and Probabilities: One Theory to Rule Them All? <i>Fabrizio Cariani</i>	11
2. Decision Theory: Yes! Truth Conditions: No! <i>Nate Charlow</i>	47
3. Linguistic and Philosophical Considerations on Bayesian Semantics <i>Daniel Lassiter</i>	82
4. Contextualism about Deontic Conditionals <i>Aaron Bronfman and J. L. Dowell</i>	117
5. Objective and Subjective ‘Ought’ <i>Ralph Wedgwood</i>	143
6. ‘Ought’: Out of Order <i>Stephen Finlay</i>	169
7. On a Shared Property of Deontic and Epistemic Modals <i>Jessica Rett</i>	200
8. Modalities of Normality <i>Seth Yalcin</i>	230
9. Extreme and Non-extreme Deontic Modals <i>Paul Portner and Aynat Rubinstein</i>	256
10. Rationalization and the Ross Paradox <i>Benj Hellie</i>	283
11. Dynamic Foundations for Deontic Logic <i>Malte Willer</i>	324
12. Dynamic Expressivism about Deontic Modality <i>William B. Starr</i>	355
13. Metanormative Theory and the Meaning of Deontic Modals <i>Matthew Chrisman</i>	395
<i>Index</i>	431

## 12

Dynamic Expressivism about  
Deontic Modality*William B. Starr*

As a first approximation, expressivism about deontic discourse holds that an utterance of *Zoyd must share* expresses certain motivational attitudes towards sharing, but without referring to and describing, or representing, those attitudes—that would amount to a form of subjectivism that makes disagreement difficult to capture (Moore, 1912).<sup>1</sup> (I can hardly disagree with your descriptions of your attitudes.) Some philosophers have sought to connect deontic discourse to motivational attitudes such as desire because they think it is the only way of explaining how deontic language and thought motivate us to do things: passive descriptions of things out there in the world lack the requisite motivational gravitas (Mackie, 1977). So, the aim of expressivism is to ground deontic discourse in motivational psychology without making its subject matter motivational psychological states.<sup>2</sup> From this perspective, expressivism is only distinctive and interesting to the degree that motivational and representational psychological states are distinguishable. I believe this presents a dilemma for expressivists. If representational thought and talk *isn't* substantially different from its motivational counterparts, then expressivism becomes elusive and uninteresting. If representational thought and talk *is* substantially different from its motivational counterparts, the existing tools of truth-conditional semantics tailored for representational purposes should not work for motivational thought and talk. After all, if motivational attitudes are quite different one would expect their kinematics, hence their expression and hence their semantics to also differ. It seems to follow that if expressivism is an interesting and distinctive position, it faces the Frege-Geach problem. It is not clear, as Frege (1923) stressed, what logical compounds of non-representational sentences mean and,

<sup>1</sup> Note that even on this approximation, expressivism should not be equated with a non-propositional semantics. A non-propositional semantics, e.g. Hamblin's (1973) for interrogatives, needn't involve motivational attitudes. Further, one could offer a propositional expressivist semantics as long as one had a story about how those contents are connected to motivational attitudes rather than representational ones. For discussion see Chrisman (2012), Schroeder (2013), Ridge (2014), and Charlow (2014).

<sup>2</sup> For a more detailed version of this narrative see Schroeder (2008a, Chapters 1–2).

as Geach (1965) stressed, it is not clear how non-representational sentences could bear relations of consequence and consistency.

One goal of this chapter is to press this dilemma against recent expressivists that use the tools of truth-conditional semantics to model expressive meaning and communication (Gibbard, 2003; Dreier, 2009; Yalcin, 2011, 2012; Silk, 2014) (§§2–4). I will argue that these accounts either rely on representational concepts that are inappropriate for expressive motivational discourse, or they end up blurring the distinctions that make expressivism a distinctive and interesting thesis. In doing so, I will propose a new way of decomposing the negation problem (§1) and argue that expressivism cannot succeed without a distinctive theory of expressive communication (§2). My other goal is more positive. I will assume for the purposes of the argument set out that representational language inherits its distinctive semantic properties from the function of representational attitudes: to track how the world is. Further, I will assume that motivational attitudes have a very different function, namely to motivate choice and action, and that motivational language inherits rather different semantic properties from this motivational function. This is not so much because I am confident in these assumptions, but because they provide the most interesting and challenging setting to defend expressivism. Against this backdrop, I will use dynamic semantics, a relatively recent semantic framework discussed in computer science and linguistics, to offer a model of expressive and representational language where both kinds of discourse are seamlessly integrated, and where uniform explanations of compositionality and logical relations are forthcoming (§§5–6). Dynamic semantics provides an algebra of processes rather than contents. I will show how this embodies a more general conception of meaning and logic that yields motivational and representational meaning as special cases.<sup>3</sup>

Here at the outset, I should clarify that I will not be offering a positive linguistic or philosophical argument in favor of expressivism. Yet, linguists and empirically-minded philosophers of language may find something of interest here. The empirical investigation of 'expressive meaning' is a burgeoning field, whose object of study has remained somewhat difficult to classify and integrate in compositional systems.<sup>4</sup> The perspective on expressivism, communication and compositional semantics offered here should be useful for these purposes. The semantics developed here also makes interestingly different predictions than more orthodox approaches (e.g. Kratzer, 1991), thereby connecting to a broader debate over non-truth-conditional semantics

<sup>3</sup> This emphasis on processes (acceptance, rejection) rather than content is a theme in Gibbard (2003), but is not embraced in Gibbard's formalism. Dynamic semantics provides the tools for doing so. Alwood (2015) pursues a similar dynamic approach and situates it in more detail with work in metaethics. I aim to provide a more thorough technical implementation here and connect such an approach to the issue of communication in expressivism. One virtue of this implementation is not having to distinguish, as Alwood (2015) must, between two kinds of negation. I will offer one semantics for negation which models both propositional and expressive negation.

<sup>4</sup> E.g. evidentiality (Murray, 2010; Faller, 2002), attention (Bittner, 2001), narrative structure (Hobbs, 1990), slurs (Anderson and Lepore, 2013) and other phenomena (Potts, 2007; McCready, 2012).

(Veltman, 1996, p. 235; Portner, 2009, p. 104)<sup>5</sup> and expressivist pragmatics (Yalcin, 2011, 2012) for modality. Throughout, I will be assuming that expressivism is a semantic thesis about the linguistic meaning of sentences which challenges the idea that all communication proceeds by referring to things and describing them. Along the way (§4), I will present difficulties for the idea that expressivism is merely a pragmatic thesis (Yalcin, 2011, 2012). Unfortunately, I will not have space to discuss important recent work that pitches expressivism as a metasemantic thesis (Chrisman, 2012; Ridge, 2014) about how certain expressions get their traditional meanings, or what it means to ascribe those traditional meanings to those expressions (Carballo, 2014). It is interesting to note, however, that the linguistic meanings ascribed in my dynamic approach are abstract objects that model the processes which Lewis (1969) and Grice (1968) took to endow sentences with their contents. However, in taking these processes themselves to be the meanings the dynamic approach arrives at additional semantic resources for solving the Frege-Geach problem.<sup>6</sup> Only further, more general, work on metasemantics will reveal the comparative virtues of semantic and metasemantic approaches.

## 1 The Virtues of Dynamic Expressivism: a Sketch

Subsequent sections present dynamic expressivism and its alternatives in full technical detail. This section offers a less formal sketch of the positions and issues with the hope that this will bring order to the wealth of detail that is to follow. I will start with a simple model of the role representational and motivational states play in rational agency. Then I will use this framework to articulate deontic expressivism and show how it leads to the Frege-Geach problem. This will allow me to say what dynamic expressivism is and how it is better suited to solve this problem than accounts which use the tools of truth-conditional semantics.

According to a familiar, if limited, model of rational agency an agent's beliefs are modeled with a probability space and their values with a utility function. The probability space measures how likely they take certain propositions to be and the utility function measures how valuable they take certain states of affairs to be. Agents make rational choices by selecting those actions that lead to outcomes with the highest 'expected utility': the value of an outcome weighted by how likely the agent takes it to be. It is crucial to note that beliefs and values play quite distinct theoretical roles in this model: agents' decisions are primarily driven by what they value, and tempered only by what they believe. An agent may strive for a very unlikely outcome if it is far more valuable than the other options. Further, if they had no values they would have

<sup>5</sup> In an earlier version of this work I adopted a semantics much closer to Veltman (1996, p. 235) and Portner (2009, p. 104). The new semantics is now a bit of a hybrid between that semantics, one from dynamic logic (van Benthem and Liu, 2007) and Starr's (2013) dynamic semantics for imperatives.

<sup>6</sup> This approach has much in common with Millikan (1984, 2005).

no basis for choice, but if they had no beliefs they could still (naively) choose actions. These different theoretical roles justify rather different assumptions about probabilities and utilities. If an agent has credence 0.6 that Zoyd shared, then that agent is assumed to have credence 0.4 that Zoyd didn't share. By contrast, if Zoyd's sharing has utility 2 for an agent, it is not assumed that Zoyd's sharing gets utility  $t - 2$  where  $t$  is the total amount of utility actions can have for the agent. The assumption for probabilities is reasonable since they capture how the agent represents the world: in the world either Zoyd shared or he didn't. This assumption is not plausible for utilities because they don't represent the world.

Against this backdrop, an expressivist might say that utterances of *Zoyd must share* express utility functions. Perhaps functions where Zoyd's sharing is assigned a high utility relative to some standard. In simple cases, this might serve to bring one agent's values in line with another's. By contrast, the subjectivist would say that deontic utterances of *Zoyd must share* express propositions about utility functions. In simple cases, these utterances would serve primarily to align agent's beliefs (i.e. probability spaces). Though this might result in agents' aligning their values, that's not the primary function of deontic discourse according to the subjectivist. A third position, call it *simple naturalism*, maintains with subjectivism that deontic utterances express beliefs about the world. But it rejects that these beliefs are *about* utility-relative facts: they are simple beliefs about the utility-independent world. Seeing how the Frege-Geach problem arises in this setting will allow me to say what dynamic expressivism is and how it differs from other accounts.

Frege (1879, 1923) famously proposed the standard explanation of how representational content behaves in a compositional way. But, less famously, Frege (1923) also contended that those explanations do not work for non-representational meaning. The argument was simple: the meaning of negation is exhausted by its role of turning a true content into a false one and vice versa. This explains why any sentence and its negation are inconsistent, i.e. cannot both be true. Here's the problem for expressivists: if (1a) does not represent the world, then how does negating it do anything meaningful, let alone inconsistent with (1a)?

- (1) a. Zoyd must share.
- b. It's not the case that Zoyd must share.

This is the Frege-Geach Problem, whose title reflects its application by Geach (1965) to non-representational theories of moral language.<sup>7</sup> To solve this problem, it seems that expressivists must say what non-representational attitude (1a) expresses, how the attitude expressed by (1b) is generated by negating the expression of an attitude by (1a) and what it is for two non-representational attitudes to be inconsistent. Parallel issues arise with other connectives and explaining entailment patterns.

<sup>7</sup> See also Schroeder (2008d) and references therein.

One can see the challenge more precisely from within the simple model of rational agency above. There is a relationship between a rational agent’s credence in *Zoyd shared* and *Zoyd didn’t share*: no credence should be assigned to the possibility of both being true. This, as mentioned above, follows from background assumptions about truth, representation and the functional role of beliefs—more on that in §3. It is less obvious that such a relationship exists between how much a rational agent values *Zoyd sharing* and how much they value *Zoyd not sharing*. One can certainly imagine a scenario where both outcomes have a high utility for the agent. Further, the relationship between how an agent values *Zoyd sharing* and how they value *Zoyd not sharing* speaks to the consistency of (2a) and (2b), rather than (1a) and (1b).<sup>8</sup>

- (2) a. *Zoyd must share.*
- b. *Zoyd must not share.*

Although (2b) intuitively entails (1b), the opposite is not true. Note that it is only consistent to follow up (1b) with *but Zoyd may share*. This leads to an important observation: there are really two ‘negation problems’ for the expressivist. Suppose one *can* successfully analyze (2a) and (2b) as expressing ‘motivationally conflicted’ values: being in favor of sharing and being in favor of not sharing. This would be a solution to the *internal negation problem*. Since one can endorse (1b) without endorsing (2b), an analysis of (2b) is still needed. This leaves open the possibility that a different account of the conflict between (1a) and (1b) will be needed. Call this the *external negation problem*. I believe the two problems require different solutions. While this point will be argued for more precisely later, the basic idea can be articulated with the resources currently on hand. It will then serve as a basis for contrasting dynamic expressivism with others (Gibbard, 2003; Dreier, 2009; Yalcin, 2011, 2012; Silk, 2014).

Even if (2a) and (2b) express motivationally conflicted states of mind, it is hard to see how the same is true of (1a) and (1b). Initially, this is because it is hard to see how a sentence with external negation expresses a particular motivational attitude at all. If one thinks of *Zoyd must share* as serving to instill pro-sharing utilities in the hearer, then *It is not true that Zoyd must share* rejects instilling those utilities. But what particular utilities does this rejection promote? I believe that focusing on this question gets to the heart of the problem for expressivists. Previous attempts to say what attitude is expressed have failed to render (1a) and (1b) inconsistent in a way that is clearly non-representational.

Gibbard (2003, pp. 71–5) offers the following expressivist model:

- (3) a. ‘You must share’ expresses: Agreeing with sharing
- b. ‘You must not share’ expresses: Disagreeing with sharing
- c. ‘It is not the case that you must share’ expresses: Disagreeing with agreeing with sharing

<sup>8</sup> Unwin (2001) highlights the importance of this distinction in the context of expressivism.



This model does not say anything more about what disagreeing and agreeing are which explains why agreeing with  $X$  and disagreeing with  $X$  are in rational tension. In the setting above, this amounts to saying that *You must share* and *It is not the case that you must share* express conflicting utility functions without saying what utility functions these are and how exactly those functions or their acceptance conflict. Some have worried whether such a semantic theory is adequately explanatory (e.g. Schroeder, 2008a, d). §2 will press that concern in another way by arguing that expressivists need to offer a positive account of how expressive communication works and how it differs from representational communication. It is hard to see how this could work without characterizing representational and expressive communication as importantly different processes.

Building on Dreier (2006, 2009), Silk (2014) develops what might be called **preference expressivism**, which formulates a deontic semantics in terms of preference orderings. In Starr (2013) I pursue roughly the same strategy to explain the incompatibility of contrary imperatives such as *Dance!* and *Don't dance!*, without assuming that imperatives are representational. This aims to fill the gaps in Gibbard's approach. The basic idea is that deontic sentences express preferences. Much like utilities, the function of preferences is to motivate choice.<sup>9</sup> Crucially, this means that preferences which *prevent* a choice will be dysfunctional. The idea is then that two deontic sentences can be inconsistent if they express preferences which, taken together, are dysfunctional in this way. A theory along these lines offers an interesting expressivist solution of the internal negation problem. But I do not think it can solve the external negation problem.

A (strict) preference for you sharing could be modeled as an ordering which ranks any world  $w_1$  where you share over every world  $w_0$  where you don't:  $w_1 \succ w_0$ . Now, consider a semantics on which:

- (4) a. 'You must share' expresses:  $w_1 \succ_1 w_0$
- b. 'You must not share' expresses:  $w_0 \succ_2 w_1$

For now, set aside how the meaning for (4b) is determined from the meaning of (4a) and a semantics for negation. Here is how one might explain the inconsistency of these two sentences on such a semantics.<sup>10</sup> The first step is to determine what it is to take these two preference orderings 'together'. *Unioning* them into a bigger preference ordering is the most natural idea:  $w_1 \succ_3 w_0, w_0 \succ_3 w_1$ .<sup>11</sup> It is a platitude about choice that you shouldn't choose  $A$  if you strictly prefer  $B$  to  $A$  and you can choose  $B$


<sup>9</sup> The earlier utility expressivist proposal was that **Must(A)** expresses utility functions which assign a utility to **A**-worlds that is high according to some standard. By contrast, preference expressivism says: the utility of **A**-worlds is greater than the utility of  $\neg$ **A**-worlds.

<sup>10</sup> Neither Dreier (2006, 2009) nor Silk (2014) explicitly formulates the decision-theoretic constraints on preference and choice that they take to constitute preferential coherence. Following Starr (2013), I assume in the main text above a non-dominance principle about choice and tie this to the need for acyclic preferences. I do not know whether this is what Dreier (2006, 2009) and Silk (2014) had in mind.

<sup>11</sup> Note that  $\succ_1$  is the set  $\{(w_1, w_0)\}$  and  $\succ_2$  is the set  $\{(w_0, w_1)\}$ . The fact that union is the operation appropriate to combining preferences is discussed by van Benthem and Liu (2007). This is fairly clear from thinking about how to combine the orderings from **Must(A)** and **Must(B)**. Intersection predicts that  $A \wedge B$ -worlds,  $\neg A \wedge B$ -worlds and  $A \wedge \neg B$ -worlds are best, while union predicts that only  $A \wedge B$ -worlds are best.

instead. This means that any cyclic strict preference ordering is irrational. Something is preferred to every alternative, so no alternative can be rationally chosen. This is exactly the situation in  $\succ_3$ . Since it fails to motivate a choice, it is motivationally defective. This seems like exactly the kind of explanation of inconsistency that expressivists seek. But, I claim, it does not work for the external negation problem.

Following standard accounts of modality, Silk (2014) proposes that  $\text{Must}(A)$  requires all the most preferred worlds to be  $A$ -worlds, while  $\neg\text{Must}(A)$  requires some of the most preferred worlds to be  $\neg A$ -worlds. The issue is that taking together—unioning—two preference orderings that meet these respective descriptions does not always yield a dysfunctional one. To illustrate this, consider a model in which we are concerned only with  $A$  and  $B$  worlds. Using a capital letter to indicate truth and a lowercase letter for falsity,  $\text{Must}(A)$  can express an ordering such as  $\succ_1$  while  $\neg\text{Must}(A)$  can express an ordering such as  $\succ_2$ :  $w_{AB}$  and  $w_{Ab}$  are the most preferred worlds according to  $\succ_1$ , and  $w_{aB}$  is among the most preferred worlds according to  $\succ_2$ . As the graphs illustrate, there is nothing dysfunctional about taking these preferences together. Together, they will motivate the agent to choose  $w_{AB}$  or  $w_{Ab}$ . These combined preferences are not motivationally dysfunctional.



From the dynamic perspective, the preference expressivist went wrong in two assumptions: (i) accepting  $\neg\text{Must}(A)$  is the same kind of process as accepting  $\text{Must}(A)$  and (ii) consistency is a relationship between contents rather than processes. On the dynamic account developed in §6, the meaning of an expressive sentence is characterized in terms of how accepting it changes a preference ordering.  $\text{Must}(A)$  introduces a strict preference for all live  $A$ -worlds over  $\neg A$ -worlds, and tests that doing so makes  $A$  a practical necessity: all the choices it motivates result in  $A$ -worlds. (The success of this test will depend on which preferences have been previously introduced.)  $\neg\text{Must}(A)$ , on the other hand, does *not* introduce a preference at all, so its inconsistency with  $\text{Must}(A)$  is not an inconsistency of preference.  $\neg\text{Must}(A)$  changes a preference ordering by removing any preferences that would be added by  $\text{Must}(A)$ .  $\text{Must}(A)$  and  $\neg\text{Must}(A)$  are *dynamically inconsistent*: since one sentence adds what the other removes there is no ordering which contains the effects both have on a preference ordering. But that is crucially different from saying that they lead to orderings which, combined in the standard way, are rationally defective.  $\text{Must}(A)$  and

$\neg\text{Must}(A)$  are inconsistent in the same way that certain actions are. Consider a robot that rolls around a bar to each table. It fills an empty shot glass with vodka and it vacuums up the contents of non-empty shot glasses. Then it cruises on. The action of pouring a shot and the action of vacuuming one are incompatible in the sense that no state of a given glass that directly results from the robot's actions contains the effects of both actions. But this is not because the effects themselves are incompatible: the empty space in the shot glass and the shot of vodka are clearly compatible in that one can fill the other. As will be discussed, a non-dynamic account does not make room for the idea that  $\phi$  and  $\neg\phi$  have radically different communicative effects. On a non-dynamic account, negation takes one content (e.g. preference-ordering or proposition) and yields another of the same kind. Their communicative effects are modeled by a single pragmatic process which combines contents of that kind (the previously accepted contents incremented with the new content).

The dynamic approach to the *external* negation problem combines elegantly with the preference expressivists' approach to the *internal* negation problem. In §6, I define dynamic preferential inconsistency as follows:  $\phi$  and  $\psi$  are dynamically preferentially inconsistent just in case there is no single practically rational preference ordering that contains the effects on preferences that both would have. In the case of  $\text{Must}(A)$  and  $\text{Must}(\neg A)$ , there exists a preference ordering that contains preferences both introduce, but it is irrational because those preferences are cyclic.  $\text{Must}(A)$  and  $\neg\text{Must}(A)$  are different: there is no preference ordering which contains the effects on preferences that each sentence produces. If an ordering has the preferences  $\text{Must}(A)$  adds, then updating with  $\neg\text{Must}(A)$  would remove them. If an ordering lacks the preferences  $\neg\text{Must}(A)$  would remove, then updating with  $\text{Must}(A)$  would add them.

Formalizing the dynamic account requires a logic of *actions* rather than *contents*, one which speaks to questions of how connectives combine actions and how actions can bear logical relations to each other. Fortunately, logics of this kind—called *dynamic logics* (Harel et al., 2000)—have been extensively developed in computer science and have inspired new semantic approaches to natural language (Heim, 1983; Kamp, 1981; Groenendijk et al., 1996). The basic idea of this approach is to formally describe how accepting each type of sentence changes an agent's state of mind. Surprisingly, one can reformulate classical logic in this way and then add non-classical operators to produce operations that do not amount to operations on content (§5). Hence, this is not an alternative approach to semantics, but a more general one. By importing the dynamic semantics for disjunction from Starr (2013) below, I will also be able to navigate a challenge presented by Schroeder (2015) for previous dynamic analyses and analyses such as Yalcin's (2011, 2012): they incorrectly predict that believing  $\text{Must}(A) \vee B$  entails either believing  $\text{Must}(A)$  or believing  $B$ .<sup>12</sup> As I will emphasize, parallel observations hold for what disjunctions *communicate*. This connection between an expressivist

<sup>12</sup> Schroeder (2015) makes a more general observation: this prediction holds even if both disjuncts are expressive. The semantics in §6 also blocks this result.

account of communication and solving the negation and disjunction problems is surprising. As I've argued above, solving the negation problem requires  $\text{Must}(A)$  and  $\neg\text{Must}(A)$  to do very different communicative jobs. This is unlike  $A$  and  $\neg A$ , so it seems to follow that the viability of expressivism relies crucially on the idea that expressive communication works differently to representational communication. In the following section, I argue that the need for a distinctive expressivist model of communication follows from more general features of the expressivist program. This will also allow me to consider what is distinctive about expressive communication.

## 2 Expressivism and Communication

Philosophers who maintain that some bit of discourse does not work by describing distinctions in the world often rally under the label *expressivism*, while those who disagree rally under *descriptivism*. Defining expressivism could easily occupy a chapter of its own, so I will not authoritatively impose such a definition on the discussion here. Instead, I wish to consider three characteristic expressivist theses:<sup>13</sup>

### Expressivist Theses

1. *Communication*: 'To express a state of mind is not to say that one is in it' (Gibbard, 1986, p. 473).
2. *Explanation*: 'The semantic properties of sentences are to be explained, fundamentally, in terms of properties of the attitudes conventionally expressed by utterances of those sentences' (Silk, 2014, §1).
3. *Non-representation*: The states of mind expressed by sentences are non-representational, and, more specifically, motivational.

As suggested above, the point of the explanatory thesis is to replace distinctions 'in the world' with psychological distinctions, all while keeping the subject matter of deontic discourse non-psychological. Viewed in this light, a surprising fact becomes apparent: the explanatory expressivist thesis assumes the one about communication. Suppose the communication thesis did not hold, and consider whether the explanatory one could. Without the communication thesis, expressing one's preference for sharing is to describe oneself as preferring sharing. Then one's attitude of preference is involved in the same way as books are involved in the sentence *Books are made from paper*. The semantic fact that *You must share* and *You must not share* are inconsistent is no more explained by the properties of your preferences than the parallel inconsistency between *Books are made from paper* and *Books are not made from paper* is explained by the properties of books. In both cases, one needs things around—books, preferences—and distinctions between them. So vindicating the explanation thesis is at best a partial

<sup>13</sup> The quotation from Silk (2014, §1) nicely captures the operative expressivist thesis at play in much recent work including Gibbard (2003, p. 7), Schroeder (2008c, p. 576) and Dreier (2006, 2009, p. 97).

vindication of expressivism. Contemporary expressivists have done little to meet this other important objective.<sup>14</sup>

As summarized in §1, the account of Gibbard (2003) declines to positively characterize the content of  $\neg\text{Must}(A)$ . But without such a characterization, it is not possible to positively characterize how an agent's state of mind changes when someone communicates  $\neg\text{Must}(A)$  to them. What are the prospects for supplementing analyses such as those found in Gibbard (2003) and Silk (2014) with an expressivist pragmatics that clearly distinguishes expressing attitudes from expressing representational contents about those attitudes? Yalcin (2012, §10) and Silk (2014) propose similar deontic semantics and Yalcin (2012, §10) sketches an expressivist pragmatics for his account. On that account, expressive communication works in a way that is formally parallel to representational communication. §4 will present a challenge for this account. As suggested in §1, to solve the negation problem one must reject the idea that expressive sentences convey content in the same way as representational sentences. This is because motivational contents combine and conflict in different ways than their representational counterparts. But that suggests that a model which fits the expression of motivational attitudes into the representational model is not the model expressivists need. It is more naturally interpreted as a model of how one conveys information *about* motivational attitudes. This is no accident, I claim. On a robust representational interpretation of truth-conditional semantics, there is an essential reliance on reference. All contents are therefore abstractions of referential relations. This means that one cannot employ a robust interpretation of truth-conditional semantics, and use truth-conditional semantics to either solve the Frege-Geach problem or give an expressivist theory of communication.

Although the expressivist communication thesis is my focus, a word about the explanatory thesis is in order.<sup>15</sup> It is hard to say whether that thesis applies to the view offered here since this view does not proceed by assigning sentences to the attitudes, or states of mind, they express. Further, the semantic properties of *all* sentences could be said to be explained in terms of properties of the mental states they update. For example, in one sense of *inconsistency*, the inconsistency of two descriptive sentences is explained by the fact that an update with both leads to a state of mind whose function is to represent the world, but fails to do so.<sup>16</sup> The inconsistency of two expressive deontic sentences is ultimately explained by the fact that an update with both requires a state of preference whose function is to motivate agents to choose one option over another, but fails to do so. One can recover something in the vicinity of the explanatory thesis

<sup>14</sup> Schroeder (2008b) also highlights this gap and explores several expressivist accounts of 'expression' which pair a sentence with the attitude it expresses. The expressivist accounts I'll consider, like Yalcin (2011), do not make use of such a pairing, and focus on the interpersonal relation of communication rather than the individualistic notion of expression. So while Schroeder (2008b) makes important and related points, its criticism is complementary to that presented here.

<sup>15</sup> See also Silk (2013) and Charlow (2014) who discuss the explanatory thesis.

<sup>16</sup> I will say more about what I take 'explaining inconsistency' to be in §3.

by replacing talk of sentences expressing mental states with updating mental states, and focusing on explanations which appeal to the non-representational dimensions of those states (e.g. attention, preference).

**Dynamic Expressivist Explanatory Thesis** Semantic properties of expressive sentences are explained, fundamentally, in terms of the way they update non-representational dimensions of language users' mental states.

I've said a bit above about why expressivism is a thesis worthy of attention in metaethics, but I think it is also important to say why it is worthy of attention in the study of natural language meaning.

Expressivism may seem like an esoteric philosophical doctrine, but its communication thesis is familiar from commonsense observations. Suppose we're looking at the same field of wild flowers. I've noticed a bee fly on a nearby wild flower and I'm attending to it. I point at it and say, *that's what a bee fly looks like*. I was attending to the bee fly and you are now attending to the bee fly. I've communicated my state of attention to you. But I did not do so by referring to my state of attention and stating a fact about it, i.e. I did not say something like *My state of attention is thus-and-so*. The way I communicated my *state* of attention was very different than the way I communicated the *object* of my attention, to which I referred with a gesture. Further, one can imagine two agents' having identical visual representations of a scene but their attention to be centered on different objects. Since such a difference is likely to impact actions where they coordinate, it seems that successful communication will involve agents attuning their state of attention, despite the fact that they are not referring to those states and talking about them. Indeed, given that agents seem capable of communicating certain features of their mental states without even being aware of them, expressivism in these domains seems inevitable. Attention, of course, is not the only expressive dimension where this difference seems to matter. When I tell you that bee flies lay their eggs in beetle nests and that after hatching their larvae feed on their hosts' eggs, you might express disgust at this parasitism with a grimace. You've thereby communicated your disgust to me.<sup>17</sup> But you did not do so by referring to your state of disgust and stating a fact about it. You did not say anything like *My state of disgust is thus-and-so*. Nothing in your facial expression seems to have brought your affective state into the conversation in the way my pointing brought the bee fly into the conversation, nor does anything in your facial expression describe your affective state. These two examples illustrate the expressivist communication thesis at work in a more general setting. In these cases it seems plausible enough. But neither example involved essentially deontic language, or language at all: gestures were the expressive

<sup>17</sup> Add to the story that I intended to do so if you are inclined to the Gricean orthodoxy that requires an effect to be intended for it to be communicated. Note that the cost of such orthodoxy is positing a higher-order mental state which involves the speaker attending to their state of attention. Expressivism is naturally allied with the various attempts to do without the Gricean intellectualization of communication and convention (e.g. Burge, 1975).

signals. Thus, an expressivist must say how this kind of communication operates in the linguistic realm in general, and the deontic realm in particular. Note, however, that it is not problematic for the expressivist to admit that deontic discourse does not appear to be psychological. On an expressivist view, psychological states are behind the scene but not on stage. In fact this is likely another advantage to be celebrated by the expressivist. The proper characterization of the psychological states relevant to the interpretation of modals requires theoretical sophistication beyond our everyday grasp of our own minds, e.g. orderings or probability spaces.

I will now proceed by presenting an entirely standard truth-conditional semantics for propositional logic and highlighting how its approach to logic, compositionality and communication depends on reference (§3). I then critically examine attempts to use an extension of this framework to solve the negation problem and to provide an expressivist account of communication in §4. That discussion will propel us towards a dynamics semantics. I will first present a dynamic semantics for propositional logic to illustrate the basic ways in which it departs from classical semantics (§5). I will then, in §6, provide a dynamic semantics for deontic discourse and explain how dynamic expressivism tackles the difficulties for other approaches.

### 3 Logical Semantics, Classically

Logical semantics has provided a model of linguistic meaning that is useful in three ways. First, it accounts for how the meanings of complex sentences are built up from the meanings of their parts. Secondly, it provides an account of why some sentences follow from others and why some are inconsistent with others. Thirdly, it provides an account of the basic information communicated by any use of the sentence. All three of these applications rely on exclusively referential semantic concepts. That is, they rely on the idea that linguistic expressions function by referring to and making distinctions between things in the world. To be concrete but keep the formalism simple, I will illustrate this with an intensional (possible-worlds) semantics for a propositional logic.

In intensional logical semantics the meaning of a sentence is usually modeled as a set of possible worlds, namely the ways the world could be if the sentence is true. Call such a set a *proposition*.<sup>18</sup> The aim of a compositional semantics is then to assign each sentence  $\phi$  to a proposition,  $\llbracket \phi \rrbracket$ . But before turning to those details, consider the more intuitive and pertinent account of communication built on this semantics. An idealized but useful model of how representational language is used to communicate is familiar from Stalnaker (1978, 1984): it provides information by excluding possibilities. One agent has some information they want to convey to another. This information can be thought of as ruling out some possibilities  $w_1, \dots, w_n$ . The first agent can then utter

<sup>18</sup> This would be more clear in a first-order system where the 'referent' of an atomic sentence is determined by the referents of its terms and predicate. But that would introduce unnecessary complexities.

a sentence  $\phi$  that excludes those possibilities, i.e.  $w_1, \dots, w_n \notin \llbracket \phi \rrbracket$ , and by combining this information with their  $i$ , the second agent can rule out some possibilities, i.e.  $i \cap \llbracket \phi \rrbracket$ . The important part is that on this view, gaining information amounts to pointing to some ways the world could be and saying that the world isn't those ways. Or, to use the exact phrasing from earlier: it refers to a distinction, say  $\mathfrak{p}$  or  $\neg\mathfrak{p}$ , and describes the world as being on one side of that distinction. An expressivist cannot endorse this account. After all, they hold that deontic discourse communicates those states of mind without referring to them.

The compositional semantics employs a space  $W$  of possible worlds. Though a simplification, one can equate a possible world with an atomic valuation, i.e. a function assigning exactly one truth value (1, 0) to every atomic sentence—thinking of a truth-table containing a column for every atomic sentence, each row specifies a possible world.<sup>19</sup> Atomics then denote the set of worlds where they are true, and the Boolean connectives are treated as complementation, intersection and union, respectively.

**Definition 1 (Compositional Semantics).**

- |   |   |
|---|---|
| (1) $\llbracket \mathfrak{p} \rrbracket = \{w \in W \mid w(\mathfrak{p}) = 1\}$                           | (2) $\llbracket \neg\phi \rrbracket = W - \llbracket \phi \rrbracket$                                   |
| (3) $\llbracket \phi \wedge \psi \rrbracket = \llbracket \phi \rrbracket \cap \llbracket \psi \rrbracket$ | (4) $\llbracket \phi \vee \psi \rrbracket = \llbracket \phi \rrbracket \cup \llbracket \psi \rrbracket$ |

What is the intuitive picture of meaning behind this formalism? An atomic sentence  $\mathfrak{p}$  points to a region of logical space by referring to a particular distinction between worlds,  $w(\mathfrak{p}) = 1$  or  $w(\mathfrak{p}) = 0$ , and describing that world as being a certain way:  $w(\mathfrak{p}) = 1$ . Connectives make it possible to point to regions of logical space that lone atomic sentences cannot reach. In this framework, the only way for the language to express distinctions between psychological states is for it to refer to them. The basic notion of meaning at work is thoroughly referential, whether one is talking about a sentence pointing at a region of logical space or whether one is talking about the distinctions between worlds that make that possible. The compositional semantics for the connectives is not really separable from this representational function. If one introduced operators that did something non-referential, Definition 1 would provide no clue what Boolean combinations would mean—that was Frege's (1963) point. This makes it clear how non-trivial it is to add an operator with a non-referential function to this framework. This compounds the challenge for an expressivist account of communication. If distinctions between mental states are tracked compositionally in the fashion of Definition 1, then they must be treated exactly like the distinctions that are being referred to. This makes it very hard to see how one can take a content  $\llbracket \phi \rrbracket$  and *separate* those distinctions that are being referred to and those that are not. It is this tension that I will develop into a more detailed criticism of pragmatic expressivism. Before developing that criticism, I will turn to logical relations.

<sup>19</sup> This allows me to compress the distinction between worlds and models, or between ways the world could be and ways the atomic expressions could refer.



The standard definitions of logical relations also rely essentially on reference, though that is less than clear on a first pass.<sup>20</sup>

**Definition 2 (Consequence).**

$$\phi_1, \dots, \phi_n \models \psi \Leftrightarrow (\llbracket \phi_1 \rrbracket \cap \dots \cap \llbracket \phi_n \rrbracket) \subseteq \llbracket \psi \rrbracket$$

**Definition 3 (Consistency).**

$$\phi_1, \dots, \phi_n \text{ are consistent} \Leftrightarrow (\llbracket \phi_1 \rrbracket \cap \dots \cap \llbracket \phi_n \rrbracket) \neq \emptyset$$

On the referential picture of meaning sentences point to regions of logical space by pointing at some distinction between possible worlds and describing the points in the space as being on a particular side of that distinction. The consequence relation requires  $\phi_1, \dots, \phi_n$  to collectively point to a subregion of what  $\llbracket \psi \rrbracket$  points to. This means that the distinctions made by  $\phi_1, \dots, \phi_n$  guarantee that the distinction made by  $\llbracket \psi \rrbracket$  has already been made. Similarly, consequence requires that  $\phi_1, \dots, \phi_n$  collectively point to a non-empty region. This means that the distinctions they describe can be satisfied by at least one world.

As before, there is a difficulty for an expressivist who would like to assign a meaning  $\llbracket \phi \rrbracket$  which expresses psychological states without referring to them. Because both definitions are entirely moored in the referential framework, it is simply unclear what they would mean when extended to a variety of non-referential meaning. Consider consistency. If one isn't trying to refer to distinctions in worlds is it actually problematic if there is not a world which bears all those distinctions simultaneously? This draws out a feature that will be important below, and is often suppressed in discussions about what an expressivist must do to explain the inconsistency of expressive sentences. Even in entirely standard truth-conditional semantics, there is a teleological dimension to the explanations of logical relations.

One might think that if an expressivist could assign denotations to non-referential expressive sentences, then Definition 2 could be simply applied to them. However, such a minimal answer is arguably unsatisfactory even for non-expressive sentences. When one seeks an *explanation* of  $A$  and  $\neg A$ 's inconsistency, one often wants an explanation of why an agent cannot believe/assert both without changing their mind and an explanation of why two agents which believe/assert these contraries are disagreeing. Simply saying that the denotation of the two sentences has an empty overlap does not fully address this question. What is wrong with simultaneously asserting/believing sentences whose denotations have an empty overlap? Answering this question requires saying something about what denotation, assertion and belief are such that empty overlap results in some kind of rational tension for an agent. For example, one might maintain that the function of belief and assertion is to represent the world (Stalnaker, 1978, 1984), and the denotation of a sentence captures how it represents the world to be. Then, semantically inconsistent sentences represent the world as being some way

<sup>20</sup> Geach (1965) stresses the difficulty of capturing inferences involving expressive sentences.

no world can be, so they are dysfunctional things to simultaneously believe or assert. This explanation combines three resources.

**Explaining Inconsistency** Representational explanations of inconsistency explain the defective nature of simultaneously believing/asserting two sentences by:

1. Construing the denotation relation as the representation relation
2. Construing the function of asserting/believing those sentences as representing the world
3. A definition of semantic consistency as empty overlap, and an explanation of why that tracks dysfunctional assertions/beliefs.

Agreement and disagreement can then be thought of as consistency between the beliefs/assertions of two agents. This broader understanding of what it means to explain inconsistency clarifies both the challenge and prospects for an expressivist. The prospect is clear: they need an alternative conception of the denotation relation, a different construal of the function sentences serve, a formal definition of consistency and an explanation of how that formal definition tracks dysfunctional assertions/beliefs.<sup>21</sup> Now for the challenge. The expressivist can semantically treat the sensitivity to psychological states exactly as one would treat referential discourse, but then they will face two challenges. They must have a different account of why empty overlaps in denotation are bad, since they believe the function of expressive discourse is non-representational. Indeed, it seems that anything in the neighborhood of the explanatory expressivist thesis commits them to appealing here in an essential way to the non-representational attitudes that are supposed to explain the semantic properties of expressive sentences. Additionally, the expressivist must explain how that meaning, which is formally modeled with means suited to representational meaning, is used in a non-representational form of communication. In the next section I will argue that recent expressivist analyses (Silk, 2014; Yalcin, 2011, 2007, 2012) appear to explain inconsistency, but it is not actually clear that they meet the first challenge. Further, I will argue that they do not meet the second.

#### 4 Preference Expressivism and Pragmatic Expressivism

Silk (2014, p. 16) and Yalcin (2012, §10) offer quite similar semantic analyses of deontic modals, both of which resemble the account of expressive content sketched by Gibbard

<sup>21</sup> Charlow (2014) construes explaining consistency more narrowly: deriving a contradiction in the meta-language from the assumption that the sentences are jointly true/satisfied. Charlow (2014, §5.2) argues that a psychologist semantics is not entitled to assume in such a derivation that irrational psychological states do not exist, since they are logically possible. I'm taking a broader view here, which begins with the fact that content is functionally grounded. Setting aside which approach is better, note that an account like the above can happily admit the existence of irrational states of mind and still explain inconsistencies. They are essential to that explanation!

(2003).<sup>22</sup> Both approaches adopt a truth-conditional semantics, but one with a format that is different from the above: it assigns a sentence to a truth-value relative to a world  $w$  and, crucially, a *practical parameter*. For Silk (2014), the practical parameter  $\succsim$  is a weak-preference (pre)ordering (reflexive, transitive) on possible worlds, i.e.  $w_0 \succsim w_1$  means  $w_0$  is preferred at least as much as  $w_1$ , and when  $w_1 \succsim w_0$  also holds  $w_0$  and  $w_1$  are equally preferred. Yalcin (2012, §10) stays closer to Gibbard (2003), and uses ‘hyperplans’.<sup>23</sup> I’ll focus on Silk’s (2014) semantics, since it will be closer to the view I develop later. Both analyses follow Gibbard (2003) in proposing truth-conditions that do not place any constraint on the world parameter, only the practical parameter. For Silk (2014), this means placing a constraint only on the  $\succsim$  parameter: there is no mention of the world of evaluation  $w$  on the right-hand side of the semantic clauses. Deontic modals constrain  $\succsim$  by requiring something of the worlds that are *best* according to  $\succsim$ , i.e. worlds that are at least as preferable as every other world.

**Definition 4 (Preference Expressivism, Silk 2014, 16).**

Where  $Best(\succsim) = \{w \mid \forall w': w \succsim w' \text{ if } w' \succsim w\}$ :

1.  $\llbracket \text{Must}(\phi) \rrbracket_{w, \succsim} = 1 \iff \forall w' \in Best(\succsim): \llbracket \phi \rrbracket_{w', \succsim} = 1$ 
  - $\text{Must}(\phi)$  is true at  $w, \succsim$  iff all the worlds best according to  $\succsim$  are  $\phi$ -worlds.
2.  $\llbracket \text{May}(\phi) \rrbracket_{w, \succsim} = 1 \iff \exists w' \in Best(\succsim): \llbracket \phi \rrbracket_{w', \succsim} = 1$ 
  - $\text{May}(\phi)$  is true at  $w, \succsim$  iff some world best according to  $\succsim$  is a  $\phi$ -world.

This semantics predicts that  $\text{Must}(A)$  and  $\neg \text{Must}(A)$  cannot both be true. Given the standard semantics for negation,<sup>24</sup> the truth of both  $\text{Must}(A)$  and  $\neg \text{Must}(A)$  requires:  $\forall w' \in Best(\succsim): \llbracket A \rrbracket_{w', \succsim} = 1$  and  $\exists w' \in Best(\succsim): \llbracket A \rrbracket_{w', \succsim} \neq 1$ . In words, the truth of both requires all  $\succsim$ -best worlds to be  $A$ -worlds and some  $\succsim$ -best world to be a  $\neg A$ -world. There is no set of worlds, let alone one generated from  $\succsim$ , that satisfies that requirement. This semantics clearly predicts the inconsistency of contrary expressive sentences. But does it predict it in a way that the expressivist can endorse? Silk moves quickly here, contending that it does:

In response to the negation problem, we can say that  $\text{MUST}$  and  $\text{NOT MUST}$  are inconsistent because their conjunction expresses an incoherent set of weak preferences. It is incoherent to have all of one’s most weakly preferred alternatives be ones where  $[A]$  and some of one’s most weakly preferred alternatives be ones where  $[\neg A]$ . There is no (non-empty) preorder that

<sup>22</sup> I am tempted, but hesitant, to categorize the account proposed by Charlow (2013, 2014) along with the pragmatic expressivists. The view developed there has a more nuanced and developed metasemantics that make this categorization unclear. How exactly the points made in this section relate to these voices will have to be left for another day.

<sup>23</sup> Unlike Gibbard (2003), Yalcin (2012, §10) constructs each hyperplan  $h$  out of possibilities:  $h$  is a function that maps a set of worlds  $s$  to one of its subsets, namely those worlds permissible according to the hyperplan given that one is in circumstances  $s$ .

<sup>24</sup>  $\llbracket \neg \phi \rrbracket_{w, \succsim} = 1 \iff \llbracket \phi \rrbracket_{w, \succsim} \neq 1$

represents such a body of preferences. Preferential incoherence appears to be the right kind of inconsistency in attitude to explain the inconsistency between normative sentences ‘ $\phi$ ’ and ‘ $\neg\phi$ ’.  
(Silk, 2014, 16)

The key idea is that the two sentences express preferences that are incoherent. It is crucial for the expressivist that this incoherence is non-representational. This is where I think there is room for skepticism. The kinds of non-representational coherence constraints placed on preferences are like the one discussed in §1: no cyclic orderings. That is because such orderings exist and would fail to motivate choice. There is no parallel coherence constraint on preferences in decision theory that says one cannot have all of one’s most weakly preferred alternatives be ones where **A** and some of one’s most weakly preferred alternatives be ones where  $\neg\mathbf{A}$ . It’s not that there is a preference ordering which does have these two properties but would fail to motivate choice. It’s that no preference ordering meets the two descriptions at all. It is no more a practical constraint on preference orderings than the fact that they cannot both range over ten alternatives and twelve alternatives. The nonexistence of such a preference ordering does not follow from any background theory about how preference motivates choice, but from the fact that they are inconsistent ways of *representing preferences*. More generally, if there is a genuine appeal to preferential incoherence and not representational inconsistency here, then one should be able to construct an ordering which has these two properties and show how it fails to motivate choice. Since the descriptions are logically incompatible, I do not see how this is possible.

Another way of illustrating the problem here is to return to the observations of §1. There I motivated the idea that the natural way to combine preferences is to union them. This means that the motivational consistency of two preference orderings amounts to whether their union can serve its function of motivating choices. Now, consider two orderings  $\succsim_1$  and  $\succsim_2$  over one **A**-world  $w_1$  and one  $\neg\mathbf{A}$ -world  $w_0$ . Suppose then that  $w_1 \succsim_1 w_0$  and  $w_0 \not\succeq_1 w_1$ , while the other ordering is just the opposite:  $w_0 \succsim_2 w_1$  and  $w_1 \not\succeq_2 w_0$ . **Must(A)** then expresses  $\succsim_1$  while  $\neg\mathbf{Must(A)}$  expresses  $\succsim_2$ . If the two orderings are combined into one, call it  $\succsim_3$ , a perfectly rational ordering results where one is indifferent between **A** and  $\neg\mathbf{A}$ :  $w_1 \succsim_3 w_0$  and  $w_0 \succsim_3 w_1$ . Indeed, since  $\succsim_2$  is also expressed by **Must( $\neg\mathbf{A}$ )** the same problem arises for explaining the rational tension produced by committing oneself to both **Must(A)** and **Must( $\neg\mathbf{A}$ )**.<sup>25</sup> Yet,  $\succsim_3$  is a perfectly rational state of preference. There are no background decision-theoretic principles the preference expressivist can appeal to here to render this ordering incoherent. There is, however, an important lesson one can take away from this discussion. Solving the external negation problem by appealing to incoherent preferences may not be possible.

<sup>25</sup> To some degree, this latter point is a bug of Silk’s formalization using weak preferences. Weak preferences are ill-suited to the job here because it is impossible to distinguish between an irrational symmetric strict preference and indifference when you take weak preference as basic and define strict preference as an asymmetric preference. For this reason, I use strict preferences in §6.

Let me now turn to the question of how deontic modals with the semantics in Definition 4 could be used to communicate preferences without referring to them and describing them, i.e. representing them. This is where Yalcin (2012, §10) makes a valuable proposal.<sup>26</sup> Beginning with the dynamics of language use, the view is that the common ground of a conversation consists not just of mutually supposed information  $s = \{w_0, \dots, w_n, \dots\}$  but also a set of preferences  $P = \{\succsim_0, \dots, \succsim_n, \dots\}$  (Yalcin, 2012, p. 150). While descriptive discourse constrains the world and eliminates worlds from  $s$ , deontic language constrains preferences and eliminates preference orderings from  $P$ . In brief: each type of sentence affects the common ground in a different way.<sup>27</sup> Since  $s$  is a representational state and  $P$  is a motivational state this distinction appears apt for distinguishing representational from motivational discourse. Yalcin’s (2012, p. 150) sketch suggests a formalization in the following territory:<sup>28</sup>

#### Pragmatic Expressivist Discourse Model

1. A descriptive sentence  $\phi$  expresses a constraint on how the world can be, i.e. the set of left-indices that make it true:  $\langle\langle\phi\rangle\rangle = \{w \mid \exists \succsim: \llbracket\phi\rrbracket_{w, \succsim} = 1\}$
2. An expressive deontic sentence  $\phi$  expresses a constraint on preferences, i.e. the set of right-indices that make it true:  $\langle\rangle\phi\rangle\rangle = \{\succsim \mid \exists w: \llbracket\phi\rrbracket_{w, \succsim} = 1\}$
3. Common ground:  $C = \langle s_C, P_C \rangle$ , where  $s_C = \{w_0, \dots, w_n\}$ ,  $P_C = \{\succsim_0, \dots, \succsim_m\}$ 
  - Background assumption:  $s_C$  is the ways the world could be given what the agents are taking for granted
  - Background assumption:  $P_C$  is the preferences compatible with the situations the agents’ are in favor of
4. Representational communication:  $C \upharpoonright \phi = \langle s_C \cap \langle\langle\phi\rangle\rangle, P_C \rangle$
5. Expressive communication:  $C \upharpoonright \phi = \langle s_C, P_C \cap \langle\rangle\phi\rangle\rangle \rangle\rangle$

How does this allow one to distinguish referring to and describing preferences from expressing preferences? The idea seems to be that referring to and describing preferences eliminate worlds while expressing preferences eliminates orderings. The problem is that this interpretation of the model is entirely optional. Since the dynamics of expressive and representational content are identical, the system could be reformulated in terms of updating a set of world-ordering pairs, or centered-worlds, which distinguish ways the world could be from the perspective of a preference ordering. Then one can easily describe the system as behaving as referring to and describing preferences, just as worlds centering on individuals are said to refer to and describe an individual (Egan, 2007). The only way to clearly distinguish expressive

<sup>26</sup> See Rothschild (2012); Swanson (forthcoming); and Moss (2015) for related work. Ninan (2005) presents a view that is ambivalent between semantic and pragmatic expressivism. The examination of pragmatic expressivism through examining embedded *must* claims is a strategy Ninan (2005, §6.1) already highlights as an important way to probe these views. See note 53 in §7 for more discussion of Ninan (2005).

<sup>27</sup> The terminology of expressing *constraints* comes from Swanson (2006; forthcoming).

<sup>28</sup> The notation of  $\langle\langle\phi\rangle\rangle$  and  $\langle\rangle\phi\rangle\rangle$  is novel here, but iconic:  $\llbracket\phi\rrbracket$  is a set of pairs  $\langle w, \succsim \rangle$  while  $\langle\langle\phi\rangle\rangle$  is the set of left members (worlds) and  $\langle\rangle\phi\rangle\rangle$  is the set of right members (orderings).

from representational communication is to motivate and adopt a model where the two kinds of discourse have clearly different dynamics, say one where expressive sentences do not refer to orderings and place constraints on them. It is hard to see how one can do this without departing from truth-conditional semantics. As discussed above, this difference also seems necessary for a solution to the negation problem. It turns out there is a rather similar issue with disjunction.

Mixed disjunctions provide a hard challenge for the expressivist approach above. It is hard to see how they can be explained while maintaining that expressive and representational communication have distinct communicative effects. The challenge is related to Schroeder’s (2015) concerns about mixed disjunctions in belief ascriptions. But my focus will be on what mixed disjunctions communicate.<sup>29</sup> Suppose you know that exactly two things make Alex happy, namely Pat being in town or Alex receiving permission to skateboard, and I ask *Why is Alex happy?* you might answer:

- (5) Either Pat is in town or Alex may skateboard.

Consider how this sentence would update the common ground on the pragmatic expressivist model above. The standard semantics for disjunction gives us the following truth-conditions and corresponding constraints:

- (6)  $\llbracket P \rrbracket_{w, \succ} = 1$  or  $\llbracket \text{May}(S) \rrbracket_{w, \succ} = 1$   
 (7)  $\langle\langle P \vee \text{May}(S) \rangle\rangle = \{w \mid \exists \succ: \llbracket P \rrbracket_{w, \succ} = 1 \text{ or } \llbracket \text{May}(S) \rrbracket_{w, \succ} = 1\}$   
 (8)  $\langle\langle P \vee \text{May}(S) \rangle\rangle = \{\succ \mid \exists w: \llbracket P \rrbracket_{w, \succ} = 1 \text{ or } \llbracket \text{May}(S) \rrbracket_{w, \succ} = 1\}$

The problem for the world constraint is that the deontic disjunct will allow every world to meet this condition. There is a preference ordering that makes  $\text{May}(S)$  true and if one world-ordering pair featuring that ordering makes it true, then any other pair with that ordering makes it true: deontics do not constrain the world parameter. The same problem arises for the preference constraint, since there is a world that makes  $P$  true and  $P$  does not constrain the ordering parameter. This predicts that a mixed disjunction cannot communicate anything, contrary to our intuitions about (5).<sup>30</sup>

There is an alternative analysis of disjunction that pragmatic expressivists could adopt to capture this data. Suppose we took common grounds to be sets of the old common grounds:  $\mathcal{C} = \{\langle s_0, P_0 \rangle, \dots\}$ . For any non-disjunctive sentence, proceed as before, only applying the update to each information state or preference set. But for disjunctions, don’t use truth conditions, but rather perform a special update process. Take the result of updating  $\mathcal{C}$  with the left disjunct and union it with the result of

<sup>29</sup> Mixed disjunctions are a focal point in the arguments I’ve offered for dynamic analyses of imperatives (Starr, 2010, 2013).

<sup>30</sup> Silk (2014, §5) offers a modified semantics designed to solve Schroeder’s (2015) disjunction problem. But this results in practical sentences that constrain the world parameter and hence update  $S$ , and representational ones that constrain the preference parameter and hence update the preference set. This solves the problem at the expense of having a way of neatly distinguishing expressive from representational communication.

updating with the right disjunct. To illustrate how that would work, consider a simple common ground where there is one pair:  $\mathcal{C} = \{\{\{w_1, w_0\}, \{\lesssim_1, \lesssim_0\}\}\}$ . Suppose  $w_1$  is a  $\mathbf{P}$ -world,  $w_0$  is a  $\neg\mathbf{P}$ -world,  $\lesssim_1$  makes  $\mathbf{May}(\mathbf{S})$  true and  $\lesssim_0$  makes  $\mathbf{May}(\mathbf{S})$  false. Then updating with (5) will produce  $\mathcal{C}' = \{\{\{w_1\}, \{\lesssim_1, \lesssim_0\}\}, \{\{w_1, w_0\}, \{\lesssim_1\}\}\}$ . We may interpret these new common grounds as ‘imprecise common grounds’, on analogy with imprecise credences. The agents’ states of mind are not precise enough to distinguish which state they are in.<sup>31</sup> This revised analysis can be interpreted in two ways. Perhaps disjunction does not have a truth-conditional semantics, but only a distinctive pragmatic update rule that is different from that which applies to every other kind of sentence. This assumption conflicts with the received methodology of semantics and pragmatics. No connective has *only* a pragmatics. It does, however, suggest a more radical approach: what if all sentence forms were given update rules as their semantics? This more radical approach is simply *dynamic semantics*, and will be pursued in the next section.

## 5 Logical Semantics, Dynamically

On a dynamic semantics, the meaning of a sentence can be modeled as a function  $[\phi]$  which specifies how accepting it transforms an agent’s state of mind  $s$  into another  $s'$ . One specifies the semantics with equations such as  $s[\phi] = s'$ , which say how applying  $[\phi]$  to  $s$  results in some other state  $s'$ . Since meaning resides in how sentences change states, it is said to be dynamic. By contrast, truth-conditional semantics is static: it specifies a relation between symbols and the world rather than specifying how those symbols are used. To illustrate this new approach in familiar territory, I will provide a semantics for propositional logic consisting of equations with the form  $s[\phi] = s'$ . While many have proposed systems in the vicinity of this, including Kamp (1981), Heim (1982), and Groenendijk and Stokhof (1991), what will follow is basically the system presented in Veltman (1996).<sup>32</sup> The intuitive idea is that using  $\phi$  consists in moving from state  $s$  to state  $s'$ . With programs, the states were machine states, but what are states now? The goal here is to use the formal language of propositional logic to model the use of natural language. I will take  $s$  to be the *state of information* in the conversation or inquiry:  $s$  is the set of worlds compatible with what the agents’ are assuming for the purpose of the exchange. Formally, filling out an equation of the form  $s[\phi] = s'$  consists in saying for any given set of worlds  $s$ , the set of worlds  $s'$  that

<sup>31</sup> This interpretation and solution are a wink to Rothschild (2012) who, essentially, uses lifted states like this to address a related problem for pure expressive disjunctions. Mixed disjunctions, however, are not modeled there. This is the only reason Rothschild (2012) can do without a special disjunction pragmatic update rule. For more on the motivations for imprecise credences and pointers to that literature see Halpern (2003, §2.3).

<sup>32</sup> Unlike Veltman (1996, p. 228) but like Groenendijk et al. (1996, §3) I will treat conjunction as sequential update. Unlike Groenendijk et al. (1996, §3) but like Veltman (1996, p. 228) I will treat disjunction as forming the union of parallel updates to the initial state.

result from interpreting  $\phi$ . Intuitively, this means saying how processing  $\phi$  requires the agents’ to adjust the information they are assuming in their exchange. To save a few words I will say that  $s'$  is the result of *updating*  $s$  with  $\phi$ .

**Definition 5 (Compositional Semantics).**

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

Atomic sentences eliminate possibilities incompatible with their truth. For complex sentences, the equations specify how the interpretation of  $\phi$  depends on interpreting its constituents. Negation eliminates the possibilities compatible with its scope. Conjunctions update with each of their conjuncts in sequence. Disjunctions update with each of their disjuncts in parallel and form the union of each result. The clause for atomic sentences directly appeals to truth/reference in specifying the interpretation procedure for an atomic sentence. But note that unlike in the classical semantics the connectives here are not given a representational interpretation. Intuitively, they do not work by increasing the precision with which one can refer to distinctions between possible worlds. They increase the space of transitions between states the language can traverse. This different understanding of expressive power is the crucial difference between dynamic and classical semantics. It allows one to conceive of two different ways to make the language more expressive: add operators which discern finer distinctions in the space of possible worlds, and add operators which trace more transitions through the space of states. The key in understanding the relationships between the two forms of semantics comes to this: what kinds of transitions amount to drawing distinctions between possible worlds and what kinds of transitions amount to something else? One way of answering this in the present context is considering which operators exploit operating on a whole set of worlds rather than considering each world in isolation. In systems such as that above where  $s[\phi] \subseteq s$ , this formal difference amounts to whether or not the update is distributive.<sup>33</sup>

**Definition 6 (Distributive).**  $s[\phi]$  is distributive just in case  $s[\phi] = \bigcup\{w[\phi] \mid w \in s\}$ .

The semantics above is distributive, and so the underlying language does not ‘do anything’ which does not amount to making distinctions between possible worlds. This is what I will eventually give up. The possibility of doing so hangs crucially on the fact that dynamic semantics furnishes logical concepts more general than those employed in truth-conditional semantics.

A more general conception of meaning is molded by a more general conception of logic. Instead of truth, dynamic semantics builds its logic on **support**.

**Definition 7 (Support, Truth in  $w$ ).**

- (1) **Support**  $s \models \phi \Leftrightarrow s[\phi] = s$       (2) **Truth in  $w$**   $w \models \phi \Leftrightarrow w[\phi] = \{w\}$

<sup>33</sup> This property is sometimes also called *continuity*, and has been isolated as a key feature by van Eijck and Visser (2012, §6.2), Muskens et al. (1997, p. 595) and van Benthem (1989, p. 364).



Generally, a state  $s$  supports  $\phi$  just in case interpreting  $\phi$  requires no change in  $s$ . Where  $s$  is an information state this means that interpreting  $\phi$  requires no change in that information. (Equivalently:  $[\phi]$  is a fixed point with respect to  $s$ .) Support is more general than truth: truth in a world is a special case of support.  $\phi$  is true in  $w$  just in case accepting  $\phi$  requires no change to perfect information about  $w$ :  $\{w\}$ .<sup>34</sup> This definition of truth makes plain that non-distributive dynamic meanings are the non-classical ones: they are the ones where support, rather than truth, is necessary.

Support is the basic logical concept in the sense that consequence and consistency are defined in terms of it.  $\psi$  a consequence of  $\phi$  just in case updating any state with  $\phi$  results in a state that supports  $\psi$ .

**Definition 8 (Dynamic Consequence).**

$$\phi_1, \dots, \phi_n \models \psi \Leftrightarrow \forall s: s[\phi_1] \cdots [\phi_n] \models \psi$$

The basic idea is that any agent who accepts the premises is committed to accepting the conclusion (or rejecting those premises) because accepting the conclusion in the resulting state requires no more information than is already possessed; they are making no additional commitments about what the world is like. Consistency is also defined in non-propositional terms: there is a non-absurd state of information which supports each of  $\phi_1, \dots, \phi_n$ . That is, there is a single informational perspective which supports all of the sentences.

**Definition 9 (Dynamic Consistency).**

$$\phi_1, \dots, \phi_n \text{ are consistent} \Leftrightarrow \exists s \neq \emptyset: s \models \phi_1, \dots, s \models \phi_n$$

Throughout, I’ve emphasized that the dynamic system presented is a generalization of the classical one. It is worth making this more explicit. On the dynamic semantics, one can still define propositional content.

**Definition 10 (Propositional Content).**  $\llbracket \phi \rrbracket = \{w \mid w \models \phi\}$

Taken with the dynamic compositional semantics from Definition 5, this implies the classical compositional semantics (Veltman, 1996).

**Corollary 1 (Possible Worlds Semantics).**

- |   |   |
|---|---|
| (1) $\llbracket p \rrbracket = \{w \in W \mid w(p) = 1\}$   | (2) $\llbracket \neg \phi \rrbracket = W - \llbracket \phi \rrbracket$                                  |
| (3) $\llbracket \phi \wedge \psi \rrbracket = \llbracket \phi \rrbracket \cap \llbracket \psi \rrbracket$ | (4) $\llbracket \phi \cup \psi \rrbracket = \llbracket \phi \rrbracket \cup \llbracket \psi \rrbracket$ |

Thus, even though propositions do not directly figure in the compositional clauses and connectives are not propositional operators (they actually combine update procedures), the dynamic semantics behaves exactly *as if* these classical ideas were operative. Propositions do not directly figure in the definitions of consequence or consistency either. Just as truth is a special case of support, classical consequence and consistency are special cases of the dynamic notions.

<sup>34</sup> This definition is mentioned, but not adopted, by Muskens et al. (1997, p. 594).

**Definition 11 (Classical Consequence).**

$$(\llbracket \phi_1 \rrbracket \cap \dots \cap \llbracket \phi_n \rrbracket) \subseteq \llbracket \psi \rrbracket \Leftrightarrow \forall w: \{w\}[\phi_1] \dots [\phi_n] \models \psi$$

**Definition 12 (Classical Consistency).**

$$(\llbracket \phi_1 \rrbracket \cap \dots \cap \llbracket \phi_n \rrbracket) \neq \emptyset \Leftrightarrow \exists w: \{w\} \models \phi_1, \dots, \{w\} \models \phi_n$$

In full generality, dynamic consequence and consistency are not equivalent to the classical notions. But restricted to the language of propositional logic with the dynamic semantics of Definition 5, they are. This follows rather directly from the fact that all updates in the present semantics amount to information gain: eliminating possibilities.

**Fact 1 (Update is Propositional Communication).**

Given the semantics of Definition 5:  $\forall \phi: s[\phi] = s \cap \llbracket \phi \rrbracket$

This fact not only means that the notions of consequence and consistency are equivalent to their classical counterparts. It also means that the account of communication is equivalent: communication works by providing information, i.e. referring to some aspect of the world and stating a fact about it. As long as this equivalence holds, the difference between dynamic and classical semantics is purely conceptual. This conceptual difference resides in the centrality of support, rather than truth, to dynamic semantics. This conceptual difference makes room for a non-equivalent account which preserves the virtues of the classical account, albeit with non-classical definitions of compositional meaning, consequence, consistency and communication. But the classical logician cannot charge the dynamic logician with a change of topic. The classical definitions are a special case of the non-classical ones. The disagreement is thus not one of subject matter, but of breadth. It is easy to see how this disagreement is relevant to the debate about expressivism. Frege and Geach took the narrow classical concepts as given. But what happens when the more general concepts are embraced?

## 6 Dynamic Expressivism about Deontic Modality

The equivalence between the dynamic and classical definitions in §5 can be broken. In this section I will provide a dynamic semantics where the updates assigned to *must* claims have the distinctive expressivist property: they do not communicate by providing information, i.e. they do not refer to some part of the world and provide information about it. Instead, they promote certain preferences and then test that those preferences make the scope of the modal a practical necessity/possibility. I will first present the semantics and then turn back to the key issues from the discussion of expressivism above. The first step will be to model states that reflect not just the agents' information, but their motivations.

### 6.1 What are States?

Here is my answer to the question:

**Definition 13 (States).** A state  $S = \{s_1^{\succsim_1}, \dots, s_n^{\succsim_n}\}$  is a set of *substates*  $\langle s_i, \succsim_j \rangle$ .

- (1) Each  $s_i$  is a set of worlds compatible with what the agents are mutually supposing, and is competing to be *the* information they are mutually supposing.
- (2) Each **preference frame**  $\succsim_j = \langle \succ_j, \sim_j \rangle$  consists of two relations,  $\succ_j$  and  $\sim_j$ , on  $W$ .
  - Where  $w \succ_j w'$  just in case  $w$  is strictly more preferred than  $w'$  and  $\succ_j$  is competing to guide the agents
  - Where  $w \sim_j w'$  just in case  $w$  is equally preferred to  $w'$  and  $\sim_j$  is competing to guide the agents
- (3) Each **substate**  $\langle s_i, \succsim_j \rangle$  consists of information  $s_i$  and a preference frame  $\succsim_j$ .
- (4) *Notation:*  $s^{\succsim} := \langle s, \succsim \rangle$  and any set theoretic notation applied to ‘ $s^{\succsim}$ ’ is understood as being applied to  $s$ , e.g.  $s^{\succsim} \subseteq W$  means that  $s \subseteq W$ .

Of course, it may be helpful to hear why each of these components is here and how states are being interpreted. Those who wish to skip those niceties, proceed to §6.2, but have a peek at Definition 14 along the way.

I will follow a developed tradition in decision theory and model an agent’s motivational state as an ordering of worlds, where  $w \succ w'$  means that  $w$  is strictly preferred to  $w'$ , and  $w \sim w'$  to mean that  $w$  and  $w'$  are equally preferred.<sup>35</sup> While it is common to just use a weak preference ordering  $\succsim$  and define equal and strict preference in terms of it, this will not work for my purposes. My analysis of inconsistency of  $\text{Must}(\phi)$  and  $\text{Must}(\neg\phi)$  will require modeling the irrational state of strictly preferring  $\phi$ -worlds to  $\neg\phi$ -worlds and vice versa. On a weak preference analysis, this cannot be distinguished from the perfectly rational state of  $\phi$  and  $\neg\phi$ -worlds being equally preferred. At the same time, I cannot do with just strict preference. I will need equal preference to model a state where  $\neg\text{Must}(\phi)$  has been accepted, but  $\text{Must}(\neg\phi)$  has not: since the agents find  $\phi$  and  $\neg\phi$  equally preferable they reject that they must  $\phi$ , but also reject that they must  $\neg\phi$ .<sup>36</sup> So I propose to capture the motivational component of states using a *preference frame*, which is a pair consisting of these two relations on worlds:  $\langle \succ, \sim \rangle$ . To simplify notation, I’ll refer to a frame as a whole with ‘ $\succsim$ ’ but reserve ‘ $\succ$ ’ or ‘ $\sim$ ’ for use between worlds.

As has already surfaced, rational preferences have certain properties, e.g. strict ones are acyclic, equal ones are reflexive. But I’ve already said that I am not limiting this model to rational states. Some take these properties to be definitional of preferences so it is important for me to articulate a different view. I do not take preference to

<sup>35</sup> Some gloss  $w \sim w'$  as *indifference* but this gives the misleading impression that the agent doesn’t care about  $w$  and  $w'$  at all.

<sup>36</sup> Appropriately capturing the distinction between  $\neg\text{Must}(\phi)$  and  $\text{Must}(\neg\phi)$  is a well-documented difficulty for expressivism (Unwin, 2001; Schroeder, 2008c; Dreier, 2009), and so one this analysis must solve.

be reducible to choice *behavior*, but I will assume that preferences have a function: to motivate choice. This functional account assumes that rational preferences have the properties they do because those are the properties they must have to achieve their function of motivating choice. This idea about the function of preferences will be important when I discuss consistency, and it should recall the discussion of consistency at the end of §3.

Even though I will not require preferences to be rational, their rationality will nonetheless be relevant. It is important to be clear about how. Their first relevance will be in the pragmatics: I will assume that a speech act which brings about irrational preferences is unstable: the discourse participants will not be satisfied with it. Their second relevance is in the semantics for deontic modals. For example,  $\text{Must}(\phi)$  will test whether the alternatives which may be rationally chosen are the  $\phi$ -alternatives. If the underlying preferences are not rational, this test is sure to fail since no alternatives may be rationally chosen from irrational preferences. This means that I must say more about how alternatives are rationally chosen, and which properties this requires preferences to have. This is obviously an enormous and delicate issue, but all I will need below is two plausible claims about rational choice: (i) rational agents can choose any alternative which is not dominated by—strictly less preferred than—an alternative; (ii) if  $w$  can be rationally chosen and  $w \sim w'$ , then  $w'$  can be chosen.<sup>37</sup> The set of rational alternatives to choose from a space of worlds  $s$  can then be defined as follows.

**Definition 14 (Choice Possibilities).**

$\text{Choice}(s, \succ) := \{w \in s \mid w \in Nd \text{ or } \exists w' \in Nd: w \sim w'\}$

- $Nd(s, \succ) := \{w \in s \mid \nexists w' \in s: w' \succ w\}$
- Choice possibilities in  $s$  are either non-dominated, or equally preferable to some  $w \in s$  that is.

On this conception of choice, a cyclic strict preference ordering such as  $w \succ w', w' \succ w$  has the function of promoting the choice of both  $w$  and  $w'$ , but discounts both. It is thereby irrational. There are many other properties rational preference must have on this definition, but asymmetry is the only one that will be crucial below.

I seem to have arrived at a good model of states: a space of worlds and a preference frame  $\langle s, \succ \rangle$ . But, for the reasons discussed at the end of §4, disjunctions featuring deontic modals challenge such a model. At least in the static setting this challenge motivated a higher type for states, namely a *set* of these pairs. I will also adopt this higher type. Here's why. Suppose  $\langle s, \succ \rangle$  was the form states take. A standard idea is that  $\text{Must}(\phi)$  tests that the 'choice' worlds in  $\langle s, \succ \rangle$  are the  $\phi$ -worlds.<sup>38</sup>

<sup>37</sup> See Hansson and Grüne-Yanoff (2011, §3.2).

<sup>38</sup> This is a deontic adaptation of Veltman's (1996, p. 235) *presumably*; see also Portner (2009, p. 104), definition 146.2.

**Definition 15 (Simple Test Semantics).**

$$\langle s, \succ \rangle[\text{Must}(\phi)] = \begin{cases} \langle s, \succ \rangle & \text{if } \langle s, \succ \rangle[\phi] = \langle \text{Choice}(s, \succ), \succ \rangle \\ \langle \emptyset, \succ \rangle & \text{otherwise} \end{cases}$$

Consider a scenario where Pat is coming to visit Alex, who lives in a neighborhood where parking is allowed only on one side of the street, depending on the day. Alex knows that today, parking on the left is required, so he might tell Pat *You must park on the left*. The idea, as Veltman (1996, pp. 235–6) puts it, is for Pat to test whether the information and preferences in play have the property in question: are the choice worlds the live worlds where Pat parks on the left? If so, then proceed as before with  $\langle s, \succ \rangle$ , otherwise  $\langle \emptyset, \succ \rangle$  spurs Pat to adjust  $\succ$  so that the test is passed. Since the preferences in this case are new to Pat, he must adjust the preferences he takes to be at play. Importantly, this adjustment is a pragmatic process on this analysis.

Focusing now on disjunction, one must first generalize the semantics from Definition 5.4 to the new form of states. The simplest idea is to union the information provided by each disjunct and keep the input ordering:

$$(9) \quad \langle s, \succ \rangle[\phi \vee \psi] = \langle s_\phi \cup s_\psi, \succ \rangle, \text{ where } \langle s, \succ \rangle[\phi] = \langle s_\phi, \succ \rangle \text{ and } \langle s, \succ \rangle[\psi] = \langle s_\psi, \succ \rangle.$$

Consider a similar scenario in which Pat is coming to visit Alex, but Alex cannot remember which day goes with which side. Alex might then tell Pat:

- (10) You must park on the left side of the street or you must park on the right side of the street (I can't remember which).

As before, the preferences are new to Pat, so the disjunction in (10) leads to the state  $\langle \emptyset, \succ \rangle$ . Pat must revise the preferences, but how should we represent the state he enters when he does so? Switching to a strict ordering  $\succ_L$  which ranks parking-on-the-left worlds over parking-on-the-right-worlds, will lead to the state  $\langle s, \succ_L \rangle$ . But this means that Pat would know what to do: park on the left. The same goes for an ordering that prefers parking-on-the-right worlds. There is no clear candidate to represent the uncertainty Alex expresses. What is needed is for the pragmatic process of adjusting  $\succ$  to happen semantically, when each disjunct is interpreted, and then to somehow combine the results into one state. The first task is tractable, and will be carried out in §6.2, Definition 17.4.<sup>39</sup> The hard question is how to combine those two augmented strict orderings  $\succ_L$  and  $\succ_R$  into one that captures an undecided state. Forming the union of the two orderings produces an irrational set of symmetric strict preferences where parking on the left is strictly preferred to parking on the right and vice versa. Moving to a frame where parking on the left and on the right are equally preferred is also problematic. Alex has pretty explicitly expressed that the two options aren't

<sup>39</sup> Veltman's (1996, p. 235) *normally* works this way, but I will pursue yet another way.

equally good, just that she can't say which is better. Whatever she is conveying, it seems to leave both orderings open. A simple move can capture this idea: states are sets of  $\langle s_1, \succsim \rangle$ .<sup>40</sup> So interpreting Alex's disjunction results in the state  $\{\langle s_1, \succsim_L \rangle, \langle s_1, \succsim_R \rangle\}$ . The idea is not that agents possess a representation of their preferences which fails to distinguish these two options. Instead, both preferences are competing for control of their choices. Given her epistemic state, all Alex can do is induce this competition among Pat's preferences.<sup>41</sup>

## 6.2 The Semantics

Recall the guiding slogan for the expressivist:  $\text{Must}(\phi)$  promotes certain motivational attitudes towards  $\phi$ .<sup>42</sup> The difficult part for semantically implementing this slogan has been 'promotion'. The essential idea of the semantics I will offer is that this promotion consists in adding a strict preference for  $\phi$ -worlds. More precisely,  $\text{Must}(\phi)$  extends each existing strict preference ordering by ranking each live  $\phi$ -world over each live  $\neg\phi$ -world. This preference update is defined below as  $\phi^+(\succsim)$ . It simply extends  $\succ$  with a strict preference for each  $\phi$ -world over each  $\neg\phi$ -world and leaves  $\sim$  as it was. Crucially, the semantics also tests, for every substate, that this augmented ordering makes  $\phi$  a practical necessity: the choice worlds are the  $\phi$ -worlds.<sup>43</sup> This test relates the preference update to the information supposed in the interaction.<sup>44</sup>

**Definition 16 (Dynamic Expressivist Semantics).**

$$S[\text{Must}(\phi)] = \begin{cases} \{s^{\phi^+(\succsim)} \mid s^{\succsim} \in S\} & \text{if } \forall s^{\succsim} \in S: \text{Choice}(s^{\phi^+(\succsim)}) = s^{\succsim}_{\phi} \\ \{\emptyset^{\phi^+(\succsim)} \mid s^{\succsim} \in S\} & \text{otherwise} \end{cases}$$

- $s^{\succsim}_{\phi} := \bigcup(\{s^{\succsim}\}[\phi])$
- $s^{\succsim}_{\phi}$  is the set of  $\phi$ -worlds in  $s$ .

<sup>40</sup> This allows variation along both axes:  $\{\langle s_1, \succsim_1 \rangle, \langle s_1, \succsim_2 \rangle\}$ ,  $\{\langle s_1, \succsim_1 \rangle, \langle s_2, \succsim_1 \rangle\}$ , and  $\{\langle s_1, \succsim_1 \rangle, \langle s_2, \succsim_2 \rangle\}$  are all states. Variation along the informational axis is needed for mixed disjunctions such as  $\mathbf{A} \vee \text{Must}(\mathbf{B})$ .

<sup>41</sup> This idea of various states competing for control of an agent's actions is widespread in artificial intelligence, in particular in the use of evolutionary algorithms to model adaptive problem solving. See Franklin (1995, Chapter 9) for a helpful overview. This seems in the spirit of Blackburn's (1988) being 'tied to a tree.'

<sup>42</sup> Ninan (2005) is a clear predecessor of this view, but is ambivalent about whether the ordering-effect is semantic. Making this effect semantic is essential for expressivist purposes: it makes  $\text{Must}(\mathbf{A})$  and  $\neg\text{Must}(\mathbf{A})$  inconsistent in a non-representational way.

<sup>43</sup> I here define  $\text{May}(\phi)$  as the dual of  $\text{Must}(\phi)$ , but a less conservative option is worth noting:  $\text{May}(\phi)$  creates a new substate in which  $\phi$ -worlds are preferred to  $\neg\phi$ -worlds, and tests that  $\phi$  is consistent with the choice worlds in each substate. Given the logical definitions to follow this predicts free-choice permission patterns:  $\text{May}(\mathbf{A}) \vee \text{May}(\mathbf{B}) \models \text{May}(\mathbf{A})$  and  $\text{May}(\mathbf{A} \vee \mathbf{B}) \models \text{May}(\mathbf{A})$  while  $\neg\text{May}(\mathbf{A} \vee \mathbf{B}) \models \neg\text{May}(\mathbf{A}) \wedge \neg\text{May}(\mathbf{B})$ .

<sup>44</sup> Note that this account captures both performative and descriptive uses of deontic modals within the same semantics: descriptive ones are those issued in a state that already contains the preference it promotes and performative ones are those issued in a state that previously lacked that preference. Of course, pragmatic social conditions concerning authority govern when performative uses are felicitous.

- $\phi^+(\succ) := \langle \{ \langle w, w' \rangle \in s \times s \mid w \succ w' \text{ or, } w \in s_{\phi}^{\succ} \ \& \ w' \in s_{\neg\phi}^{\succ} \}, \sim \rangle$ 
  - $\phi^+(\succ)$  adds to  $\succ$  a preference for each  $w \in s_{\phi}^{\succ}$  over each  $w' \in s_{\neg\phi}^{\succ}$ .
- $\text{May}(\phi) := \neg\text{Must}(\neg\phi)$

After all, this is a model of declarative modal sentences and interacting with information is the function of declaratives. The fact that they have some relation to information makes them suitable to negate and embed under *believe*. But they hardly work by providing information, and they certainly do not eliminate worlds by referring to a distinction in each world. Their key contribution is the promotion of preferences and this seems on track for expressivist aims. It is rather crucial for these aims that the test component of the meaning is not redundant. After all, that would mean that an update with any deontic modal claim will succeed. To see why the preference update does not guarantee that the test will be passed, consider the fact that the incoming preferences may be inhospitable. The test asks whether the alternatives a rational agent could choose according to the augmented preferences are  $\phi$ -alternatives. But if the augmentation leads to an irrational body of preferences, then the choice set will be empty. (That’s why those preferences are irrational.) To illustrate, consider a simple scenario where the agents have narrowed in on two worlds, an **A**-world  $w_1$  and a  $\neg$ **A**-world  $w_0$ . Further suppose the agents’ only strict preference is for  $w_0$  over  $w_1$  and the agents only equally prefer each world to itself:  $\succ_0 = \{ \langle w_0, w_1 \rangle \}$  and  $\sim_0 = \{ \langle w_0, w_0 \rangle, \langle w_1, w_1 \rangle \}$ . The agents’ state is then modeled as  $\{ \langle w_1, w_0 \rangle \}^{\sim_0}$ . The first thing that an update with  $\text{Must}(\mathbf{A})$  will do is create a new preference ordering where  $w_1$  is also strictly preferred to  $w_0$ :  $\mathbf{A}^+(\succ_0) = \langle \{ \langle w_0, w_1 \rangle, \langle w_1, w_0 \rangle \}, \sim_0 \rangle$ . But such a preference ordering is irrational. Its symmetric preference guarantees that there are no non-dominated worlds, thus:  $\text{Choice}(\{w_1, w_0\}^{\mathbf{A}^+(\succ_0)}) = \emptyset$ . That means  $\text{Must}(\mathbf{A})$  is testing that there are no **A**-worlds. Since  $w_1$  is an **A**-world, this test is failed, resulting in the state  $\emptyset^{\mathbf{A}^+(\succ_0)}$ . The information  $\emptyset$  in this result state reflects that it is not stable and the preferences explain why:  $\mathbf{A}^+(\succ_0)$  is irrational. The agents then face a choice of which preference to jettison. I do not have space to elaborate how I see that pragmatic process unfolding, but suffice it to say that it may involve debate about the particular preferences, or less conscious deference based on social influences.

As it turns out,  $\succ_0$  is the ordering that would have resulted from updating a state with an empty strict ordering and trivial equality ordering  $\{ \langle w_1, w_0 \rangle \}^{(\emptyset, \sim_0)}$  with  $\text{Must}(\neg\mathbf{A})$ . This allows me to sketch how I propose to explain the inconsistency of  $\text{Must}(\neg\mathbf{A})$  and  $\text{Must}(\mathbf{A})$ : there is no rational preference ordering which contains both the preferences they promote. If the state is going to contain the preferences  $\text{Must}(\mathbf{A})$  promotes, it is going to strictly prefer **A**-worlds to  $\neg$ **A**-worlds. If the state is going to contain the preferences  $\text{Must}(\neg\mathbf{A})$  promotes, it is going to strictly prefer  $\neg$ **A**-worlds to **A**-worlds. But any state containing both will be irrational in the sense that it contains dysfunctional preferences. To fill out this explanation, and treat the harder case of  $\text{Must}(\mathbf{A})$  and  $\neg\text{Must}(\mathbf{A})$ , a semantics for negation (and atomics) is needed.

Apart from negation, the semantics for the rest of the language parallels Definition 5. Atomics eliminate worlds where they are false from each substate. Conjunction is still sequential update and disjunction is again the union of parallel updates to the initial state. This semantics for disjunction does generate a new behavior, but I will return to that. Focus now on negation:  $\neg\phi$  not only removes the information  $\phi$  would provide (as before), it removes the preferences that  $\phi$  would promote.<sup>45</sup> This is done by subtracting any preferences that  $\phi$  would add to an empty ordering.

**Definition 17 (Dynamic Expressivist Boolean Semantics).**

1.  $S[p] = \{\{w \in s \mid w(p) = 1\} \succsim \mid s \succsim \in S\}$
2.  $S[\neg\phi] = \{s^{\phi^-}(\succsim) - \bigcup(\{s \succsim\}[\phi]) \mid s \succsim \in S\}$ 
  - $\phi^-(\succsim) := \langle \succ - \{\langle w, w' \rangle \in \succ_i \mid \{W^{(\emptyset,=)}\}[\phi] = \{s_0^{\succsim_1}, \dots, s_n^{\succsim_m}\} \& 1 \leq i \leq m\}, \sim \rangle$
  - $\phi^-(\succ)$  removes from  $\succ$  any pairs that updating with  $\phi$  would add to an empty ordering. For non-expressive discourse this will idle. If  $\phi = \text{Must}(\psi)$  this will extract preferences for  $\psi$ -worlds over  $\neg\psi$ -worlds.
3.  $S[\phi \wedge \psi] = S[\phi][\psi]$
4.  $S[\phi \vee \psi] = S[\phi] \cup S[\psi]$

Since descriptive discourse never adds preferences, negations of descriptive sentences will behave exactly as before. However, when an expressive sentence is negated *both* the preferences it would promote and the information it would provide are removed. Consider again the example from above where testing the state  $\{\{w_1, w_0\} \succsim^0\}$  with  $\text{Must}(A)$  failed, returning  $\emptyset^{A^+(\succsim^0)}$ . What if one updates that same state with  $\neg\text{Must}(A)$  instead? The semantics predicts that this returns the initial state  $\{\{w_1, w_0\} \succsim^0\}$  via the following process. It would first arrive at the state  $\emptyset^{A^+(\succsim^0)}$  and then find the informational difference  $\{w_1, w_0\} - \emptyset$  and the preferential difference  $\text{Must}(A)^-(A^+(\succsim^0))$ . Since  $\text{Must}(A)^-(A^+(\succsim^0)) = \succsim^0$ ,<sup>46</sup> this results in the original state  $\{\{w_1, w_0\} \succsim^0\}$ . The fact that negation dynamically manipulates preferences in this way guarantees that there will be no state, not even an irrational one, which contains the preferences both  $\text{Must}(A)$  and  $\neg\text{Must}(A)$  promote. If that state already has the preferences in it that  $\text{Must}(A)$  would add, then  $\neg\text{Must}(A)$  would remove them. Conversely, if the state already lacks the preferences  $\neg\text{Must}(A)$  would remove then  $\text{Must}(A)$  would add them. This is the gist of my expressivist explanation of why  $\text{Must}(A)$  and  $\neg\text{Must}(A)$  are inconsistent.

<sup>45</sup> This analysis of negation is inspired by the converse operator of Dynamic Logic (Harel et al., 2000, p. 177), which serves as a variety of program negation. While there is a growing literature on using dynamic logic for deontic reasoning, it does not directly speak to the key issues here.

<sup>46</sup> Recall that  $\succsim^0 = \{\langle w_0, w_1 \rangle\}$ , so  $A^+(\succsim^0) = \{\langle w_1, w_0 \rangle, \langle w_0, w_1 \rangle\}$ . Then  $\text{Must}(A)^-(A^+(\succsim^0)) = \text{Must}(A)^-(\{\langle w_1, w_0 \rangle, \langle w_0, w_1 \rangle\})$ . To find  $\text{Must}(A)^-(\{\langle w_1, w_0 \rangle, \langle w_0, w_1 \rangle\})$  one removes any strict preference that ends up in the result state of  $\{W^{(\emptyset,=)}\}[\text{Must}(A)] = \{W^{\{\langle w_1, w_0 \rangle, \dots, \langle w_0, w_1 \rangle\}}\}$ . So  $\langle w_1, w_0 \rangle$  is removed from  $\{\langle w_1, w_0 \rangle, \langle w_0, w_1 \rangle\}$  leaving  $\{\langle w_0, w_1 \rangle\} = \succsim^0$ .



However, recall that  $\{\{w_1, w_0\}^{\succsim^0}\}$  is a state where the test imposed by  $\text{Must}(\neg A)$  is also successful. How exactly are  $\neg\text{Must}(A)$  and  $\text{Must}(\neg A)$  distinguished on this semantics? This is where states with equally preferred worlds are required.

Consider a state where  $w_0 \sim w_1$ , and there are no strict preferences. Intuitively, this state supports  $\neg\text{Must}(A)$ , but not  $\text{Must}(\neg A)$ . If one prefers  $A$  and  $\neg A$  equally, one definitely isn't compelled to choose only  $\neg A$ -worlds or to choose only  $A$ -worlds. The first fact means that one rejects  $\text{Must}(\neg A)$ , while the second means that one accepts  $\neg\text{Must}(A)$ . To see how the formalism above replicates this intuitive reasoning, consider a state  $\{s^{\succsim}\}$  where  $s = \{w_1, w_0\}$ ,  $\succ = \emptyset$ , and  $\sim = \{\langle w_0, w_0 \rangle, \langle w_1, w_1 \rangle, \langle w_1, w_0 \rangle, \langle w_0, w_1 \rangle\}$ . Updating  $\{s^{\succsim}\}$  with  $\text{Must}(\neg A)$  will augment  $\succ$  to  $\succ_0 = \{\langle w_0, w_1 \rangle\}$  (leaving  $\sim$  as it was), and test whether the choice worlds relative to the augmented ordering are the  $\neg A$ -worlds:  $\{w_0\}$ . The only non-dominated alternative in  $\succ_0$  is  $w_0$ , but recall that  $w_0 \sim w_1$ . Since any choice world is either non-dominated, or equally preferred to a non-dominated one, the choice worlds are  $\{w_0, w_1\}$  and so the test is failed. In the same way, the test for  $\text{Must}(A)$  will fail, and result in the state  $\{\emptyset^{\succsim^1}\}$ , where  $\succ_1 = \{\langle w_1, w_0 \rangle\}$ . But this means that updating the original state with  $\neg\text{Must}(A)$  will involve taking the informational difference  $\{w_1, w_0\} - \emptyset$  and removing the preference  $\langle w_1, w_0 \rangle$ . And so an update of  $\{s^{\succsim}\}$  with  $\neg\text{Must}(A)$  returns  $\{s^{\succsim}\}$  while an update of  $\{s^{\succsim}\}$  with  $\text{Must}(\neg A)$  returns  $\{\emptyset^{\succsim^1}\}$ . In other words:  $\{s^{\succsim}\}$  supports  $\neg\text{Must}(A)$  but does not support  $\text{Must}(\neg A)$ .

It is important to note that I do not distinguish  $\neg\text{Must}(A)$  and  $\text{Must}(\neg A)$  by appealing to a state where there are two competing substates, one where  $A$  is strictly preferred and one where  $\neg A$  is strictly preferred, e.g.  $\{\{w_1, w_0\}^{\succsim^1}, \{w_1, w_0\}^{\succsim^0}\}$ . That is the kind of state that results from updating with  $\text{Must}(A) \vee \text{Must}(\neg A)$ .<sup>47</sup> This is *not* a state which supports  $\neg\text{Must}(A)$ . Updating with  $\neg\text{Must}(A)$  after updating with the disjunction would change the preferences and lead one to rule out the substate produced by the left disjunct, thereby bringing one to a state where only the substate produced by the right disjunct remains. In other words, one would infer  $\text{Must}(\neg A)$ . This illustrates why it is so important to distinguish  $\neg\text{Must}(A)$  from  $\text{Must}(A) \vee \text{Must}(\neg A)$ . Conflating them straightforwardly leads to a collapse of  $\neg\text{Must}(A)$  into  $\text{Must}(\neg A)$  (Unwin, 2001; Schroeder, 2008c; Dreier, 2009). Mention of logical matters highlights the fact that above I have been relying on the logical notions of support and consequence roughly in the vein of §5. But it is crucial to make these notions explicit and to note that they arise in the way expressivists require. In particular it will be important to see how the semantic properties of deontic sentences are explained by properties of the attitudes they express.

<sup>47</sup> This dynamic semantics of disjunction comes from Starr (2013) where it is used to analyze disjunctions of imperatives, declaratives (and combinations thereof). It works by dynamically generalizing the idea that disjunctions introduce alternative propositions (Kratzer and Shimoyama, 2002; Simons, 2005; Alonso-Ovalle, 2006), to the idea that disjunctions introduce alternative updates.

### 6.3 Expressivism Redux

How does dynamic expressivism explain the logical relations between expressive sentences? Can one simply adopt the definitions of support, consistency and consequence from §5, or does the added dynamics of preference make other definitions suitable? Since support is the central logical concept, that is the place to start. Generalizing the previous definition of support (§5, Definition 7) would amount to saying that  $S$  supports  $\phi$  when  $S[\phi] = S$ . This is a rather strict form of support which uniformly requires no change to the information, number of substates or preferences of  $S$ . This notion of support yields an interesting, though radically non-classical, logic. However, for the purposes here it is somewhat more illuminating and conservative to consider two parallel notions of support, one specialized to information, one to preferences.<sup>48</sup>

**Definition 18 (Informational Support).**  $S \models \phi \iff i_S = i_{S[\phi]}$

- $i_S = \bigcup \{s \mid \exists \succsim: s \succsim \in S\}$

**Definition 19 (Preferential Support).**  $S \models \phi \iff Pref_S = Pref_{S[\phi]}$

- $Pref_S = \{\succsim \mid \exists s \neq \emptyset: s \succsim \in S\}$

Informational support is familiar from §5, but preferential support tracks the additional expressive dimension added by deontic modals. It requires that the preference frames in  $S$  stay the same after updating with  $\phi$  and that  $\phi$  does not add any preference frames.<sup>49</sup> This furnishes two notions of consequence and of consistency:

**Definition 20 (Informational Consequence).**  $\phi_1, \dots, \phi_n \models \psi \iff \forall S: S[\phi_1] \dots [\phi_n] \models \psi$

**Definition 21 (Preferential Consequence).**  $\phi_1, \dots, \phi_n \models \psi \iff \forall S: S[\phi_1] \dots [\phi_n] \models \psi$

**Definition 22 (Informational Consistency).**

$\phi_1, \dots, \phi_n$  are informationally consistent  $\iff \exists S: i_S \neq \emptyset \ \& \ S \models \phi_1, \dots, S \models \phi_n$

**Definition 23 (Preferential Consistency).**

$\phi_1, \dots, \phi_n$  are preferentially consistent  $\iff \exists S: Ch(S) \neq \emptyset \ \& \ S \models \phi_1, \dots, S \models \phi_n$

- Where  $Ch(S) = \bigcup \{Choice(s, \succsim) \mid s \succsim \in S\}$

<sup>48</sup> There is actually a third to consider, **choice-support**, which produces a more classical logic for **Must**:  $S$  choice-supports  $\phi$  just in case  $\bigcup \{Choice(s, \succsim) \mid s \succsim \in S\} = \bigcup \{Choice(s, \succsim) \mid s \succsim \in S[\phi]\}$ . This will predict that both **Must**( $\phi$ ) and **Must**( $\psi$ ) are consequences of **Must**( $\phi \wedge \psi$ ). However it also predicts that **Must**( $\phi \vee \psi$ ) is a consequence of **Must**( $\phi$ ). See §7 for discussion of these principles.

<sup>49</sup> It also discounts changes to the preferences that concern only problematic substates whose information is contradictory:  $\emptyset \succsim$ . Intuitively, those preferences are to be ignored since they are not constraining agents' choices:  $Choice(\emptyset, \succsim)$  will always be  $\emptyset$ . This restriction is needed to predict that **Must**(**B**) is a preferential consequence of **Must**(**A**)  $\vee$  **Must**(**B**) and  $\neg$ **Must**(**A**).

- Recall: if  $Choice(s, \succsim) = \emptyset$  then  $\succsim$  is dysfunctional, i.e. fails to motivate a choice.
  - E.g. if  $\succsim$  is cyclic over  $s$ ,  $Choice(s, \succsim) = \emptyset$

For informational consistency, I require a state that is not representationally dysfunctional to informationally support each of the sentences. By contrast, preferential consistency requires a state that is not *motivationally dysfunctional* to preferentially support each of the sentences. This allows a non-representational explanation of consistency which parallels our representational explanations. For example,  $A$  and  $\neg A$  are inconsistent because their function is to provide information to  $i_S$ ,  $i_S$  has the function of representing the world, but only a dysfunctional  $i_S$  has the information each of them provides. Similarly,  $Must(A)$  and  $Must(\neg A)$  are inconsistent because their function is to update the preferences in  $Pref_S$ ,  $Pref_S$  has the function of motivating choices but only a dysfunctional  $Pref_S$  contains the preferences each of them promotes. Recall from above that a state which contains a symmetric strict preference ordering ranking  $A$ -worlds over  $\neg A$ -worlds and  $\neg A$ -worlds over  $A$ -worlds will preferentially support both  $Must(A)$  and  $Must(\neg A)$ . But such a state is motivationally dysfunctional:  $Ch(S) = \emptyset$ . It is important to highlight that  $Must(A)$  and  $\neg Must(A)$  are inconsistent for a different reason. It is not that only an irrational state can support both. It's that there is no state which supports both because they perform converse updates on the preferences. If a state has the preferences that  $Must(A)$  would add, then  $\neg Must(A)$  would remove them. If a state lacks the preferences  $\neg Must(A)$  would remove then  $Must(A)$  would add them. This is an essentially dynamic form of inconsistency: one formula builds something which the other removes. It is precisely this resource which pragmatic expressivism lacked, and which made its explanation of the inconsistency between  $\neg Must(A)$  and  $Must(A)$  seem suspiciously representational. The beauty of the dynamic account is that one can add this new element while keeping the explanations of consistency parallel. In both the motivational and representational cases, consistency is allied with a function our mental states serve, and inconsistent sentences are explained by the fact that if they can be embraced by a single state at all, it is a dysfunctional state.

It was hard to motivate the pragmatic expressivist view that the contents communicated by deontics were non-representational. Those contents seemed open to a representational interpretation since they were reified from a representational theory of consequence and had identical dynamics. On the dynamic view, the underlying theory of consequence is not inherently representational. It tracks various fixed points in the interpretive process—in information, preferences, etc. Even though it has a representational edition—when that fixed point concerns information—the conception of logic is more general. It is precisely this more general approach to meaning and logic that is essential to capturing the expressive dimension of deontic modals.

On the semantics in Definitions 16 and 17, updates involving deontic modals do not refer to preferences and describe them. Those sentences promote and discount preferences. This variety of update differs radically from descriptive, informational

updates. Crucially, they differ not just in the contents (preferences vs propositions) they manipulate, but in *how* those contents are manipulated.<sup>50</sup> The adding and removing of preferences is quite different from the way descriptive discourse works. Descriptive discourse provides information by making distinctions between worlds. As I mentioned in §5, this way of providing information can be formally characterized as a distributive and eliminative update:  $s[\phi] = \bigcup\{\{w\}[\phi] \mid w \in s\}$ , and  $s[\phi] \subseteq s$ . One way of putting the problem for pragmatic expressivism was that it was indistinguishable from a view where expressive discourse works via a distributive and eliminative update over more fine-grained worlds such as  $\langle w, \succ \rangle$ . It then becomes hard to elucidate the expressivist communication thesis, since it looks as if expressive communication works by making distinctions between centered worlds, and one of those distinctions is made by referring to preferences. But on the dynamic view, the fact that preferences are added, while worlds are removed, guarantees that preference update and information update cannot be the same process. Indeed, on the dynamic account there are four kinds of updates: adding preferences (deontic modals), removing preferences (negation), eliminating worlds (atomics, negation) and creating substates (disjunction). Language compositionally interleaves these distinct processes rather than providing a single semantic object which is factored into distinct contributions.

Ultimately, the dynamic account makes room for the expressivist communication thesis by breaking from the traditional ‘conduit’ metaphor for communication: expressing yourself is like putting a package on a conveyor belt between minds. That metaphor assumes there is a discrete object shared and a single process by which you unpack the box and sort the contents. On the dynamic approach, a better metaphor would be two minds connected by a series of vibrating cords which encode different modalities with different frequency ranges. Language is a virtuosic tool for weaving these frequencies and modalities together. This departure from the conduit metaphor is far from a debt for expressivists. It is one urged by researchers not debating expressivism and spanning the fields of animal communication (Owren et al., 2010; Scott-Phillips and Kirby, 2013) and natural language (Murray and Starr, 2012).

## 7 A Wider View

I want to conclude this chapter by taking a wider view and being clear about what has and has not been accomplished. I claim to have offered a dynamic semantics which vindicates the expressivist communication thesis, and done so better than competing pragmatic accounts. Dynamic meanings were essential for implementing the expressive effect of deontic modals, and generalizing our familiar semantics for

<sup>50</sup> *Varieties of update* is an allusion to Murray (2014), which exploits dynamic semantics to compositionally interleave a number of distinct update processes. It is there applied to a number of constructions across languages which manipulate but differentiate at-issue and not-at-issue content.

negation so as to interact with that expressive effect. This provided a new expressivist solution to the negation problem.  $\text{Must}(A)$  and  $\neg\text{Must}(A)$  have converse dynamic effects, which predicts that they are inconsistent because there is no state that can contain the converse effects of their updates.  $\text{Must}(A)$  and  $\text{Must}(\neg A)$  are inconsistent because only irrational preferences support them. It was argued that these explanations of inconsistency perfectly parallel those for descriptive discourse. There is nevertheless much work left to do. There are still several constructions which have not been analysed here and may turn out to be problematic. For example, the interaction of modals with attitude verbs, modal questions, deontic predicates/adverbs such as *wrong/wrongly*, more complex modals such as *ought*, *should*, *could* and non-deontic uses of *must* and *may* all lie outside the scope of the analysis here.<sup>51</sup> Furthermore, at some point one would like a positive argument in favor of an expressivist semantics over the impressive achievements of its descriptive competitors such as Kratzer (1991, 2012). This must wait for another day, but some suggestive remarks on a few fronts may help chart the territory where such an argument might be found.

The compositional virtuosity of the dynamic account was essential for analyzing mixed disjunctions. It explained how communicating  $\text{Must}(A) \vee B$  was different from communicating  $\text{Must}(A)$  or communicating  $B$ . The disjunction involves a distinctive substate-creating update, where quite different updates are unleashed on each incoming substate to create two new substates. It is worth highlighting how this addresses the concern raised by Schroeder (2015) about the interaction of expressive modals and attitude verbs. That concern is whether  $\text{Bel}_X(\text{Must}(A) \vee B)$  will entail  $\text{Bel}_X(\text{Must}(A)) \vee \text{Bel}_X(B)$ .  $\text{Bel}_X(\phi)$  can be analyzed as a descriptive update requiring of a world  $w$  that the state  $S_X$  modeling  $X$ 's state of mind in  $w$  both informationally and preferentially supports  $\phi$ . But consider a world in which  $S_X$  is a state with two substates:  $S_X = \{s \sim^A, s \sim^B\}$ , the first has information  $s$  and an ordering which prefers  $A$ -worlds, and the second lacks that particular preference but has the information carried by  $B$ .  $S_X$  will preferentially and informationally support  $\text{Must}(A) \vee B$ . Yet  $S_X$  will not preferentially support  $\text{Must}(A)$ , since  $\text{Pref}_{S_X[\text{Must}(A)]}$  will lack  $\sim$ . Also,  $S_X$  will not informationally support  $B$ , since  $i_{S_X[B]}$  will exclude the  $\neg B$  worlds in  $s$ . Thus  $\text{Bel}_X(\text{Must}(A) \vee B)$  will not informationally entail  $\text{Bel}_X(\text{Must}(A)) \vee \text{Bel}_X(B)$ . The success of this analysis would not surprise Schroeder (2015). He acknowledges that if belief ascriptions mix information and motivation, and there is a special sort of state for believing disjunctions, then the entailment can be blocked. What he doubts is a philosophically plausible interpretation of these constructs. Even the small amount I have said here suggests a way of resisting this doubt. A substate is an informational and motivational perspective, and a state containing multiple substates captures the perspectives competing for control over the agent's actions. There seems to be no good

<sup>51</sup> I think there is some hope for analysing  $?\text{Must}(A)$  as  $S[\text{Must}(A)] \cup S[\neg\text{Must}(A)]$ .

philosophical or psychological motivation for assuming that beliefs are exclusively informational or motivational, or that different perspectives can't compete within a single agent. Assuming beliefs are 'purely informational' seems philosophically tidy, but psychologically implausible to me. Surely our actual beliefs are influenced by the motivational perspective we adopt, just as our desires are influenced by the informational perspective we adopt.<sup>52</sup>

The dynamic account allows one to articulate two components of a deontic modal's meaning: preference promotion and a test about that promotion. This leads one to wonder whether language contains operators which simply promote preferences. That is essentially Starr's (2013) analysis of imperatives:<sup>53</sup>

**Definition 24 (Imperative Semantics).**  $S[!\phi] = \{s^{\phi^+}(\tilde{s}) \mid s\tilde{s} \in S\}$

This illuminates one difference between deontic modals and imperatives. Unlike deontic modals, negation does not scope over imperatives in natural language (Han, 2001). On the semantics above, negating an imperative would disastrously subtract each  $s$  from itself, returning  $\emptyset$  for each. Negation thus turns any imperative into an informational contradiction. This goes some way towards explaining what many other theories treat as an arbitrary syntactic restriction: negation scoped over an imperative does not produce a useful meaning, so no human language should bother to conventionalize those syntactic structures. Explaining this fact by syntactic stipulation looks less satisfying when it becomes clear that languages allow imperatives to robustly embed under other connectives:  $!A \wedge !B, A \wedge !B, !A \vee !B, A \vee !B$  (Starr, 2013). Starr (2013) presents two related arguments that build on this observation. First, it is argued there that no plausible non-propositional analysis of imperatives can compositionally account for this behavior without dynamic meanings like those above. Secondly, it is argued that non-propositional non-dynamic theories such as Portner (2012) incorrectly predict  $!(A \wedge B)$  to be a consequence of  $(!A \vee !B) \vee !(A \wedge B)$ , but that a dynamic theory such as the above does not. This highlights that dynamic meanings provide not only new compositional resources, but logical ones too. I turn now to this logical frontier of dynamic thinking.

The deontic realm is a minefield of logical problems, but some recent work suggests that dynamic semantics provides resources for resolving them. Willer (2014), for example, shows that a dynamic non-monotonic consequence relation relieves the pressures placed on modus ponens in deontic settings by Chisholm's Paradox (Chisholm, 1963), the Miner's Paradox (Kolodny and MacFarlane, 2010) and the Gentle

<sup>52</sup> Charlow (2015) explores these issues in much more depth. He pursues an account which I learned much from and which bears at least some affinities with the account pursued here.

<sup>53</sup> Given this analysis of imperatives, one may draw a comparison between the semantics above and the analysis in Ninan (2005) where imperatives and deontic *must* have the same effect on context. On the analysis above, the semantics of deontic *must* includes that effect, but adds the crucial element of a test. This approach also predicts Ninan's (2005) key data:  $\neg A \wedge \text{Must}(A)$  will be infelicitous. Further investigation will be needed to see whether the added test element is empirically advantageous.

Murder Paradox (Forrester, 1984). Willer, (this volume, Chapter 11) elaborates on the relationship between this kind of non-monotonic approach and one based on default logic. The semantics above makes two predictions that may add momentum to this movement. Suppose a father tells his son that he must take out the trash or do the dishes. The son proceeds to the trashcan only to find that it has already been taken out. The son should infer that he must wash the dishes. However, if the son goes to do the dishes and they too have been taken care of, he may no longer sensibly think he must wash the dishes. These are facts classical theories do not predict:

- (11)  $\text{Must}(A \vee B), \neg A \vDash \text{Must}(B)$
- (12)  $\text{Must}(A \vee B), \neg A, \neg B \not\vDash \text{Must}(B)$

Just because the best worlds are all  $A \vee B$  worlds and we aren't in an  $A$ -world, it doesn't follow that all the best worlds are  $B$ -worlds. Further, the pattern in (12) seems to show exactly the kind of non-monotonicity which classical consequence relations fail to predict. By contrast, both patterns are predicted by the analysis above.

Now return to the scenario where the son has correctly inferred that he must do the dishes. Note that he should not infer that he must do the dishes or watch a movie. He also shouldn't infer that either he must do the dishes or he must watch a movie:

- (13)  $\text{Must}(B) \not\vDash \text{Must}(B \vee C)$
- (14)  $\text{Must}(B) \not\vDash \text{Must}(B) \vee \text{Must}(C)$

This is a modal version of Ross's Paradox (Ross, 1941) concerning imperatives. Their validity in standard approaches to modality has been widely recognized as problematic. Several attempts have been made to leverage pragmatic reasoning and a non-standard semantics for disjunction to explain them (Zimmermann, 2000; Simons, 2005; Geurts, 2005; Aloni, 2007). Yet Cariani (2013) and Lassiter (2011) argue that a strictly semantic approach is justified. While their semantic approach predicts (13), it does not, to my knowledge, predict (14). The analysis above predicts both.<sup>54</sup> While none of these prospects constitutes a full empirical argument for the kind of dynamic analysis pursued here, they illustrate some of the territory where those arguments may evolve. It would, of course, also be unsurprising if dynamic expressivism proved useful in the burgeoning literature on non-descriptive meaning.

## Acknowledgments

I am extremely thankful to Andrew Alwood, Nate Charlow, Matthew Chrisman, Andy Egan, Sarah Murray, Paul Portner, Angelika Kratzer, Aynat Rubenstein, Mark Schroeder, Alex Silk, Justin Snedegar, Malte Willer and audiences at Edinburgh, NYU, Rutgers and USC, some for

<sup>54</sup> Like Cariani (2013) and Lassiter (2011) the analysis here predicts that  $\text{Must}(A \wedge B) \not\vDash \text{Must}(A)$ , which is one way to explain Jackson's (1985) 'professor procrastinate' puzzle.

detailed feedback and others for brief, incisive questions. Charlow and Chrisman deserve special thanks, not only for editing this volume, but also for guiding this project towards properly engaging with the expressivism literature. I should also acknowledge a debt to Yuna Won. I have learned at least as much from advising her PhD research on deontic logic as she has learned from me.

## References

- Aloni, M. (2007) 'Free Choice, Modals and Imperatives.' *Natural Language Semantics*. 15. 1. pp. 65–94.
- Alonso-Ovalle, L. (2006) *Disjunction in Alternative Semantics*. Ph.D. thesis. UMass Amherst, Amherst, MA.
- Alwood, A. (2015) 'Nondescriptive Negation for Normative Sentences.' (To be published in *Philosophical Quarterly*.)
- Anderson, L. and Lepore, E. (2013). 'Slurring Words.' *Noûs*. 47. 1. pp. 25–48.
- Beaney, M. (ed.) (1997) *The Frege Reader*. Malden, MA: Blackwell.
- Bittner, M. (2001) 'Topical Referents for Individuals and Possibilities.' In Hastings, R., Jackson, B. and Zvolenszky, Z. (eds.) *Proceedings from Semantics and Linguistic Theory 11*. pp. 36–55. Ithaca, NY: Cornell University.
- Blackburn, S. (1988) 'Attitudes and Contents.' *Ethics*. 98. 3. pp. 501–17.
- Burge, T. (1975) 'On Knowledge and Convention.' *The Philosophical Review*. 84. pp. 249–55.
- Carballo, AP. (2014) 'Semantic Hermeneutics.' In Burgess, A. and Sherman, B. (eds.) *Metasemantics: New Essays on the Foundations of Meaning*. pp. 119–46. New York, NY: Oxford University Press.
- Cariani, F. (2013) '“Ought” and Resolution Semantics.' *Noûs*. 47. 3. pp. 534–58.
- Charlow, N. (2013) 'Logic and Semantics for Imperatives.' *Journal of Philosophical Logic*. pp. 1–48.
- Charlow, N. (2014) 'The problem with the Frege–Geach problem.' *Philosophical Studies*. 167. 3. pp. 635–65.
- Charlow, N. (2015) 'Prospects for an Expressivist Theory of Meaning.' *Philosophers' Imprint*. 15. 23. pp. 1–43. [Online] Available from: <http://hdl.handle.net/2027/spo.3521354.0015.023>. [Accessed 5 Nov 2015.]
- Chisholm, R. M. (1963) 'Contrary-to-Duty Imperatives and Deontic Logic.' *Analysis*. 24. 2. pp. 33–6.
- Chrisman, M. (2012) 'On the Meaning of “Ought”.' In Shafer-Landau, R. (ed.) *Oxford Studies in Metaethics*. Volume 7. pp. 304–32. Oxford: Oxford University Press.
- Dreier, J. (2006) 'Negation for Expressivists: a collection of problems with a suggestion for their solution.' In Shafer-Landau, R. (ed.) *Oxford Studies in Metaethics*. Volume 1. pp. 217–33. New York, NY: Oxford University Press.
- Dreier, J. (2009) 'Relativism (And Expressivism) And The Problem Of Disagreement.' *Philosophical Perspectives*. 23. 1. pp. 79–110.
- Egan, A. (2007) 'Epistemic Modals, Relativism and Assertion.' *Philosophical Studies*. 133. 1. pp. 1–22.
- Faller, M. (2002) *Semantics and Pragmatics of Evidentials in Cuzco Quechua*. Ph.D. thesis. Stanford University, Palo Alto, CA.



- Forrester, J. W. (1984) 'Gentle Murder, or the Adverbial Samaritan.' *Journal of Philosophy*. 81. pp. 193–7.
- Franklin, S. (1995) *Artificial Minds*. Cambridge, MA: MIT Press.
- Frege, G. (1879) *Begriffsschrift, eine der arithmetischen nachgebildete Formalsprache des reinen Denkens*. Halle: L. Nebert. pp. 45–78. References to M. Beaney's translation in Beaney (1997) pp. 45–78.
- Frege, G. (1923) 'Logische Untersuchungen.' *Beiträge zur Philosophie des deutschen Idealismus*. 3. pp. 36–51. References to Frege (1963).
- Frege, G. (1963) 'Compound Thoughts.' *Mind*. 72. 285. pp. 1–17. Translation of Frege (1923).
- Geach, P. T. (1965) 'Assertion.' *The Philosophical Review*. 74. 4. pp. 449–65.
- Geurts, B. (2005) 'Entertaining Alternatives: Disjunctions as Modals.' *Natural Language Semantics*. 13. 4. pp. 383–410.
- Gibbard, A. (1986) 'An Expressivistic Theory of Normative Discourse.' *Ethics*. 96. 3. pp. 472–85.
- Gibbard, A. (2003) *Thinking How to Live*. Cambridge, MA: Harvard University Press.
- Grice, H. P. (1968) 'Utterer's Meaning, Sentence-Meaning, and Word-Meaning.' *Foundations of Language*. 4. pp. 225–42.
- Groenendijk, J. and Stokhof, M. (1991) 'Dynamic Predicate Logic.' *Linguistics and Philosophy*. 14. 1. pp. 39–100.
- Groenendijk, J., Stokhof, M. and Veltman, F. (1996) 'Coreference and Modality.' In Lappin, S. (ed.) *The Handbook of Contemporary Semantic Theory*. pp. 179–213. Oxford: Blackwell.
- Halpern, J. Y. (2003) *Reasoning about Uncertainty*. Cambridge MA: MIT Press.
- Hamblin, C. L. (1973) 'Questions in Montague English.' *Foundations of Language*. 10. 1. pp. 41–53.
- Han, C. H. (2001) 'Force, negation and imperatives.' *Linguistic Review*. 18. 4. pp. 289–325.
- Hansson, S. O. and Grüne-Yanoff, T. (2011) 'Preferences.' In Zalta (ed.), *The Stanford Encyclopedia of Philosophy*, fall. [Online] Available from: <http://plato.stanford.edu/archives/fall2011/entries/preferences/>. [Accessed 5 Nov 2015.]
- Harel, D., Kozen, D. and Tiuryn, J. (2000) *Dynamic Logic*. Cambridge, MA: MIT Press.
- Heim, I. R. (1982) *The Semantics of Definite and Indefinite Noun Phrases*. Ph.D. thesis. University of Massachusetts, Amherst, Massachusetts.
- Heim, I. (1983) 'File Change Semantics and the Familiarity Theory of Definiteness.' In Bäuerle, R., Schwarze, C. and von Stechow, A. (eds.) *Meaning, Use, and Interpretation of Language*. pp. 164–89. Berlin: Walter de Gruyter.
- Hobbs, J. R. (1990) *Literature and Cognition*. Stanford, CA: Center for the Study of Language and Information.
- Jackson, F. (1985) 'On the Semantics and Logic of Obligation.' *Mind*. 94. 374. pp. 177–95.
- Kamp, H. (1981) 'A Theory of Truth and Semantic Representation.' In Groenendijk, J. A., Janssen, T. and Stokhof, M. (eds.) *Formal Methods in the Study of Language*. pp. 277–322. Dordrecht: Foris.
- Kolodny, N. and MacFarlane, J. (2010) 'Ifs and Oughts.' *Journal of Philosophy*. 107. 3. pp. 115–43.
- Kratzer, A. (1991) 'Modality.' In von Stechow, A. and Wunderlich, D. (eds.) *Semantics: An International Handbook of Contemporary Research*. pp. 639–50. Berlin: De Gruyter Mouton.
- Kratzer, A. (2012) *Modals and Conditionals: New and Revised Perspectives*. New York, NY: Oxford University Press.
- Kratzer, A. and Shimoyama, J. (2002) 'Indeterminate Pronouns: the View from Japanese.' In Otsu, Y. (ed.) *Proceedings of the Third Tokyo Conference on Psycholinguistics*. Tokyo: Hituzi Syobo.

- Lassiter, D. (2011) *Measurement and Modality: the scalar basis of modal semantics*. Ph.D. thesis. New York University.
- Lewis, D. K. (1969) *Convention: A Philosophical Study*. Cambridge, MA: Harvard University Press.
- Mackie, J. L. (1977) *Ethics: Inventing Right and Wrong*. New York: Penguin.
- McCready, E. (2012) 'Emotive equilibria.' *Linguistics and Philosophy*. 35. 3. pp. 243–83.
- Millikan, R. G. (1984) *Language, Thought and Other Biological Categories*. Cambridge, MA: MIT Press.
- Millikan, R. G. (2005). *Language: A Biological Model*. New York, NY: Oxford University Press.
- Moore, G. E. (1912) *Ethics*. Oxford: Oxford University Press.
- Moss, S. (2015) 'On the Semantics and Pragmatics of Epistemic Vocabulary.' *Semantics and Pragmatics*. 8. 5, pp. 1–81. [Online] Available from: <http://dx.doi.org/10.3765/sp8.5> [Accessed 5 Nov 2015.]
- Murray, S. E. (2010) *Evidentiality and the Structure of Speech Acts*. Ph.D. thesis. Rutgers University.
- Murray, S. E. (2014) 'Varieties of Update.' *Semantics and Pragmatics*. 7. 2, pp. 1–53. [Online] Available from: <http://dx.doi.org/10.3765/sp.7.2>. [Accessed 5 Nov 2015.]
- Murray, S. E. and Starr, W. B. (2012) 'The Structure of Communicative Acts.' Manuscript Cornell University.
- Muskens, R. van Benthem, J. and Visser, A. (1997) 'Dynamics.' In van Benthem, J. and ter Meulen, A. (eds.) *Handbook of Logic and Language*. 1st edition. pp. 587–648. Cambridge, MA: MIT Press.
- Ninan, D. (2005) 'Two Puzzles About Deontic Necessity.' In Gajewski, J., Hacquard, V. and Yalcin, S. (eds.) *New Work on Modality, MIT Working Papers in Linguistics*. Volume 51. Cambridge, MA: MIT Press.
- Owren, M. J., Rendall, D. and Ryan, M. J. (2010) 'Redefining Animal Signaling: influence versus information in communication.' *Biology and Philosophy*. 25. 5. pp. 755–80.
- Portner, P. (2009) *Modality*. New York, NY: Oxford University Press.
- Portner, P. (2012) 'Permission and Choice.' In Grewendorf, G. and Zimmermann, T. E. (eds.) *Discourse and Grammar: From Sentence Types to Lexical Categories*. Berlin: Mouton de Gruyter.
- Potts, C. (2007) 'The Expressive Dimension.' *Theoretical Linguistics*. 33. 2. pp. 165–97.
- Ridge, M. (2014) *Impassioned Belief*. New York, NY: Oxford University Press.
- Ross, A. (1941) 'Imperatives and Logic.' *Theoria*. 25. 7. pp. 53–71. References to Ross (1944).
- Ross, A. (1944) 'Imperatives and Logic.' *Philosophy of Science*. 11. 1. pp. 30–46.
- Rothschild, D. (2012). 'Expressing Credences.' *Proceedings of the Aristotelian Society*. 112. 1. pp. 99–114.
- Schroeder, M. (2008a). *Being For: Evaluating the Semantic Program of Expressivism*. New York, NY: Oxford University Press.
- Schroeder, M. (2008b). 'Expression for Expressivists.' *Philosophy and Phenomenological Research* 76. 1. pp. 86–116.
- Schroeder, M. (2008c). 'How Expressivists Can and Should Solve Their Problem with Negation.' *Noûs*. 42. 4. pp. 573–99.
- Schroeder, M. (2008d) 'What is the Frege-Geach Problem?' *Philosophy Compass*. 3. 4. pp. 703–20.
- Schroeder, M. (2013) 'Two Roles for Propositions: Cause for Divorce?' *Noûs*. 47. 3. pp. 409–30.

- Schroeder, M. (2015) Attitudes and Epistemics. *Expressing Our Attitudes: explanation and expression in ethics*. Volume II. New York, NY: Oxford University Press.
- Scott-Phillips, T. and Kirby, S. (2013) 'Information, Influence and Inference in Language Evolution.' In Stegmann, U. (ed.) *Animal Communication Theory: information and influence*. pp. 421–42. Cambridge: Cambridge University Press.
- Silk, A. (2013) 'Truth-Conditions and the Meanings of Ethical Terms.' In Shafer-Landau, R. (ed.) *Oxford Studies in Metaethics*. Volume 8. pp. 195–222. New York, NY: Oxford University Press.
- Silk, A. (2014) 'How to Be an Ethical Expressivist.' *Philosophy and Phenomenological Research*. 91. 1. pp. 47–81.
- Simons, M. (2005) 'Dividing Things Up: The Semantics of *or* and the Modal/*or* Interaction.' *Natural Language Semantics*. 13. 3. pp. 271–316.
- Stalnaker, R. C. (1978) 'Assertion.' In Cole, P. (ed.) *Syntax and Semantics 9: Pragmatics*. pp. 315–32. New York, NY: Academic Press. References to Stalnaker (1999).
- Stalnaker, R. C. (1984) *Inquiry*. Cambridge, MA: MIT Press.
- Stalnaker, R. C. (1999) *Context and Content: Essays on Intentionality in Speech and Thought*. Oxford: Oxford: Oxford University Press.
- Starr, W. B. (2010) *Conditionals, Meaning and Mood*. Ph.D. thesis. Rutgers University.
- Starr, W. B. (2013) 'A Preference Semantics for Imperatives.' Ms. Cornell University.
- Swanson, E. (2006). *Interactions with Context*. Ph.D. thesis. MIT.
- Swanson, E. (forthcoming). 'The Application of Constraint Semantics to the Language of Subjective Uncertainty.' *Journal of Philosophical Logic*. [Online] Available from: [link.springer.com/article/10.1007%27F510992-015-9367-5](http://link.springer.com/article/10.1007%27F510992-015-9367-5). [Accessed 5 Nov 2015.]
- Unwin, N. (2001) 'Norms and Negation: A Problem for Gibbard's Logic.' *The Philosophical Quarterly*. 51. 202. pp. 60–75.
- van Benthem, J. (1989) 'Semantic Parallels in Natural Language and Computation.' In Ebbinghaus, H. D., Fernandez-Prida, J., Garrido, M., Lascar, D. and Artalejo, M. R. (eds.) *Logic Colloquium '87: Proceedings of the Colloquium held in Granada, Spain*. Volume 129 of *Studies in Logic and the Foundations of Mathematics*. pp. 331–75. Amsterdam: North-Holland.
- van Benthem, J. and Liu, F. (2007) 'Dynamic Logic of Preference Upgrade.' *Journal of Applied Non-Classical Logics*. 17. 2. pp. 157–82.
- van Eijck, J. and Visser, A. (2012) 'Dynamic Semantics.' In Zalta (ed.) *The Stanford Encyclopedia of Philosophy*, winter. [Online]. Available from: <http://plato.stanford.edu/archives/win2012/entries/dynamic-semantics/>. [Accessed 5 Nov 2015.]
- Veltman, F. (1996) 'Defaults in Update Semantics.' *Journal of Philosophical Logic*. 25. 3. pp. 221–61.
- Willer, M. (2014) 'Dynamic Thoughts on Ifs and Oughts.' *Philosophers' Imprint*. 14. 28. pp. 1–30. [Online] Available from: <http://hdl.handle.net/2027/spo.3521354.0014.028>. [Accessed 5 Nov 2015.]
- Yalcin, S. (2007) 'Epistemic modals.' *Mind*. 116. 464. pp. 983–1026.
- Yalcin, S. (2011) 'Nonfactualism about Epistemic Modality.' In Egan, A. and Weatherson, B. (eds.) *Epistemic Modality*. pp. 295–332. New York, NY: Oxford University Press.
- Yalcin, S. (2012) 'Bayesian Expressivism.' *Proceedings of the Aristotelian Society Supplementary Volume*. 112. 2. pp. 123–60.
- Zimmermann, T. E. (2000) 'Free Choice Disjunction and Epistemic Possibility.' *Natural Language Semantics*. 8. 4. pp. 255–90.